UPTEC STS07 029
Examensarbete 20 p
Augusti 2007

Working memory enhancement and filtering of distractions modelled as a top-down signal from the dorsolateral prefrontal cortex

Pär Johansson

UPPSALA
UNIVERSITET

Teknisk- naturvetenskaplig fakultet
UTH-enheten

Besöksadress:
Ångströmlaboratoriet
Lägerhyddsvägen 1
Hus 4, Plan 0

Postadress:
Box 536
751 21 Uppsala

Telefon:
018 – 471 30 03

Telefax:
018 – 471 30 00

Hemsida:
http://www.teknat.uu.se/student

Abstract

Working memory enhancement and filtering of distractions modelled as a top-down signal from the dorsolateral prefrontal cortex

Pär Johansson

Working memory capacity (WMC), the amount of information a person can keep in fast and easily accessible storage, has been linked to a range of cognitive skills and diagnoses, including general intelligence and ADHD. Several studies have found an inverse relationship between WMC and distractibility, and it has been shown that the dorsolateral prefrontal cortex (dlPFC) is activated both when distractions need to be filtered away as well as when many items are held in working memory (WM). Here, we propose a dynamic mechanism by which activity in dlPFC can explain both filtering of distractions and WMC enhancement and explore this mechanism with computational modelling. In the model, a general activation of the dlPFC increases capacity of the memory store in the intraparietal sulcus (IPS) by raising its level of activity. Filtering can be achieved if a subpopulation of the dlPFC activates neurones in the IPS selective for a certain attribute. Stimuli with this attribute will be encoded, while stimuli with differing attributes will be filtered away due to increased inhibition. The dlPFC can hence provide a link between WMC and distractibility, as the amount of activity in dlPFC affects both WMC and distractibility. Three different hypotheses for how the dlPFC can be connected to IPS and a visual area (V4) have been tested, and fMRI activity was simulated for the hypotheses so that they can be differentiated experimentally. Knowledge about the rule encoding property of dlPFC can provide a framework for understanding intelligence and ADHD.

Handledare: Fredrik Edin
Ämnesgranskare: Mats Gustafsson
Examinator: Elisabet Andrésdóttir
ISSN: 1650-8319, UPTEC STS07 029

Populärvetenskaplig sammanfattning

Hjärnans arbetsminne håller information i ett snabbt och lättillgängligt förvar och kan ses som en länk mellan intryck från omgivningen och långtidsminne. Även när minnen från långtidsminnet framkallas för beslut eller tankeoperationer lagras de temporärt i arbetsminnet. Arbetsminnet kan bara hålla information en kort stund och bara ett antal saker i taget. Hur många saker en person kan komma ihåg samtidigt, arbetsminneskapacitet, varierar mellan ungefär 4-7. Arbetsminnet har en betydelsefull roll, det har visat sig att hög arbetsminneskapacitet hänger samman med hög intelligens, och låg arbetsminneskapacitet har kopplats till exempelvis ADHD. Typiskt för personer med ADHD är att de har svårt att filtrera bort distraktioner, alltså irrelevant information som egentligen kan bortses ifrån. Experiment har visat att förmågan att filtrera hänger väl samman med just arbetsminneskapaciteten.

Det är inte helt självklart vad i hjärnan det är som gör att hög arbetsminneskapacitet skulle göra det lättare att välja ut relevant information. Hur många minnen som kan lagras brukar nämligen anses bero på hur mycket aktivitet man har i ett område i den så kallade parietalloben. Filtrering av distraktioner handlar istället om att blockera intryck innan de lagras där, frågan är på vilket sätt det påverkas av kapaciteten. Teorin är att det finns ett annat område i hjärnan som kan påverka båda dessa förmågor. Detta skulle kunna vara ett område i pannloben som kallas dlPFC, det området har visat sig vara aktivt både när intryck blockeras och när man måste lagra väldigt många arbetsminnen. Man kan se dlPFC som en extra resurs som portioneras ut där det behövs beroende på vilken typ av uppgift man utför. Det är dock oklart hur detta fungerar, och målet med det här exjobbet har varit att bygga en datorbaserad modell där man kan simulera händelseförloppet för att se om teorin kan stämma.

Designen av modellen har gjorts för att efterlikna ett experiment där försökspersoner fick se röda och gula prickar på en skärm och ombads omväxlande att komma ihåg alla, och att bara komma ihåg de röda. Modellen simulerar aktivitet på cellnivå i nätverk kopplade för att efterlikna tre hjärnregioner som antas viktigast för att förklara vad som händer i experimentet, dlPFC, ett minnesområde och ett visuellt område.

Utan påverkan från dlPFC kan modellen lagra två minnen, alltså två prickar oberoende av färg. Om hela dlPFC området aktiveras skickas aktivitet till hela minnesområdet. Denna extra aktivitet förhöjer minneskapaciteten och gör att tre minnen kan lagras. För att efterlikna de fall då försökspersoner ombeds komma ihåg röda prickar men ignorera gula, aktiveras bara den del av dlPFC som ansvarar för rött. Aktiviteten i dlPFC förhöjer verksamheten i motsvarande del av minnesområdet och konkurrerar ut aktiviteten i den gula delen. När då både röda och gula prickar presenteras kommer röda prickar att lagras, men aktiviteten i det gula området är för låg för att kunna lagra dessa.

Resultaten visar att dlPFC på det här sättet kan vara ansvarigt både för förhöjd arbetsminneskapacitet och filtrering av distraherande stimuli. En lågfungerande dlPFC innebär låg arbetsminneskapacitet och oförmåga att filtrera distraktioner, typiska symptom för ADHD. En bra fungerande dlPFC innebär att arbetsminneskapaciteten kan höjas och distraktioner filtreras. En välutvecklad förmåga att uttrycka regler på det här sättet kan också vara en viktig länk till att förstå hur arbetsminneskapacitet och aktivitet i dlPFC är kopplat till intelligens.

Content

1 Introduction

Working memory (WM) refers to the ability to keep information in a fast and easily accessible storage for manipulation (Baddeley & Hitch 1974, Goldman-Rakic 1995). While storage in long-term memory occurs through long-lasting changes in synaptic connections between neurones, information in WM is kept by persistent neuronal activity (Goldman-Rakic 1995). A defining feature of WM is that only a limited amount of information can be held and only for short intervals. These limitations may constitute one of the primary bottlenecks of human cognition (Marois & Ivanoff 2005), which could be one of the reasons for the observed link between WM performance and a range of cognitive skills, including general intelligence (Kane & Engle 2002, Conway et al. 2003).

The most common way to measure WM is to determine the WM capacity (WMC), the maximum number of items a person can hold in WM, which is usually around 4-7 items (Miller 1956, Luck & Vogel 1997, Cowan 2000) depending on the characteristics of the task. By measuring WMC, Westerberg et al. (2004) found unmistakable connections between ADHD and poor WM. The neurophysiological mechanisms behind these findings have not yet been explained in a satisfactory way. However, several studies have found an inverse relationship between distractibility, the degree to which information in WM is affected by the presentation of irrelevant information, and WMC (ref), findings that fit well with the fact that children with ADHD are typically easily distracted.

A neurophysiological mechanism behind the relationship between WMC and distractibility would be of interest as it could help in understanding the role of WMC in ADHD. The covariation between WMC and the ability to resist distractions indicates that they could both depend on activity in a same certain part of the brain. However, at first glance, WMC and the ability to resist distractions do not seem to have much in common. Storage capacity is mainly determined by activity in the parietal lobe, in particular the intraparietal cortex, IPS (Todd & Marois 2004). Filtering distractions on the other hand, is usually thought of as blocking of sensory inputs during the propagation through the sensory areas towards the memory areas. It is difficult to see how the size of the storage buffer in IPS could have an effect on this block.

Another possible explanation is that the brain regions underlying WMC and distractibility share the same modulating system. A strong candidate brain region is the dorsolateral prefrontal cortex (dlPFC), a region that for a long time has been thought to play a crucial role in working memory (Goldman-Rakic, 1987, Curtis & D'Esposito 2003, Conway et al. 2003). The role of the dlPFC seems to be of high level management and strong activity in this region has in several studies been related to high WMC, resistance against distractions and intelligence (Duncan 1995, Kane & Engle 2002, Gray et al. 2003). It has been shown that the dlPFC is activated to maintain memories when the number of items held in WM, approaches the WMC (Leung 2002, Bunge & Klingberg 2000). Furthermore, Klingberg et al. (2005) showed that training of WM, which has been seen to increase brain activity in the dlPFC (Olesen et al. 2004), improved WMC as well as the ability to block out distraction (as measured with the Stroop task) in people with ADHD.

Recent theories of dlPFC function propose that dlPFC activity supports WM storage specific areas in the parietal lobe by directing top-down attention in the form of neural control signals (Curtis & D'Esposito, 2003). Even if the support for the executive nature of dlPFC is strong, the way in which it is able to modulate different operations like WMC and distractibility is not clear. Rules for where attention should be directed could be maintained by active maps of goal relevant attributes (Miller & Cohen 2001). High WMC can be thought of as a result of active attention and hence subject to the same activity as the attention control that regulates the distractibility. Kane and Engle (2003) touched upon the subject when they observed that the performance of high WMC subjects in a WM task was reduced to the same level as that of low WMC subjects when another attention demanding task was performed at the same time. The authors argued that this result indicates that high WMC subjects normally gain their extra capacity by recruiting attentional recourses, but that when attention needs to be redirected to the other task, it cannot be used to boost capacity, so in this way what is the storage capacity of IPS is measured. Looking from the other way, Vogel et al. (2005) saw that WMC predicts the efficiency of attentional control: "Individual differences in memory capacity may not simply reflect variability in available storage space, but may also be strongly constrained by the efficiency with which the available space is allocated." Since Vogel et al. recorded event-related potentials, the resolution of the data is not good enough to make detailed claims about the involved areas, but an explanation in line with the theory presented above would be that the low capacity subjects have a poorly functioning dlPFC, explaining both the low WMC and the inability to filter distractions. This way of linking WMC and filtering rests on extensive reviews (Kane & Engle 2002, Curtis & D'Esposito 2003, Conway et al. 2003), but it need to be said that not all researcher in the area would agree. Finding the mechanism that link WMC and distractibility is relevant and highly interesting as it would further the understanding of intelligence and WMC, and could be a significant piece of explanation in the comprehension of ADHD like symptoms.

In this thesis, I investigate two main questions. First, I present a hypothetical mechanism whereby the dlPFC can dynamically change its activity mode to modulate activity in posterior areas in a task-relevant way. Given this mechanism within the dlPFC, it is unclear to which posterior region it connects to in order to modulate activity. I therefore put forth three experimentally testable hypotheses for the identity of the target area of the dlPFC. All three hypotheses entail a network consisting of the dlPFC, a visual area, V4, and the IPS, in which stimuli are encoded in V4 and passed on to the IPS, where they are stored.

- H1: IPS-filtering. dlPFC connections here target the IPS. The dlPFC boosts capacity through general excitation of the memory area, and filters by sending an excitatory signal that renders the IPS unfit to store distractors
- H2: V4-bias. dlPFC connections slightly excite the neural populations in V4 that encode the relevant information. This excitatory bias serves as an advantage in the competition for access to the limited memory store in IPS. Boosting occurs through general excitation of the visual area which in turn excites the IPS.
- H3: V4-filtering. dlPFC connects to the visual area. Filtering occurs in V4 due to inhibition of the neural populations encoding distractors. Memory boosting is the same as in Visual-bias.

To test the hypotheses (which are visualized in figure 12), computational modelling will be combined with functional magnetic resonance imaging (fMRI). During the last decade, computational models of WM, incorporating results from extensive in-vitro experiments and electrophysiological single cell studies on monkeys, have been able to explain a number of experimentally observed features of WM (Compte, 2006). Models like these have been used to study WMC (Macoveanu et al. 2006, Edin et al. 2007) as well as distractibility (Compte et al. 2000, Macoveanu et al. 2007). This thesis makes a novel approach in modelling both WMC and filtering as an executive top-down effect originating from a dlPFC that can be set in two different modes. Simulated data will later be combined with brain activity measured with fMRI to investigate which hypotheses could account for the observed data. Only the modelling part will be dealt with in the thesis work. However, I will also describe those parts of the experimental setup that are needed to understand the modelling.

1.1 Neurones and the cortex

The cerebral cortex, or the grey matter of the brain, consists of interconnected neurones that transmit neural impulses called action potentials to each other. Neurones are essentially leaky electrical capacitors with a potential that depends on the currents injected into them. When the potential reaches a threshold, a neural impulse is emitted and the potential is reset to a reset potential.

There are two main types of neurones, pyramidal cells and inhibitory interneurones. Pyramidal cells connect to other cells with synapses containing two types of receptors, the NMDA receptor (NMDAR) and the AMPA receptor (AMPAR). Signalling via the AMPAR leads to currents in the recipient cell with a much faster decay time than the NMDAR, and the effect of a train of action potentials over synapses dominated by AMPAR is hence much briefer in time than when the NMDAR is involved. Both receptors are excitatory, meaning that they bring the potential of the recipient cell closer to its threshold. Since pyramidal cells signal via excitatory receptors, they are also called excitatory. Connections from inhibitory interneurones are, on the other hand, mediated by the $GABA_A$ receptors ($GABA_AR$), which is inhibitory (it brings the potential of the recipient cell further from the threshold). The kinetics of the $GABA_AR$ are fast, almost as fast as those of the AMPAR.

The activity of a neurone or a neural network is dependent on numerous properties of the involved cells and in the network case it also depends on the way cells are connected. The dynamics that these properties create can be studied with the frequency-current curve (fI-curve) (figure 1), where the activity, measured in spikes/s is dependent on the current injected into the neurones in the network.

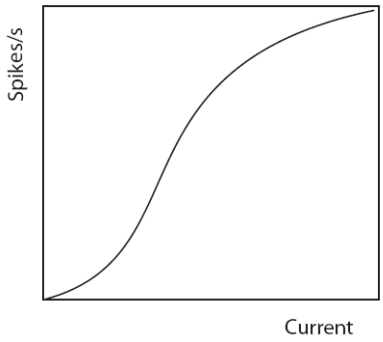

Figure 1 - A fI-curve. The single-cell or network activity is dependent on the current injected into the neurones.

1.2 Visuospatial WM

Visuospatial WM is the WM of the location of objects in visual space, and is the type of WM studied here. Not only does it correlate strong with ADHD (Westerberg et al. 2004), but it is also one of the simplest and best explored types of WM. Neurones that respond to nearby stimuli are more strongly connected and closer situated on the cortex than neurones that respond to locations further away in the visual field. This property creates what is called a retinotopic map. Retinotopic maps have been found in both visual areas and parietal as well as frontal memory areas.

Visual cues that are transmitted from the retina to the brain are first processed by the different regions of the visual cortex. The representation in the earliest regions is a two-dimensional image of the visual field, almost like a pixel based screen. Inputs are then decoded by increasingly specified areas as they are passed on to other areas along sequence of visual areas called the visual stream. Different objects and their spatial location are separated from the surrounding and processed by regions specialised for their features. As visual information proceeds along the visual stream, the effect of attention related top-down signals, which enhance the representation of the attended information, becomes more pronounced (Kastner & Ungerleider 2001). After processing in the visual cortex, inputs are passed on to WM areas in the parietal and temporal lobe where they can be held and used for more complex operations and/or passed to long term memory. It is typically thought that visuospatial information is passed to the IPS in the parietal lobe (the dorsal pathway) and object information to the temporal lobe (the ventral pathway).

1.3 Persistent activity and the bump model

Activity in the visual cortex is a direct response to visual cues, and vanishes as soon as the cues are removed. In contrast, networks in the WM regions have been shown to display stimulus-specific persistent activity, which is the prevailing theory behind WM (Goldman-Rakic 1987). Persistent activity is the ability of a group of neurones that are sensitive to the same stimulus to maintain a high firing frequency. When a group of neurones fire, the activity is fed back to the same group so that they again activate each other in a process called recurrent excitation. Since the neurones are feeding each other with excitatory currents, the activity remains even after the stimulus that first started the activity is removed, hence a memory is held. Persistent activity related to WM has been found among neurones in the parietal, temporal and prefrontal lobe.

In order to hold a memory, persistent activity needs to be concentrated in this way to only those cells encoding the stimulus, otherwise it would be impossible to identify the stimulus that gave rise to memory activity by observing the activity patterns in the network. A memory held in this way is often referred to as a “memory bump”, because the firing pattern of the involved neurones looks like a bump compared to the low frequency firing of the surrounding neurones, see figure 7. A network that can show persistent activity has non-linear properties that can be determined by the shape of the fI-curve. Two important aspects of a network that contribute to the properties of the fI-curve, and hence the behaviour of the memory bump, are the connection curve and the interplay between excitatory and inhibitory neurones. The connection curve describes how strongly neurones connect to other neurones depending on the difference in their preferred stimuli. A wide and flat connection curve, implying that neurones do not discriminate strongly between the difference in input signal, is likely to induce a relatively wide bump, whereas a networks where neurones connects much stronger to close neighbours than distant, is likely to result in a narrow bump. However, it is also important that the network has strong enough inhibition so that neurones on the side of the bump are suppressed. Without this component, bumps in networks with even the very narrow curve will grow wider and wider until they fill the whole network and the memory is lost.

2 Methods

Since the modelling study presented in this thesis is thought to be combined with a fMRI experiment conducted by other members in the developmental cognitive neuroscience group at Karolinska Institutet, the model has been built in accordance to that experiment.

2.1 fMRI experiment

Measurements of brain activity with functional magnetic resonance imaging (fMRI) and measurements of WMC were performed in 32 subjects while they performed a WM task inside a magnetic resonance scanner. There were two versions of the task. In the distraction task, subjects were instructed to remember the position of red dots and ignore yellow dots. In the no distraction task, subjects needed to remember the location of all dots. The task had five phases (figure 2). During an initial fixation phase, subjects were required to look at the cross. Following this was a task instruction phase where a symbol told the subject which of the tasks to prepare for. A triangle was the symbol for the distraction task and a square the symbol for the no distraction task. The instruction phase was followed by the memory stimuli phase, where visual cues in the form of dots were presented in a circle. Three red and two yellow dots were displayed in the distraction task, whereas three red dots or three red dots and two yellow dots were presented with equal probability in the no distraction task. During the subsequent delay phase, subjects held the location of the dots in memory. Lastly, during the response phase, a probe in the form of a question mark was presented at a location, and the subject had to respond whether or not there had been a cue presented there. The length of some of the phases was varied between trials for statistical post processing reasons. The cues were presented at an eccentricity of approximately 6° from the fixation point when subjects lied in the fMRI scanner.

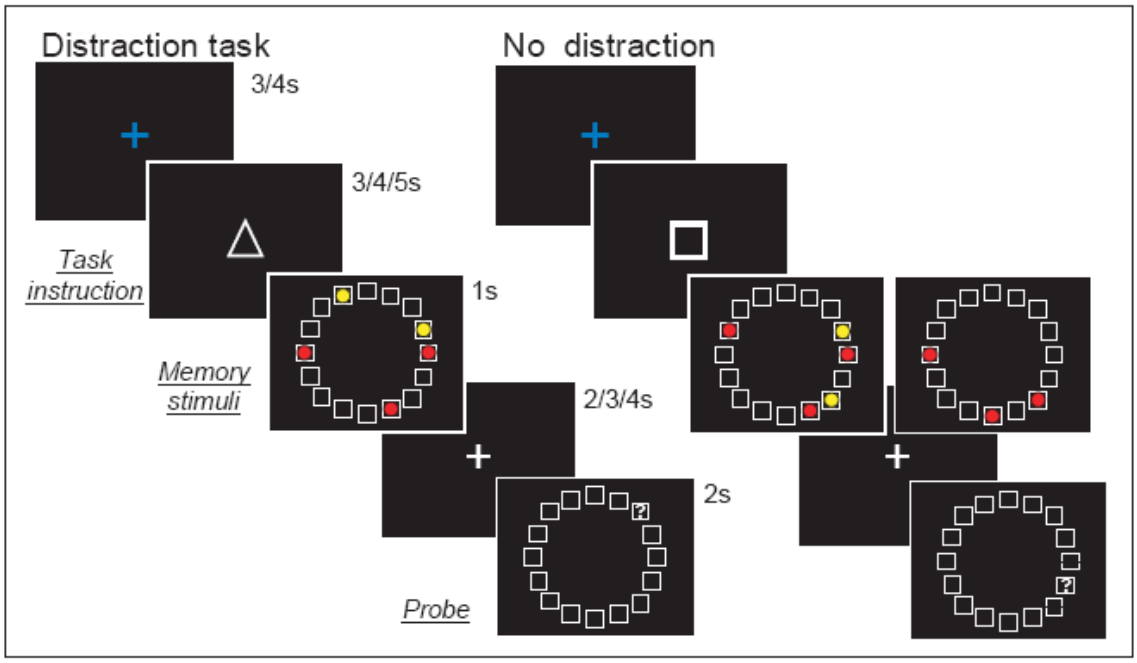

Figure 2 - The experiment consisted of a distraction and a no distraction task, which were distinguished between by a symbol during task instruction. In the distraction task, only red cues were to be memory cues, whereas all cues were to be remembered in the no distraction task. After a delay, a probe marked a location in the circle and subjects had to respond whether a memory cue had been presented there or not.

2.2 The model

2.2.1 Overall structure

Three brain regions are of particular importance for the execution of the WM task, a visual region, V4, the memory region in the IPS and the dlPFC. Visual stimuli enter V4 and are passed on via excitatory connections to the IPS, where they are stored. The dlPFC allocates attentional resources by sending excitatory signals to the other two regions in order to control the access to and capacity of WM. The way the dlPFC connects to the other two regions is different for each hypothesis and is described below. The visual cortex, being a target for top-down signals (Desimone 1998), could be a region with which the dlPFC connects to exercise top-down control, and the idea was to include the last visual area that cues are thought to pass before entering the IPS, since it has been found that top-down modulations are most pronounced there (Kastner & Ungerleider 2001). V4 was assumed to be the main visual area involved in this task, since it encodes both colour and locations and mediates spatial information to parietal areas (Gatass et al. 2005). It should be noted that it is usually thought that spatial information is passed to the parietal areas via dorsal visual areas like V6, but no colour selectivity has been reported for this area and it has been described to handle location of high eccentricities (~30˚and above, whereas the cues in the experiment are presented at 6°) (Gatass et al. 2005). Other regions are also activated (such as early visual regions), but they are not central to the execution of the task and need not be included in the model. Likewise, although the modelled regions are heavily interconnected to other regions in the brain, only those connections that are of presumed importance for the execution of the task are included in the model.

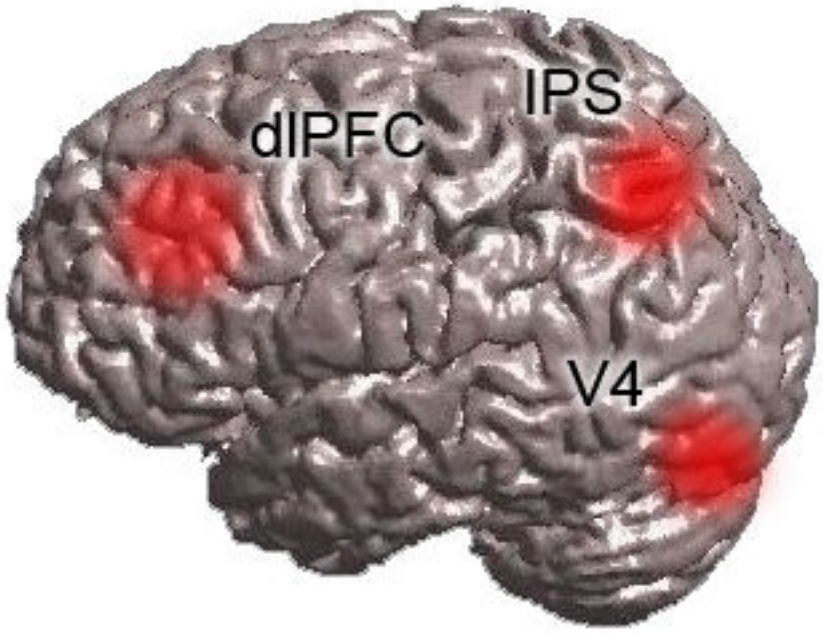

Figure 3 - Approximate locations for the regions of interest.

2.2.2 Basic structure of each local network

Each brain region in the model is built up of two local cortical networks coding for the red and yellow dots, respectively. Although the basic architecture of each network is identical, networks in the IPS and the dlPFC are parametrised to have stable persistent activity, whereas V4 lacks this ability. Each local network consists of one population of excitatory pyramidal neurones and one population of inhibitory interneurones. Both cell types are modelled as leaky integrate-and-fire neurones as in Ardid et al. (2007) and Compte et al. (2000). Cellular membrane potentials are described by the equation

$$\begin{cases} C_m \dfrac{\mathrm{d}V}{\mathrm{d}t} = I_{\mathrm{AMPAR}} + I_{\mathrm{NMDAR}} + I_{\mathrm{GABA_AR}} + g_{\mathrm{L}}(V - E_{\mathrm{L}}), & V < V_{th} \\ V = V_{reset}, & V \geq V_{th} \end{cases} \quad (1)$$

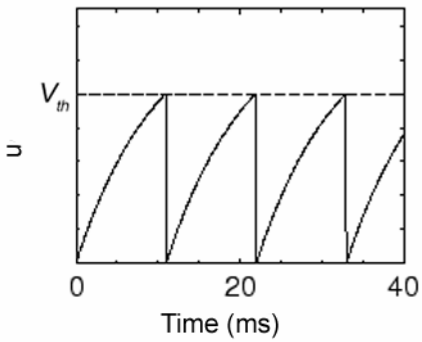

Figure 4 - The potential of an integrate-and-fire neurone (1) driven by constant input current I0 = 1.5. The voltage u(t) increases up to the threshold potential Vth, where it is said the neurone spikes. The potential is then set to the value of the reset potential Vreset.

There is also a refractory period τ_{ref}, which sets the time the potential is held at the reset potential after an action potential. Here follows an excerpt from Ardid et al. (2007)

describing the inner design of a network. For a more detailed account of the network properties, see Compte et al. 2000.

> [E]ach type of cell is characterized by six intrinsic parameters: the total capacitance C_m, the total leak conductance g_L, the leak reversal potential E_L, the threshold potential V_{th}, the reset potential V_{res}, and the refractory time τ_{ref} . The values used are C_m = 0.5 nF, g_L = 25 nS, E_L = −70 mV, V_{th} = −50 mV, V_{res} = −60 mV, and τ_{ref} = 2 ms for pyramidal cells; and C_m = 0.2 nF, g_L = 20 16 nS, E_L = −70 mV, V_{th} = −50 mV, V_{res} = −60 mV, and τ_{ref} = 1 ms for interneurons. All cells receive random background excitatory inputs. This overall external input is modelled as uncorrelated Poisson spike trains to each neuron at a rate of V_{ext} = 1800 Hz per cell (or equivalently, 1000 presynaptic Poisson spike trains at 1.8 Hz). The external input is exclusively mediated by AMPARs…
>
> Neurones receive their recurrent excitatory inputs through AMPAR and NMDAR mediated transmission and their inhibitory inputs through $GABA_A$Rs. These conductance-based synaptic responses are calibrated by the experimentally measured dynamics of synaptic currents. Thus, postsynaptic currents are modelled according to $I_{syn} = g_{syn}s(V - V_{syn})$ [where *syn* is AMPAR, $GABA_A$R or NMDAR], where g_{syn} is a synaptic conductance, *s* a synaptic gating variable, and V_{syn} the synaptic reversal potential (V_{syn} = 0 for excitatory synapses, V_{syn} = −70 mV for inhibitory synapses). AMPAR and $GABA_A$R synaptic gating variables are modelled as an instantaneous jump of magnitude 1 when a spike occurs in the presynaptic neuron followed by an exponential decay with time constant 2 ms for AMPA and 10 ms for $GABA_A$. The NMDA conductance is voltage dependent, with g_{syn} multiplied by $1/(1 + [Mg^{2+}] \exp(-0.062V_m)/3.57)$, $[Mg^{2+}]$ = 1.0 mM. The channel kinetics [for NMDA] is modelled by the following equations:

$$\frac{ds}{dt} = -\left(\frac{1}{\tau_s}\right)s + \alpha_s x(1-s) \qquad (2) \qquad\qquad \frac{dx}{dt} = -\left(\frac{1}{\tau_x}\right)x + \sum_i \delta(t - t_i) \qquad (3)$$

> where s is the gating variable, x is a synaptic variable proportional to the neurotransmitter concentration in the synapse, t_i are the presynaptic spike times, τ_s =100 ms is the decay time of NMDA currents, τ_x =2 ms controls the rise time of NMDAR channels, and α_s=0.5 kHz controls the saturation properties of NMDAR channels at high presynaptic firing frequencies. Parameters for synaptic transmission are taken from Compte et al. (2000). (Ardid et al., 2007).

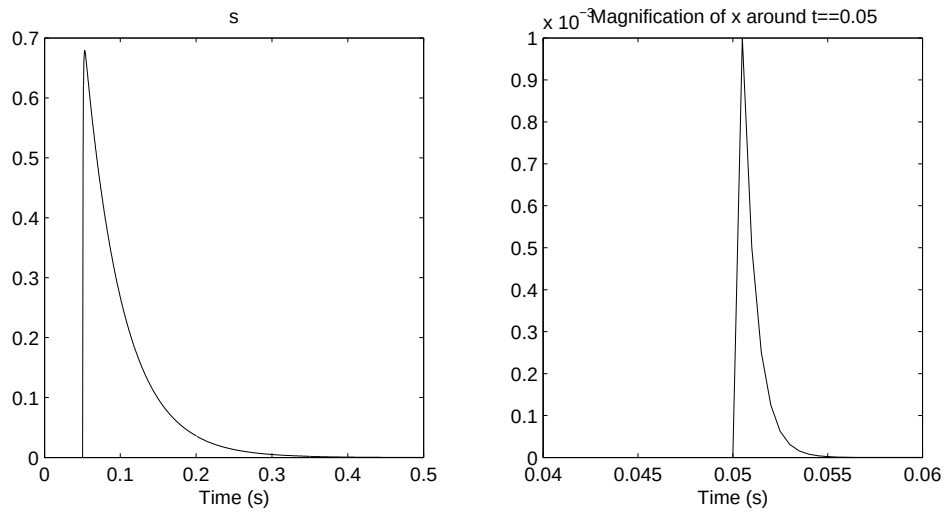

Figure 5 - The gating variable s (2), the ratio of open ion channels, for NMDA in response to a presynaptic spike. i.e. a neighbouring neurone spikes at t = 0.05, most

channels are opened instantly, but the proportion open decays with time. s depends on the synaptic variable x plotted in the right graph.

For each circuit, pyramidal cells (N_E = 4096) and interneurons (N_I =1024) are spatially distributed in a ring according to preferred cues ($0° \leq \theta_{pref} < 360°$). Connections between neurones can be structured or unstructured. Neurones connect stronger to neurones with similar cue preference than neurones with disparate if the connection are structured, and they connect equally strong to all neurones when connections are unstructured. The architecture is unchanged from that in Ardid et al. (2007), although synaptic strengths differ. The strength of recurrent connections between neuronse in the network depends on the difference in preferred cues. This is implemented by taking the synaptic conductance between neurone *i* and neurone *j* to be $g_{syn,ij} = W(\theta_i - \theta_j)G_{syn}$, where $W(\theta_i - \theta_j)$ is either a constant for unstructured connections ($W(\theta_i - \theta_j) = 1$), or the sum of a constant term plus a Gaussian function for structured connections

.

$$W(\theta_i - \theta_j) = J^- + (J^+ - J^-)\mathrm{e}^{-(\theta_i - \theta_j)^2 / 2\sigma^2} \qquad (4)$$

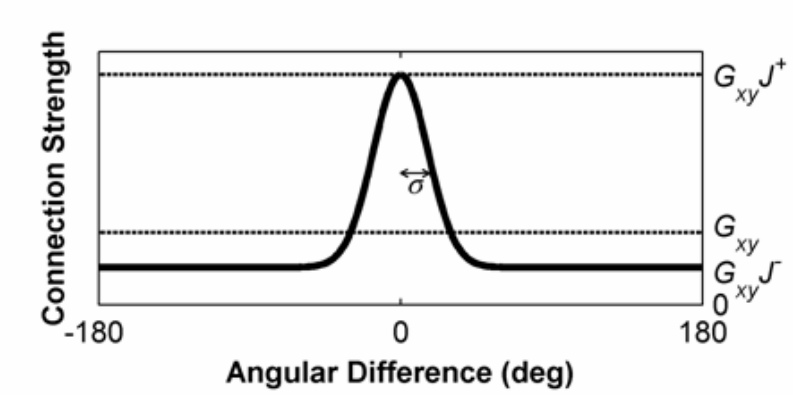

Figure 6 - Structured connection (4). Connection strength between two neurones depends on their difference in preferred cue location. Picture from Edin et al. (2007).

2.2.3 Specific tuning of the local networks in each region

Appendix A holds the values for the synaptic connectivity within the networks. The IPS and the dlPFC modules have strong recurrent excitatory connections to sustain persistent activity. The networks in V4 and IPS have structured connectivity in all but inhibitory-to-inhibitory connections. Finding appropriate values for all parameters was, because of the vast parameters space of the model, a time consuming work subject to a number of requirements. The main issue was to create an IPS module with at least a capacity of holding three memory bumps. With a trail-and-error based methodology and with guidance from Macoveanu et al. (2006), the width of the mean strength of local excitatory-to-excitatory connections ($G_{E\to E}$) was increased and the connection curve was made narrower and more peaked by decreasing σ and increasing J^+. Structure connection was also introduced for the excitatory-to-inhibitory ($G_{E\to I}$) and the inhibitory-to-excitatory connections ($G_{I\to E}$) to maintain a fixed width of the bumps regardless of the number. The network was required to have a resting state activity around 1 Hz and a mean activity in memory bumps preferably below 50 Hz to fit experimental data (Ardid et al. 2007). A functional memory further requires that bumps are relatively stable in position. As will be explained more thoroughly later on, a model network tuned to the limit of capacity will be increasingly likely to burst out into spontaneous bumps. This is considered as a highly undesirable behaviour with no

support in experimental data from monkeys or humans. The ability to maintain a stable resting state was therefore also an important requirement when capacity evaluations of the modelled network were made.

The V4 module was modified to achieve sensory bumps of approximately the same width as in the IPS. Reasonable firing rates for this area, being a late visual region, was set to 10 Hz at rest and somewhat below 100 Hz in the first bumps (Desimone & Schein1987). A distinction from earlier visual areas is the relatively strong recurrent inhibition (via E→I and I→E connections) that reduces the firing frequencies in the bumps when the number of bumps grows.

The modelled dlPFC area was tuned to encode one of two rules, capacity increase or filtering. In neither case is it required to retain spatial information, and hence the dlPFC was modelled with only unstructured connections, i.e. all neurones target all other neurones in the network equally strongly. This way, each network only has an on and an off state. Firing rates here were a few Hz at rest and around 20 Hz during persistent activity, similar to those in experiments (Funahashi et al. 1989).

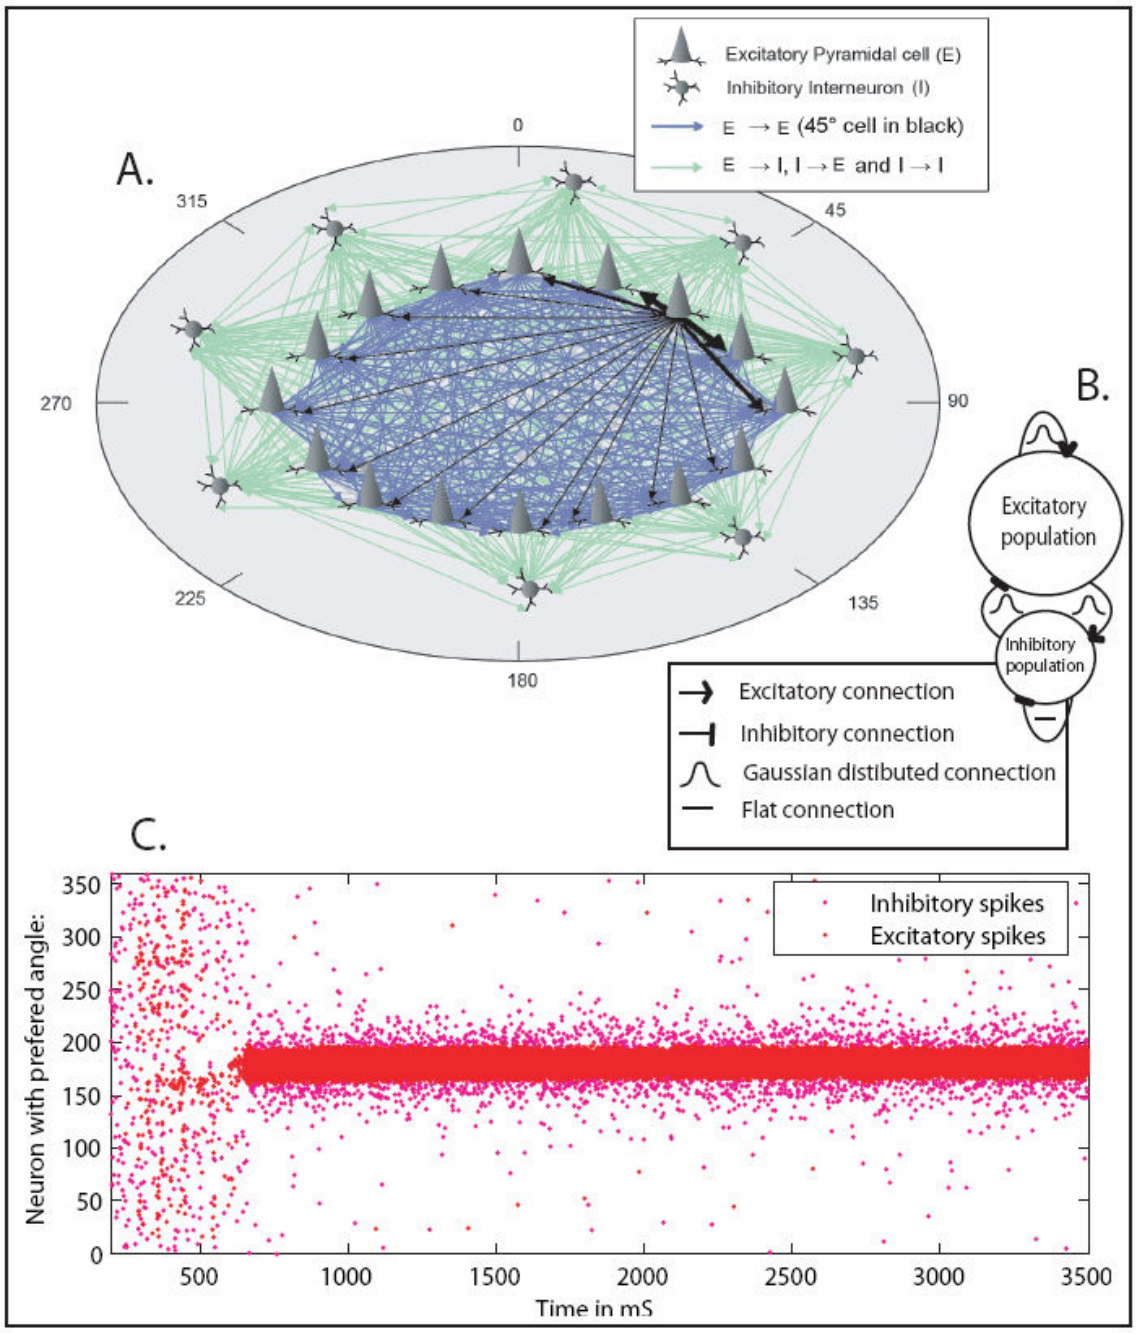

Figure 7 - Structure and example simulation of a model network. A) Connections within a network. All cells connect to all other cells. Connections can be structured or flat, where a structured connection implies that a cell connects more strongly to cells with similar cue preference than to other cells. B) Schematic way to show the connections within a network. C) Example of a memory bump. During the first 500 ms the network is in a resting state and fires spontaneously at low rate. When a cue is presented (modelled as a current) to for neurones with preferred cue angle around 180˚, these neurones starts to fire to each other at much higher rates, thus creating a memory bump. Neurones outside the side of the bump are suppressed by increased inhibition.

2.2.4 Connections of local networks into model regions

Memory dots and distracters are distinguished only by their colour, red or yellow. The representation of colours is spatially separated into local subnetworks distributed across the cortex in area V4 (Bartels & Zeki 2000). Therefore, V4 was modelled as two

spatially separated and only weakly connected local networks (connection strengths between colour networks are approximately 5% of those within a colour network, except in the V4-filtering hypotheses where they were stronger). Only two colours were modelled, since only two colours appeared in the fMRI experiment. The connections between the colour networks were modelled the same way as the internal connections of a network and are structured for all but the inhibitory to inhibitory connections.

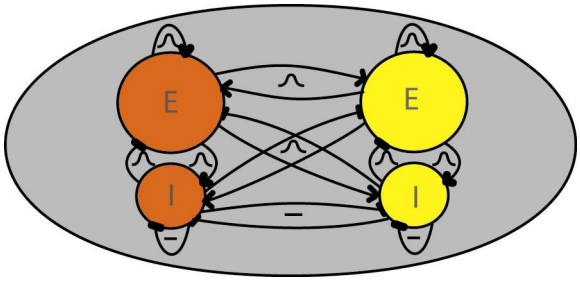

Figure 8 - One area with two colour-networks.

Colour selectivity has also been found in the parietal cortex, especially if the colour of cues are goal relevant (Toth & Assad 2002), even though this property is far less documented than for the visual cortex. There is no reason, however, to believe that it would be easier to store more locations if the dots are of different colours than if they have the same colour. This implies that inhibition between the colour-networks must be as strong as within each network. After one stimulus has been encoded in one of the networks, it is then just as difficult to store a second stimulus in either of the networks.

The dlPFC networks do not need to be thought of as explicitly colour coding, but rather as rule encoding. Being a region that encodes rules dynamically, connectivity in this region is presumed to be easily changeable (Miller & Cohen 2001). The two dlPFC networks in the model are unconnected to each other since there are no sources good enough to even know whether such a connection should be modelled as excitatory or inhibitory.

2.2.5 Connections between regions

Model brain regions are interconnected through bottom-up and top-down excitatory AMPA-mediated connections. Inter-regional connections were modelled just like the internal connections of a network. Connections from V4 to the IPS are structured, since they transmit spatial information. Connections going the opposite way, from the parietal to the visual area were omitted since they where not thought to play a significant role in the process of interest. The same was true for connections from these areas to the dlPFC. The connections from the dlPFC networks to the networks in the visual and parietal areas were unstructured since the dlPFC in the model does not hold any spatial information.

2.3 Implementing the hypotheses

The main purpose of the present study was to propose a mechanism whereby the dlPFC can dynamically switch the memory network into a capacity-enhancing or a filtering mode. To achieve this, networks in the dlPFC were connected to networks in the posterior regions representing the same colour. The dlPFC was tuned to be able to activate one or two of its networks for filtering or capacity enhancement, respectively.

By activating a single dlPFC network representing one colour, that network would send excitatory signals that would enhance activity of neurones in posterior areas coding for dots having that colour. On the other hand, by activating both regions, both networks receive excitation and neither colour has an advantage over the other. Instead, the general excitation enhances capacity by increasing external input.

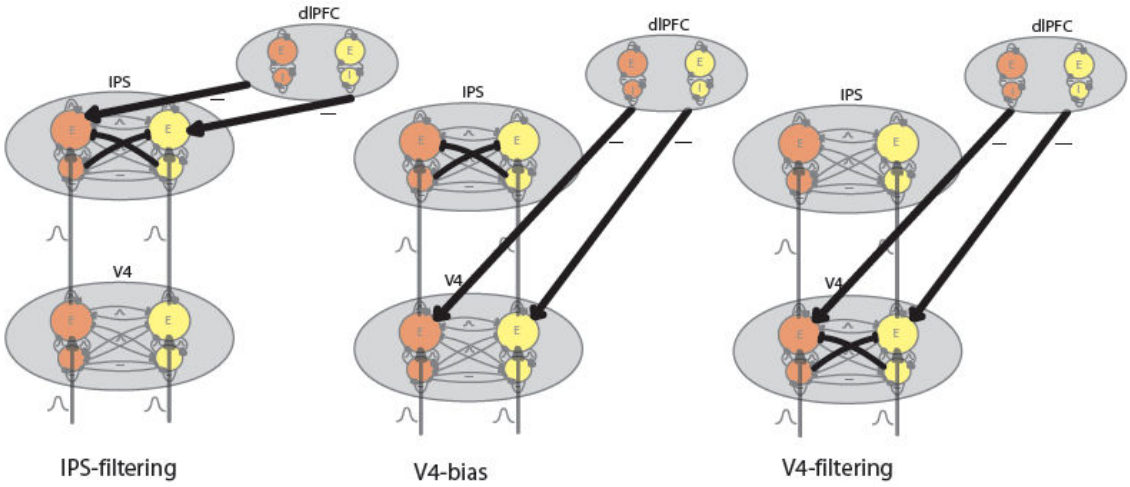

Figure 9 - The three hypotheses. Bold arrows indicate where the top-down signal targets and where the filtering is thought to take place. The Gaussian curves next to the arrows indicate structured connections, and lines indicate flat connections. IPS-filtering: Top-down signals from the dlPFC target IPS, where filtering is also carried out. V4-filtering: dlPFC targets V4, filtering is still carried out in IPS. V4-filtering: dlPFC targets V4, that also performs the filtering due to stronger inhibitory connections than the V4 areas in the former hypotheses.

Three hypotheses regarding the targets for the dlPFC activity has been put up (figure 12). In the IPS-filtering hypothesis, filtering and capacity enhancement take place in the IPS. In the visual-bias hypotheses, the biasing signal from dlPFC targets V4, but the connections between the colour-networks here are not strong enough to carry out the biased competition. The extra top-down activity is transmitted together with the stimuli to IPS, where competition or capacity enhancement occurs as in the first hypothesis. The idea for the visual filtering hypothesis is that connections are strong enough between the colour-networks for biased competition to filter distracting stimuli already at visual level. If both dlPFC networks are turned on for capacity boosting, enhanced activity in V4 should propagate to IPS and there cause the capacity increasing effect as in the V4-bias hypothesis.

2.4 Modelling Protocol

To be make prediction regarding the results from the fMRI experiment, the model needed to feature a similar protocol as in that experiment so that simulation and fMRI data are comparable.

2.4.1 Inputs

Stimulus dots where presented in a ring around a fixation point. Spatial locations for the dots enter area V4 via earlier visual areas and are here modelled by injecting external current into excitatory neurones in V4. For a single dot with angle θ_S, the current injected into a V4 neurone coding for angle θ_i is $I(\theta_i) = I_0 + I_1 \exp(\mu(\cos(\theta_i - \theta_S) - 1))$,

where I_0 = 0.4 nA and I_1 = 1.77 nA for pyramidal cells and I_0 = 0.2 nA and I_1 = 0.18 nA for interneurones. For both cell types, μ = 39 π/rad (this choice of μ gives a connectivity profile very close to a Gaussian function with a constant baseline and with the same width as V4-to-IPS connections). When two or more stimuli are visually presented, the current to V4 neurones "is the sum of the currents corresponding to the two single stimuli, normalized so that the maximal current is still $I_0 + I_1$. This normalization is derived from the observation that the maximal response of a direction-selective V1 neurone remains the same for either single motion or transparent motion stimuli (Snowden et al. 1991)" (Compte et al. 2007).

The neural interaction that leads to the onset of rule representation in dlPFC is so far unknown and it would be beyond the scope of this study to make any predictions thereof. The two colour networks of dlPFC were turned on by injecting a current I_0 = 0.058 nA into all excitatory neurones.

2.4.2 Simulations

Three different kinds of simulations were made in order to test the rule encoding role of the dlPFC; simulations with both dlPFC networks off, with one dlPFC network on and with both dlPFC networks on. The simulations included four of the five phases in the fMRI protocol, but with shorter times. Fixation phase: 0 – 100 ms, instruction phase. 100 – 600 ms, during which currents were applied to dlPFC in those simulations that required so. Cue phase: 600 – 1600 ms, during which currents were injected into the V4 network corresponding to two red and one yellow dot. 1600 – 3500 ms, delay phase. There was no need for a response phase since the number of stored memories can easily be read out at the end of the delay period (see figure 10).

The stability of the memory networks turned out to be sensitive to the number of neurones used in the simulation. If too few were used the spontaneous state behaved unstably. To accommodate for this, the networks used to evaluate the hypotheses were simulated with 4096 neurones in each colour-network of V4 and IPS, and 2048 in the networks of dlPFC, 20480 neurones in total. However, the network used to produce figure 11 only included 512. The reason for the unstable behaviour is that the IPS network turned to high capacity by the dlPFC (figure 10) is very close to the right bifurcation point (figure 12). As the numbers of neurones is increased the input to each neurone is at the same level, but since it receives the input from more sources, the variability is reduced. It is then possible for the network to be closer to the unstable state without being pushed over the limit by random fluctuations. Tens to hundreds of thousands of neurones is a more realistic number, but such a large number of neurones is impossible to simulate.

The code for the simulations was written in *C++*, and the equations where integrated using a second-order Runge-Kutta algorithm with a time step of Δt = 0.02 ms.

3 Results

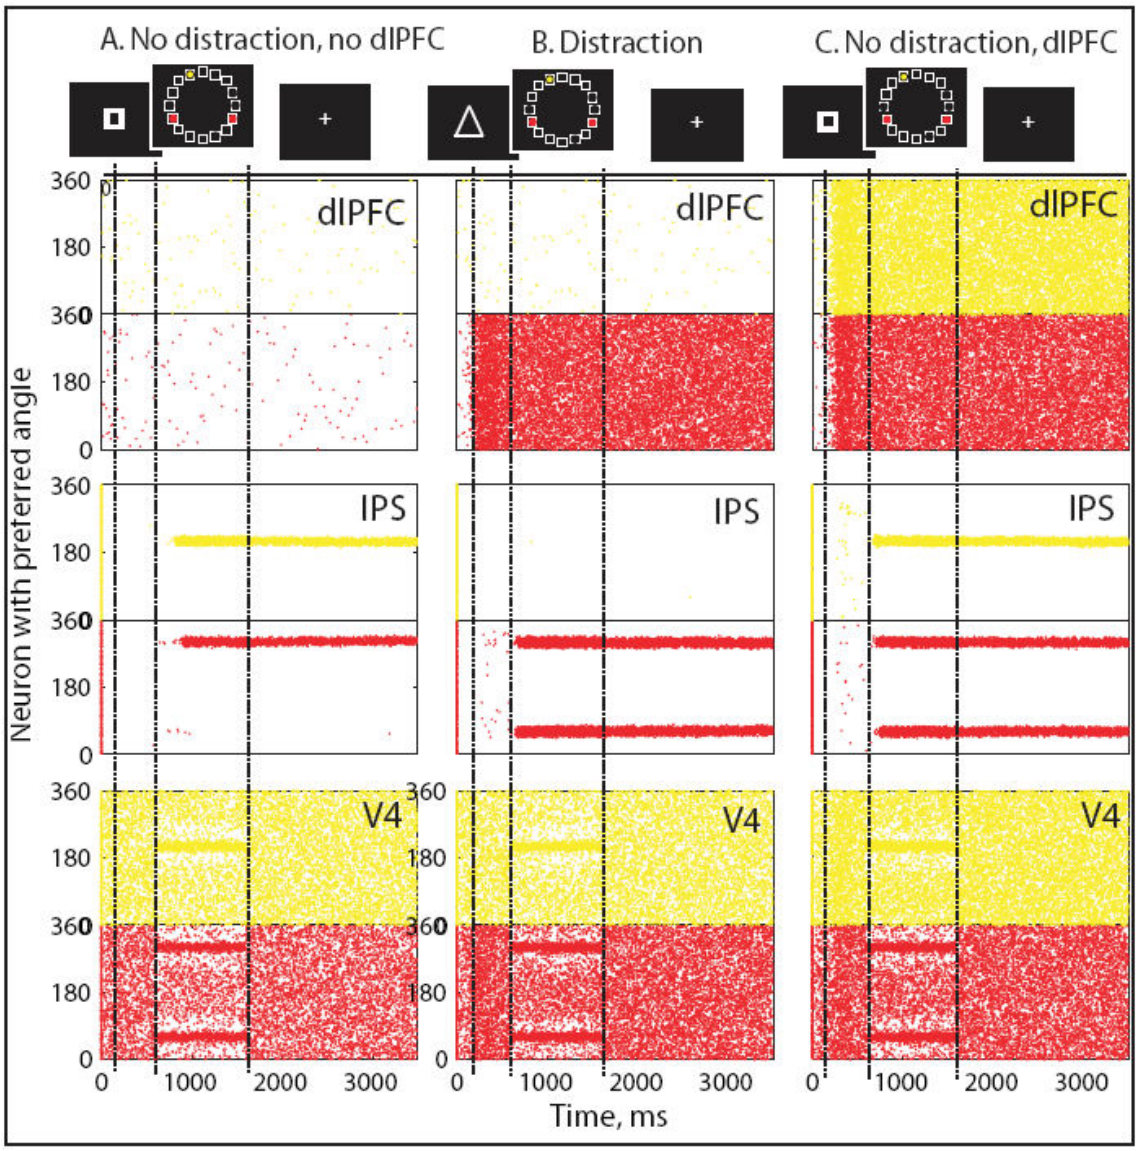

Figure 10 - Example simulations of the WM task for the visual bias hypothesis. The figures at the top correspond to what subjects where shown in the fMRI experiment. Cells are aligned according to their preferred locations. Each dot represents an action potential. A) The no distraction task performed with a spontaneously firing dlPFC, only two of three stimuli got stored. Reference simulation to compare the others with. B) The distraction task. The network successfully filtered the yellow distractor when the red dlPFC network was turned on. C) The network stored all three stimuli in the no distraction task when both colour networks of dlPFC were turned on.

Figure 10 shows example rastergrams for the V4-bias hypothesis in accordance with the experiment protocol. To recall, the assumption is that the dlPFC is responsible for filtering as well as capacity support at high loads. First consider the reference network figure 10.A which shows a no distraction task with a silent dlPFC. During the first 100 ms all areas fire at spontaneous rates. No stimulation was given to the dlPFC in the instruction phase in this simulation and the spontaneous firing continues. Stimulus currents to V4 are turned on after 600 ms and drive activity in V4 for 1 s, after which they are turned off. A few hundred ms after the onset, the IPS networks responds to the

stimuli from V4 and two of the cues get encoded. This is the capacity limit of the network, and the third cue does only create a short and weak response. In this example one cues of each colour was encoded, but this differs from time to time. The encoded cues remain in IPS throughout the delay period. Thus, without the dlPFC activated, the subject could only remember 2 stimuli, 1 red and 1 yellow.

The simulation presented in figure 10.B shows the distraction task. During the instruction phase, the presentation of the instruction to filter leads to an input current into the red network in the dlPFC that triggers persistent activity of the whole network, whereas the yellow network remains at spontaneous state. Activity of the red dlPFC causes a slight increase in activity in the red V4, which in turn leads to an increase in the activity of the red IPS network and a slight decrease in the yellow IPS network. When the visual cues enter V4, only the red cues are encoded in the IPS. Successful filtering has been accomplished.

Figure 10.C is again a no distraction task, but this time, both dlPFC networks are activated during the instruction phase. Persistent elevated firing in the dlPFC increases firing rates in the other regions of the model. All cues get encoded and remain throughout the delay period. The signal from the dlPFC has increased the capacity in IPS regardless of stimulus colour.

As shown by the example simulation, the dlPFC was able devote attentional resources to either filtering out distracters or increasing memory capacity by activating differently depending on task demands.

3.1 Modulating capacity

As was shown above, the model is capable of dynamically increasing the capacity of the IPS network by activating the dlPFC. The increased activity in dlPFC leads to increased excitation and higher firing rate of IPS neurones, but it is not clear exactly how increased firing rate results in higher capacity. Simulations were made with a single local IPS network to see how they are linked. Figure 11.A shows how capacity and rate covary as the conductance of the external driving of the excitatory population, $g_{X\rightarrow E}$ is varied. Increasing the excitation increases the firing rate. By so doing, it was possible to regulate the network from a capacity of 0 up to 3, which was the maximal capacity tested. Many parameters, several of which were mentioned in the Methods section, interact in setting the upper limit. The major limiting factor is the gradual loss of stability of the resting state which occurs at capacities of around 2.5. It should be noted however that the risk of spontaneous bumps, and so the actual capacity, to a certain degree is a result of input current variability, which can be greatly reduced by increasing the number of neurones in the simulation. A network with a larger amount of neurones will handle small deviations better and be more stable. Normally, however, it is not feasible to run simulations with a physiological number of neurones since they take too much time.

Figure 11.B is included to show that the presented relationship can also be found by varying parameters within the IPS network that control the firing rate. The curves in 8.B were obtained by varying the leak potential, V_L, of the excitatory population.

Having seen how capacity in the IPS can be regulated by varying the degree of excitatory input, it is easy to imagine how another area like the dlPFC (figure 11.C) can

control the capacity of IPS by regulating the degree of excitation of the IPS. This situation was implemented in 8.D., where part of the external excitation to the IPS was modelled as a flat connection from the dlPFC. The other area needs to have persistent activity in order to maintain the raised capacity for longer times. The connection from the dlPFC to the IPS was tuned so that the memory area performs at capacity 2 when the other area is in resting state, but is raised to capacity 3 when it is turned on.

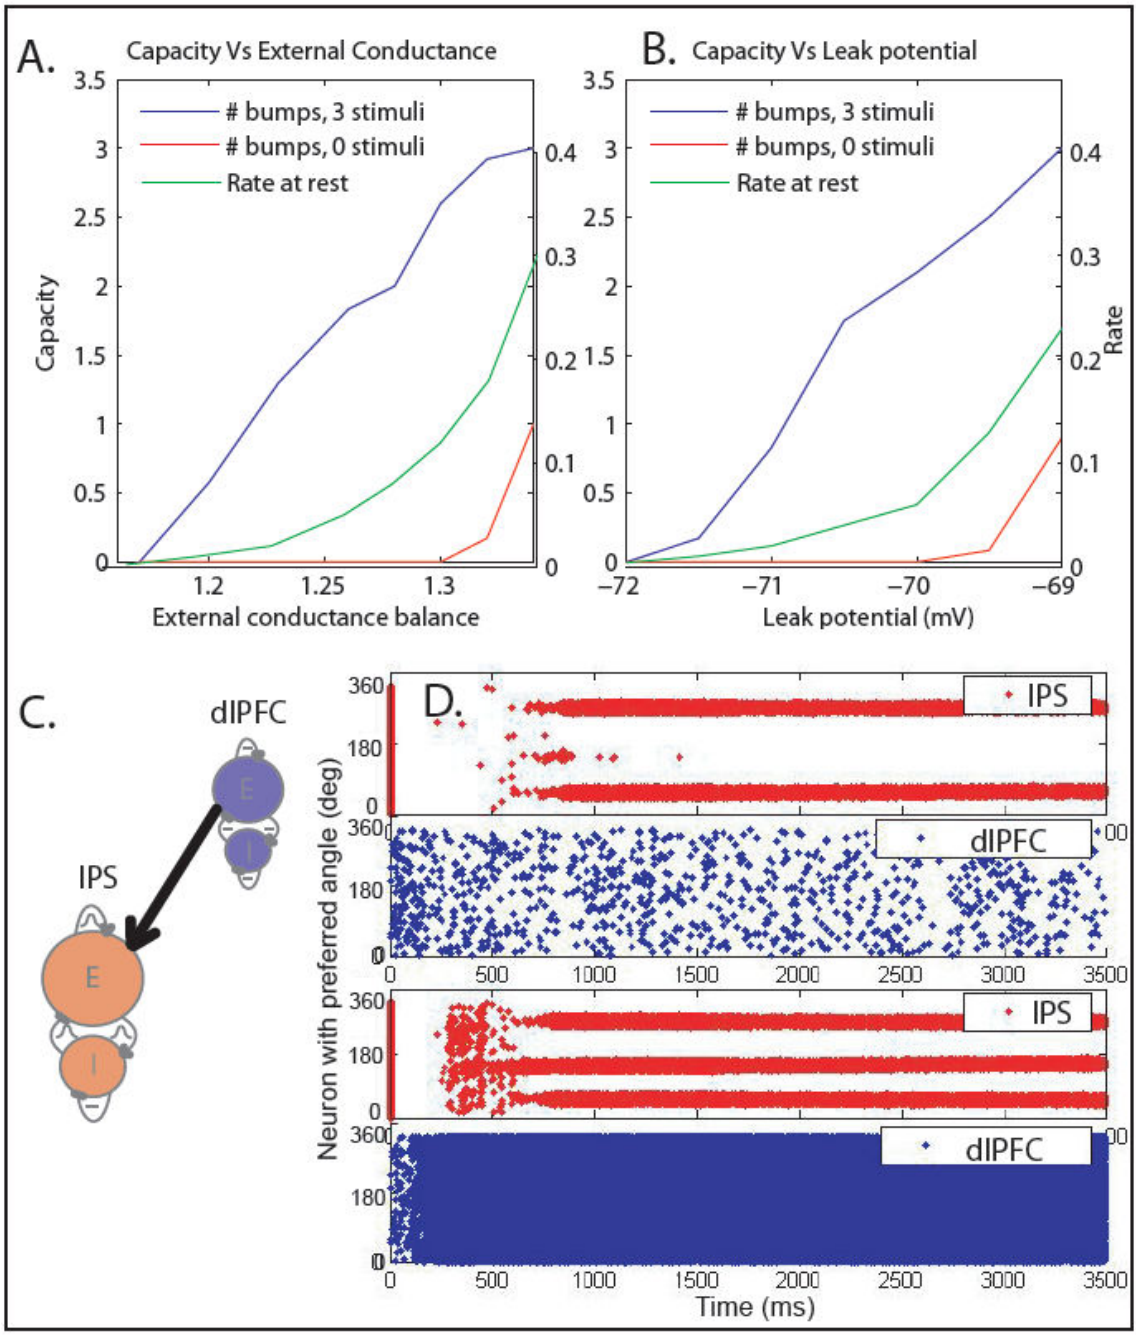

Figure 11 - Capacity is increased by increasing the firing rate of the memory region. A) The network is sensitive to the external inputs to the excitatory population. Increasing the conductance for the excitatory population increases firing rate (green line) and hence the capacity (blue line). Increasing it too much will destabilise the resting state by increasing the risk of spontaneous bumps (red line). B) This effect is not exclusively found by varying the external input. Graph B shows similar curves obtained by varying an internal excitation controlling parameter for the excitatory population, the leak

potential. C) External excitation can be controlled dynamically by inputs from another area. D) The top panel shows the response of the IPS to three memory stimuli when the dlPFC is firing at a spontaneous rate. When the dlPFC is activated as in the lower panel, all three stimuli can be stored. See figures 6 and 7 for an explanation of symbols.

3.2 Bifurcation

What does the extra excitation actually do? To understand the dynamics behind the behaviour of the memory network it helps to consulting a bifurcation graph like the schematic figure 12. A bifurcation graph is used to describe how a non-linear network changes state as a parameter is varied. In this case, the bifurcation graph is a type of fI-curve describing the stable firing rate of neurones in a part of the network as a function of the input current it receives from outside the network as well as other parts of the network. When the input current is low, there is only one stable firing rate, the resting state firing rate. As the input current increases, the subnetwork can also enter the memory state with the higher firing rate. Finally, when the input current is high, then the resting state becomes unstable and only the memory state is stable. An example helps to illustrate the graph further. Following the figures in graph 12.A, think of it as the two capacity network of the upper panel in figure 11.D. Position 1 in graph 12.A corresponds to the resting state of the first few hundred milliseconds. When the subnetwork forms a bump in response to a stimulus, the increased rate in these neurones will reach that of position 2. As the neurones in the bump increases their firing rate, the inhibitory neurones respond by activating leading to other neurones of the network receiving increased inhibition, which means that they will be pushed down to position 3. A second stimulus would bring about a second bump. The two bumps can coexist, but because of the inhibition they generate, both bumps will have a firing rate as in 4 and suppress the neurones not participating in a bump to the rate in position 5. Neurones in position 5 that get stimulated by a third stimulus will not be able to form an extra bump since the current will be too weak. The limit where no more bumps can be formed is a bifurcation point and corresponds to the point where the top solid line turns and becomes dashed.

The next part of figure 12 corresponds to the bottom pane of figure 11.D, where the network of the pane above has been boosted to a three capacity network by an external current (red arrow in figure 12.D). If the network starts further to the right in the bifurcation graph, it is possible to fit another bump before reaching the left bifurcation point. One problem that has been mentioned before is that of spontaneous bumps in the resting state. This behaviour corresponds to the right bifurcation point, beyond this point the networks loses its stability. The conclusion is that a given network can only be boosted with the implemented method if it is not already tuned close to the right bifurcation point.

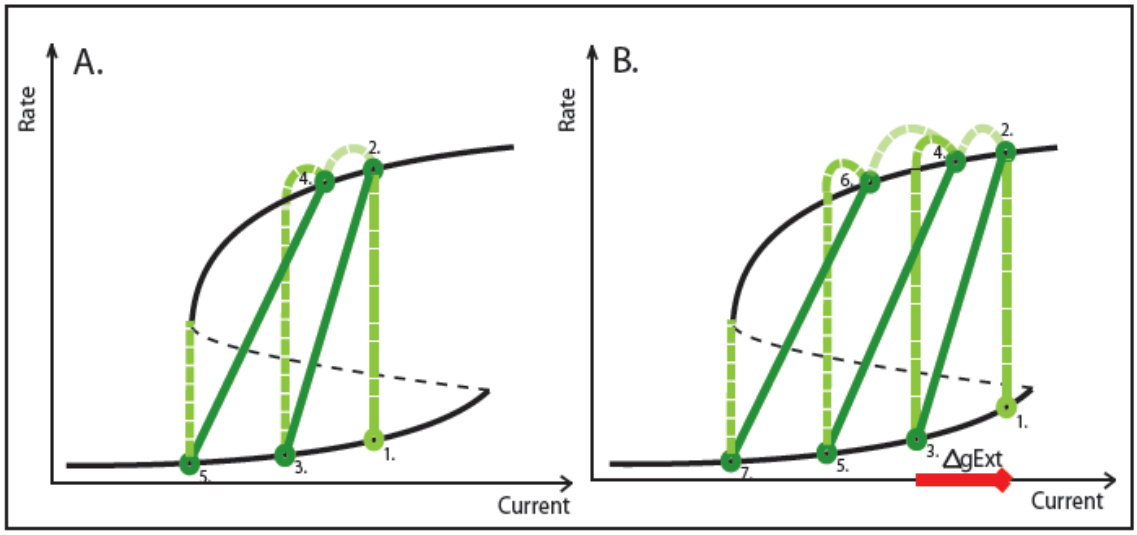

Figure 12 - Bifurcation graphs, example of a capacity two and a capacity three network. A) The lower line represents those neurones that do not participate in a bump and the top line those that do. The network at rest is in position 1. When a bump is formed, the neurones in the bump will be firing at a rate as in 2 and the rest of the network will be in 3. When a second bump is formed, the firing rate it of that one is determined by location straight above 3. However, both bumps will adjust to a mean firing rate, 4. B) The red arrow shows how capacity in this network can be increased by an increase in the external excitation of the excitatory population.

3.2.1 Capacity modulations in the hypotheses networks

Simulations of the IPS-filtering and the V4-bias hypotheses showed that both hypotheses can modulate capacity. Mean capacity is 1.9 and 1.8 respectively when the dlPFC is firing spontaneously. When both networks of the dlPFC were triggered to persistent activity, firing rates increased and mean capacity was raised to 2.9. In the IPS-filtering hypothesis, the dlPFC increases the firing rate in the IPS directly, whereas in the V4-bias hypothesis, the dlPFC increases the firing rate in V4 which in turn increases the rate in the IPS.

3.3 Blocking distractions

Further simulations showed that filtering takes place in the IPS for both the IPS-filtering and the V4-bias hypotheses. This turned out to be an emergent property of the strong inhibition between the two local networks in the IPS. The strong inhibition was originally needed to make sure that capacity was the same regardless of whether stimuli were of a single or multiple colours, but here, the same connections facilitate filtering. If the firing rate in one of the colour-networks is increased due to external excitation, it generates a net inhibition that suppresses the rate in the other network. Enough inhibition can lower the rates below the left bifurcation point of figure 12, hence shutting off bistability. For example, when the dlPFC is in resting state, the firing rate of the IPS in the V4-bias hypothesis is 0.02 Hz. When the dlPFC enters filtering mode, the rate of the favoured colour is increased to 0.19 Hz and the other network is suppressed to silence.

3.4 Results for the hypotheses

Simulations of the IPS-filtering and the V4-bias hypotheses showed that they can both handle the capacity modulation as well as filtering (table 1). In IPS-filtering, each

colour-network in dlPFC connects directly to its corresponding colour-network in the IPS. Turning on one of the colour-networks in the dlPFC-area causes the IPS to filter stimuli of the other colour just as explained above. When the whole dlPFC is turned on, capacity is increased for cues regardless of colour.

Instead of connecting directly to the IPS, the V4-bias hypothesis proposes that the dlPFC connects to V4. The increased firing rate in V4 that follows from the top-down signal from dlPFC propagates to IPS, which from V4 receives the same inputs as the IPS in the IPS-filtering hypothesis receives from V4 and dlPFC together. The actual mechanisms for capacity increase and filtering are hence the same in the two hypotheses and what differs is really the way the signal is transmitted.

A finding of both the presented hypotheses was that when only one area of the dlPFC was turned on, as in filtering mode, capacity was enhanced just as much for that favoured colour. Implications of this finding for the model are considered in the discussion.

Table 1. Simulation results

	Simulations	IPS Rate, Hz			Capacity	
Network		dlPFC off	Filtering (red/yellow)	Boosting	dlPFC off	Boosting
IPS-filtering	8	0.02	0.20/0.01	0.17	1.9	2.9
V4-bias	8	0.02	0.19/0.00	0.16	1.8	2.9

Table 1 - Simulation results for the two working hypotheses.

3.4.1 The difficulties of building a V4-filtering network

V4-filtering is different from the other two hypotheses. Although dlPFC inputs target V4 as in the V4-bias hypothesis, filtering occurs in V4 rather than IPS. V4 is a visual area and what needs to be filtered away is the bump of an externally driven stimulus. The direct input of the stimulus in V4 is relatively strong and so are the bumps that are formed. Suppressing these bumps is not as easy as removing the possibility for bistability as in the case of IPS, where the firing rate is just pushed below the bifurcation point. Compared with filtering in the IPS, the inhibitory connections between the colour-networks in V4 must be much stronger than they are in IPS when IPS is doing the filtering.

Simulations showed that it is difficult to realise the hypothesis as it has been defined. The problem is that the inhibition needs to be really strong between the colour-networks in order to even make a noticeable difference in firing rate when filtering. But with that strong inhibition, it is difficult to increase activity when both networks are turned on for the boosting mode, which causes IPS to receive too little excitation to increase capacity significantly. Adding a connection from dlPFC to IPS could be a way of overcoming this problem. Solving the problem of low rates, when both colour-networks are turned on, with direct connection from dlPFC to IPS poses a new difficulty though. Filtering in V4 will, as mentioned above, bring one of the colour-networks down and the other up. When the increased activity of the favoured colour enters an IPS that also receives inputs from the dlPFC, the activity easily adds up to be pushed over the right bifurcation point and the network loses stability. It would be too easy to say these results prove that V4-filtering is impossible to model in this way, but if the solution is there to be found, it involves times consuming fine-tuning. Attempts were made to get there, but the time was too short to solve the problems that mounted.

3.5 Predictions for fMRI comparison

The two working hypotheses, IPS-filtering and V4-bias can be told apart with fMRI. In order for this to be possible, predicted fMRI activity for two hypotheses are here presented as simulated data translated into blood oxygenation level dependent (BOLD) signal, the kind of data generated by fMRI (figure 13). The layout of figure 13 corresponds to figure 10. The dlPFC is not turned on in the simulations of the left column, which is thought to correspond to instructions for a no distraction task conducted with low-capacity subjects. It could also be instructions for an easy no distraction task with a high capacity subject. One network of the dlPFC is turned on in the middle column, corresponding to filtering for a high WMC subject. Both networks are turned on in the right column, corresponding to the no distraction task, also for high capacity subjects.

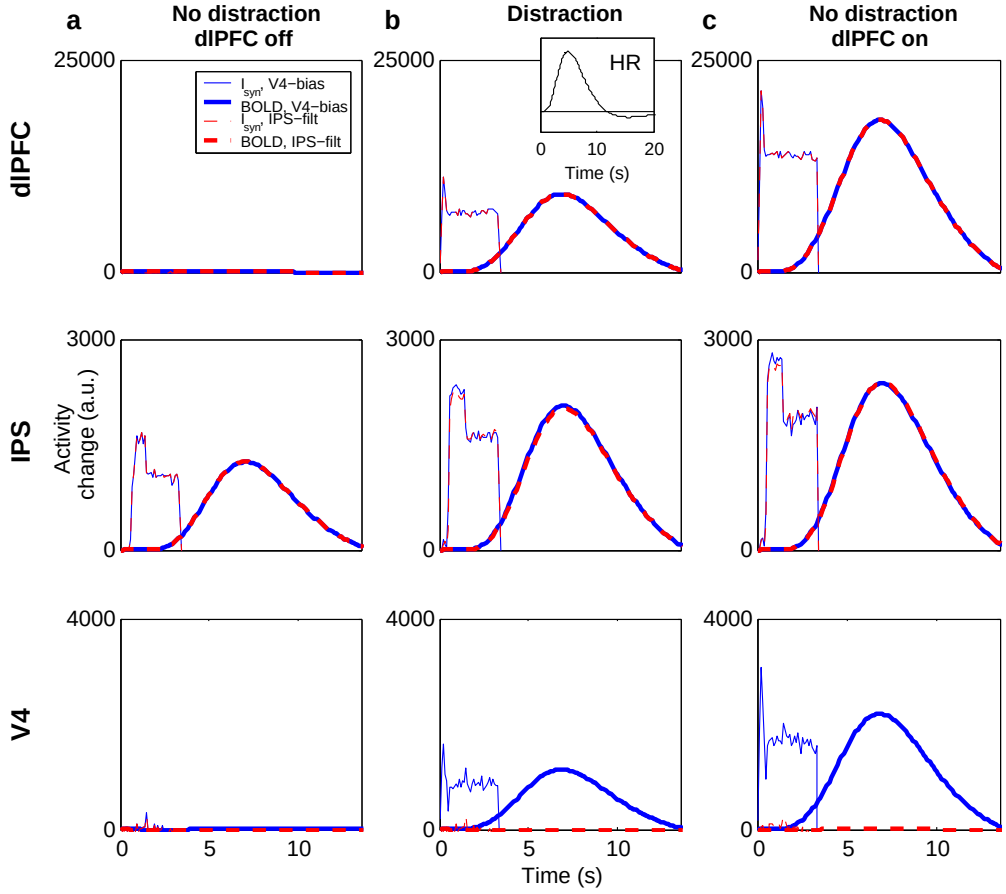

Figure 13 - Predictions for BOLD-responses. Blue lines correspond to the V4-bias hypothesis and red dashed lines to IPS-filtering. Thin lines show population activity as the summed synaptic currents. Thick lines are the BOLD-responses, the population activity convolved with the homodynamic response function showed in the inset. A) No distraction task, dlPFC only firing at spontaneous level. B) Distraction task simulated with one dlPFC area tuned on. C) No distraction task where both areas of dlPFC were turned on. The BOLD-response in V4 is differs in the two hypotheses. A.u.: arbitrary units.

The only area where the hypotheses differed in activity was V4. Thus, what should be looked for in an fMRI experiment is whether the activity of V4 is stable over different task settings or whether it covaries with the activity of the dlPFC. A stable behaviour is

an indication for the IPS-filtering hypothesis, whereas increasing activity in V4 is an indication for the V4-bias hypothesis.

4 Discussion

In this study, we sought to investigate two main questions. The primary question was whether activity in the dlPFC could explain the link between WMC and distractibility. Secondarily, we investigated how the dlPFC connects to the parietal and occipital areas presumably responsible for memory storage and filtering of distracters. The results show that if the part of dlPFC responsible for one colour is turned on, the model stores only stimuli of that colour and filter distractors. On the other hand, if both areas are turned on, the capacity of the whole memory area is increased regardless of stimulus colour. In this way, the dlPFC serves as the link between WMC and distractibility. With a low functioning dlPFC, capacity remains at a low level and distractors are stored just as easily as task relevant cues. A high functioning dlPFC on the other hand, can enhance the WMC and also make sure that distractors are effectively filtered. Capacity is enhanced by increased excitation to the whole memory area, and filtering is achieved by biased competition when only favoured attributes of the memory area are targeted by the top-down signal.

Successful simulations were made for the IPS-filtering and the visual bias hypotheses. That is, the top-down signal can target either a late stage of the visual cortex or the memory area directly. In both cases, the filtering by biased competition happens in the IPS. The different hypotheses predict different BOLD-responses, and support for either can hence be sought for in fMRI data from an experiment using a similar protocol. As has been mentioned, such an experiment is being conducted but data is not yet available.

I never succeeded in creating a V4-filtering network, and the problems with filtering in the visual area might be explained by the results of Kastner & Ungerleider (2001), that spatial filtering in visual areas is relatively local and a matter of a percentage change, whereas it has been found that filtering in later regions can be total. “In prefrontal cortex, filtering of ignored locations is strong, early and spatially global” (Everling et al. 2002). The reason why it is relatively easy to cause a total filtering in later areas is probably because these are bistable memory areas. The non-linearity of a bistable area implies that when it is inhibited below the left bifurcation point (figure 12), it will lose all its capacity to hold memories. The area can hence be shut off by a relatively small change in activity. The absence of non-linearities could explain the difficulties of filtering effectively in the visual areas. All activity that needs to be reduced needs to be so actively. This requires very strong inhibition and makes the filtering like a balance scale, when one is pushed down the other goes up.

The results in this study reproduce the results of Vogel et al. (2005), who used event-related potentials to show that subjects with high WMC filter stimuli more efficiently than those with low WMC. We hypothesize that the dlPFC of subjects in that study was in capacity-enhancing mode (figure 10.C) when their WMC was measured, and that the dlPFC was in filtering mode (figure 10.B) when their distractibility was measured. Thus, the strength of dlPFC activity explains the findings by Vogel et al. (2005) that persons with low WMC cannot filter distractions while persons with high WMC can. The way in which the dlPFC was modelled to maintain different goal relevant rule

representations depending on the task goes hand in hand with the strong link between dlPFC and intelligence (Conway 2003, Kane & Engle 2002). The dlPFC has connections to many brain regions and its importance to intelligence seem to stem from allocation of its resources in accordance with dynamically changeable representations, which is exactly the abilities feature by the dlPFC in the present model (Engle et al. 1999, Miller & Cohen, 2001). The rule encoding representations held in this task were only based on colours, but it is thought that there are many dimensions by which rules can be integrated and mapped, and as they become more complex, the dlPFC activation seem to increase (Kruger et al. 2002). The common dependency of dlPFC activity that has here been observed for WMC and filtering according to rules could provide a framework for understanding intelligence.

Even if a main result is that dlPFC activity can explain WMC, it should be noted that it is not implied that it explains all variations in WMC. In figure 11 it was shown that the same dependency for capacity could be achieved by controlling parameters within the IPS. But it is not clear if the internal capacity determining properties can be modulated dynamically the way dlPFC can modulate external input into IPS. It is also hard to see how internal properties could explain the link between capacity and WMC. Nevertheless, it is worth making the point that some interpersonal variation in WMC can be due to differences in the IPS.

The internal and external contributions to WMC play a role also for ADHD. It has previously been found that visuospatial WMC is a sensitive measure of the severity of the ADHD (Westerberg et al. 2004). Following the reasoning from above, this correlation is most likely driven by the part of the WMC that rely on the activity in dlPFC. We argue that it is the attention part of the WM and not the actual storage capacity in the IPS that constitutes the link between WMC and ADHD. In order to get even more reliable and better explaining measures of ADHD, it is therefore important to develop methods to asses WMC in a way that distinguishes between the storage capacity of the IPS and the influence of dlPFC. One such possibility would be a test battery with subtasks that are likely to activate the dlPFC to different degrees.

4.1 Limitations of the model

The individual networks of the model is of the same type as been used in previous studies with this model that have generated interesting and testable findings and predictions, but it experimental studies need to conducted in order to verify the assumptions of the model. However, the connections of the full network used here are obviously very tentative. The assumption that visual stimuli enter via V4 is maybe the most striking, since it has been observed that spatial information often enters the parietal areas through the dorsal pathway. Nevertheless, it is unclear if it would have been necessary to build the model in a significantly different way if other areas were to constitute the main entry point of visual information into the IPS. There are studies showing that rules might be maintained in the parietal cortex, This could mean that there are other areas that interact with those modelled here that need to be incorporated into the model to really understand what is happening. In order to be a manageable project, what has been modelled is probably not the full story of the connections between WMC and distractibility, but part of an explanation.

Interactions between the areas have solely been implemented as excitatory-to-excitatory connections. Inhibitory interactions could make filtering more easily explained, but it would on the other hand make it more difficult to propose hypotheses in line with the increased excitation in IPS needed for capacity enhancement.

4.1.1 Other hypotheses

Except for the restriction to only model top-down signal as excitatory connection, further restrictions were made that reduced the number of hypotheses to three. Whereas all three hypotheses were based on a dlPFC that connects to the excitatory cells of either the visual area or the memory area, it is plausible that the reality is more of a mix of all the hypotheses, but the priority was to set up hypotheses different enough to be meaningfully tested against each other. Since the modelling results are supposed to later be compared with fMRI data, the hypotheses were chosen in a way that they were thought to generate different BOLD activity in different brain regions, so that the hypotheses can be told apart.

An important component left out of the model is neuromodulators, in particular dopamine, which in monkey studies has been shown to have an effect on WM performance (Sawahuchi & Goldman-Rakic 1994) as well as dlPFC activity (Williams 1995). Interestingly, the dopamine system is a major target of the ADHD drug, Ritalina, which decreases distractibility. It might be that the relationship between dopamine on the one hand and WMC and distractibility on the other is caused by its influence on the activity of dlPFC, as suggested by this study. It might also be that other aspects of the dopamine system are important for this relationship. Looking at the bifurcation graph in figure 12 for example, it is possible to think of other ways that other areas could affect the performance of the WM system. Changes in the span between the bifurcation points, or modulation of the width of the connection-curve (Durstewitz & Seamans 2002) are possible ways by which dopamine might regulate WMC. Also, a recent study has shown that dopamine signals to the basal ganglia and prefrontal regions can block distractors by “locking the gates to WM” in a task where distractors were presented during the delay phase (Gruber et al. 2006). Although separate, these two actions of dopamine could provide another explanation for the link between WMC and distractibility.

4.2 Suggestions for future studies

The assumption when the modelling started was that there would be a trade-off between filtering and capacity at high memory loads (deFockert et al. 2001). In other words, when attentional resources are devoted to filtering out distracters, the capacity would be lower than when they are used to boost capacity. However, my results show that capacity for the favoured colour in filtering mode is the same as the total capacity in capacity enhancing mode. It is unclear if this is a realistic situation or not. As it is now, it does not cause any disagreement with the findings of Vogel et al (2005). But that study was only carried out for one level of difficulty. It would be interesting to see the same study conducted for higher loads as well. The model still holds if it would turn out that high capacity subjects maintain their filtering abilities when the WM load approaches their WMC. But if it turns out that the subjects store distractors at high loads, the model would somewhat faulty because then the filtering signal should not increase capacity as well. If our original assumptions were correct, then filtering or boosting becomes a choice of strategy in the face of limited attentional resources: remember only the relevant stimuli but not all, or remember as many stimuli as possible, but with the risk of including distractors. It would be interesting to test

whether such a choice of strategy actually takes place. If so, more or less efficient choices of strategy could be a further difference between people with good and bad WMC, intelligence and resistance to distraction.

It has been seen in this study, that in order for the IPS to benefit from top-down signaling from the dlPFC, it needs to be in an activity state where the extra excitation does not push the activity above the right bifurcation point (figure 12), which would cause spontaneous bumps. This imposes causes a restriction on how large the effect of the dlPFC on the memory area can be. As the dlPFC is turned on during the instruction cue, it directly starts to excite the resting IPS with full power. It would be interesting to investigate the possibility that the dlPFC is turned on gradually to fulfil a demand from IPS. From the resting state, the activation by the instruction cue could turn it in a low persistent activity state. When cues enter the IPS, the extra activity here could via a bottom-up connection to dlPFC push this area to a higher activation level, which in turn is fed back to the IPS. Done this way, the dlPFC could maintain an optimal spontaneous firing rate in IPS even as the WM load increases and the capacity limit could exceed that set by the right bifurcation point.

5 Conclusions

This study has shown, by computational modelling, that rule encoding activity in the dlPFC could explain the link between WMC and distractibility. It has been found that dlPFC can target either the memory area IPS or the visual area V4. Using the predicted fMRI activity, the hypotheses presented in this study can be tested by an fMRI experiment.

Acknowledgement

This thesis has been made as a part of a bigger project in the developmental cognitive neuroscience group at Karolinska Institutet. The results that I have generated will be used to analyse, and be compared with, experimental data from Fiona McNabb. I thank Fredrik Edin, Jesper Tegnér, Torkel Klingberg and Albert Compte that have all been involved in the design of the study. The model I have created builds on code that Albert Compte and Salava Ardid have written for earlier studies (Compte et al. 2000, Ardid et al. 2007).

References

Ardid S., Wang X.J., Compte A., (2007) “An integrated microcircuit model of attentional processing in the neocortex.” Journal of Neuroscience, 8, 27, 32, 8486-8495

Baddeley, A.D., Hitch, G.J., Bower, G.H., (1974). The psychology of learning and motivation. Academic Press, New York.

Bartels, A., Zeki, I., (2000). “The architecture of the colour centre in the human visual brain: new results and a review.” European Journal of Neuroscience, 12, 172-193.

Braver, T.S., Cohen, J.D., Nystrom, L.E., Jonides, J., Smith, E.E. and Noll, D.C., (1997). “A parametric study of prefrontal cortex involvement in human working memory.” NeuroImage 5, 49–62.

Bunge, S.A., Klingberg, T, Jacobsen, R.B., Gabrieli J.D.E., “A resource model of the neural basis of executive working memory., Proceedings of the National Academy of Sciences, 97, 7, 3573-3578.

Compte, A., Brunel, N., Goldman-Rakic, P.S., Wang, X.J., (2000). “Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model.” Cerebral Cortex 10, 910–923.

Compte, A., (2006). “Computational and in vitro studies of persistent activity: edging towards cellular and synaptic mechanisms of working memory.” Neuroscience, 139, 1, 135-51.

Conway, A.R., Kane, R.W., Engle, M.J., (2003). “Working memory capacity and its relation to general intelligence.” Trends Cognitive Science, 7, 12, 547-552.

Curtis, C.E., D’Esposito, M., (2003). “Persistent activity in the prefrontal cortex during working.” Trends in Cognitive Sciences, 2003, 7, 9, 414-423.

De Fockert, J. W., Rees, G., Frith, C. D., Lavie, N., (2001). The role of working memory in visual selective attention. Science, 291, 1803-1806.

Desimone, R., (1998). “Visual attention mediated by biased competition in extrastriate visual cortex.” Philosophical Transactions of the Royal Society London B Biological Science, 353, 1373, 1245–1255.

Desimone, R., Schein, S.J., (1987). “Visual properties of neurons in area V4 of the macaque: sensitivity to stimulus form.” Journal of Neurophysiology, 57, 3, 835-68.

Durstewitz, D., Seamans, J.K., (2002). “The computational role of dopamine D1 receptors in working memory.”, Neural Networks, 15, 561-572.

Duncan, J., (1995). “Attention, intelligence, and the frontal lobes.” In M. S. Gazzaniga (Ed.), The cognitive neurosciences, 721–733. Cambridge, MA, MIT Press.

Edin, F., Macoveanu, J., Olesen, P., Tegner, J., Klingberg, T., (2007). “Stronger Synaptic Connectivity as a Mechanism behind Development of Working Memory-related Brain Activity during Childhood.” Journal of Cognitive Neuroscience, 19, 750-760.

Engle, R.W., Tuholski, S.W., Laughlin, J.E., Conway, A.R., (1999). “Working memory, short-term memory, and general fluid intelligence: a latent-variable approach.” Journal of Experimental Psychology: General, 128, 3, 309-331.

Funahashi, S., Bruce, C.J., Goldman-Rakic, P.S., (1989). “Mnemonic coding of visual space in the monkey’s dorsolateral prefrontal cortex.” Journal of Neurophysiology, 61, 331–349.

Gattass, R., Nascimento-Silva, S., Soares, J.G.M., Lima, B., Jansen, A.K., Cinira A., Diogo, M., Farias, M.F., Marcondes, M., Botelho, E.P., Mariani, O.S., Azzi J., Fiorani, M., (2005). “Cortical visual areas in monkeys: location, topography, connections, columns, plasticity and cortical dynamics.” Philosophical Transactions of the Royal Society London B, 709–731.

Goldman-Rakic, P.S., (1987). “Circuitry of primate prefrontal cortex and regulation of behavior by representational memory.” In Handbook of Physiology, the Nervous System, Higher Functions of the Brain (Plum, F., ed.), 373–417, American Physiological Society

Gray, J.R., Chabris, C.F., Braver T.S., (2003). “Neural mechanisms of general fluid intelligence.” Nature, 6, 3, 316-322.

Gruber, A.J., Dayan P., Gutkin B.S., Solla, S.A. (2006) “Dopamine modulation in the basal ganglia locks the gate to working memory”, Jurnal of Computational Neuroscience 20, 153–166.

Kane, M.J., Engle R.W. (2002). “The role of prefrontal cortex in working-memory capacity, executive attention, and general fluid intelligence: An individual-differences perspective.” Psychonomic Bulletin & Review, 9, 4, 637-671.

Klingberg, T., Fernell, E., Olesen, P., Johnson, M., Gustafsson, P., Dahlström, K., Gillberg, C., Forssberg, H., Westerberg, H., (2005). ”Computerized Training of Working Memory in Children With ADHD-A Randomized, Controlled Trial.” Journal of the American Academy of Child & Adolescent Psychiatry, 44, 2, 177-186.

Kroger, J.K., Sabb, F.W., Fales, C.L., Bookheimer, S.Y., Cohen, M.S., Holyoak, K.J., (2002), “Recruitment of Anterior Dorsolateral Prefrontal Cortex in Human Reasoning: a Parametric Study of Relational Complexity.” Cerebral Cortex, 12, 5, 477-485.

Leung, H.C., Gore, J.C., Goldman-Rakic, P.S., (2002) “Sustained mnemonic response in the human middle frontal gyrus during on-line storage of spatial memoranda.” Journal of Cognitive Neuroscience, 14, 4, 659–671.

Luck, S.J., Vogel, E.K., (1997), “The capacity of visual working memory for features and conjunctions.”, Nature, 390, 279-281.

Marois, R., Ivanoff, J., (2005). “Capacity limits of information processing in the brain.” Trends in Cognitive Sciences, 9, 296-305.

Macoveanu, J., Klingberg, T., Tegner, J., (2006), “A biophysical model of multiple-item working memory: a computational and neuroimaging study.” Neuroscience, 2006, 141, 3, 1611-1618.

Macoveanu, J., Tegner, J., Klingberg, T., (2006), “Neuronal firing rates account for distractor effects on mnemonic accuracy in a visuo-spatial working memory task.” Biological Cybernetics, 96, 4, 407-419.

Miller, G.A., (1956). “The Magical Number Seven, Plus Or Minus Two: Some Limits on Our Capacity for Processing Information.” The Psychological Review, 63, 81-97.

Miller, E.K., Cohen, J.D., (2001). “Integrative theory of prefrontal cortex function.” Annual Review of Neuroscience, 24: 167-202.

Sawahuchi, T., Goldman-Rakic, P.S., (1994). “The Role of Dl-Dopamine Receptor in Working Memory: Local Injections of Dopamine Antagonists Into the Prefrontal Cortex of Rhesus Monkeys Performing an Oculomotor Delayed-Response Task.” Journal of neurophysiology, 71, 2, 515-528.

Snowden, R.J., Treue, S., Erickson, R.G., Andersen, R.A., (1991). ”The response of area MT and V1 neurons to transparent motion.” Journal of Neuroscience, 11, 2768–2785.

Todd, J.J., Marois, R., (2005). “Posterior parietal cortex activity predicts individual differences in visual short-term memory.” Cognitive, Affective, & Behavioural Neuroscience, 5, 2, 144-155.

Toth, L.J., Assad, J.A., (2002). ”Dynamic coding of behaviourally relevant stimuli in parietal cortex.” Nature, 415, 165-168.

Westerberg, H., Hirvikoski, T., Forssberg, H., Klingberg, T., (2004). ”Visuo-Spatial Working Memory Span: A Sensitive Measure of Cognitive Deficits in Children With ADHD”, Child Neuropsychology, 10, 3, 155-161.

Williams, G.V., Goldman-Rakic, P.S., (1995), “Modulation of memory fields by dopamine Dl receptors in prefrontal cortex.” Nature, 376, 572-575.

APPENDIX A

This appendix specifies all values used in the simulations. The equations have been presented in the Methods section, but are here shown again to make the tables easier to read.

$$\begin{cases} C_m \dfrac{\mathrm{d}V}{\mathrm{d}t} = I_{\mathrm{AMPAR}} + I_{\mathrm{NMDAR}} + I_{\mathrm{GABA_AR}} + g_{\mathrm{L}}(V - E_{\mathrm{L}}), & V < V_{th} \\ V = V_{reset}, & V \geq V_{th} \end{cases} \quad (1)$$

$$I_{syn} = g_{syn} s (V - V_{syn}) \quad (2) \qquad g_{syn,ij} = W(\theta_i - \theta_j) G_{syn} \quad (3)$$

All parameters of the integrate-and-fire equation (1) are specified below. It should be noted that there is also refractory period τ_{ref} which sets the time the potential is held at the reset potential after an action potential. In (2) and (3), *syn* is AMPAR, $GABA_AR$ or NMDAR. $g_{syn,ij}$ is the conductance between neurone i and neurone j, which depends on the difference in preferred cue, set by $W(\theta_i - \theta_j)$.

$$W(\theta_i - \theta_j) = J^- + (J^+ - J^-)\mathrm{e}^{-(\theta_i - \theta_j)^2 / 2\sigma^2} \quad (4) \qquad \frac{1}{360}\int_0^{360} W(\theta_i - \theta_j) d\theta_j = 1 \quad (5)$$

For unstructured connections, where all neurones connects equally strong to each other $W(\theta_i - \theta_j)=1$. Structured connections are determined by (4), normalized as in (5). The normalization imposes a functional relationship between the parameters, thus only J^+ and σ needs to be specified.

AMPAR and $GABA_AR$ synaptic gating variables, s, are modelled as an instantaneous jump of magnitude 1 when a spike occurs in the presynaptic neurone followed by an exponential decay with time constant τ_s ms. The NMDA conductance is voltage dependent, with g_{syn} multiplied by $1/(1 + [Mg^{2+}] \exp(-0.062V_m)/3.57)$, $[Mg^{2+}] = 1.0$ mM. The channel kinetics is modelled with (6) and (7).

$$\frac{ds}{dt} = -\left(\frac{1}{\tau_s}\right)s + \alpha_s x(1 - s) \quad (6) \qquad \frac{dx}{dt} = -\left(\frac{1}{\tau_x}\right)x + \sum_i \delta(t - t_i) \quad (7)$$

Excitatory cells		**Inhibitory cells**	
C_m, nF	0.5	C_m, nF	0.2
g_L, nS	25	g_L, nS	2
E_L, mV	-70	E_L, mV	-70
V_{res}, mV	-60	V_{res}, mV	-60
V_{th}, mV	-50	V_{th}, mV	-50
τ_{ref}, ms	2	τ_{ref}, ms	1
τ_{s_AMPA}, ms	2	τ_{s_GABA}, ms	10
α_{s_AMPA}, kHz	1	α_{s_GABA}, kHz	1

τ_{s_NMDA}, ms	100	α_{Ext}, kHz	1
α_{s_NMDA}, kHz	0.5	$\tau_{\_Ext}$, ms	2
τ_{x_NMDA}, ms	2	$rate_{ext}$, kHz	1.8
α_{x_NMDA}, kHz	1		
α_{Ext}, kHz	1		
$\tau_{\_Ext}$, ms	2		
$rate_{ext}$, kHz	1.8		

Table A - Parameter values used for all cells in the model. The total capacitance Cm, the total leak conductance g_L, the leak reversal potential E_L, the threshold potential V_{th}, the reset potential V_{res} and the refractory time τ_{ref} . τ_s is the exponential decay time, αs controls the saturation properties of synaptic currents, τ_x controls the rise time of NMDAR channels. Both excitatory and inhibitory cells receive external excitation mediated by AMPAR at a rate of rateext.

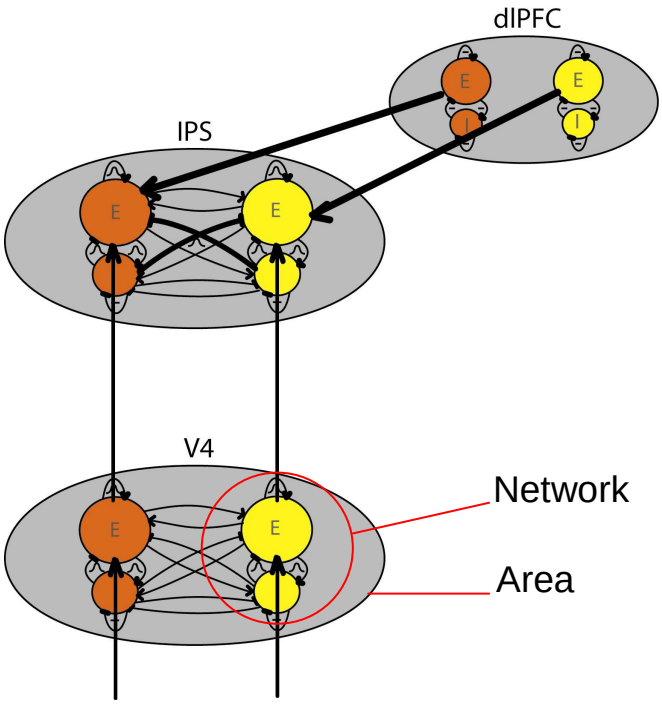

Figure A -Schematic layout of the model. In each time step, each cell receives input currents summed from all cells in the network, the other network in the area and networks in other areas. All connections are specified for the different hypotheses in the tables below.

	IPS	V4		dlPFC
Hypotheses	All hypotheses	IPS-filtering & V4-bias	V4-filtering	All hypotheses
Neuronal connections within the network				
$G_{E\rightarrow E\_AMPA}$, nS	0.03515	0.00475	0.00475	0
$G_{E\rightarrow E\_NMDA}$, nS	0.6429	0.1225	0.1225	0.983
$G_{E\rightarrow I\_AMPA}$, nS	0	0.00475	0.00475	0
$G_{E\rightarrow I\_NMDA}$, nS	0.6336	0.19	0.19	0.74
$G_{I\rightarrow E}$, nS	0.863	0.323	0.323	0.937

$G_{I\rightarrow I,}$ nS	0.6857	0.08265	0.08265	0.725
$\sigma_{E\rightarrow E,}$ deg	9.4	10.1	10.1	-
$\sigma_{E\rightarrow I,}$ deg	32.4	32.4	32.4	-
$\sigma_{I\rightarrow E,}$ deg	32.4	32.4	32.4	-
$\sigma_{I\rightarrow I,}$ deg	-	-	-	-
$J^{+}_{E\rightarrow E}$	4.1	5.15	5.15	1.5
$J^{+}_{E\rightarrow I}$	1.5	1.58	1.58	1.5
$J^{+}_{I\rightarrow E}$	1.5	1.5	1.5	1.5
$J^{+}_{I\rightarrow I}$	1.5	1.5	1.5	1.5
Neuronal connections between networks in an area				
$G_{E\rightarrow E2\_AMPA,}$ nS	0.00285	0.00025	0.00025	0
$G_{E\rightarrow E2\_NMDA,}$ nS	0.0521	0.0065	0.0065	0
$G_{E\rightarrow I2\_AMPA,}$ nS	0	0.00025	0.00025	0
$G_{E\rightarrow I2\_NMDA,}$ nS	0.0514	0.01	**0.18**	0
$G_{E2\rightarrow E\_AMPA,}$ nS	0.00285	0.00025	0.00025	0
$G_{E2\rightarrow E\_NMDA,}$ nS	0.0521	0.0065	0.0065	0
$G_{E2\rightarrow I\_AMPA,}$ nS	0	0.00025	0.00025	0
$G_{E2\rightarrow I\_NMDA,}$ nS	0.0514	0.01	**0.18**	0
$G_{I\rightarrow E2,}$ nS	0.1399	0.017	0.017	0
$G_{I\rightarrow I2,}$ nS	0.0556	0.00435	0.00435	0
$G_{I2\rightarrow E,}$ nS	0.1399	0.017	0.017	0
$G_{I2\rightarrow I,}$ nS	0.0556	0.00435	0.00435	0
$\sigma_{E\rightarrow E2,}$ deg	9.4	14.4	14.4	0
$\sigma_{E\rightarrow I2,}$ deg	32.4	-	-	0
$\sigma_{E2\rightarrow E,}$ deg	9.4	14.4	14.4	0
$\sigma_{E2\rightarrow I,}$ deg	32.4	-	-	-
$\sigma_{I\rightarrow E2,}$ deg	32.4	-	-	-
$\sigma_{I\rightarrow I2,}$ deg	-	-	-	-
$\sigma_{I2\rightarrow E,}$ deg	32.4	-	-	-
$\sigma_{I2\rightarrow I,}$ deg	-	-	-	-
$J^{+}_{E\rightarrow E2}$	4.1	5.0	5.0	0
$J^{+}_{E\rightarrow I2}$	1.5	1.58	1.58	0
$J^{+}_{E2\rightarrow E}$	4.1	5.0	5.0	0
$J^{+}_{E2\rightarrow I}$	1.5	1.58	1.58	0
$J^{+}_{I\rightarrow E2}$	1.5	1.5	1.5	0
$J^{+}_{I\rightarrow I2}$	1.5	1.5	1.5	0
$J^{+}_{I2\rightarrow E}$	1.5	1.5	1.5	0
$J^{+}_{I2\rightarrow I}$	1.5	1.5	1.5	0
External input				
$G_{E\_EXTERNAL}$	0.00186	0.015	0.015	0.003
$G_{I\_EXTERNAL}$	0.001725	**0.045**	**0.026**	0.00238

Table B - Table for neuronal connections in the model. E referrers to the excitatory population and I the inhibitory, E→E is connections from excitatory neurones to other excitatory neurones, E→I is excitatory to inhibitory neurones, etc. In the sections of connections between the networks in an area, E→E2 should be read as the excitatory neurones in one of the networks to the excitatory neurones in the other network. A dash instead of sigma value indicates unstructured connection. Some figures are written in bold to highlight the differences between the hypotheses.

	IPS-filtering	V4-bias	V4-filtering
Neuronal connections between areas			
$G_{V\rightarrow P,}$ nS	0.088	0.088	0.088
$G_{D\rightarrow P,}$ nS	**0.0115**	-	-
$G_{D\rightarrow V,}$ nS	-	**0.13**	**0.07**
$\sigma_{V\rightarrow P}$, deg	11.2	11.2	11.2
$\sigma_{D\rightarrow P}$, deg	-	-	-
$\sigma_{D\rightarrow V}$, deg	-	-	-
$J^{+}_{V\rightarrow P}$	5.2	5.2	5.2
$J^{+}_{D\rightarrow P}$	0.79	-	-
$J^{+}_{D\rightarrow V}$	-	0.79	0.79

Table C - Connections between areas are all excitatory-to-excitatory AMPA mediated. V = V4, P = IPS, D = dlPFC. A dash indicates that the areas are unconnected.