

Available online at www.sciencedirect.com

SciVerse ScienceDirect

www.elsevier.com/locate/brainresBRAIN
RESEARCH

Research Report

Cortical attractor network dynamics with diluted connectivity

Edmund T. Rolls^{a, b, *}, Tristan J. Webb^c^aOxford Centre for Computational Neuroscience, Oxford, UK^bUniversity of Warwick, Department of Computer Science, Coventry CV4 7AL, UK^cUniversity of Warwick, Department of Computer Science and Complexity Science Centre, Coventry CV4 7AL, UK

ARTICLE INFO

Article history:

Accepted 2 August 2011

Available online 7 August 2011

Keywords:

Decision-making

Noise

Fano factor

Stochastic neurodynamics

Neural encoding

Attractor network

Diluted connectivity

Variability of firing

ABSTRACT

The connectivity of the cerebral cortex is diluted, with the probability of excitatory connections between even nearby pyramidal cells rarely more than 0.1, and in the hippocampus 0.04. To investigate the extent to which this diluted connectivity affects the dynamics of attractor networks in the cerebral cortex, we simulated an integrate-and-fire attractor network taking decisions between competing inputs with diluted connectivity of 0.25 or 0.1, and with the same number of synaptic connections per neuron for the recurrent collateral synapses within an attractor population as for full connectivity. The results indicated that there was less spiking-related noise with the diluted connectivity in that the stability of the network when in the spontaneous state of firing increased, and the accuracy of the correct decisions increased. The decision times were a little slower with diluted than with complete connectivity. Given that the capacity of the network is set by the number of recurrent collateral synaptic connections per neuron, on which there is a biological limit, the findings indicate that the stability of cortical networks, and the accuracy of their correct decisions or memory recall operations, can be increased by utilizing diluted connectivity and correspondingly increasing the number of neurons in the network, with little impact on the speed of processing of the cortex. Thus diluted connectivity can decrease cortical spiking-related noise. In addition, we show that the Fano factor for the trial-to-trial variability of the neuronal firing decreases from the spontaneous firing state value when the attractor network makes a decision.

This article is part of a Special Issue entitled "Neural Coding".

© 2011 Elsevier B.V. All rights reserved.

1. Introduction

Many processes in the brain are influenced by the noise or variability of neuronal spike firing (Deco et al., 2009; Faisal et al., 2008; Rolls and Deco, 2010). The action potentials are generated in a way that frequently approximates a Poisson process, in which the spikes for a given mean firing rate occur at times that are essentially random (apart from a small effect of the refractory period), with a coefficient of variation of the

interspike interval distribution (CV) near 1.0 (Rolls and Deco, 2010; Softky and Koch, 1993). The sources of the noise include quantal transmitter release, and noise in ion channel openings (Faisal et al., 2008; Ribault et al., 2011). The membrane potential is often held close to the firing threshold, and then small changes in the inputs and the noise in the neuronal operations cause spikes to be emitted at almost random times for a given mean firing rate. Spiking neuronal networks with balanced inhibition and excitation currents and associatively modified

* Corresponding author at: Oxford Centre for Computational Neuroscience, Oxford, UK.

E-mail address: Edmund.Rolls@oxcns.org (E.T. Rolls).

URL: <http://www.oxcns.org> (E.T. Rolls).

recurrent synaptic connections can be shown to possess a stable attractor state where neuron spiking is approximately Poisson too (Amit and Brunel, 1997; Miller and Wang, 2006). The noise caused by the variability of individual neuron spiking which then affects other neurons in the network can play an important role in the function of such recurrent attractor networks, by causing for example an otherwise stable network to jump into a decision state (Deco and Rolls, 2006; Rolls and Deco, 2010).

If such an attractor network is provided with two or more inputs, as illustrated in Figs. 1a and b, each biasing an attractor, then this forms a biased competition model of decision-making in which a high firing rate of one of the possible attractor states represents a decision (Deco and Rolls, 2006; Wang, 2002, 2008). The noise in the operation of the system makes the decision-making process non-deterministic, with the system choosing one of the attractor states with a probability that depends on the relative strengths of the different input

biases λ_1, λ_2 , etc. (see Experimental procedures) (Deco and Rolls, 2006; Wang, 2002). The randomness or stochasticity in the operation of the system can be advantageous (Rolls and Deco, 2010). It provides a basis for probabilistic decision-making in which each decision will be sampled in a way that depends on the relative strengths of the inputs. It is adaptive in memory recall which by being probabilistic allows different memories to be recalled from occasion to occasion, helping with creative thought processes. The stochasticity is also useful in signal detection which can become more sensitive than a fixed threshold system in the process known as stochastic resonance (Rolls and Deco, 2010). We note that there are less biologically plausible race or accumulator models of decision-making that involve linear accumulation of noisy inputs until some threshold is reached (Palmer et al., 2005; Ratcliff and McKoon, 2008; Smith and Ratcliff, 2004). Although those and other models can describe decision time distributions and related phenomena (Braun and Mattia, 2010; Bressloff, 2010; Carpenter et al., 2009; Gigante et al., 2009; Miller, 2006; Ratcliff et al., 1999; Sakai et al., 2006; Wong and Huk, 2008; Wong and Wang, 2006) and some aspects of experimentally investigated decision-making (Beck et al., 2008; Ditterich, 2006; Miller and Katz, 2010; Resulaj et al., 2009; Roitman and Shadlen, 2002; Wong et al., 2007), that is not the aim of the work described here, which is instead to examine how diluted connectivity affects the dynamics of a much more biologically plausible neuronal network model of decision-making, which in turn has implications for understanding disorders of brain function (Loh et al., 2007; Rolls, 2005, 2011b; Rolls and Deco, 2010; Rolls


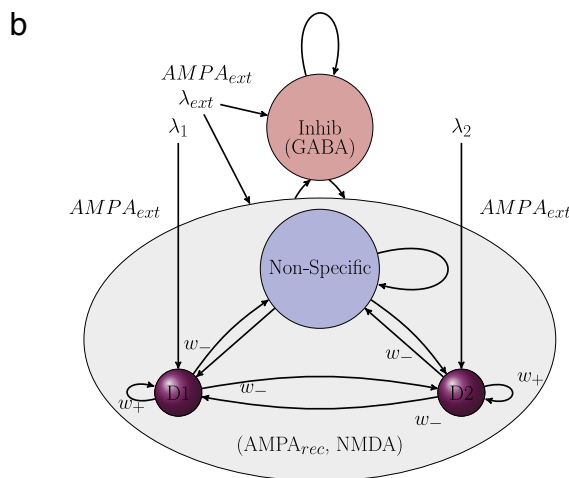
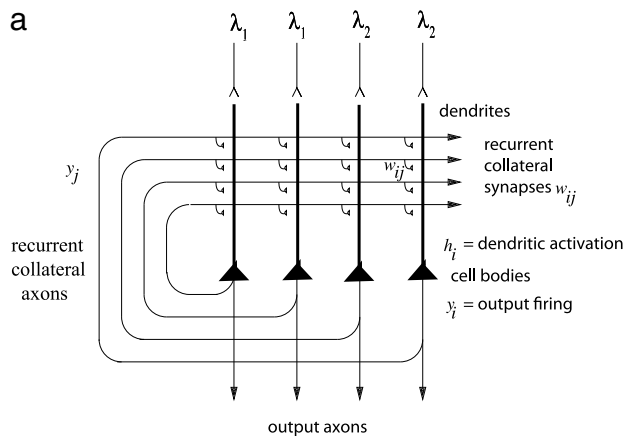


Fig. 1 – (a) Attractor or autoassociation single network architecture for decision-making. The evidence for decision 1 is applied via the λ_1 inputs, and for decision 2 via the λ_2 inputs. The synaptic weights w_{ij} have been associatively modified during training in the presence of λ_1 and at a different time of λ_2 . When λ_1 and λ_2 are applied, each attractor competes through the inhibitory interneurons (not shown), until one wins the competition, and the network falls into one of the high firing rate attractors that represents the decision. The noise in the network caused by the random spiking times of the neurons (for a given mean rate) means that on some trials, for given inputs, the neurons in the decision 1 (D1) attractor are more likely to win, and on other trials the neurons in the decision 2 (D2) attractor are more likely to win. This makes the decision-making probabilistic, for, as shown in (c), the noise influences when the system will jump out of the spontaneous firing stable (low energy) state S, and whether it jumps into the high firing state for decision 1 (D1) or decision 2 (D2). (b) The architecture of the integrate-and-fire network used to model decision-making (see text). (c) A multistable ‘effective energy landscape’ for decision-making with stable states shown as low ‘potential’ basins. Even when the inputs are being applied to the network, the spontaneous firing rate state is stable, and noise provokes transitions from the low firing rate spontaneous state S into the high firing rate decision attractor state D1 or D2. If the noise is greater, the escaping time to a decision state, and thus the decision or reaction time, will be shorter (see Rolls and Deco, 2010).

et al., 2008a,b). A possible link to the more biological type of model we investigate here is that at the bifurcation (decision) point, the attractor neuronal network model network can be described as operating according to a diffusion process, although unlike the race or accumulator models, the diffusion process may be non-linear (Deco et al., 2007; Marti et al., 2008; Rolls and Deco, 2010).

For these advantageous stochastic processes to be realized in the brain, the amount of noise must be significant. One factor that affects the amount of noise is the number of neurons in the fully connected network. As the number of neurons approaches infinity, the noise or statistical fluctuations caused by the relatively random spike timing of the neurons for a given mean firing rate decrease to zero (Mattia and Del Giudice, 2004; Rolls and Deco, 2010) (if the firing of the different neurons is independent; Webb et al., *in press*), and the mathematically convenient mean-field approximation holds, allowing many properties of the system to be calculated analytically (Brunel and Wang, 2001; Deco and Rolls, 2006; Rolls and Deco, 2010; Wang, 2002). Using integrate-and-fire attractor network simulations of decision-making which include the spiking-related noise, we have shown that the stochastic fluctuations in a finite-sized system do indeed decrease as the number of neurons in the network is increased, but are still a significant influence to produce probabilistic decision-making with fully connected networks with 4096 neurons and 4096 synapses per neuron (Deco and Rolls, 2006). This is biologically relevant in that neocortical neurons are likely to have this order (4000–9000) of recurrent collateral excitatory connections from other pyramidal cells (Abeles, 1991; Braitenberg and Schütz, 1991; Elston et al., 2006; Rolls, 2008). Further, the stochasticity is greater if the graded firing rate distributions found biologically (Franco et al., 2007; Rolls, 2008; Rolls and Tovee, 1995; Rolls and Treves, *in press*; Rolls et al., 1997b; Treves et al., 1999) are incorporated into the model (Webb et al., *in press*), relative to the more usually modeled binary firing rate distributions.

In this paper we investigate the effects of dilution of the connectivity on the noise in the network. In integrate-and-fire simulations, it is convenient to analyze the performance with full connectivity. Reasons for this include the fact that there is then a mean-field equivalent, which there is not with asymmetrically diluted connectivity; that the simulations are easier to code and run much faster with full connectivity; and that the application of approaches using energy minima applied to symmetric fully connected networks (Hopfield, 1982) no longer holds formally with asymmetrically diluted connectivity, though extensions of the approach do show that networks with diluted connectivity trained with an associative rule are still likely to perform with many of the same properties as a fully connected network (Bovier and Gayraud, 1992; Perez Castillo and Skantzos, 2004; Rolls and Treves, 1998; Rolls et al., 1997a; Treves, 1991; Treves and Rolls, 1991). For example, with diluted connectivity, the number of patterns that can be stored and recalled correctly is still determined mainly by the number of connections onto each neuron received from the other excitatory neurons in the network (Bovier and Gayraud, 1992; Perez Castillo and Skantzos, 2004; Rolls and Treves, 1998; Rolls et al., 1997a; Treves, 1991; Treves and Rolls, 1991).

However, attractor networks in the neocortex and the hippocampus have diluted connectivity. An estimate for the rat hippocampus is 4% (12,000 connections C per CA3 neuron, and 300,000 CA3 neurons). An estimate for neocortex might be 10%. (For the neocortex, assuming that there are 10,000 recurrent collaterals per pyramidal cell, that the density of pyramidal cells is 30,000/mm³ (Rolls, 2008), that the radius of the recurrent collaterals is 1 mm, and that we are dealing with the superficial (or deep) layers of the cortex with a depth of approximately 1 mm, the connectivity between the superficial (or deep) pyramidal cells would be approximately 10%. Electrophysiological data consistent with this estimate have been found (Perin et al., 2011).) What impact does this diluted connectivity have on the dynamics of the operation of cortical networks? Does it make the decision or memory recall times slower, due to a potentially longer path for every neuron to reach every other neuron functionally? If so, how much slower? Does the diluted connectivity make the diluted networks more or less stable, both in terms of spiking noise-provoked jumps from the spontaneous state to a high firing rate state; and when in a high firing rate state implementing a short-term memory, out of that state? We study this in networks in which we keep the number of excitatory recurrent collateral connections to a neuron constant, for this is the leading factor in determining the memory capacity of the network, that is the number of memories that can be stored and successfully retrieved (Rolls, 2008; Rolls and Treves, 1998; Treves, 1991). We assume that the brain is designed to make the capacity as large as possible, and that in the order of 10,000 recurrent collateral connections per neuron is about as high as can easily be produced biologically. The finite-size noise is influenced by the number of neurons N in a fully connected network, and this is also C , the number of recurrent collateral associatively modifiable connections per neuron in a fully connected network (Rolls and Deco, 2010). In a diluted network, it may be N which influences the finite-size noise, and this is one issue we investigate.

The issue of how diluted connectivity affects the stochastic dynamics of networks in the brain is important, for as just described, connectivity in the hippocampus and locally in the neocortex is diluted (with the dilution being of the order of 0.04 and 0.1 respectively), and any changes in the noise may be very important in understanding decision-making, memory recall, short-term memory, creativity (Rolls and Deco, 2010) and neuropsychiatric disorders such as schizophrenia and obsessive-compulsive disorder (Loh et al., 2007; Rolls, 2005, 2011b; Rolls and Deco, 2010, 2011; Rolls et al., 2008a,b). As far as we know, this is the first investigation of how dilution of the connectivity affects the stochastic dynamics of neural networks in the brain.

We investigate these issues in networks with 800 recurrent collateral synapses per neuron, comparing networks with dilutions of 0.25 and 0.1 with networks with full connectivity. We also extend the analyses to larger networks.

2. Results

Decision-making was investigated in the integrate-and-fire neuronal attractor network model of cortical decision-making

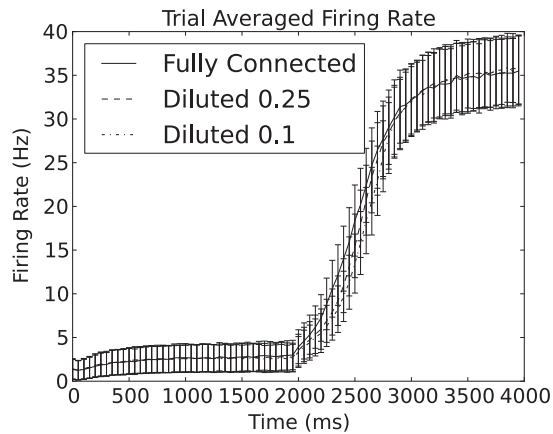


Fig. 2 – Mean firing rates over 1200 trials for the winning pool in fully connected networks and with networks with connectivity diluted to 0.25 and to 0.1. The first 2 s is the period of spontaneous firing before the decision cues were applied at time 2 s. The winning pool is chosen to be the pool with average firing rate of 10 Hz greater than the competing pool in the last 1 s of the simulation. The small differences in mean rate at the decision time after $t=2$ s are due to the small increase in the decision times with dilution.

mechanisms illustrated in Fig. 1 and described in the *Experimental procedures*. The decision-making was probabilistic because of the stochastic fluctuations produced by the almost random (Poisson) firing times for a given mean firing rate of the neurons in the network (Rolls and Deco, 2010). Fig. 2 shows that we had achieved changes, in the networks with connections between the neurons diluted to 0.25 or 0.1, that did not affect the mean firing rates of the decision populations of neurons. The first 2 s of the simulation are spontaneous firing before the decision cues are applied starting at time=2000 ms. The firing rates of the neurons in the winning decision-making pool or population of neurons are shown. The losing population maintained its activity close to the spontaneous level of firing during the decision period from time=2000 to 4000 ms. The performance of the network with a diluted connectivity=0.25 is described especially in Sections 2.1–2.4, and with dilution=0.1 elsewhere and also where stated. The size of the network was set as shown in Table 2 with $C_{\text{fromDpool}}=80$ synapses per neuron within a decision pool, except where otherwise stated.

2.1. Decision time

With diluted connectivity, it is possible that the attractor network might settle more slowly towards a decision attractor state, due to the longer time that it might take any ongoing changes to reach all neurons, and because the connections are no longer symmetric between pairs of neurons. With the method of diluting the connectivity, reciprocal connections were present only by chance; but if present the weights would be equally strong due to the Hebbian learning. Indeed, analytically with asymmetric diluted connectivity, there is no longer an energy minimum of the type defined by Hopfield (1982), though it has been shown that such a network still has many of the same properties, including storage capacity, completion, etc. (Treves, 1991; Treves and Rolls, 1991).

However, the dynamics of the settling into the attractor has not been investigated previously.

The decision time distributions for the fully connected network and the networks with dilutions of 0.25 are shown in Fig. 3. We found that the decision times were a little slower, approximately 50 ms slower, with diluted than with complete connectivity (952 ms for the diluted case vs 894 ms for the fully connected case, $p<0.0002$). Factors that influence the shape of the distribution are considered in the *Discussion*.

An important measure of the noise in the system is the escaping time of the system after the decision cues are applied from the spontaneous state to a decision state (Rolls and Deco, 2010). The implication is that the dilution makes the system operate less noisily. The smaller standard deviations of the decision times with dilution (Fig. 3) are also consistent with less noise in the diluted networks.

2.2. Decision accuracy

The accuracy of the correct decisions (with $\Delta\lambda=6.4$) was 64.3% in the fully connected case and 75.7% in the 0.25 diluted case. This was calculated over 1200 trials apart from those that were rejected from the analysis due to instability of the spontaneous state (that is, instability in the 2 s period before the decision cues were applied). The greater accuracy in the diluted case is probably related to the fact that there were more neurons in the attractor for each decision pool: 320 neurons for the 0.25 dilution case, and 80 for the fully connected case.

2.3. Stability of the spontaneous state

Noise and the positive feedback in this system can cause the network to jump into a decision state from the spontaneous state even before the decision cues are applied (at $t=2$ s in our simulations). We analyzed the stability for the diluted vs complete connectivity cases by measuring the percentage of trials on which these simulations transited into or towards a high firing rate decision state before the decision cues were applied at $t=2$ s. The parameters for the fully connected simulation had been set with the mean field analysis so that the mean spontaneous firing rate should be 3 spikes/s. The criterion for instability of the spontaneous state was that the mean rate of either decision pool exceeded 5 spikes/s in the 250 ms before the decision cues were applied.

It was found that the stability of the spontaneous state was increased (keeping the number of connections per neuron constant) in the diluted connectivity networks compared to the fully connected case, as shown in Fig. 4.

This effect was not accounted for by any difference in the mean spontaneous firing rates of the decision pool neurons in the diluted vs fully connected cases, which remained at a mean value of approximately 3 Hz, as shown in Fig. 2.

2.4. Sparseness

The sparseness of the representation was measured during the last 0.5 s period of the simulation in which the network was firing stably in one of the attractors. The sparseness a calculated in Eq. (2) was 0.39 for the fully connected network and 0.43 for the network with connectivity diluted to 0.25. (It is just

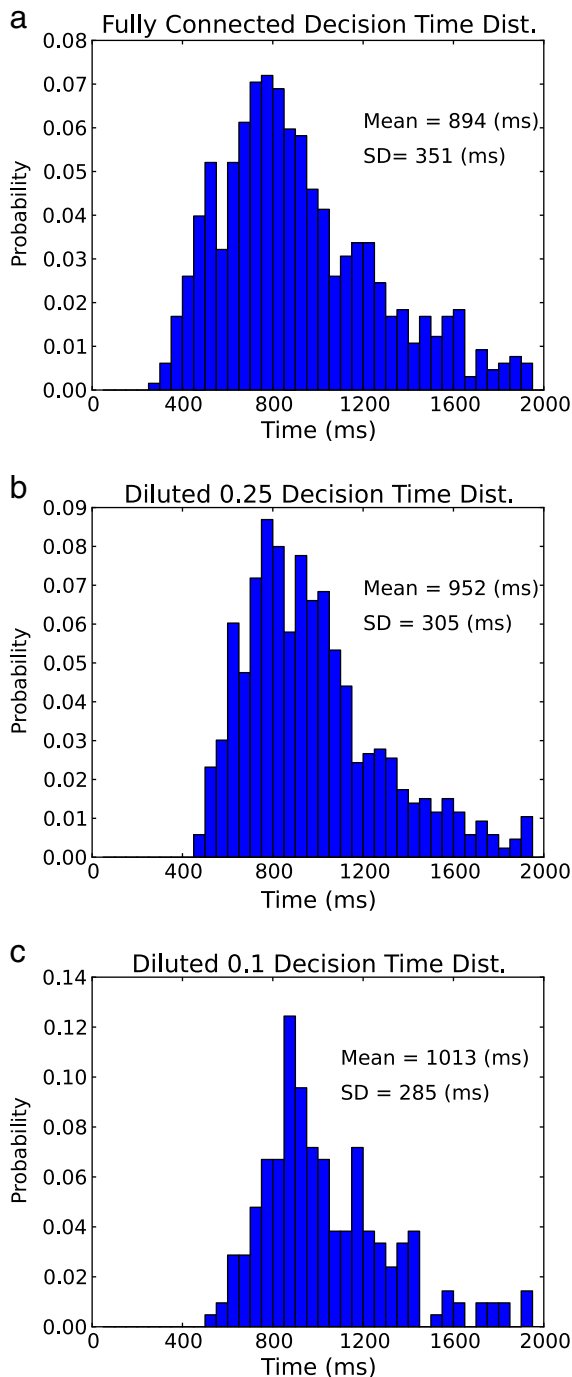


Fig. 3 – Histograms of decision times for 1200 trials with (a) full (b) 0.25 diluted connectivity and (c) 0.1 diluted connectivity. The criterion for a decision time was that the average firing rate of one decision pool should be 25 Hz higher than that of the other decision pool for three consecutive 50 ms periods. $p < 0.0002$ for (a) vs (b) using Kolmogorov–Smirnov tests, t-tests, and Mann–Whitney U tests of the two distributions. The smaller standard deviations of the decision times with dilution are consistent with less noise in the diluted networks.

the small amount of spontaneous firing in the non-selective and losing pools of approximately 3–10 spikes/s that makes these values a little higher than the value that would be obtained with

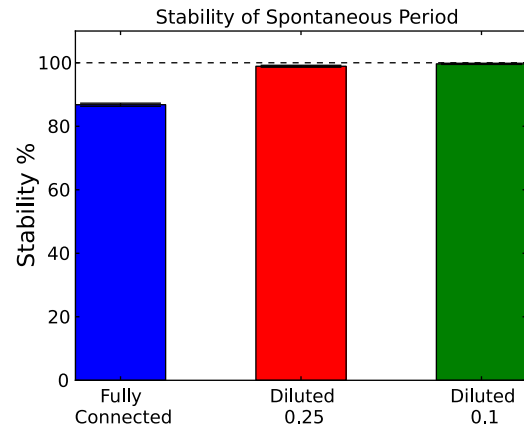


Fig. 4 – Stability of the spontaneous firing state before the decision cues are applied. The percentage of stable trials is shown for networks with full connectivity, and with dilutions to connectivities of 0.25 and 0.1. The network size was as shown in Table 2. The error bars (barely visible) are the standard deviations. There was a significant difference between the fully connected and 0.25 dilution cases ($p < 0.0001$, chi-square test, number of trials 5000).

a binarized firing rate. If the rates below 15 spikes/s were set to zero to exclude spontaneous activity (producing a measure that might be called the sparseness of the population response, where $\text{response} = \text{firing rate} - \text{spontaneous rate}$; cf. Rolls and Tovee, 1995) the sparseness values from the simulations were then very close to the expected values shown in Table 2, namely 0.10 for the fully connected case for which the expected value was 0.1, and 0.25 for the network with 0.25 dilution for which the expected sparseness value was 0.25).

2.5. Performance with connectivity diluted to 0.1

We repeated the investigations just described with a greater level of dilution, 0.1, as this is the value of an estimate of the dilution in the neocortex (0.1), and closer to an estimate of the dilution in the CA3 recurrent collateral connections of the hippocampus (0.04). The size of the network was 2240 excitatory neurons, with 80 connections within each excitatory decision pool, and 800 neurons in each excitatory pool (see Table 2). The stability of the spontaneous state was even greater (see Fig. 4). The spontaneous rate was 2.14 spikes/s (measured in the time period 1–2 s which is the last 1 s of the spontaneous period), and the rate when in the decision attractor was 33.6 spikes/s. (On incorrect trials the mean rate was 1.91 spikes/s. For the fully connected simulation, the mean spontaneous firing rate in the spontaneous period was 2.28 spikes/s.)

We found that the decision times were a little slower, approximately 145 ms slower, with connectivity diluted to 0.1 than with complete connectivity (1013 ms with 0.1 dilution vs 894 ms for the fully connected network, $p < 10^{-7}$, Fig. 3). The decision time was also slower for the 0.1 level of dilution (1013 ms) than for the 0.25 level of dilution (952 ms, $p = 0.001$). The smaller standard deviations of the decision times with dilution of 0.1 are also consistent with less noise in the diluted networks.

The accuracy of the correct decisions (with $\Delta\lambda=6.4$ Hz) was 64.3% in the fully connected case and 90.3% in the 0.1 diluted case. This was calculated over 351 trials in the diluted case. The greater accuracy in the diluted case is probably related to the fact that there were more neurons in the attractor for each decision pool: 800 neurons for the diluted case, and 80 for the fully connected case. The accuracy was intermediate for the 0.25 dilution case (75.7%).

2.6. Performance as a function of the input bias $\Delta\lambda$

The results presented so far have been with $\Delta\lambda=6.4$ Hz. We now explore the effect of different values of $\Delta\lambda$ in diluted networks. Decision times as a function of $\Delta\lambda$ for networks with full connectivity, and with dilutions to connectivities of 0.25 and 0.1, are shown in Fig. 5a. (The network size was twice that

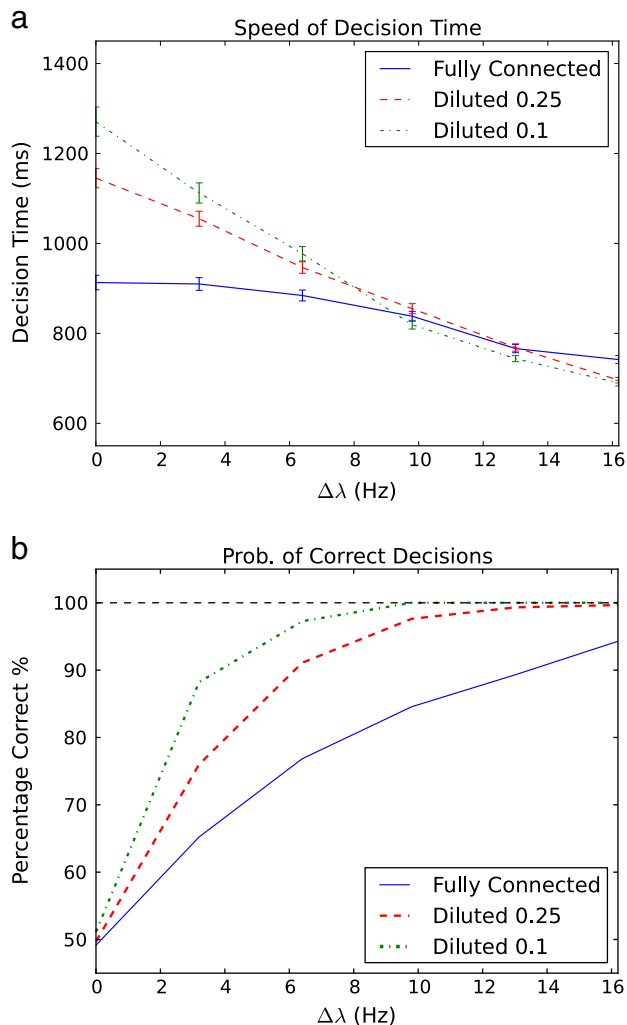


Fig. 5 – (a) Decision time as a function of $\Delta\lambda$ for networks with full connectivity, and with dilutions to connectivities of 0.25 and 0.1. The standard errors are shown as calculated across 600 trials. (b) Percentage correct as a function of $\Delta\lambda$ for networks with full connectivity, and with dilutions to connectivities of 0.25 and 0.1. The network size was $C=1600$ synapses per neuron, with the other parameters $C_{\text{fromDpool}}$ and N_{Dpool} shown in Table 2 also doubled.

of previous simulations, partly to explore these effects with much larger networks. C was 1600 synapses per neuron, with the other parameters $C_{\text{fromDpool}}$, N_{Dpool} and N_E shown in Table 2 altered accordingly. Thus the number of neurons within a decision pool N_{Dpool} for a connectivity of 0.1 was 1600 neurons for the simulations shown in Fig. 5a.) Fig. 5a shows that with low values for the bias $\Delta\lambda$ the fully connected network was fastest, with slower decision times for greater dilutions. A reason for the longer decision times with diluted connectivity is the slower percolation of the effects of activity in any given neuron to all neurons in the population with diluted connectivity. At higher values of $\Delta\lambda$ (10 Hz and above) the bias $\Delta\lambda$ is sufficiently strong in pushing the network towards a decision that the effects of the dilution and percolation effects become no longer relevant to the decision time.

Fig. 5b shows the percentage correct as a function of $\Delta\lambda$ for networks with full connectivity, and with dilutions to connectivities of 0.25 and 0.1. The percentage correct at most values of $\Delta\lambda$ is better for the diluted networks than for the fully connected network. The interpretation is that there is less noise within a decision pool as the number of neurons in the decision pool increases, in this case from 160 neurons in the fully connected case to 640 neurons with 0.25 dilution and 1600 neurons with 0.1 dilution.

2.7. Noise in the system: the Fano factor

Another measure of the noise in the system is the variance of the firing rates of the neurons in a decision pool during decision-making. If the firing of the neurons in a pool has more variance, that pool may be more likely to cross a bifurcation from the spontaneous firing rate state and to enter a decision state without any decision cue, or to make a decision more rapidly after the decision cues have been applied (cf. Fig. 1).

Fig. 6 shows the Fano factor for the neurons in the winning pool, where the Fano factor was defined as the variance across trials/the mean firing rate across trials using time windows of 50 ms. Although the Fano factors are in general similar for the 0.25 diluted and full connectivity cases, there is more variability in the Fano factor in the different 50 ms time windows in Fig. 6 in the pre-decision cue spontaneous firing period from 0 to 2 s in the fully connected compared to the diluted connectivity cases. An implication is that with the Fano factor, and the variance, being calculated over 320 neurons in the diluted case compared to 80 neurons in the full connectivity case, there are more statistical fluctuations, at least in different time windows, in the case with fewer neurons in a decision pool, the fully connected case.

In addition, the elevation in the Fano factor at the time that the decision is taken in the fully connected case is related to the fact that on some trials the decision is fast, and on others slower, so that on some trials the rate is higher, and on others still low, soon after the decision cues are applied. This greater variability in the decision time with the fully connected network is very consistent with the other evidence that there is greater noise in the fully connected than the diluted network, so that the variability in the decision times is greater in the fully connected case. Further factors involved are discussed elsewhere (Churchland et al., 2011; Renart et al., 2007).

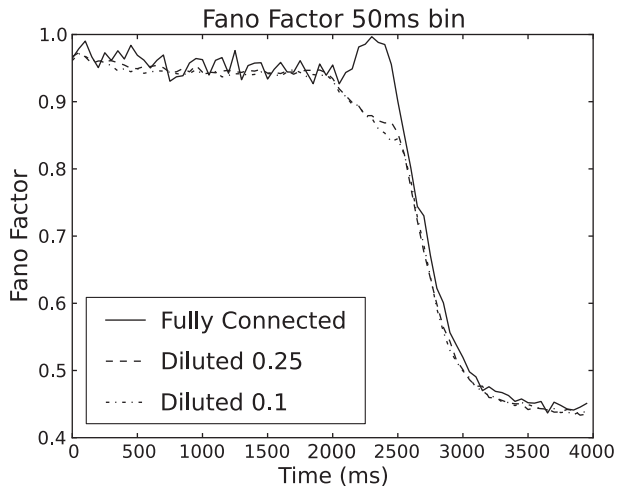


Fig. 6 – The Fano factors for the fully connected network and for the networks with diluted connectivity of 0.25 and 0.1. The Fano factor for the neurons in the winning pool is shown, where the Fano factor was defined as the variance across trials/the mean firing rate across trials for the time windows of 50 ms. The small increase in the Fano factor at the time of the decision in the fully connected case is related to the differences in the decision times on different trials. The decision cues were presented at time = 2 s.

The implication is that an advantage of diluted connectivity (if it is implemented by increasing the number of neurons in a network, as in the brain, while keeping the number of connections C onto each neuron constant), is that the diluted connectivity case will, with more neurons, be more stable with fewer fluctuations of the type illustrated in the fully connected network in the spontaneous firing rate period. The relevance of this is that the spontaneous firing rate period is the state when the decision cues are applied, and the state from which a decision is taken. This smaller stochastic fluctuation with the diluted connectivity appears to be an adaptive situation for the brain, in that too much noise can lead to instability of cortical systems, and associated disorders such as epilepsy, and neuropsychiatric disorders such as schizophrenia (Loh et al., 2007; Rolls, 2005, 2011b; Rolls and Deco, 2010; Rolls et al., 2008b).

The decrease in all networks in the Fano factor from approximately 1.0 during the spontaneous period of firing to a value of approximately 0.45 after the decision has been taken and the system is in a high firing rate state for one of the decision attractor pools is considered in the Discussion.

The Fano factor with 0.1 dilution shown in Fig. 6 was similar to that of the 0.25 diluted case, and different in the ways described above from the fully connected case.

2.8. Noise in the system: within-trial variability

To assess the variability of the firing within a trial, as a measure of the internal noise in the system, we calculated the coefficient of variation of the firing in the last 1 s of the spontaneous period, using 50 ms bins each containing the

number of spikes from a single neuron. This coefficient of variation was then averaged across all the neurons for each trial, and then the mean value across trials was calculated (ensuring that only trials with stability maintained throughout the spontaneous period were included) to ensure that the estimated coefficient of variation was reliable. The CV measures the variability in the firing of single neurons on individual trials, and is a useful measure of the noise in the system. The CV was 2.444 (± 0.010 se) for the diluted network, and 2.459 (± 0.016 se, $p < 0.05$) for the fully connected case, for the spontaneous, that is, pre-cue period. There was thus less stochasticity in the diluted network than in the fully connected network. This was related to the larger number of excitatory neurons in the diluted connectivity, and was found even though the number of connections onto each neuron was identical in the diluted and fully connected networks. This finding is consistent with the finding that there was more stability in the diluted networks in the spontaneous period, in that with less noise, there was less tendency to jump because of the noise to a high firing rate attractor state.

We also measured the CV in the period when the networks had fallen into a high firing rate decision attractor. For the final 1 s period the same calculations showed that the CV was 0.16618 for the 0.25 diluted network, and 0.27889 for the fully connected case, again indicating less noise in the diluted network.

3. Discussion

In integrate-and-fire simulations of an attractor decision-making network, we have shown that the noise is smaller for diluted connectivity than for full connectivity. This was shown by the greater stability in the spontaneous period of the diluted simulations, and by the slower decision times with diluted than with full connectivity (Figs. 3 and 4). The slower decision times in the diluted case reflect less noise, as it is the noise that with the parameters chosen is required for the network to escape from the spontaneous state into a decision state. Further evidence is that the Fano factor, reflecting the trial by trial variation in the firing rate, was more steady for the diluted network as a function of time within a trial (that is in the different 50 ms bins in each trial), as shown in Fig. 6. Further, the measure of the within-trial variability of the firing rate, the coefficient of variation, was also smaller for the diluted than for the fully connected network, reflecting less noise in the diluted network. Moreover, the accuracy of the decisions was also greater with the diluted network, reflecting fewer noise-induced errors. The finding that for more diluted networks performance and also decision time both increase can be described as a speed-accuracy tradeoff due to a reduction of the noise.

The exact distribution of the decision times shown in Fig. 3 is not a focus of the results presented here. Elsewhere we have described part of a mechanism by which the decision time distribution is affected by error trials, and indeed may be longer for error compared to correct trials (Rolls et al., 2010b). This models the effects that are found in human performance (Vickers and Packer, 1982), especially with difficult decisions, (Luce, 1986; Welford, 1980), and that have been modeled

(Carpenter et al., 2009; Moreno-Bote, 2010; Rolls et al., 2010b). There are a number of reasons why errors can occur, which affect decision times and their distributions, including a failure to maintain attention, distraction, etc., and the decision times on such trials might be short or long. Slower reaction times on error trials are usually observed in humans when there is no time pressure, while with time pressure the pattern is the opposite (Ratcliff and Rouder, 1998).

These results were obtained in a regime where the number of excitatory recurrent collateral connections per neuron was held constant (at 80 synapses for the neurons in each of the two specific pools, and a total of 800 synapses on each neuron). The reason that this number was held constant is that this is the leading term in the capacity of attractor networks (Rolls and Treves, 1998; Treves, 1991), and this is a crucial factor that affects the biological utility of attractor networks in the brain (Rolls, 2008, 2010). The interesting finding presented here is that when forming diluted connectivity by increasing the number of neurons in the network, one obtains greater stability of the network against erroneous transitions into high firing rate states, and the result may be that attractor networks can more easily be used up to capacity (set by the number of recurrent synapses per neuron) when they have diluted connectivity, because they are more stable, i.e. less noisy.

Our hypotheses for how diluted connectivity decreases the statistical fluctuations, the noise, in an attractor network are as follows. First, it is known that in a fully connected network, the noise decreases according approximately to the square root of the number of neurons in the network (Deco and Rolls, 2006; Mattia and Del Giudice, 2002, 2004; Rolls and Deco, 2010). As the number of neurons approaches infinity, the noise approaches zero, and this is the mean field limit (Rolls and Deco, 2010). The concept is that as more and more independent Poisson processes reflecting the spike times of different neurons contribute to the mean value, the mean value shows smaller and smaller statistical fluctuations. In an integrate-and-fire attractor network, the firing rate fluctuations in a single neuron make only a small contribution to the overall population noise, which provides the recurrent input back to each neuron in the network. This effect operates in our diluted connectivity network, in that as the dilution increases, so the number of neurons in a selective population or pool increases (as C , the number of recurrent connections per neuron, is held constant). Second, any correlation of the firing of the neurons in an integrate-and-fire attractor network that is related to the noise (Moreno-Bote and Parga, 2006; Renart et al., 2007) is fed back by the recurrent collaterals, and effectively amplified. This amplification of noise-related effects will be reduced by diluted connectivity, because the dilution will tend to reduce the correlation between the neuronal firing felt by a receiving neuron. It is this effect that enables diluted attractor networks to show a somewhat higher memory capacity (the number of patterns that can be stored and correctly retrieved) than fully connected networks, as we have shown (Rolls and Treves, 1998; Treves and Rolls, 1991).

We note that memory systems with a large capacity are likely to be selected for in evolution, so that animals and humans can remember for example many separate events, places, episodic memories, facts, etc. (Rolls, 2008). We believe

that neurons set the number of recurrent excitatory synapses to be rather large, up to 12,000 in the rat CA3 hippocampal network for example (Rolls, 2010; Treves and Rolls, 1994), so that the memory capacity is as large as possible. It may be that with the precision of neurons, it is difficult to utilize more than approximately 20,000 recurrent excitatory connections onto each neuron. In this situation, we argue, diluted networks in the brain are built to keep the number of synapses per neuron high, and this means that the dilution must be obtained by increasing the number of neurons in the network. This has the advantageous effect of increasing the stability of the network, and increasing its accuracy in the face of spiking-induced noise, while making the decision times only a little slower.

In addition to an effect of dilution of connectivity in increasing the stability of attractor networks as described here, there may be other advantages to dilution of connectivity in the cerebral cortex. One is that in competitive neuronal networks, which appear to be being used in the cerebral cortex to build new representations (Rolls, 2008), the diluted connectivity in the feed-forward synaptic connections may help to stabilize the ‘feature analyzer’ properties of the neurons that are built, by enabling each neuron to learn to respond to a subset of the possible input feature combinations (Rolls, 2008). This may help to break the symmetry, and may be being used in the entorhinal grid cell to hippocampal place cell transform (Rolls et al., 2006). In a fully connected competitive network there is also for example the potential problem that if the inputs are gradually drifting or changing over time, the output neurons may gradually change the feature combinations to which they respond (Carpenter, 1997; Hertz et al., 1991; Rolls, 2008). Dilution in the connectivity is likely we suggest to help minimize such drift effects, and to keep the outputs provided by cortical competitive networks constant.

It was also of interest that in the diluted case, the accuracy of the decisions was higher. The reason for this is that with less noise in diluted networks, there are fewer errors induced by the spiking-related noise producing decision transitions against the bias $\Delta\lambda$ providing the evidence for the decision (Rolls et al., 2010a,b). This is another advantage of the diluted connectivity (in the condition that the number of synapses to other neurons in the same decision pool is kept constant).

Another advantage of the diluted connectivity is that it can increase the storage capacity of autoassociation (attractor) networks (when the sparseness of the representation is not very low, when the advantage disappears) (Treves and Rolls, 1991). Another advantage of diluted connectivity is that it reduces the probability of multiple synaptic connections between pairs of neurons, which can greatly degrade the memory capacity of autoassociation networks (Rolls, 2011a).

We emphasize that the mean firing rates of the decision populations of neurons are very similar for the fully connected and diluted cases (Fig. 2). We found more within-trial variation in the spontaneous period (measured by the CV described above) in the fully connected than in the diluted case, which again is consistent with less noise in the diluted connectivity case. The implication is that another advantage of dilution in the connectivity is that once in an attractor, which might be used to implement short-term memory, the memory is more stable in the diluted connectivity case.

The decrease in the Fano factor of the firing rates shown in Fig. 6 as the simulation moves from spontaneous firing in the period time = 0–2000 ms, where the Fano factor is close to 1, to lower values when the network is being driven by inputs, is also found experimentally with neuronal activity in a number of different cortical areas (Churchland et al., 2010). Our simulation thus provides an account in an integrate-and-fire attractor decision-making network of the decrease in the variability of neuronal activity recorded in the brain as a decision is being reached (Churchland et al., 2010). It is also of interest that the coefficient of variation of the interspike interval distribution of a neuron on each trial was lower than would be expected of close to Poisson firing when the network was in a high firing rate attractor state. One factor that may contribute to the high variability in the Fano factor of the firing rate in the spontaneous period is that statistical fluctuations can be great when the numbers of spikes involved are low, as occurs with spontaneous activity (Rolls and Deco, 2010). This may help produce the variability that is reflected in a Fano factor close to 1. Another factor may be that the NMDA receptors are less engaged at low (spontaneous) firing rates than at high firing rates. A consequence of this may be that when operating with spontaneous activity, there is effectively more of a single, and short, excitatory time constant, that of the AMPA receptors (2 ms), which is close to the short time constant of the inhibitory system (in the order of 10 ms). With higher rates and more strong driving of the neurons, a greater proportion of NMDA receptor activation may increase the time constant of the system, but also introduce a long excitatory time constant to add to a short excitatory time constant (of 100 ms), and this may produce less variability of the spiking (cf. Wang, 1999). It is of interest that we found in addition that the Fano factor for the winning population of neurons is lower when $\Delta\lambda$ is larger, when the system is being driven into high firing rate states by strong inputs. The CV may also decrease in the high firing rate attractor state (Renart et al., 2007). We note that the somewhat higher values of the Fano factor found with neurons in cortical areas (Churchland et al., 2010) than those described here may be accounted for by noise that is additional to the internal spiking-related noise generated by the neurons within the network, that is if there is in addition noise in the external inputs to the network, produced for example by different amounts of arousal, attention, motivation, and signal on different trials.

In conclusion, we have investigated the effects on the dynamics of dilution of connectivity in a recurrent spiking neural network attractor model of decision-making, which applies also to memory recall. We showed that diluted connectivity (for the same number of connections per neuron and thus memory capacity) increases the stability and accuracy of the network as there is less noise in the diluted network, with little cost in increased decision times. We emphasize that it is important to understand the effects of noise in networks in the brain, and its implications for the stability of neuronal networks in the brain. For example, a stochastic neurodynamical approach to schizophrenia holds that there is less stability of cortical attractor networks involved in short-term memory and attention due to reduced functioning of the glutamate system, which decreases the firing rates of neurons in the prefrontal cortex, and therefore, given the spiking-related

noise that is present, the depth of the basins of attraction. This it is suggested contributes to the cognitive changes in schizophrenia, which include impaired short-term memory and attention (Loh et al., 2007; Rolls and Deco, 2011; Rolls et al., 2008b). In another example, a stochastic neurodynamical approach to obsessive-compulsive disorder holds that there is over-stability in some networks in the prefrontal cortex and connected areas due to hyperglutamatergia (Rolls, 2011b; Rolls et al., 2008a). In both these cases, and also in normal brain function in relation to decision-making, memory recall, etc., it is important to know to what extent noise contributed by randomness in the spiking times of individual neurons for a given mean rate contributes to stochastic effects found in the brain which affect decision-making, stability, and which may if the stability is disturbed contribute to neuropsychiatric disorders. In this context, the findings described in this paper are important for understanding normal and disordered brain function.

4. Experimental procedures

4.1. The integrate-and-fire attractor neuronal network model of decision-making

The probabilistic decision-making network we use is a spiking neuronal network model with a mean-field equivalent (Wang, 2002), but instead set to operate with parameters determined by the mean-field analysis that ensure that the spontaneous firing rate state is stable even when the decision-cues are applied, so that it is only the noise that provokes a transition to a high firing rate attractor state, allowing the effects of the noise to be clearly measured (Deco and Rolls, 2006; Rolls and Deco, 2010).

The network consists of separate populations of excitatory and inhibitory neurons as shown in Fig. 1. Two sub-populations of the excitatory neurons are referred to as decision pools, 'D1' and 'D2'. The decision pools each encode a decision to one of the stimuli, and receive as decision-related inputs λ_1 and λ_2 . The remaining excitatory neurons are called the 'non-specific' neurons, and do not respond to the decision-making stimuli used, but do allow a given sparseness of the representation of the decision-attractors to be achieved. (These neurons might in the brain respond to different stimuli, decisions, or memories.)

In our fully connected network simulations, the network contained $N=1000$ neurons, with $N_E=0.8N$ excitatory neurons, and $N_I=0.2N$ inhibitory neurons. The two decision pools are equal size sub-populations with the proportion of the excitatory neurons in a decision pool, or the sparseness of the representation with binary encoding, $f=0.1$, resulting in the number of neurons in a decision pool $N_E f=80$. The neuron pools are non-overlapping, meaning that the neurons in each pool belong to one pool only.

We structure the network by establishing the strength of interactions between pools to take values that could occur through a process of associative long-term potentiation (LTP) and long-term depression (LTD). Neurons that respond to the same stimulus, or in other words neurons that are in the same decision pool, will have stronger connections. The connection strength between neurons will be weaker if they respond to

different stimuli. The synaptic weights are set effectively by the presynaptic and post-synaptic firing rate reflecting associative connectivity (Rolls, 2008). Neurons in the same decision pool are connected to each other with a strong average weight w_+ , and are connected to neurons in the other excitatory pools with a weak average weight w_- . All other synaptic weights are set to unity. Using a mean-field analysis by Deco and Rolls (2006), we chose w_+ to be 2.1, and w_- to be 0.877 to achieve a stable spontaneous state (in the absence of noise) even when the decision cues were being applied, and stable high firing rate decision states. In particular, $w_- = \frac{0.8-f_{s1}w_+}{0.8-f_{s1}}$ (Brunel and Wang, 2001; Deco and Rolls, 2006; Loh et al., 2007; Rolls and Deco, 2010; Wang, 2002). The mean field analysis predicts the performance of a very large integrate-and-fire neuronal network (Deco and Rolls, 2006). The significance of this is that the spiking-related noise which causes statistical fluctuations in the decision populations is necessary for the transitions from the spontaneous state to a decision state when the decision cues are applied (Rolls and Deco, 2010). (The mean field analysis performed applies to a fully connected network; Rolls and Deco, 2010). The diluted network we studied operated in a very similar scenario, as shown by the properties of the integrate-and-fire network when the connectivity was diluted. We note that mean field analyses of networks with highly diluted connectivity are possible, as we have described previously (Treves, 1991; Treves and Rolls, 1991).

4.2. Neuron model

Neurons in our network use Integrate-and-Fire (IF) dynamics (Brunel and Wang, 2001; Burkitt, 2006; Deco and Rolls, 2006; Knight, 2000; Rolls and Deco, 2010; Wang, 2002) to describe the membrane potential of neurons. We chose biologically realistic constants to obtain firing rates that are comparable to experimental measurements of actual neural activity. IF neurons integrate synaptic current into a membrane potential, and then fire when the membrane potential reaches a voltage threshold. The equation that governs the membrane potential of a neuron V_i is given by:

$$C_m \frac{dV_i(t)}{dt} = -g_m(V_i(t)-V_L) - I_{syn}(t), \tag{1}$$

where C_m is the membrane capacitance, g_m is the leak conductance, V_L is the leak reversal potential, and I_{syn} is the total synaptic input. A spike is produced by a neuron when its membrane potential exceeds a threshold $V_{thr} = -50$ mV and its membrane potential is reset to a value $V_{reset} = -55$ mV. Neurons are held at V_{reset} for a refractory period τ_{rp} immediately following a spike.

4.3. Synapses

The synaptic current flowing into each neuron is described in terms of neurotransmitter components. The four families of receptors used are GABA, NMDA, AMPA_{rec}, and AMPA_{ext}. The neurotransmitters released from a presynaptic excitatory neuron act through AMPA and NMDA receptors, while inhibitory neurons activate ion channels through GABA receptors.

Each neuron in the network has $C_{ext} = 800$ external synapses that deliver input information and background spontaneous firing from other parts of the brain. Each neuron receives via each of these 800 synapses external inputs a spike train modeled by a Poisson process with rate 3.0 Hz, making the total external input 2400 Hz per neuron.

The current at each synapse is produced by synaptically activated ion channels that alter their conductances with particular time constants and that depend on the membrane potential of the neuron and the reversal potential of the currents that pass through each ion channel. These currents are summed to describe the form of the post-synaptic potentials (PSPs) at the neuron's cell body. AMPA is modeled as a fast receptor, and NMDA as a slow receptor. Synaptic current flowing into a neuron is given by the set of ordinary differential equations:

$$I_{syn}(t) = I_{GABA}(t) + I_{NMDA}(t) + I_{AMPA,rec}(t) + I_{AMPA,ext}(t)$$

in which

$$I_{AMPA,ext}(t) = g_{AMPA,ext}(V(t)-V_E) \sum_j^{C_{ext}} s_j^{AMPA,ext}(t)$$

$$I_{AMPA,rec}(t) = g_{AMPA,rec}(V(t)-V_E) \sum_j^{C_E} w_j s_j^{AMPA,rec}(t)$$

$$I_{NMDA,rec}(t) = \frac{g_{NMDA,rec}(V(t)-V_E)}{1 + e^{\beta \gamma V(t)}} \sum_j^{C_E} w_j s_j^{NMDA}(t)$$

$$I_{GABA,rec}(t) = g_{GABA}(V(t)-V_I) \sum_j^{C_I} s_j^{GABA}(t),$$

where V_E and V_I are reversal potentials for excitatory and inhibitory PSPs, the g terms represent synaptic conductances, s_j are the fractions of open synaptically activated ion channels at synapse j , and weights w_j represent the structure of the synaptic connections.

Post-synaptic potentials are generated by the opening of channels triggered by the action potential of the presynaptic neuron. As mentioned above, the dynamics of these channels are described by the gating variables s_j . The dynamics of these variables are given by:

$$\frac{ds_j^{AMPA}(t)}{dt} = -\frac{s_j^{AMPA}(t)}{\tau_{AMPA}} + \sum_k \delta(t-t_j^k)$$

$$\frac{ds_j^{NMDA}(t)}{dt} = -\frac{s_j^{NMDA}(t)}{\tau_{NMDA,decay}} + \alpha x_j(t) (1-s_j^{NMDA}(t))$$

$$\frac{dx_j(t)}{dt} = -\frac{x_j(t)}{\tau_{NMDA,rise}} + \sum_k \delta(t-t_j^k)$$

$$\frac{ds_j^{GABA}(t)}{dt} = -\frac{s_j^{GABA}(t)}{\tau_{GABA}} + \sum_k \delta(t-t_j^k)$$

where the sums over k represent a sum over spikes formulated as δ -peaks ($\delta(t)$) emitted by presynaptic neuron j at time t_j^k .

The constants used in the simulations are shown in Table 1.

4.4. Diluted connectivity

In the full connectivity case, the network we have analyzed extensively (Deco and Rolls, 2006; Rolls et al., 2010a,b) contains N neurons typically set to 500 or 1000, but systematically investigated (Deco and Rolls, 2006), with $N_E = 0.8N$ excitatory

Table 1 – Simulation constants.

Global constants			
$V_l = -70$ mV	$V_{thr} = -50$ mV	$V_{reset} = -55$ mV	$V_l = -70$ mV
$V_E = 0$ mV	$\alpha = 0.5$ ms ⁻¹	$\beta = 0.062$	$\gamma = 0.2801$
Inhibitory neuron constants			
$C_m = 0.2$ nF	$g_m = 20$ nS	$\tau_{rp} = 1$ ms	$\tau_m = 10$ ms
$g_{AMPA, ext} = 1.62$ nS	$g_{AMPA, rec} = 0.162$ nS	$g_{NMDA} = 0.516$ nS	$g_{GABA} = 1.946$ nS
$\tau_{AMPA} = 2$ ms	$\tau_{NMDA, decay} = 100$ ms	$\tau_{NMDA, rise} = 2$ ms	$\tau_{GABA} = 10$ ms
Excitatory neuron constants			
$C_m = 0.5$ nF	$g_m = 25$ nS	$\tau_{rp} = 2$ ms	$\tau_m = 20$ ms
$g_{AMPA, ext} = 2.08$ nS	$g_{AMPA, rec} = 0.208$ nS	$g_{NMDA} = 0.654$ nS	$g_{GABA} = 2.5$ nS
$\tau_{AMPA} = 2$ ms	$\tau_{NMDA, decay} = 100$ ms	$\tau_{NMDA, rise} = 2$ ms	$\tau_{GABA} = 10$ ms

neurons, and $N_i = 0.2N$ inhibitory neurons. The two decision pools are equal in size of sub-populations with the proportion of the excitatory neurons in a decision pool, or the sparseness of the representation with binary encoding, $f = 0.1$, resulting in the number of neurons in a decision pool N_{df} . The neuron pools are non-overlapping, meaning that the neurons in each pool belong to one pool only. In the investigations described here, we chose to set the number of excitatory recurrent collateral connections C per neuron to 80 in a decision-making pool, the number in a standard fully connected network with $N_E = 800$ excitatory neurons, as with a network this size there is still with a standard set of parameters some instability of the spontaneous firing rate state, so that stability can be investigated. The number of non-specific excitatory neurons in this network with $N = 1000$ is 640. These neurons simulate the effects of noise from other networks in the system than those involved in the decision-making. The sparseness of the representation, the proportion of excitatory neurons with high activity when an attractor wins, is $80/800 = 0.1$ for the fully connected network. Given that there were two decision pools each with 80 neurons in the fully connected network, there were 80 connections onto each neuron from each decision pool.

In setting up the diluted connectivity, we aimed to keep the total number of excitatory synapses devoted to the recurrent collateral synapses fixed, as this is likely to be a biological parameter that cannot be increased without limit, and that number was 800, as noted above. The number of synapses for each decision (or specific) pool was also maintained at the default value for this size of network, 80 synapses per neuron. We kept the number of inhibitory neurons constant at 200 for this size of network. To dilute the

connectivity within each specific decision-making pool of neurons to 0.25, we increased the number of neurons in each specific pool from 80 in the fully connected case to 320 in the diluted case. The 80 synapses onto each neuron were randomly chosen from the 320 neurons in each decision pool to achieve a dilution of the connectivity in each decision pool of 0.25.

This meant that the population sparseness was 0.25 (320 neurons/1280 excitatory neurons) for the representation for each specific decision-making pool for the dilution to 0.25 case. Considering each synaptic weight vector on each excitatory neuron, the number of synapses for a specific decision-making pool (e.g. pool 1) was 800, and of these 80 were intrapool synapses for decision 1, so the sparseness of a representation seen by each neuron for decision 1 was held constant at $80/800 = 0.1$, independently of the degree of dilution. (The same was true for the other decision pool, pool 2.) Given that there were two decision pools each with 80 synapses per neuron in a decision pool, there were 80 connections onto each neuron from each decision pool just as in the fully connected case above. The numbers are summarized in Table 2. In all cases, the number of non-specific neurons, i.e. those not in a specific decision-making pool, was kept constant at 640. The total number of excitatory neurons in each network is shown by N_E in Table 2.

The above values for the parameters were for connectivity diluted to 0.25. We also ran further simulations with the connectivity diluted to 0.1. The parameters in this case were as shown in Table 2. In the Connectivity column, 1 refers to full connectivity, and 0.25 to a network with the connectivity diluted to 0.25 using an algorithm which ensured that for each of the 80 synapses on a neuron for a given decision-

Table 2 – Parameters for the full and diluted connectivity simulations C = the number of excitatory connections received by an excitatory neuron. $C_{fromDpool}$ = the number of excitatory connections received from each decision-making pool. N_{Dpool} = the number of excitatory neurons in a decision-making pool. N_E = the total number of excitatory neurons in the simulations, in which the number of nonspecific neurons was 640. a = the population sparseness of the representation, the proportion of excitatory neurons active for any one decision.

Connectivity	C	$C_{fromDpool}$	N_{Dpool}	N_E	Sparseness a
1	800	80	80	800	0.1
0.25	800	80	320	1280	0.25
0.1	800	80	800	2240	0.357

making pool, the neuron in the decision pool to which to connect was chosen randomly, with no double connections allowed to a neuron.

4.5. Sparseness

The population sparseness a of a binary representation is the proportion of neurons active to represent any one stimulus or decision in the set. The sparseness can be generalized to graded representations as shown in Eq. (2):

$$a = \frac{\left(\sum_i^{N_E} r_i\right)^2}{\sum_i^{N_E} r_i^2}, \quad (2)$$

where r_i is the firing rate measured for neuron i in the population of N_E excitatory neurons in the network (Franco et al., 2007; Rolls, 2008; Rolls and Treves, 1990; Treves and Rolls, 1991). We note that this is the sparseness of the representation measured for any one stimulus over the population of excitatory neurons (Franco et al., 2007; Rolls, 2008). For the sparseness values shown in this paper, the population sparseness on each trial was calculated using the time-averaged firing rate of all excitatory neurons, and then this population sparseness was averaged across trials.

4.6. Simulation regime

The network was simulated numerically using a second order Runge–Kutta algorithm step with an integration step $dt=0.02$ ms for a time period of 4 s. First there was a 2 s baseline period of spontaneous activity in which $\lambda_1=3.0$ Hz for all external synapses onto every neuron in pool i , where the pools were inhibitory, decision-making pool 1, decision-making pool 2, and the non-specific excitatory pool. There was then a 2 s decision period in which the decision stimuli were applied by increasing the firing rates for the 800 external input synapses on each of the neurons in the two decision pools so that the mean of λ_1 and of $\lambda_2=3.04$ Hz per synapse (an extra 32 Hz per neuron, given the 800 external synapses onto each neuron). During the decision period, the noise in the network, and the increased firing rate bias applied as a decision cue to each decision pool of neurons, cause one of the decision populations of neurons to jump to a high firing rate attractor state with the assistance of the positive feedback in the recurrent collaterals, and this high firing inhibits through the inhibitory interneurons the other decision population of neurons. There is thus a single winning population on each trial, and which of the two populations wins on a particular trial is determined by the statistical fluctuations in the firing rates of the neurons in each decision population, and the difference in the two inputs λ_1 and λ_2 , i.e. $\Delta\lambda$. In the simulations we describe, unless otherwise stated, $\Delta\lambda=6.4$ Hz per neuron, which corresponds to a firing rate of 3.044 Hz per synapse being applied as λ_1 and 3.036 Hz per synapse being applied as λ_2 as the external inputs to the two decision-making pools during the decision-making period. With $\Delta\lambda=6.4$ Hz, $\Delta\lambda/\lambda$, the Weber fraction (Deco and Rolls, 2006), $=6.4/2432=0.00263$.

Acknowledgments

The research was supported by the Oxford Centre for Computational Neuroscience. For advice and helpful discussions we thank Professor Alessandro Treves of SISSA, Trieste, Italy, and Professor J. Feng of the Department of Computer Science, University of Warwick, UK.

REFERENCES

- Abeles, M., 1991. *Corticonics: Neural Circuits of the Cerebral Cortex*. Cambridge University Press, Cambridge.
- Amit, D.J., Brunel, N., 1997. Dynamics of a recurrent network of spiking neurons before and following learning. *Network* 8, 373–404.
- Beck, J.M., Ma, W.J., Kiani, R., Hanks, T., Churchland, A.K., Roitman, J., Shadlen, M.N., Latham, P.E., Pouget, A., 2008. Probabilistic population codes for Bayesian decision making. *Neuron* 60, 1142–1152.
- Bovier, A., Gayrard, V., 1992. Rigorous bounds on the storage capacity of the dilute Hopfield model. *J. Stat. Phys.* 69, 597–627.
- Braitenberg, V., Schütz, A., 1991. *Anatomy of the Cortex*. Springer-Verlag, Berlin.
- Braun, J., Mattia, M., 2010. Attractors and noise: twin drivers of decisions and multistability. *Neuroimage* 52, 740–751.
- Bressloff, P.C., 2010. Metastable states and quasicycles in a stochastic Wilson–Cowan model of neuronal population dynamics. *Phys. Rev. E Stat. Nonlin. Soft Matter Phys.* 82, 051903.
- Brunel, N., Wang, X.J., 2001. Effects of neuromodulation in a cortical network model of object working memory dominated by recurrent inhibition. *J. Comput. Neurosci.* 11, 63–85.
- Burkitt, A.N., 2006. A review of the integrate-and-fire neuron model: I. Homogeneous synaptic input. *Biol. Cybern.* 95, 1–19.
- Carpenter, G.A., 1997. Distributed learning, recognition and prediction by ART and ARTMAP neural networks. *Neural Netw.* 10 (8), 1473–1494.
- Carpenter, R.H., Reddi, B.A., Anderson, A.J., 2009. A simple two-stage model predicts response time distributions. *J. Physiol.* 587, 4051–4062.
- Churchland, M.M., Yu, B.M., Cunningham, J.P., Sugrue, L.P., Cohen, L.R., Corrado, G.S., Newsome, W.T., Clark, A.M., Hosseini, P., Scott, B.B., Bradley, D.C., Smith, M.A., Kohn, A., Movshon, A., Armstrong, K.M., Moore, T., Chang, S.W., Snyder, L.H., Lisberger, S.G., Priebe, N.J., Finn, I.M., Ferster, D., Ryu, S.I., Santhanam, G., Sahani, M., Shenoy, K.V., 2010. Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nat. Neurosci.* 13, 369–378.
- Churchland, A.K., Kiani, R., Chaudhuri, R., Wang, X.J., Pouget, A., Shadlen, M.N., 2011. Variance as a signature of neural computations during decision making. *Neuron* 69, 818–831.
- Deco, G., Rolls, E.T., 2006. A neurophysiological model of decision-making and Weber's law. *Eur. J. Neurosci.* 24, 901–916.
- Deco, G., Scarano, L., Soto-Faraco, S., 2007. Weber's law in decision making: integrating behavioral data in humans with a neurophysiological model. *J. Neurosci.* 27, 11192–11200.
- Deco, G., Rolls, E.T., Romo, R., 2009. Stochastic dynamics as a principle of brain function. *Prog. Neurobiol.* 88, 1–16.
- Ditterich, J., 2006. Evidence for time-variant decision making. *Eur. J. Neurosci.* 24, 3628–3641.
- Elston, G.N., Benavides-Piccione, R., Elston, A., Zietsch, B., Defelipe, J., Manger, P., Casagrande, V., Kaas, J.H., 2006. Specializations of the granular prefrontal cortex of primates: implications for cognitive processing. *Anat. Rec. A Discov. Mol. Cell. Evol. Biol.* 288, 26–35.

- Faisal, A., Selen, L., Wolpert, D., 2008. Noise in the nervous system. *Nat. Rev. Neurosci.* 9, 292–303.
- Franco, L., Rolls, E.T., Aggelopoulos, N.C., Jerez, J.M., 2007. Neuronal selectivity, population sparseness, and ergodicity in the inferior temporal visual cortex. *Biol. Cybern.* 96, 547–560.
- Gigante, G., Mattia, M., Braun, J., Del Giudice, P., 2009. Bistable perception modeled as competing stochastic integrations at two levels. *PLoS Comput. Biol.* 5, e1000430.
- Hertz, J., Krogh, A., Palmer, R.G., 1991. *Introduction to the Theory of Neural Computation*. Addison Wesley, Wokingham, U.K.
- Hopfield, J.J., 1982. Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci. U.S.A.* 79, 2554–2558.
- Knight, B., 2000. Dynamics of encoding in neuron populations: some general mathematical features. *Neural Comput.* 12, 473–518.
- Loh, M., Rolls, E.T., Deco, G., 2007. A dynamical systems hypothesis of schizophrenia. *PLoS Comput. Biol.* 3, e228. doi:10.1371/journal.pcbi.0030228.
- Luce, R.D., 1986. *Response Time: Their Role in Inferring Elementary Mental Organization*. Oxford University Press, New York.
- Marti, D., Deco, G., Mattia, M., Gigante, G., Del Giudice, P., 2008. A fluctuation-driven mechanism for slow decision processes in reverberant networks. *PLoS One* 3, e2534. doi:10.1371/journal.pone.0002534.
- Mattia, M., Del Giudice, P., 2002. Population dynamics of interacting spiking neurons. *Phys. Rev. E* 66, 051917.
- Mattia, M., Del Giudice, P., 2004. Finite-size dynamics of inhibitory and excitatory interacting spiking neurons. *Phys. Rev. E* 70, 052903.
- Miller, P., 2006. Analysis of spike statistics in neuronal systems with continuous attractors or multiple, discrete attractor states. *Neural Comput.* 18, 1268–1317.
- Miller, P., Katz, D.B., 2010. Stochastic transitions between neural states in taste processing and decision-making. *J. Neurosci.* 30, 2559–2570.
- Miller, P., Wang, X.J., 2006. Power-law neuronal fluctuations in a recurrent network model of parametric working memory. *J. Neurophysiol.* 95, 1099–1114.
- Moreno-Bote, R., 2010. Decision confidence and uncertainty in diffusion models with partially correlated neuronal integrators. *Neural Comput.* 22, 1786–1811.
- Moreno-Bote, R., Parga, N., 2006. Auto- and crosscorrelograms for the spike response of leaky integrate-and-fire neurons with slow synapses. *Phys. Rev. Lett.* 96, 028101.
- Palmer, J., Huk, A.C., Shadlen, M.N., 2005. The effect of stimulus strength on the speed and accuracy of a perceptual decision. *J. Vis.* 5, 376–404.
- Perez Castillo, I., Skantzos, N.S., 2004. The Little–Hopfield model on a sparse random graph. *J. Phys. A Math. Gen.* 37, 9087–9099.
- Perin, R., Berger, T.K., Markram, H., 2011. A synaptic organizing principle for cortical neuronal groups. *Proc. Natl. Acad. Sci. U.S.A.* 108, 5419–5424.
- Ratcliff, R., McKoon, G., 2008. The diffusion decision model: theory and data for two-choice decision tasks. *Neural Comput.* 20, 873–922.
- Ratcliff, R., Rouder, J.F., 1998. Modeling response times for two-choice decisions. *Psychol. Sci.* 9, 347–356.
- Ratcliff, R., Zandt, T.V., McKoon, G., 1999. Connectionist and diffusion models of reaction time. *Psychol. Rev.* 106, 261–300.
- Renart, A., Moreno-Bote, R., Wang, X.J., Parga, N., 2007. Mean-driven and fluctuation-driven persistent activity in recurrent networks. *Neural Comput.* 19, 1–46.
- Resulaj, A., Kiani, R., Wolpert, D.M., Shadlen, M.N., 2009. Changes of mind in decision-making. *Nature* 461, 263–266.
- Ribault, C., Sekimoto, K., Triller, A., 2011. From the stochasticity of molecular processes to the variability of synaptic transmission. *Nat. Rev. Neurosci.* 12, 375–387.
- Roitman, J.D., Shadlen, M.N., 2002. Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J. Neurosci.* 22, 9475–9489.
- Rolls, E.T., 2005. *Emotion Explained*. Oxford University Press, Oxford.
- Rolls, E.T., 2008. *Memory, Attention, and Decision-Making. A Unifying Computational Neuroscience Approach*. Oxford University Press, Oxford.
- Rolls, E.T., 2010. A computational theory of episodic memory formation in the hippocampus. *Behav. Brain Res.* 215, 180–196.
- Rolls, E.T., 2011a. Diluted Attractor Networks in the Brain.
- Rolls, E.T., 2011b. Glutamate, obsessive-compulsive disorder, schizophrenia, and the stability of cortical attractor neuronal networks. *Pharmacol. Biochem. Behav.* Epub 23 June.
- Rolls, E.T., Deco, G., 2010. *The Noisy Brain: Stochastic Dynamics as a Principle of Brain Function*. Oxford University Press, Oxford.
- Rolls, E.T., Deco, G., 2011. A computational neuroscience approach to schizophrenia and its onset. *Neurosci. Biobehav. Rev.* 35, 1644–1653.
- Rolls, E.T., Tovee, M.J., 1995. Sparseness of the neuronal representation of stimuli in the primate temporal visual cortex. *J. Neurophysiol.* 73, 713–726.
- Rolls, E.T., Treves, A., 1990. The relative advantages of sparse versus distributed encoding for associative neuronal networks in the brain. *Network* 1, 407–421.
- Rolls, E.T., Treves, A., 1998. *Neural Networks and Brain Function*. Oxford University Press, Oxford.
- Rolls, E.T., Treves, A., in press. Information encoding in the brain. *Progress in Neurobiology*.
- Rolls, E.T., Treves, A., Foster, D., Perez-Vicente, C., 1997a. Simulation studies of the CA3 hippocampal subfield modelled as an attractor neural network. *Neural Netw.* 10, 1559–1569.
- Rolls, E.T., Treves, A., Tovee, M., Panzeri, S., 1997b. Information in the neuronal representation of individual stimuli in the primate temporal visual cortex. *J. Comput. Neurosci.* 4, 309–333.
- Rolls, E.T., Stringer, S.M., Elliot, T., 2006. Entorhinal cortex grid cells can map to hippocampal place cells by competitive learning. *Netw. Comput. Neural Syst.* 17, 447–465.
- Rolls, E.T., Loh, M., Deco, G., 2008a. An attractor hypothesis of obsessive-compulsive disorder. *Eur. J. Neurosci.* 28, 782–793.
- Rolls, E.T., Loh, M., Deco, G., Winterer, G., 2008b. Computational models of schizophrenia and dopamine modulation in the prefrontal cortex. *Nat. Rev. Neurosci.* 9, 696–709.
- Rolls, E.T., Grabenhorst, F., Deco, G., 2010a. Choice, difficulty, and confidence in the brain. *Neuroimage* 53, 694–706.
- Rolls, E.T., Grabenhorst, F., Deco, G., 2010b. Decision-making, errors, and confidence in the brain. *J. Neurophysiol.* 104, 2359–2374.
- Sakai, Y., Okamoto, H., Fukai, T., 2006. Computational algorithms and neuronal network models underlying decision processes. *Neural Netw.* 19, 1091–1105.
- Smith, P., Ratcliff, R., 2004. Psychology and neurobiology of simple decisions. *Trends Neurosci.* 23, 161–168.
- Softky, W.R., Koch, C., 1993. The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J. Neurosci.* 13, 334–350.
- Treves, A., 1991. Dilution and sparse coding in threshold-linear nets. *J. Phys. A Math. Gen.* 24 (1), 327–335.
- Treves, A., Rolls, E.T., 1991. What determines the capacity of autoassociative memories in the brain? *Network* 2, 371–397.
- Treves, A., Rolls, E.T., 1994. A computational analysis of the role of the hippocampus in memory. *Hippocampus* 4, 374–391.
- Treves, A., Panzeri, S., Rolls, E.T., Booth, M., Wakeman, E.A., 1999. Firing rate distributions and efficiency of information transmission of inferior temporal cortex neurons to natural visual stimuli. *Neural Comput.* 11, 601–631.
- Vickers, D., Packer, J., 1982. Effects of alternating set for speed or accuracy on response time, accuracy and confidence in a

- unidimensional discrimination task. *Acta Psychol.* 50, 179–197.
- Wang, X.J., 1999. Synaptic basis of cortical persistent activity: the importance of NMDA receptors to working memory. *J. Neurosci.* 19, 9587–9603.
- Wang, X.J., 2002. Probabilistic decision making by slow reverberation in cortical circuits. *Neuron* 36, 955–968.
- Wang, X.J., 2008. Decision making in recurrent neuronal circuits. *Neuron* 60, 215–234.
- Webb, T.J., Rolls, E.T., Deco, G., Feng, J., in press. Noise in attractor networks in the brain produced by graded firing rate representations. *PLoS One*.
- Welford, A.T. (Ed.), 1980. *Reaction Times*. Academic Press, London.
- Wong, K.F., Huk, A.C., 2008. Temporal dynamics underlying perceptual decision making: insights from the interplay between an attractor model and parietal neurophysiology. *Front. Neurosci.* 2, 245–254.
- Wong, K., Wang, X.J., 2006. A recurrent network mechanism of time integration in perceptual decisions. *J. Neurosci.* 26, 1314–1328.
- Wong, K.F., Huk, A.C., Shadlen, M.N., Wang, X.J., 2007. Neural circuit dynamics underlying accumulation of time-varying evidence during perceptual decision making. *Front. Comput. Neurosci.* 1, 6.