Negatively-Correlated Firing: The Functional Meaning of Lateral Inhibition Within Cortical Columns

Simon Durrant[†]   Jianfeng Feng[‡]

[†]Department of Informatics, Sussex University

Brighton BN1 9QH, UK

[‡]Department of Computer Science, Warwick University

Coventry CV4 7AL, UK

Abstract

Lateral inhibition is a well documented aspect of neural architecture in the main sensory systems. Existing accounts of lateral inhibition focus on its role in sharpening distinctions between inputs that are closely related. However, these accounts fail to explain the functional role of inhibition in cortical columns, such as those in V1, where neurons have similar response properties. In this paper, we outline a model of position tracking using cortical columns of integrate-and-fire neurons which respond optimally to a particular location, to show that negatively-correlated firing patterns arise from lateral inhibition in cortical columns, and that this provides a clear benefit for population coding in terms of stability, accuracy, estimation time and neural resources.

Keywords: inhibition; population coding; cortical column; position tracking;

1

1 Overview

1.1 Introduction

Understanding the functional meaning of particular aspects of neural architecture is a central objective of computational neuroscience. Inhibitory interneurons are very common in the neocortex, and lateral inhibition has been shown to play an important role in sharpening the distinctions between similar inputs, where such inputs would otherwise invoke nearly the same response in neurons with only slightly different response properties. However, it is our belief that inhibition also plays an important role in population coding, stabilising the mean field potential and greatly improving the ability of a group of neurons to accurately represent a given stimulus, by creating negatively-correlated firing patterns. In this paper, we will describe the principle on which this improvement is based, showing how it extends the benefit of population coding. We will then present a model and a set of experiments using this model, which demonstrates that pools of neurons operating with inhibitory connections perform better on a stimulus-tracking task than the same model with no inhibitory connections. The model is designed to show how this can work in principle in neural sensory coding, incorporating a number of realistic aspects such as the use of spiking neurons, online estimation and a simple filtering mechanism, but it is not intended to be a biophysically realistic model, which would reduce the clarity of this initial demonstration. We will, however, include a brief discussion on how our model is related to the circuits found in the primary visual cortex.

1.2 Inhibitory mechanisms in neocortex

Sensory neural processing is most often thought of in terms of many interconnected circuits using excitatory connections to propagate signals through layers of increasingly abstract representation. Whilst there is some truth in this necessarily simplified image, it is also the case that half of all neurons

in the brain are inhibitory. Inhibitory neurons can operate in a variety of different ways, such as feedforward inhibition, where the inhibitory connection comes from the same area that gives the excitatory connections to other neurons in the current area, and feedback inhibition where the excitatory neurons in the current area in turn suppress the activity of other neurons through the use of local inhibitory interneurons. Much of the existing work on inhibitory mechanisms in the context of cortical columns has focused on the role of lateral inhibition in sharpening distinctions between neurons with slightly different response properties [8, 10]. However, the existence of local inhibitory circuits [12, 13] and the fact that horizontal connections lead to inhibitory post-synaptic potentials (IPSPs) [6], suggest that inhibitory circuits are likely to operate within cortical columns, as well as across them. Clearly, inhibitory connections within cortical columns do not have the same role as those that operate across different columns. Lateral inhibition within columns means neurons with the same response properties inhibiting each other. What could be the functional role for this local inhibitory mechanism? In this paper, we propose that this local inhibition improves the performance of pooled neurons, by exploiting a property of the central limit theorem applicable to population coding.

1.3   Population Coding and the Central Limit Theorem

Many areas of the neocortex show a columnar structure, in which all of the neurons in a given layer within the column (sometimes called a pool of neurons) have essentially the same response properties. Somatosensory cortex and primary visual cortex (V1) are two prominent examples of this. Cortical columns are an important structural unit in the brain, and operate on the principle of population coding, where the activity of the pool as a whole is taken to be the signal, rather than the firing rates of the individual neurons. This mean field potential can be formally defined as [5]:-
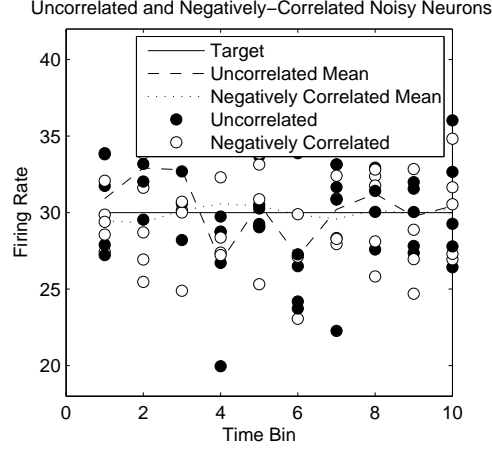
Figure 1: The central limit effect for noisy neurons. The underlying signal value for neural firing to convey is shown, along with the mean (pooled) activity for a set of five uncorrelated neurons and a set of five negatively correlated neurons. The individual neuron firing rates are also shown (circles). It can be seen that when the neurons are negatively correlated, the result is that the mean field potential gives a more reliable representation of the underlying signal.

$$A(t) = \lim_{\Delta t \to 0} \frac{1}{\Delta t} \frac{\text{Spikes}_N(t, t + \Delta t)}{N} = \frac{1}{N} \sum_{i=1}^{N} \sum_{s} \delta(t - t_i^s) \qquad (1)$$

The advantage of the this scheme for neural processing is that it allows variations in firing rates of individual neurons, due to extrinsic or intrinsic noise, whilst maintaining the signal-to-noise ration (SNR) through the cancellation of the noise when the firing rates are pooled together. This is an effect of the central limit theorem, and can be seen in figure 1. When a group of random elements, for example neural firing rates, are summed together, they give a mean value (mean field potential). For a system in which the random elements consist of signal plus zero-mean gaussian noise, which is a typical assumption for most noisy systems including neurons, the mean value will tend to be closer to the signal value because the noise cancels out (central limit effect). More elements pooled together results in greater noise cancel-

4

lation. There is, however, another way to accelerate the central limit effect, and this is to have a noise component that is negatively-correlated. It can be seen in figure 1 that negatively-correlated noise cancels out much more quickly and effectively than independent noise, because corresponding elements tend to be on opposite sides of the mean (the space-filling effect of negative correlation). In order to take advantage of this fact, however, a system has to have a means to influence the correlation of the noise component, or at least the effect of the noise component.

In neural systems (and more generally in all threshold systems), a centrally important fact is that a neuron which spikes will be on average more likely to have a positive noise component than a negative one (because positive noise components lead to an increased membrane potential, which in turn increases the probability of spiking). In order to simulate the effect of negatively-correlated noise with the other neurons (without actually effecting the extrinsic or intrinsic noise directly, which is by definition beyond control), this neuron must reduce the membrane potential of the other neurons in the pool. This can be achieved through the inhibitory connections identified in section 1.2, and the result is that neurons within a pool should have firing patterns that are negatively correlated with each other. We believe that this will result in a more stable system, with a mean field potential closer to that of the underlying signal, which will give an improved performance on whatever task the neural pool undertakes. In the following sections, we test this hypothesis on the task of tracking a moving stimulus using groups of pooled neurons.

2 Methods

2.1 Network Model

Our model consists of a set of 10 cortical columns, each containing 100 integrate-and-fire neurons, which are either unconnected (control condition;

an effective synaptic strength of zero), or have a full set of inhibitory connections within each column, with a synaptic strength of -1 to all other neurons in the column (there were no recurrent connections back to the originating neuron itself). There are no connections across columns. Each neuron is parameterised by v_{thre} (firing threshold), v_{rest} (the resting and reset potential), γ (the membrane time constant), and x_{col} (maximal response stimulus position, which is the same for all neurons within a column). In the experiments presented here, v_{rest} was always set to 0 (using a standard rescaling approach from realistic values, for ease of analysis) and γ was always set to 20. The value of v_{thre} was varied in the experiments (details are given in sections 2.4 and 3.3). x_{col} took evenly spaced values between 0 and 10 (inclusive), which were chosen to coincide with the limits of the stimulus range (in order that the model could handle all stimulus positions, but also that the full range of stimulus positions that could be handled by the model were available to be tested).

At each simulation step, external inputs (determined by stimulus position and a random element; see section 2.3) were presented to each of the neurons, and the membrane potential v updated according to the following leaky integrator equation:-

$$\frac{dv(t)}{dt} = -\frac{v(t)}{\gamma} + I_E(t) + I_S(t) \tag{2}$$

Here, $I_E(t)$ is the external input from the stimulus, and $I_S(t)$ is the spiking input from other neurons; in the control condition, $I_S(t)$ will always be zero. In order to account for shunting inhibition (where the hyperpolarising effect of inhibitory input is related to the degree of depolarisation caused by excitatory inputs); see [1]), we have also included a half-wave rectification which

6

ensures that neurons do not become hyperpolarised beyond resting potential. Neurons which reached the threshold level v_{thre} produced a spike and had their membrane potential set to v_{rest} for the remainder of the simulation step.

2.2 Decoding Strategy

In order to effectively evaluate the performance of a network model on a given task, it is necessary to be able to convert the network's output back into the domain of the task. In the case of our model using spiking neurons to perform the tracking task, this means we have to address the issue of neural decoding. In their model which performed the same tracking task, [11] explored two statistical decoding methods, one based on moment estimate, and an optimal unbiased strategy using a censored maximum likelihood approach. The latter of these provides a useful benchmark on the best decoding performance possible given a set of particular conditions (including balanced inputs). For our model, however, we prefer to use a decoding strategy that can be interpreted in terms of a neural model. There are two aspects to our decoding strategy:-

1. All of the spikes in the estimation window for each column are simply summed and normalised, and combined with their respective column position.

2. The resultant set of outputs (one for each column) are filtered using a simple gaussian filter. The stimulus position is estimated to be the mean of the filtered output.

The first stage represents the projection of activity from the columns to the next stage of cortical processing. The second stage represents a relative heightening of the more active outputs compared to the other outputs. Lateral inhibition across columns has been widely interpreted as having this

effect; feedback connections, which are prominent between V1 and the LGN, could also act as this sort of filter. The filter is set up in advance, and is identical for all experiments and conditions outlined in the paper. Its equation is as follows:-

$$\frac{1}{\sqrt{2\pi\sigma^2}} \exp(\frac{-(x_{dis} - x_{win})^2}{2\sigma^2}) \tag{3}$$

x_{dis} is the distance from the column with the strongest response, x_{pos}; σ is the standard deviation of the filter, which was fixed at a value of 1.5, determined by numerical experiments to give a reliable performance). It can be seen that both stages in our decoding strategy have biological plausibility and do not rely upon any abstract, complex statistical calculation. Also in keeping with biological plausibility, estimation was performed online, with a position estimate given at each simulation step, rather than simply at the end of an estimation period or after the simulation is complete. This was achieved using a sliding estimation window representing the memory of the estimation system. A rectangular sliding window was used in the experiments presented here, although alternative schemes can also be implemented in the model. The size of the estimation window is a parameter that was varied during the experiments (details given in the next section), in order to assess the interaction of estimation time and inhibition.

2.3 Input Stimuli

In order to facilitate comparison with the optimal statistical approach outlined in [11], we adopt the same approach to input stimuli. In the experiments presented in the next section, the stimulus position was either held constant throughout, or moved instantaneously every 100ms, creating a steplike signal.

As the former case can be contained within the latter allowing the updated position to be the same at each step, the stimulus position is given as [11]:-

$$x(t) = \sum_k \xi_k \chi(t \in \{(k-1)T_W, kT_W\}) \tag{4}$$

Here, ξ_k are independent, uniformly-distributed random variables, between limits $[X, Y]$. For a constant input position, $X = Y = Position$. For the step inputs, the limits are $[0, L]$, where L is the highest maximal-response position for any neuron in the network model. Given the stimulus position, $x(t)$, a gaussian input rate λ is created for each neuron:-

$$\lambda(t) = \lambda_{core} + c\lambda_{core} \exp(-\frac{(x(t) - x_{pos})^2}{2\sigma^2}) \tag{5}$$

There are three constants in this stimulus creation. σ is the spatial resolution of each neuron, and was set to 1 for all neurons across all experiments. λ_{core} provides both a constant input component, and effectively scales the random element in the balanced input; this was fixed at 3 for all of the experiments. Finally, c scales the stimulus intensity, effectively controlling the extent to which the input is position dependent; this was fixed to 10 for all of the experiments. Given the input parameter $\lambda(t)$, the external input to each neuron is given by a gaussian process using a balanced input [2]:-

$$I_E(t) = \mu(\lambda(t)) + \sigma(\lambda(t))B \tag{6}$$

$$\mu(\lambda(t)) = a\lambda(t)(1 - r) \tag{7}$$

$$\sigma(\lambda(t)) = a\sqrt{\lambda(t)(1+r)} \tag{8}$$

$$r(\lambda(t)) = 1 - \frac{v_{thre}}{\lambda(t)a\gamma} \tag{9}$$

B is a standard Brownian motion, providing a random element in the input, $r(\lambda(t))$ provides the balancing condition, a is the magnitude of the postsynaptic potentials (set to 0.5 for all experiments here), γ is the membrane time constant for the network neurons, v_{thre} is the firing threshold for the network neurons. It is important to understand the dynamics of the input behaviour. Given the balanced condition, μ provides a constant input to all neurons irrespective of stimulus position. σ effectively sets the variance of the Brownian motion according to how close the stimulus is to the column centre of a given neuron. The result of this is that the magnitude of the input to neurons for which the stimulus is closer to their maximal response position will be greater than that for other neurons. The mean input will remain the same because the Brownian motion remains centred at zero, but the effect of stochastic resonance [4] is that neurons with a greater magnitude will typically have a greater firing rate. In this way, neurons can receive balanced inputs while responding with different firing rates.

2.4 Experiments

A number of experiments have been conducted, to examine the effect of inhibitory connections in the network on performance on the stimulus tracking task. Basic performance was measured using the mean squared error (MSE), of the position estimate and the actual stimulus position, at the end of each estimation window period. Although an online sliding estimation window was used to produce a continual position estimate, in calculating the MSE we use the position after a period of time equivalent to a complete estimation window has passed in order to ensure that there are no contamination effects

from the previous stimulus position remaining in the estimation window. In addition to examining basic performance, correlation measures (both across and within columns, and estimate autocorrelation), and firing rates have been shown where appropriate. We also vary a number of other parameters in the experiments, including:-

- The stimulus type: both constant and steplike signals are used.

- The size of the estimation window: values from 100ms right down to just 10ms are tested.

- The number of neurons in each column: set to 100 by default, but varied in one experiment to examine the central limit effect of neuron numbers.

- The firing threshold V_{thre}: set to 5 by default for the network with lateral inhibitory input, and 20 for the network without, to ensure similar firing rates and therefore a balanced input for both, but varied in one experiment to have the same threshold value of 20 for each.

- The inhibitory connection strength (set to -1 by default, but all values from 0 to -1 tested in one experiment).

The results of all of these experiments are presented in the next section.

3   Results

3.1   Basic Performance

The performance of the network was first tested on a stimulus with a constant position of 4.5. It can be seen in figure 2 that when lateral inhibitory connections are used, the MSE is much lower, and the network's online position estimate is much more stable. As the estimation window size becomes
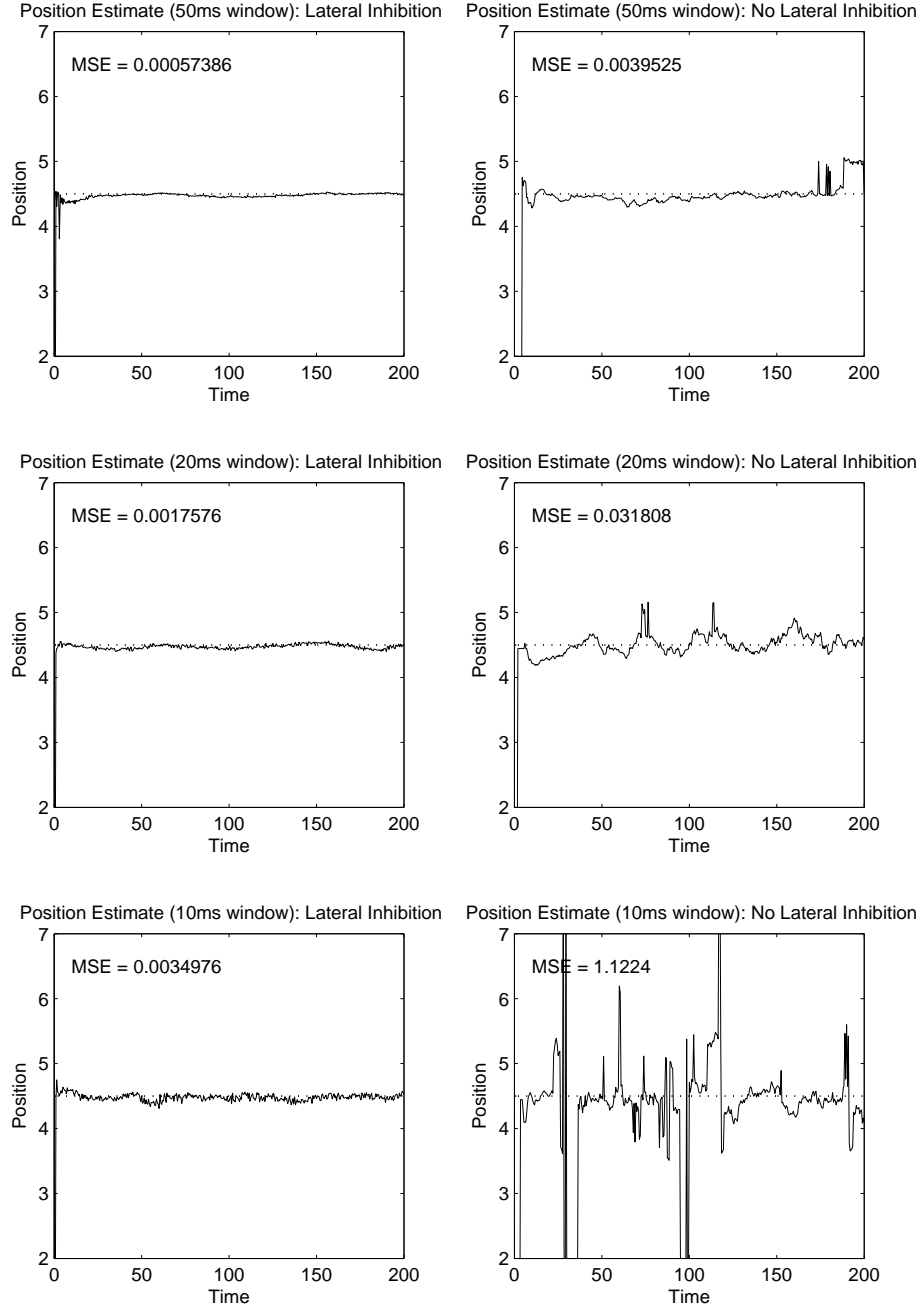
Figure 2: Position targets (dotted line) and estimates (solid line) for a constant signal, across 50ms, 20ms and 10ms estimate window sizes (shown here in different rows). *Left column*: Lateral inhibition used. *Right column*: Lateral inhibition not used. It can be seen that the MSE is lower when lateral inhibition is present, across all window sizes. The greater reliability of the position estimate when neurons have inhibitory connections is particular apparent for smaller estimation window sizes, when the estimation task becomes more difficult.
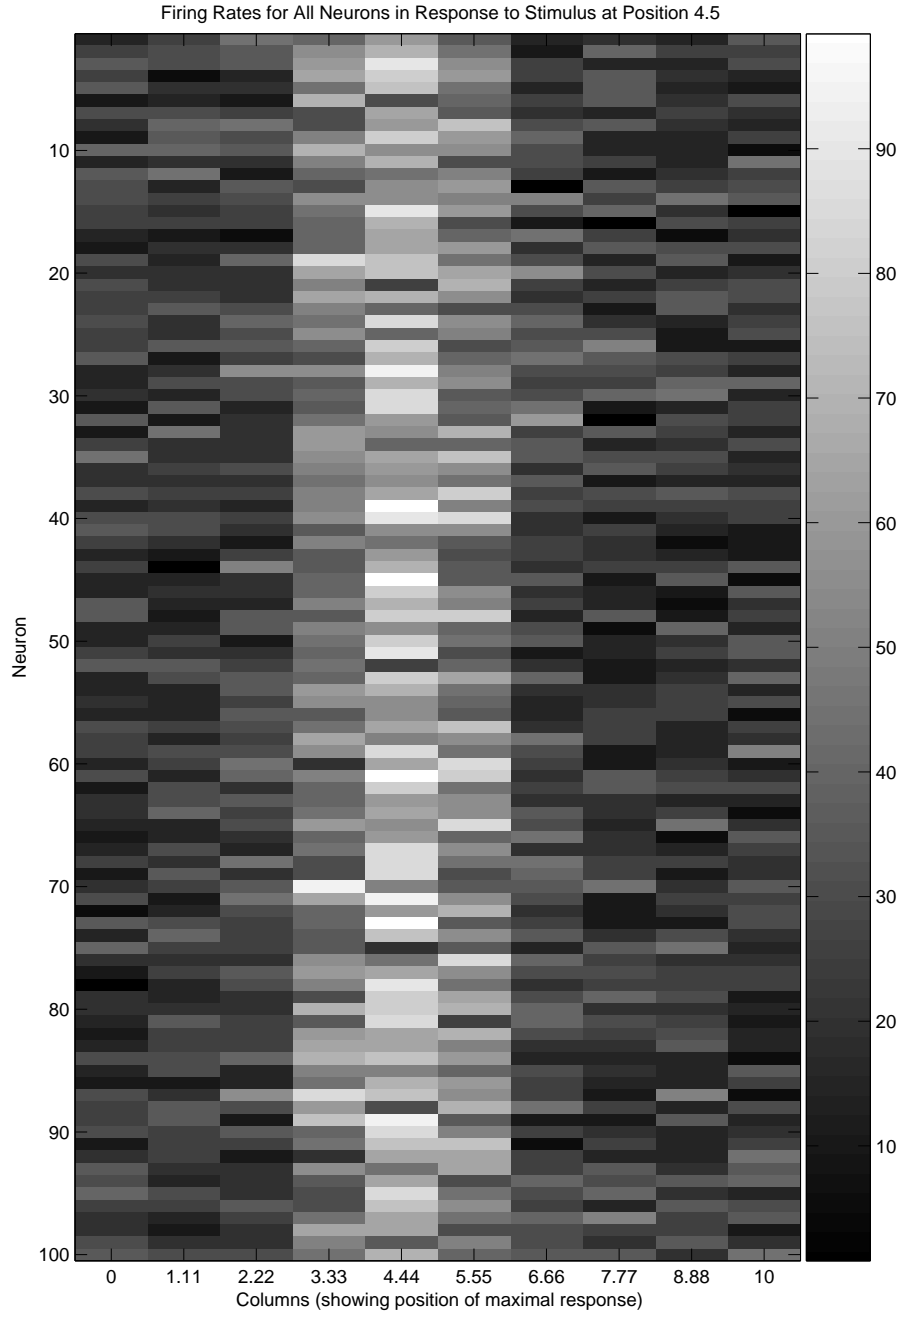
Figure 3: The mean firing rate for each neuron, shown here using a grayscale index, with higher firing rates shown by lighter shades. The neurons in the columns which have maximal response near to the constant stimulus position of 4.5 show higher firing rates, as expected. The variability in individual firing rates even within the same column is also apparent here.

smaller, the stimulus tracking task becomes more difficult, and this is reflected in the increased MSE for the smaller window sizes. It is notable, however, that even for small window sizes, the network with inhibitory connections performs quite well, giving a performance with a 10ms estimation window slightly better than that of the network with no inhibitory connections and a 50ms estimation window. This highlights the very significant benefits that inhibitory connections within a pool of neurons can bring in situations where fast estimation is required, which would be expected for an evolved sensory processing system. Figure 3 shows the firing rates for all of the neurons in the network (in this case with inhibitory connections), demonstrating that neurons in the columns whose maximal response properties are closest to the stimulus position have the highest firing rates, but also emphasising the noisy nature of neural firing while performing the tracking task.

Figures 4 and 5 show the same results for a random steplike signal which moves position instantaneously every 100ms. As for the constant stimulus, the network with inhibition significantly outperforms the network with no lateral connections, and the difference, not that apparent at 100ms, becomes increasingly obvious as the estimation window size decreases. The relationship between estimation window size and MSE for networks with and without inhibitory connections is shown explicitly in figure 6.

3.2   Correlation Measures

It is useful to look more closely at one specific simulation, and in particular several different measures of correlation. Figure 7 gives the results of a simulation (of both a network with inhibitory connections, and one without) using default values and a 50ms estimation window. The top row of Figure shows that, in keeping with the results of the previous section, the MSE for a network with inhibition is much lower than that of a network with no inhibition. The difference is qualitatively characterised by a smoother estimation curve
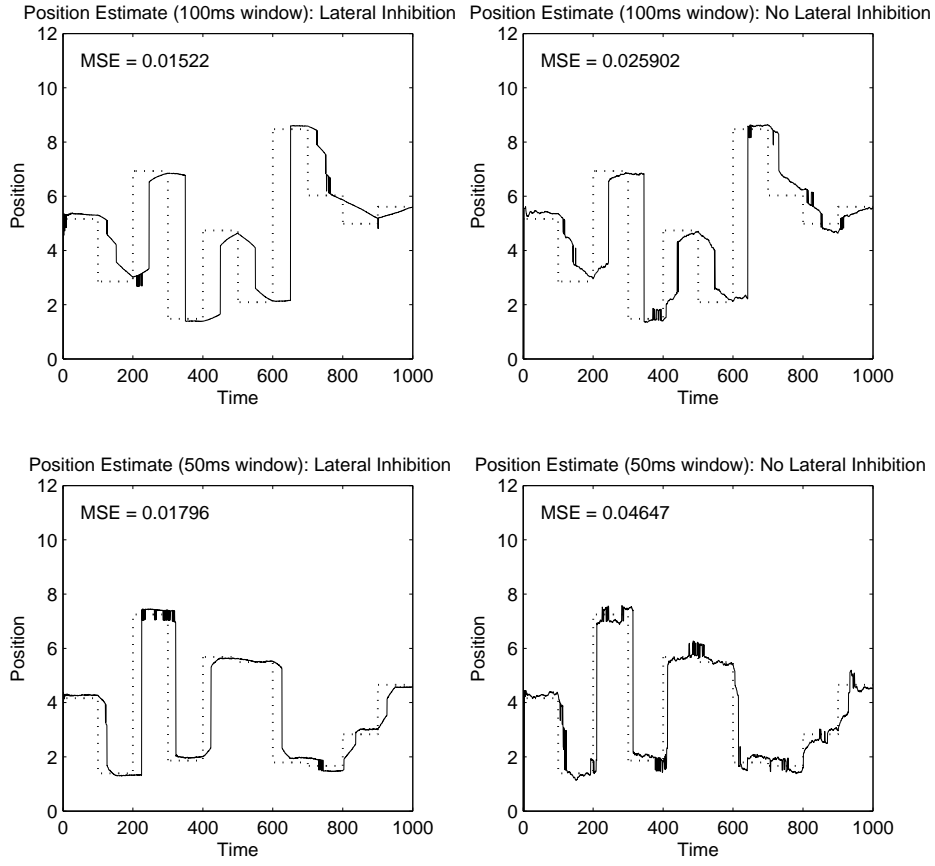
Figure 4: Position targets (dotted line) and estimates (solid line) for a random steplike signal, across 100ms and 50ms estimate window sizes (shown here in different rows). *Left column*: Lateral inhibition used. *Right column*: Lateral inhibition not used. It can be seen that the MSE is lower when lateral inhibition is present, across all window sizes.
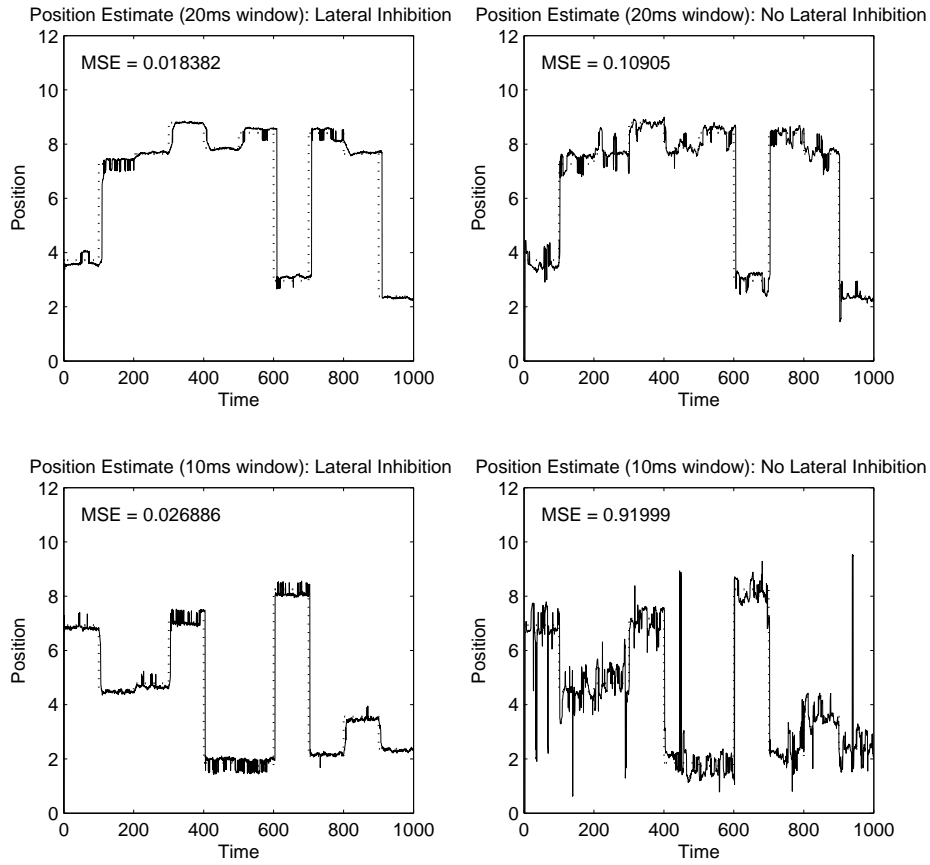
Figure 5: Position targets (dotted line) and estimates (solid line) for a random steplike signal, across 20ms and 10ms estimate window sizes (shown here in different rows). *Left column*: Lateral inhibition used. *Right column*: Lateral inhibition not used. It can be seen that the MSE is lower when lateral inhibition is present, across all window sizes. The difference is greater for these smaller estimation windows than for the larger estimation windows in the previous figure.
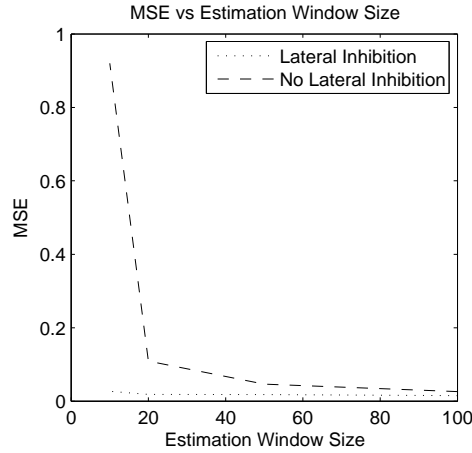
Figure 6: The effect on the MSE of estimation window size. The increased difficulty of the estimation task given less time can be seen here. Performance for the network without inhibitory connections drops quite considerably for small window sizes, while the presence of inhibitory connections helps the network to maintain a robust performance at these smaller sizes.

for the inhibitory network, reflecting its greater reliability. This difference can measured quantitatively by autocorrelation at low lags, which are shown in the left-middle graph, the greater autocorrelation reflected the smoother, more reliable, curve.

Of central importance to the understanding of lateral inhibitory connections in a column of neurons is the concept of firing rate correlation. This is the level of synchronisation (or anti-synchronisation) of the spike trains of neurons within a column. It was earlier shown (figure 1) that when the firing rates of neurons were negatively correlated (that is, when one neuron fired more and other neurons correspondingly tended to fire less), the central limit benefit of pooled activity was enhanced. One mechanism for these firing rates to be achieved is if the spike trains of the neurons themselves are negatively correlated, which means that when one neuron fires, it reduces the probability of other neurons firing. Clearly, lateral inhibition might be expected to lead to this behaviour. The right-middle graph of figure 7 shows correlation curves for the networks with and without lateral inhibition. These are
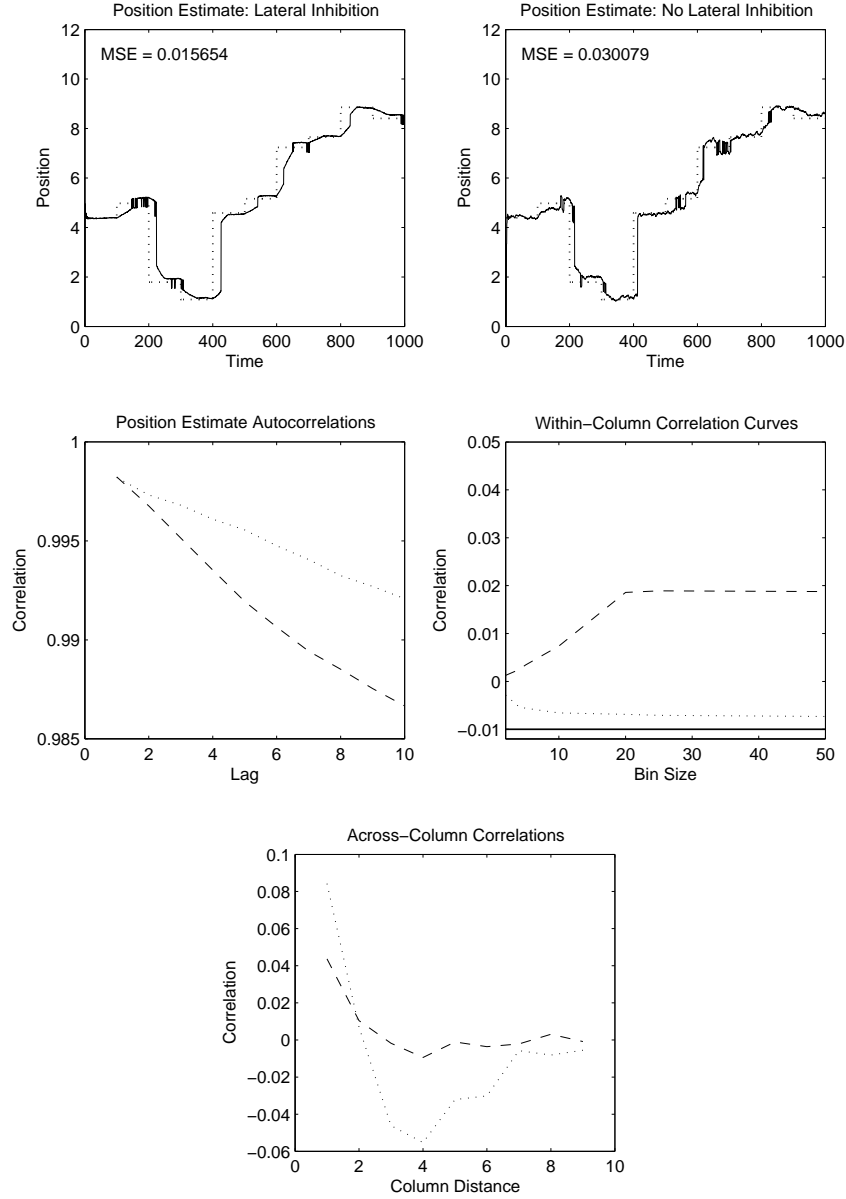
17

Figure 7: A detailed example using a 50ms estimate window. *Top*: Position targets (dotted line) and estimates (solid line) for a random steplike signal. The better performance with lateral inhibition present can be seen. *Middle left*: Position estimate autocorrelation. These give an indication of the reliability of the signal by measuring short fluctuations in the estimate, and show the greater reliability of the estimate when lateral inhibition is used (dotted line). *Middle right*: Correlation curves of firing patterns for neurons in the same column. The use of inhibitory connections results in negatively-correlated firing patterns (dotted line). *Bottom*: Mean correlation across columns as a function of column distance. This shows that neurons in columns close to each other will tend to respond in a similar way as expected. Neurons in columns at medium distances tend to respond in quite different ways, giving a 'mexican hat' shape, which is accentuated by the presence of inhibitory connections (dotted line).

calculated by measuring the number of spikes within short time bins (and shown across a range of bin sizes in order to ensure that the results are not an artefact of bin size), and evaluating the correlation of these with each other; the mean correlation value is shown. The graph also shows a baseline value (solid line), which is the lowest possible mean correlation value, a constraint resulting from the number of neurons in a column. It is clear from the graph that inhibitory connections have resulted in negatively correlated firing patterns emerging from the network, and without these, the firing patterns are slightly positively correlated.

In addition to measuring the correlation of neurons within columns, it is of interest to evaluate the correlation across columns, given that column location is directly related to neural response properties (neurons in columns close to each other have similar response properties). The mean correlation between columns is shown at the bottom of figure 7. It is clear from this that neurons in columns close to each other tend to respond in similar ways as expected, and also that neurons in very distant columns tend to be independent of each other. Interestingly, neurons a medium distance away show a negative correlation, indicating that they tend to respond at different times. We believe that this is because neurons at longer distances are actually have a negative correlation component when either are particularly active, but a positive correlation component when columns centrally in between them are more active and as a result they are both less active at the same time. Columns a middle distance apart do not have as strong a positive correlation component (since there are fewer columns in between), and as a result appear to be more negatively correlated overall. It is also interesting to observe that the correlation curve is more accentuated for the network with inhibitory connections than the one without, indicating a smoother harmonic behaviour by that network.

3.3 Threshold, Firing Rate and Inhibitory Connection Strength

As described in section 2.3, the inputs have been set to balanced, in order to allow comparison with [11]. However, this is not in itself needed for our network, since we are not using the analytical expression for neural firing rates dependent on this, given in [2]. It is also the case that in the network with inhibitory connections obviously has an additional source of input to each neuron (lateral inhibitory inputs from other spiking neurons) beyond the external stimulus input. As a result, given the same input strength and firing threshold, neurons in the inhibitory network will have a lower firing rate. This is shown in the top row of figure 8. It should be noted that in this situation, the improved performance of the network with inhibitory connections is evident. In order to rule out the difference in firing rates as a possible factor, and to simulate balanced inputs, the firing rate threshold of neurons in the inhibitory network was reduced to a level that would allow a firing rate approximately the same as that of the non-inhibitory network (the default situation). It should be noted that this is simply a scaling operation, equivalent to increasing the scale of the inputs or simulating background excitatory connections, for example, and in no way changed the dynamics of the model. The middle row of figure 8 confirms that the inhibitory network still has a much better performance than the non-inhibitory network, showing that irrespective of whether the firing threshold or the firing rates are the same, the inhibitory connections still allow the network to perform better.

Closely related to the relationship between inhibitory connections and firing rates is the relationship between inhibitory connections and within-column correlation. We have already seen in section 3.2 that inhibitory connections lead to negatively correlated firing patterns. The bottom of 8 shows that as the connections become more inhibitory, the correlation duly becomes more negative, as we would expect if the inhibitory connections were responsible
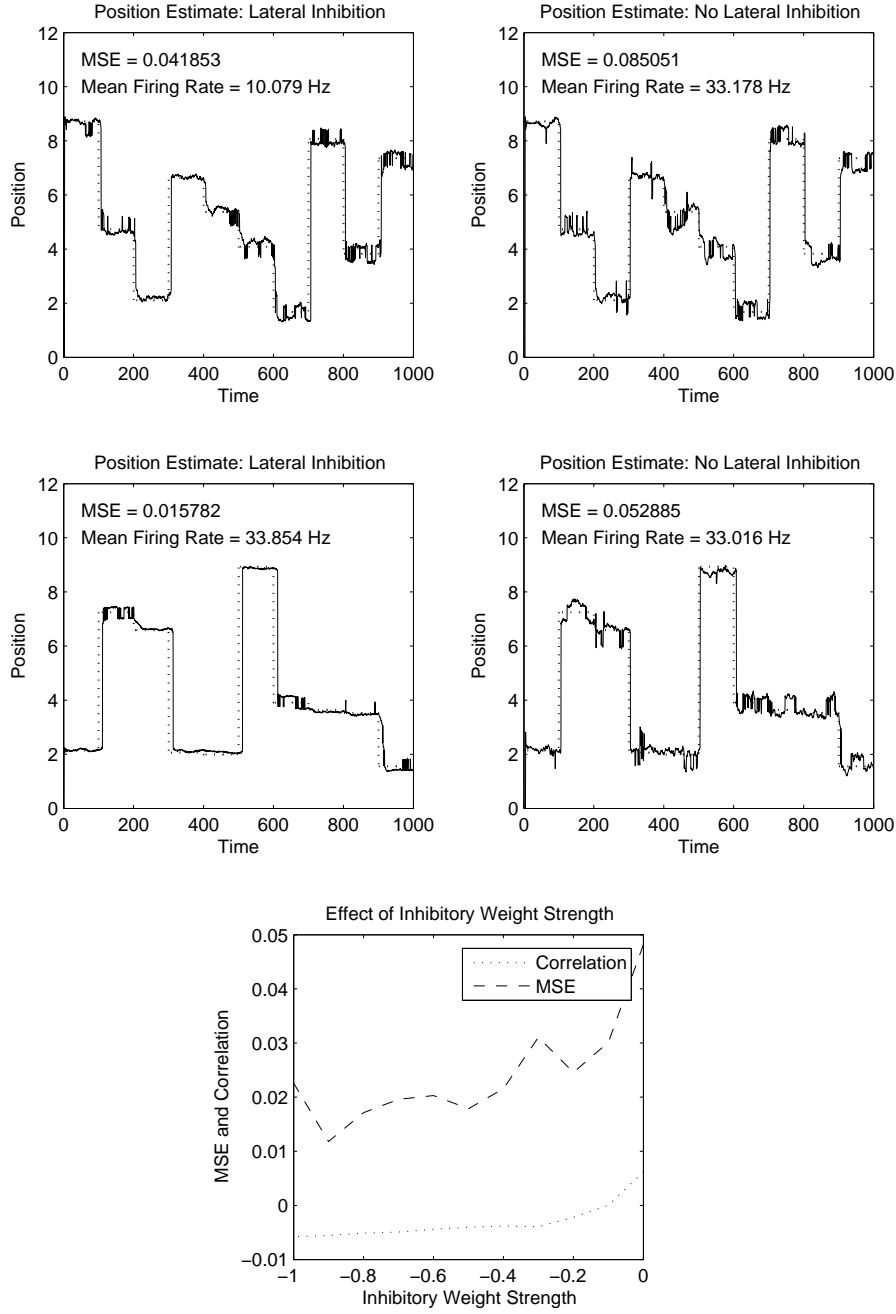
20

Figure 8: The effects of threshold and firing rate. *Top*: Position targets (dotted line) and estimates (solid line) for a random steplike signal using a 20ms estimate window, with the same firing threshold for each ($v_{thre} = 20$). *Middle*: Position targets (dotted line) and estimates (solid line) for a random steplike signal using a 20ms estimate window, with approximately the same firing rate. *Bottom*: The effect of changing the inhibitory connection strength on MSE and correlation. Increasing the strength of the inhibitory connections decreases the MSE, and reduces the correlation in the firing patterns of neurons within columns.

for the negative correlation. It also shows that the MSE tends to decrease as the connections become more inhibitory, which conforms with our earlier results.

3.4   Number of Neurons

Our hypothesis, described earlier in section 1.3, is that population coding benefits from the central limit theorem, and that by having negatively correlated firing patterns, that benefit is enhanced. We have already seen that negatively correlated firing patterns arise from inhibitory connections and improve performance. In order to assess the initial central limit aspect, we tested our model with different numbers of neurons. Figure 9 top row shows the performance of networks with and without inhibitory connections, with 100 neurons in each column. The middle row shows the performance with just 10 neurons per column. It can clearly be seen that with fewer neurons, the performance is much worse. The bottom graph shows the central limit effect clearly across a range of neuron population sizes. It also shows that across all sizes, the presence of inhibitory connections, with the associated negative correlation, improves performance.

4   Discussion

4.1   Mechanism

Current theories of inhibitory mechanisms in the sensory cortex focus on lateral inhibition between neurons with slightly different response properties. These theories do not give an explanation of lateral inhibition between neurons with the same response properties, such as neurons in the same cortical columns in the primary visual cortex. We have presented a model here which demonstrates that the functional explanation of these local inhibitory mechanisms is that they improve the accuracy of the mean field potential by
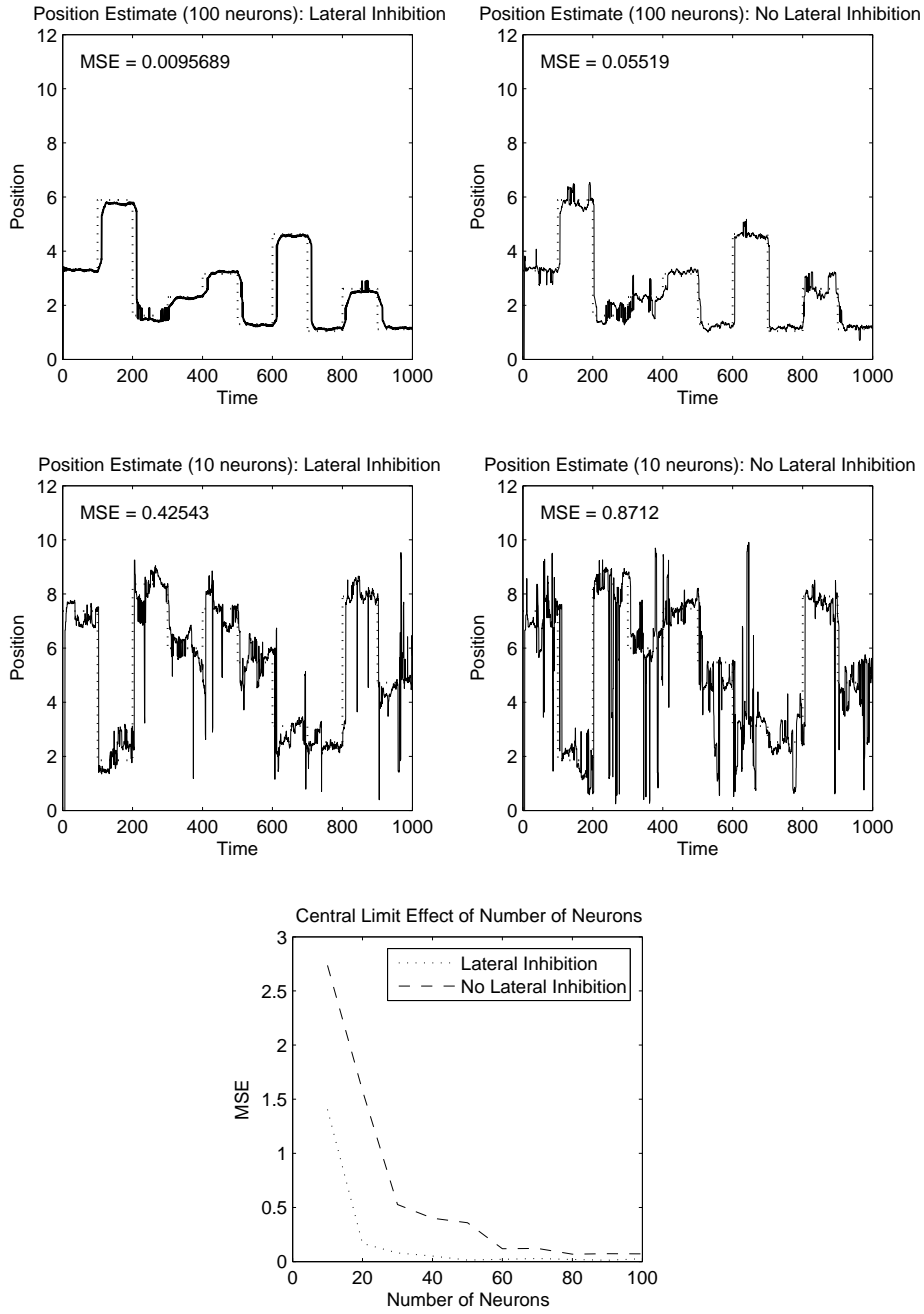
Figure 9: The effect of neuron numbers. *Top*: Position targets (dotted line) and estimates (solid line) for a random steplike signal using a 25ms estimate window, with 100 neurons in each column. *Middle*: Position targets (dotted line) and estimates (solid line) for a random steplike signal using a 25ms estimate window, with 10 neurons in each column. *Bottom*: The effect of changing the number of neurons on the MSE for when inhibitory connections are present (dotted line) and not present (dashed line). The central limit effect can be seen clearly in both cases, where increasing the number of neurons reduces the MSE. In addition the MSE is consistently lower when inhibitory connections are present irrespective of the number of neurons.

reducing the effects of noise. The explanation is as follows:-

1. Inhibition between neurons in the same column ensures that when one neuron fires, the probability of other neurons firing is decreased. This gives negatively-correlated firing patterns.

2. Negatively-correlated firing patterns mean that one neuron with a higher than average firing rate will result in reduced firing rates of other neurons in the pool. This is in effect the centre-surround space-filling property of negatively-correlated data.

3. The combined activity of all neurons in the pool (mean field potential) is stabilised by the centre-surround property, ensuring that the effects of random noise are minimised in comparison to the same pool of neurons operating without inhibitory connections.

4. The reduced-noise pooled activity is carried through decoding and leads directly to an improved performance on the given task of the pool of neurons.

In the case of our model, the improved performance of the network on the stimulus tracking task was evident in each experiment outlined. It is important to note that the by effectively making the noise negatively correlated, inhibitory connections are actually minimising the negative effects of noise without actually needing to reduce the intrinsic or extrinsic noise of individual neurons, the positive aspects of noise such as stochastic resonance, which our simulation actually exploited, remain intact.

4.2 Multi-Electrode Data

Given that our model has produced negatively-correlated firing patterns, and shown the benefit of these, we might expect actual neural firing in the brain to be negatively-correlated, at least in areas where the type of population

coding we have used is in force. Historically, data taken from single-electrode recordings has not been able to give the firing patterns of neurons close to each other during the same time period, meaning that it has not been possible to accurately assess whether or not neural firing in the brain is in fact negatively correlated. Recently, however, data from the olfactory bulb obtained using a multi-electrode array, coupled with an advance in spike sorting [7], has allowed us to examine this question, and significantly it appears that neural firing patterns are indeed negatively correlated, as our model predicts.

4.3   Further Developments

Our model outlines the functional meaning local inhibitory connections, and demonstrates the benefits on a specific task. While being careful to ensure that the model has clear biophysical correlates as far as possible, it is not intended in itself to be a biophysically accurate model. We are currently working on a more biophysically realistic model, taking into account the specific neural architecture of orientation columns in the primary visual cortex, in order to see how the principles outlined in our current model will translate into this specific neural environment.

Related to a more biophysically-inspired model is the important question of how existing descriptions of lateral inhibition, and in particular how inhibition between neurons with similar but not identical response properties fits into the picture. We believe that the negatively-correlated firing patterns seen between pooled neurons in our current model will also exist between neurons across columns where lateral inhibition is in effect, but the functional meaning of this type of inhibition may be quite different.

Finally, further exploration of the implications of the current model in terms of other theories could yield fruitful results. For example, it has been sug-

gested that the brain employs sparse codes [3, 9]. Sparse coding actually arises naturally as a consequence of negatively correlated neural firing patterns, where the negative correlation ensures that only a small number of neurons are actively involved in representing a given input at any one time. Could the important benefits of local inhibitory mechanisms outlined here be responsible for the brain's choice of neural code? This is one of the questions that we will seek to answer in continued investigations into the benefits of local inhibitory mechanisms.

References

[1] Andersen, P., Dingledine, R., Gjerstad, L., Langmoen, I.A. and Laursen, A.M. (1980): Two different responses of hippocampal pyramidal cells to application of gamma-aminobutyric acid *J.Physiol* 307:279-296

[2] Feng, J.F. and Ding, M.Z. (2004): Decoding spikes in a spiking neuronal network *J.Phys.A* 37:5713-5727

[3] Field, D.J. (1995): Visual Coding, Redundancy and 'Feature Detection' *The Handbook of Brain Theory and Neural Networks* ed. Arbib, M., MIT Press, 1012-1016

[4] Gammaitoni, L., Hänggi, P., Jung, P. and Marchesoni, F. (1998): Stochastic Resonance *Rev.Mod.Phys.* 70:223-287

[5] Gerster, W. and Kistler, W. (2002): *Spiking Neuron Models* CUP

[6] Hirsch, J.A. and Gilbert, C.D. (1991): Synaptic physiology of horizontal connections in the cat's visual cortex *J.Neurosci* 11:1800-1809

[7] Horton, P. and Feng, J.F. (2005): (Spike Sorting Paper Title Here) *submitted*

[8] Martin, K.A.C. 1984: Neuronal circuits in cat striate cortex *Cerebral Cortex, Vol.2, Functional Properties of Cortical Cells* ed. Jones, E. and Peters, A., Plenum, 241-284

[9] Olshausen, B.A. and Field, D.J. (1997): Sparse coding with an overcomplete basis set: A strategy employed by V1? *Vision Research* 37:3311-3325

[10] Ratliff, F. (1972): Contour and contrast *Scientific American* 226,6:901-10

[11] Rossoni, E. and Feng, J.F. (2005): Decoding spike ensembles: tracking a moving stimulus *submitted*

[12] Tucker, T.R. and Katz, L.C (2003): Spatiotemporal patterns of excitation and inhibition evoked by the horizontal network in layer 2/3 of ferret visual cortex *J.Neurophysiol* 89:488-500

[13] Tucker, T.R. and Katz, L.C (2003): Recruitment of Local Inhibitory Networks by Horizontal Connections in Layer 2/3 of Ferret Visual Cortex *J.Neurophysiol* 89:501512