

SIMULATION OF A PROPOSED BINDING MODEL

L. Andrew Coward
Department of Computer Science
Australian National University
ACT 0200
Australia
andrewc@cs.anu.edu.au

Abstract

A model for cortex information processing is described which depends upon a combination of population, rate and temporal coding for action potential spikes. In this model, binding of information derived from one visual attention object occurs because attention causes a slight (~ 1 millisecond) shift in each spike in the sensory inputs derived from the object towards the nearest peak to the spike in a 40 Hz modulation frequency. This frequency modulation results in preferential processing of the information derived from the attention object. Simulations of populations of leaky integrator neurons with both excitatory and inhibitory connectivity demonstrate that this preferential processing occurs with physiologically reasonable synaptic integration times, and allows object categorization on time scales consistent with human cognitive processing.

Introduction

An early proposal for the role of temporal synchronization of neuron action potentials or spikes in the cortex was that such synchronization could be used to tag the activity of neurons with a signal indicating the perceived object to which its activity relates. This temporal binding hypothesis [von der Malsburg 1981] has been seen as requiring neurons capable of detecting coincidences of spikes to within a few milliseconds, and has therefore led to controversy over whether coincidence detection is a viable model for cortex neurons [Konig et al 1996]. The hypothesis has been severely criticized by Shadlen and Movshen [1999] on a number of grounds, including the incompleteness of the theory in that it describes the indicator of binding without describing how binding is computed; the presence of the proposed binding indicator in areas of the visual cortex where binding is not cognitively present; and the lack of evidence for the required coincidence detection neurons.

Another aspect of the debate has been discussion of the information content of action potentials. There have been two major positions on this issue: rate coding and temporal coding [Gautrais and Thorpe 1998]. In rate coding, information is represented by average firing rates, consistent with the observations that firing rates in sensory neurons correlate with the intensity of an encoded feature [van Rossum, Turrigiano, and Nelson 2002]. Synchrony is not a factor with the model, and may even compromise rate coding [Reyes 2003]. An issue with the rate coding view is that the actual generation of action potentials by biological neurons is very erratic, and it has therefore been argued that 10 – 50 millisecond integration times would be required in each neuron layer to read out rate coded signals [Gautrais and Thorpe 1998; Shadlen and Newsome 1998]. Such integration

times would be in conflict with the human ability to categorize a complex scene in ~150 milliseconds [van Rossum, Turrigiano, and Nelson, 2002].

In temporal coding, some aspect of the relative spike times plays a crucial role in the information content of the spikes. One example is the proposal that information is contained in the order in which neurons spike [Gautrais and Thorpe, 1998]. The temporal binding hypothesis is another example.

Simulations of the responses of cortex neurons using models with a range of biological realism indicate that in layered networks of such neurons, propagation of activity across multiple layers tends either to become synchronized or to be extinguished [Diesmann, Gewaltig, and Aertsen, 1999; van Rossum, Turrigiano, and Nelson, 2002; Reyes 2003].

The way in which information is coded in the specific population of neurons generating action potentials has received somewhat less attention, partly because the response fields of neurons do not appear to correlate with cognitively significant conditions or objects [Tanaka 1993]. A model which makes use of population coding in a way that incorporates this lack of correlation with cognitive conditions is the recommendation architecture [Coward 1990; 2000; 2001]. In this model, cortex neurons detect groups of similar conditions, but for reasons connected with a combination of the need to conserve information recording resources and the need to learn without interference with prior learning, the group detected by one neuron can expand but not contract [Coward 2001]. Groups therefore cannot be converged to correspond exactly with cognitive conditions, and the presence of conditions within one group can only be interpreted as a set of recommendations that a range of cognitive conditions are present. The actual detection of a specific cognitive condition results from summation of all the currently present recommendations and selection of the strongest. However, in the recommendation architecture model, information of specific other types can also be encoded in both average spike generation rates and partial spike synchronizations [Coward 2000], and the combination of population, rate and temporal coding makes it possible to avoid the issues created by length of integration times needed for rate coding, and the apparent absence of the synchronization detectors needed for temporal coding.

This paper has four objectives. The first objective is to describe a partial synchronization approach called frequency modulation in which spikes are shifted towards peaks in a modulation frequency. It will be demonstrated by simulations that shifts of the order of one millisecond applied with a modulation frequency of 40 Hz to the inputs to a layered network of leaky integrator neurons greatly increases activity

propagation across the layers, for synaptic integration times ranging from 2 milliseconds to 25 milliseconds. The second objective is to describe a proposed model for binding which uses a combination of population, rate and temporal coding and addresses all the issues raised by Shadlen and Movshen [1999]. A third objective is to report the results of simulations of this binding model. The simulations employ networks of leaky integrator neurons made up of columns of devices across sequences of layers. Connectivity is randomly defined, but with a bias within columns in favour of excitatory inputs which tend to be active at the same time and in favour of inhibitory inputs from different columns. These biases are typical results of recommendation architecture learning. The simulations demonstrate that the simultaneous presentation of modulated inputs derived from one object and unmodulated inputs derived from another object results in network activity characteristic of just the modulated object, as required by the binding model. The fourth objective is to demonstrate that the discrimination in favour of the modulated object occurs within the observed timescales for cognitive processing.

Neuron Network Dynamics

There have been a number of studies of the dynamic behaviour of cortical neurons in vivo [e.g. Shadlen and Newsome 1998], in electronic simulation [e.g. Diesmann, Gewaltig, and Aertsen, 1999], and even using one biological neuron to model many different neurons in a network [Reyes 2003].

A typical cortex neuron receives 3000 – 10,000 synaptic contacts, of which 85% are believed to be excitatory. The post synaptic potential generated by one input action potential depolarizes the membrane by about 3 – 10%, so neglecting inhibition a neuron would be expected to produce a spike whenever 10 – 40 input spikes arrive within 10 – 20 milliseconds [Shadlen and Newsome 1998]. Average neuron firing rates can be up to 100 spikes/second, but within an average rate actual spike generation timing is very irregular [Tomko and Crapper 1974].

Simulations have typically employed the leaky integrator model neuron, in which a spike input generates a post synaptic potential which then decays with some time constant. Examples of such simulations include Brody [1998], Shadlen and Newsome [1998], Diesmann, Gewaltig, and Aertsen [1999] and van Rossum, Turrigiano, and Nelson [2002]. The simulated neural networks have generally been successive feedforward layers of neurons with connectivity between layers all to all (e.g. Diesmann et al [1999]) or randomly selected more sparse (e.g. 10% in Reyes [2003]). In general, only excitatory connectivity has been included.

One general result is that synchrony tends to develop across layers. This synchrony has been labeled pulse packets by Diesmann et al [1999], and it appears that a pattern of input activity either propagates synchronously across a series of layers or dies away [Reyes 2003]. Another interesting result is that a relatively slow interaction between neurons can result in synchronization of their outputs on a much shorter timescale [Brody 1998].

Information coding in the recommendation architecture cognitive model

In the connectionist recommendation architecture [Coward 1990; 2000], the pyramidal neuron is modeled as shown in figure 1. Such a neuron detects any occurrences of a set of similar conditions, where a condition is defined as a combination of inputs of a condition defining type. The learning algorithm for this device is such that new conditions can be added to the set, but once a condition has been added it cannot be removed. Management of the circumstances under which conditions can be added is therefore of critical importance [Coward 2001]. These neurons are arranged in

layers, with columns defined across multiple layers. The condition defining inputs to devices in one layer of a column are mainly derived from the preceding layer of the same column. As illustrated in figure 1 and discussed in detail in Coward [2001], devices also have inputs which excite recording of additional conditions and inputs which inhibit such recording. Outputs from other layers and columns indicating the general level of activity in those layers and columns are the sources for these change management inputs. A device has groups of provisional inputs, where groups are preassigned to be similar to previously recorded conditions. One such group is illustrated in figure 1. If inputs exciting the recording of additional conditions are active and inputs inhibiting such recording are inactive, an additional condition is defined as the largest currently active subset of available provisional input groups.

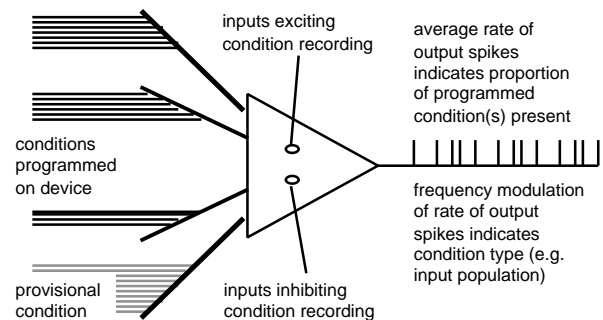


Figure 1. A condition recording device. These devices have different groups of inputs defining similar but different conditions. The presence of one or more recorded conditions activates the device. Other inputs determine whether or not a device will record an additional condition at any point in time, but these inputs do not form parts of conditions. An activated device generates a series of voltage spikes. The average spike rate indicates the number of recorded conditions currently present, and the phase of a frequency modulation imposed on the average rate indicates the input space within which the conditions have been detected.

The outputs of a column are the outputs of devices in a specific layer of the column. These outputs indicate the presence of the conditions detected by these devices. The learning algorithms result in a column detecting a set of fairly similar conditions, and an array of columns detecting different sets of conditions within the same input space. Because the sets can expand but not contract, a column set will not correspond exactly with a cognitive feature. However, the learning algorithms result in column outputs which have better discrimination between such features than column inputs. In other words, if an input can typically occur in instances of N cognitive features and an output in n such features, then $n \ll N$ [Coward 2001]. A feature will therefore be indicated by the presence of any large subset of a limited group of columns. There will be some overlap between the groups for different features. The output from one column can be regarded as a recommendation in favour of the presence of a small number of features, and the largest total recommendation weight found by summation of the recommendation weights for all active columns across all recommended features indicates the feature currently present. This summation and comparison takes place in subcortical structures [Coward 2001].

As illustrated in figure 1, device outputs are sequences of spikes. Any output indicates the current presence of conditions within the neuron set, and the average spike generation rate indicates the number of such conditions. A frequency modulation of the rate indicates the input population within which the conditions were detected. Frequency modulation means that although the average rate

over some period of time is constant, spikes are more likely to occur bunched near peaks in some modulation signal as illustrated in figure 2. Strong bunching is equivalent to packetization as defined by Deisman et al [1999]. If such a modulation were imposed on the inputs to a layer of neurons, then for higher thresholds the activity in the layer will be much higher than in the absence of modulation as illustrated in figure 2. Furthermore, as demonstrated by the simulations discussed below, the difference in activity level between modulated and unmodulated inputs would be maintained or even amplified in subsequent layers, because of the tendency for neurons to produce outputs in phase with the modulation.

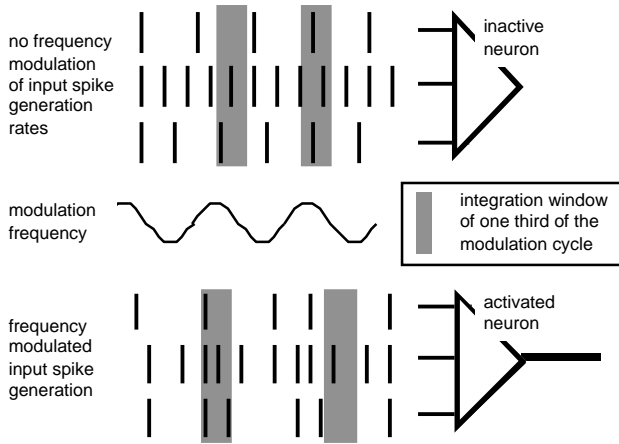


Figure 2 Frequency modulation slightly shifts each individual action potential spike towards the nearest peak in a modulation signal. If three inputs to a neuron are unmodulated, the number of spikes within one integration window varies from two to four, but from one to five with modulation. If the neuron threshold was equivalent to five spikes, only the modulated inputs would generate any output.

The functional effect of modulation can be understood as follows. Suppose that inputs were derived from a visual field, and within that field was an object defined by a continuous closed boundary. Suppose further that detection of the presence of such a boundary could trigger imposition of a frequency modulation only on inputs from within the boundary. The result would be preferential detection of conditions within the visual object relative to the rest of the visual field. Independent detection of simple conditions of different types (colour, shape etc.) could be associated by a common frequency modulation phase and therefore used to detect combination conditions limited to information derived from the one object.

The recommendation architecture approach thus employs a combination of population, rate and temporal coding to perform information processing. The frequency modulation coding has some resemblances with von der Malsburg's original proposal that temporal coherence in neuron spike generation could contribute to identifying populations of activity related to individual objects, but there are some important differences. Von der Malsburg and later workers [see Shadlen and Movshon 1999 for extensive references] have argued that synchronicity between spikes within a few milliseconds is the determining factor indicating neuron activity relating to a particular visual object, and the 40 Hz signal found in the EEG gamma band is a manifestation of such synchronization processes [Llinas et al 1994]. In the model proposed here, spike synchronization can be much less precise.

The Von der Malsburg temporal binding model has been strongly criticized by Shadlen and Movshon [1999] on a number of grounds. Firstly, they point out that in the theory as proposed, temporal correlation is not a theory about how binding is computed, only a theory of how binding is

signaled. In other words, it does not indicate how the visual system decides which elements belong to different objects. They suggest that the binding problem could more readily be solved by successive elaboration of progressively more complex representations of visual scenes as explored by Tsotsos [1995]. Secondly, they point out that the study of visual deficits indicates that binding only exists in the higher levels of the visual cortex hierarchy, in which case the observations of synchrony in the primary visual cortex is unexpected. Thirdly they argue that the implied coincidence detector model for cortex neurons is implausible, and that given a more plausible integration window of 5 - 10 milliseconds, all spikes are in synchrony with other spikes. There has been extensive debate over whether cortical neurons are integrators or coincidence detectors. Konig et al [1996] point out that the essential difference between these views is the period of time over which synaptic inputs are integrated. This period of time is of the order of the synaptic potential decay time, or comparable with the membrane constant which is estimated to be in the range 8 - 16 milliseconds [Konig et al, 1996], much longer than the few milliseconds suggested for temporal coincidence detectors.

Simulation of the effect of synaptic integration time on frequency modulation mechanism

In order to investigate the requirements placed on neuron parameters such as synaptic integration times by the frequency modulation model, the mechanism has been tested using a simulated population of neurons. Six hundred neurons in three layers of two hundred were simulated by a software implementation. Two hundred externally generated signals were inputs to the first neuron layer, outputs from the first layer were inputs to the second layer and so on. The externally generated signals were structurally identical with neuron outputs, being made up of sequences of spikes modulated and delayed in various ways. For each run the performance of the model was simulated over a period of 500 milliseconds broken down into 0.33 millisecond time slots. The status of each neuron was determined from its past condition, its current inputs, and some noise related uncertainties once every time slot.

Each neuron had a set of randomly selected inputs from the appropriate input population, and the size of that input set was randomly selected in the range 40 - 60 for each neuron. Duplicate inputs to one neuron were not permitted. Each input from the appropriate population had the same chance of being selected for any neuron. The externally generated inputs to this model had spike rates for which the average interval between spikes was randomly selected for each input within the range 20 - 100 milliseconds, corresponding with rates in the range 10 - 50 spikes/second. Each of the 241 possible intervals at 0.33 millisecond resolution has the same selection probability. The overall average of spike rates across a set of 200 inputs was 20 spikes/second.

A spike occupied one time slot of 0.33 milliseconds. Once the interval had been selected for an input, the initial spike was delayed from the start of the 500 millisecond simulation period by a randomly selected phase of up to one interval. Generation rates were noisy in the sense that the time slot for each spike varied by a randomly selected interval of zero to two time slots from the slot defined by the average rate. The noisy average spike generation rates were then modulated at 40 Hz. This modulation was implemented by defining a phase for the 40 Hz signal relative to simulation start. Each spike was then shifted towards the nearest modulation peak. This shift was randomly selected for each spike, with a 50% probability of selecting a primary modulation shift, a 40% probability of selecting one time slot less than that shift, and a 10% chance of two slots less. The primary value could be varied in the range 3 to 5 time slots or 1 - 1.67 milliseconds, and the average shift (or the strength of modulation) therefore varied in the range 0.8 - 1.5

milliseconds. In addition, it was possible to introduce delays which could be randomly selected for each input in the range 0 - 2 milliseconds. The purpose of the delays was to test the tolerance of the process for inconsistent transmission delays. The various random noise factors test the robustness of the observed phenomena.

A neuron had a set of inputs, and detected any significant total activity in those inputs. In information terms the set of the condition sets indicated by its inputs. A single spike generated a post synaptic potential which varied as Ate^{-at} where t is the time after arrival of the spike and A and a are constants which were selected so that the peak in the post synaptic potential curve was 100 arbitrary units and the time from arrival of the spike until the potential declined to 33% of its peak value was a controllable parameter called the integration time.

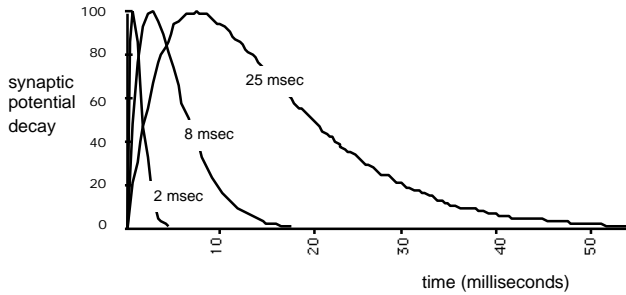


Figure 3 Synaptic Potential Decay Curves. Examples of synaptic potential decay curves used for neuron models, for a range of integration time parameters

The values of this integration time were selected in the range 2 - 25 milliseconds, to cover the periods postulated for synchronization detection (a few milliseconds) through the observed membrane constants in cortical neurons (8 - 16 milliseconds, see Konig et al [1996]), up to the period defined by the modulation frequency (i.e. 25 milliseconds for 40 Hz). Examples of the synaptic potential decay curves used are shown in figure 3. The neuron model employed was thus a conventional leaky integrate and fire mechanism [Dayan and Abbott 2001].

A neuron generated an action potential or spike if at any point in time the total potential generated by past input spikes exceeded a threshold T , provided that no output spike had been generated within an immediately preceding absolute refractory period set at 3 milliseconds to correspond with such periods observed in cortex neurons [Dayan and Abbott 2001]. The output spike was delayed by a randomly selected interval of 0.33 - 1 millisecond after threshold was exceeded to model the observed reliability of spike timing in neocortical neurons [Mainen and Sejnowski 1995]. When a spike was generated, the post synaptic potential was reduced to zero. The value of the threshold T was the same for all neurons in one layer, but set at different values for each synaptic integration time. The values were set so that a modulated set of inputs with an average spike rate of 20 spikes/second generated a similar level of output activity in each layer.

A number of different system models were generated and a number of different randomly constructed input states were applied to these models. Each model construction resulted in a different pattern of connectivity, each input state construction resulted in different spike rates, phases, noise, delays and modulation shifts, and each presentation of the same input state to the same model would generate slightly different results because of random delays in spike generation. The results described in this section are from over 100 different runs with different parametric settings.

The spikes in each timeslot in the input and in each layer in the first 40 milliseconds of a typical simulation with modulated inputs is shown in figure 4, with an unmodulated

input for comparison. As can be seen from the diagram, response is packetized in phase with the modulation signal, and packets are shorter and slightly delayed (by an interval of the order of the time between arrival of spike and peak in synaptic potential) as they progress through the layers. Response to an unmodulated signal also packetized, but was smaller and unrelated to a modulation signal.

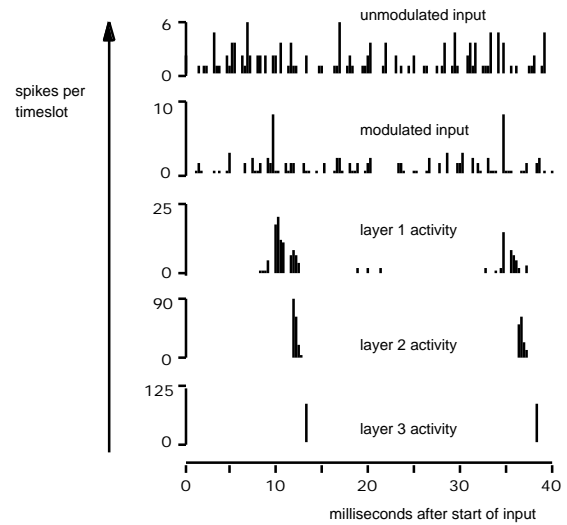


Figure 4 Total spikes per timeslot in a typical unmodulated and corresponding modulated input, and activity in the three layers in response to the modulated input. Response to an unmodulated input was also packetized but with lower total packet sizes.

The ratio of the sums of the average spike rates of all the neurons in each layer between when the same external inputs are unmodulated and modulated is illustrated in figure 5, with a modulation strength of 0.8 milliseconds and as the integration time varied over the range 2 - 25 milliseconds. The variability between runs for layer 2 is also illustrated. Random delays of up to 2 milliseconds in external inputs, had very little effect on the activity ratios. The variation in activity ratio with modulation strength in the range 0.8 - 1.5 milliseconds is shown for an integration time of 8 milliseconds in figure 6.

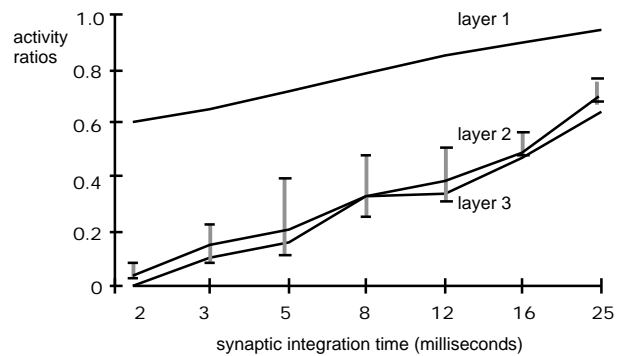


Figure 5 Effects of modulation on neuron activity levels as a function of synaptic integration time. The average ratios of activity of neurons between when the input layer 1 is presented an unmodulated set of inputs and when the same input rates are modulated at 40Hz with an average modulation shift of 0.8 milliseconds. The bars indicate the variation in the ratio found for layer 2 between different runs.

A number of implications follow from these results. Firstly, consistent modulation at 40 Hz across a set of inputs

considerably enhances the activity of the neurons targeted by those inputs. Although the enhancement is greater for integration times in the range 2 - 3 milliseconds (i.e. the "coincidence detector" realm), it remains strong at much longer integration times, even times comparable with the modulation period. Secondly, the enhancement is transmitted between layers, even though the modulation is only applied directly to the inputs and the neurons have no special tuning for 40 Hz. Thirdly, small, inconsistent delays of less than 2 milliseconds in the arrival times of the external inputs reduce but do not eliminate the enhancement. Fourthly, the enhancement is present despite random delays of the order of 1 millisecond in the generation of a spike following achievement of threshold. Fifthly, because for all spike generation rates there is a tendency for spikes to occur shifted towards the nearest modulation peak, the generation of spike trains will appear irregular, as observed for in vitro cortex neurons [Stevens and Zador, 1998].

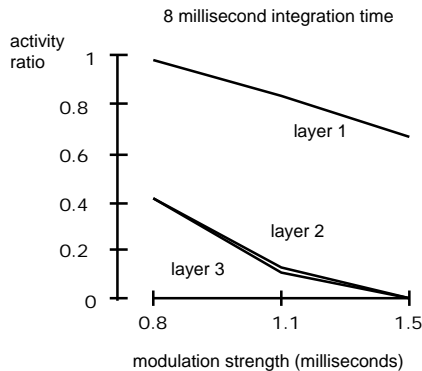


Figure 6 Effect of degree of modulation on neuron population activity for an 8 millisecond integration time.

The simulations thus demonstrate that if the inputs to a population of leaky integrator neurons are given a consistent frequency modulation of 40 Hz with a modulation amplitude which shifts spikes towards the peaks in a modulation frequency by a time ~ 1 millisecond, the activity of the population increases substantially, and the modulation is propagated across the population. The effect exists for a wide range of synaptic integration times and is robust against a range of noise parameters.

Proposed binding model

As indicated earlier, the model for binding of visual objects is as follows. A visual domain corresponding with an object defined as the area within a continuous enclosing boundary is selected within the total visual field. The definition of this attention domain would be the result of recommendations for domain changes generated by conditions including those corresponding with the existence of boundary elements. The outputs from the domain will in general be at many different average spike generation rates and phases. The selection of the attention domain results in all the rates from within that domain being modulated at a common frequency and phase.

Outputs from the entire visual field are then processed, but there is a bias in favor of detecting conditions derived only from information within the attention domain, and these conditions would be tagged by the same spike rate frequency modulation because the synaptic potentials of the devices tend to exceed threshold in phase with the modulation. The set of domain outputs could be processed separately, for example to detect conditions correlating with color and shape, and when the conditions came together they would be tagged by a common frequency modulation phase which would favor detection of higher level conditions derived from the original attention domain. Because the higher level

activity populations indicate conditions predominantly derived for the object within the attention domain, the corresponding behavioral recommendations will be predominantly appropriate to that object.

This binding model has additional operational advantages. There are cognitive processes which require behavior to be generated with respect to groups of objects. These processes require information derived from the presence of all members of the group required for high integrity behavior, even though attention must be directed to different members of the group in sequence. An example could be reading, where sequences of words make up phrases and sequences of phrases make up sentences. Information from all the words in the phrase should be present before higher level conditions corresponding with the meaning of the phrase are detected. A similar example is mental arithmetic processing as discussed in Coward [2001]. The generation of appropriate recommendations could be managed as follows. The attention domain is directed in sequence at the several objects in the group, and for each object a population of higher level conditions is maintained active, but each population with a different phase of frequency modulation so that few conditions at even higher level are detected. Once conditions derived from all the members of the group are present the modulation phases are synchronized and combination conditions corresponding with behavioral recommendations with respect to the group are generated.

This binding model addresses the issues raised by Shadlen and Movshon [1999]. It is a model for how binding is computed, it explains why a binding related synchronization should be present at levels where binding is not detected cognitively, and it does not require unrealistically narrow integration windows.

Simulation of the binding model

The system model was three layers of leaky integrator neurons, organized into four columns. Each column contained 100 neurons in each layer and detected different types of conditions within external input states. There were external excitatory inputs to the first layer from 200 separate sequences of action potential spikes. The excitatory inputs to the second layer were outputs from the first layer and so on.

Four types of input state were defined by the most probable range of values for the average spike generation rate of each of the 200 external inputs. For one input state type, each input was randomly assigned an integer bias parameter in the range 1 to 10, with equal selection probability for each value. A bias parameter value had an associated relative probability for each possible average spike generation rate. A bias parameter of 1 corresponded with a high probability of a high average spike rate, 5 with a high probability of an intermediate spike rate, and 10 with a high probability of a low or zero spike rate. Different instances of an input state type were constructed by random selection of average spike rates for each input, with the selection biased by the corresponding bias parameter. Instances were then normalized to an overall average of 20 spikes/second. The presence of a type is therefore only indicated by a trend across all inputs.

Neurons were the leaky integrators described earlier, except that each neuron also had inhibitory inputs derived from the other neurons in the same layer. The synaptic potential decay curve in response to both excitatory and inhibitory spikes was for an 8 millisecond synaptic decay time as illustrated in figure 3. However, for inhibitive inputs, the curve could be scaled so that the maximum synaptic potential was adjusted by a factor from 0 (i.e. no inhibition) to 8. The time at which a spike was generated was a randomly selected interval in the range 1 - 5 milliseconds after post synaptic potential exceeded threshold. The threshold for all devices in all layers was the same, but could

be varied in the range 20 - 40 times the peak magnitude of the excitatory synaptic potential in response to one spike.

Recommendation architecture learning results in inputs to first layer neurons in a column tending to be inputs which have had high activity at similar times in the past, with a corresponding bias placed on second and third layer devices in the column. Detailed algorithms have been described and simulated in Coward [2001]. Learning was not modeled directly for the current simulations, but the results of such learning were modeled by biasing the selection of external inputs to the first layer of each column in favour of inputs with a low bias parameter (i.e. a high probability of high activity) in a different input state type. Excitatory inputs to the second layer were biased in favour of first layer devices often active at the same time and so on. One effect of the bias was that 90% of excitatory connectivity was within columns. The number of excitatory inputs to a device was randomly selected in the range 40 - 60.

The inhibitory connectivity to a device from other devices in the same layer was randomly selected with a 3:1 bias in favour of inputs from outside the column of the device, and set at 15% of the number of excitatory inputs. The intended effect of the inhibitive connectivity was that if spike packets of comparable size were initiated in more than one column, the earliest packet would suppress the packets in other columns.

Instances of one input state type would typically generate the strongest activity in the column corresponding with the type. In a practical learning situation there would be many columns corresponding with different condition types, each defined by different groups of system inputs which tended to have high activity at the same times. A cognitive feature would be indicated by significant activity in many columns in a particular subset, different features being indicated by different (but overlapping) subsets. However, to simplify the description of the simulations and results it will be assumed that the input state types correspond with different cognitive objects.

The scenario was that two cognitive objects were within the visual field, and attention was focused on one of the objects. The attention focus was assumed to result in outputs from one object being frequency modulated at 40 Hz. The modulation of inputs was imposed by shifting each spike towards the nearest modulation peak. The average shift was a controllable parameter varying between 1 and 3.7 milliseconds. The actual shift for a spike was varied randomly from the selected average by ± 0.7 millisecond. Eight sets of instances were generated, each made up of one instance of each input state type. Mixed inputs for each of the twelve possible combinations of modulated instance of one type and unmodulated of the other type were constructed as a 200 input mixed set by adding the modulated spikes for one (attention) input state to the unmodulated spikes for a different type of input state. A supplementary modulation, in phase with the input modulation, could be applied to the first level neurons. This modulation was intended to test if such a modulation could be used to adjust to different overall levels of input activity.

Simulations were performed for 200 millisecond input period with a timeslot of 0.3 milliseconds. The total activity in each layer of each column was determined over the period, and also the timing of the spikes in each column.

Simulation Results

Spike counts were measured in each layer of each column. Spikes generated by layers were again packetized. The first packet in layer 1 of the column corresponding with the modulated input was usually generated ~ 8 milliseconds after the start of inputs, in response to the first modulation peak in the inputs. Packets were initiated in layer 1 of other columns, particularly in the column corresponding with the

unmodulated input, with a tendency to be initiated later than in the column corresponding with the modulated input.

In the absence of inhibition, a packet with less than a minimum number of spikes in layer 1 decayed through the other layers. The minimum depended on threshold. At the highest thresholds, packets in all columns tended to be extinguished by the third layer. For low thresholds there tended to be activity in the third layer of the columns corresponding with both the modulated and the unmodulated input or even other columns at the lowest threshold levels.

For intermediate threshold levels, there was a strong tendency for the packet in the column corresponding with the modulated input to be earlier and larger than packets in any other column in layer 1. Often the smaller packets decayed in later levels. When inhibition was introduced, the decay of packets in other than the column corresponding with the modulated inputs was more frequent and faster. However, at certain threshold levels the introduction of inhibition would result in decay of all packets by layer 3.

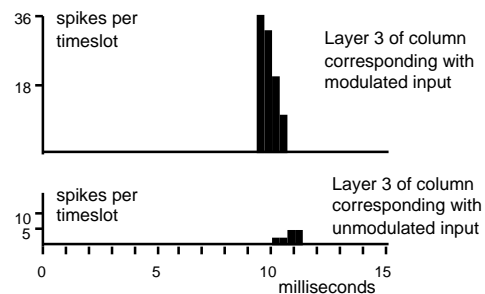


Figure 7 Typical spike generation by layer 3 neurons in two columns in response to a mixed input with modulated spikes for the condition corresponding with one column and unmodulated spikes for the condition corresponding with the other column. Activity in the rest of the columns was negligible.

If the packet in layer 3 of the column corresponding with the modulated condition was earlier than and had at least twice the number of spikes as any packets in any other columns, system response was regarded as correct. A typical example of correct spike activity in layer 3 in the two columns corresponding with the modulated and unmodulated condition is illustrated in figure 7. Under these conditions activity in layer 3 of other columns was minimal.

Table 1 gives the number out of the 12 possible modulated/unmodulated instance combinations of a set of four instances, one of each input state type in which spike activity was in the correct column. The variables in table 1 are the degree of input modulation, device threshold, inhibition level and first level modulation. The variation for different sets of instances and different system connectivity selections was of the order of ± 1 . For low thresholds, errors were activity in the wrong column, for high thresholds, errors were lack of activity in any column.

A number of conclusions can be drawn from figure 7 and table 1. Firstly, column activity discriminates between modulated and unmodulated inputs with a high degree of accuracy, reaching over 90% for the highest modulation level and optimized threshold range. Secondly, first level modulation changes the threshold range within which the highest accuracy is achieved, but does not affect the level achieved. This means that for a fixed threshold level, first level modulation could be increased until there was column activity, and at this point the accuracy would be high. Thirdly, accuracy is improved by inhibition in the threshold ranges in which peak accuracy is achieved, but the improvement in accuracy is not sensitive to the degree of inhibition in the range studied. At higher threshold levels inhibition sometimes made a lower accuracy worse, because

the effect of inhibition was to eliminate activity in all columns by layer 3. Fourthly, figure 7 indicates that discrimination is achieved within ~ 10 milliseconds for 3 layers and requires an extra 2 milliseconds per layer, consistent with observed human categorization times ~ 150 milliseconds.

Inhibition	First level modulation = 0				First level modulation = 4			First level modulation = 8		
	0	2	4	6	2	4	6	2	4	6
Average modulation = 3.7 msec										
Threshold										
3000	0	2	8	8	0	4	3	0	0	2
3250	1	10	11	11	2	8	9	0	4	7
3500	6	12	11	10	9	12	11	3	8	9
3750	9	7	6	6	12	11	11	10	11	12
4000	0	0	1	1	8	6	6	11	10	10
Average modulation = 2.3 msec										
Threshold										
3000	1	8	9	8	2	6	6	0	3	3
3250	10	10	9	8	8	9	9	2	6	6
3500	7	5	4	4	10	9	7	5	10	10
3750	0	0	0	0	5	5	3	10	9	9
4000	0	0	0	0	0	0	0	4	4	3
Average modulation = 1.0 msec										
Threshold										
3000	7	8	5	3	4	7	7	2	4	6
3250	0	0	0	0	6	6	4	4	7	7
3500	0	0	0	0	0	0	0	9	5	5
3750	0	0	0	0	0	0	0	0	0	0
4000	0	0	0	0	0	0	0	0	0	0

Table 1 Typical number correct out of 12 possible modulated/unmodulated combinations of one instance of each condition. Grey indicates bands in which maximum accuracy occurs.

Conclusions

Simulations have demonstrated that if the inputs to a population of leaky integrator neurons are given a consistent frequency modulation at 40 Hz which shifts individual action potential spikes by ~ 1 millisecond from the position defined by an average generation rate, the activity of neurons in the population is considerably increased, and the modulation is transmitted across the population. The enhancement has been demonstrated to be significant for synaptic potential integration times over a range from 2 milliseconds to 25 milliseconds. The enhancement is therefore not dependent on the existence of coincidence detection type mechanisms with integration times of the order of a few milliseconds.

The enhancement provided by frequency modulation has been used as the basis for a proposed computational model for object binding which results in a high proportion of the activity at higher levels being derived from one object that is the subject of an attention focus, even though other objects are present in the overall visual field. The model uses population, rate and temporal coding of action potential spikes to communicate different types of information. The model indicates how binding is computed, it explains why a binding related synchronization should be present at levels where binding is not detected cognitively, and the required neuron model has physiologically plausible parameters.

The binding mechanism has been tested by simulations of populations of model leaky integrator neurons organized into columns in which each column learns a different, cognitively ambiguous condition in accordance with the algorithms of a proposed recommendation architecture cognitive model. Results of these simulations demonstrate that discrimination of cognitive objects can occur with high integrity on timescales consistent with human cognitive processing times.

References

- Brody, C. D. (1998). Slow Covariations in Neuronal Resting Potentials Can Lead to Artefactually Fast Cross-Correlations in Their Spike Trains. *Journal of Neurophysiology* 80, 3345 – 3351.
- Coward, L. A. (1990). *Pattern Thinking*, New York: Praeger.
- Coward, L. A. (2000). A Functional Architecture Approach to Neural Systems. *International Journal of Systems Research and Information Systems*, 9, 69 - 120.
- Coward, L. A. (2001). The Recommendation Architecture: lessons from the design of large scale electronic systems for cognitive science. *Journal of Cognitive Systems Research* 2(2), 111-156.
- Dayan, P. and Abbott, L.F. *Theoretical Neuroscience*. MIT Press: Cambridge, Mass.
- Diesmann, M., Gewaltig, M-O., and Aertsen, A. (1999). Stable propagation of synchronous spiking in cortical neural networks. *Nature* 402, 529 – 533.
- Gautrais, J. and Thorpe, S. (1998). Rate coding versus temporal order coding: a theoretical approach. *Biosystems* 48, 57 – 65.
- Konig, P., Engel, A. K., and Singer, W. (1996). Integrator or Coincidence Detector? The role of the cortical neuron revisited. *Trends in Neurosciences* 19(4), 130 - 137.
- Llinas, R., Ribary, U., Joliot, M., and Wang, X.-J. (1994). Content and Context in Temporal Thalamocortical Binding, in Buzsaki G. et alii (eds), *Temporal Coding in the Brain*, Berlin: Springer.
- Mainen, Z., and Sejnowski, T. (1995). Reliability of Spike Timing in Neocortical Neurons. *Science* 268 (5216), 1503 - 1506.
- Reyes, A. D. (2003). Synchrony-dependent propagation of firing rate in iteratively constructed networks in vitro. *Nature Neuroscience* 6(6), 593 – 599.
- Shadlen, M. N. and Movshon, J. (1999). Synchrony Unbound: a critical evaluation of the temporal binding hypothesis. *Neuron* 24, 67 - 77.
- Shadlen, M. N. and Newsome, W. T. (1998). The Variable Discharge of Cortical Neurons: Implications for Connectivity, Computation, and Information Coding. *Journal of Neuroscience* 18(10), 3870 – 3896.
- Stevens, C.F. and Zador, A.M. (1998). Input synchrony and the irregular firing of cortical neurons. *Nature Neuroscience* 1 (3), 210 - 217.
- Tanaka, K. (1993). Neuronal Mechanisms of Object Recognition. *Science* 262, 685 – 688.
- Tomko, G. J. and Crapper, D. R. (1974). Neuronal Variability: non-stationary responses to identical visual stimuli. *Brain Research* 79, 405 – 418.
- Tsotsos, J. K. (1995). Towards a Computational Model of Visual Attention. In *Early Vision and Beyond*, T. V. Pappathomas ed. Cambridge MA: MIT Press.
- Van Rossum, M. C. W., Turrigiano, G. G. and Nelson, S. B. (2002). Fast Propagation of Firing Rates through Layered Networks of Noisy Neurons. *Journal of Neuroscience* 22(5), 1956 – 1966.
- Von der Malsburg, C. (1981). *The Correlation Theory of Brain Function*. Internal Report 81-2, Max Planck Institute for Biophysical Chemistry, Göttingen, Germany.
- Von der Malsburg, C. (1999). The What and Why of Binding: the modeler's perspective. *Neuron* 24, 95 - 104.