

Computation with Spikes in a Winner-Take-All Network

Matthias Oster

mao@ini.phys.ethz.ch

Rodney Douglas

rjd@ini.phys.ethz.ch

Shih-Chii Liu

shih@ini.phys.ethz.ch

Institute of Neuroinformatics, Uni-ETH Zurich, CH-8057 Zurich, Switzerland

The winner-take-all (WTA) computation in networks of recurrently connected neurons is an important decision element of many models of cortical processing. However, analytical studies of the WTA performance in recurrent networks have generally addressed rate-based models. Very few have addressed networks of spiking neurons, which are relevant for understanding the biological networks themselves and also for the development of neuromorphic electronic neurons that communicate by action potential like address-events. Here, we make steps in that direction by using a simplified Markov model of the spiking network to examine analytically the ability of a spike-based WTA network to discriminate the statistics of inputs ranging from stationary regular to nonstationary Poisson events. Our work extends previous theoretical results showing that a WTA recurrent network receiving regular spike inputs can select the correct winner within one interspike interval. We show first for the case of spike rate inputs that input discrimination and the effects of self-excitation and inhibition on this discrimination are consistent with results obtained from the standard rate-based WTA models. We also extend this discrimination analysis of spiking WTAs to nonstationary inputs with time-varying spike rates resembling statistics of real-world sensory stimuli. We conclude that spiking WTAs are consistent with their continuous counterparts for steady-state inputs, but they also exhibit high discrimination performance with nonstationary inputs.

1 Introduction ---

The winner-take-all (WTA) computation is an intrinsic property of recurrent networks, which abound in cortex. Several studies have discussed the computational power a WTA network offers (Riesenhuber & Poggio, 1999; Yuille & Geiger, 1998; Maass, 1999, 2000) and its role in cortical processing models, for example, a hierarchical model of vision in cortex (Riesenhuber

& Poggio, 1999) or a model of selective attention and recognition processes (Itti, Koch, & Niebur, 1998).

This computation is thought to be an intrinsic decision component of the cortical microcircuit. Decision processes in the brain are not localized in one specific region, but evolve in a distributed manner when different brain regions cooperate to reach a consistent interpretation. The winner-take-all circuit with both cooperative and competitive properties is a main building block that contributes to this distributed decision process (Amari & Arbib, 1982; Douglas, Koch, Mahowald, Martin, & Suarez, 1995).

Because of these properties, WTA networks have been of great interest to researchers. Yuille and Grzywacz (1989) and Ermentrout (1992) are classical references to theoretical analyses. In these early models, the neurons are nonspiking, that is, they receive an analog input and have an analog output. The analog WTA computation can be efficiently implemented in very-large-scale integrated (VLSI) transistor circuits. With initial circuits described in Lazzaro, Ruckebusch, Mahowald, and Mead (1989), a whole series of analog models (e.g., Kaski & Kohonen, 1994; Barnden & Srinivas, 1993; Hahnloser, Sarpeshkar, Mahowald, Douglas, & Seung, 2000) and implementations has been developed (He & Sanchez-Sinencio, 1993; Starzyk & Fang, 1993; Serrano-Gotarredona & Linares-Barranco, 1995; Kincaid, Cohen, & Fang, 1996; Indiveri, 1997, 2001; Moller, Maris, Tomes, & Mojaev, 1998; Hahnloser, Sarpeshkar, Mahowald, Douglas, & Seung, 2000; Liu, 2000, 2002).

In the past decade, spiking neuron models and their electronic counterparts have gained increasing interest. Spike-based networks capture the asynchronous and time-continuous computation inherent in biological nervous systems. Neuron models with analog inputs and analog outputs can be converted into models with spiking output if a thresholding operating is introduced to the neuron. Coultrip, Granger, and Lynch (1992) is an early theoretical analysis, with further descriptions in Jin and Seung (2002) and Yu, Giese, and Poggio (2002) and VLSI implementations in Chicca, Indiveri, and Douglas (2004), Abrahamsen, Häfliger, and Lande (2004), and Oster, Wang, Douglas, and Liu (2008).

The next theoretical consideration are models with both spiking input and spiking output. Previous theoretical studies focused on population models (e.g., Lumer, 2000), where the population firing represents a graded analog value. Indiveri, Horiuchi, Niebur, and Douglas (2001) and Chicca, Lichtsteiner, Delbruck, Indiveri, and Douglas (2006) are VLSI implementations that use the firing rate as an analog input and output encoding. A special case is presented in Carota, Indiveri, and Dante (2004) and Bartolozzi and Indiveri (2004), in which the core of the winner-take-all is analog but the signals are converted to spike rates for communication with the outside world. Other theoretical studies consider alternate neuron models, such as oscillatory neurons (Wang, 1999) or differentiating units (Jain & Wysotzki, 2004).

No analysis until now has considered the effect of single spikes and spike timing on the winner-take-all computation with spiking inputs and outputs. Gautrais and Thorpe (1998) start their analysis from a similar point of view, that is, how the network decides which of two input spike rates is higher, but they do not consider sampling of this estimation in the output spikes (this analysis could be classified as spiking input and analog output).

The emergence of multichip spiking systems that incorporate the WTA as part of their decision process (Serrano-Gotarredona et al., 2005; Choi, Merolla, Arthur, Boahen, & Shi, 2005; Chicca, Lichtsteiner, Delbruck, Indiveri, & Douglas, 2006; Vogelstein, Mallik, Culurciello, Cauwenberghs, & Etienne-Cummings, 2007) highlights the necessity for theoretical quantification of a WTA network based on different network parameters (e.g., CAVIAR), especially if these systems are to be used in different applications. To address this need, we develop a framework for quantifying the performance of a spiking WTA network with spiking inputs based on the network parameters and input statistics. We start with the definition of the hard WTA architecture in section 2 and treat the network as an event-based classifier. The network has to decide which neuron, the “winner,” receives the strongest input after a certain time interval, and it indicates its decision with an output spike. In a spiking system, the term *strongest input* has to be defined: How is the input signal encoded in the spikes? What are the statistical properties of the spike trains? We consider the following cases: stationary inputs of regular frequencies in section 3, stationary inputs of Poisson distribution, effects of self-excitation and inhibition on the WTA performance in section 4, and a model of nonstationary Poisson inputs in section 5 for two cases: where the strongest input switches between two neurons and where the input activity travels across the neurons of the network. The latter resembles the statistics of real-world sensory stimuli. This formalism has been applied during the programming of the WTA module in the CAVIAR multichip asynchronous vision system (Oster, Wang, Douglas, & Liu, 2008; Serrano-Gotarredona et al., 2005).

2 Winner-Take-All Network Connectivity

We assume a network of integrate-and-fire neurons that receives input spikes through excitatory or inhibitory connections. The WTA operation is implemented by having these neurons compete against one another through inhibition. In biological networks, excitation and inhibition are specific to the neuron type. Excitatory neurons make only excitatory connections to other neurons, and inhibitory neurons make only inhibitory connections. This specificity is obeyed in a typical WTA network where inhibition between excitatory neurons is mediated by populations of inhibitory interneurons (see Figure 1a).

To adjust the amount of inhibition between the neurons (and thereby the strength of the competition), both types of connections could be modified:




Figure 1: Simplification of connectivity (a) with a global inhibitory neuron and (b) with direct inhibitory connections between neurons. (Reproduced from Figure 1 in Oster et al., 2008.)

the excitatory connections from excitatory neurons to interneurons and the inhibitory connections from interneurons to excitatory neurons. In our model, we assume the forward connections between the excitatory and the inhibitory neurons to be strong, so that each spike of an excitatory neuron triggers a spike in the global inhibitory neurons. The amount of inhibition between the excitatory neurons is adjusted by tuning the connections from the global inhibitory neurons to the array neurons. This configuration allows the fastest spreading of inhibition through the network.

With this configuration, we can simplify the network by replacing the global inhibitory interneurons with full inhibitory connectivity between the excitatory neurons (see Figure 1b). This simplification is used only for the analysis. In an electronic implementation, the configuration with the global interneurons is more suitable.

While a cortical neuron can receive possibly correlated inputs from up to 10,000 synapses, we assume that the external inputs to the neurons are uncorrelated and that all the inputs can be summed together into one spiking input. This assumption is valid especially for Poisson-distributed input, since summing spike trains with Poisson statistics results again in a spike train with Poisson statistics.

In addition to the excitation from the external inputs and inhibition from other neurons, each neuron has a self-excitatory synapse that facilitates its selection as the winner in the next cycle once it has been chosen as the winner. In this analysis, we do not consider external inhibitory input spikes even though this could easily be done by assigning a sign to every input spike. Since the statistical properties stay the same, we disregard the external inhibitory input for clarity.

2.1 Neuron Model. We assume a network of N integrate-and-fire neurons with excitatory and all-to-all inhibitory connections. The dynamics

of the membrane potential of the i th neuron, V_i , follows the equation of a nonleaky integrate-and-fire neuron model with nonconductance-based synapses,

$$\frac{dV_i}{dt} = V_E \sum_n \delta(t - t_E^{(n)}) - V_I \sum_{\substack{j=1 \\ j \neq i}}^N \sum_m \delta(t - t_j^{(m)}), \quad (2.1)$$

where we assume unit capacitance. The first term on the right describes the excitation from an external spike source (E) as the weight, V_E , times a sum of delta functions reflecting the spike train from this source. The second term on the right describes the inhibition from other neurons as the weight, V_I , times a sum of delta functions, which reflects the spike train from a particular neuron. In this analysis, we neglect the dynamics of the synaptic currents and the delay in the transition of the spikes. The neuron spikes when its membrane potential exceeds a threshold V_{th} , that is, $V_i \geq V_{th}$, and is then reset to $V_i = 0$. Immediately afterward, it receives a jump of V_{self} in its membrane potential due to self-excitation. The membrane potential satisfies the constraint, $0 \leq V_i \leq V_{th}$, that is, an inhibitory spike cannot drive V_i below 0.

3 Stationary Inputs of Regular Rates

We first discuss how we set up the network connectivity for the case of the hard WTA mode where the neurons receive external input spike trains of regular frequency. In this mode, a winning neuron is selected after it has received a predetermined number of input spikes, n , needed to generate an output spike. This neuron is the one whose external input spike train has the smallest interspike interval. We now consider how to set up the constraints on V_E and V_I so that even though other neurons beside neuron k might first produce transient output spikes because of initial conditions, these neurons will not fire again once neuron k spikes.

We assume that all neurons $i \in 1, \dots, N$, receive an external input spike train of constant frequency r_i , and neuron k receives an external input with the highest frequency ($r_k > r_i; \forall i \neq k$). We now describe the conditions for fulfilling the constraints for this hard WTA mode (see Figure 2):

1. Neuron k (the winning neuron) spikes after receiving $n_k = n$ input spikes that cause its membrane potential to exceed threshold if it is initially discharged. After every spike, the neuron is reset to V_{self} :

$$V_{self} + n_k V_E \geq V_{th}. \quad (3.1)$$

2. As soon as neuron k spikes once, no other neuron $i \neq k$ should be able to spike because of the inhibitory spike from neuron k . Another neuron can receive up to n spikes even if its input spike frequency is lower than that of neuron k because the neuron is reset to V_{self} after a




Figure 2: Membrane potential of the winning neuron k (a) and another neuron in the network (b). Black bars show the times of input spikes. Traces show the changes in the membrane potential caused by the various synaptic inputs. Black dots show the times of output spikes of neuron k . (Reproduced from Figure 2 in Oster et al., 2008).

spike, as illustrated in Figure 2. Hence, the inhibitory weight should be such that

$$V_I \geq n_k V_E \geq n_i V_E. \tag{3.2}$$

3. If a neuron j other than neuron k spikes in the beginning, there will be some time in the future when neuron k spikes and becomes the winning neuron. From then on, conditions 1 and 2 hold, so a neuron $j \neq k$ can generate a few transient spikes, but neuron k wins.

To see this, let us assume that the external inputs to neurons j and k spike with almost the same frequency (but $r_k > r_j$). For the interspike intervals $\Delta_i = 1/r_i$, this means $\Delta_j > \Delta_k$. Since the spike trains are not synchronized, the corresponding input spikes to both neurons k and j have a changing temporal offset ϕ . At every output spike of neuron j , the input spike temporal offset decreases by $n_k(\Delta_j - \Delta_k)$ until $\phi < n_k(\Delta_j - \Delta_k)$. When this happens, neuron k receives $(n_k + 1)$ input spikes before neuron j spikes again and crosses threshold:

$$(n_k + 1)V_E \geq V_{th}. \tag{3.3}$$

We can choose $V_{self} = V_E$ and $V_I = V_{th}$ to fulfill inequalities 3.1 to 3.3. V_E is adjusted to achieve the desired n .

Case 3 happens only under certain initial conditions, for example, when $V_k \ll V_j$ or when neuron j initially receives an external input spike frequency that is higher than that to neuron k . A leaky integrate-and-fire model will ensure that all membrane potentials are discharged ($V_i = 0$) at the onset of a stimulus. The network will then select the winning neuron after receiving a predetermined number of input spikes, and this winner will have the first output spike. Even if the conditions above are fulfilled, the neuron with the smaller input could stay as the winner for a long time before the neuron with the larger input takes over, as the switching dynamics depends on the initial conditions of the network.

If the rate is regular, the information about the strongest input is already contained in one interspike interval. If $V_{th}/2 < V_E < V_{th}$ and $V_{th}/2 < V_{self} < V_{th}$ are chosen, the network performs the selection in one interspike interval. We call this an optimal decision in the sense that the network will always choose the winning neuron as the one receiving the highest input frequency assuming perfect network homogeneity.

3.1 Multistability. A WTA network can exhibit multistability or hysteresis, that is, dependent on initial conditions, a neuron with a lower input frequency can remain the winner. By choosing $V_{self} > V_E$, a neuron k that is initially the winner will remain the winner even if another neuron j receives the stronger input. The increase in input frequency that is needed for neuron j to take over as the winner will depend on V_{self} .

WTA models that exploit firing rate thresholds are normally dependent on the number of neurons in the network. The WTA computation that we describe here exploits the spike timing; hence, the mechanism of competition is independent of the number of neurons. The network can be scaled to any size as long as the inhibitory neuron can still completely inhibit the excitatory neurons with one output spike.

4 Stationary Poisson Rates

For stationary Poisson-distributed input spikes, we first consider a network of two neurons, labeled 0 and 1, with the connectivity shown in Figure 3. Because of the Poisson statistics of the inputs, the probability of selecting the correct winner depends on the ratio of their input Poisson rates ν and the number of input spikes, n , needed for the neuron to reach threshold.

We assume that the neurons are initially completely discharged ($V = 0$). The probability that a neuron receives n input spikes in time t from a Poisson spike source with rate ν is

$$P(\nu t, n) = \frac{(\nu t)^n}{n!} \exp(-\nu t).$$




Figure 3: Simplified network architecture. Two neurons labeled 0 and 1 receive Poisson-distributed input spike trains with frequencies ν_0 and ν_1 , respectively. The synaptic strengths are V_E for the excitatory external input spikes, V_I for the cross-inhibition spikes between the neurons, and V_{self} for the self-excitation spikes.

A neuron crosses threshold at time t (thus producing an output spike) if it receives the n th spike exactly at time t and $n - 1$ input spikes in $[0; t]$. Neuron 0 has the first output spike of the network if it reaches threshold at time t (probability density π_0), while neuron 1 is still below threshold (probability P_1):

$$\pi_0(t) = \nu_0 P(\nu_0 t, n - 1)$$

$$P_1(t) = \sum_{i=0}^{n-1} P(\nu_1 t, i).$$

We integrate over all times $t = 0, \dots, \infty$ to compute the probability that neuron 0 has the first output spike:

$$P_{0out1} = \int_0^\infty P(\nu_0 t, n - 1) \cdot \nu_0 \cdot \left(\sum_{i=0}^{n-1} P(\nu_1 t, i) \right) dt. \tag{4.1}$$

For $n = 1$, every input spike elicits an output spike. The probability that the first output spike of the network comes from neuron 0 is the higher rate normalized to the sum of both rates:

$$P_{0out1} \Big|_{n=1} = \int_0^\infty \nu_0 e^{-\nu_0 t} e^{-\nu_1 t} dt = \nu_0 \int_0^\infty e^{-(\nu_0 + \nu_1)t} dt = \frac{\nu_0}{\nu_0 + \nu_1}. \tag{4.2}$$




Figure 4: Filter model of the WTA network. The WTA (black box) receives a Poisson input spike train of total rate ν . Every input spike encodes either symbol 0 or 1, with the probability of a spike coding for 0 of P_{0in} . The output of the WTA is a spike train with the probability of a spike coding for 0 of P_{0out} . The WTA ensures $P_{0out} > P_{0in}$ in the case where $P_{0in} > P_{1in}$ and amplifies the difference between the input probabilities.

For $n \rightarrow \infty$ or very large n , the neurons integrate over many input spikes. The Poisson distributed rates can then be approximated by the mean rate, and the decision is deterministic:

$$P_{0outI} \Big|_{n \rightarrow \infty} \rightarrow \begin{cases} 1 & : \nu_0 > \nu_1 \\ 0 & : \nu_0 < \nu_1 \end{cases} \tag{4.3}$$

This case is of course not practical because the WTA network would never produce an output spike since it has to integrate over an infinite number of input spikes.

In general, the probability that the first spike of the network comes from the neuron receiving the higher input rate increases as the neurons integrate over more spikes. Equation 4.1 cannot be solved easily in closed form, but it can be integrated numerically.

In the analysis, we assume that the neurons were initially discharged in the derivation of P_{0outI} . This probability is not necessarily the same as the general case of P_{0out} , which describes the probability that any output spike of the network comes from neuron 0. However, in the case of strong inhibition ($V_I = V_{th}$) and no self-excitation ($V_{self} = 0$), both neurons are also completely discharged after any output spike so $P_{0out} = P_{0outI}$. The network has no memory of its past input, and integration will start again with the initial conditions.

4.1 Filter Model. An alternate way of viewing this WTA process is to consider the network as a filter that increases the percentage of spikes encoding for neuron 0. Figure 4 details this interpretation as a filter model in signal theory. Interestingly, P_{0out} is independent of the total input rate, ν . To show this, we replace $\nu_0 \rightarrow P_{0in}\nu$ (rate of input spikes to neuron 0) and $\nu_1 \rightarrow (1 - P_{0in})\nu$ (rate of input spikes to neuron 1) with $\nu = \nu_0 + \nu_1$ in

equation 4.1. We substitute $\nu t \rightarrow t'$ and $\nu dt \rightarrow dt'$. The integration limits $(0; \infty)$ do not change:

$$P_{0out} = \int_0^\infty P(P_{0in}t', n-1) \cdot P_{0in} \cdot \left(\sum_{i=0}^{n-1} P((1-P_{0in})t', i) \right) dt'. \quad (4.4)$$

The output rate ζ is the sum of the output rates of each neuron. The sum is weighted by the probability that neurons 0 and 1 generate an output spike (P_{0out} and $1 - P_{0out}$):

$$\zeta = P_{0out} \frac{\nu P_{0in}}{n} + (1 - P_{0out}) \frac{\nu(1 - P_{0in})}{n} \quad (4.5)$$

$$= \frac{\nu}{n} (1 - P_{0in} - P_{0out} + 2P_{0in}P_{0out}) \quad (4.6)$$

The output spike train is not Poisson distributed, but will tend toward regular rates, since the neurons integrate over n input spikes for each output spike.

We quantify the effect of the WTA on the input and output probabilities for different values of n in Figure 5. The optimal classification would result in a step function: for any $P_{0in} > 0.5$, the output probability P_{0out} is 1. The results show that the output performance improves as the neurons integrate over more input spikes.

4.2 General Case. In the analysis above, we assume strong inhibition and no self-excitation. In the general case, $V_I < V_{th}$ and $V_{self} \neq 0$. Only the neuron that spiked is reset to ground, and, in addition, it receives self-excitation V_{self} . This neuron will reach threshold again with m spikes:

$$m = \left\lceil \frac{V_{th} - V_{self}}{V_E} \right\rceil \leq n. \quad (4.7)$$

The neuron that did not spike is inhibited by the spiking neuron, which lowers its membrane potential by V_I . It will reach threshold again with p spikes, dependent on the membrane potential at the time (t^-) of inhibition:

$$p = \left\lceil \frac{V_{th} - \max(V(t^-) - V_I, 0)}{V_E} \right\rceil \leq n. \quad (4.8)$$

$V(t^-)$ can take any value between ground and just below the threshold voltage.

To analyze the general case, we model the output of the network by a Markov model with inhomogeneous transition times with two states




Figure 5: Probability of neuron 0 producing output spikes of the network in the case of strong inhibition and no self-excitation. P_{0in} is the probability of a spike encoding for neuron 0 in the input rate; P_{0out} is the probability of a spike encoding for neuron 0 in the output rate. The different curves show the performance dependence on the number of input spikes n that the neurons have to integrate before producing a spike. The probability that the first output spike of the network comes from the neuron that receives the higher rate increases if the neurons integrate over more input spikes. The output probabilities are independent of the input rate v . Continuous curves: numerical evaluation of equation 4.4, overlaid data points: simulation results (error bar corresponds to 10 trials with 10,000 output spikes each).

(see Figure 6). The states determine which of the two neurons last fired. Each transition from state i to j has assigned a transition probability p_{ij} and a transition time t_{ij} . In contrast to a Markov chain, the states are sampled only when the network spikes and the transition times between the states are not constant. The transition probabilities are given by

$$p_{00} = \int_0^\infty v_0 P(v_0 t, m - 1) \cdot \sum_{i=0}^{p-1} P(v_1 t, i) dt \tag{4.9}$$

$$p_{01} = \int_0^\infty \sum_{i=0}^{m-1} P(v_0 t, i) \cdot v_1 P(v_1 t, p - 1) dt. \tag{4.10}$$

The transition probabilities p_{11} and p_{10} are the same as p_{00} and p_{01} , respectively, with v_0 and v_1 interchanged. The average transition time is the




Figure 6: Markov model of the WTA network with inhomogeneous transition times: the states determine which of the two neurons fires. The transition probabilities p_{ij} together with the transition time t_{ij} determine the sequence of firing.

number of spikes a neuron needs to reach threshold divided by its input rate:

$$\bar{t}_{00} = \frac{m}{v_0}, \quad \bar{t}_{01} = \frac{p}{v_1}, \quad \bar{t}_{10} = \frac{p}{v_0}, \quad \bar{t}_{11} = \frac{m}{v_1}. \tag{4.11}$$

The process reaches equilibrium if transitions between the states happen with the same probability:

$$p_{01} P_{0out} = p_{10} P_{1out}. \tag{4.12}$$

With $P_{1out} = 1 - P_{0out}$, we get

$$P_{0out} = \frac{p_{10}}{p_{01} + p_{10}}. \tag{4.13}$$

This is the probability that if a spike is emitted by the network, it is emitted by neuron 0. If $m = p$, then $p_{01} = p_{11} = 1 - p_{10}$, so $P_{0out} = p_{10}$ is equal to P_{0out1} in equation 4.1.

The mean output interspike interval $\langle \Delta \rangle$ is given by the weighted sum of the transition times:

$$\langle \Delta \rangle = P_{0out} p_{00} \frac{m}{v_0} + P_{1out} p_{10} \frac{p}{v_0} + P_{0out} p_{01} \frac{p}{v_1} + P_{1out} p_{11} \frac{m}{v_1},$$

with the corresponding mean output rate of the network:

$$\zeta = \frac{1}{\langle \Delta \rangle} = \frac{v_0 v_1 (p_{10} + p_{01})}{v_0 p_{01} (p_{10} p + p_{11} m) + v_1 p_{10} (p_{01} p + p_{00} m)}.$$

4.3 Effect of Self-Excitation. The Markov model allows us to quantify the effect of the self-excitation on the WTA performance. Here, we assume




Figure 7: Effect of self-excitation: the probability of a correct decision P_{0out} increases as the self-excitation V_{self} is increased. $m(p)$ is the number of spikes a neuron needs to reach threshold after self-excitation (global inhibition). For $m = p$, the curve is equal to the one shown in Figure 5 for $n = 10$. For $p = 10$, m is varied from $1, \dots, 10$. As the self-excitation is increased, the probability of a correct decision increases. Numerical evaluation of equation 4.13.

$V_I = V_{th}$, so p is independent of the membrane potential at the time of the inhibition (see equation 4.8). Strong self-excitation improves the probability that an output spike of the network is emitted from the neuron with the higher firing rate (see Figure 7). Interestingly, the output rate for a network that integrates over many input spikes and has strong self-excitation is similar to that of a network that integrates over fewer input spikes and has no self-excitation.

4.4 Effect of Inhibition. We now look at the performance in the case where the inhibition is weak, that is, $V_I < V_{th}$. From equation 4.8, the neuron that did not spike needs p spikes to reach threshold after being inhibited, while the neuron that last spiked needs m spikes to reach threshold. If $p < m$, a neuron that did not spike before will have a greater probability of spiking next. Starting from the same initial state for both neurons, the neuron that receives the strongest input is the most probable one to spike. After this output spike, the network will then select a neuron that did not spike before, and the probability of a correct decision of the network will decrease correspondingly. Hence, lowering the strength of the inhibition will lead to a decrease in performance.

Our description does not make any predictions about the membrane voltage V_{E1}^- of the non-winning neuron before the output spike, except $0 < V_{E1}^- < V_{th}$. Let us assume that the membrane was sitting close to threshold before the neuron receives inhibition V_I . After inhibition, $V_{E1}^+ = V_{th} - V_I$.

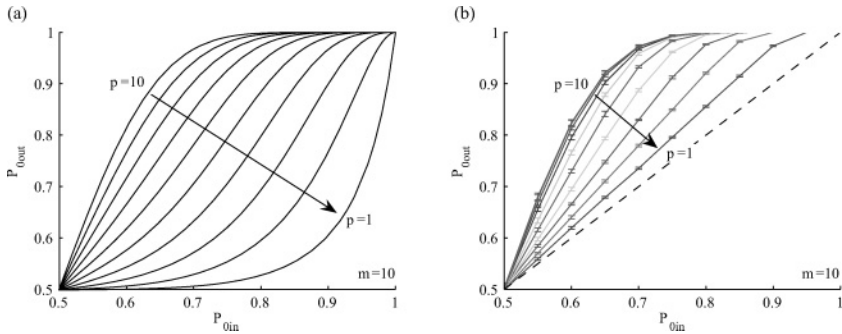


Figure 8: Effect of inhibition. Both graphs show the performance of the network if the strength of the inhibition is $V_I = p/nV_{th}$. In each graph, p is varied from 10 to 1 (top left to bottom curve, $m = 10$). (a) Numerical evaluation of the functional description that assumes that the membrane potential of the non-winning neuron is close to threshold before inhibition. (b) Simulation results. As can be seen by the simulation results, the assumption severely underestimates the performance of the network.

The non-winning neuron will then reach threshold with $p \approx V_I/V_E$ spikes. With this assumption, weakening the inhibition leads to a drastic decrease in performance (see Figure 8a). But depending on the history of the spike input to the non-winning neuron, the membrane potential will be significantly lower than V_{th} before the inhibition, that is, less inhibition is needed to achieve the same effect. We can address this in simulation (see Figure 8b), showing that the network performance does not decrease that rapidly. For $m = 10$, inhibition can be decreased to about 70% of V_{th} before the network shows a significant decrease in performance.

We conclude that weakening the strong inhibition always leads to a decrease in performance. For weak inhibition, the functional description underestimates the performance of the network, while simulation results show that inhibition can be decreased to about 70% of V_{th} before the network shows a significant decrease in performance.

4.5 Generalization to N Neurons. We extend the simplified case of two neurons to a network with N neurons. The probability of a neuron k generating an output spike starting from the initial condition where all neurons are reset is given by

$$P_{koutI} = \int_0^\infty v_k P(v_k t, n - 1) \cdot \prod_{\substack{j=1 \\ j \neq k}}^N \left(\sum_{i=0}^{n-1} P(v_j t, i) \right) dt. \tag{4.14}$$

The transitions can be described by the matrices (\mathbf{p}, \mathbf{t}) with the transition probabilities p_{kl} ($k, l \in 1, \dots, N$):

$$p_{kk} = \int_0^\infty v_k P(v_k t, m - 1) \cdot \prod_{\substack{j=1 \\ j \neq k}}^N \left(\sum_{i=0}^{p-1} P(v_j t, i) \right) dt$$

$$p_{kl} = \int_0^\infty v_l P(v_l t, p - 1) \cdot \sum_{i=0}^{m-1} P(v_k t, i) \cdot \prod_{\substack{j=1 \\ j \neq k, l}}^N \left(\sum_{i=0}^{p-1} P(v_j t, i) \right) dt,$$

with average transition times

$$\overline{t_{kk}} = \frac{m}{v_k}, \quad \overline{t_{kl}} = \frac{p}{v_l}.$$

The probability of an output spike of the network originating from a neuron k is then given by a vector \mathbf{P} . In the equilibrium state, \mathbf{P}^* is the first eigenvector of matrix \mathbf{p} .

5 Time-Varying Firing Rates

We extend our analysis of stationary inputs to inputs with time-varying firing rates in two computational tasks that involve the WTA. In the first task, we evaluate the discrimination performance of a neuron to rate changes in Poisson-distributed inputs (see section 5.1), in particular, we look at the effect of hysteresis from self-excitation, and the subsequent impact of self-excitation on the switching time of the neuron to the input rate change. In the second task, we look at the ability of the network to accurately reconstruct the position of a traveling object that generates an input wave across the neurons of the network (see section 5.2). In this case, each neuron sees a time-varying Poisson input rate due to the movement of the object. For such inputs, each neuron will receive only a finite number of input spikes for a finite duration of time. As a result, the number n of input spikes that a neuron can integrate to make a decision is constrained by the dynamics of the input to the network.

In addition, because of the Poisson statistics of the inputs, one cannot make assumptions on when events occur in a given time interval, but one can make assumptions about the number of spikes within this interval. This value is given by the Poisson rate of the input. Hence, we can replace the number of spikes vt in equation 4.1 with the integral $\int_0^t v(t') dt'$ that describes the area under the input spike rate function over time. This substitution can also be done for the generalization case in section 4.3, and

equation 4.14 becomes

$$P_{koutI} = \int_0^\infty P\left(\int_0^t v_k(t') dt', n-1\right) \cdot v_k(t) \cdot \prod_{\substack{j=0 \\ j \neq k}}^N \left(\sum_{i=0}^{n-1} P\left(\int_0^t v_j(t') dt', i\right)\right) dt. \tag{5.1}$$

5.1 Switching of Stronger Input. We first discuss the case where the input with the higher rate switches between the two neurons in a network with self-excitation. Using again a two-neuron network as in Figure 6, we assume that neuron 1 initially receives the higher rate input, that is, $v_1 > v_0$. At time $t = 0$, neuron 0 now receives the higher rate input, $v_0 > v_1$. We also assume that the duration of the inputs is long enough before the switching so that the network has a stable output, that is, the output spikes of the network come from neuron 1.

We then compute how many spikes neuron 1 will emit after time $t = 0$ before neuron 0 will spike. The probability of having exactly k subsequent spikes from neuron 1 is the probability of staying k times in state 1 (see Figure 6) before transitioning to state 0:

$$(p_{11})^k \cdot p_{10}. \tag{5.2}$$

Summing over all possible k results in the average number of spikes k_1 that neuron 1 emits before neuron 0 spikes:

$$k_1 = p_{10} \sum_{k=1}^\infty k (p_{11})^k = p_{10} \frac{p_{11}}{(1 - p_{11})^2} = \frac{p_{11}}{p_{10}}. \tag{5.3}$$

For every output spike, neuron 1 integrates m input spikes that arrive with rate v_1 . At the moment when the switch from neuron 1 to neuron 0 occurs, neuron 0 has to integrate additional p input spikes with rate v_0 . The corresponding average time taken for the winner to switch from neuron 1 to neuron 0 is

$$\bar{t}_{10} = \frac{p_{11}}{p_{10}} \frac{m}{v_1} + \frac{p}{v_0}. \tag{5.4}$$

The average switching time also increases with the self-excitation.

5.1.1 Classification Performance. To quantify the classification performance of the WTA network, we consider the first spike of neuron 0 as

the criterion for a correct detection of the changed inputs. This decision is described by two variables:

- $P_{TP}(t)$ —the probability that neuron 0 spikes in time t after the inputs are switched, that is, the probability that the network indicates a correct decision (*true positive*)
- $P_{FP}(t)$ —the probability that neuron 0 spikes in time t if the rates did not switch, that is, the probability that the network indicates a switch even though the rates did not change (*false positive*).

With $v_0 > v_1$ after the switch, the fraction of a true positive after time t can be calculated from the probability that no switch is detected and only neuron 1 is firing. In time t , neuron 1 receives on average $v_1 t$ input spikes and makes on average $v_1 t/m$ output spikes. This is equivalent to the case that the network stays $v_1 t/m$ times in state 1 before transitioning to state 0:

$$P_{TP}(t) = 1 - (p_{11})^{\frac{v_1 t}{m}}. \quad (5.5)$$

A false positive is a state transition from state 0 to state 1 if we keep the convention that $v_0 > v_1$. We model this case again by the probability of the contrary event that no transition occurs from the neuron with the stronger input (neuron 0):

$$P_{FP}(t) = 1 - (p_{00})^{\frac{v_0 t}{m}}. \quad (5.6)$$

Figure 9 shows the true and false-positive rates similar to the receiver operating characteristic curve (ROC). The ROC curves quantify the dependence of the classifier performance on a parameter. Varying this parameter determines the true and false-positive probabilities of the classifier. For our classifier, this parameter is time, and it corresponds to the different values of P_{TP} and P_{FP} . The area under the ROC curve quantifies the classifier performance. We follow the engineering definition of the discrimination performance, which is the area under the ROC curve measured between the curve and chance level (dashed line).

The dependence of the discrimination performance and average switching time as a function of n and m for the two-neuron network is shown in Figure 10. In general, the discrimination performance increases if n increases or if self-excitation is increased for a fixed value of n . However, the switching time also increases in both cases. Interestingly, for the same average switching time, a network with a larger n and no self-excitation has a better discrimination performance than a network with a smaller n and self-excitation.

5.2 Traveling Input. Real-time systems that interact with the world usually have to process nonstationary input, for example, retina spikes from an object moving in space or auditory spikes from music and speech. Here, we




Figure 9: Classifier characteristics: Probability of true positives versus false positives, similar to the common receiver-operating-characteristics (ROC) curves. The network has to detect a switch in the neuron that receives the stronger input. This input has a spike rate that is 50% higher than the other input. The different curves corresponding to different amounts of self-excitation show the classification performance dependent on the time when the network spikes first after the switch. The area between curve and chance level (dashed line) quantifies the classifier performance, the discrimination performance.

show the analysis of a situation where a network has to determine with high accuracy the position of a ball traveling at a constant speed. We model the input by a gaussian-shaped wave traveling across the neurons. This simplifying description of the input is a good description of the output spikes from one of the feature extraction stages of a multichip asynchronous vision system (Oster, Douglas, & Liu, 2007). It can also be a general description of the outputs of temporal feature filters.

We analyze how we can set n , the number of input spikes that a neuron needs to reach threshold, so that we obtain optimal discriminability performance of the WTA network. In this scenario, this discriminability measure corresponds to a neuron firing only once and exactly at the time when the maximum of the gaussian wave moves over that neuron. Since the input scenario corresponds to the overlap of gaussian inputs in space and time, the discrimination analysis is related to the discriminability measure used in psychophysics.

We assume that the neurons are arranged in a one-dimensional array with numbers $i = -\infty, \dots, \infty$ (see Figure 11). Without loss of generality, our analysis focuses on neuron 0. The temporal distance between the neurons




Figure 10: Classifier characteristics: Discrimination performance versus the average switching time. The discrimination performance is computed from Figure 9. Each curve corresponds to a particular number of spikes needed for the neurons to reach threshold. Along each curve, self-excitation is increased by increasing m . The left-most data point of each curve corresponds to zero self-excitation ($m = n$). Each subsequent data point along the curve corresponds to an increase of self-excitation by one, meaning that the winning neuron needs one fewer spike to reach threshold ($m = n, \dots, 1$). For about the same average switching time, the discrimination performance is better if more spikes are needed to reach threshold and the neuron receives less self-excitation (see the indicated points). At the point ($n = 4, m = 3$), the neurons need four spikes to reach threshold and receive self-excitation equivalent to one spike. At the point ($n = 5, m = 5$), the neurons need five spikes to reach threshold and do not receive self-excitation. Similar comparisons can be made for any other pair of data points of (no self-excitation/higher n) versus (self-excitation/smaller n).

is d , that is, the time the center of the ball needs to travel from one neuron to the next. The time-dependent input to a neuron i , $v_i(t)$, is a gaussian function centered on that neuron,

$$v_i(t) = \frac{v_{\max}}{2\pi\sigma^2} \exp\left(-\frac{(t - di)^2}{2\sigma^2}\right), \tag{5.7}$$

where v_{\max} is the maximum spike rate, σ is the variance of the gaussian, and di is the temporal distance of neuron i from neuron 0. The spike rate is normalized so that the neuron receives the rate v_{\max} when the input wave is aligned to the neuron.

We can compare the temporal distance d to the discriminability measure \hat{d} used in psychophysics: it measures the distance between the center of two gaussian distributions that are normalized to $\sigma = 1$. \hat{d} is the overlap between the distributions, which reflects the difficulty in separating the two distributions. In our example, d is the distance between two neurons that receive spike input consecutively. Nevertheless, d is also a measure of the difficulty in classifying the ball position like the discriminability measure in psychophysics.

The performance of the WTA is best if the output is as sparse as possible. Since we require that each neuron makes, on average, one spike when the object passes over it, the average integration time of the WTA decision is then the time needed for the object to move from one neuron to the next, d . It is natural to center this interval on the neuron, for neuron 0 at $t = 0$. In this time, the neuron receives a higher firing than all its neighbors (see Figure 11). Integration of inputs to the WTA starts then at the intersection of two neighboring input firing rates at $t = -d/2$. We use the variable T to indicate the integration time ($T = t + d/2$). To start the analysis, we assume that all neurons are reset at this time.

We define an auxiliary function for the area under the input spike rate function for each neuron i in the integration time T :

$$R_i(T) = \int_{-d/2}^{T-d/2} v_i(t) dt = \frac{v_{\max}}{2\pi\sigma^2} \int_{-d/2}^{T-d/2} \exp\left(-\frac{(t-di)^2}{2\sigma^2}\right) dt. \quad (5.8)$$

On average, the winner should spike at $T = d$ after receiving n spikes:

$$n = R_0(d) = \frac{v_{\max}}{2\pi\sigma^2} \int_{-d/2}^{d/2} \exp\left(-\frac{t^2}{2\sigma^2}\right) dt. \quad (5.9)$$

We will refer most of our analysis to n , instead of v_{\max} , so we can make comparisons to section 3.

We rewrite equation 5.1 for neuron 0 in a line of neurons $i = -\infty, \dots, \infty$ and with the corresponding area under the input spike rate function $R_i(T)$:

$$P_{0out|I} = \int_0^\infty P(R_0(T), n-1) \cdot v_0(T-d/2) \prod_{\substack{k=-\infty \\ k \neq 0}}^{+\infty} \left(\sum_{i=0}^{n-1} P(R_k(T), i) \right) dT. \quad (5.10)$$

This is the probability that neuron 0 makes the first output spike after a reset at time $t = -d/2$.




Figure 11: Model of a traveling input. An object moves across a line of neurons of the WTA network, producing a wave of activity. The neurons are equidistantly spaced: assuming a constant velocity of the object, this results in the temporal distance d . The neurons are numbered from $-\infty$ to ∞ . The diagram shows the spike rate evolving over time. At a certain point in time (thick dashed line), neuron 0 receives an input spike rate indicated by the curve that peaks at its position. Integration of the WTA network starts at the intersection of two neighboring input curves. At an integration time T , the neuron 0 aligned to the ball position receives a number of spikes equivalent to the light gray area under its spike input curve, whereas a neighboring neuron receives a number of spikes equivalent to the dark gray area.

The rate of neuron 0 making the first output spike of the network after $d/2$ is equivalent to the rate of a correct decision:

$$r_{correct}(T) = P(R_0(T), n - 1) \cdot v_0(T - d/2) \cdot \prod_{\substack{k=-\infty \\ k \neq 0}}^{+\infty} \left(\sum_{i=0}^{n-1} P(R_k(T), i) \right). \tag{5.11}$$

The rate of an incorrect decision of the network is

$$r_{false}(T) = \sum_{\substack{j=-\infty \\ j \neq 0}}^{+\infty} P(R_j(T), n-1) \cdot v_j(T-d/2) \cdot \prod_{\substack{k=-\infty \\ k \neq j}}^{+\infty} \left(\sum_{i=0}^{n-1} P(R_k(T), i) \right). \quad (5.12)$$

How can we infer the ability of the network to reconstruct the position of the stimulus from the probability of a correct output spike and the output rate? The network discretizes the stimulus position in time and space: in time, because the output of the network is event or spike based, and in space, since the network can report the stimulus position only at every neuron. The ideal result of this discretization is a staircase function. This is equivalent to decoding the output of the network with a memory of one variable, the position, which is updated at every spike. Of course, one could think of a more elaborate processing of the network output such as spatial and temporal interpolation. But this interpolation is limited by the accuracy of the network output, which is exactly what we analyze here.

Deviation of the stimulus position reconstructed from network output from the ideal case results from two types of errors: jitter in the timing of the output spikes and spikes from neurons that are not aligned to the stimulus position. We will summarize the effect of both by defining the total position error e as the area between the stimulus position reconstructed from the network output and the ideal case. We norm this area to the number of neurons, so that e denotes an average deviation of the reconstructed stimulus position per neuron.

The jitter error is caused by variation of the time when the first neuron of the network reaches threshold. We define e_{jitter} as the time between the first output spike of the network and the average time the network needs to make an output spike if the ball moves from one position to the next. e_{jitter} is normalized to the temporal distance of the neurons d :

$$e_{jitter} = \frac{1}{d} \int_{-\infty}^{+\infty} |T-d| (r_{correct}(T) + r_{false}(T)) dT. \quad (5.13)$$

Spikes from neurons that are not aligned to the stimulus position contribute to the classification error e_{class} . We assume that if the network decides on an incorrect stimulus position, this position will be kept until the next spike of the network at time d later. We define the error as the area between the correct position 0 and j , weighted by the probability P_{jout} that neuron

j makes the first output spike of the network:

$$\begin{aligned}
 e_{class} &= \sum_{j=-\infty}^{+\infty} |j| P_{jout} \\
 &= \sum_{j=-\infty}^{+\infty} |j| \int_0^{\infty} P(R_j(T), n-1) \cdot v_j(T-d/2) \\
 &\quad \cdot \prod_{\substack{k=-\infty \\ k \neq j}}^{+\infty} \left(\sum_{i=0}^{n-1} P(R_k(T), i) \right) dT.
 \end{aligned} \tag{5.14}$$

We can add these two errors together to obtain the total position error e as defined before:

$$e \approx e_{jitter} + e_{class} \tag{5.15}$$

However, there are two problems in this approximation: the jitter area and the classification error overlap, which means that the same area is accounted for twice. This is dominant when both jitter and classification errors are large for small n . The second problem is our assumption of the start of the integration at the intersection of the input spike rate functions of two neighboring neurons ($t = -d/2$). Due to the jitter in the first output spike, the next integration period will actually start earlier or later.

Results of the position error obtained by evaluation of the functional description therefore cannot be directly compared to simulation data. Nevertheless, the data fit qualitatively (see Figure 12).

6 Discussion

We present an analytical study for quantifying the discrimination performance of a spike-based WTA network based on different network parameters and spiking input statistics. Our work is in contrast to previous analyses of WTA behavior that considered only analog or spike-rate coded inputs. Rate coding schemes are useful for stationary inputs but fail to describe the dynamic input signals that occur in real-time systems for time-varying stimuli. In particular, these studies do not consider the transition region between single spike and spike rate coding in a WTA network.

We assume a nonleaky integrate-and-fire neuron model in our analysis in part because we could formulate closed-form solutions for the input statistics considered and also in part because the VLSI implementation follows this model. The effect of adding leakage would be rather small, at least if we added linear leakage. Adding leakage with an exponential dependency of the membrane voltage (the same effect as conductance-based synapses) would lead to a more complex neuron behavior exhibiting




Figure 12: Area errors dependent on the number of spikes n to reach threshold. Shown are the jitter error (variation of the timing of the output spikes, diamonds) and the classification error (output spikes from neurons not aligned to the stimulus position, circles), from numerical evaluations of equations 5.13 and 5.15. For comparison, the total area error is shown from simulation results (crosses). In simulations, the complete line of neurons is processed, so that the start of the integration period depends on the time of the last output spike, whereas in the functional description, the integration always starts at the optimal intersection of two neighboring input waveforms. For low n , jitter and classification error include the same area twice; therefore the sum of both would exceed the error obtained in the simulation results. Error bars of the simulation data denote one standard deviation ($d = 1, \sigma = 1$, every data point result of 100 trials with 20 neurons).

random walk, as discussed before. Unfortunately, an analysis with closed-form equations and inequalities would not be feasible in this case, but the network would have to be simulated.

We have also not considered the influence of noise on the membrane because of the following two reasons. First, we consider only excitatory external input, and, second, our neuron model does not include conductance-based synapses. Therefore, the membrane potential does not have a stable intermediate voltage as in a random walk case where the neurons can quickly reach threshold with any additional input, including

noise input. Since the considered inhibition in our network is strong, the other neurons besides the winner are discharged to $V = 0$. The condition that all neurons (except the winner) have a membrane potential of zero will therefore happen frequently—after every output spike of the network. This behavior is specific to the chosen model of neuron and synapses, which was a prerequisite to a closed formulation of the winner-take-all behavior.

Our analysis is centered on spike numbers and spike times, thus making this analysis extendable for any size network. This analysis is especially relevant for the development of neuromorphic electronic spiking networks used in engineering applications. For example, we used this theoretical work to predict the outputs of our hardware VLSI implementation of the WTA network in an asynchronous multichip spiking system called CAVIAR (Serrano-Gotarredona et al., 2005). The CAVIAR system consists of a temporal contrast retina (Lichtsteiner, Pösch, & Delbrück, 2008), a bank of spike-based convolution filters (Serrano-Gotarredona, Serrano-Gotarredona, Acosta-Jimenez, & Linares-Barranco, 2006), a winner-take-all network (Oster et al., 2008), and a learning module (Häfliger, 2007) that classifies trajectories of moving objects.

The WTA module in CAVIAR computes the best feature represented in the image based on the spike outputs of the convolution filters and also computes the locations of this best feature. When the system sees a moving object, the input to the WTA network is represented by a spike wave of gaussian shape that travels across the neurons. We show that the output of the higher stages of CAVIAR can be well approximated by Poisson statistics (Oster et al., 2007), although the retina and convolution chips are completely deterministic. This is the same input that we used for our analysis of WTA behavior in section 5.2, and we can compare the predicted performance of our theoretical model with the performance of the implemented model in CAVIAR. The chip's performance follows the theoretical prediction for inputs of constant currents, regular spike rates, and Poisson spike trains (Oster et al., 2008). The achieved performance, that is, the detection of the position of a moving object by the WTA module, is close to optimal in the case of nonstationary Poisson statistics (Oster et al., 2007).

Fast detection tasks that can be performed by a system like CAVIAR with a finite number of neurons ($\approx 32,000$ neurons) do not allow time for a system to encode inputs as spike rates. Single spikes should be considered instead. Our analytical approach can be used to evaluate systems whose outputs range from single spike codes to spike rate coding by quantifying the performance of the classifier dependent on the number the input spikes the network needs to reach a decision, that is, to elicit an output spike. It is also applicable to systems that do not make any assumptions about the specific coding scheme, for example, in the visual information processing system, SpikeNET, in which signals are encoded in a rank-order code (Van Rullen & Thorpe, 2001, 2002).

7 Conclusion

Spike-based networks that capture the asynchronous and time-continuous computation inherent in biological nervous systems can provide a powerful alternate technology to today's digital processors. Applying the principles of biological processing to engineering applications requires a thorough understanding of the underlying computation: the processing architecture, the range of network and circuit parameters, and the resulting performance.

We made steps in that direction by using a simplified Markov model of the spiking network to examine analytically the ability of a spike-based WTA network to discriminate the statistics of inputs ranging from stationary regular to nonstationary Poisson events. Our work extends previous theoretical results showing that a WTA recurrent network receiving regular spike inputs can select the correct winner within one interspike interval. We showed that for the case of Poisson spike inputs, the discrimination performance of the network (i.e., the probability of the network making a correct decision) increases as self-excitation is increased, but as expected, the self-excitation leads to hysteresis, which means that the switching time for the network to detect a change in input rates increases. We find that weak inhibition primarily decreases the network performance.

We also extended this discrimination analysis of spiking WTAs to nonstationary inputs with time-varying spike rates resembling statistics of real-world sensory stimuli. We used this analysis to predict the performance of a WTA chip in a large-scale, multichip, asynchronous vision system. We conclude that spiking WTAs also exhibit high discrimination performance with nonstationary inputs.

Acknowledgments

This work was supported in part by EU grants IST-2001-34124 and FP6-2005-015803. We acknowledge Sebastian Seung for discussions on the winner-take-all mechanism. We also acknowledge Prashanth D'Souza and YingXue Wang for a careful reading of the draft and the anonymous referees for their helpful comments.

References

- Abrahamsen, J. P., Häfliger, P., & Lande, T. S. (2004). A time domain winner-take-all network of integrate-and-fire neurons. In *Proceedings of the 2004 IEEE International Symposium on Circuits and Systems* (vol. 5, pp. 361–364). Piscataway, NJ: IEEE Press.
- Amari, S., & Arbib, M. (1982). Competitive and cooperative aspects in dynamics of neural excitation and self-organization. In S. Amari & M. Arbib (Eds.), *Competition and cooperation in neural networks* (pp. 1–28). New York: Springer-Verlag.

- Barnden, J., & Srinivas, K. (1993). Temporal winner-take-all networks: A time-based mechanism for fast selection in neural networks. *IEEE Transactions on Neural Networks*, 4(5), 844–853.
- Bartolozzi, C., & Indiveri, G. (2004). A neuromorphic selective attention architecture with dynamic synapses and integrate-and-fire neurons. In *Proceedings of Brain Inspired Cognitive Systems (BICS 2004)* (Vol. BIS2.2, pp. 1–6). Stirling: University of Stirling.
- Carota, L., Indiveri, G., & Dante, V. (2004). A software-hardware selective attention system. *Neurocomputing*, 58(60), 647–653.
- Chicca, E., Indiveri, G., & Douglas, R. (2004). An event-based VLSI network of integrate-and-fire neurons. In *Proceedings of the 2004 IEEE International Symposium on Circuits and Systems* (Vol. 5, pp. 357–360). Piscataway, NJ: IEEE Press.
- Chicca, E., Lichtsteiner, P., Delbruck, T., Indiveri, G., & Douglas, R. J. (2006). Modeling orientation selectivity using a neuromorphic multi-chip system. In *Proceedings of the 2006 IEEE International Symposium on Circuits and Systems* (pp. 1235–1239). Piscataway, NJ: IEEE Press.
- Choi, T., Merolla, P., Arthur, J., Boahen, K., & Shi, B. (2005). Neuromorphic implementation of orientation hypercolumns. *IEEE Transactions on Circuits and Systems I*, 52(6), 1049–1060.
- Coultrip, R., Granger, R., & Lynch, G. (1992). A cortical model of winner-take-all competition via lateral inhibition. *Neural Networks*, 5(1), 47–54.
- Douglas, R., Koch, C., Mahowald, M., Martin, K., & Suarez, H. (1995). Recurrent excitation in neocortical circuits. *Science*, 269(5226), 981–985.
- Ermentrout, B. (1992). Complex dynamics in winner-take-all neural nets with slow inhibition. *Neural Networks*, 5(3), 415–431.
- Gautrais, J., & Thorpe, S. (1998). Rate coding versus temporal order coding: A theoretical approach. *Biosystems*, 48(1), 57–65.
- Häfliger, P. (2007). Adaptive WTA with an analog VLSI neuromorphic learning chip. *IEEE Transactions on Neural Networks*, 18(2), 551–572.
- Hahnloser, R. H. R., Sarpeshkar, R., Mahowald, M. A., Douglas, R. J., & Seung, S. (2000). Digital selection and analogue amplification coexist in a cortex-inspired silicon circuit. *Nature*, 405(6789), 947–951.
- He, Y., & Sanchez-Sinencio, E. (1993). Min-net winner-take-all CMOS implementation. *Electronics Letters*, 29(14), 1237–1239.
- Indiveri, G. (1997). Winner-take-all networks with lateral excitation. *Analog Integrated Circuits and Signal Processing*, 13(1), 185–193.
- Indiveri, G. (2001). A current-mode hysteretic winner-take-all network with excitatory and inhibitory coupling. *Analog Integrated Circuits and Signal Processing*, 28(3), 279–291.
- Indiveri, G., Horiuchi, T., Niebur, E., & Douglas, R. (2001). A competitive network of spiking VLSI neurons. In F. Rattay (Ed.), *World Congress on Neuroinformatics* (pp. 443–455). Vienna, Austria: ARGESIM/ASIM Verlag.
- Itti, L., Koch, C., & Niebur, E. (1998). A model of saliency-based visual attention for rapid scene analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(11), 1254–1259.
- Jain, B., & Wysotzki, F. (2004). Discrimination networks for maximum selection. *Neural Networks*, 17(1), 143–154.

- Jin, D. Z., & Seung, H. S. (2002). Fast computation with spikes in a recurrent neural network. *Phys. Rev. E*, 65(5), 051922-1–051922-4.
- Kaski, S., & Kohonen, T. (1994). Winner-take-all networks for physiological models of competitive learning. *Neural Networks*, 7(6-7), 973–984.
- Kincaid, T., Cohen, M., & Fang, Y. (1996). Dynamics of a winner-take-all neural network. *Neural Networks*, 9(7), 1141–1154.
- Lazzaro, J., Ryckebusch, S., Mahowald, M. A., & Mead, C. A. (1989). Winner-take-all networks of O(n) complexity. In D. Touretzky (Ed.), *Advances in neural information processing systems*, 1 (pp. 703–711). San Mateo, CA: Morgan Kaufmann.
- Lichtsteiner, P., Posch, C., & Delbrück, T. (2008). A 128×128 120dB 15 μ s latency asynchronous temporal contrast vision sensor. *IEEE Journal of Solid-State Circuits*, 43(2), 566–576.
- Liu, S.-C. (2000). A winner-take-all circuit with controllable soft max property. In S. A. Solla, T. K. Leen, & K.-R. Müller (Eds.), *Advances in neural information processing systems*, 12 (pp. 717–723). Cambridge, MA: MIT Press.
- Liu, S.-C. (2002). A normalizing aVLSI network with controllable winner-take-all properties. *Analog Integrated Circuits and Signal Processing*, 31(1), 47–53.
- Lumer, E. (2000). Effects of spike timing on winner-take-all competition in model cortical circuits. *Neural Computation*, 12, 181–194.
- Maass, W. (1999). Neural computation with winner-take-all as the only nonlinear operation. In M. Kearns, S. Solla, & D. Cohn (Eds.), *Advances in neural information processing systems*, 11 (pp. 293–299). Cambridge, MA: MIT Press.
- Maass, W. (2000). On the computational power of winner-take-all. *Neural Computation*, 12, 2519–2535.
- Moller, R., Maris, M., Tomes, J., & Mojaev, A. (1998). A strong winner-take-all neural network in analogue hardware. In L. Smith & A. Hamilton (Eds.), *Neuromorphic systems: Engineering silicon from neurobiology* (pp. 178–190). Singapore: World Scientific.
- Oster, M., Douglas, R., & Liu, S.-C. (2007). Quantifying input and output statistics of a winner-take-all network in a vision system. In *Proceedings of the 2007 IEEE International Symposium on Circuits and Systems* (pp. 853–856). Piscataway, NJ: IEEE Press.
- Oster, M., Wang, X., Douglas, R., & Liu, S.-C. (2008). Quantification of a spike-based winner-take-all VLSI network. *IEEE Transactions on Circuits and Systems I*, 55(10), 3160–3169.
- Riesenhuber, M., & Poggio, T. (1999). Hierarchical models of object recognition in cortex. *Nature Neuroscience*, 2(11), 1019–1025.
- Serrano-Gotarredona, T., & Linares-Barranco, B. (1995). A modular current-mode high-precision winner-take-all circuit. *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, 42(2), 132–134.
- Serrano-Gotarredona, R., Oster, M., Lichtsteiner, P., Linares-Barranco, A., Paz-Vincent, R., Gomez-Rodríguez, F., et al. (2005). AER building blocks for multi-layer multi-chip neuromorphic vision systems. In Y. Weiss, B. Schölkopf, & J. Platt (Eds.), *Advances in neural information processing systems*, 18 (pp. 1217–1224). Cambridge, MA: MIT Press.
- Serrano-Gotarredona, R., Serrano-Gotarredona, T., Acosta-Jimenez, A., & Linares-Barranco, B. (2006). A neuromorphic cortical-layer microchip for spike-based

- event processing vision systems. *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications*, 53, 2548–2566.
- Starzyk, J., & Fang, X. (1993). CMOS current mode winner-take-all circuit with both excitatory and inhibitory feedback. *Electronics Letters*, 29(10), 908–910.
- Van Rullen, R., & Thorpe, S. (2001). Rate coding versus temporal order coding: What the retinal ganglion cells tell the visual cortex. *Neural Computation*, 13(6), 1255–1283.
- VanRullen, R., & Thorpe, S. (2002). Surfing a spike wave down the ventral stream. *Vision Research*, 42(23), 2593–2615.
- Vogelstein, R., Mallik, U., Culurciello, E., Cauwenberghs, G., & Etienne-Cummings, R. (2007). A multichip neuromorphic system for spike-based visual information processing. *Neural Computation*, 19, 2281–2300.
- Wang, D. (1999). Object selection based on oscillatory correlation. *Neural Networks*, 12(4), 579–592.
- Yu, A., Giese, M., & Poggio, T. (2002). Biophysiological plausible implementations of the maximum operation. *Neural Computation*, 14(12), 2857–2881.
- Yuille, A., & Geiger, D. (1998). *The handbook of brain theory and neural networks*. Cambridge, MA: MIT Press.
- Yuille, A., & Grzywacz, N. (1989). A winner-take-all mechanism based on presynaptic inhibition feedback. *Neural Computation*, 1, 334–347.

Received July 22, 2008; accepted February 19, 2009.