



CONTRIBUTED ARTICLE

Networks of Spiking Neurons: The Third Generation of Neural Network Models

WOLFGANG MAASS

Institute for Theoretical Computer Science, Technische Universität Graz

(Received 27 March 1996; accepted 10 November 1996)

Abstract—The computational power of formal models for networks of spiking neurons is compared with that of other neural network models based on McCulloch Pitts neurons (i.e., threshold gates), respectively, sigmoidal gates. In particular it is shown that networks of spiking neurons are, with regard to the number of neurons that are needed, computationally more powerful than these other neural network models. A concrete biologically relevant function is exhibited which can be computed by a single spiking neuron (for biologically reasonable values of its parameters), but which requires hundreds of hidden units on a sigmoidal neural net. On the other hand, it is known that any function that can be computed by a small sigmoidal neural net can also be computed by a small network of spiking neurons. This article does not assume prior knowledge about spiking neurons, and it contains an extensive list of references to the currently available literature on computations in networks of spiking neurons and relevant results from neurobiology. © 1997 Elsevier Science Ltd. All rights reserved.

Keywords—Spiking neuron, Integrate-and-fire neuron, Computational complexity, Sigmoidal neural nets, Lower bounds.

1. DEFINITIONS AND MOTIVATIONS

If one classifies neural network models according to their computational units, one can distinguish three different generations. The *first generation* is based on *McCulloch–Pitts neurons* as computational units. These are also referred to as perceptrons or threshold gates. They give rise to a variety of neural network models such as multilayer perceptrons (also called threshold circuits), Hopfield nets, and Boltzmann machines. A characteristic feature of these models is that they can only give *digital* output. In fact they are *universal* for computations with digital input and output, and every boolean function can be computed by some multilayer perceptron with a single hidden layer.

The *second generation* is based on computational units that apply an “activation function” with a *continuous* set of possible output values to a weighted sum (or polynomial) of the inputs. Common activation functions are the *sigmoid function* $\sigma(y) = 1/(1 + e^{-y})$ and the linear

saturated function π with $\pi(y) = y$ for $0 \leq y \leq 1$, $\pi(y) = 0$ for $y < 0$, $\pi(y) = 1$ for $y > 1$. Besides piecewise polynomial activation functions we consider in this paper also “piecewise exponential” activation functions, whose pieces can be defined by expressions involving exponentiation (such as the definition of σ). Typical examples for networks from this second generation are feedforward and recurrent sigmoidal neural nets, as well as networks of radial basis function units. These nets are also able to compute (with the help of thresholding at the network output) arbitrary boolean functions. Actually it has been shown that neural nets from the second generation can compute certain boolean functions with *fewer* gates than neural nets from the first generation (Maass, Schnitger, & Sontag, 1991; DasGupta & Schnitger, 1993). In addition, neural nets from the second generation are able to compute functions with *analog* input and output. In fact they are *universal* for analog computations in the sense that any continuous function with a compact domain and range can be approximated arbitrarily well (with regard to uniform convergence, i.e., the L_∞ – norm) by a network of this type with a single hidden layer. Another characteristic feature of this second generation of neural network models is that they support learning algorithms that are based on gradient descent such as backprop.

Acknowledgements: I would like to thank Eduardo Sontag and an anonymous referee for their helpful comments. Written under partial support by the Austrian Science Fund.

Requests for reprints should be sent to W. Maass, Institute for Theoretical Computer Science, Technische Universität Graz, Klosterwiesgasse 32/2, A-8010, Graz, Austria; tel. +43 316 873-5822; fax: +43 316 873-5805; e-mail: maass@igi.tu-graz.ac.at




FIGURE 1. Simultaneous recordings (over 4 sec) of the firing times of 30 neurons from monkey striate cortex by Krüger & Aiple (1988). Each firing is denoted by a short vertical bar, with a separate row for each neuron. For comparison we have marked the length of an interval of 100 msec by two vertical lines. This time span is known to suffice for the completion of some complex multilayer cortical computations.

For a biological interpretation of neural nets from the second generation one views the output of a sigmoidal unit as a representation of the current firing rate of a biological neuron. Since biological neurons, especially in higher cortical areas, are known to fire at various intermediate frequencies between their minimum and maximum frequency, neural nets from the second generation are, with regard to this “firing rate interpretation”, biologically more realistic than models from the first generation.

However, at least with regard to *fast* analog computations by networks of neurons in the cortex, the “firing rate interpretation” itself has become questionable. Perrett, Rolls, and Caan (1982) and Thorpe and Imbert (1989) have demonstrated that visual pattern analysis and pattern classification can be carried out by humans in just 100 msec, in spite of the fact that it involves a minimum of 10 synaptic stages from the retina to the temporal lobe (see Figure 1.) The same speed of visual processing has been measured by Rolls and Tovee (1994) in macaque monkeys. Furthermore, they have shown that a single cortical area involved in visual processing can complete its computation in just 20–30 msec (Rolls, 1994; Rolls & Tovee, 1994). On the other hand, the firing rates of neurons involved in these computations are usually below 100 Hz, and hence at least 20–30 msec would be needed just to sample the current

firing rate of a neuron. Thus a coding of analog variables by firing rates seems quite dubious in the context of fast cortical computations.

On the other hand, experimental evidence has accumulated during the last few years which indicates that many biological neural systems use the timing of single action potentials (or “spikes”) to encode information (Abeles, 1991; Abeles, Bergman, Margalit, & Vaadia, 1993; Aertsen, 1993; Arbib, 1995; Bair, Koch, Newsome, & Britten, 1994; Bialek & Rieke, 1992; Ferster & Spruston, 1995; Hopfield, 1995; Kempter, Gerstner, van Hemmen, & Wagner, 1996; Lestienne, 1996; Rieke, Warland, van Stevenick, & Bialek, 1996; Sejnowski, 1995; Singer, 1995; Softky, 1994; Thorpe & Imbert, 1989).

These experimental results from neurobiology have lead to the investigation of a *third generation* of neural network models which employ *spiking neurons* (or “integrate-and-fire neurons”) as computational units. Recently, one has also started to carry out experiments with related new types of electronic hardware such as pulse stream VLSI (see, e.g., DeYong, Findley, & Fields, 1992; Douglas, Koch, Mahowald, Martin, & Suarez, 1995; Horinchi, Lazzaro, Moore, & Koch, 1991; Jahnke, Roth, & Klar, 1996; Jiu & Leong, 1996; Mahowald, 1992, 1994; Mead, 1989; Meador, Wu, Cole, Nintunze, & Chintrakulchai, 1991; Murray &

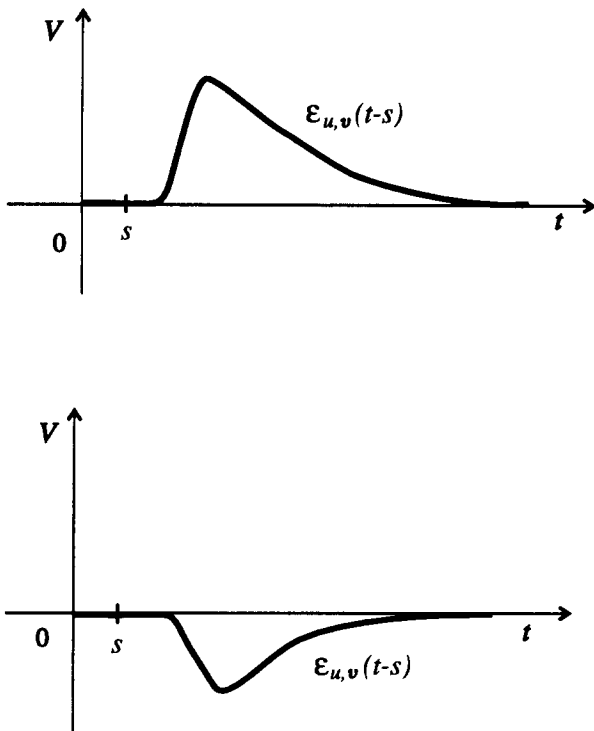


FIGURE 2. Typical shape of response functions (EPSP and IPSP) of a biological neuron.

Tarassenko, 1994; Northmore & Elias, 1996; Pratt, 1989; Zaghloul, Meador, & Newcomb, 1994). In these new chips one can encode analog variables by *time differences* between pulses, which has practical advantages over other encoding methods. The goal of understanding the capabilities and limitations of this new type of analog neural hardware provides additional motivation for theoretical investigation of the third generation of neural network models.

One may also view threshold circuits (i.e., neural nets from the first generation) as abstract models for digital computation on networks of spiking neurons, where the bit 1 is coded by the firing of a neuron within a certain short time window, and 0 by the non-firing of this neuron within this time window (see e.g., Valiant, 1994). However, under this coding scheme a threshold circuit provides a reasonably good model for a network of spiking neurons only if the firing times of all neurons that provide the input bits for another spiking neuron are *synchronized* (up to a few msec). Apparently such strongly synchronized activity does occur in biological neural systems (see Abeles et al., 1993; Bair et al., 1994) but many argue that it is not their typical mode of operation.

Mathematical models for “integrate-and-fire neurons” (or “spiking neurons” as they have been called more recently) can be traced back to Lapique (1907) (see Tuckwell, 1988). There exist a number of variations of this model, which are described and compared in a recent survey (see Gerstner, 1995). With regard to the relationship of these mathematical models

to the known behaviour of biological neurons we refer to Abeles (1991); Aertsen (1993); Arbib (1995); Bower and Beeman (1995); Churchland and Sejnowski (1993); Hopfield (1995); Johnston and Wu (1995); Rieke et al. (1996); Shepherd (1990, 1994); Tuckwell (1988); and Taylor and Alavi (1993). These mathematical models for spiking neurons do not provide a complete description of the extremely complex computational function of a biological neuron. Rather, like the computational units of the previous two generations of neural network models, these are *simplified* models that focus on just a few aspects of biological neurons. However, in comparison with the previous two models they are substantially more realistic. In particular, they describe much better the actual *output* of a biological neuron, and hence they allow us to investigate on a theoretical level the possibilities of using *time* as a resource for computation and communication. Whereas the timing of computation steps is usually “trivialized” in the models from the preceding two generations (either through an assumed synchronization, or through an assumed stochastic asynchronicity), the timing of individual computation steps plays a key role for computations in networks of spiking neurons. In fact, the output of a spiking neuron v consists of the set $F_v \subseteq \mathbf{R}^+$ of points in time when v “fires” (where $\mathbf{R}^+ = \{x \in \mathbf{R} : x \geq 0\}$).

In the simplest (deterministic) model of a spiking neuron one assumes that a neuron v fires whenever its “potential” P_v (which models the electric membrane potential at the “trigger zone” of neuron v) reaches a certain threshold Θ_v . This potential P_v is the sum of so-called excitatory postsynaptic potentials (“EPSPs”) and inhibitory postsynaptic potentials (“IPSPs”), which result from the firing of other neurons u that are connected through a “synapse” to neuron v . The firing of a “presynaptic” neuron u at time s contributes to the potential P_v at time t an amount that is modelled by the term $w_{u,v} \epsilon_{u,v}(t-s)$, which consists of a “weight” $w_{u,v} \geq 0$ and a *response function* $\epsilon_{u,v}(t-s)$. Biologically realistic shapes of such response functions are indicated in Figure 2.

The “weight” $w_{u,v} \geq 0$ in the term $w_{u,v} \epsilon_{u,v}(t-s)$ reflects the “strength” (called “efficacy” in neurobiology) of the synapse between neuron u and neuron v . In the context of *learning* one can replace $w_{u,v}$ by a *function* $w_{u,v}(t)$. In addition it has been conjectured that rapid changes of the value of $w_{u,v}(t)$ are also essential for *computations* in biological neural systems. However for simplicity we view here $w_{u,v}$ just as a constant.

The restriction of $w_{u,v}$ to non-negative values is motivated by the assumption that a biological synapse is either “excitatory” or “inhibitory”, and that it does not change its “sign” in the course of a “learning process”. In addition, for most biological neurons u , either all response functions $\epsilon_{u,v}(t-s)$ for postsynaptic neurons v are “excitatory” (i.e., positive), or all of them are “inhibitory” (i.e., negative). Obviously these constraints

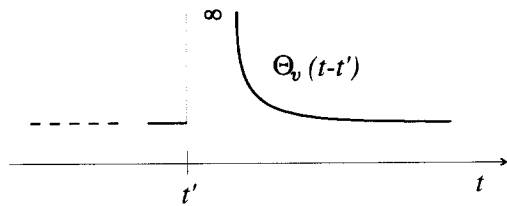


FIGURE 3. Typical shape of the threshold function of a biological neuron.

have basically no impact on theoretical complexity investigations (just consider pairs of excitatory and inhibitory neurons instead of single neurons), unless one cares about small constant factors in the size of networks, or one wants to model the actual architecture of cortical circuits (see Douglas et al., 1995; Shepherd, 1990).

It is mathematically more convenient to assume that the potential P_v has value 0 in the absence of postsynaptic potentials, and that the threshold value Θ_v is always > 0 . In a “typical” biological neuron the resting membrane potential is around -70 mV, the firing threshold of a “rested” neuron is around -50 mV, and a postsynaptic potential (i.e., EPSP or IPSP) changes the membrane potential temporarily by at most a few mV.

If a neuron v has fired at time t' , it will not fire again for a few msec after t' , no matter how large its current potential $P_v(t)$ is (“absolute refractory period”). Then for a few further msec it is still “reluctant” to fire, i.e., a firing requires a larger value of $P_v(t)$ than usual (“relative refractory period”). Both of these refractory effects are modelled by a suitable “threshold function” $\Theta_v(t - t')$, where t' is the time of the most recent firing of v . In the deterministic (i.e., noise-free) version of the spiking neuron model one assumes that v fires whenever $P_v(t)$ crosses from below the function $\Theta_v(t - t')$. A typical shape of the function $\Theta_v(t - t')$ for a biological neuron is indicated in Figure 3. We assume that $\Theta_v(t - t') = \Theta_v(0)$ for large values of $t - t'$. We will consider in this article only computations in models for networks of spiking neurons where can assume that each neuron v did not fire for a while (i.e., $t - t'$ is large); hence, its threshold function has returned to its “resting value” $\Theta_v(0)$. Therefore, the shape of Θ_v is not relevant for these arguments, provided that $\Theta_v(x) = \Theta_v(0)$ for sufficiently large x .

A formal *Spiking Neuron Network* (SNN)—which was introduced in Maass (1995b, 1996a)—consists of a finite set V of *spiking neurons*, a set $E \subseteq V \times V$ of *synapses*, a weight $w_{u,v} \geq 0$ and a *response function* $\varepsilon_{u,v}: \mathbf{R}^+ \rightarrow \mathbf{R}$ for each synapse $\langle u, v \rangle \in E$ (where $\mathbf{R}^+ := \{x \in \mathbf{R}: x \geq 0\}$), and a *threshold function* $\Theta_v: \mathbf{R}^+ \rightarrow \mathbf{R}^+$ for each neuron $v \in V$.

If $F_u \subseteq \mathbf{R}^+$ is the set of *firing times* of a neuron u , then the *potential* at the trigger zone of neuron v at time t is given by

$$P_v(t) := \sum_{u: \langle u, v \rangle \in E} \sum_{s \in F_u: s < t} w_{u,v} \cdot \varepsilon_{u,v}(t - s).$$

In a noise-free model a neuron v fires at time t as soon as

$P_v(t)$ reaches $\Theta_v(t - t')$, where t' is the time of the most recent firing of v .

For some specified subset $V_{\text{input}} \subseteq V$ of *input neurons* one assumes that the firing times (“spike trains”) F_u for neurons $u \in V_{\text{input}}$ are not defined by the preceding convention, but are given from the outside. The firing times F_v for all other neurons $v \in V$ are determined by the previously described rules, and the output of the network is given in the form of the spike trains F_v for the neurons v in a specified set of *output neurons* $V_{\text{output}} \subseteq V$.

Experiments have shown that in vitro biological neurons fire with slightly varying delays in response to repetitions of the same current injection (Aertsen, 1993). Only under certain conditions neurons are known to fire in a more reliable manner (Mainen & Sejnowski, 1995). Therefore one also considers the *stochastic* or *noisy* version of the SNN model (Maass, 1996b), where the difference $P_v(t) - \Theta_v(t - t')$ just governs the *probability* that neuron v fires at time t . The choice of the exact firing times is left up to some unknown stochastic processes, and it may for example occur that v does *not* fire in a time interval I during which $P_v(t) - \Theta_v(t - t') > 0$, or that v fires spontaneously at a time t when $P_v(t) - \Theta_v(t - t') < 0$.

The previously described noisy version of the SNN model is basically identical with the *spike response model* in Gerstner (1995) (see also Gerstner & van Hemmen, 1994), and with the other common mathematical models for networks of spiking neurons (see, e.g., Abeles et al., 1993; Arbib, 1995; Tuckwell, 1988). Subtle differences exist between these models with regard to their treatment of the refractory effects and the “reset” of the membrane potential after a firing. But these differences will be irrelevant for the results that are considered in this article.

For theoretical results about stable states, synfire chains, associative memory, etc. in networks of spiking neurons we refer to Abeles (1991); Aityuan and Barrow (1993); Bienenstock (1995); Crair and Bialek (1990); Gerstner (1991); Gerstner and van Hemmen (1994); Gerstner, Ritz, and van Hemmen (1993); Herrmann, Hertz, and Prügel-Bennett (in press); Hopfield and Herz (1995); Ritz, Gerstner, Fuentes, and van Hemmen (1994). Results about computations with stochastic spiking neurons in firing rate coding can be found in Koch and Poggio (1992); Shawe-Taylor, Jeavons, and Van Daalen (1991), and results about the information transmitted by spiking neurons in Stevens and Zador (1996). Computations with a somewhat different model of a stochastic spiking neuron are studied in Judd and Aihara (1993) (see also the discussion in Maass, 1996a; Shawe-Taylor et al., 1991; Zhao, 1995). The possible use of phases of periodically firing neurons for the dynamic binding of variables is investigated in Shastri and Ajanagadde (1993).

We use in this article the terms *analog*, *numerical* and *real-valued* interchangeably to denote variables that range over \mathbf{R} or an interval of \mathbf{R} . For simplicity we

assume that all neural nets from the first two generations that are considered in the following have a feedforward architecture.

2. SIMULATION AND SEPARATION RESULTS

The mathematically simplest one within the range of SNN models is the one where the firing is deterministic, and both the response functions and the threshold functions are *piecewise constant* (i.e., “step functions”) as indicated in Figure 4. In the following we refer to this version as *type A*. This version of the SNN model actually captures quite well the intended capabilities of artificial spiking neurons in pulse stream VLSI.

We will later also discuss SNN models of *type B*, where we assume that response and threshold functions are *continuous and piecewise linear*. Examples for the simplest non-trivial response functions of type B are indicated in Figure 5. By using four or five linear segments one can approximate quite well the response and threshold functions of biological neurons with continuous piecewise linear functions (and hence with spiking neurons of type B).

2.1. Computation of Boolean Functions

We first observe that for the case of *boolean* input this model is computationally at least as powerful as neural nets from the first generation. We assume that n input bits x_1, \dots, x_n are given to the SNN via n input neurons a_1, \dots, a_n , where a_i fires at a specific time T_{input} if $x_i = 1$, and a_i does not fire at all if $x_i = 0$. We assume that the output bit of the SNN is given by the firing or non-firing of a specified *output neuron* during some specified time window. One can then simulate any layered feedforward neural net \mathcal{N} from the first generation by an SNN \mathcal{N}' of type A which has basically the same architecture as \mathcal{N} .

Only if one wants to respect in \mathcal{N}' the biologically motivated constraint that each neuron in \mathcal{N}' should only trigger EPSPs, or only IPSPs, then each gate of \mathcal{N} has to be simulated by a *pair* consisting of an excitatory and an inhibitory spiking neuron that both get the same input. In \mathcal{N}' one need not make use of the possibility to assign different values to the delays $\Delta_{u,v}$ of a neuron v (which model the time that passes until a firing of u has an effect on $P_v(t)$; see Figure 4) for different neurons u with $\langle u, v \rangle \in E$. For a biological neuron, these delays $\Delta_{u,v}$ may very well be different, depending on the length of the axon of u and the distance from the synapse to the trigger zone of v , but also on the distribution of ion channels in the dendritic tree of v . In fact, it is frequently assumed that the delays $\Delta_{u,v} = \Delta_{u,v}(t)$ are parameters that are tuned by some learning algorithm in biological neural systems (see, e.g., Kempter et al., 1996). Recent theoretical results (Maass & Schmitt, 1997) indicate that the expressive power of a neuron of type A with n variable delays is larger than that of a neuron of type A with n variable weights: its VC-dimension is $\Theta(n \log n)$ in the former case, but only $\Theta(n)$ in the latter case.

If one makes use of the possibility to employ for certain neurons v *different* delays $\Delta_{u,v}$ for different neurons u , then one can show that an SNN of type A is in fact computationally *more powerful* than neural nets of the same or similar size from the first or second generation. For that purpose we consider the concrete boolean function $CD_n: \{0, 1\}^{2n} \rightarrow \{0, 1\}$, which is defined by

$$CD_n(x_1, \dots, x_n, y_1, \dots, y_n) = \begin{cases} 1, & \text{if } x_i = y_i = 1 \\ & \text{for some } i \in \{1, \dots, n\} \\ 0, & \text{otherwise.} \end{cases}$$

This function appears to be relevant in a biological context, since it formalizes some form of pattern-matching, respectively, *coincidence-detection*.

A single spiking neuron v of type A (or of any other “reasonable” type) can compute CD_n . One just has to choose the delays to v from the input nodes a_1, \dots, a_n (for x_1, \dots, x_n) and the input nodes b_1, \dots, b_n (for y_1, \dots, y_n) in such a way that $\Delta_{a_i, v} = \Delta_{b_i, v}$ for $i = 1, \dots, n$, and $\Delta_{a_j, v}$ is so much larger than $\Delta_{a_i, v}$ for $j > i$ that the non-zero parts of the response functions $\epsilon_{a_j, v}$ and $\epsilon_{a_i, v}$ do not overlap if a_j and a_i fire simultaneously. All weights can be chosen equal to 1.

On the side, we would like to point out that a single spiking neuron of type A (or of type B) can compute this function CD_n in a *noise-robust* fashion, where small deviations in the firing times of the input neurons a_1, \dots, a_n , in the delays from these input neurons, in the weights or in the firing threshold do not affect the correctness of the output. To achieve this, it suffices to assign to the firing threshold $\Theta_v(0)$ of the spiking neuron a value such as $1.5 \cdot (\text{maximal value of an EPSP})$.

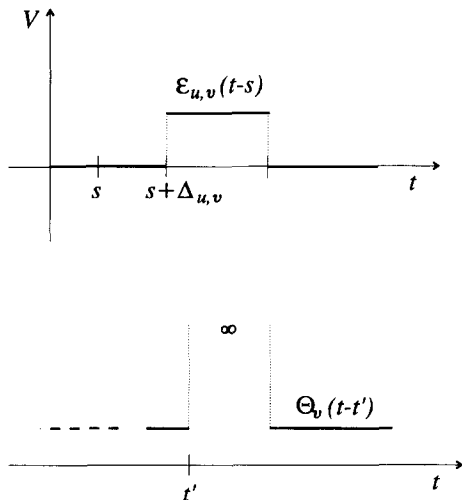


FIGURE 4. Response and threshold functions of a spiking neuron of type A.

THEOREM 1. 1. Any threshold circuit \mathcal{N} that computes CD_n has at least $n/\log(n + 1)$ gates.

2. Any sigmoidal neural net \mathcal{N} with piecewise polynomial activation functions that computes CD_n has $\Omega(n^{1/2})$ gates. For the case of piecewise exponential activation functions (such as σ) one gets a lower bound of $\Omega(n^{1/4})$.

Proof. Let $a_1, \dots, a_n, b_1, \dots, b_n$ be the input nodes of \mathcal{N} where it receives the values $x_1, \dots, x_n, y_1, \dots, y_n$ of its $2n$ input variables. We show in fact a slightly stronger result than claimed: The lower bounds hold already for the numbers of those gates in \mathcal{N} that have a direct edge from at least one of the input nodes b_1, \dots, b_n . Thus in the case of *layered neural nets* these are lower bounds for the number of gates on the first hidden layer.

We consider computations of \mathcal{N} where some “fixed” vector $\underline{q} \in \{0,1\}^n$ is assigned to the input nodes b_1, \dots, b_n , so that the output of \mathcal{N} may be viewed as a function of the assignments to the input nodes a_1, \dots, a_n . We only consider the set S of those n assignments $\underline{e}_1, \dots, \underline{e}_n \in \{0,1\}^n$ to a_1, \dots, a_n where exactly one of the n input variables has the value 1. Since \mathcal{N} computes CD_n , it is obvious that for the 2^n different choices of $\underline{q} \in \{0,1\}^n$ the network computes 2^n different functions from S into $\{0,1\}$.

For the proof of *Part 1* we fix a linear order $<$ on the computation nodes in \mathcal{N} so that each computation node g receives (apart from input nodes a_1, \dots, a_n and b_1, \dots, b_n) only edges from other computation nodes in \mathcal{N} that precede g in this linear order. Consider some arbitrary computation node g in \mathcal{N} , and a set Q of assignments $\underline{q} \in \{0,1\}^n$ to b_1, \dots, b_n so that every computation node before g computes a function from S into $\{0,1\}$ (with regard to assignments of inputs from S to the input nodes a_1, \dots, a_n), which is the same for each of the assignments $\underline{q} \in Q$ to b_1, \dots, b_n . Note that for the first computation node in \mathcal{N} we can set $Q := \{0,1\}^n$.

Then for assignments from S to a_1, \dots, a_n , the values received by gate g from other computation nodes do not depend on the chosen assignment $\underline{q} \in Q$ to b_1, \dots, b_n . Hence, the weighted sum of the values received by g via direct edges from the input nodes a_1, \dots, a_n , and from computation nodes that precede g in $<$, assumes at most n different values $r_1 \leq \dots \leq r_n$ for the n different assignments from S to a_1, \dots, a_n and arbitrary assignments from Q to b_1, \dots, b_n . Obviously the output of g depends only on the value of this weighted sum and on the weighted sum r of those values that g receives via direct edges from input nodes b_1, \dots, b_n . If Θ is the threshold of the threshold gate g , then the minimal i such that $r_i + r \geq \Theta$ can assume at most $n + 1$ different values (including the value $i = n + 1$ if $r_n + r < \Theta$). Consequently, with different fixed assignments of $\underline{q} \in Q$ to b_1, \dots, b_n the node g can compute at most $n + 1$ different functions from S into $\{0,1\}$. This yields a partition of Q into $n + 1$ equivalence classes, and one can apply the same argument—for each of these equivalence

classes—to the *next* node in \mathcal{N} (with regard to the linear order $<$).

If one starts this construction with $Q = \{0,1\}^n$ for the first computation node in \mathcal{N} , after the k^{th} node one gets a partition of Q into at most $(n + 1)^k$ equivalence classes. On the other hand the fact that \mathcal{N} computes CD_n implies that the output node of \mathcal{N} computes for each assignment to b_1, \dots, b_n a *different* function from S into $\{0,1\}$, i.e., it partitions $\{0,1\}^n$ into 2^n different equivalence classes Q . Hence, the number s of computation nodes in \mathcal{N} that have a direct edge from at least one of the input nodes b_1, \dots, b_n satisfies $(n + 1)^s \geq 2^n$, i.e., $s \geq n/\log(n + 1)$.

In the proof of *Part 2* we construct from \mathcal{N} a related sigmoidal neural net \mathcal{N}' for which we can show that it has “high” VC-dimension, and hence must contain a substantial number of sigmoidal gates. Such proof structure was first used by Koiran (1996), in a somewhat different context.

If one considers just a_1, \dots, a_n as input nodes of \mathcal{N} , then different fixed assignments to b_1, \dots, b_n can only shift the threshold of those s computation nodes in \mathcal{N} that have direct edges from b_1, \dots, b_n . We now consider a variation \mathcal{N}' of \mathcal{N} where the input nodes b_1, \dots, b_n are deleted, and the *thresholds* of the abovementioned s gates in \mathcal{N} are viewed as the only “programmable parameters” (or “weights”) in the usual sense of VC-dimension theory for neural networks (for a brief survey see Maass (1995a)). The fact that \mathcal{N} computes CD_n implies that \mathcal{N}' shatters S (with regard to different assignments to these s programmable parameters). Thus, \mathcal{N}' has a VC-dimension of at least n . On the other hand, the results of Goldberg & Jerrum (1995) and Karpinski & Macintyre (in press) imply that in this case the number s of programmable parameters in \mathcal{N} satisfies $n = O(s^2)$ in the case of piecewise polynomial activation functions, respectively $n = O(s^4)$ in the case of piecewise exponential activation functions.

2.2. Computation of Functions with Analog Input and Boolean Output

We have already shown that for *boolean* inputs a network of spiking neurons of type A has the full computational power of a neural net from the first generation of similar size, and is in fact more powerful. However, neural nets from all three generations are also able to process *numerical* inputs from \mathbf{R}^n or $[0,1]^n$, instead of just boolean inputs from $\{0,1\}^n$. For networks of spiking neurons it is natural to encode a numerical input variable $x_i \in \mathbf{R}$ by the firing time $T_{\text{input}} - x_i \cdot c$ of input neuron a_i (see also Hopfield, 1995), where $c > 0$ is some constant and T_{input} is a parameter that depends on the time when the input arrives, but not on the values of the input variables x_i . Similarly one expects that a numerical output $y \in \mathbf{R}$ is realized in an SNN by the firing of a certain “output neuron” at time $T_{\text{output}} - y \cdot c$ where $T_{\text{output}} > T_{\text{input}}$ is independent from the values x_1, \dots, x_n of the input

variables. We will refer to this method of encoding analog variables by the timing of single spikes as “linear temporal coding”. For the computation of functions with *boolean* output one can either employ the same output convention as before, or apply rounding (i.e., one considers a firing of the output neuron *before* a certain fixed time T as an output of “1”).

A concrete example for an interesting function with analog input and boolean output is the “*element distinctness function*” $ED_n : (\mathbf{R}^+)^n \rightarrow \{0,1\}$ defined by

$$ED_n(x_1, \dots, x_n) = \begin{cases} 1, & \text{if } x_i = x_j \text{ for some } i \neq j \\ 0, & \text{if } |x_i - x_j| \geq 1 \text{ for all } i, j \text{ with } i \neq j \\ \text{arbitrary,} & \text{otherwise.} \end{cases}$$

If one encodes the value of input variable x_i as the firing time $T_{\text{input}} - x_i \cdot c$ (of input neuron a_i), then for sufficiently large values of the constant $c > 0$ a single spiking neuron v can compute ED_n (even with $\Delta a_{i,v} = \Delta a_{j,v}$ for all $i, j \in \{1, \dots, n\}$). This holds for any reasonable type of response function, e.g., type A, or the type B considered below.

We also would like to point out that ED_n can be computed by a single spiking neuron in a very *noise-robust* fashion. Let ε_{\max} be the maximal value that is assumed by an EPSP, and let $\varepsilon(c)$ be the maximal value that can be achieved by the sum of two EPSPs that arrive with a temporal difference of at least c . By choosing the value $\Theta_v(0) = \frac{2 \cdot \varepsilon_{\max} + \varepsilon(c)}{2}$ for the firing threshold of a “rested” neuron v one achieves that v definitely fires if $x_i = x_j$ for some $i \neq j$, and that it does definitely not fire if $|x_i - x_j| \geq 1$ for any two different inputs x_i, x_j given in temporal coding. In addition with this choice of $\Theta_v(0)$ the neuron v gives the correct output even if its membrane potential, its firing threshold, and the arrival times of its input-EPSPs are subject to noise. Furthermore, its “safety margin” of $\frac{2 \cdot \varepsilon_{\max} - \varepsilon(c)}{2}$ can be increased up to the value $\frac{\varepsilon_{\max}}{2}$ if c is chosen so large that $\varepsilon(c) = \varepsilon_{\max}$.

This noise-robust computation of ED_n by a spiking neuron is made possible through the way in which this function ED_n is defined: if $\min\{|x_i - x_j| : i \neq j\}$ has a value *between* 0 and 1 for some input $\langle x_1, \dots, x_n \rangle \in \mathbf{R}^n$, then it does not matter whether the neuron fires or not. Thus, the clause “arbitrary” in the definition of ED_n makes sure that “hair-trigger situations” can be avoided by a spiking neuron that computes ED_n .

THEOREM 2. Any layered threshold circuit \mathcal{N} that computes ED_n has $\Omega(n \cdot \log n)$ gates on its first hidden layer.

Proof. Let k be the number of gates in \mathcal{N} on the first hidden layer. The corresponding k halfspaces partition the input space \mathbf{R}^n into at most 2^k different polytopes (i.e., intersections of halfspaces) so that \mathcal{N} gives the same output for all inputs from the same polytope. For

this consideration one has to allow polytopes that are intersections of closed and open halfspaces.

We now consider those $n!$ inputs $\underline{x}_\pi = \langle \pi(1), \dots, \pi(n) \rangle \in \{1, \dots, n\}^n$ that represent all $n!$ permutations π of $\{1, \dots, n\}$. It suffices to show that each \underline{x}_π lies in a different polytope, since this implies that $2^k \geq n!$. Thus assume for a contradiction that two permutations \underline{x}_π and $\underline{x}_{\hat{\pi}}$ lie in the same polytope P . By construction the threshold circuit \mathcal{N} gives the same output for all $\underline{x} \in P$. Since P is convex, \mathcal{N} gives not only the same output for \underline{x}_π and $\underline{x}_{\hat{\pi}}$, but also for all points on the line L that connects these two points. This yields a contradiction, since $ED_n(\underline{x}_\pi) = ED_n(\underline{x}_{\hat{\pi}}) = 0$, but $ED_n(\underline{x}) = 1$ for some point \underline{x} on this line L .

In order to analyse the complexity of functions with *boolean* output on sigmoidal neural nets, one needs to fix a suitable convention for *rounding* the real-valued output of such nets. In order to make our subsequent lower bound result as strong as possible, one may assume here the *weakest* possible rounding convention, where for some arbitrary parameter Θ the real-valued output r of the output node of the net is rounded to 1 if $r \geq \Theta$. No separating interval is required between outputs that are rounded to 0 respectively 1.

In the same way as for CD_n one can show that any neural net from the second generation that computes ED_n needs to have $\Omega(n^{1/4})$ gates. This lower bound will be improved to $(n - 1)/4$ in the following theorem. The proof of this stronger separation result exploits, instead of a bound for the VC-dimension, Sontag’s better upper bound of $2w + 1$ (Sontag, 1997) for the maximal number d such that *every* set of d different inputs in general position can be shattered by a sigmoidal neural net with w programmable parameters. In order to apply his result in our lower bound argument one has to construct from an arbitrary sigmoidal neural net which computes ED_n a related net that shatters *every* set of $n - 1$ inputs.

THEOREM 3. Any sigmoidal neural net \mathcal{N} that computes ED_n has at least $\frac{n-4}{2} - 1$ hidden units.

Proof. Let \mathcal{N} be an arbitrary sigmoidal neural net with k gates that computes ED_n .

Consider any set $S \subseteq \mathbf{R}^+$ of size $n - 1$. Let $\lambda > 0$ be sufficiently large so that the numbers in $\lambda \cdot S$ have pairwise distance ≥ 2 . Let A be a set of $n - 1$ numbers $> \max(\lambda \cdot S) + 2$ with pairwise distance ≥ 2 .

By assumption \mathcal{N} can decide for n arbitrary inputs from $\lambda \cdot S \cup A$ whether they are all different. Let \mathcal{N}_λ be a variation of \mathcal{N} where all weights on edges from the first input variable are multiplied by λ . Then by assigning suitable fixed sets of $n - 1$ pairwise different numbers from $\lambda \cdot S \cup A$ to the other $n - 1$ input variables, \mathcal{N}_λ computes any characteristic function over S .

Thus, if one considers as *programmable* parameters of \mathcal{N} the $\leq k$ weights on edges from the first input variable of \mathcal{N} and the $\leq k$ thresholds of gates that are connected to

some of the other $n - 1$ input variable, then \mathcal{N} shatters S with $2k$ programmable parameters. Actually in the more general setting of the subsequent argument we have only $k + 1$ programmable parameters, since the k occurrences of the factor λ in the weights may be counted as a *single* programmable parameter.

Since the set $S \subseteq \mathbf{R}^+$ of size $n - 1$ was chosen arbitrarily, we can now apply the result from Sontag (1997), which implies that $n - 1 \leq 2(k + 1) + 1$, hence $k \geq (n - 4)/2$. Thus, \mathcal{N} has at least $(n - 4)/2$ computation nodes, and therefore at least $(n - 4)/2 - 1$ hidden units.

REMARK 4.

1. The lower bound of $\Omega(n)$ in Theorem 3 is the largest lower bound for the size of sigmoidal neural nets that has so far been achieved (not just for ED_n but for any concrete function). The best previously known lower bound was $\Omega(n^{1/4})$ for some other function, due to Koiran (1996).
2. The result of Section 4 in Sontag (1997) implies that his upper bound, and hence the lower bound of the preceding Theorem 3, remain valid if the neural net \mathcal{N} computing ED_n employs both sigmoidal gates and threshold gates.

Apparently for most neurons v in the cortex it is not likely that the “weights” $w_{u,v}$ of its synapses are large enough such that just two synchronous EPSPs suffice to increase the potential P_v over the firing threshold $\Theta_v(0)$ of a “rested” neuron v . In that regard the common mathematical model for a spiking neuron “overestimates” the computational capabilities of a biological neuron. It is more realistic to assume that six simultaneously arriving EPSPs can cause a neuron to fire (see the discussion in Valiant, 1994). Therefore, we consider the following variation $\widetilde{ED}_n : (\mathbf{R}^+)^n \rightarrow \{0, 1\}$ of the function ED_n :

$$\widetilde{ED}_n(x_1, \dots, x_n) = \begin{cases} 1, & \text{if there exists some } k \geq 1 \text{ such that} \\ & x_1, x_2, x_3, x_{3k+1}, x_{3k+2}, x_{3k+3} \text{ all have} \\ & \text{the same value} \\ 0, & \text{if every interval } I \subseteq \mathbf{R}^+ \text{ of length 1} \\ & \text{contains the values of at most} \\ & 3 \text{ input variables } x_i \\ \text{arbitrary,} & \text{otherwise.} \end{cases}$$

In the common model of a spiking neuron the membrane potential $P_v(t)$ is assumed to be a linear sum of the post-synaptic potentials. This is certainly an idealization, since isolated EPSPs that arrive at synapses far away from the trigger zone (which is located at the beginning of the axon) are subject to an exponential decay on their

way to the trigger zone. Hence, such isolated EPSPs have hardly any impact on the membrane potential $P_v(t)$ at the trigger zone. On the other hand, EPSPs that arrive *synchronously* at *adjacent* synapses are “boosted” at “hot spots” of the dendritic tree, and hence may have a significant impact on the membrane potential $P_v(t)$ at the trigger zone (Shepherd, 1994). We have defined \widetilde{ED}_n in such a way that, in spite of these nonlinear effects in the integration of EPSPs, it is quite plausible that a biological neuron can compute \widetilde{ED}_n in temporal coding for a fairly large value of n . A neuron computing \widetilde{ED}_n needs to fire only when two “blocks” consisting of three *adjacent* synapses all receive *synchronous* EPSPs. Furthermore, a “hair-trigger” situation is avoided, since *no* requirements are made for the case when the neuron receives just four or five synchronous (or almost synchronous) EPSPs. Non-firing is required only in the case when the neuron receives at most three EPSPs during any time interval of length c .

In order to prove a lower bound for the number of hidden units in arbitrary neural nets \mathcal{N} that compute \widetilde{ED}_n with sigmoidal and threshold gates, one proceeds as in the proof of Theorem 3. One now considers arbitrary sets $S \subseteq \mathbf{R}^+$ of size $\lfloor (n - 3)/3 \rfloor$ and divides the remaining $n - 3$ input variables into $\lfloor (n - 3)/3 \rfloor$ blocks of three variables that always receive the same input value. Let \mathcal{N}_λ be a variation of \mathcal{N} which identifies the first three input variables, and multiplies all their weights by a common factor λ . Since \mathcal{N} computes \widetilde{ED}_n , the network \mathcal{N}_λ with k computation nodes shatters S with the help of $k + 1$ programmable parameters. Hence, Sontag’s result (Sontag, 1997) yields $\lfloor (n - 3)/3 \rfloor \leq 2(k + 1) + 1$, i.e., $k \geq (n - 15)/6$.

If one plugs in a common estimate for the number n of synapses at a biological neuron, such as $n = 10000$, the preceding inequality yields a lower bound of 1663 for the number $k - 1$ of hidden units in \mathcal{N} . Hence, even if one prefers to plug in somewhat different values for some of the abovementioned constants, the preceding proof for \widetilde{ED}_n (respectively, for a variation of \widetilde{ED}_n that reflects different choices of the parameters involved) still yields a lower bound of several hundreds for the minimal size of a sigmoidal neural net which computes the same function. Thus we have demonstrated a *substantial* difference between the computational power of *biological neurons* and *sigmoidal “neurons”* (i.e., computational units from the second generation).

For *numerical* inputs our previously sketched simulation of threshold circuits (i.e., neural nets from the first generation) by a network of spiking neurons of type A fails. More surprisingly, one can *prove* that no such simulation is possible. Let $f: \mathbf{N} \rightarrow \mathbf{N}$ be any function. Then for numerical inputs there *exists no way* of simulating an arbitrary threshold circuit with s gates by a network of $f(s)$ spiking neurons of type A. Consider a threshold circuit that outputs 1 for inputs $x_1, x_2, x_3 \in [0, 1]$ if $x_1 + x_2 = x_3$, and 0 else. Obviously this can be

achieved by a circuit with just three threshold gates: the circuit outputs 1 if $(x_1 + x_2 \leq x_3 \text{ AND } x_1 + x_2 \geq x_3)$. However, it has been shown that this function from $[0, 1]^3$ into $\{0, 1\}$ (as well as any restriction to $[0, \gamma]^3$ for some $\gamma > 0$) cannot be computed by any network of spiking neurons of type A, no matter how many neurons and how much computation time it employs. This follows from a general characterization of the computational power of networks of spiking neurons of type A for numerical inputs in terms of the computational power of a restriction called $N = \text{RAM}$ of the common model of a random access machine (RAM) that is given in Maass & Ruf (1995).

Thus, we have arrived here at a limit of the computational power of spiking neurons of type A for numerical inputs. The question arises whether this limitation indicates a weakness of spiking neurons in general, or just a weakness of the extremely simple response and threshold functions of type A. For answering this question let us consider spiking neurons with *continuous piecewise linear* (instead of piecewise constant) response and threshold functions, to which we refer as *spiking neurons of type B*. Examples for the simplest non-trivial response functions for the type A spiking neuron are indicated in Figure 5.

With regard to the computational power of spiking neurons of type B it does not make much difference whether one allows here piecewise constant, piecewise linear, or more general types of threshold functions Θ_u , as long as we consider only feedforward computations and the threshold functions Θ_v have the value “ ∞ ” for small arguments. In addition the, concrete shape of the response functions of type B will be irrelevant in the following.

One can show that in contrast to the abovementioned negative result about neural nets of type A, a network of $O(1)$ spiking neurons with response functions of type B (e.g., as indicated in Figure 5) can simulate any threshold gate even for n real-valued input variables. This simulation exploits an important effect of spiking neurons of type B that cannot be realized with spiking neurons of type A: incoming EPSPs and IPSPs can *shift* the firing time of a neuron in a *continuous* manner (Maass, 1997). More precisely, for a certain range of the parameters involved, the firing time t_v of a neuron v in response to the firings of presynaptic neurons u at times $T_{\text{input}} - x_u \cdot c$ can be written in the form

$$t_v = T_{\text{output}} - \sum_{u:(u,v) \in E} \text{sign}(\epsilon_{u,v}) \cdot w_{u,v} \cdot x_u \quad (1)$$

where T_{output} does not depend on the values of the x_u , and where $\text{sign}(\epsilon_{u,v}) = -1$ in the case of an EPSP and $\text{sign}(\epsilon_{u,v}) = 1$ in the case of an IPSP. Thus, neuron v outputs the *weighted sum*

$$\sum_{u:(u,v) \in E} \text{sign}(\epsilon_{u,v}) \cdot w_{u,v} \cdot x_u$$

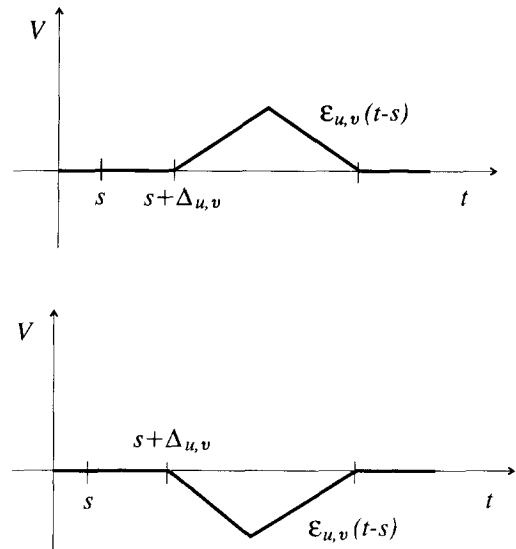


FIGURE 5. Response functions (EPSP and IPSP) of a spiking neuron of type B. The particular shape of the ‘triangle’ is not important for results in this article.

in *temporal coding* (in response to analog inputs x_u given in temporal coding).

Equation (1) reveals the somewhat surprising fact that, in the context of *temporal coding*, the “weights” $w_{u,v}$ of synapses of spiking neurons are able to play the same role as those of computational units of the first two generations of neural network.

All subsequent layers after the first hidden layer in a layered neural net from the first generation receive just *boolean* inputs, even if the network inputs are real-valued. Hence, these subsequent layers can easily be simulated by spiking neurons of type A (as indicated before). However, a subtle but serious problem arises if one wants to simulate threshold circuits with *boolean* inputs and outputs (or any other type of boolean circuit) with spiking neurons of type B, e.g., with response functions as in Figure 5, which are substantially closer to the biological prototypes in Figure 2 than response functions of type A. It is obvious that a spiking neuron of type B can simulate a *boolean* gate only if it receives *synchronized* input spikes. The problem is that even if a layer of spiking neurons of type B receives boolean input via *synchronized* input spikes (e.g., in a coding where a spike corresponds to “1” and no spike corresponds to “0”), the neurons on this layer will not fire in a *synchronized* manner, but at slightly different times that depend on their concrete input “bits”. The root of this problem (which does not arise for spiking neurons of type A) is the fact that a potential $P_v(t)$ that is the sum of several EPSPs and IPSPs of type B will itself be continuous and piecewise linear, and that the slopes of its linear pieces will depend in particular on the *number* of EPSPs that it receives simultaneously (hence on the concrete “boolean” input in our interpretation). Thus, the precise time when $P_v(t)$ crosses the threshold $\Theta_v(0)$ will in

general depend on the “boolean” input of the spiking neuron. This causes a serious problem for the simulation of multilayer threshold circuits (or other multilayer boolean circuits) by SNNs of type B, because if those neurons v on the considered layer that are firing (and hence represent a “1” in the simulation of a boolean circuit) do not fire in a *synchronized* manner, the simulation of threshold gates, or even of simpler boolean gates (such as AND), by the *next* layer of spiking neurons of type B becomes impossible.

THEOREM 5. *Any threshold circuit of s gates having real-valued inputs from $[0, 1]^n$ can be simulated by a network of $O(s)$ spiking neurons of type B.*

Proof. Consider first an arbitrary threshold gate G with inputs $\langle x_1, \dots, x_n \rangle$ from $[0, 1]^n$ that outputs 1 if $\sum_{i=1}^n \alpha_i x_i \geq \alpha_0$, and 0 otherwise. We show that G can be simulated by a network having a constant number (i.e., $O(1)$) of spiking neurons of type B with regard to temporal coding of network inputs x_1, \dots, x_n (for a sufficiently small value of the constant c). One employs here the same construction as for the simulation of a linear (respectively sigmoidal) gate given in Maass (1997), which yields a spiking neuron v whose firing time represents the weighted sum $\sum_{i=1}^n \alpha_i x_i$ in temporal coding. In particular v fires at or before a fixed time T (which does not depend on x_1, \dots, x_n) if $\sum_{i=1}^n \alpha_i x_i \geq \alpha_0$, and after time T otherwise. We arrange that the resulting EPSP from v arrives at a subsequent spiking neuron v' , which receives in addition an EPSP from an auxiliary spiking neuron whose firing time depends on T_{input} , but not on x_1, \dots, x_n . With a suitable choice of weights and delays for v' , the neuron will fire if and only if v fires at or before time T .

Obviously one can simulate in the same way the whole first layer of any given threshold circuit C . In order to simulate the subsequent layers of C with spiking neurons of type B, one can employ the construction from Maass (1996a). The previously described spiking neurons v' represent the outputs of gates from the first layer of C by firing if and only if the corresponding gate in C outputs 1. However, the precise time at which v' fires in this case depends on x_1, \dots, x_n . Hence, before one can use the, “boolean” outputs of these gates v' as inputs for other spiking neurons of type B to simulate the subsequent layers of C according to the construction in Maass (1996a), one has to employ a synchronization module as constructed in the proof of Theorem 2.1 in Maass (1996a).

2.3. Further Results for Networks of Spiking Neurons of Type B

We have shown in the preceding section that in contrast to SNNs of type A, networks of spiking neurons of type B can simulate neural nets from the *first* generation even for the case of *real-valued* network input. Hence, the

question arises whether networks of spiking neurons of type B can also simulate (respectively approximate) neural nets from the *second* generation which have real-valued input *and* output. This question is answered affirmatively in Maass (1997), by showing that, with regard to temporal coding of real-valued variables x , any continuous function $F: [0, 1]^n \rightarrow [0, 1]^k$ can be approximated arbitrarily closely (with regard to uniform convergence, i.e., L_∞) by a one hidden layer network of spiking neurons of type B.

In fact, this result holds not just for the simple scheme of linear temporal coding described at the beginning of Section 2.2, but also for any other scheme of coding analog variables by the timing of single spikes that is “continuously related” to this scheme. Thus for example, it also holds if a neuron that fires at time $T - x \cdot c$ does not encode the analog number x , but instead e^{-x} or x^3 .

In addition there exists evidence that many practically relevant analog function F can be approximated by *small* networks of spiking neurons of type B. A large number of results regarding practical applications of learning with backprop on sigmoidal neural nets suggest that the relevant target functions F for these applications can be learned (and hence approximated) by sigmoidal neural nets with a rather small number of sigmoidal gates. Additional empirical evidence suggests that the precise form of the sigmoidal activation function is not important for the number of sigmoidal gates that are needed. Thus one can argue that the target functions $F: [0, 1]^n \rightarrow [0, 1]^k$ that arise in application problems can in general be approximated quite well by sigmoidal neural nets with a small number s of sigmoidal units that employ the following *linear saturated activation function* π :

$$\pi(y) = \begin{cases} 0, & \text{if } y < 0 \\ y, & \text{if } 0 \leq y \leq 1 \\ 1, & \text{if } y > 1. \end{cases}$$

The approximation result of Leshno, Lin, Pinkus, and Schocken (1993) implies that in this case F can also be approximated quite well by a network of $O(s)$ *spiking neurons* of type B.

Thus, one may say that with regard to circuit complexity for computing analog functions, networks of spiking neurons of type B are at least as powerful as neural nets from the second generation. Furthermore, our previously described *lower* bounds for the size of neural nets from the first two generations (for nets that compute the functions CD_n , ED_n or \widetilde{ED}_n) imply that networks of spiking neurons of type B are in fact *strictly more powerful* than neural nets from the first two generations: in order to achieve separation results between SNNs of *type B* and neural nets from the first two generations it just remains to verify that instead of a single spiking neuron of type A also a single spiking neuron of type B can compute CD_n , ED_n and \widetilde{ED}_n .

We refer to Maass (1996a, 1997) for details of the proofs of the abovementioned simulation results. It can be seen from these proofs that—for positive results about the computational power of SNNs of type B—they do not actually require that the response or threshold functions are piecewise linear (i.e., of type B). Rather it suffices to assume that EPSPs have some small linearly increasing segment and IPSPs have some small linearly decreasing segment. These properties are approximately satisfied by EPSPs and IPSPs of biological neurons (see Figure 2). In Maass (1995a, c) a complete characterization of the computational power of SNNs of type B is given in terms of a restriction (called N-RAM) of the familiar model of a random access machine.

In addition it is shown in Maass (1997) that the simulation of sigmoidal neural nets by SNNs can also be carried out with the biologically more realistic model of a *stochastic* or *noisy* spiking neuron. It is easy to see that the functions CD_n , ED_n and \widetilde{ED}_n considered here, can be computed by a single *noisy* spiking neuron of type A or B. Furthermore, it is shown in Maass (1996b) that even with *very noisy* spiking neurons of type A or B one can in principle carry out arbitrary *digital* computations with any desired degree of reliability. However, noise certainly affects the computational power of networks of spiking neurons for analog input, and we refer to Maass and Orponen (1997) with regard to limits of the computational power of networks of *noisy* spiking neurons with analog input.

3. CONCLUSIONS

We have analysed in this article the computational power of networks of spiking neurons with regard, to temporal coding with single spikes. It turns out that this computational model has at least the same computational power as neural nets from the first two generations (i.e., multi-layer perceptions and sigmoidal neural nets) of a similar size. Furthermore we have exhibited concrete functions which require for their computation significantly *fewer* neurons in a network of *spiking* neurons.

The proof of Theorem 3 appears to be of independent interest in the theory of sigmoidal neural nets, since it provides the strongest lower bound result for sigmoidal neural nets that is currently known. It improves the largest previously known lower bound $\Omega(n^{1/4})$ (Koiran, 1996) to $\Omega(n)$. This new lower bound result is also of interest from the technical point of view, since it provides the first known application of recent results Sontag (1997) about the “Sontag dimension” of neural nets. This is a new notion of a “dimension” for a neural net that is in a certain sense dual to the familiar concept of the Vapnik–Chervonenkis dimension of a neural net (one replaces “there exists a set S of d inputs...” by “for all sets S of d inputs...” in the definition of the dimension).

As the references in this article indicate, the theoretical investigation of networks of spiking neurons is not a new

research topic. In fact it has a long tradition in theoretical neurobiology, biophysics, and theoretical physics. However, a mathematically rigorous analysis of the *computational power* of networks of spiking neurons has so far been missing. We believe that such analysis will be helpful in understanding the organization of computations in complex biological neural systems.

In addition such analysis appears to be helpful for evaluating the potential capabilities of various designs of “artificial networks of spiking neurons”, in particular of silicon implementations of integrated circuits that compute with pulses (DeYong et al., 1992; Douglas et al., 1995; Horinchi et al., 1991; Jahnke et al., 1996; Jiu & Leong, 1996; Mahowald, 1994; Mead, 1989; Meador et al., 1991; Murray & Tarassenko, 1994; Northmore & Elias, 1996; Pratt, 1989; Zaghoul et al., 1994; Zhao, 1995). For example, the results of this article and those in Maass and Ruf (1995) show that there exist drastic differences between the computational capabilities of networks of spiking neurons that operate with rectangular pulses (i.e., type A) and those that operate with triangular pulses (i.e., type B).

REFERENCES

- Abeles, M. (1991). *Corticonics: Neural circuits of the cerebral cortex*. Cambridge: Cambridge University Press.
- Abeles, M., Bergman, H., Margalit, E., & Vaadia, E. (1993). Spatio-temporal firing patterns in the frontal cortex of behaving monkeys. *Journal of Neurophysiology*, 70, 1629–1638.
- Aertsen, A. (Ed.) (1993). *Brain theory: spatio-temporal aspects of brain function*. Elsevier.
- Aityan, S. K., & Barrow, D. L. (1993). Paradigm, logical performance, and training of recurrent refractory neural networks. *Neural, Parallel & Scientific Computations*, 1, 3–28.
- Arbib, M. A. (1995). *The handbook of brain theory and neural networks*. Cambridge: MIT Press.
- Bair, W., Koch, C., Newsome, W., & Britten, K. (1994). Reliable temporal modulation in cortical spike trains in the awake monkey. In *Proceedings of the Symposium on Dynamics of Neural Processing*. Washington, DC.
- Bialek, W., & Rieke, F. (1992). Reliability and information transmission in spiking neurons. *Trends in Neuroscience*, 15, 428–434.
- Bienenstock, E. (1995). A model of neocortex. *Network: Computation in Neural Systems*, 6, 179–224.
- Bower, J. M., & Beeman, D. (1995). *The book of GENESIS: exploring realistic neural models with the General Neural Simulation System*. New York: Springer.
- Churchland, P. S., & Sejnowski, T. J. (1993). *The computational brain*. Cambridge: MIT Press.
- Crair, M. C., & Bialek, W. (1990). Non-Boltzmann dynamics in networks of spiking neurons. In *Advances in neural information processing systems*, Vol. 2 (pp. 109–116). San Mateo: Morgan Kaufmann.
- DasGupta, B., & Schnitger, G. (1993). The power of approximating: a comparison of activation functions. In *Advances in neural information processing systems*, Vol. 5 (pp. 615–622). San Mateo: Morgan Kaufmann.
- DeYong, M. R., Findley, R. L., & Fields, C. (1992). The design, fabrication, and test of a new VLSI hybrid analog-digital neural processing element. *IEEE Transcripts on Neural Networks*, 3, 363–374.

- Douglas, R. J., Koch, C., Mahowald, M., Martin, K. A. C., & Suarez, H. H. (1995). Recurrent excitation in neocortical circuits. *Science*, 269, 981–985.
- Ferster, D., & Spruston, N. (1995). Cracking the neuronal code. *Science*, 270, 756–757.
- Gerstner, W. (1991). Associative memory in a network of ‘biological neurons’. In *Advances in neural information processing systems*, Vol. 3 (pp. 84–90). San Mateo: Morgan Kaufmann.
- Gerstner, W. (1995). Time structure of the activity in neural network models. *Physics Review E*, 51, 738–758.
- Gerstner, W., & van Hemmen, J. L. (1994). How to describe neuronal activity: spikes, rates, or assemblies. In *Advances in neural information processing systems*, Vol. 6 (pp. 463–470). San Mateo: Morgan Kaufmann.
- Gerstner, W., Ritz, R., & van Hemmen, J. L. (1993). A biologically motivated and analytically soluble model of collective oscillations in the cortex: I. Theory of weak locking. *Biological Cybernetics*, 68, 363–374.
- Goldberg, P. W., & Jerrum, M. R. (1995). Bounding the Vapnik–Chervonenkis dimension of concept classes parameterized by real numbers. *Machine Learning*, 18, 131–148.
- Herrmann, M., Hertz, J. A., & Prügel-Bennett, A. (in press). Analysis of synfire chains. *Nordita Preprint*.
- Hopfield, J. J. (1995). Pattern recognition computation using action potential timing for stimulus representations. *Nature*, 376, 33–36.
- Hopfield, J. J., & Herz, A. V.M. (1995). Rapid local synchronization of action potentials: towards computation with coupled integrate-and-fire neurons. *Proceedings of the National Academy of Science*, 92, 6655–6662.
- Horinchi, T., Lazzaro, J., Moore, A., & Koch, C. (1991). A delay-line based motion detection chip. In *Advances in neural information processing systems*, Vol. 3 (pp. 406–412). San Mateo: Morgan Kaufmann.
- Jahnke, A., Roth, U., & Klar, H. (1996). A SIMD/dataflow architecture for a neurocomputer for spike-processing neural networks (NESPINN). *MicroNeuro*, 232–237.
- Jiu, C. T., & Leong, P. H. W. (1996). An analog VLSI time-encoded pattern classifier. In *Proceedings of the 7th Australian Conference on Neural Networks* (pp. 212–215). Canberra.
- Johnston, D., & Wu, S. M. (1995). *Foundations of cellular neurophysiology*. Cambridge: MIT Press.
- Judd, K. T., & Aihara, K. (1993). Pulse propagation networks: a neural network model that uses temporal coding by action potentials. *Neural Networks*, 6, 203–215.
- Karpinski, M., & Macintyre, A. (in press). Polynomial bounds for VC-dimension of sigmoidal and general Pfaffian neural networks. *Journal of Computer and System Sciences*.
- Kempler, R., Gerstner, W., van Hemmen, J. L., & Wagner, H. (1996). Temporal coding in the sub-millisecond range: model of barn owl auditory pathway. In *Advances in neural information processing systems*, Vol. 8 (pp. 124–130). Cambridge: MIT Press.
- Koch, C., & Poggio, T. (1992). Multiplying with synapses and neurons. In T. McKenna, J. Davis, & S. F. Zornetser (Eds.), *Single neuron computation* (pp. 315–346). Boston: Academic Press.
- Koiran, P. (1996). VC-dimension in circuit complexity. In *Proceedings of the Conference on Computational Complexity* (pp. 81–85).
- Krüger, J., & Aiple, F. (1988). Multielectrode investigation of monkey striate cortex: spike train correlations in the infragranular layers. *Journal of Neurophysiology*, 60, 798–828.
- Lapique, L. (1907). Recherches quantitatives sur l’excitation électrique des nerfs traitée comme une polarisation. *Journal of Physiology and Pathology*, 9, 620–635.
- Leshno, M., Lin, V. Y., Pinkus, A., & Schocken, S. (1993). Multilayer feedforward networks with a nonpolynomial activation function can approximate any function. *Neural Networks*, 6, 861–867.
- Lestienne, R. (1996). Determination of the precision of spike timing in the visual cortex of anaesthetised cats. *Biological Cybernetics*, 74, 55–61.
- Maass, W. (1995a). Vapnik–Chervonenkis dimension of neural nets. In M. A. Arbib (Ed.), *The handbook of brain theory and neural networks* (pp. 1000–1003). Cambridge: MIT Press.
- Maass, W. (1995b). On the computational complexity of networks of spiking neurons. In *Advances in neural information processing systems*, Vol. 7 (183–190). Cambridge: MIT Press.
- Maass, W. (1995c). Analog computations on networks of spiking neurons. In *Proceedings of the 7th Italian Workshop on Neural Nets*. World Scientific Press, 99–104.
- Maass, W. (1996a). Lower bounds for the computational power of networks of spiking neurons. *Neural Computation*, 8(1), 1–40.
- Maass, W. (1996b). On the computational power of noisy spiking neurons. In *Advances in neural information processing systems*, Vol. 8 (pp. 211–217). Cambridge: MIT Press.
- Maass, W. (1997). Fast sigmoidal networks via spiking neurons. *Neural Computation*, 9, 279–304.
- Maass, W., & Orponen, P. (1997). On the effect of analog noise in discrete-time analog computations. *Advances in neural information processing systems*, Vol. 9. Cambridge: MIT Press.
- Maass, W., & Ruf, B. (1995). On the relevance of the shape of post-synaptic potentials for the computational power of spiking Neurons. *Proceedings of the International Conference on Artificial Neural Networks, ICANN’95* (pp. 515–520). Paris: EC2&Cic.
- Maass, W., & Schmitt, M. (1997). On the complexity of learning for a spiking neuron. Proc. of the 10th Conference on Computational Learning Theory 1997, ACM-Press, New York. Forthcoming.
- Maass, W., Schnitger, G., & Sontag, E. (1991). On the computational power of sigmoid versus boolean threshold circuits. In *Proceedings of the 32nd Annual IEEE Symposium on Foundations of Computer Science* (pp. 767–776).
- Mahowald, M. (1992). *VLSI. Analogs of neuronal visual processing: a synthesis of form and function*. Ph.D. dissertation, California Institute of Technology.
- Mahowald, M. (1994). *An analog VLSI system for stereoscopic vision*. Boston: Kluwer.
- Mainen, Z. F., & Sejnowski, T. J. (1995). Reliability of spike timing in neocortical neurons. *Science*, 268, 1503–1506.
- Mead, C. (1989). *Analog VLSI and neural systems*. Reading: Addison-Wesley.
- Meador, J. L., Wu, A., Cole, C., Nintunze, N., & Chintrakulchai, P. (1991). Programmable impulse neural circuits. *IEEE Transcripts on Neural Networks*, 2, 101–109.
- Murray, A. and Tarassenko, L. (1994). *Analog neural VLSI: a pulse stream approach*. Chapman and Hall.
- Northmore, D. P., & Elias, J. G. (1996). Discrimination of spike patterns by dendritic processing in a network of silicon neuromorphs. In *Proceedings of the 5th Annual Conference on Computational Neuroscience*. San Diego: Academic Press.
- Perrett, D. I., Rolls, E. T., & Caan, W. C. (1982). Visual neurons responsive to faces in the monkey temporal cortex. *Experimental Brain Research*, 47, 329–342.
- Pratt, G. A. (1989). *Pulse computation*. Ph.D. thesis, MIT, Cambridge.
- Rieke, F., Warland, D., van Stevenick, R., & Bialek, W. (1996). *SPIKES: exploring the neural code*. Cambridge: MIT Press.
- Ritz, R., Gerstner, W., Fuentes, U., & van Hemmen, L. (1994). A biologically motivated and analytically soluble model of collective oscillations in the cortex: II. Applications to binding and pattern segmentation. *Biological Cybernetics*, 71, 49–358.
- Rolls, E. T. (1994). Brain mechanism for invariant visual recognition and learning. *Behavioural Processes*, 33, 113–138.
- Rolls, E. T., & Tovee, M. J. (1994). Processing speed in the cerebral cortex, and the neurophysiology of visual backward masking. *Proceedings of the Royal Society of Britain*, 257, 9–15.
- Sejnowski, T. J. (1995). Time for a new neural code?. *Nature*, 376, 21–22.
- Shastri, L., & Ajanagadde, V. (1993). From simple associations to systematic reasoning: a connectionist representation of rules, variables and dynamic bindings using temporal synchrony. *Behavioural and Brain Sciences*, 16, 417–494.

- Shawe-Taylor, J., Jeavons, P., & Van Daalen, M. (1991). Probabilistic bit stream neural chip: theory. *Connection Science*, 3, 317–328.
- Shepherd, G. M. (Ed.) (1990). *The synaptic organization of the brain* (3rd ed.). New York: Oxford University Press.
- Shepherd, G. M. (1994). *Neurobiology* (3rd ed.). New York: Oxford University Press.
- Singer, W. (1995). Synchronization of neuronal responses as a putative binding mechanism. In M. A. Arbib (Ed.), *The handbook of brain theory and neural networks* (pp. 960–964). Cambridge: MIT Press.
- Softky, W. (1994). Sub-millisecond coincidence detection in active dendritic tree. *Neuroscience*, 58, 13–41.
- Sontag, E. D. (1997). Shattering all sets of k points in "general position" requires $(k - 1)/2$ parameters. *Neural Computation*, 9, 337–348.
- C. F. Stevens, & Zador, A. (1996). Information through a spiking neuron. In *Advances in neural information processing systems*, Vol. 8 (pp. 75–81). Cambridge: MIT Press.
- Taylor, J. G., & Alavi, F. N. (1993). Mathematical analysis of a competitive network for attention. In J. G. Taylor (Ed.), *Mathematical approaches to neural network* (pp. 341–382). Amsterdam: North Holland.
- Thorpe, S. T., & Imbert, M. (1989). Biological constraints on connectionist modelling. In R. Pfeifer, Z. Schreter, F. Fogelman-Soulié, & L. Steels (Eds.), *Connectionism in perspective* (pp. 63–92). Amsterdam: Elsevier, North Holland.
- Tuckwell, H. C. (1988). *Introduction to theoretical neurobiology*. Vols. 1 and 2. Cambridge: Cambridge University Press.
- Valiant, L. G. (1994). *Circuits of the mind*. Oxford University Press.
- Zaghloul, M. L., Meador, J. L., & Newcomb, R. W. (Eds.) (1994). *Silicon implementations of pulse coded neural network*. Kluwer.
- Zhao, J. (1995). *Stochastic bit stream neural networks: theory, simulations and applications*. Ph.D. thesis, University of London, London.