

Spatio-temporal Spike Pattern Classification in Neuromorphic Systems

Sadique Sheik, Michael Pfeiffer, Fabio Stefanini, and Giacomo Indiveri

Institute of Neuroinformatics, University of Zurich and ETH Zurich,
Wintherthurerstrasse 190, 8057 Zurich, Switzerland

Abstract. Spike-based neuromorphic electronic architectures offer an attractive solution for implementing compact efficient sensory-motor neural processing systems for robotic applications. Such systems typically comprise event-based sensors and multi-neuron chips that encode, transmit, and process signals using spikes. For robotic applications, the ability to sustain real-time interactions with the environment is an essential requirement. So these neuromorphic systems need to process sensory signals continuously and instantaneously, as the input data arrives, classify the spatio-temporal information contained in the data, and produce appropriate motor outputs in real-time. In this paper we evaluate the computational approaches that have been proposed for classifying spatio-temporal sequences of spike-trains, derive the main principles and the key components that are required to build a neuromorphic system that works in robotic application scenarios, with the constraints imposed by the biologically realistic hardware implementation, and present possible system-level solutions.

1 Introduction

Spiking neural networks represent a promising computational paradigm for solving complex pattern recognition and sensory processing tasks that are difficult to tackle using standard machine vision and machine learning techniques [1, 2].

Most research on spiking neural networks is done using software simulations [3]. For large networks such approaches require high computational power and need to be run on bulky and power-hungry workstations or distributed computing facilities. Alternative approaches, for example based on full-custom implementations of spiking neural networks built using analog-digital Very Large Scale Integration (VLSI) circuits, can lead to compact and real-time solutions that are optimally suited for robotic applications. Systems built following this approach would involve multiple neuromorphic VLSI sensory and processing devices interfaced among each other to perform complex cognitive tasks. For example, they would use event-based sensors [4] for producing sensory signals in the form of continuous and asynchronous streams of spikes, and spiking multi-neuron chips [5, 6, 7] for implementing state dependent neural processing, learning mechanisms, and generation of appropriate behaviors.

Within this context, one of the basic requirements for exhibiting complex behaviors is the system's ability to classify spike patterns generated by the sensors. In real-world application scenarios, these spike patterns would be typically spread across space and time reflecting the *spatio-temporal* properties of the stimuli they are encoding. Several experiments in various animal species have shown that behaviorally relevant information is encoded in spatio-temporal spike patterns, and used for decision making. Barn owls for example rely on auditory cues to capture the prey, by accurately determining the spatial location of the sound source. This is done by computing the temporal delay in the spikes originating at the two ears of the owl [8, 9]. Remarkably precise temporal codes have been found in bush crickets [10] in early auditory inter-neurons, even in the presence of strong natural background noise, and such precise temporal coding could reduce the uncertainty involved in decision making.

The timing of spikes does not necessarily have to encode only temporal features of the input signal, but can also carry information about other aspects. For example, in the human peripheral nervous system, Johansson et al. [11] found that the response of afferent projections from touch receptors in the fingertips carried significant information about pressure in the relative timing of the first spike.

To utilize information encoded in spike times, temporal codes typically require a signal that indicates the onset of the stimulus. In the cortex there is evidence for an alternative model of temporal coding in which information is encoded by the phase relative to one or several background oscillations [12], rather than relative to an external input signal. As an example, hippocampal place cells in rodents encode the spatial location of the animal by their spike timing in relation to the phase of gamma band oscillation [13]. A more detailed analysis of the role of temporal coding in biological systems is discussed e.g. in [14]. It is clear therefore that biological systems employ multiple mechanisms and resources for efficiently processing and classifying spatio-temporal spike patterns. If we want to build biologically inspired neuromorphic systems that utilize similar concepts for robotic applications, we therefore first have to answer the question: *what are the essential components that enable an electronic neuromorphic system to accomplish this task?*

Here we will address this question by evaluating the state-of-the-art spike-based spatio-temporal processing models proposed in the literature and determining their common underlying principles. We will then highlight the general principles that can be applied to design efficient spatio-temporal spike pattern classification mechanisms with electronic neuromorphic components, and present examples of possible hardware solutions for real-time robotic applications.

2 Single Neuron Models

A wide range of models has been proposed in the past that employ single neurons to classify spatio-temporal spike patterns [15, 16, 17, 18, 19]. All these models employ some form of spike based learning, such as Spike Timing Dependent

Plasticity (STDP), on the afferent input synapses that converge onto a single neuron to adjust their synaptic weights. These neurons are then typically exposed to multiple trains of input spike patterns, which contain a specific temporal pattern of spikes presented embedded in them, and presented repeatedly during the training phase. Eventually the synapses settle to a set of weights that make the neuron respond selectively to presentations of the embedded spike patterns.

A detailed analysis of what type of patterns neurons can learn through STDP is presented in [20]. The authors show that the synaptic weights converge to the *perceptron* learning rule if the input spike trains have Poisson distributions. But it does not address the cases in which the input spike trains have temporal structures that can be critical for classification tasks.

The type of features that can be detected with such types of patterns depends on the exact form of learning rule, and the underlying neuron model. We will discuss two of the most commonly used neuron models, the Poisson and the Leaky Integrate and Fire (LIF) neuron models.

2.1 Coincidence Detection

Let us first consider a Poisson neuron that has a probability of spiking proportional to its instantaneous input firing rate, which can be derived from the current membrane potential, given by the sum of all input Excitatory Post-Synaptic Potential (EPSP)s. For constant mean firing rates of input spike trains the post-synaptic neuron's membrane potential fluctuates around a mean value. In this condition synchronous inputs increase the instantaneous input firing rate, and therefore increase the probability of a post-synaptic spike. This model can therefore be used directly to detect synchronicity [21].

Coherent spiking activity of a population of neurons is easily detectable with LIF neurons [22], although the analysis of the firing properties of LIF model is not as straight forward as in the Poisson case. In [21] the authors show how the threshold and membrane time constants of a LIF neuron determine its ability to act as a coincidence detector. They demonstrate that even with constant mean input rate the mean output rate varies as a function of *temporal structure* of the input. They define “temporal structure” as the amount of spikes in the input spike trains that are phase locked to a periodic stimulus. They argue that a “small” membrane time constant is essential for a coincidence detector. The small membrane time constant should be chosen such that the mean membrane potential, under uniform input spike rate, is just below the neuron's firing threshold. Under these conditions, the neuron is sensitive to the fluctuations in the membrane potential, and can detect coincidences in its input spike trains in the same way that the Poisson neuron model described above does. It is this coincidence-detection ability of the LIF model that allows single-neuron training schemes to learn to respond selectively to specific spatio-temporal spike train sequences [23].

But to what extent is coincidence detection in spike-patterns enough to distinguish patterns with complex spatio-temporal structure?

To answer this question, let us consider a simple spatio-temporal pattern recognition task that involves patterns of spike-trains. Take for example those produced by an event-based (spiking) vision sensor such as the Dynamic Vision Sensor (DVS) [24], in which every pixel responds only to temporal changes in light intensity in its visual field. When observing a visual stimulus moving from left-to-right and right-to-left, under ideal conditions, the left-ward movement of the visual stimulus would be equivalent to a temporal reversal of the right-ward movement. The spike patterns produced in response to the stimulus moving in one direction would then be equivalent to the time-reversed version of the patterns produced by the stimulus moving in the opposite direction.

We can transpose this example to a more generic case: let's consider for example five specific randomly generated spatio-temporal spike patterns and their time-reversed versions, as in the case of the visual stimulus moving in two opposite directions. By using a spike-based STDP learning rule, such as the one proposed in [25], we can train six Poisson neurons arranged in a Winner-take-all (WTA) network to recognize all five patterns (presenting repeatedly only the original non-reversed patterns), as well as the background activity when none of the five patterns is present. After training, five of the six neurons fire selectively in response to one of the five input patterns, and a 6th neuron fires in response to the background activity. We trained such a network, and present the response of the neurons during presentation of their corresponding preferred spike patterns in Fig. 1. Each spike pattern and the corresponding neuronal response is marked with a distinct color.

Figure. 1 however shows also the activity of the same set of neurons in response to the time-reversed versions of the original spike patterns. As is evident from these results, the output neurons respond selectively not only to the trained spike patterns but also to their time-reversed versions. This is because even when the spike patterns are reversed the key feature the learning mechanism captures is *coincidences* (or near-coincidences) in the spike pattern. The pattern of coincidences to which a neuron responds is encoded in the vector of weights, and thus such neurons cannot distinguish between patterns that share the same sets of coincident inputs. Similar results can be obtained using analogous learning algorithms, such as the “tempotron” model [18], which is designed to capture coincidences hidden in one class of patterns and not in the other.

2.2 Beyond Coincidence Detection

In order to train neurons to distinguish more complex temporal patterns, without relying solely on coincidences it is necessary to introduce additional mechanisms that make the neuron sensitive to the causality of the input pattern, e.g., by responding differently to reversed patterns. Multi-compartmental models, conductance-based neurons [26], short-term synaptic plasticity, and dendritic computation [27, 28, 29, 30, 31] are all examples of strategies that allow such properties.

In [32], the authors demonstrate a neuromorphic hardware implementation where single neurons with short term depression can demonstrate *Stimulus-Specific*

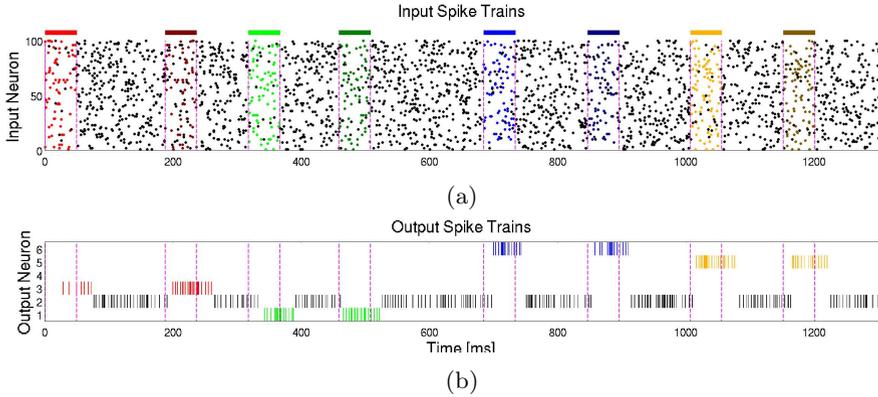


Fig. 1. a Raster plot of input spike pattern. The lighter colors, red, green, blue and yellow mark 4 of the 5 spike patterns that were used to train the network. The darker colored patterns, following each of the above colored spike patterns are time reversed spike patterns corresponding to their color. Spikes shown in black are background noise. b Response of output Poisson neurons on presentation of the input spike patterns. After training with a set of spike patterns, the neurons respond selectively to both the original and to time-reversed version of these spike patterns.

Adaptation, a phenomenon observed in neural systems. They show how short-term depression can enable neurons to adapt to specific stimuli based on the frequency of their occurrence and therefore show higher response to rare events. This behavior can be attributed to the fact that short-term plasticity lasts over hundreds of milliseconds, and carries information about synaptic inputs in the recent past. This mechanism has also been exploited in [33], where it is demonstrated that temporal information can be decoded using short-term plasticity. The author even hypothesizes that manipulations that eliminate short-term plasticity will produce deficits in temporal processing and therefore deem it a critical computational component for temporal processing.

3 Network Models

An alternative strategy, that does not require more elaborate neuron models and dedicated circuitry, is to exploit the dynamics of a neural *network* to obtain the needed sensitivity. This approach has led to the development of several network models employed in spatio-temporal pattern recognition tasks.

Several models based on networks of recurrently coupled LIFs neurons have been proposed to classify spatio-temporal patterns. For example it has been shown how in a Liquid-State Machine (LSM) [34] that the low-pass filtered activity of a recurrently connected network (the reservoir) at any point in time reflects the temporal evolution of the input stimulus. By training a perceptron on snapshots of the neural activity in the reservoir, the liquid states, it is therefore possible to classify input patterns that have complex temporal evolution [35].

The parameters needed to obtain optimal performance are such that the input activity reverberates over a time-scale of several neuronal time-constants. With a careful choice of those parameters, the network is able to correlate events that are distant in time and that have a certain causality relationship. Neuromorphic circuits with slow dynamics for neuron and synapses can be used to exploit the reverberating properties of the network similarly to a LSM to process real-world data from spiking sensors.

An analogous approach is described in [36], where a recurrently-connected network of spiking neurons implemented in neuromorphic hardware is trained to differentiate between a sequence of spikes regularly distributed in time and its reversed version. It has been shown that this type of computation is involved in the recognition of tone sweeps in the thalamo-cortical pathway in the auditory stream [37]. The neural network consists of a rack of coincidence detectors receiving the spikes directly from the input layer and from delay lines propagating the activity of the other neurons in the network. Synapses are modified through STDP such that only the synapses capturing the temporal correlations typical of the learned stimulus survive. A similar principle could be exploited in polychronous networks [38], where random axonal delays in a recurrent network can be used for spike pattern classification.

These models are closely related to the general concept of the Reichardt detector [39]. In a Reichardt detector two or more receptors respond at different times to the input stimuli and project to a non-linear unit, such as a neuron. Each of these projections incurs some amount of delay due to low-pass filtering. If these delays match the activation delays caused by the input stimulus, the signals from the receptors are coincident at their destination and consequently activate the nonlinear unit. Models of this kind have the advantage that they explicitly encode time within specific delay elements and lend themselves to unsupervised online learning methods. Depending on the task, such networks can be optimized to be implemented with a smaller number of neurons as compared to LSMs.

4 Neuromorphic Building Blocks for Temporal Processing

In the light of the above analysis we identified the basic computational operators required for processing of spatio-temporal spiking patterns in hardware implementations of spiking neural network. These computational operators include voltage-gating mechanisms that can make neurons sensitive to coincidences [21, 14, 28, 40], nonlinear operators such as short-term plasticity [33], low-pass filters [39] and delay elements [41].

Complex neural processing systems are likely to require large numbers of such operators. It is therefore important to design efficient hardware implementations using compact and low-power neuromorphic circuits [5]. Spiking neural networks built using such neuromorphic circuits comprise massively parallel arrays of processing nodes with memory and computation co-localized. Given their

architectural constraints, these neural processing systems cannot process signals using the same strategies used by conventional von Neumann computing architectures that exploit time-domain multiplexing of small numbers of highly complex processors at high clock-rates, and operate by transferring the partial results of the computation from and to external memory banks. The synapses and neuron circuits in neuromorphic architectures have to process input spikes and produce output responses in real-time at the rate of the input data. It is not possible to virtualize time and transfer partial results in memory banks outside the architecture core, at higher rates, but it is necessary to employ resources that compute with time-constants that are well matched to those of the signals they are designed to process. Therefore, to interact with the environment and process signals with biological time-scales efficiently, hardware neuromorphic systems need to be able to compute using biologically realistic time constants. In this way, they are well matched to the signals they process, and are inherently synchronized with the real world events.

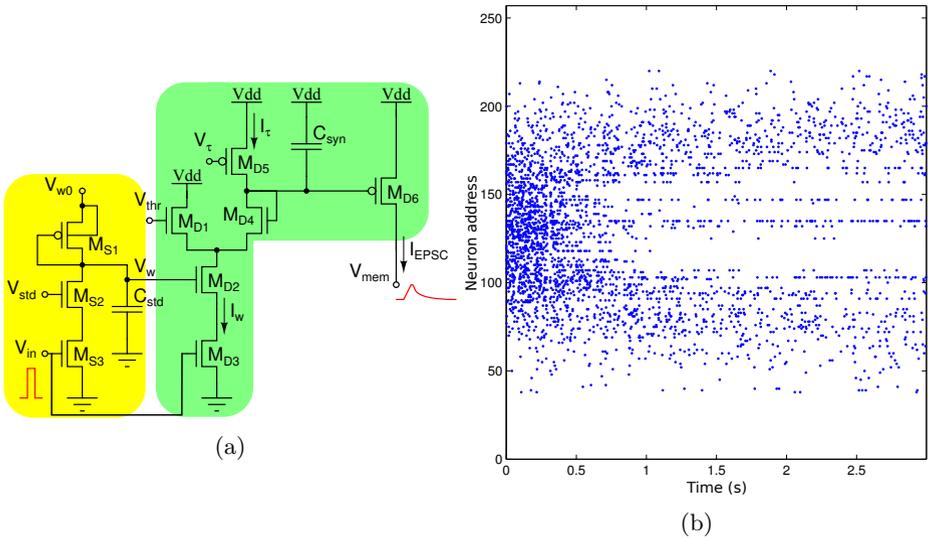


Fig. 2. a Analog synapse circuit that can implement biologically realistic synaptic dynamics and short-term depression. b Raster plot: response of a population of silicon neurons with short-term depressing synapses to a Gaussian profile of Poisson input spike trains.

Implementing biologically realistic time constants in analog VLSI technology is a non-trivial task. Conventional analog design techniques, in which typical currents are in the range of micro- or milliampere would require extremely large capacitors to implement time constants of the order of milliseconds, therefore losing the ability to integrate very large numbers of synapses in small areas. An elegant and efficient solution is that of using analog circuits that operate in the “weak inversion regime” [42], in which currents are on the order of pico-amperes.

An analog circuit of this type is the Differential-Pair Integrator (DPI), originally proposed in [43]. This circuit implements a low-pass filter which faithfully reproduces synaptic dynamics. An example of such synapse circuit is shown in Fig. 2a. It has been shown [44], by log-domain circuit analysis techniques [45, 46], that the transfer function of the part of the circuit composed by transistors $MD1 - MD6$ is:

$$\tau \frac{d}{dt} I_{EPSC} + I_{EPSC} = \frac{I_w I_{thr}}{I_\tau} \quad (1)$$

under the assumption that the input current I_w is larger than the leak current I_τ . The term τ represents the circuit time constant, defined as $\tau = CU_T/\kappa I_\tau$. The term U_T represents the thermal voltage and κ the sub-threshold slope factor [42]. The term I_{thr} represents a virtual p-type sub-threshold current that is not generated by any p-FET in the circuit, and is defined as $I_{thr} = I_0 e^{-\frac{\kappa(V_{thr} - V_{dd})}{U_T}}$. This circuit therefore implements an extremely compact and efficient model of synaptic dynamics. The transistors $MS1 - MS3$ of the DPI circuit of Fig. 2a implement a very useful non-linearity for spatio-temporal processing, that models the short-term depression properties of biological synapses [47]. This circuit is fully characterized in [48]. In [49] the authors show how this circuit can faithfully reproduce the behavior of real short-term depressing synapses, and demonstrate the equivalence with the computational model of short-term depression described in [50]. We present quantitative experimental data from this circuit in [44]. We show in Fig. 2b measurements from a population of silicon neurons that have these short-term depressing synapse circuits. The raster plot of Fig. 2b shows the response of 256 neurons, arranged along a line. The input pattern is a Gaussian profile of Poisson spike trains, with maximum firing rates centered around the address 128. As shown, high firing rate inputs, at the center of the array, are suppressed or strongly attenuated by the short-term depression mechanism of the input synapses, while low firing rate inputs at the periphery are transmitted without attenuation.

While short-term depression can be very useful for processing temporal patterns of spike trains [33, 51], additional essential elements are temporal delays and delay lines. Explicit delay circuits can be implemented in multi-compartmental neuron models to carry out spatio-temporal processing. For example Wang and Liu [28] demonstrate a VLSI neuron chip with programmable dendritic compartments and delay elements, showing how different spatio-temporal input patterns have different effects on the evoked dendritic integration. Analogous approaches have been proposed using floating-gate structures [30, 52]. But simpler Integrate and Fire (IF) neuron circuits can also implement delay elements without having to implement complex dendritic spatial structures. In [53] the authors show how slow dynamics of DPI synapses, combined with the variability and in-homogeneity “features” of neuromorphic analog circuits, can be used to generate a range of temporal delays. The low-pass filtering properties of the DPI synapses can be effectively used to create delay-elements, and these can be integrated in multi-chip networks of spiking neurons, feed-forward or recurrent, to increase their temporal retention capability. Specifically, in [53] the authors show a thalamo-cortical

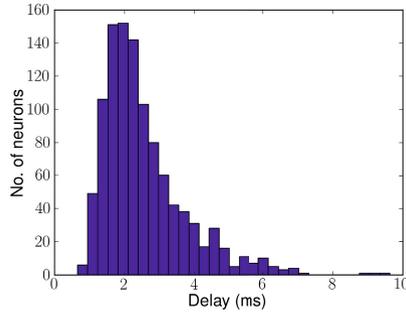


Fig. 3. Histogram of times used by post-synaptic neurons to fire after the arrival of a single pre-synaptic input spike (*delay*), across a population of neurons on an analog VLSI neuromorphic chip

neural network model that can learn to recognize specific spatio-temporal patterns of input spikes, by exploiting the mismatch present in the transistors of the DPI circuits giving rise to a wide range of delays (see also Fig. 3).

5 Conclusion

Spatio-temporal patterns produced by neuromorphic sensors reflect the time scales of events coming from the environment in real-time. In this paper we highlight the role of slow components in neural network models in order to classify spatio-temporal patterns produced by these sensors. We show a case in which IF neurons are capable of detecting coincidences and are incapable of differentiating the temporal order in the input spike pattern. We argue that this limitation is common to all single point-neuron models with first-order dynamics. We further review possible mechanisms that can be introduced to enable spatio-temporal pattern discriminability at the network level and capture the causality of the input stimuli. We present analog VLSI circuits for synaptic dynamics that have been included in recent neuromorphic implementations and can be exploited, together with the variability of the neuromorphic hardware, to classify spatio-temporal spike patterns.

Acknowledgements. This work was supported by the EU ERC Grant “neuroP” (257219).

References

- [1] Maass, W., Sontag, E.: Neural systems as nonlinear filters. *Neural Computation* 12(8), 1743–1772 (2000)
- [2] Belatreche, A., Maguire, L.P., McGinnity, M.: Advances in design and application of spiking neural networks. *Soft Computing* 11(3), 239–248 (2006)

- [3] Brette, R., Rudolph, M., Carnevale, T., Hines, M., Beeman, D., Bower, J., Diesmann, M., Morrison, A., Goodman, P.H.J.F., Zirpe, M., Natschläger, T., Pecevski, D., Ermentrout, B., Djurfeldt, M., Lansner, A., Rochel, O., Vieville, T., Muller, E., Davison, A., El Boustani, S., Destexhe, A.: Simulation of networks of spiking neurons: A review of tools and strategies. *Journal of Computational Neuroscience* 23(3), 349–398 (2007)
- [4] Liu, S.C., Delbruck, T.: Neuromorphic sensory systems. *Current Opinion in Neurobiology* 20(3), 288–295 (2010)
- [5] Indiveri, G., Linares-Barranco, B., Hamilton, T., van Schaik, A., Etienne-Cummings, R., Delbruck, T., Liu, S.C., Dudek, P., Häfziger, P., Renaud, S., Schemmel, J., Cauwenberghs, G., Arthur, J., Hynna, K., Folowosele, F., Saighi, S., Serrano-Gotarredona, T., Wijekoon, J., Wang, Y., Boahen, K.: Neuromorphic silicon neuron circuits. *Frontiers in Neuroscience* 5, 1–23 (2011)
- [6] Choudhary, S., et al.: Silicon neurons that compute. In: Villa, A.E.P., Duch, W., Érdi, P., Masulli, F., Palm, G. (eds.) *ICANN 2012, Part I. LNCS*, vol. 7552, pp. 121–128. Springer, Heidelberg (2012)
- [7] Yu, T., Park, J., Joshi, S., Maier, C., Cauwenberghs, G.: 65k-neuron integrate-and-fire array transceiver with address-event reconfigurable synaptic routing. In: *Biomedical Circuits and Systems Conference (BioCAS)*, pp. 21–24. IEEE (November 2012)
- [8] Carr, C.E., Konishi, M.: Axonal delay lines for time measurement in the owl's brainstem. *Proceedings of the National Academy of Sciences* 85(21), 8311–8315 (1988)
- [9] Carr, C.E., Konishi, M.: A circuit for detection of interaural time differences in the brain stem of the barn owl. *The Journal of Neuroscience* 10(10), 3227–3246 (1990)
- [10] Pfeiffer, M., Hartbauer, M., Lang, A.B., Maass, W., Römer, H.: Probing real sensory worlds of receivers with unsupervised clustering. *PloS One* 7(6), e37354 (2012)
- [11] Johansson, R., Birznieks, I.: First spikes in ensembles of human tactile afferents code complex spatial fingertip events. *Nature Neuroscience* 7(2), 170–177 (2004)
- [12] Singer, W.: Time as coding space? *Current Opinion in Neurobiology* 9(2), 189–194 (1999)
- [13] O'Keefe, J., Burgess, N.: Geometric determinants of the place fields of hippocampal neurons. *Nature* 381(6581), 425–428 (1996)
- [14] Stiefel, K.M., Tapson, J., van Schaik, A.: Temporal order detection and coding in nervous systems. *Neural Computation* 25(2), 510–531 (2013)
- [15] Dhoble, K., Nuntalid, N., Indiveri, G., Kasabov, N.: Online spatio-temporal pattern recognition with evolving spiking neural networks utilising address event representation, rank order, and temporal spike learning. In: *International Joint Conference on Neural Networks, IJCNN 2012*, pp. 554–560. IEEE (2012)
- [16] Nessler, B., Pfeiffer, M., Maass, W.: Stdp enables spiking neurons to detect hidden causes of their inputs. In: Bengio, Y., Schuurmans, D., Lafferty, J., Williams, C.I., Culotta, A. (eds.) *Advances in Neural Information Processing Systems*, vol. 22, pp. 1357–1365 (2009)
- [17] Masquelier, T., Guyonneau, R., Thorpe, S.J.: Spike timing dependent plasticity finds the start of repeating patterns in continuous spike trains. *PLoS One* 3(1), e1377 (2008)
- [18] Gütig, R., Sompolinsky, H.: The tempotron: a neuron that learns spike timing-based decisions. *Nature Neuroscience* 9, 420–428 (2006)

- [19] Thorpe, S., Delorme, A., Van Rullen, R., et al.: Spike-based strategies for rapid processing. *Neural Networks* 14(6-7), 715–725 (2001)
- [20] Legenstein, R., Nager, C., Maass, W.: What can a neuron learn with spike-timing-dependent plasticity? *Neural Computation* 17(11), 2337–2382 (2005)
- [21] Kempter, R., Gerstner, W., Van Hemmen, J.L.: How the threshold of a neuron determines its capacity for coincidence detection. *Biosystems* 48(1), 105–112 (1998)
- [22] Gerstner, W., Kistler, W.: *Spiking Neuron Models. In: Single Neurons, Populations, Plasticity.* Cambridge University Press (2002)
- [23] Masquelier, T., Guyonneau, R., Thorpe, S.J.: Competitive stdp-based spike pattern learning. *Neural Computation* 21(5), 1259–1276 (2009)
- [24] Lichtsteiner, P., Posch, C., Delbruck, T.: A 128×128 120dB 30mW asynchronous vision sensor that responds to relative intensity change. In: 2006 IEEE ISSCC Digest of Technical Papers, pp. 508–509. IEEE (February 2006)
- [25] Nessler, B., Pfeiffer, M., Maass, W.: Bayesian computation emerges in generic cortical microcircuits through spike-timing-dependent plasticity. *PLoS Computational Biology* (2013)
- [26] Gutig, R., Sompolinsky, H.: Time-warp-invariant neuronal processing. *PLoS Biology* 7(7), e1000141 (2009)
- [27] Koch, C., Poggio, T., Torre, V.: Nonlinear interactions in a dendritic tree: Localization, timing, and role in information processing. *Proceedings of the National Academy of Sciences of the USA* 80, 2799–2802 (1983)
- [28] Wang, Y., Liu, S.C.: Multilayer processing of spatiotemporal spike patterns in a neuron with active dendrites. *Neural Computation* 8, 2086–2112 (2010)
- [29] Arthur, J., Boahen, K.: Recurrently connected silicon neurons with active dendrites for one-shot learning. In: *IEEE International Joint Conference on Neural Networks*, vol. 3, pp. 1699–1704 (July 2004)
- [30] Ramakrishnan, S., Wunderlich, R., Hasler, P.: Neuron array with plastic synapses and programmable dendrites. In: *Biomedical Circuits and Systems Conference (BioCAS)*, pp. 400–403. IEEE (November 2012)
- [31] Rasche, C., Douglas, R.: Forward- and backpropagation in a silicon dendrite. *IEEE Transactions on Neural Networks* 12, 386–393 (2001)
- [32] Mill, R., Sheik, S., Indiveri, G., Denham, S.: A model of stimulus-specific adaptation in neuromorphic aVLSI. In: *Biomedical Circuits and Systems Conference (BioCAS)*, pp. 266–269. IEEE (2010)
- [33] Buonomano, D.: Decoding temporal information: A model based on short-term synaptic plasticity. *The Journal of Neuroscience* 20, 1129–1141 (2000)
- [34] Maass, W., Natschlager, T., Markram, H.: Real-time computing without stable states: A new framework for neural computation based on perturbations. *Neural Computation* 14(11), 2531–2560 (2002)
- [35] Maass, W., Natschlager, T., Markram, H.: Fading memory and kernel properties of generic cortical microcircuit models. *Journal of Physiology – Paris* 98(4-6), 315–330 (2004)
- [36] Sheik, S., Coath, M., Indiveri, G., Denham, S., Wennekers, T., Chicca, E.: Emergent auditory feature tuning in a real-time neuromorphic VLSI system. *Frontiers in Neuroscience* 6(17) (2012)
- [37] Coath, M., Mill, R., Denham, S.L., Wennekers, T.: Emergent Feature Sensitivity in a Model of the Auditory Thalamocortical System. *Advances in Experimental Medicine and Biology*, vol. 718, pp. 7–17. Springer, New York (2011)
- [38] Izhikevich, E.M.: Polychronization: Computation with spikes. *Neural Computation* 18(2), 245–282 (2006)

- [39] Reichardt, W.: Autocorrelation, a principle for the evaluation of sensory information by the central nervous system. *Sensory Communication*, 303–317 (1961)
- [40] Wyss, R., König, P., Verschure, P.F.: Invariant representations of visual patterns in a temporal population code. *Proceedings of the National Academy of Sciences* 100(1), 324–329 (2003)
- [41] Jeffress, L.A.: A place theory of sound localization. *J. Comp. Physiol. Psychol.* 41(1), 35–39 (1948)
- [42] Liu, S.C., Kramer, J., Indiveri, G., Delbruck, T., Douglas, R.: *Analog VLSI: Circuits and Principles*. MIT Press (2002)
- [43] Bartolozzi, C., Indiveri, G.: Synaptic dynamics in analog VLSI. *Neural Computation* 19(10), 2581–2603 (2007)
- [44] Bartolozzi, C., Mitra, S., Indiveri, G.: An ultra low power current–mode filter for neuromorphic systems and biomedical signal processing. In: *Biomedical Circuits and Systems Conference (BioCAS)*, pp. 130–133. IEEE (2006)
- [45] Drakakis, E., Payne, A., Toumazou, C.: “Log-domain state-space”: A systematic transistor-level approach for log-domain filtering. *IEEE Transactions on Circuits and Systems II* 46(3), 290–305 (1999)
- [46] Frey, D.: Log-domain filtering: An approach to current-mode filtering. *IEE Proceedings G: Circuits, Devices and Systems* 140(6), 406–416 (1993)
- [47] Markram, H., Tsodyks, M.: Redistribution of synaptic efficacy between neocortical pyramidal neurons. *Nature* 382, 807–810 (1996)
- [48] Rasche, C., Hahnloser, R.: Silicon synaptic depression. *Biological Cybernetics* 84(1), 57–62 (2001)
- [49] Boegerhausen, M., Suter, P., Liu, S.C.: Modeling short-term synaptic depression in silicon. *Neural Computation* 15(2), 331–348 (2003)
- [50] Varela, J., Sen, K., Gibson, J., Fost, J., Abbott, L., Nelson, S.: A quantitative description of short–term plasticity at excitatory synapses in layer 2/3 of rat primary visual cortex. *The Journal of Neuroscience* 17, 7926–7940 (1997)
- [51] Mill, R., Sheik, S., Indiveri, G., Denham, S.: A model of stimulus-specific adaptation in neuromorphic analog VLSI. *Transactions on Biomedical Circuits and Systems* 5(5), 413–419 (2011)
- [52] Basu, A., Ramakrishnan, S., Petre, C., Koziol, S., Brink, S., Hasler, P.: Neural dynamics in reconfigurable silicon. *IEEE Transactions on Biomedical Circuits and Systems* 4(5), 311–319 (2010)
- [53] Sheik, S., Chicca, E., Indiveri, G.: Exploiting device mismatch in neuromorphic VLSI systems to implement axonal delays. In: *International Joint Conference on Neural Networks, IJCNN 2012*, pp. 1940–1945. IEEE (2012)