

# Event-based attention and tracking on neuromorphic hardware

Alpha Renner

Institute of Neuroinformatics  
University of Zurich and ETH Zurich  
alpren@ini.uzh.ch

Matthew Evanusa

University of Maryland, MD, USA  
matthew.evanusa@gmail.com

Yulia Sandamirskaya

Institute of Neuroinformatics  
University of Zurich and ETH Zurich  
ysandamirskaya@ini.uzh.ch

## Abstract

*We present a fully event-driven vision and processing system for selective attention and tracking, realized on a neuromorphic processor Loihi interfaced to an event-based Dynamic Vision Sensor DAVIS. The attention mechanism is realized as a recurrent spiking neural network that implements attractor-dynamics of dynamic neural fields. We demonstrate capability of the system to create sustained activation that supports object tracking when distractors are present or when the object slows down or stops, reducing the number of generated events.*

## 1. Introduction

Event-based sensors send out data packages, or events, from each pixel asynchronously when the pixel detects a local brightness change, rather than reading all pixels and sending out frames at a constant rate. Such event-based sensing allows us to perform some vision tasks extremely efficiently, reducing the amount of required computation, transmitted data, and power consumption. Many event-based vision pipelines and architectures have been developed over the last decade [1, 2], which address such vision tasks as stereo vision [3], 3D pose estimation [4, 5], or optical flow [6]. These event-based pipelines are typically implemented on conventional Von Neumann computer architectures. While such implementations try to make the best out of the event-driven nature of the sensor output, they cannot fully utilize its advantages: the clocked, sequential operation of a CPU, as well as separation between memory and processor stay in contrast to the highly parallel asynchronous temporal stream of events coming from an event-based sensor.

Neuromorphic hardware, in contrast to the conventional

CPU, offers a massively parallel computing substrate that is inherently event-based, which matches the processing paradigm of event-based sensors [7, 8, 9, 10]. We aim to develop neuronal architectures that solve different tasks on such neuromorphic devices taking full advantage of the event-driven computation. In this work, we target a vision task of object tracking; in particular we show how a simple attention network can be configured on a neuromorphic hardware to select one object in an event-based input stream and to track this object in presence of equally salient distractors. While similar principles to the one realized here have been used in early days of neuromorphic engineering to design a dedicated spiking neuromorphic chip for attention [11], here we present its realization on a generic neuromorphic device Loihi that can also support other vision, cognitive, and motor control tasks [8].

Neuromorphic processors emulate dynamics of biological spiking neurons in hardware and thus allow us to run spiking neural network architectures in real-time, with a small energy footprint, and in small form-factor devices, making them a promising computing platform for event-based vision [7, 8]. However, most neuromorphic devices target offline computation, with applications in either computational neuroscience [10, 12] or data processing [9]. The Kapoho Bay – a USB-stick form-factor version of the Intel’s latest neuromorphic research platform Loihi [8] – is among neuromorphic systems that offer a direct Address-Event Representation (AER) interface to event-based sensors [13]. This allows us to build a setup, in which a neuromorphic camera DAVIS [14] can be directly interfaced to Loihi, stimulating on-chip neurons configured in a network that can solve vision tasks. Here, we focus in particular on the task of object-centered attention and tracking.

While the event-based output of a DAVIS camera singles out a fast-moving object easily, two capabilities need addi-

tional processing: the ability to suppress distractors even if their salience changes and at times surpasses that of the target object and; keeping track of an object if it slows down or even stops. To gain these capabilities, the system needs a mechanism to hold a memory of the object’s location in the field of view. While such a memory mechanism can be realized on a conventional computing system, doing so would alleviate the advantages of the low-power event-based computing. Here we explore a setup in which an event-based sensor is interfaced to an event-based processor running a recurrent neural network that is capable of creating memory states based on the incoming events.

Feed-forward artificial neural networks (ANNs) are stateless – they merely transfer inputs to outputs and if they do not receive input, the activity in the network fades away. In a spike-based network with integrate-and-fire neurons, the activity decays with a time-constant of neuronal dynamics, in a conventional ANN, even more radically, on the next clock cycle. In order to create a memory state, recurrence in the network is needed. A well-known model for working memory that has been studied in computational neuroscience and cognitive science is a Dynamic Neural Field (DNF) [15] – a neural population-based model. The DNF is a dynamical system that can be realized as a recurrent neural network with attractor dynamics, created by configuring a population of neurons with a winner-take-all connectivity [16]. In this connectivity pattern, neurons that encode similar values have an excitatory connection and neurons that encode different values – an inhibitory one. This simple connectivity pattern performs a selective amplification of a noisy input and, in an extreme case of strong interaction, can create sustained activation patterns that are kept active even if the initial input ceases completely. This property has been used as a model of working memory [15].

DNFs have been used previously to realize object tracking with on SCAMP – a smart camera with an in-focal-plane processor array [17]. Here, we demonstrate object tracking with DNFs in an event-based setting using a spiking neuromorphic device Loihi.

## 2. The hardware setup

### 2.1. Neuromorphic Device Loihi

Intel Neuromorphic Computing Lab designed the neuromorphic research chip Loihi, in which spiking neural network models can be simulated in real-time efficiently [8]. The chip consists of a mesh of 128 neuromorphic cores, three embedded x86 processor cores, and an off-chip communication interface that allows to scale up architectures to multiple Loihi devices. Compartments are the main building blocks used to configure both single- and multi-compartment neurons. In this work we only use single-compartment neurons.

The external input to a network on Loihi is provided through spike generators. Spike generators are ports connected to compartments that can generate spikes at precise time-steps. Loihi provides an instrument for measuring the variables and sending them off the chip using “probes”. For compartments it is possible to define probes to measure spike events, neuron’s membrane voltage and input current. For the connections, probes can measure multiple synaptic variables, including weight, pre-synaptic and post-synaptic traces. Probing, however, affects the performance of the chip.

Loihi’s Python NxSDK-0.8.0 API allows us to implement SNNs on the chip [18]. The NxNet API provides ways to define a graph of neurons and synapses and configure their parameters (such as decay time constants, spike impulse values, synaptic weights, refractory delays, spiking thresholds), inject external stimulus into the network, implement custom learning rules, and monitor and modify the network during runtime.

### 2.2. Dynamic Vision Sensor DAVIS

In this work, we used a Dynamic Vision Sensor type of a camera, the DAVIS240C [14]. The DAVIS camera emulates the dynamics of biological retinal cells in silicon using mixed-signal analog/digital technologies. There are  $240 \times 180$  pixels integrated on the chip. Each pixel independently detects the brightness change in a small area of the visual scene and emits an event if the brightness change passes a positive (“on” event) or a negative (“off” event) threshold. Each event is a digital data packet that carries the address of the pixel, the polarity of the detected brightness change, and the time of the event using Address Event Representation. Due to its high dynamic range, the sensor captures moving objects in its visual field in a wide range of lighting conditions.

To connect DAVIS to Loihi, the direct parallel AER interface on the device can be used. The events are captured and distributed to neurons on the neuromorphic cores through the embedded FPGA and x86 processors. For reproducibility, in this work, we use recorded spikes that we feed into Loihi using a spike generator from the NxNet API to generate some of the plots. The network was also tested with the direct AER interface between the DAVIS and Loihi.

## 3. Attractor Dynamics for Object Tracking: the Dynamic Neural Fields

A Dynamic Neural Field is a mathematical model that was derived to describe activity of large homogeneous populations of biological neurons [15]. The connectivity pattern in such a neuronal population is shown in Fig. 1. First, neurons are “aligned” according to a feature which they are

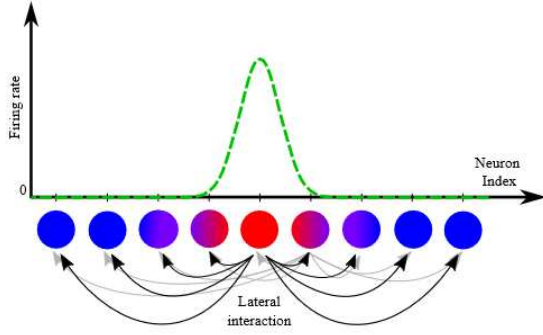


Figure 1. Schematic representation of a 1-dimensional winner-take-all dynamic neural field. The lateral connections are all-to-all and the synaptic weights are defined by the kernel function that depends on the distance between the pre- and post-synaptic neurons.

sensitive to (here, the position in the camera’s field of view). Second, neurons that are sensitive to similar values of the feature (i.e. are close to each other on this behavioral space) are connected via excitatory (positive weight) connections, and neurons that are sensitive to dissimilar features inhibit each other (negative-weight connections). This soft winner-take-all connectivity pattern, combined with a non-linearity of the neuronal activation function, leads to formation of an attractor state in a DNF neuronal population. In particular, DNFs form a so-called bump-attractor – a localized activity peak centred over a salient value of the behavioral variable (green line in Fig. 1). Such bump-attractor networks have been used in the past both to explain activity patterns in biological neural networks and to build artificial cognitive systems for robot control [19, 15]. To realize DNF dynamics on Loihi, we create a group of compartments that are connected with synapses with weights matching the “Mexican-hat” connectivity kernel.

The output of the DNF population is computed as a population vector using the instantaneous firing rate of neurons that is inversely proportional to the inter-spike intervals.

## 4. Results

### 4.1. Attractor dynamics on a neuromorphic Chip

First, we demonstrate the properties of a DNF realized in a spiking neural network on Loihi that are used in the object-tracking application. We have configured a small population of 12 neurons in a winner-take-all fashion, as shown in Fig. 1. We used two sets of parameters to demonstrate two dynamical configurations of the DNF: a selective input-driven regime, shown in Fig. 3a, b, c and a self-sustained regime shown in Fig. 3d. In each subfigure, the upper plot shows spikes from the spike generators that send input to Loihi, the middle plot shows the output spikes from the DNF population, and the lower plot shows population




Figure 2. Performance evaluation.

activity vector of the DNF population (the position of the mean of neuronal activity) over time of the experiment.

Fig. 3a shows that a DNF configured in a selective regime selects one of the input bumps in the case of a bimodal input distribution. For each pair of input bumps with equal strength (average firing rate), the DNF selects one of them randomly. Fig. 3b shows behavior of the same DNF population for an input sequence, in which one of the input bumps arrives first, is selected and then stabilized by the lateral interactions in the neuronal population. In this configuration, the second input bump is rejected by the DNF population and does not lead to any activity of the neurons in the respective region.

Fig. 3c and d contrast behavior of the DNF configured to be in an input-driven regime (c) versus in a self-sustained regime (d). For the same input, the input-driven DNF (Fig. 3c) follows the input activity, only rejecting noise around the activity bumps. The self-sustained DNF (Fig. 3d) keeps the position of the selected object and ignores input at different locations unless it is spatially proximal to the current activity bump. Thus, it shows tracking behavior.

To configure the two DNFs, we used the parameters on Loihi listed in the Table 1. In particular, we use Gaussian-shaped connectivity profile for lateral connections in the WTA (of amplitude “Excitatory weight” and spread “Connectivity kernel  $\sigma$ ”) and a direct global inhibition. Input- and background-noise are generated as Poisson spikes.

### 4.2. Performance evaluation

Fig. 2 shows the duration of simulation on Loihi per time step depending on the number of neurons in a 2D DNF population with a single self-sustained bump without spike generators (apart from the initial one creating the bump) or probing. Neurons are distributed over the 2x128 cores. One can observe that the simulation time per time step stays on the level of  $20\mu s$  for network sizes above 1000 neurons<sup>1</sup>.

<sup>1</sup>Although quite fast already, the simulation time depends on the details of the network implementation and can be further optimized.




Figure 3. The plots show activity of spiking neurons on the Loihi chip, configured as a DNF in an input-driven (a, b, c) and self-sustained (d) regime. In each subfigure, the top plot shows input spikes generated on the computer, the middle plot shows output of the DNF population on Loihi, and the bottom plot shows the population vector (mean of the activity) of the DNF population.

### 4.3. Event-based attention and tracking

The tracking network consists of two two-dimensional WTA/DNF layers with 64x64 neurons each. The first layer has excitatory connections at a level at which peaks are self sustained and weak global inhibition so that multiple bumps can form. Bumps form at the locations of the on-event input and follow the on-events when objects move over the field of view of the DVS. The off-events inhibit neurons in this population, decreasing activity in the bumps. This facilitates fast moving bumps avoiding a tail in the activation

pattern. The second WTA layer receives input from the first layer through one-to-one connections. This layer has strong self-excitation and strong global inhibition which lead to the selection of a single bump. To this layer, we provide an excitatory initial input, cuing one of the objects in the beginning of the DVS stimulation; feature-based cues can also be used here [15] (Chapter 5). The WTA forms an activity bump over the selected object, which is moved by the excitatory input from the input layer when the selected object moves. In these experiments, parameters listed in

Parameter	Input driven	Self-sustained
Voltage threshold	$3000 * 2^6$	$3000 * 2^6$
Voltage decay time constant	150 ts	150 ts
Current decay time constant	10 ts	10 ts
Connectivity kernel $\sigma$	1.5	1.5
Self excitation	no	yes
Excitatory weight	200	150
Global inhibitory weight	-160	-75
Input weight	200	200
Input firing rate (Poisson)	60 Hz	60 Hz
Noise firing rate (Poisson)	2 Hz	2 Hz

Table 1. DNF parameters used to produce plots in Fig. 3

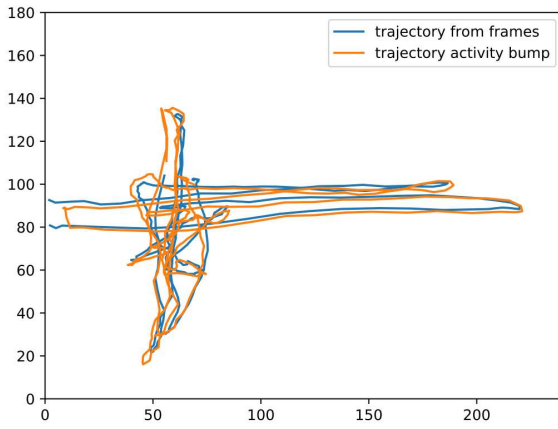


Figure 4. Trajectory of the selected object (star), obtained from the activity bump in the DNF on Loihi and from frames, used as ground truth. The average distance between the points in the two trajectories over all frames is 3.5 DAVIS pixels. The distance was calculated with an offset between frames- and events-trajectory of 15 ms, where the distance is minimal.

Table 2 were used. In particular, both DNF layers are self-connected using a Mexican hat connectivity kernel specified by a difference of two Gaussians (with amplitude and  $\sigma$  of the excitatory and inhibitory kernels); a global inhibitory group of neurons is used to provide global inhibition. Every neuron in the inhibitory group has a given probability to be connected to any neuron in the excitatory layer.

Fig. 5 shows performance of the two-dimensional DNF on the Loihi chip in a tracking experiment. Here, the shapes.translation dataset [20] is used that contains a number of objects drawn on a wall in front of a moving DVS. Plots in Fig. 5b show input that is sent to Loihi over the course of the experiment: the DVS events are emitted at the edges of the objects.

Fig. 5c shows activity of the first, non-selective DNF layer: activity in this layer forms peaks over all objects, stabilizing this activity in moments with reduced motion and

Parameter	value
Voltage threshold	$640 * 2^6$
Voltage threshold global inhibition	$896 * 2^6$
Voltage decay time constant	20 ts
Current decay time constant	20 ts
Input connectivity kernel $\sigma$	1.5
Excitation kernel $\sigma$	2
Inhibition kernel $\sigma$	4
Excitatory weight non-selective	152
Inhibitory weight non-selective	-41
Excitatory weight selective	230
Inhibitory weight selective	-41
Excitatory weight to global inhibition	5
Global inhibitory weight non-selective	-20
Global inhibitory weight selective	-90
Excitatory weight non-selective to selective (1:1 connectivity)	740
Input weight on events	70
Input weight off events	-50
Refractory period	12 ts
Refractory period global inhibition	7 ts
Connection probability to/from global inhibition	0.6
Number of global inhibitory neurons	40

Table 2. DNF parameters used to produce plots in Fig. 5 and 6, if layer is not specified, the parameter applies to both.

weaker DVS output. Fig. 5d shows activity of the output layer of the tracking DNF. Here, one of the objects (the star) is selected (by a local boost to this layer in the beginning of the simulation) and is tracked throughout the experiment, despite presence of the distractors.

To obtain the spike plots, the spikes were filtered with a 50ms rectangular filter and snapshots were taken at regular intervals within the simulation of 6500 time steps. DAVIS input events to the first layer were down-sampled and binned into 1ms per time step (events that exceeded one event per bin, were discarded), i.e. the 6500 time steps correspond to 6.5s of DAVIS input.

Fig. 4 shows the trajectory of the selected object that is extracted from the activity of the output layer of the DNF model and the ground truth extracted from the input frames. The blue trace shows the center of the star object extracted from the DAVIS frames by thresholding (ground truth). The orange trace shows mean (i.e. population vector) of the instantaneous firing rate of neurons in the second (output) layer of the network (i.e. the middle of the tracked bump). Instantaneous firing rates are estimated based on the interspike intervals.

The DVS input to the network was down-sampled to 64x64 neurons, the trajectory was up-sampled to the DAVIS resolution of 240x180. The mean error was calculated as the mean of all distances between the locations of the bump ac-



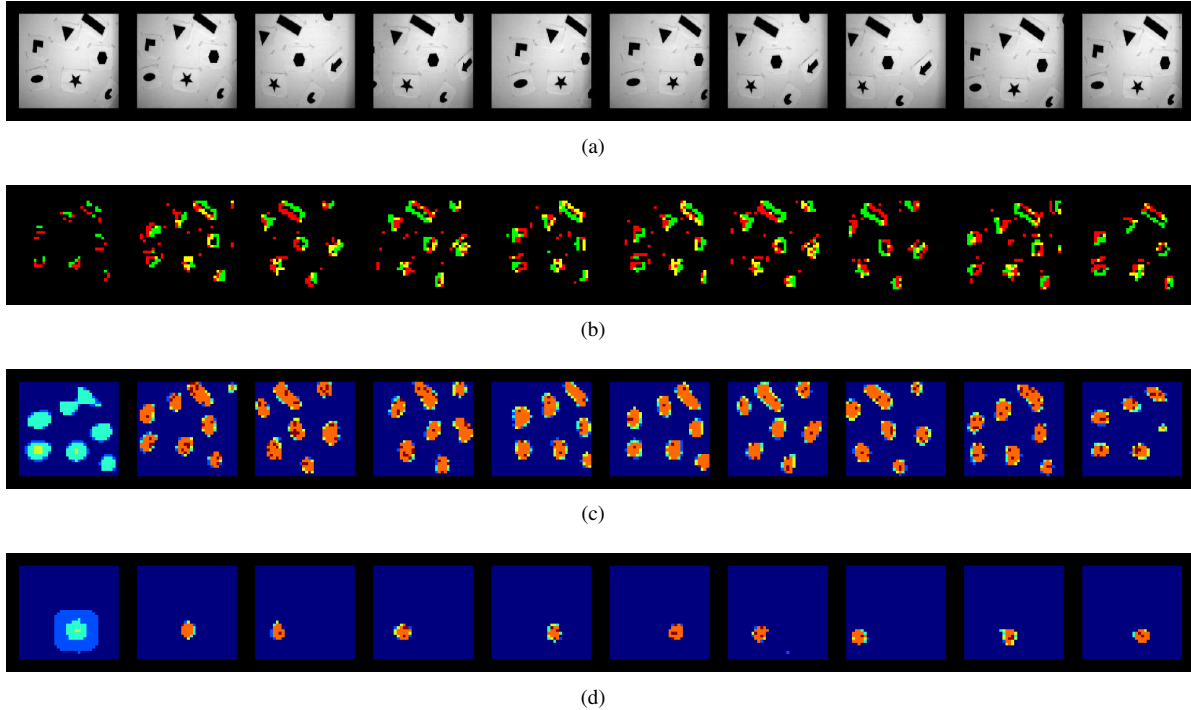


Figure 5. Object tracking experiment: (a) snapshots of input DAVIS frames (top); (b) DAVIS on (green) and off (red) events binned into 10ms frame (second from top); (c) Firing rate of first non-selective WTA layer on Loihi (third from top); and (d) second, selective WTA layer on Loihi (bottom).

tivity and the locations of the frame based extraction at the timesteps of the DAVIS frames and amounts to 3.5 DAVIS pixels.<sup>2</sup>

Fig. 6 shows our second tracking experiment. Here, a ring with five identical objects is rotating in front of the DAVIS camera. The first layer of the tracking SNN architecture creates activity bumps for all five objects, while the second layer (bottom plots) only tracks the single object, selected by a localized activity boost in the beginning of the experiment (first pane in the plot). The same parameters were used here as for Fig. 5. The length of simulation on Loihi was 3000 timesteps here. DAVIS input events were binned into 0.5ms per timestep, i.e. the simulation corresponds to 1.5s of DAVIS input.

## 5. Conclusion

In this work, we have shown for the first time a setup that interfaces an event-based camera DAVIS with the spiking neuromorphic system Loihi, creating a purely event-driven sensing and processing system. We have implemented a simple attention and tracking network on Loihi that allows to select a single object out of a number of moving objects

<sup>2</sup>The Figure was generated using a Brian2 simulation of the equations implemented in Loihi, as it is currently impossible to probe a large network for that many time steps, however, performance of the network can be observed in a live demo.

in the visual field and track this object, even in cases when movement stops and the event stream is interrupted. Full evaluation of the system in terms of tracking speed and quality, as well as power efficiency and robustness is target of our current work and will be reported shortly, while functioning of the system can be observed in a live demonstration during the workshop.

## Acknowledgements

We would like to thank Dr. Julien Martel and the Intel Neuromorphic Computing Lab for their help with the hardware and software setup used in this work. This project has received funding from the SNSF project Ambizione (PZ00P2\_168183) and a ZNZ Fellowship from the Neuroscience Center Zurich.

## References

- [1] X. Lagorce and R. Benosman, “STICK: Spike Time Interval Computational Kernel, a Framework for General Purpose Computation Using Neurons, Precise Timing and Synchrony,” *Neural computation*, vol. 27, pp. 2261–2317, 2015.
- [2] X. Lagorce, G. Orchard, F. Gallupi, B. E. Shi, and R. Benosman, “HOTS: A Hierarchy Of event-based Time-Surfaces for pattern recognition,” *IEEE Transactions on*




Figure 6. Object tracking experiment 2: DAVIS on (green) and off (red) events binned into 10ms frames (top); firing rate of the first non-selective WTA layer on Loihi (middle); output of the second, selective WTA layer on Loihi (bottom).

- Pattern Analysis and Machine Intelligence*, vol. 8828, no. c, pp. 1–1, 2016. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=7508476> 1
- [3] M. Osswald, S.-H. Ieng, R. Benosman, and G. Indiveri, “A spiking neural network model of 3D perception for event-based neuromorphic stereo vision systems,” *Scientific Reports*, vol. 7, no. December 2016, p. 40703, 2017. [Online]. Available: <http://www.nature.com/articles/srep40703> 1
- [4] D. R. Valeiras, G. Orchard, S. H. Ieng, and R. B. Benosman, “Neuromorphic Event-Based 3D Pose Estimation,” vol. 9, no. January, pp. 1–15, 2015. 1
- [5] J. Carneiro, S. H. Ieng, C. Posch, and R. Benosman, “Event-based 3D reconstruction from neuromorphic retinas,” *Neural Networks*, vol. 45, pp. 27–38, 2013. [Online]. Available: <http://dx.doi.org/10.1016/j.neunet.2013.03.006> 1
- [6] R. Benosman, S.-H. Ieng, C. Clercq, C. Bartolozzi, and M. Srinivasan, “Asynchronous frameless event-based optical flow,” *Neural networks : the official journal of the International Neural Network Society*, vol. 27, pp. 32–7, mar 2012. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/22154354> 1
- [7] G. Indiveri and S.-C. Liu, “Memory and information processing in neuromorphic systems,” *Proceedings of the IEEE*, vol. 103, no. 8, pp. 1379–1397, 2015. 1
- [8] M. Davies, N. Srinivasa, T.-h. Lin, G. China, Y. Cao, S. H. Choday, G. Dimou, P. Joshi, N. Imam, S. Jain, Y. Liao, C.-k. Lin, A. Lines, R. Liu, D. Mathaikutty, S. Mccoy, A. Paul, J. Tse, G. Venkataramanan, Y.-h. Weng, A. Wild, Y. Yang, and H. Wang, “Loihi: a Neuromorphic Manycore Processor with On-Chip Learning,” *IEEE Micro*, vol. 38, no. 1, pp. 82–99, 2018. 1, 2
- [9] P. A. Merolla, J. V. Arthur, R. Alvarez-Icaza, A. S. Cassidy, J. Sawada, F. Akopyan, B. L. Jackson, N. Imam, C. Guo, Y. Nakamura, B. Brezzo, I. Vo, S. K. Esser, R. Appuswamy, B. Taba, A. Amir, M. D. Flickner, W. P. Risk, R. Manohar, D. S. Modha, C. Eliasmith, T. C. Stewart, X. Choo, T. Bekolay, T. DeWolf, Y. Tang, D. Rasmussen, M. Mishkin, L. G. Ungerleider, K. A. Macko, B. V. Benjamin, M. Mahowald, R. Douglas, G. Indiveri, R. H. Hahnloser, R. Sarpeshkar, M. A. Mahowald, R. J. Douglas, H. S. Seung, S.-C. Liu, T. Delbruck, J. Backus, R. J. Douglas, K. A. Martin, S. B. Laughlin, T. J. Sejnowski, V. B. Mountcastle, D. H. Hubel, T. N. Wiesel, T. Ohno, T. Hasegawa, T. Tsuruoka, K. Terabe, J. K. Gimzewski, M. Aono, M. M. Shulaker, G. Hills, N. Patil, H. Wei, H. Y. Chen, H. S. Wong, S. Mitra, D. S. Modha, R. Singh, S.-B. Paik, D. L. Ringach, L. A. Palmer, T. L. Davis, E. Painkras, L. A. Plana, J. Garside, S. Temple, F. Galluppi, C. Patterson, D. R. Lester, A. D. Brown, S. B. Furber, P. Merolla, J. Arthur, R. Alvarez, J.-M. Bussat, K. Boahen, A. G. Andreou, and K. A. Boahen, “Artificial brains. A million spiking-neuron integrated circuit with a scalable communication network and interface.” *Science (New York, N.Y.)*, vol. 345, no. 6197, pp. 668–73, 2014. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/25104385> 1
- [10] S. B. Furber, D. R. Lester, L. A. Plana, J. D. Garside, E. Painkras, S. Temple, and A. D. Brown, “Overview of the SpiNNaker System Architecture,” *IEEE Transactions on Computers*, vol. 62, no. 12, pp. 2454–2467, 2012. 1
- [11] G. Indiveri, “Modeling Selective Attention Using a Neuromorphic Analog VLSI Device,” *Neural Computation*, vol. 12, no. 12, 2000. [Online]. Available: <http://dx.doi.org/10.1162/089976600300014755> 1
- [12] S. A. Aamir, Y. Stradmann, P. Muller, C. Pehle, A. Hartel, A. Grubl, J. Schemmel, and K. Meier, “An accelerated LIF neuronal network array for a large-scale mixed-signal neuromorphic architecture,” *IEEE Transactions on Circuits and*

*Systems I: Regular Papers*, vol. 65, no. 12, pp. 4299–4312, 2018. 1

- [13] K. A. Boahen, “Point-to-Point Connectivity Between Neuromorphic Chips using Address-Events,” *Ieee Transactions on Circuits & Systems*, vol. 47, no. 5, pp. 416–434, 1999. [Online]. Available: <https://web.stanford.edu/group/brainsinsilicon/pdf/00{-}journ{-}IEEEts{-}Point.pdf> 1
- [14] C. Brandli, R. Berner, M. Yang, S. C. Liu, and T. Delbruck, “A 240 180 130 dB 3  $\mu$ s latency global shutter spatiotemporal vision sensor,” *IEEE Journal of Solid-State Circuits*, vol. 49, no. 10, pp. 2333–2341, 2014. 1, 2
- [15] G. Schöner and J. P. Spencer, Eds., *Dynamic Thinking: A Primer on Dynamic Field Theory*. Oxford University Press, 2015. [Online]. Available: <https://books.google.ch/books?id=iLVPAAQBAJ{&}printsec=frontcover{&}hl=de{&}source=gbs{-}ge{-}summary{-}r{&}cad=0> 2, 3, 4
- [16] Y. Sandamirskaya, “Dynamic Neural Fields as a Step Towards Cognitive Neuromorphic Architectures,” *Frontiers in Neuroscience*, vol. 7, p. 276, 2013. 2
- [17] J. N. P. Martel and Y. Sandamirskaya, “A neuromorphic approach for tracking using dynamic neural fields on a programmable vision-chip,” in *Proceedings of the 10th International Conference on Distributed Smart Camera (ICDSC)*, ACM, 2016. 2
- [18] T.-H. Lin, G. N. Chinya, M. Davies, C.-K. Lin, A. Wild, and H. Wang, “Mapping spiking neural networks onto a many-core neuromorphic architecture,” *ACM SIGPLAN Notices*, vol. 53, no. 4, pp. 78–89, 2018. 2
- [19] Y. Sandamirskaya, S. K. U. Zibner, S. Schneegans, and G. Schöner, “Using Dynamic Field Theory to extend the embodiment stance toward higher cognition,” *New Ideas in Psychology*, vol. 31, no. 3, pp. 322–339, 2013. 3
- [20] E. Mueggler, H. Rebecq, G. Gallego, T. Delbruck, and D. Scaramuzza, “The event-camera dataset and simulator: Event-based data for pose estimation, visual odometry, and SLAM,” *International Journal of Robotics Research*, vol. 36, no. 2, pp. 142–149, 2017. 5