

Neuromorphic Computing and Sensing in Space

Dario Izzo^{*†}, Alexander Hadjiivanov^{*†}, Domink Dold^{*†}
Gabriele Meoni^{*†‡}, Emmanuel Blazquez^{*†}

1 Introduction

The term “neuromorphic” refers to systems that are closely resembling the architecture and/or the dynamics of biological neural networks [1, 2, 3]. Typical examples would be novel computer chips designed to mimic the architecture of a biological brain, or sensors that get inspiration from, e.g., the visual or olfactory systems in insects and mammals to acquire information about the environment. This approach is not without ambition as it promises to enable engineered devices able to reproduce the level of performance observed in biological organisms – the main immediate advantage being the efficient use of scarce resources, which translates into low power requirements. Nowadays, the neuromorphic approach is mostly investigated at two levels (i) algorithmic and (ii) hardware. On the algorithmic level, it leverages spike-based processing and training [2] to build novel machine learning pipelines able to process data efficiently. At the hardware level, the neuromorphic approach is pursued in designing novel analog and digital circuits and computer chips inspired by biological neural systems. This results in novel sensing devices believed to produce particularly good candidates to emulate biological vision, as well as in the design of computer chips dedicated to efficiently implement the spike-based systems just introduced. In fact, due to the discontinuous nature of spike-based communication and the temporal dynamics of spiking neurons, simulating the behaviour of a whole network of spiking neurons on conventional computer hardware is computationally – and thus energy-wise – very inefficient. This has now also a close relative in the field of artificial intelligence (AI) where Geoffrey Hinton has recently introduced the concept of “Mortal Computation” [4]: a form of computing where no separation between software and hardware exists. In “Mortal Computation”, neural network solutions are uniquely tied to their

^{*}Advanced Concepts Team (ACT), European Space Research & Technology Centre (ESTEC), Keplerlaan 1, 2200 AG Noordwijk (Netherlands)

[†]All authors have contributed equally to this work.

[‡]Φ-lab, European Space Research Institute (ESRIN), Via Galileo Galilei, 1, 00044 Frascati RM (Italy)

underlying (analogue) hardware substrate, which Hinton argues (according to us correctly) might be the only way of obtaining large-scale neural networks that are energy-efficient – a concept that is closely following the neuromorphic paradigm.

The emphasis on low power and energy efficiency of neuromorphic devices is a perfect match for space applications. Spacecraft – especially miniaturized ones – have strict energy constraints as they need to operate in an environment which is scarce with resources and extremely hostile [5]. Numerous works have been investigating different energy-efficient solutions, especially leveraging commercial off-the-shelf (COTS) devices, aiming at optimizing model performance and energy usage trade-offs [5, 6, 7, 8, 9] much less have investigated neuromorphic devices. Early work [10, 11] performed in 2010 at the Advanced Concepts Team (ACT) suggested considering a neuromorphic approach for onboard spacecraft applications. Focusing on optic flow detection [12, 13], these preliminary works showed the possibility to safely land a spacecraft on an unknown planetary surface, assuming a neuromorphic approach to sensing based on the Elementary Motion Detector (EMD) [14], a device inspired by the visual system in flying insects. More recently, with the availability of new neuromorphic sensors such as the Dynamic Vision Sensor [15, 16] and chips such as Loihi [17], TrueNorth [18], Akida and others, the interest on neuromorphic architectures for spacecraft missions grew considerably.

Another significant potential advantage of neuromorphic hardware for space applications concerns radiation. Earth’s atmosphere and magnetic field protects us from a lot of the cosmic radiation, but this poses a considerable problem even in relatively low orbits. While in many instances it can damage the actual hardware, radiation can also interfere with its operation (for instance, by flipping bits in memory), leading to software failure. Neuromorphic hardware can potentially mitigate these issues since, apart from intermittent spikes, it is in fact mostly silent. [19].

In addition to processors, also event-based cameras have contributed to the growing interest in the field of neuromorphic engineering. Event-based vision sensors are well equipped for operation in space: they have a very high dynamic range (on the order of $120dB$), respond only to moving segments in the visual field with very low latency and, most importantly, consume very little power due to their sparse output (on the order of mW). Naturally, these advantages come at a price – for instance, very high noise in very dark environments and proportionally low fidelity for slow-moving objects in the visual field.

In Section 6, we present several research lines that aim to harness these advantages and mitigate the downsides of event-based vision. Before diving into neuromorphic sensing though, we will first give a brief introduction to neuromorphic algorithms (Sections 2 and 3) and neuromorphic hardware platforms (Section 4) to then discuss past and current research mainly con-

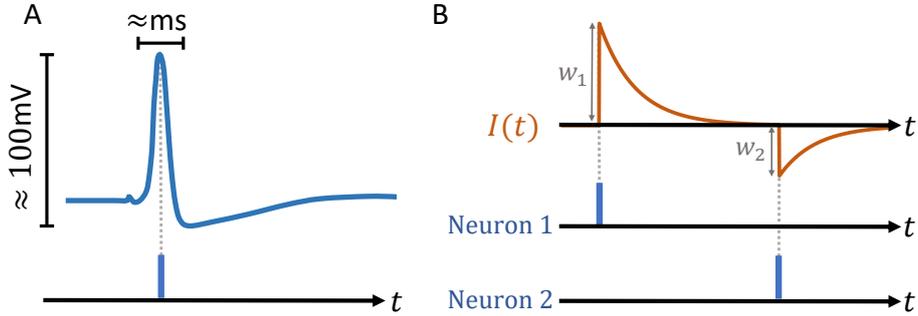


Figure 1: **(A)** Illustration of an action potential or sodium spike [20], an electrical pulse that is sent through the axon, which acts as the ‘output cable’ of a neuron, as a signal to other neurons. For computational purposes, it can be reduced to simply the time of its occurrence (bottom). **(B)** Presynaptic spikes trigger currents $I(t)$ at the postsynaptic neuron, with strength and direction depending on the interaction strength w_i (commonly referred to as synaptic *weight* in NN research).

ducted by the ACT on evaluating the feasibility of a neuromorphic approach for onboard AI applications (Section 5). We hope that this chapter will stimulate further research pursuing a neuromorphic approach to spacecraft onboard computation and sensing.

2 Spiking neural networks

Arguably, the feature which is found most often in modern neuromorphic algorithms is spike-based communication. In the mammalian brain, neurons communicate with each other using action potentials (‘spikes’ or ‘events’) – electrical pulses with a stereotypical shape where only the time at which the spike occurs carries information (Fig. 1A). This realizes a highly sparse, and hence energy efficient, computing paradigm, as neurons only actively change their internal state when excited by an incoming spike, remaining passive otherwise.

Mathematical models of spiking neurons differ substantially in terms of computational complexity and biological realism. A widely used model (both in terms of algorithms and neuromorphic hardware) is the Leaky Integrate-and-Fire (LIF) neuron model, which adequately balances complexity and realism. The LIF model represents a biological neuron as an RC-circuit¹ with capacitance C and resistance R . The capacitor maintains an electric voltage u (the ‘membrane potential’) and is charged by an electric current I on a characteristic time scale $\tau = R \cdot C$ (the ‘membrane time constant’):

$$\tau \frac{d}{dt} u(t) = -u(t) + R \cdot I(t). \quad (1)$$

¹I.e., an electrical circuit with a resistor and capacitor coupled in parallel.

More specifically, $u(t)$ represents the potential difference between the inside and outside of the neural cell caused by different ion concentrations. Via R , ions can ‘leak’ through the membrane until an equilibrium potential is reached (here, the equilibrium potential is set to 0 for simplicity). The current $I(t)$ is caused by the spikes of source (‘presynaptic’) neurons (labelled with indices $i \in \mathbb{N}$ here), which connect to the target (‘postsynaptic’) neuron through synapses with interaction strengths $w_i \in \mathbb{R}$. In the absence of spikes, $I(t)$ decays to 0 and increases or decreases only whenever a presynaptic spike arrives, with strength and polarity depending on the value and sign of w_i (Fig. 1B). The postsynaptic neuron emits a spike at time t_s when the membrane potential crosses a threshold value² ϑ , $u(t_s) = \vartheta$ (Fig. 2A). In biological neurons, the ability to spike is diminished for a short period of time (known as the ‘refractory period’) immediately after spiking, as the mechanism responsible for creating action potentials has to recover first. This is often modelled by clamping the membrane potential to a reset value for the refractory period (Fig. 2A, yellow shaded areas). Without the leak term $-u(t)$, the model is reduced to the Integrate-and-Fire (IF) neuron model.

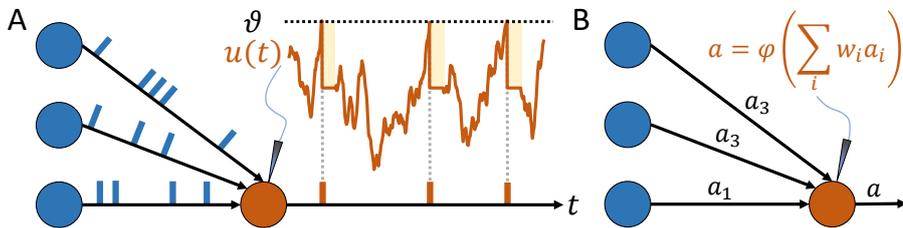


Figure 2: Comparison of SNNs and ANNs. **(A)** In a LIF-based SNN, presynaptic spikes (left, blue) lead to fluctuations in the membrane potential of the postsynaptic neuron (right, orange). When the membrane u reaches a threshold value ϑ , a spike is emitted and the membrane potential is clamped to a reset value for a period of time during which subsequent spiking is impossible (area highlighted in yellow). **(B)** In an ANN, real-valued activations a_i are multiplied by weights w_i and summed up at the postsynaptic neuron. The output a of the postsynaptic neuron is a non-linear function $\varphi(\cdot)$ of this sum.

The realism of the LIF model can be further increased by extending Eq. (1) with additional terms. For instance, the Adaptive Exponential LIF model (AdEx) [22] further adds action potential shapes to $u(t)$ as well as spike-rate adaptation, where prolonged tonic spiking (i.e., permanent activity) of a neuron is preceded by a period of increased or decreased activity. The temporal dynamics of AdEx neurons is capable of replicating a broad

²This threshold is not physically manifested in biological neurons, but it is a reasonably good approximation for modelling their response behaviour [21].

range of dynamics observed in biological neurons. In particular, it has been found that adaptation is employed by neurons for solving temporal tasks such as sequence prediction (see, e.g., [23, 24]).

In principle, every architecture that exists for ANNs can be converted into a spiking neural network (SNN) by replacing artificial neurons with spiking neurons, e.g., LIF neurons. However, ANNs and SNNs are different in two key aspects: (i) as discussed above, spiking neurons are characterised by a dynamic internal state, while artificial neurons possess no intrinsic state or dynamics, and therefore have no intrinsic ‘awareness’ of time; and (ii) spiking neurons interact by triggering currents at specific times through spikes, while artificial neurons communicate directly with each other via real-valued activations (Fig. 2B). Even though the event-based nature of SNNs makes them potentially more energy efficient than ANNs, it is still an open question whether it also provides other advantages (e.g., in terms of performance, robustness to noise or training time). A key challenge hereby is identifying how information can be efficiently encoded in the temporal domain of spikes [25]. A selection of proposed spike-based encoding schemes addressing this question is summarized in Table 1 and illustrated in Fig. 3.

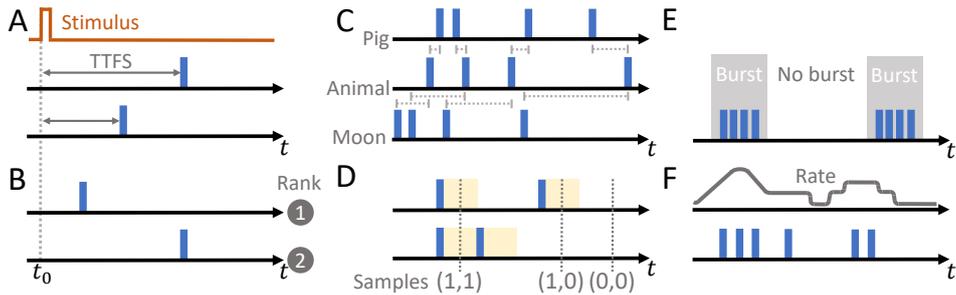


Figure 3: An illustration of spike-based encoding schemes. **(A)** TTFS encoding. **(B)** Rank order encoding with a population of two neurons. **(C)** Example of how concepts can be encoded and decoded (via similarity) in spike trains. Here, the spike train representing ‘Pig’ and ‘Animal’ are similar, while ‘Moon’ and ‘Animal’ are dissimilar. For clarity, the differences between spike times are shown in gray. **(D)** Encoding as random samples. At every point in time, a sample can be read out from the network, with neurons being in state ‘1’ while refractory (yellow area) and in state ‘0’ otherwise. **(E)** Encoding with bursts, short and intense periods of spike activity. **(F)** Rate-based encoding.

In general, for machine learning applications it would be preferred to have a learning algorithm that automatically finds the optimal (combination of) encoding schemes. Although training SNNs has been a daunting task for a long time, recent progress in terms of both theory and software infrastructure has enabled exactly such end-to-end learning for SNNs, opening novel opportunities for building highly efficient and powerful spike-based AI systems.

Encoding scheme	Description
Time-To-First-Spike (TTFS) [26]	Given a stimulus onset at time t_0 , information is encoded in the time required for <i>any</i> neuron in a population of neurons to emit its first spike (Fig. 3A). This encoding scheme allows fast processing (i.e. low latency) and is highly energy efficient since each involved neuron spikes only once at the most in order to solve a given task.
Rank order [26]	Information is encoded by the order in which spikes occur in a population of neurons given the onset of a stimulus (Fig. 3B). Thus, in contrast to TTFS, the exact spike time becomes irrelevant, but information can still be processed using only a low number of spikes depending on how many of the earliest spikes are used for encoding the input.
Spike patterns	Instead of only using single spikes to represent information, whole spike patterns of individual neurons or populations can be used [27, 28]. For instance, in [29], symbolic information (e.g., abstract concepts) is represented by spike trains, and relatedness between concepts is encoded in the similarity of spike trains (Fig. 3C).
Sampling	In neural sampling, neurons represent sampled values of binary random variables (refractory after spiking = 1, non-refractory = 0), allowing them to encode probability distributions (Fig. 3D). This is suitable for approximating, e.g., sampling-based deep learning architectures such as restricted Boltzmann machines [30, 31].
Bursts	Information can be encoded in bursts, i.e., short periods of high spike activity (Fig. 3E). For instance, in [32], bursts have been used to propagate information through SNNs that guides learning.
Rates	An instantaneous spike rate obtained by averaging over spiking activity (either over time or populations) carries the information (Fig. 3F). Since the spiking rate is a continuous variable, it encodes information in a similar way to ANNs.

Table 1: Different ways of encoding information using spikes.

3 Learning algorithms for SNNs

One of the main enablers of the incredible success of deep learning in recent years is the error backpropagation algorithm (‘backprop’). However, for a long time, this success did not pass over to SNNs – mostly due to the threshold mechanism of spiking neurons that leads to vanishing gradients at all times except at the time of threshold crossing [33] (cf. Fig. 4 for an illustration of the concept of surrogate gradients, which present one solution to this problem). Instead, to apply SNNs to a variety of machine learning applications, it was customary to convert the parameters of a trained ANN to a SNN [34]. SNNs ‘trained’ this way typically show a significant degradation in performance compared to the original ANN. In addition, the mapping

promotes purely rate-based encoding in the SNNm which has been found to reduce the energy efficiency of SNNs [35, 36]. An alternative approach is to use biologically inspired learning rules, such as Spike-Time-Dependent Plasticity (STDP) [37] or variants thereof, although they do not scale well beyond shallow networks³.

Recently, several approaches have been found for successfully unifying SNNs and backprop, which have become the *de facto* state of the art for training SNNs on machine learning tasks and are currently being investigated at the ACT. For instance, one approach is based on SNNs using neuron models where the TTFS is both analytically calculable and differentiable, allowing gradients to be calculated exactly without having to deal with the discrete nature of the threshold mechanism [39, 40, 41, 42], although this method is limited to neurons that only spike once.

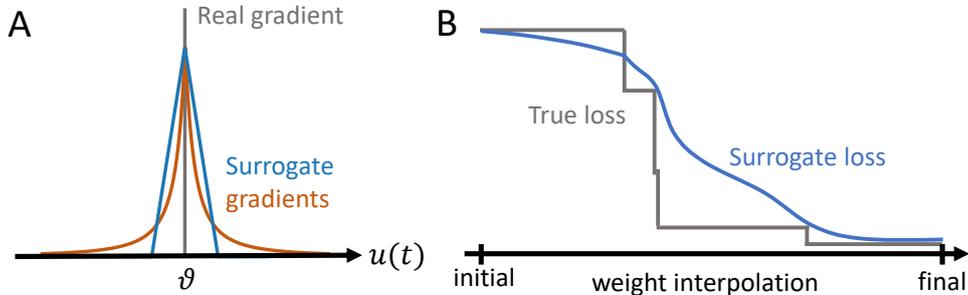


Figure 4: Illustration of the surrogate gradient method. **(A)** The gradient of the threshold function is a Dirac delta peak, which is infinite at $u = \theta$ and vanishes otherwise (gray). In the surrogate gradient approach, this gradient is replaced by a well-behaved function, such as a mirrored Lorentz function (orange) or a triangular function (blue). **(B)** In a SNN, the landscape of the loss function has discontinuities due to the spiking mechanism (gray). Especially long plateaus with constant loss inhibit learning. Using surrogate gradients ‘smoothens’ the loss landscape (blue trace). Here, the loss landscape is illustrated by interpolating between two sets of weights of a SNN. Reproduced based on [33].

At the time of writing, the most widely used approach is the so-called “Surrogate Gradient” method, which is applicable to all kinds of spiking neuron models [33, 43]. Here, the gradient of the threshold function is replaced with a surrogate function that has non-zero parts for membrane potential values away from the threshold (Fig. 4A). This ‘softening’ of the spike threshold gradient function allows gradient information to flow continuously through the network (Fig. 4B), enabling end-to-end training of SNNs capable of utilizing the temporal domain of spikes to find, e.g., highly sparse and energy efficient solutions. For instance, in [44] SNNs are trained

³As of late, several biologically plausible learning rules that are applicable to deep neural architectures have been proposed [38]. However, they lack the flexibility and theoretical guarantees of backprop, i.e., they are not proven to minimize a task-dependent (and customizable) loss function. Nevertheless, they present intriguing alternatives for training SNNs.

that reach a competitive classification accuracy ($\approx 1.7\%$ test error) on the MNIST handwritten digits dataset [45] with, on average, only 10 to 20 spikes per inference. Surrogate gradients can be used to optimise not only the weights but all parameters that influence the dynamic behaviour of spiking neurons, such as time constants and spike thresholds of individual neurons – with potential benefits for the robustness and expressiveness of SNNs [24, 46]. A major downside of the surrogate gradient method is that currently there is no theoretical framework for choosing the shape of the surrogate gradient function. Nevertheless, initial evidence suggests that the approach is relatively robust with respect to this choice [44]. In the future, this downside could be alleviated through exact derivations of gradient-based learning rules for SNNs (see, e.g., [47]). Additional information on the current standards for training SNNs can be found in [33, 48].

To summarize, the surrogate gradient method enables end-to-end training of SNNs using error backpropagation. Several open-source packages are available that standardize and ease the implementation and training of SNNs by utilising existing libraries for automatic differentiation. Currently, four general SNN libraries based on pyTorch are being developed: Norse [49], spikingjelly [50], snnTorch [48] and BindsNET [51]. The first three adopt both the workflow and class structure of pyTorch, effectively extending it to support SNNs. The last one is geared towards developing machine learning algorithms that take inspiration from biology. Even though it does not support gradient-based learning, it contains a larger variety of neuron models and biologically inspired learning rules than the other three packages. In general, the emerging landscape of SNN libraries greatly reduces the development time of SNN-based algorithms, painting a promising picture for further exploration of the capabilities and potential benefits of SNNs. In addition, it increases the accessibility of SNNs to researchers outside the field of neuromorphic computing and computational neuroscience by making it possible to seamlessly exchange or interweave ANNs and SNNs in a common framework. Thus, in the coming years we can expect to see an increasing number of contributions from the aerospace field attempting to incorporate spike-based algorithms and hardware onboard spacecraft.

4 Existing hardware

Although end-to-end training of SNNs is possible nowadays, exploring their capabilities remains challenging due to the computationally demanding nature of simulating the internal neuron dynamics of spiking neurons. In contrast to artificial neurons, the dynamics of spiking neurons have to be solved using numerical solvers for ordinary differential equations, introducing a significant overhead. Thus, on conventional hardware systems, the benefit of spike-based coding (in the absence of spikes, no active computations are

performed) is not immediately apparent.

An emerging technology that is capable of harnessing the potential of spike-based information processing is neuromorphic hardware, which explores novel computing architectures and paradigms that closely emulate how the brain processes information. Standard computer architectures separate processing units and memory (storing data and instructions for the processor), introducing a bottleneck due to the constant flow of data between memory and processor (the so-called ‘von Neumann bottleneck’). Instead, neuromorphic hardware follows several design philosophies that can be found in the mammalian brain.

- No separation between processing and memory. Thus, a processing unit (e.g., a neuron) can only access information that is locally available (e.g., synaptic weights or activity of other neurons that it is connected to);
- Large-scale parallel computing;
- Asynchronous instead of clocked computations;
- Event-based information processing (i.e., using spikes instead of continuous values);
- (Re)programming the chip consists of (re)mapping networks or adapting network parameters through learning;
- Time-continuous and locally constrained learning rules.

Broadly speaking, neuromorphic hardware comes in two flavours: digital and analogue. Both designs are transistor-based (i.e., CMOS), but the transistors are used in different operating regimes. In simple terms, digital means calculating with discrete-valued bits (i.e., transistors take the states ‘on’ and ‘off’), while analogue means calculating with continuous-valued currents and voltages. Digital chips provide stable solutions that are closer to commonly used hardware chips. Analogue chips have to deal with noise (e.g., manufacturing noise of components, cross-talk between electrical components, temperature dependence) and have a longer development cycle, but provide intriguing advantages such as high energy efficiency and potentially accelerated (compared to biology) emulation of neuron and network dynamics. In addition to CMOS, memristor-based circuits⁴ are being developed that open novel opportunities for energy-efficient and adaptive neuromorphic devices.

Examples of digital neuromorphic systems are Loihi (1 and 2) (developed by Intel) [17], TrueNorth (developed by IBM)[18], Akida (developed by

⁴A memristor is a fundamental electrical element that displays a hysteretic resistance profile where the resistance depends on the recent history of current that has passed through the element.

Brainchip), SpiNNaker(1 and 2) [52, 53], Darwin [54] and Tianjic [55]. Examples of analogue chips are BrainScaleS (1 and 2) [56] and Spikey [57], DYNAPs [58], ROLLS [59] and Neurogrid [60]. A thorough review of state-of-the-art neuromorphic platforms at the time of writing can be found in [1]. Two of these platforms (Loihi by Intel and BrainScaleS-2 by the University of Heidelberg) are discussed in more detail below to further illustrate the difference between digital and analogue platforms.

First, Loihi’s (Fig. 5A) main components are neuromorphic cores, which are processing units with custom circuitry and small amount of memory necessary to simulate a population of LIF neurons and their plasticity (i.e., learning rules). Spikes are exchanged between neuromorphic cores using a routing grid. In contrast, the BrainScaleS-2 chip (Fig. 5B) realizes physical AdEx neurons and synapses to emulate biology – in other words, neurons and synapses are not *simulated* but rather *implemented* directly as analogue circuits. Hence, no simulation time step exists and the system evolves continuously. Although neurons and synapses are implemented using analogue circuits, spikes are transmitted digitally (also known as “mixed-signal”). Through a synaptic crossbar, connections can be set flexibly and are not pre-wired.

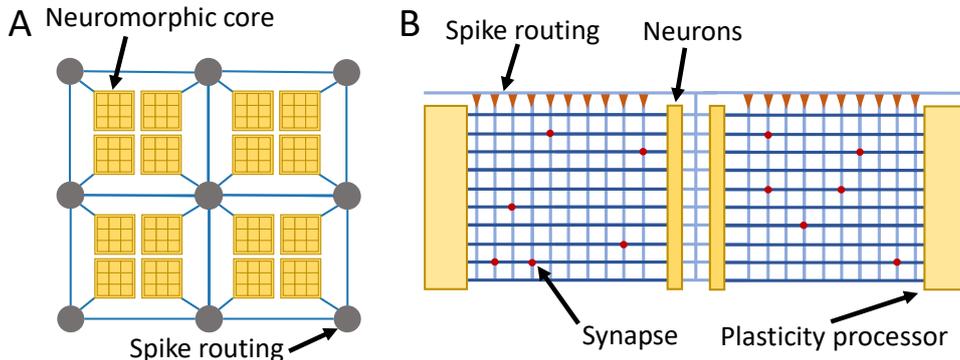


Figure 5: (A) Simplified schematic of the Loihi chip [61]. (B) Simplified schematic of the BrainScaleS-2 chip. Spikes emitted from the neurons enter the synaptic crossbar through synapse drivers (orange triangles) and, in essence, arrive at the target neuron when a connection is set (red dot). Each horizontal line in the crossbar is an input line for a neuron [56].

The main potential of neuromorphic hardware lies in enabling the deployment of energy-efficient low-latency AI systems suitable for edge applications. First glimpses into this potential can already be obtained nowadays: in Table 2, we list a few neuromorphic platforms that provide both performance and energy benchmarks for MNIST classification, with most systems reaching an energy footprint on the order of μJ per inference. We chose MNIST here as it is the most commonly used benchmark and allows for easy comparison across different (non-)neuromorphic platforms. However,

it has been noted that SNNs perform best on data with temporal structure, meaning that MNIST is less suited as a benchmark for showing the advantages of SNNs as well as neuromorphic hardware [44, 62]. In fact, we argue that in order to better assess the potential of neuromorphic devices, standardized and application-specific benchmarks are required to fairly evaluate both algorithms and hardware platforms.

A recent feature of neuromorphic chips that is especially interesting for space applications is on-chip learning, found in chips like Loihi, BrainScaleS-2 and SPOON [63]. On-chip learning enables continuous learning directly integrated in the neuromorphic chip with a low energy footprint, suitable for the distributed design philosophy of neuromorphic devices. In the future, such ‘intelligent’ neuromorphic devices could be especially useful for retraining or fine-tuning onboard models during deep space missions to autonomously adjust and adapt to previously unknown or unexpected circumstances.

Platform	Type	Tech	Model	MNIST	Energy / Image
Unnamed (Intel) [64]	digital	10nm	LIF	97.70%	1.7 μ J
Intel Loihi [17]	digital	14nm	LIF	96.40%	n.a.* [65]
				94.70%	2.47mJ [66]
IBM TrueNorth [18]	digital	28nm	LIF ¹	92.70%	0.268 μ J
				99.42%	108 μ J
Brainchip Akida [67]	digital	28nm	IF	99.20%	n.a.*, ²
SpiNNaker-2 [53]	digital	28nm	– ³	96.60%	23 μ J [68]
SPOON [63]	digital	28nm	– ⁴	97.50%	0.3 μ J ⁵
BrainScales-2 [56]	analogue	65nm	AdEx	96.90%	8.4 μ J [42]
SpiNNaker [52]	digital	130nm	– ³	95.00%	3.3mJ [69]

* Not available.

¹ [18] states that TrueNorth implements “(...) a dual stochastic and deterministic neuron based on an augmented integrate-and-fire (IF) neuron model [70]”.

² https://doc.brainchipinc.com/zoo_performances.html. Accessed: 2022-11-07.

³ SpiNNaker is an ARM-based processor platform and can therefore support arbitrary neuron models as long as a software implementation for SpiNNaker is available.

⁴ SPOON is an “event-driven convolutional neural network (eCNN) for adaptive edge computing” where “TTFS encoding” is used “in the convolutional layers” [63].

⁵ Pre-silicon simulation result.

Table 2: A list of neuromorphic platforms benchmarked on MNIST. As a baseline, [42] provides values for a convolutional ANN implemented on a nVidia Tesla P100, reaching 99.2% accuracy with an average energy / image of 852 μ J. It should be noted that all neuromorphic platforms here process images sequentially, while the GPU uses batching. For sequential processing, [66] reports an accuracy of 98.90% with an average energy / image of 37mJ (Intel i7 8700) and 16mJ (nVidia RTX 5000). Of the results shown above, only TrueNorth, Akida and SPOON implement a convolutional SNN. Table adapted from [42].

5 Application to onboard processing for Earth observation

As previously mentioned, the potential advantages of SNNs (e.g., in terms of energy efficiency) are promising for data processing, especially onboard miniaturized satellites. However, when applied to static data, as is the case in the classification of Earth observation images, it is not clear how SNNs compare to ANNs in terms of energy, latency and performance trade-offs [36]. For instance, some previous works [71, 72] showcase that when rate coding is used, the advantages of SNNs in terms of energy/performance decrease for classification datasets which contain complex features. This makes satellite data a challenging task for SNNs because of the complexity of their features [36]. For that reason, ESA’s ACT, in collaboration with ESA’s Φ -lab, is currently benchmarking SNNs and ANNs for onboard scene classification. The latter was chosen as a target application due to the abundance of benchmark datasets such as EuroSAT [73, 74] which was also the dataset chosen by the ACT for this project. EuroSAT is a 10-classes land-user and land-cover classification dataset consisting of images captured by the Sentinel-2 satellite. For the study, the RGB version of the images was used. We are mostly investigating different information encoding solutions (including rate encoding and spike-time encoding) and their impact on the energy and accuracy that could be expected in performing onboard inference using an end-to-end neuromorphic approach. The results of past and ongoing efforts are presented in Sections 5.0.1 and 5.1, respectively.

5.0.1 Rate-based SNNs for onboard Earth observation

Rate-based models were explored in a previous work by the ACT [36], where the SNN model was trained by exploiting the approximate equivalence between spike rates of IF neurons and ReLU activations of artificial neurons, fundamentally using a weight conversion technique as discussed in Section 3 with limited loss in accuracy [36]. The methodology is summarized in Fig. 6.

As a first step, we started from a VGG16 ANN model [75] pretrained on ILSVRC-2012 (ImageNet) [76]. We replaced the last three dense layers with a cascade of average pooling and dense layers. Furthermore, we replaced the max pooling with average pooling layers that do not require lateral connections and offer an easier implementation for spiking models [36]. The VGG16 network was then trained on the EuroSAT RGB dataset. To convert the trained ANN model, we replaced the ReLU activation function with a cascade of IF neurons and post-synaptic filters, whose dynamic behaviour is described in [36]

$$y(t) := \left(1 - e^{-\Delta t/\tau}\right) \cdot x(t) + e^{-\Delta t/\tau} \cdot x(t-1), \quad (2)$$

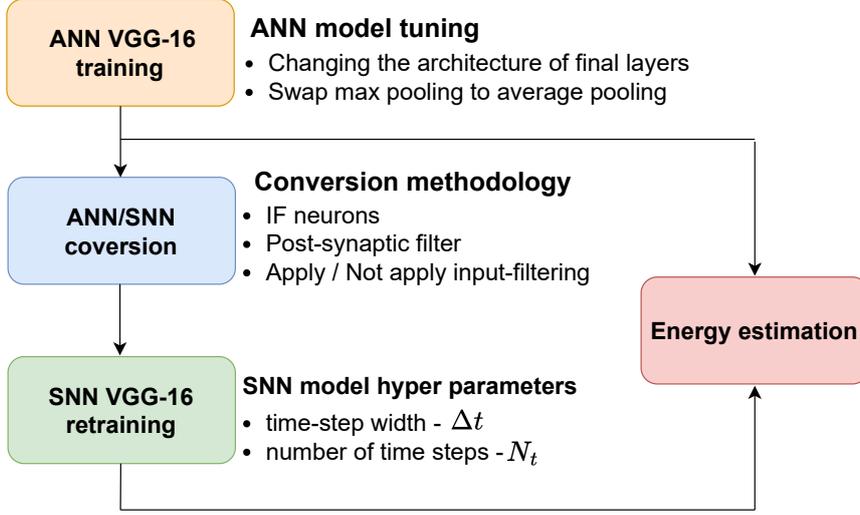


Figure 6: Rate-based SNN vs ANN performance/energy benchmark methodology used in our previous work [36].

where $x(t)$ is the output spike train of a neuron at time t , $y(t)$ is the output of the postsynaptic filter, and Δt is the timestep width. Finally, the converted SNN model is retrained to optimize the performance. To estimate and compare the energy consumption of the SNN and ANN models, we exploited the methodology implemented in *KerasSpiking*. In particular, more than performing an accurate estimation of the energy consumption, the methodology used aimed at allowing a relative comparison [36]. By considering a single spiking layer L , the used methodology assumes that the energy consumption during the inference is mostly due to two contributions: energy dissipated due to the synaptic activity E_s and energy usage due to neuron updates E_n . E_s can be calculated as

$$E_s = E_o \cdot \sum_N S_N \cdot f_{\text{in}} \cdot N_t \cdot \Delta t, \quad (3)$$

where E_o is the energy for a single synaptic operation and depends on the hardware used, S_N is the number of synapses for every neuron, f_{in} is the average neuron spiking rate, N_t is the number of timesteps in the simulation and Δt is the timestep width. For artificial models, $f_{\text{in}} = \frac{1}{\Delta t}$ and $N_t = 1$. Furthermore, E_n can be calculated as

$$E_n = E_u \cdot N_n \cdot N_t, \quad (4)$$

where E_u is the energy for a neuron update and depends on the hardware (similarly to E_o), and N_n is the number of neurons in a layer. The total energy consumption for a model is then given by the sum of the contributions of the different layers. Since for the SNN models, the energy depends

Model	Accuracy [%]	T	Δt [s]	Energy [J] on GPU	Energy [J] on Loihi
ANN	95.07	1	-	0.06996	0.00636
ANN + Prewitt	90.19	1	-	0.06996	0.00636
SNN (Best accuracy)	85.11	4	0.0381	-	0.00444
SNN + Prewitt (Lowest energy)	87.89	4	0.0403	-	0.00205
SNN + Prewitt (Highest accuracy)	85.07	1	0.0813	-	0.00476

Table 3: Summary of the results shown in Table 2 of [36].

on N_t and Δt , different SNN models were trained to explore various energy/accuracy trade-offs. To that end, we trained different models by setting the number of timesteps N_t ranging from 1 to 32 and making Δt a learnable parameter. The constant τ of every neuron post-synaptic filter was also trained to optimize the dynamic response of each neuron. In addition, since E_s is proportional to f_{in} . Therefore, in order to reduce the average rate for each neuron, the input was processed with a Prewitt filter, whose dynamics is given by

$$\mathbf{X}' := \max \left(c \sqrt{(\mathbf{G} \circledast \mathbf{X})^2 + (\mathbf{G}^\top \circledast \mathbf{X})^2}, 0.0078 \cdot \mathcal{K} \right), \quad (5)$$

where

$$\mathbf{G} := \begin{pmatrix} 1 & 1 & 1 \end{pmatrix}^\top \begin{pmatrix} 1 & 0 & -1 \end{pmatrix},$$

and \circledast represents the two-dimensional convolution operator, \mathcal{K} is a unitary tensor having the same shape as the input \mathbf{X} , and c is a normalization constant. The effect of the Prewitt filter is to concentrate the input spikes mostly around the boundaries of different crops while keeping the input color information. Training was performed with and without applying the Prewitt filter to test its effects. Since the energy depends on the hardware used through the constants E_o and E_u , to estimate the energy consumption we tested various models on different hardware, including on desktop CPU (Intel i7-4960X), one embedded processor (ARM Cortex-A), one GPU (nVidia GTX Titan Black) for ANN models and SpiNNaker [52], SpiNNaker 2 [53], and Loihi [17] neuromorphic processors for both the ANN and SNN networks. The test results are summarized in Table 3.

The maximum accuracy of the SNN model was 85.11% and 87.89% with and without Prewitt filtering, respectively [36]. The ANN models reached

95.07% and 90.19% with and without Prewitt filter, respectively. In addition, the input filter had a strong effect on the tradeoff between energy consumption and accuracy for SNN models, whereby the energy consumption dropped by half while maintaining similar accuracy or the accuracy increased by over 2% with a similar energy consumption. The drop in accuracy of the SNN + Prewitt (highest accuracy) and SNN + Prewitt (lowest energy) was 1.18% and 10%, respectively, which was significant. However, when compared to the GPU implementation of the ANN, inference with SNN + Prewitt (highest accuracy) and SNN + Prewitt (lowest energy) on Loihi required 14.70 and 34.13 times less energy, respectively. Since this improvement was partially because of the higher efficiency of Loihi compared to the GPU, we also tested the inference performance of the ANN model on Loihi to eliminate the effects of the hardware. Even in this case, the energy consumption was 3.1 and 1.3 times lower compared to the SNN + Prewitt (lowest energy) and SNN + Prewitt (highest accuracy) models. However, to obtain more realistic estimates, future work will aim at more hardware-aware proxies that also take into account other effects, such as on-chip saturation effects and power-consumption due to off-chip accesses, which might affect the estimates presented here. In addition, the results seem to confirm trends described in [71, 72], according to which longer simulation leads to better accuracy at the expense of energy consumption when rate coding is used. Because of that, as described in Section 5.1, different information encoding solutions shall be explored.

5.1 Spike-time-based SNNs for onboard Earth observation

Additional encoding schemes, such as temporal encoding, and surrogate gradient training approaches are currently being explored that could yield better trade-offs for Earth observation. By encoding information in spike times instead of rates, information can be processed with a much lower number of spikes, leading to higher sparsity in the network activity and consequently lower energy consumption. We are focusing on two approaches: (i) training SNNs that solely use TTFS encoding, either trained with exact gradients or surrogate gradients, and (ii) training SNNs end-to-end using surrogate gradients without imposing any specific encoding scheme. Our evaluation of models is two-fold: performance and network metrics (for instance, the average number of spikes per inference) are investigated in simulation using SNN frameworks such as Norse. In addition, we are mapping parts of these models to the neuromorphic chips Loihi and BrainScaleS to obtain an estimate for the energy efficiency of a fully neuromorphic implementation. As this is still an ongoing study, only a few preliminary results will be discussed here.

Performance-wise, we found that convolutional architectures trained end-to-end show the highest accuracy. We implemented a spike-based convolu-

tional neural network (CNN) with 4 layers (3 convolutional, 1 dense layer) that takes the pixel values of images (rescaled to be between 0 and 1) as input currents. In this architecture, the output layer consists of LIF neurons without a threshold mechanism, i.e., without spiking. All other layers consist of LIF neurons. A classification result is given by the index of the neuron with the maximum membrane potential value during the whole simulation time – in simpler terms, the label neuron that obtained most votes from the last hidden layer provides the classification outcome. Trained end-to-end, this model achieves competitive results ($\approx 91 - 92\%$ test accuracy on EuroSAT RGB) with low spike activity (on average $\approx 2 - 3$ spikes per neuron during inference).

Moreover, we are investigating SNNs based on TTFS encoding for all layers (input, hidden and output). Thus, in the label layer, the classification outcome is given by the label neuron that spikes first. We are currently exploring both models that can be trained using exact gradients and ones that can be trained using surrogate gradients. Preliminary results show that these models reach slightly lower accuracy, although they also require fewer spikes per inference than models without pre-imposed encoding. In general, these initial results already hint at a confirmation that direct training of SNNs using surrogate gradients leads to better results than (rate-based) weight conversion for Earth observation image processing tasks.

6 Neuromorphic sensing

In addition to *processing*, neuromorphic computing has also entered the domain of *perception*, most notably in the area of computer vision. Biological systems provide a rich source of inspiration due to their energy efficiency, strong data compression and feature extraction properties as well as remarkable adaptability. A widely used example of biological computations is the visual system in flying insects, which has been deeply researched in the last decades (e.g. [77, 78]). Research on the neuronal systems in insects allows for single-neuron analysis and thus provides very detailed insights on the internal working mechanisms of insect vision. The primary source of motion information in insects is the so-called EMD - Elementary Motion Detector (e.g. [14]). EMDs are a minimalistic neural architecture able to extract important motion information from only two photoreceptors. Preliminary studies performed at the European Space Agency on their applications to spacecraft landing [10, 11, 12, 13] have suggested the possibility to use EMDs as the only needed sensor to successfully land spacecraft in specific scenarios. A second example of highly efficient and adaptive biological sensor is the mammalian retina, which is composed of several layers of neurons specialised in processing visual information. The mammalian visual system is important to understand the biological inspiration behind the concept

of event-based vision, which we later focus on, thus the following section provides a brief overview of its structure and function.

6.1 Structural and functional organisation of the retina

The structural and functional organisation of the mammalian retina has been explored in great detail over the past few decades. The elaborate 1965 study by Hubel and Wiesel [79] was the first to present a collective summary of many fundamental properties of the retina, such as the hierarchy of retinal neurons, their receptive fields (RFs) and orientation selectivity, and so forth. The mammalian retina consists of several layers containing nerve cells and receptors that perform highly specialised functions [80]. The following is a brief overview of the main types of cells found in the mammalian retina, together with their main distinctive features.

Receptors The outer nuclear layer (ONL) contains the actual photo-sensitive receptors: rods, which are responsible for scotopic vision (low-light conditions), and cones, which are responsible for photopic vision (in well-lit conditions). Furthermore, rods function in low-light conditions but do not distinguish colour, whereas cones are active in display colour specialisation (e.g., responding to light in the red, green or blue part of the visible spectrum). Receptors are excitatory cells.

Horizontal cells The signal produced by receptors in the ONL is modulated by horizontal cells, which provide local inhibitory signal informed by the immediate neighbourhood of the receptor. Horizontal cells are responsible for local brightness adaptation (especially in the foveal region) as well as the formation of ON/OFF RFs in ganglion cells [81]).

Bipolar cells The raw visual input normalised through feedback from horizontal cells is then fed into bipolar cells, which effectively split the visual channel into two pathways (ON and OFF) [81]. This means that ON bipolar cells respond to an increase in brightness, whereas OFF bipolar cells respond to a *decrease* in brightness. The separate pathways are preserved all the way to the visual cortex and have been shown to be asymmetric in terms of their behaviour – for instance, the OFF pathway reacts faster and sends stronger signals to the cortex compared to the ON pathway [82]. In addition, while horizontal cells provide the retina with spatial brightness adaptation capabilities, bipolar cells are responsible for temporal adaptation and filtering [83, 84].

Amacrine cells The excitatory signals produced by bipolar cells are combined with inhibitory feedback from amacrine cells in many different configurations. There are more than 30 types of amacrine cells, whose primary

function is to combine the signals arriving from bipolar cells in many different configurations that define, modulate and control the RFs of ganglion cells via inhibitory synapses [84] (however, there is a possibility that they also provide excitatory signals via gap junctions [85]).

Ganglion cells Amacrine and bipolar cells form different wiring patterns providing input to the ganglion cell layer (GCL). Different ganglion cells have different RFs, which determine their *selectivity* to certain patterns (both static and dynamic, such as orientation, motion or combinations of both) in the visual field [81]. For instance, some ganglion cells respond to edges, others are highly sensitive to motion relative to the background in the visual field, and yet others respond to approaching motion while remaining silent in response to lateral motion [84]. While ganglion cells are officially considered to be the first neurons in the retina that produce spikes [86] (the other cells have graded responses), there are indications that at least bipolar cells also produce a limited number of spikes [87] that are phase-locked to visual stimuli.

Counterintuitively, the ONL is the *innermost* layer of the retina, which means that light has to travel through all the other layers before it reaches the receptors. This and other structural properties of cellular layers in the retina have several functional implications.

- The GCL contains a lot of cell bodies as well as axons which relay spikes via the optic nerve to the visual cortex. Since the GCL is on top of the receptor layer, the part of the retina where ganglion cell axons interface with optic nerve is devoid of receptors, giving rise to the infamous *blind spot*.
- A small area of the retina known as the *fovea*, which is located at the centre of the retina directly behind the lens, contains the highest concentration of receptors. To maximise the resolution in the fovea, the axons of ganglion cells in the fovea are highly stretched out, forming a dense circular bunch of axon bodies around the fovea, where the resolution starts to drop sharply.
- At least in humans (and many other mammals), most of the colour-sensitive receptors are in the foveal region. Away from the fovea, vision is dominated by rods, which are much more sensitive (down to a single photon [88]) but not colour-specific. This is easily observed in low-light conditions, where our visual perception shifts towards detect only silhouettes).
- The RFs of ganglion cells responsible for encoding signals from the fovea are connected to as few as a single receptor, while the RF size increases further away from the fovea. This eccentricity-dependent

RF size has been found to have a number of interesting consequences for scene recognition, including cortical magnification, scale invariance and object recognition [89, 90].

6.2 Adaptation and homeostasis in the retina

Adaptation is a hallmark of most types of neurons and circuits found in the retina, a lot of research effort has been dedicated to unveiling the underlying principles of visual adaptation [91, 92]. A number of different homeostatic processes have been observed in practically all layers in the retina. For instance, cones have a very large dynamic range and are also extremely sensitive to small changes in contrast, both in terms of response time and contrast range. In comparison, rods can detect very low levels of illumination (as low as individual quanta of light [93]), with the trade-off that they adapt very slowly and become saturated by large sudden changes in light intensity [93]. ON-type (OFF-type) bipolar cells have a light center/dark surround (dark center/light surround) and respond strongly to positive (negative) contrast. The ON/OFF differentiation is maintained in retinal ganglion cells (RGCs), which receive input from one of the two types of BCs. The RFs of both ON and OFF RGCs independently cover almost the entire retina [94]. The retina uses adaptation mechanisms in BCs and RGCs to adapt to both spatial and temporal variations in illumination [95, 96]. Both BCs and RGCs homeostatically regulate the intensity of their output depending on the magnitude of the stimulus, allowing the retina to adapt to changes in luminance that span 9 to 12 orders of magnitude over a 24-hour period [97, 98].

Importantly, there is overwhelming evidence that RGCs adapt their activity to both the mean and the variance of the input [95, 97, 99]. In addition to temporal adaptation, both ON and OFF BCs also maintain a fast ‘push-pull’ mechanism that balances the activation of the center relative to the surround for the purpose of locally enhancing the contrast within the RF [100, 101]. There is even some evidence that RGCs also employ a gain control mechanism to change the effective *size* of the RF based on the illumination [102].

Finally, the eye is not a static sensor – it is constantly performing a combination of motions at various spatial and temporal scales – ocular tremors, drift, microsaccades and saccades [103]. These different types of movements serve different purposes: for instance, small random movements (such as drift and microsaccades), which are outside voluntary control and occur even while fixating on a target, prevent the phenomenon of visual fading (also known as retinal fatigue), in which the rapid adaptation of the retinal ganglion cells to a fixed visual scene leads to the complete elimination of neural response (the eye effectively becomes blind to anything that is not moving) [103, 104]. The concept of a vision sensor modelled around the ret-

ina emerged more than three decades ago by building upon detailed research on the neuroscience of the mammalian retina. The following section presents a brief overview of the major milestones since the conception of the earliest retinomorphonic circuits to present-day commercial event-based cameras.

6.3 Retinomorphonic models

Neuromorphic systems that specifically deal with modelling the mammalian visual system (mostly the retina) are known as ‘retinomorphonic’ systems [105]. While some the operating principles of the circuits differ substantially from that of biological retinal cells, they are in many respects *functionally* identical (or at least very similar) to biological retinal cells.

To optimize the power consumption and response speed of sensors, most recent research on event-based vision is heavily oriented towards hardware solutions, which are difficult to customise for the purpose of implementing novel computer vision algorithms and processing paradigms. In addition, existing software libraries are geared towards processing and emulating the existing hardware rather than providing a generic base for implementing novel algorithms.

The ACT is thus pursuing research on retinomorphonic models to explore the question of whether sophisticated visual processing can be performed *on the sensor*, rather than being *learned* by a downstream algorithm. A prime example can be given with specialised ganglion cells which are sensitive to certain types of motion in their RF, to the exclusion of all other types of motion. For instance, there are ganglion cells that are sensitive to approaching motion while being insensitive to approaching or receding motion [84, 106], which can be useful for applications such as spacecraft landing, where the landscape is perceived as approaching from the perspective of the camera.

To maximise the efficiency and speed of onboard processing, it would be beneficial to design sensors capable of obtaining insight about motion and features in the visual field, such as edges, depth or optical flow. We know that biological ganglion cells with different RFs are sensitive only to very specific features and types of motion in the visual field (such as approaching motion or egomotion) while largely remaining silent in the presence of other types of motion (such as receding or lateral motion). It would therefore make sense to be able to model the function of these cells in a convenient and reproducible way that could ultimately inform the design of novel dedicated hardware sensors.

Therefore, the ACT is developing a simulation framework that can model as many features of the mammalian retina as possible in an efficient manner in order to facilitate rapid prototyping of retinomorphonic architectures. In particular, our research and development efforts have been geared towards implementing support for all of the following features:

- All major types of cells in the retina (cf. 6.1);
- Spatial and temporal adaptation and homeostasis (cf. 6.2);
- Complex RFs (e.g., centre/surround);
- Foveation and eccentricity-dependent RFs;
- Sparsification of cells towards the edges of the retina (i.e., cell distribution becomes increasingly sparse towards the edges of the retina, in proportion to increasing RF size).
- Saccadic movements (microsaccades, drift, etc.);

As outlined above in Section 6.1, the retina consists of several layers, whereby cells in each layer (except the photoreceptor layer) receive input from a certain local neighbourhood of cells in the preceding layer(s). This mechanism is behind the inspiration for CNNs, which have occupied a central role in computer vision domain of AI research [107, 108]. The successful application of CNNs on a wide range of tasks, combined with the availability of several software packages that offer highly optimised CNN implementations out of the box, has largely dampened any incentive to explore alternative architectures or mechanisms for implementing convolution. However, incorporating all of the above features into a coherent framework is not straightforward, specifically with respect to saccadic movements and eccentricity-dependent RFs. This requires rethinking how convolution is implemented.

6.3.1 Sparse convolution

Convolution is a costly operation to implement, and state-of-the-art CNNs rely on a number of tricks to speed up the process. N -dimensional convolution is usually implemented using a number of convolutional kernels (represented as a single $(N + 2)$ -dimensional tensor, where N is the number of dimensions of the input and two extra dimensions representing the number of kernels and the batch size). Therefore, without loss of generality, we focus on $2D$ convolution, unless explicitly stated otherwise.

The generic expression for convolving a batch of N $2D$ inputs (each with dimensions H and W) using C_{in} convolutional kernels can be given as follows:

$$Conv2d(N, C_{out}, H, W) = \sum_{i=0}^{C_{in}} W_{(C_{out}, i)} \star Input(N, i) \quad (6)$$

The dimensions H_{out} and W_{out} of the output channels are computed as

$$H_{out} = \left\lfloor \frac{H_{in} + 2p_H - d_H(k_H - 1) - 1}{s_H} + 1 \right\rfloor \quad (7)$$

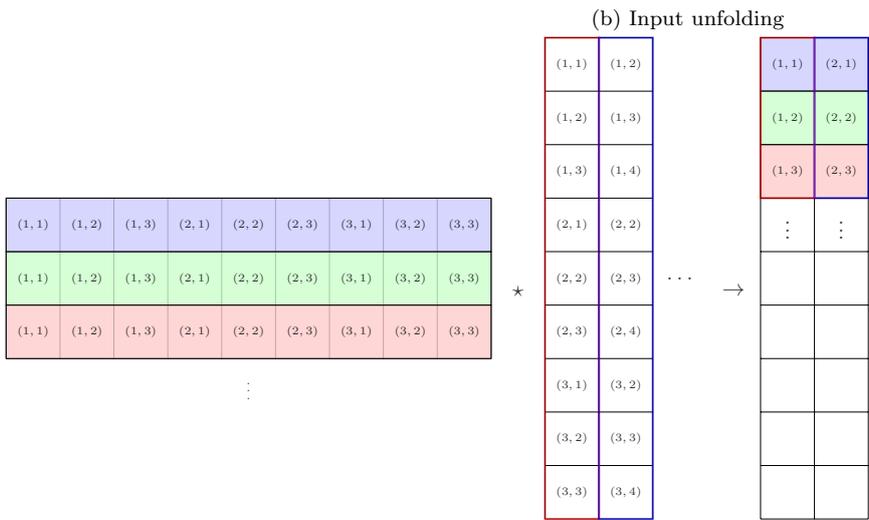
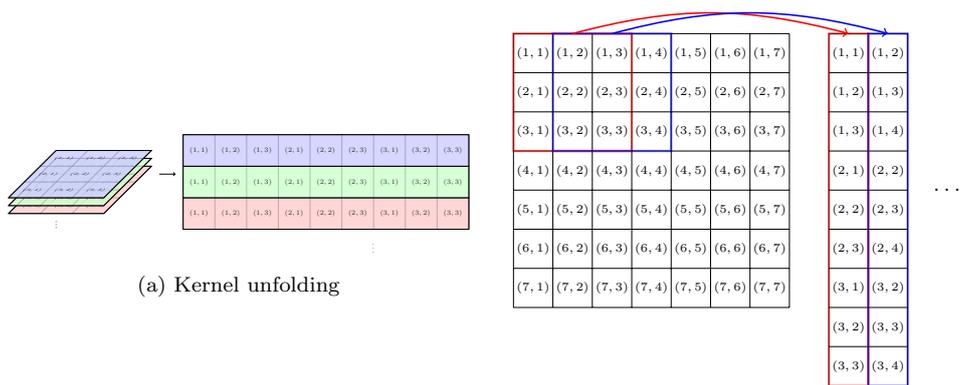
$$W_{out} = \left\lfloor \frac{W_{in} + 2p_W - d_W(k_W - 1) - 1}{s_W} + 1 \right\rfloor, \quad (8)$$

Here, p , d , k and s (with indices H and W) represent the padding, dilation, kernel size and stride of the convolution operation (in the direction of H and W , respectively). These four parameters impose constraints on each other depending on the input since H_{out} and W_{out} in Eqs. (7) and (8) must be integer. For instance, convolving an image of size 32×32 with a kernel of size 5×5 , a stride of 2 and dilation and padding of 0 means that the convolution would not be symmetrical at the edges of the image since ensuring that H_{out} and W_{out} are integer enforces the convolution operation to stop short of covering the last column and row of the image. Adding a padding of 2 around the image mitigates this problem and ensures that the convolution is symmetric. We refer the reader to the relevant literature for an in-depth overview of the meaning of these parameters [109].

In practical implementations, convolution relies on a crucial but often understated preprocessing step. This operation, known as ‘image-to-column’ (or *im2col* for short) [110] involves stretching out the convolutional kernels and arranging them into rows of a matrix, and stretching *all* kernel-sized subsections of the input into one-dimensional column vectors and arranging them as columns of another matrix (hence the term *im2col*, since $2D$ convolutions are most commonly applied to images) (Fig. 7). In this way, the convolution operation is reduced to a single dense matrix-matrix multiplication.

It is immediately clear that while it improves the efficiency of convolution, this approach introduces a number of cascading limitations that impede the effort to implement the desired features outlined in 6.3:

- All kernels must be the same size, which precludes foveation.
- An inability to convolve the input with kernels of different size limits the type of RFs that can be implemented; specifically, centre/surround RFs are difficult to implement with a fixed kernel size, especially in combination with foveation.
- Sparsification is impossible to model since the operation is by definition **dense**, so it applies the kernel to all segments of the input as defined in Eq. (6).
- It is not trivial to implement saccadic movements as this would entail performing the *im2col* preprocessing step for each movement. Furthermore, as an attention shifting mechanism, saccadic movements are much more useful in combination with foveation.



(c) Convolution operation.

Figure 7: Default Im2col operation.

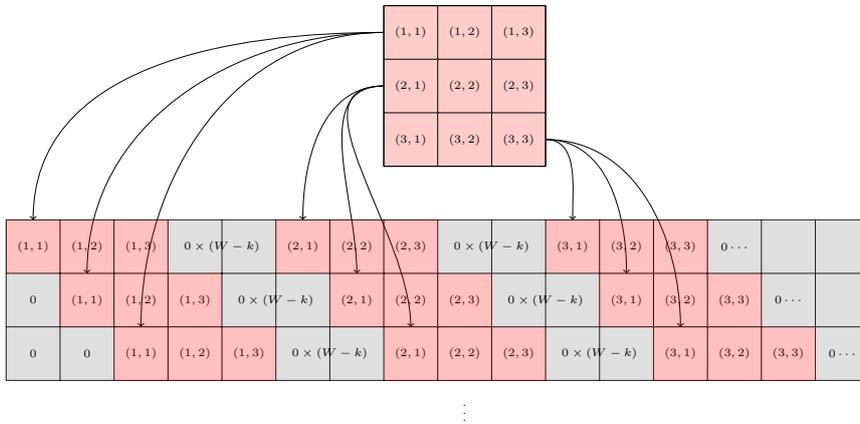


Figure 8: Unfolding a kernel into a sparse matrix. The space redundancy introduced in this operation is offset by the elimination of the *im2col* operation.

- Without saccades, adaptation (specifically temporal adaptation) as observed in the retina quickly erases features in stationary parts of the image.

To address these limitations, we have developed a straightforward alternative to the default convolution operation. Rather than unfolding the input, this alternative convolutional method relies on a *kernel* unfolding step (tentatively named *kernel2row*, or *k2row* for short) that consists of unfolding the kernels, splitting them into continuous chunks and arranging those chunks into a *sparse* matrix, while the input is unfolded into a single large column vector. The kernel chunks are arranged in such a way that the result of multiplying the sparse kernel matrix and the stretched input vector is identical to that obtained with the above dense convolution method. The *k2row* operation is outlined in Fig. 8.

The most notable advantage of performing convolution in this way is that the kernel size can be decoupled from the implementation constraints, allowing for kernels of different sizes to be used within the same convolution matrix.

Most of the desired features for simulating a retinal model as outlined in 6.3 can be implemented trivially with this sparse convolution method. For instance, foveation can be implemented by using smaller kernels towards the centre of the image than towards the edges. Similarly, a separate kernel ‘instance’ can be used to convolve each pixel of a $2D$ image, which allows us to implement complex RFs (such as ones relying on the centre-surround mechanism) as a difference of Gaussians (‘Mexican hat’ filter) represented by the difference of two convolutions with kernels of different size. Finally, saccadic movements can be implemented in combination with foveation as an attention shifting mechanism by simply zero-padding the input image in

a way that offsets the centre of the resulting padded image by the desired amount equal to the distance travelled by the ‘fovea’ during the saccadic movement.

To demonstrate the flexibility of this method, we implemented a pipeline emulating all retinal layers (see §6.1). In this model, the input to excitatory layers (receptors, bipolar cells and ganglion cells) is modulated by feedback from inhibitory layers (horizontal and amacrine cells). Cells in the excitatory layers have eccentricity-dependent RFs (smaller towards the centre of the image and gradually growing in size towards the edges). A local spatial normalisation map is implemented using the horizontal layer, whereas the bipolar layer implements a temporal filter as well as the split into separate ON and OFF pathways. The temporal filter was implemented as an exponential running mean which was updated iteratively as follows at every frame:

$$\mu_{t+1} = \alpha x + (1 - \alpha)\mu_t, \quad (9)$$

where μ_t is the mean at time step t , α is the ‘forgetting rate’ that determines how much of the preceding input is remembered, and x is the input frame. The mean is computed in this way for each pixel.

The ganglion layer preserves the separation of these two paths and implements two types of ganglion cells with ON-centre / OFF-surround and OFF-centre / ON-surround RFs, respectively. For the sake of brevity, we have omitted the demonstration of saccadic movements, which can be used for inducing events from an otherwise static input. The parameters for this model are given in Table 4. We use part of the Perseverance landing sequence as recorded with the rover’s onboard camera minutes before touchdown [111].

Several interesting observations can be made about the results in Fig. 9. First, despite the simplicity of the simulated retina, it can clearly perform certain types of preprocessing, such as local brightness normalisation at the horizontal layer, edge extraction (e.g., at the bipolar layer) and event generation at the ganglion layer. It is also noteworthy that ON/OFF and OFF/ON ganglion cells highlight the two sides (brighter and darker) of edges in the visual scene. It is noteworthy that the RF organisation in this case successfully suppresses noise in the output of the ganglion layer. This demonstration of effects observable in biological retinas is encouraging as we are working towards implementing more sophisticated processing of the incoming visual information, such as optical flow reconstruction and detection of approaching motion for applications such as obstacle avoidance and landing.

The retinal simulation code is part of a library named *pyrception*, which

Table 4: Parameters for the retinomorph model in 6.3.1

Cell type	Parameter	Value
Receptors	Input type	Grayscale
Horizontal, bi-polar, amacrine, ganglion	RF size distribution	Gaussian ¹
	RF size distribution: mean	$(H/2, W/2)^2$
	RF size distribution: SD	$(H/3, W/3)$
	RF type	Proportional ³
Horizontal, bi-polar, amacrine	Min. / max. RF size	$1 \times 1 / 4 \times 4$
Ganglion	Min. / max. RF size (centre)	$1 \times 1 / 4 \times 4$
	Min. / max. RF size (surround)	$4 \times 4 / 9 \times 9$
Bipolar cells	α^4	0.95
Ganglion cells	Threshold ⁵	15

¹ For all cell types except horizontal cells, the Gaussian was *inverted*, meaning that the RF size increased with distance from the centre. For horizontal cells, the RF size was *larger* at the centre and decreased towards the edges. This is consistent with the fact that horizontal cells are concentrated around the foveal region [81], where they facilitate the mechanism of adaptation to the local brightness level. A larger RF size for horizontal cells in the foveal region leads to a sharper image in the fovea, consistent with observations from biology.

² In this table, H and W denote the height and width of the input image, respectively.

³ Proportional here means that the value of each element of the kernel is the inverse of the kernel size (e.g., $1/9$ for a 3×3 kernel).

⁴ α is the ‘forgetting rate’ of the exponential running average used as a temporal filter in the bipolar layer.

⁵ The threshold indicates the difference between the intensity at the centre vs. the surround of the receptive field. An event is produced if this threshold is crossed.

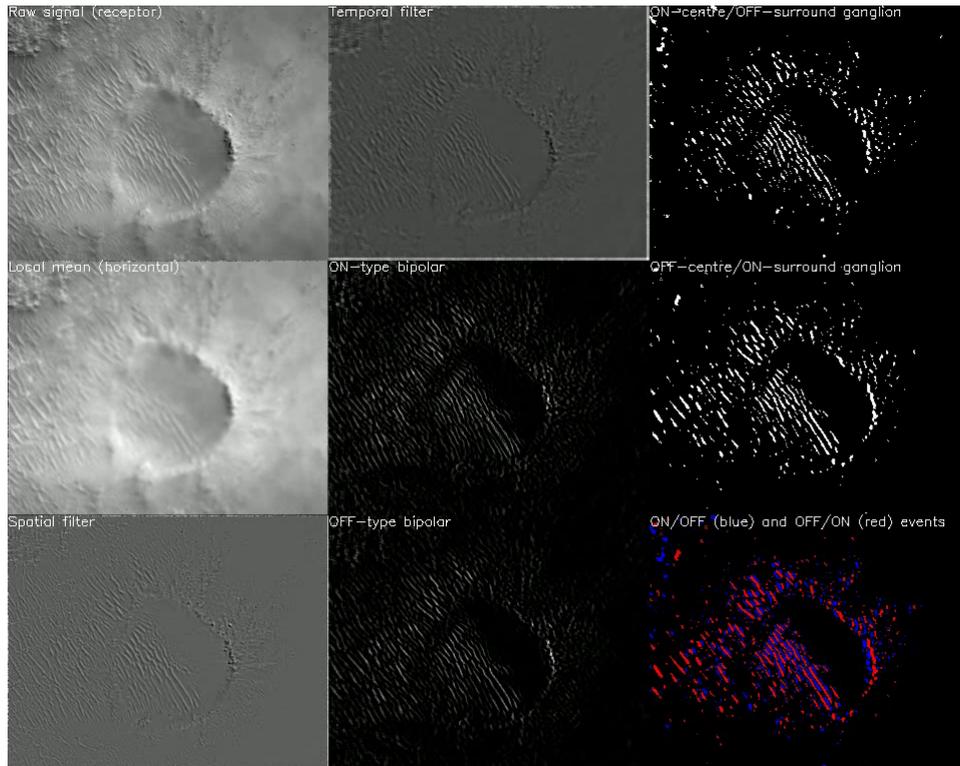


Figure 9: A snapshot of the state of each layer and the combination of ON/OFF and OFF/ON events. (a) Raw input (converted to grayscale). (b) Simulated horizontal cell layer computing the local mean illumination. (c) The raw signal with the mean illumination from (b) subtracted. (d) Temporal filter (implemented as a running average of the normalised signal in (c)). (e) ON and (f) OFF bipolar cells computing positive and negative deviation from the temporal mean, respectively (i.e., parts that are brighter and darker than the mean). Ganglion cells with (g) ON/OFF and (h) OFF/ON centre/surround RFs. In ON/OFF ganglion cells, the centre is driven by ON bipolar cells and the surround by OFF bipolar cells, and vice versa for OFF/ON ganglion cells. A ganglion cell produces a binary spike (event) when the difference between the centre and surround crosses a predefined threshold (a model parameter). Finally, (i) shows the combined activity of ON/OFF (blue) and OFF/ON (red) ganglion cells.

in turn is part of the ACT open-source ecosystem ⁵. The ultimate goal of the *pyrception* library is to provide convenient interfaces for converting input from multiple sensory modalities (such as auditory and olfactory) into spike trains, which can be used downstream for spike-based multi-modal learning tasks.

6.4 Existing silicon retinas

The understanding that the retina deals with the inherent unreliability of ‘wetware’ by responding to relative intensity changes (rather than the absolute magnitude) in the visual field [112] has been central to the development of silicon analogues to the retina. One of the earliest circuits that modelled the adaptive behaviour of retinal cells was a photoreceptor that mimicked some homeostatic processes in the retina. Specifically, the receptor could detect *changes* in illumination rather than the absolute illumination level [113, 114]. The idea behind this circuit was inspired by the operation of receptors and bipolar cells (see §6.1). It employed a feedback mechanism that compared the response to the changes in intensity to a filtered version of the output of the receptor, with an added hysteretic element that provided a ‘memory’ of recent illumination. Importantly, the circuitry emulated the way that biological photoreceptors adapt to changes in light intensity on a logarithmic rather than linear scale (a phenomenon known as Weber’s law). In other words, if the output of the receptor is given as $V \propto \log(I)$, then the change in that response would be logarithmic as well ($\delta V \propto \delta I/I$).

A prominent milestone along this line of research was the development of a ‘silicon retina’ [115], where each photoreceptor could adapt to the local light intensity through feedback from simulated horizontal cells and the adapted signal was further amplified by silicon analogues of bipolar cells. The silicon retina demonstrated how a globally connected matrix of elements can still adapt to local fluctuations in light intensity, producing a familiar output comparable to that of an actual biological retina. Ultimately, a circuit was demonstrated that could model the behaviour of all five types of retinal cells, including ganglion cells [116].

Subsequently, the Address Event Representation (AER) [117, 118] was developed with a slightly different objective, namely efficient communication of a continuous signal in the form of multiplexed spike trains. The generic nature of AER was used to demonstrate how a circuit could perform spatial convolution (using, for instance, Gabor filters), allowing for feature extraction and tracking with high temporal resolution.

⁵<https://gitlab.com/cantordust/pyrception>

6.4.1 The dynamic vision sensor (DVS)

One of the main problems with the early circuits described above was that they were built primarily as proof-of-concept devices – they were never meant to be used as actual camera sensors. This changed with the development of the dynamic vision sensor (DVS) [15, 16].

In many ways, DVS is a simplification on the original approach of modelling the retina holistically (i.e., all the layers separately). For instance, pixels in the DVS camera are entirely independent of all other pixels, allowing for a completely asynchronous operation that employs AER to output events (x and y coordinates, timestamp t and polarity p). DVS pixels are relatively simple and implement the entire pipeline from receptors through bipolar cells and ganglion cells within a single circuit. The main principle of event generation and adaptation remains similar to that in earlier research: each pixel detects changes in the brightness (log intensity) in either positive or negative direction relative to a baseline. Once the derivative of the logarithmic-intensity crosses a fixed threshold Θ , an event is produced and the brightness value in correspondence of the threshold plus some delay time set as the new baseline. In this way, rapid changes in brightness result in a large number of events, whereas slow fluctuations produce few events. Importantly, since the baseline is *reset independently for each pixel* to the level of the threshold each time the respective pixel produces an event, the sensor as a whole has an exceedingly large dynamic range – on the order of $140dB$. This means that stark contrast does not saturate the sensor because the response of a pixel on the dark side of an edge does not depend on the response on the bright side of the same edge; it only depends on the *relative* changes in brightness (Fig. 10).

The DAVIS sensor [15] adds active pixel sensor circuitry to the core DVS logic in order to be able to keep information about the absolute brightness in the scene while retaining the advantages of the DVS sensor, such as high dynamic range, low latency and sparse activity. This allows a conventional frame-based readout to be obtained simultaneously with the fast asynchronous event-based output of the DVS sensor.

DVS pixels also show a characteristic junction leakage current, which depends strongly on the temperature [119]. This leakage current is the source of sporadic (noisy) ON events, with the noise level becoming more prominent at high temperature and in dark environments.

The DVS pixel circuits enabled event-based sensors to bridge the gap between proof of concept and usability, leading to several commercial implementations (e.g., by Inivation⁶ and Prophesee⁷).

There are also some notable points where the DVS departs from the operating principles of the mammalian retina. For instance, the DVS does

⁶<https://inivation.com/>

⁷<https://www.prophesee.ai/>

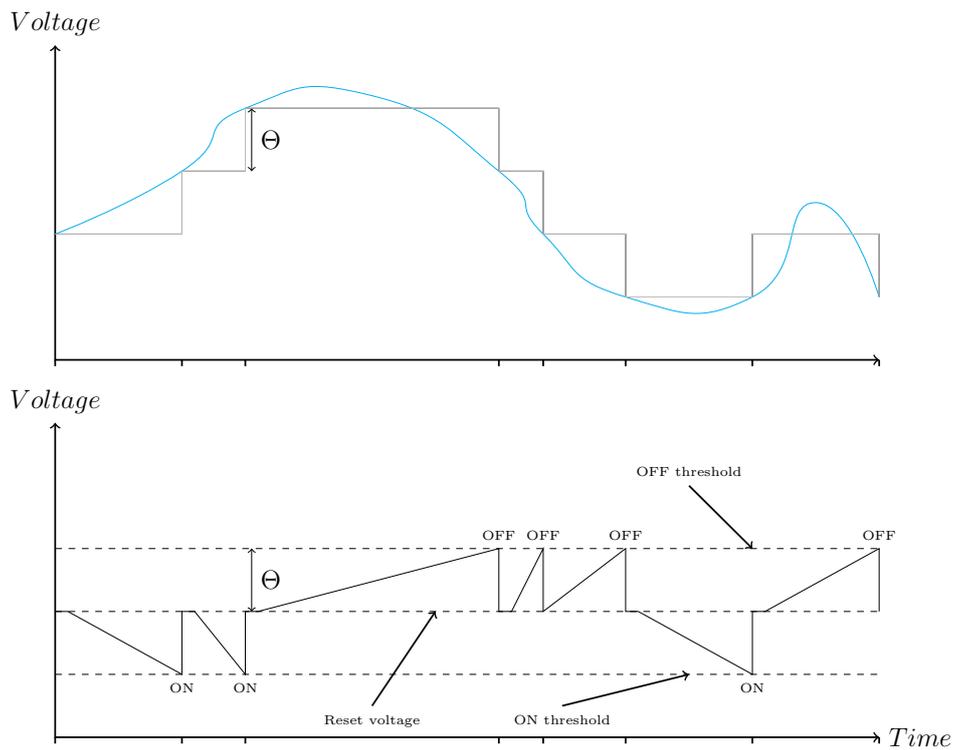


Figure 10: Operating principle of a DVS pixel. Every time the brightness of the area of the visual scene covered by the pixel increases (decreases) by a fixed amount (quantified by the threshold Θ), the pixel generates an ON (OFF) event. The AER representation is commonly used for encoding the output of an event sensor, whereby an AER event consists of a timestamp, the x and y coordinates of the pixel that generated the event and the polarity of the output (+1 for ON and -1 for OFF events).

not fully preserve the ON/OFF pathway separation since a single pixel can produce both ON and OFF events (whereas in biological systems these pathways are in fact kept separate all the way to the visual cortex [82]). Also, due to the lack of cross-talk between pixels, it is impossible for DVS cameras to implement the centre/surround RFs of bipolar and ganglion cells (although recent research aimed to introduce RFs [120], mainly with the goal of suppressing noise). Other features of the retina and the eye in general, such as microsaccades [121] and foveation [122], have also been attracting the attention of event-based vision researchers in recent years. In parallel, there are also efforts to develop software that can emulate, simulate or otherwise complement the capabilities offered by hardware event-based sensors. For instance, the development of the video2events (v2e) framework [123] and the ESIM library [124] was driven by the objective to harness existing conventional video datasets for event-based vision research.

6.5 Why in space?

In 2021, the DAVIS240 sensor became the first neuromorphic device to be launched in space and is, at the time of writing, also onboard the ISS launched as part of the ISS resupply payload in the Falcon Neuro Project [125]. This shows the great interest aroused by neuromorphic sensing on the aerospace experts. The understanding and quantification of the advantages of event based cameras in the context of space missions is, though, only at its beginning [126, 127, 128, 129, 130]. Besides the discussed benefits of pursuing a neuromorphic approach whenever low-resource requirements are driving the design of a mission, additional advantages over a conventional camera could derive from its higher dynamic range, achieving superior performance in high-contrast scenes as well as allowing a relaxation on the constraints imposed by light entering the field of view of the sensor. The high temporal resolution and compression qualities of an event sensor are also of interest, as well as its recently proved resilience to different types of radiation (such as wide-spectrum neutrons) that are often present in space operating environments [131]. Promising scenarios where event based cameras could make a difference in a space context are landing [126, 127], pose estimation [129], planetary surface mapping, astronomy [128, 130], monitoring of specific events (e.g., explosive events, pebble dynamics around asteroids, etc.). In all these cases it can be argued that the neuromorphic nature of the camera (combining high power efficiency and other beneficial properties), could bring advantages with respect to conventional systems, especially when a neuromorphic approach also used from the algorithmic side to tackle the interpretation of the produced events [126]. However, an open question that remains in this field is the quantification of the domain gap introduced by the use of synthetic data produced by video-to-event converters (a classical example being v2e [123]). A recent work [129] addresses this issue by pro-

posing and testing an algorithm (trained on a synthetic dataset) on event streams created by a realistic mock-up experiment performed in a lab.

6.6 Landing with events

The ACT has recently demonstrated the suitability of event-based vision for autonomous planetary landing operations by inferring and processing optical flow measurements to predict time-to-contact (TTC) and divergence. TTC is defined as:

$$\text{TTC} = -\frac{z}{v_z} \quad (10)$$

where z is the altitude of the spacecraft with respect to the planetary surface and v_z its vertical velocity, so that a descending spacecraft will always satisfy $z > 0$ and $v_z < 0$. Previous work has already demonstrated the suitability of TTC feedback for controlling a spacecraft during flight operations and landing scenarios by using visual information from traditional frame-based sensors onboard micro air vehicles [132, 133]. Our recent studies have shown that event data could also be effective for reconstructing on-board TTC, therefore bridging the gap from on-board event data streams to real-time control [126, 127].

In these studies, event streams representative of relevant ventral landing scenarios were constructed synthetically by converting sequences of frames similar to those provided by traditional frame-based sensors to events. The PANGU (Planet and Asteroid Natural Scene Simulation Utility)[134] software was used to generate proxy models of the Moon surface and render frames representative of the output of onboard vision instrumentation. From the generated synthetic events, TTC was estimated by identifying features in the event stream corresponding to static structures in the planetary body (such as craters and boulders) and computing their rate of perceived expansion or optical flow divergence. The reconstructed TTC from the event stream was fed back into a closed-loop control system to simulate real-time autonomous ventral landing by enforcing constantly decreasing TTC. Simulations were performed for different lunar terrains characterized by diverse crater and boulder distributions. Results showcased the suitability of the solution as a proof of concept for real-time onboard processing of events in planetary landing scenarios, where the main areas with room for improvement that could benefit from future research are the inference of divergence and the inclusion of more accurate noise models for event modelling based on known properties of DVS hardware.[126]

In recent work, the divergence estimation procedure was further improved by using a contrast maximisation [135, 136] formulation for event-based divergence estimation, demonstrating its applicability on both synthetic and real event datasets.[127] Spacecraft position and velocity vectors were obtained for several planetary ventral landing profiles computed via

indirect optimisation methods. A UR5 robotic arm with a Prophesee GEN 4 event sensor was employed to replicate these ventral trajectories using 2D and 3D printed planetary surfaces. Some of the generated trajectories were used to render 3D landing reconstructions in PANGU, which were then passed through the v2e [123] pipeline to generate synthetic event streams. A mathematical formulation for the event-based radial flow was proposed, and a GPU-accelerated optimisation procedure was used to maximise the contrast of the resulting flow-compensated event images and estimate the divergence of event batches. Comparisons to other state-of-the-art divergence estimation methods showcased the accuracy and stability of the proposed procedure across a wide variety of event streams, with competitive run times achieved owing to the GPU-accelerated implementation.

6.7 From events to egomotion

Events as produced by the DVS or some other device based on retinomorph models contain information on the perceived scene motion as projected in the camera plane. But what can be learned about the spacecraft egomotion (or pose) using only this information? It is conceivable that events can be used for reconstructing the motion field, whereby the problem of motion field interpretation has already been discussed in the literature [137]. A number of researchers have commented on this issue, starting with the work by Longuet Higgins [138], which addressed the question for the case of retinal vision, and subsequent works that lay down the mathematical structure of the problem, mapping it into a linear problem and thus marking it as “solved” [137]. More recently, substantial work in the area of drones and robotic vision has been produced based on these fundamental results established in the 80s and 90s.

Let us consider the familiar equations [137, 138]:

$$\begin{cases} u = (v_{xc} + xv_{zc})h(x, y) + q + ry - pxy + qx^2 \\ v = (v_{yc} + yv_{zc})h(x, y) + rx - p - py^2 + qxy \end{cases}$$

where u, v are the velocities of the feature in x, y on the camera plane, p, q, r the angular velocity components of the camera, v_{xc}, v_{yc}, v_{zc} the velocity of the camera center of mass and $h(x, y)$ the inverse of the depth map. From these equations, if $h, v_{xc}, v_{yc}, v_{zc}$ constitutes a solution, so does $\frac{h}{k}, kv_{xc}, kv_{yc}, kv_{zc}$; it is thus impossible to invert the motion field to absolute values of all motion parameters, as only relative depth can be distinguished. Approximating the surface as an infinite plane, one can show that necessarily:

$$h(x, y) = \alpha x + \beta y + \gamma$$

where $H^2(\alpha^2 + \beta^2 + \gamma^2) = 1$, having denoted the spacecraft altitude with H . Let us analyse the structure of the resulting equations. If we assume to know the motion field $u_i, v_i, \quad i = 1..n$ at n different points, we can write a system

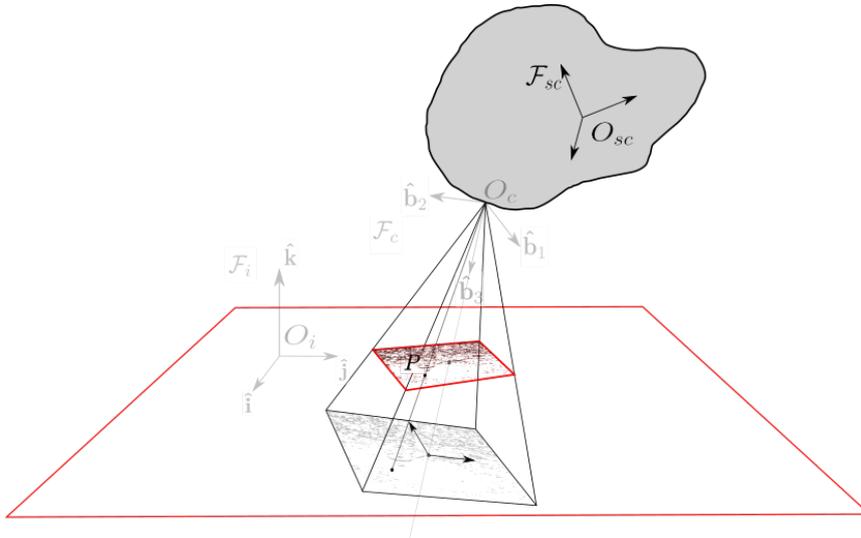


Figure 11: Camera geometry during planetary landing. The three reference systems, namely planet (i), camera (c) and spacecraft (sc), are shown as well as the camera plane projection of a generic surface feature assumed to be stationary.

of $2n$ equations in the unknowns $v_{xc}, v_{yc}, v_{zc}, p, q, r, \alpha, \beta, \gamma$. As before, only relative depth can be estimated, hence one can set $H = 1$ and consider $\gamma^2 = 1 - \alpha^2 - \beta^2$ to conclude that in theory $n = 4$ points of the motion field are necessary and sufficient to fully determine the spacecraft egomotion in relative terms to the altitude H . This conclusion does not consider the effect of measurement noise on the resulting estimates, a problem that is highly dependent on the sensor used (i.e., the quality of the events produced and their translation into a motion field). To conclude, it is worth mentioning the recent work from Jawaid *et al.*[129] on pose estimation from events. In a docking scenario (both simulated and reproduced in a lab mockup), the approach proposed in that work is able to determine from event data also the absolute values of mutual distances. This is not in contradiction to what is claimed above, as the algorithm is not exploiting depth cues from a reconstructed motion field but rather cues learned directly from a predefined dataset of known poses where the absolute dimensions appear explicitly.

References

- [1] C. Frenkel, D. Bol and G. Indiveri. ‘Bottom-up and top-down neural processing systems design: Neuromorphic intelligence as the convergence of natural and artificial intelligence’. *arXiv preprint arXiv:2106.01288* (2021).

- [2] K. Roy, A. Jaiswal and P. Panda. ‘Towards spike-based machine intelligence with neuromorphic computing’. *Nature* 575.7784 (2019), pp. 607–617.
- [3] G. Indiveri and T. K. Horiuchi. *Frontiers in neuromorphic engineering*. 2011.
- [4] G. Hinton. ‘The Forward-Forward Algorithm: Some Preliminary Investigations’ (2022).
- [5] G. Furano, G. Meoni et al. ‘Towards the use of artificial intelligence on the edge in space systems: Challenges and opportunities’. *IEEE Aerospace and Electronic Systems Magazine* 35.12 (2020), pp. 44–56.
- [6] F. C. Bruhn, N. Tsog et al. ‘Enabling radiation tolerant heterogeneous GPU-based onboard data processing in space’. *CEAS Space Journal* 12.4 (2020), pp. 551–564.
- [7] G. Giuffrida, L. Fanucci et al. ‘The Phi-Sat-1 Mission: The First On-Board Deep Neural Network Demonstrator for Satellite Earth Observation’. *IEEE Transactions on Geoscience and Remote Sensing* 60 (2022), pp. 1–14. DOI: [10.1109/TGRS.2021.3125567](https://doi.org/10.1109/TGRS.2021.3125567).
- [8] G. Labrèche, D. Evans et al. ‘OPSSAT Spacecraft Autonomy with TensorFlow Lite, Unsupervised Learning, and Online Machine Learning’. *2022 IEEE Aerospace Conference*. 2022.
- [9] A. S. Danielsen, T. A. Johansen and J. L. Garrett. ‘Self-Organizing Maps for Clustering Hyperspectral Images On-Board a CubeSat’. *Remote Sensing* 13.20 (2021), p. 4174.
- [10] V. Medici, G. Orchard et al. *Neuromorphic Computation of Optic Flow Data*. Tech. rep. 08/6303. Available online at www.esa.int/act. European Space Agency, the Advanced Concepts Team, 2010.
- [11] F. Valette, F. Ruffier and S. Viollet. *Neuromorphic Computation of Optic Flow Data*. Tech. rep. 08/6303b. Available online at www.esa.int/act. European Space Agency, the Advanced Concepts Team, 2009.
- [12] F. Valette, F. Ruffier et al. ‘Biomimetic optic flow sensing applied to a lunar landing scenario’. *2010 IEEE International Conference on Robotics and Automation*. IEEE. 2010.
- [13] D. Izzo, N. Weiss and T. Seidl. ‘Constant-optic-flow lunar landing: Optimality and guidance’. *Journal of Guidance, Control, and Dynamics* 34.5 (2011), pp. 1383–1395.
- [14] M. Frye. ‘Elementary motion detectors’. *Current Biology* 25.6 (2015), R215–R217.

- [15] C. Brandli, R. Berner et al. ‘A 240×180 130 dB 3 Ms Latency Global Shutter Spatiotemporal Vision Sensor’. *IEEE Journal of Solid-State Circuits* 49.10 (2014), pp. 2333–2341. ISSN: 1558-173X. DOI: [10.1109/JSSC.2014.2342715](https://doi.org/10.1109/JSSC.2014.2342715).
- [16] G. Taverni, D. Paul Moeys et al. ‘Front and Back Illuminated Dynamic and Active Pixel Vision Sensors Comparison’. *IEEE Transactions on Circuits and Systems II: Express Briefs* 65.5 (2018), pp. 677–681. ISSN: 1549-7747, 1558-3791. DOI: [10.1109/TCSII.2018.2824899](https://doi.org/10.1109/TCSII.2018.2824899). (Visited on 17/11/2022).
- [17] M. Davies, N. Srinivasa et al. ‘Loihi: A neuromorphic manycore processor with on-chip learning’. *IEEE Micro* 38.1 (2018), pp. 82–99.
- [18] F. Akopyan, J. Sawada et al. ‘Truenorth: Design and tool flow of a 65 mw 1 million neuron programmable neurosynaptic chip’. *IEEE transactions on computer-aided design of integrated circuits and systems* 34.10 (2015), pp. 1537–1557.
- [19] Z. Ye, R. Liu et al. ‘Evaluation of Radiation Effects in RRAM-Based Neuromorphic Computing System for Inference’. *IEEE Transactions on Nuclear Science* 66.1 (2019), pp. 97–103. ISSN: 1558-1578. DOI: [10.1109/TNS.2018.2886793](https://doi.org/10.1109/TNS.2018.2886793).
- [20] A. L. Hodgkin and A. F. Huxley. ‘Action potentials recorded from inside a nerve fibre’. *Nature* 144.3651 (1939), pp. 710–711.
- [21] W. Gerstner and R. Naud. ‘How good are neuron models?’ *Science* 326.5951 (2009), pp. 379–380.
- [22] R. Brette and W. Gerstner. ‘Adaptive exponential integrate-and-fire model as an effective description of neuronal activity’. *Journal of neurophysiology* 94.5 (2005), pp. 3637–3642.
- [23] G. Bellec, D. Salaj et al. ‘Long short-term memory and learning-to-learn in networks of spiking neurons’. *Advances in neural information processing systems* 31 (2018).
- [24] B. Yin, F. Corradi and S. M. Bohtë. ‘Accurate and efficient time-domain classification with adaptive spiking recurrent neural networks’. *Nature Machine Intelligence* 3.10 (2021), pp. 905–913.
- [25] F. Zenke, S. M. Bohtë et al. ‘Visualizing a joint future of neuroscience and neuromorphic engineering’. *Neuron* 109.4 (2021), pp. 571–575.
- [26] S. Thorpe, D. Fize and C. Marlot. ‘Speed of processing in the human visual system’. *nature* 381.6582 (1996), pp. 520–522.
- [27] R. Gütig and H. Sompolinsky. ‘The tempotron: a neuron that learns spike timing–based decisions’. *Nature neuroscience* 9.3 (2006), pp. 420–428.

- [28] D. Dold, J. Soler Garrido et al. ‘Neuro-symbolic computing with spiking neural networks’. *Proceedings of the International Conference on Neuromorphic Systems 2022*. 2022.
- [29] D. Dold. ‘Relational representation learning with spike trains’. *2022 International Joint Conference on Neural Networks (IJCNN)*. 2022. DOI: [10.1109/IJCNN55064.2022.9892829](https://doi.org/10.1109/IJCNN55064.2022.9892829).
- [30] L. Buesing, J. Bill et al. ‘Neural dynamics as sampling: a model for stochastic computation in recurrent networks of spiking neurons’. *PLoS computational biology* 7.11 (2011), e1002211.
- [31] M. A. Petrovici, J. Bill et al. ‘Stochastic inference with spiking neurons in the high-conductance state’. *Physical Review E* 94.4 (2016), p. 042312.
- [32] A. Payeur, J. Guerguiev et al. ‘Burst-dependent synaptic plasticity can coordinate learning in hierarchical circuits’. *Nature neuroscience* 24.7 (2021), pp. 1010–1019.
- [33] E. O. Neftci, H. Mostafa and F. Zenke. ‘Surrogate gradient learning in spiking neural networks: Bringing the power of gradient-based optimization to spiking neural networks’. *IEEE Signal Processing Magazine* 36.6 (2019), pp. 51–63.
- [34] M. Pfeiffer and T. Pfeil. ‘Deep learning with spiking neurons: opportunities and challenges’. *Frontiers in neuroscience* (2018), p. 774.
- [35] S. Davidson and S. B. Furber. ‘Comparison of artificial and spiking neural networks on digital hardware’. *Frontiers in Neuroscience* 15 (2021), p. 651141.
- [36] A. S. Kucik and G. Meoni. ‘Investigating spiking neural networks for energy-efficient on-board ai applications. a case study in land cover and land use classification’. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2021.
- [37] G.-q. Bi and M.-m. Poo. ‘Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type’. *Journal of neuroscience* 18.24 (1998), pp. 10464–10472.
- [38] B. Illing, J. Ventura et al. ‘Local plasticity rules can learn deep representations using self-supervised contrastive predictions’. *Advances in Neural Information Processing Systems* 34 (2021), pp. 30365–30379.
- [39] H. Mostafa. ‘Supervised learning based on temporal coding in spiking neural networks’. *IEEE transactions on neural networks and learning systems* 29.7 (2017), pp. 3227–3235.

- [40] I. M. Comsa, K. Potempa et al. ‘Temporal coding in spiking neural networks with alpha synaptic function’. *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE. 2020.
- [41] S. R. Kheradpisheh and T. Masquelier. ‘Temporal backpropagation for spiking neural networks with one spike per neuron’. *International Journal of Neural Systems* 30.06 (2020), p. 2050027.
- [42] J. Göltz, L. Kriener et al. ‘Fast and energy-efficient neuromorphic deep learning with first-spike times’. *Nature machine intelligence* 3.9 (2021), pp. 823–835.
- [43] F. Zenke and S. Ganguli. ‘Superspike: Supervised learning in multilayer spiking neural networks’. *Neural computation* 30.6 (2018), pp. 1514–1541.
- [44] F. Zenke and T. P. Vogels. ‘The remarkable robustness of surrogate gradient learning for instilling complex function in spiking neural networks’. *Neural computation* 33.4 (2021), pp. 899–925.
- [45] Y. LeCun and C. Cortes. ‘MNIST handwritten digit database’ (2010).
- [46] N. Perez-Nieves, V. C. Leung et al. ‘Neural heterogeneity promotes robust learning’. *Nature communications* 12.1 (2021), pp. 1–9.
- [47] T. C. Wunderlich and C. Pehle. ‘Event-based backpropagation can compute exact gradients for spiking neural networks’. *Scientific Reports* 11.1 (2021), pp. 1–17.
- [48] J. K. Eshraghian, M. Ward et al. ‘Training spiking neural networks using lessons from deep learning’. *arXiv preprint arXiv:2109.12894* (2021).
- [49] C. Pehle and J. E. Pedersen. *Norse - A deep learning library for spiking neural networks*. Version 0.0.7. Documentation: <https://norse.ai/docs/>. January 2021. DOI: [10.5281/zenodo.4422025](https://doi.org/10.5281/zenodo.4422025).
- [50] W. Fang, Y. Chen et al. *SpikingJelly*. 2020. URL: <https://github.com/fangwei123456/spikingjelly> (visited on 09/11/2022).
- [51] H. Hazan, D. J. Saunders et al. ‘BindsNET: A Machine Learning-Oriented Spiking Neural Networks Library in Python’. *Frontiers in Neuroinformatics* 12 (2018), p. 89. DOI: [10.3389/fninf.2018.00089](https://doi.org/10.3389/fninf.2018.00089).
- [52] S. B. Furber, D. R. Lester et al. ‘Overview of the SpiNNaker system architecture’. *IEEE transactions on computers* 62.12 (2012), pp. 2454–2467.
- [53] C. Mayr, S. Hoepfner and S. Furber. ‘Spinnaker 2: A 10 million core processor system for brain simulation and machine learning’. *arXiv preprint arXiv:1911.02385* (2019).

- [54] J. Shen, D. Ma et al. ‘Darwin: a neuromorphic hardware co-processor based on spiking neural networks’. *Science China Information Sciences* 59.2 (2016), pp. 1–5.
- [55] J. Pei, L. Deng et al. ‘Towards artificial general intelligence with hybrid Tianjic chip architecture’. *Nature* 572.7767 (2019), pp. 106–111.
- [56] C. Pehle, S. Billaudelle et al. ‘The BrainScaleS-2 accelerated neuromorphic system with hybrid plasticity’. *Frontiers in Neuroscience* 16 (2022).
- [57] T. Pfeil, A. Grübl et al. ‘Six networks on a universal neuromorphic computing substrate’. *Frontiers in neuroscience* 7 (2013), p. 11.
- [58] S. Moradi, N. Qiao et al. ‘A scalable multicore architecture with heterogeneous memory structures for dynamic neuromorphic asynchronous processors (DYNAPs)’. *IEEE transactions on biomedical circuits and systems* 12.1 (2017), pp. 106–122.
- [59] N. Qiao, H. Mostafa et al. ‘A reconfigurable on-line learning spiking neuromorphic processor comprising 256 neurons and 128K synapses’. *Frontiers in neuroscience* 9 (2015), p. 141.
- [60] B. V. Benjamin, P. Gao et al. ‘Neurogrid: A mixed-analog-digital multichip system for large-scale neural simulations’. *Proceedings of the IEEE* 102.5 (2014), pp. 699–716.
- [61] *Loihi - Intel*. URL: <https://en.wikichip.org/wiki/intel/loihi> (visited on 14/11/2022).
- [62] M. Davies. ‘Benchmarks for progress in neuromorphic computing’. *Nature Machine Intelligence* 1.9 (2019), pp. 386–388.
- [63] C. Frenkel, J.-D. Legat and D. Bol. ‘A 28-nm convolutional neuromorphic processor enabling online learning with spike-based retinas’. *2020 IEEE International Symposium on Circuits and Systems (ISCAS)*. IEEE. 2020.
- [64] G. K. Chen, R. Kumar et al. ‘A 4096-neuron 1M-synapse 3.8-pJ/SOP spiking neural network with on-chip STDP learning and sparse weights in 10-nm FinFET CMOS’. *IEEE Journal of Solid-State Circuits* 54.4 (2018), pp. 992–1002.
- [65] C.-K. Lin, A. Wild et al. ‘Programming spiking neural networks on Intel’s Loihi’. *Computer* 51.3 (2018), pp. 52–61.
- [66] A. Shrestha, H. Fang et al. ‘In-hardware learning of multilayer spiking neural networks on a neuromorphic processor’. *2021 58th ACM/IEEE Design Automation Conference (DAC)*. IEEE. 2021.

- [67] A. Vanarse, A. Osseiran et al. ‘Application of Neuromorphic Olfactory Approach for High-Accuracy Classification of Malts’. *Sensors* 22.2 (2022), p. 440.
- [68] C. Liu, G. Bellec et al. ‘Memory-efficient deep learning on a SpiN-Naker 2 prototype’. *Frontiers in neuroscience* 12 (2018), p. 840.
- [69] E. Stromatias, D. Neil et al. ‘Scalable energy-efficient, low-latency implementations of trained spiking deep belief networks on spinnaker’. *2015 International Joint Conference on Neural Networks (IJCNN)*. IEEE. 2015.
- [70] A. S. Cassidy, P. Merolla et al. ‘Cognitive computing building block: A versatile and efficient digital neuron model for neurosynaptic cores’. *The 2013 International Joint Conference on Neural Networks (IJCNN)*. IEEE. 2013.
- [71] M. Bouvier, A. Valentian et al. ‘Spiking neural networks hardware implementations and challenges: A survey’. *ACM Journal on Emerging Technologies in Computing Systems (JETC)* 15.2 (2019), pp. 1–35.
- [72] B. Han, A. Ankit et al. ‘Cross-layer design exploration for energy-quality tradeoffs in spiking and non-spiking deep artificial neural networks’. *IEEE Transactions on Multi-Scale Computing Systems* 4.4 (2017), pp. 613–623.
- [73] P. Helber, B. Bischke et al. ‘EuroSAT: A novel dataset and deep learning benchmark for land use and land cover classification’. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* (2019).
- [74] P. Helber, B. Bischke et al. ‘Introducing EuroSAT: A Novel Dataset and Deep Learning Benchmark for Land Use and Land Cover Classification’. *IGARSS 2018-2018 IEEE International Geoscience and Remote Sensing Symposium*. IEEE. 2018.
- [75] K. Simonyan and A. Zisserman. ‘Very deep convolutional networks for large-scale image recognition’. *arXiv preprint arXiv:1409.1556* (2014).
- [76] J. Deng, W. Dong et al. ‘Imagenet: A large-scale hierarchical image database’. *2009 IEEE conference on computer vision and pattern recognition*. Ieee. 2009.
- [77] N. Franceschini, J.-M. Pichon and C. Blanes. ‘From insect vision to robot vision’. *Philosophical Transactions of The Royal Society Of London. Series B: Biological Sciences* 337.1281 (1992), pp. 283–294.
- [78] Y. M. Song, Y. Xie et al. ‘Digital cameras with designs inspired by the arthropod eye’. *Nature* 497.7447 (2013), pp. 95–99.

- [79] D. H. Hubel and T. N. Wiesel. ‘Receptive Fields and Functional Architecture in Two Nonstriate Visual Areas (18 and 19) of the Cat’. *Journal of Neurophysiology* 28.2 (1965), pp. 229–289. ISSN: 0022-3077, 1522-1598. DOI: [10.1152/jn.1965.28.2.229](https://doi.org/10.1152/jn.1965.28.2.229). (Visited on 19/11/2022).
- [80] E. N. Grigoryan. ‘Self-Organization of the Retina during Eye Development, Retinal Regeneration In Vivo, and in Retinal 3D Organoids In Vitro’. *Biomedicines* 10.6 (2022), p. 1458. ISSN: 2227-9059. DOI: [10.3390/biomedicines10061458](https://doi.org/10.3390/biomedicines10061458).
- [81] R. H. Masland. ‘Cell Populations of the Retina: The Proctor Lecture’. *Investigative Ophthalmology & Visual Science* 52.7 (2011), pp. 4581–4591. ISSN: 1552-5783. DOI: [10.1167/iovs.10-7083](https://doi.org/10.1167/iovs.10-7083).
- [82] R. Mazade, J. Jin et al. ‘Functional Specialization of ON and OFF Cortical Pathways for Global-Slow and Local-Fast Vision’. *Cell Reports* 27.10 (2019), 2881–2894.e5. ISSN: 22111247. DOI: [10.1016/j.celrep.2019.05.007](https://doi.org/10.1016/j.celrep.2019.05.007). (Visited on 29/05/2020).
- [83] T. Euler, S. Haverkamp et al. ‘Retinal Bipolar Cells: Elementary Building Blocks of Vision’. *Nature Reviews Neuroscience* 15.8 (2014), pp. 507–519. ISSN: 1471-003X, 1471-0048. DOI: [10.1038/nrn3783](https://doi.org/10.1038/nrn3783). (Visited on 10/11/2021).
- [84] T. Gollisch and M. Meister. ‘Eye Smarter than Scientists Believed: Neural Computations in Circuits of the Retina’. *Neuron* 65.2 (2010), pp. 150–164. ISSN: 08966273. DOI: [10.1016/j.neuron.2009.12.009](https://doi.org/10.1016/j.neuron.2009.12.009). (Visited on 26/11/2021).
- [85] R. H. Masland. ‘The Tasks of Amacrine Cells’. *Visual Neuroscience* 29.1 (2012), pp. 3–9. ISSN: 0952-5238, 1469-8714. DOI: [10.1017/S0952523811000344](https://doi.org/10.1017/S0952523811000344). (Visited on 26/11/2021).
- [86] E. J. Chichilnisky and R. S. Kalmar. ‘Temporal Resolution of Ensemble Visual Motion Signals in Primate Retina’. *Journal of Neuroscience* 23.17 (2003), pp. 6681–6689. ISSN: 0270-6474, 1529-2401. DOI: [10.1523/JNEUROSCI.23-17-06681.2003](https://doi.org/10.1523/JNEUROSCI.23-17-06681.2003). pmid: [12890760](https://pubmed.ncbi.nlm.nih.gov/12890760/). (Visited on 14/11/2022).
- [87] T. Baden, F. Esposti et al. ‘Spikes in Retinal Bipolar Cells Phase-Lock to Visual Stimuli with Millisecond Precision’. *Current Biology* 21.22 (2011), pp. 1859–1869. ISSN: 09609822. DOI: [10.1016/j.cub.2011.09.042](https://doi.org/10.1016/j.cub.2011.09.042). (Visited on 14/11/2022).
- [88] J. N. Tinsley, M. I. Molodtsov et al. ‘Direct Detection of a Single Photon by Humans’. *Nature Communications* 7.1 (2016), p. 12172. ISSN: 2041-1723. DOI: [10.1038/ncomms12172](https://doi.org/10.1038/ncomms12172). (Visited on 04/11/2021).

- [89] Y. Han, G. Roig et al. ‘Scale and Translation-Invariance for Novel Objects in Human Vision’. *Scientific Reports* 10.1 (2020), p. 1411. ISSN: 2045-2322. DOI: [10.1038/s41598-019-57261-6](https://doi.org/10.1038/s41598-019-57261-6). (Visited on 29/11/2021).
- [90] R. Pramod, H. Katti and S. Arun. ‘Human Peripheral Blur Is Optimal for Object Recognition’. *Vision Research* 200 (2022), p. 108083. ISSN: 00426989. DOI: [10.1016/j.visres.2022.108083](https://doi.org/10.1016/j.visres.2022.108083). (Visited on 31/07/2022).
- [91] H. B. Barlow and P. Földiák. ‘Adaptation and Decorrelation in the Cortex’. *The Computing Neuron*. Ed. by R. Durbin, C. Miall and G. Mitchison. Addison-Wesley, 1989, pp. 54–72. ISBN: 0-201-18348-X.
- [92] R. Shapley and C. Enroth-Cugell. ‘Visual Adaptation and Retinal Gain Controls’. *Progress in Retinal Research* 3 (1984), pp. 263–346. ISSN: 02784327. DOI: [10.1016/0278-4327\(84\)90011-7](https://doi.org/10.1016/0278-4327(84)90011-7). (Visited on 04/11/2021).
- [93] M. E. Burns and T. D. Lamb. ‘Visual Transduction by Rod and Cone Photoreceptors’. *The Visual Neurosciences*. Ed. by J. S. Werner and L. M. Chalupa. MIT Press, 2004, p. 19.
- [94] C. P. Ratliff, B. G. Borghuis et al. ‘Retina Is Structured to Process an Excess of Darkness in Natural Scenes’. *Proceedings of the National Academy of Sciences* 107.40 (2010), pp. 17368–17373. ISSN: 0027-8424, 1091-6490. DOI: [10.1073/pnas.1005846107](https://doi.org/10.1073/pnas.1005846107).
- [95] D. K. Freeman, G. Graña and C. L. Passaglia. ‘Retinal Ganglion Cell Adaptation to Small Luminance Fluctuations’. *Journal of Neurophysiology* 104.2 (2010), pp. 704–712. ISSN: 0022-3077, 1522-1598. DOI: [10.1152/jn.00767.2009](https://doi.org/10.1152/jn.00767.2009).
- [96] F. Rieke. ‘Temporal Contrast Adaptation in Salamander Bipolar Cells’. *The Journal of Neuroscience* 21.23 (2001), pp. 9445–9454. ISSN: 0270-6474, 1529-2401. DOI: [10.1523/JNEUROSCI.21-23-09445.2001](https://doi.org/10.1523/JNEUROSCI.21-23-09445.2001).
- [97] A. T. Rider, G. B. Henning and A. Stockman. ‘Light Adaptation Controls Visual Sensitivity by Adjusting the Speed and Gain of the Response to Light’. *PLOS ONE* 14.8 (2019). Ed. by S. Barnes, e0220358. ISSN: 1932-6203. DOI: [10.1371/journal.pone.0220358](https://doi.org/10.1371/journal.pone.0220358).
- [98] J. T. Pearson and D. Kerschensteiner. ‘Ambient Illumination Switches Contrast Preference of Specific Retinal Processing Streams’. *Journal of Neurophysiology* 114.1 (2015), pp. 540–550. ISSN: 0022-3077, 1522-1598. DOI: [10.1152/jn.00360.2015](https://doi.org/10.1152/jn.00360.2015).
- [99] F. Rieke and M. E. Rudd. ‘The Challenges Natural Images Pose for Visual Adaptation’. *Neuron* 64.5 (2009), pp. 605–616. ISSN: 08966273. DOI: [10.1016/j.neuron.2009.11.028](https://doi.org/10.1016/j.neuron.2009.11.028).

- [100] M. Meister and M. J. Berry. ‘The Neural Code of the Retina’. *Neuron* 22.3 (1999), pp. 435–450. DOI: [10.1016/S0896-6273\(00\)80700-X](https://doi.org/10.1016/S0896-6273(00)80700-X).
- [101] M. Van Wyk, H. Wässle and W. R. Taylor. ‘Receptive Field Properties of ON- and OFF-ganglion Cells in the Mouse Retina’. *Visual Neuroscience* 26.3 (2009), pp. 297–308. DOI: [10.1017/S0952523809990137](https://doi.org/10.1017/S0952523809990137).
- [102] H. B. Barlow, R. Fitzhugh and S. W. Kuffler. ‘Change of Organization in the Receptive Fields of the Cat’s Retina during Dark Adaptation’. *The Journal of Physiology* 137.3 (), pp. 338–354. DOI: [10.1113/jphysiol.1957.sp005817](https://doi.org/10.1113/jphysiol.1957.sp005817).
- [103] S. Martinez-Conde and S. L. Macknik. ‘Unchanging Visions: The Effects and Limitations of Ocular Stillness’. *Philosophical Transactions of the Royal Society B: Biological Sciences* 372.1718 (2017), p. 20160204. ISSN: 0962-8436, 1471-2970. DOI: [10.1098/rstb.2016.0204](https://doi.org/10.1098/rstb.2016.0204). (Visited on 21/10/2021).
- [104] R. G. Alexander and S. Martinez-Conde. ‘Fixational Eye Movements’. *Eye Movement Research*. Ed. by C. Klein and U. Erttinger. Springer International Publishing, 2019, pp. 73–115. ISBN: 978-3-030-20083-1. DOI: [10.1007/978-3-030-20085-5_3](https://doi.org/10.1007/978-3-030-20085-5_3). (Visited on 18/11/2022).
- [105] K. Boahen. ‘Retinomorph Vision Systems’. *Proceedings of Fifth International Conference on Microelectronics for Neural Networks*. Proceedings of Fifth International Conference on Microelectronics for Neural Networks. 1996. DOI: [10.1109/MNNFS.1996.493766](https://doi.org/10.1109/MNNFS.1996.493766).
- [106] T. R. Appleby and M. B. Manookin. ‘Selectivity to Approaching Motion in Retinal Inputs to the Dorsal Visual Pathway’. *eLife* 9 (2020), e51144. ISSN: 2050-084X. DOI: [10.7554/eLife.51144](https://doi.org/10.7554/eLife.51144). (Visited on 29/10/2022).
- [107] Y. LeCun, B. Boser et al. ‘Backpropagation Applied to Handwritten Zip Code Recognition’. *Neural Computation* 1.4 (1989), pp. 541–551. ISSN: 0899-7667, 1530-888X. DOI: [10.1162/neco.1989.1.4.541](https://doi.org/10.1162/neco.1989.1.4.541). (Visited on 19/11/2022).
- [108] A. Krizhevsky, I. Sutskever and G. E. Hinton. ‘ImageNet Classification with Deep Convolutional Neural Networks’. Conference on Neural Information Processing Systems. Ed. by P. L. Bartlett, F. C. N. Pereira et al. 2012.
- [109] I. Goodfellow, Y. Bengio and A. Courville. *Deep Learning*. MIT Press, 2016.
- [110] K. Chellapilla, S. Puri and P. Simard. ‘High Performance Convolutional Neural Networks for Document Processing’ (2006), p. 7.
- [111] *Perseverance rover landing on Mars recorded with onboard camera*. URL: <https://mars.nasa.gov/mars2020/multimedia/videos/?v=461> (visited on 23/11/2022).

- [112] C. Mead. ‘Adaptive Retina’. *Analog VLSI Implementation of Neural Systems*. Ed. by C. Mead and M. Ismail. Red. by J. Allen. Vol. 80. Springer US, 1989, pp. 239–246. ISBN: 978-1-4612-8905-0. DOI: [10.1007/978-1-4613-1639-8_10](https://doi.org/10.1007/978-1-4613-1639-8_10). (Visited on 17/11/2022).
- [113] C. Mead. ‘A Sensitive Electronic Photoreceptor’. *1985 Chapel Hill Conference on Very Large Scale Integration*. 1985, pp. 463–471. ISBN: 978-0-88175-103-1.
- [114] T. Delbrück and C. A. Mead. ‘An Electronic Photoreceptor Sensitive to Small Changes in Intensity’. *Proceedings of the 1st International Conference on Neural Information Processing Systems*. 1988.
- [115] M. A. Mahowald and C. Mead. ‘The Silicon Retina’. *SCIENTIFIC AMERICAN* (1991), p. 8.
- [116] K. A. Zaghoul and K. Boahen. ‘A Silicon Retina That Reproduces Signals in the Optic Nerve’. *Journal of Neural Engineering* 3.4 (2006), pp. 257–267. ISSN: 1741-2560, 1741-2552. DOI: [10.1088/1741-2560/3/4/002](https://doi.org/10.1088/1741-2560/3/4/002). (Visited on 16/11/2022).
- [117] T. Serrano-Gotarredona, A. Andreou and B. Linares-Barranco. ‘AER Image Filtering Architecture for Vision-Processing Systems’. *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications* 46.9 (1999), pp. 1064–1071. ISSN: 1558-1268. DOI: [10.1109/81.788808](https://doi.org/10.1109/81.788808).
- [118] R. Serrano-Gotarredona, T. Serrano-Gotarredona et al. ‘A Neuromorphic Cortical-Layer Microchip for Spike-Based Event Processing Vision Systems’. *IEEE Transactions on Circuits and Systems I: Regular Papers* 53.12 (2006), pp. 2548–2566. ISSN: 1057-7122. DOI: [10.1109/TCSI.2006.883843](https://doi.org/10.1109/TCSI.2006.883843). (Visited on 16/11/2022).
- [119] Y. Nozaki and T. Delbruck. ‘Temperature and Parasitic Photocurrent Effects in Dynamic Vision Sensors’. *IEEE Transactions on Electron Devices* 64.8 (2017), pp. 3239–3245. ISSN: 1557-9646. DOI: [10.1109/TED.2017.2717848](https://doi.org/10.1109/TED.2017.2717848).
- [120] T. Delbruck, C. Li et al. *Utility and Feasibility of a Center Surround Event Camera*. 26th February 2022. arXiv: [2202.13076](https://arxiv.org/abs/2202.13076). URL: <http://arxiv.org/abs/2202.13076> (visited on 29/10/2022).
- [121] E. Hajizada, P. Berggold et al. ‘Interactive Continual Learning for Robots: A Neuromorphic Approach’. *Proceedings of the International Conference on Neuromorphic Systems 2022*. ICONS: International Conference on Neuromorphic Systems. 2022. ISBN: 978-1-4503-9789-6. DOI: [10.1145/3546790.3546791](https://doi.org/10.1145/3546790.3546791). (Visited on 29/10/2022).
- [122] A. Gruel, D. Hareb et al. *Stakes of Neuromorphic Foveation: A Promising Future for Embedded Event Cameras*. preprint. In Review, 2022. DOI: [10.21203/rs.3.rs-2120721/v1](https://doi.org/10.21203/rs.3.rs-2120721/v1). (Visited on 18/11/2022).

- [123] Y. Hu, S.-C. Liu and T. Delbruck. ‘v2e: From Video Frames to Realistic DVS Events’. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. 2021. DOI: [10.1109/CVPRW53098.2021.00144](https://doi.org/10.1109/CVPRW53098.2021.00144).
- [124] H. Rebecq, D. Gehrig and D. Scaramuzza. ‘ESIM: An Open Event Camera Simulator’. *Conference on Robot Learning*. 2018.
- [125] M. G. McHarg, R. L. Balthazor et al. ‘Falcon Neuro: an event-based sensor on the International Space Station’. *Optical Engineering* 61.8 (2022), p. 085105.
- [126] O. Sikorski, D. Izzo and G. Meoni. ‘Event-based spacecraft landing using time-to-contact’. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. 2021. DOI: [10.1109/CVPRW53098.2021.00222](https://doi.org/10.1109/CVPRW53098.2021.00222).
- [127] S. McLeod, G. Meoni et al. ‘Globally Optimal Event-Based Divergence Estimation for Ventral Landing’. *2022 European Conference on Computer Vision (ECCV)*. 2022.
- [128] T.-J. Chin, S. Bagchi et al. ‘Star tracking using an event camera’. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*. 2019.
- [129] M. Jawaid, E. Elms et al. ‘Towards Bridging the Space Domain Gap for Satellite Pose Estimation using Event Sensing’. *arXiv preprint arXiv:2209.11945* (2022).
- [130] N. O. Ralph, A. Marcireau et al. ‘Astrometric Calibration and Source Characterisation of the Latest Generation Neuromorphic Event-based Cameras for Space Imaging’. *arXiv preprint arXiv:2211.09939* (2022).
- [131] S. Roffe, H. Akolkar et al. ‘Neutron-Induced, Single-Event Effects on Neuromorphic Event-Based Vision Sensor: A First Step and Tools to Space Applications’. *IEEE Access* 9 (2021), pp. 85748–85763.
- [132] D. Izzo and G. de Croon. ‘Landing with Time-to-Contact and Ventral Optic Flow Estimates’. *Journal of Guidance, Control, and Dynamics* 35.4 (2012), pp. 1362–1367. DOI: [10.2514/1.56598](https://doi.org/10.2514/1.56598).
- [133] H. W. Ho, G. C. de Croon and Q. Chu. ‘Distance and velocity estimation using optical flow from a monocular camera’. *International Journal of Micro Air Vehicles* 9.3 (2017), pp. 198–208. DOI: [10.1177/1756829317695566](https://doi.org/10.1177/1756829317695566).
- [134] *Planet and Asteroid Natural Scene Generation Utility product website*. URL: <https://pangu.software/> (visited on 14/11/2022).

- [135] G. Gallego, H. Rebecq and D. Scaramuzza. ‘A Unifying Contrast Maximization Framework for Event Cameras, with Applications to Motion, Depth, and Optical Flow Estimation’. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2018. DOI: [10.1109/CVPR.2018.00407](https://doi.org/10.1109/CVPR.2018.00407).
- [136] T. Stoffregen and L. Kleeman. ‘Event Cameras, Contrast Maximization and Reward Functions: An Analysis’. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2019. DOI: [10.1109/CVPR.2019.01258](https://doi.org/10.1109/CVPR.2019.01258).
- [137] N. C. Gupta and L. N. Kanal. ‘3-D motion estimation from motion field’. *Artificial Intelligence* 78.1-2 (1995), pp. 45–86.
- [138] H. C. Longuet-Higgins and K. Prazdny. ‘The interpretation of a moving retinal image’. *Proceedings of the Royal Society of London. Series B. Biological Sciences* 208.1173 (1980), pp. 385–397.

