



OPEN Photonic neuromorphic accelerators for event-based imaging flow cytometry

I. Tsilikas^{1,2,5}, A. Tsirigotis^{1,5}, G. Sarantoglou¹, S. Deligiannidis³, A. Bogris³, C. Posch⁴, G. Van den Branden⁴ & C. Mesaritakis¹✉

In this work, we present experimental results of a high-speed label-free imaging cytometry system that seamlessly merges the high-capturing rate and data sparsity of an event-based CMOS camera with lightweight photonic neuromorphic processing. This combination offers high classification accuracy and a massive reduction in the number of trainable parameters of the digital machine-learning back-end. The event-based camera is capable of capturing 1 Gevents/sec, where events correspond to pixel contrast changes, similar to the retina's ganglion cell function. The photonic neuromorphic accelerator is based on a hardware-friendly passive optical spectrum slicing technique that is able to extract meaningful features from the generated spike-trains using a purely analogue version of the convolutional operation. The experimental scenario comprises the discrimination of artificial polymethyl methacrylate calibrated beads, having different diameters, flowing at a mean speed of 0.1 m/sec. Classification accuracy, using only lightweight digital machine-learning schemes has topped at 98.2%. On the other hand, by experimentally pre-processing the raw spike data through the proposed photonic neuromorphic spectrum slicer at a rate of 3×10^6 images per second, we achieved an accuracy of 98.6%. This performance was accompanied by a reduction in the number of trainable parameters at the classification back-end by a factor ranging from 8 to 22, depending on the configuration of the digital neural network. These results confirm that neuromorphic sensing and neuromorphic computing can be efficiently merged to a unified bio-inspired system, offering a holistic enhancement in emerging bio-imaging applications.

Imaging flow cytometry (IFC) is the physical evolution of conventional flow cytometry (FC)¹ that strengthens the light-scatter recording capabilities of FC with accurate recording of the morphological features of the particles under investigation. IFC significantly enhances the detection capabilities of typical cytometers². Furthermore, visualization of spatial information adds an additional layer of qualitative information and allows the simultaneous recording of both brightfield and darkfield images^{3,4} without hindering typical fluorescent-based analysis. In this context, the stark difference between FC and IFC lies in the nature of the signals they capture. In FC, the signal consists of single-pixel time traces, acquired through photodiodes or photomultipliers, which provide only statistical information about the interaction between the particle and incident photons. In contrast, IFC employs a detailed 2D imaging scheme, enabling the detection of the particle's actual spatial features (see Fig. 1a). Figure Furthermore, the high throughput of IFC schemes minimizes the lab-to-diagnose time to few minutes, unlocks the use of ultra-low sample volumes, whereas more importantly allows the detection of very rare cell populations⁵, such as cancer cells in blood flow⁶, profile complex phenotypes or capture the associated dynamic of cell development phases^{1,7}. Exploiting these merits, IFC has started infiltrating a broad range of applications among which the most prominent are new drug discovery⁸, personalized medicine, DNA sequencing and rapid disease diagnostics^{3,9,10}. In principle, the above applications can be partially addressed by conventional high-throughput FC, but in most cases, a fluorescent agent is required that in turn can affect the molecule/cell's chemical or biological properties and thus can compromise the detection sensitivity or render it difficult to be applied. Furthermore, conventional FCs are mechanically complex and costly, demand high sample volume due to the fact that their output is based on statistics, mandate the use of elaborate microfluidic systems and last but not least can be operated only from trained personnel. This last aspect is of paramount

¹Department of Information and Communication Systems Engineering, University of the Aegean, Palama 2, 83100 Samos, Greece. ²Department of Physics, School of Applied Mathematical and Physical Sciences, Zografou Campus, 157 80 Athens, Greece. ³Department of Informatics and Computer Engineering, University of West Attica, Ag. Spyridonos, Egaleo, Greece. ⁴Prophesee Metavision, Rue du Faubourg Saint-Antoine 74, 75012 Paris, France. ⁵These authors contributed equally: I. Tsilikas and A. Tsirigotis. ✉email: cmesar@aegean.gr

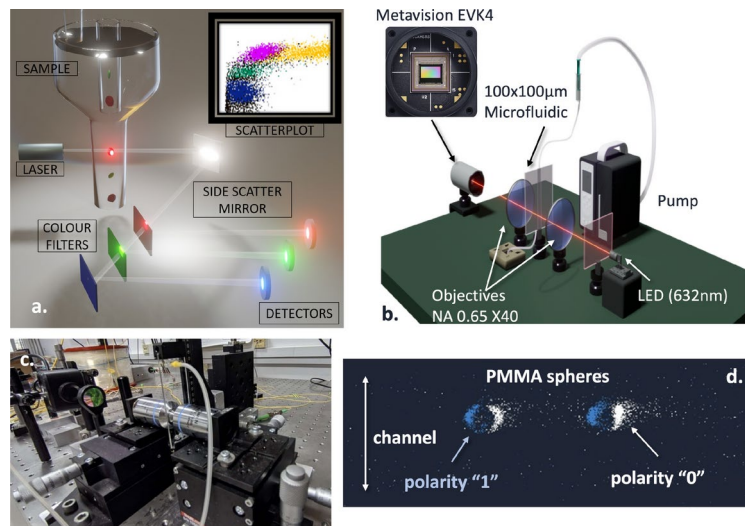


Fig. 1. (a) Basic concept of a typical multispectral FC, where statistics of scattered photons generate a particle signature. (b) Schematic of the proposed IFC: the neuromorphic camera used is a Prophesee Gen4 sensor¹⁷. (c) The installed IFC. (d) Typical recorded synthetic frame depicting 20 µm PMMA spheres flowing within the microfluidic channel: polarity “1” and “0” corresponds to pixel contrast increase and decrease respectively.

importance, as easy-to-handle home-care diagnostics are a rapidly growing necessity in modern societies with an aging population.

As a response, during the last years, a growing number of IFC modalities have emerged, that try to simultaneously address two contradicting requirements: namely high-speed operation and crisp spatial feature capturing. The contradiction is based on the fact that conventional capturing devices, experience blur effects, when high frame rate is requested, due to their limited bandwidth and high latency. Aiming to amend this, a wide pallet of technologies has been proposed, each offering vastly different performance metrics, in terms of the maximum particle flow that they can handle, footprint and wavelength range that they can operate. Broadly, IFC technologies can be categorized into two main categories: Single detector/pixel IFCs and 2-dimensional (2D) IFCs. The first category is dominated by time-stretched based approaches, where spatial information is projected in time and it is detected through a single high-speed photodiode. In this case, spatial resolution is linked to temporal resolution and thus it can proliferate by the high-speed capabilities of photonic technology^{11,12}, theoretically offering detection of particle flow $> 10^6$ particles/s. Following a similar concept, dual comb microscopy schemes, projects spatial information to the consecutive radio-frequency (RF) beating of two optical combs with different free spectral range^{13,14}. Despite the discrete merits of both techniques, they both mandate complex optical setups, costly optical sources, are restricted to the near-infrared (NIR) wavelengths, while both, in practice, generate massive amount of data during acquisition. In the 2D IFC category, time delay and integration (TDI) cameras allow motion-blur free images but at the cost of low particle-speed (< 3000 particle/s), this trade off stems from the necessity of high-integration time to unlock sufficient gain at the detector¹⁵. Alternatively, schemes based on strobe-photography can unlock high rates exceeding 50000 particles/s but at the cost of using sophisticated microfluidic schemes so as to enable precise control of the particle motion (trajectory and speed)¹⁶. Similarly, to the 1D case, 2D modalities also in typical cases generate an increased volume of data.

Recently, an alternative type of camera has been proposed for IFC, which follows a bio-inspired acquisition principle that relies on autonomous and asynchronously recording contrast detection events, instead of the camera periodically transmitting every pixel's intensity in the form of a frame¹⁷. In detail, event-based cameras (neuromorphic cameras) detect the contrast variation among adjacent pixels and transmit, in an asynchronous manner electrical spikes, similar to the retina-ganglion cells in mammals. Through this technique, time-continuous data transmission is limited to pixels detecting temporal contrast events in their field of view, thus massively reducing transmission bandwidth requirements, while offering wide dynamic range acquisition and radically increased capturing rate (i.e. temporal resolution) to IFC compatible levels^{2,18,19}. More importantly, these types of cameras rely on standard CMOS technology, thus are relative low-cost, have limited footprint and do not require complex optical systems or sophisticated control electronics. A critical difference in this IFC modality is that bio-inspired sparsity (spike encoding) can potentially reduce the volume of data generated during measurement.

Another critical step in IFC involves the data processing techniques employed after data capture. In this context, requirements vary depending on the specific IFC installation. For example, generic microfluidics permit multiple particles in the field of view, necessitating the use of object tracking algorithms²⁰. In addition, the acquisition speed and resolution of the generated signal (frame rate and pixel count), impose additional restriction on the size of the following neural network and on the associated processing latency. A plethora of neural network architectures have been proposed so as to efficiently tackle the aforementioned requirements, ranging from convolutional neural networks (CNNs) to recurrent neural networks (RNNs) and deep neural

networks (DNNs)^{21,22}. Each of these neural network implementations is adapted to the IFC, offering a diverse mix of advantages and disadvantages in terms of complexity, accuracy, and latency. However, critical features such as a limited number of floating-point operations (FLOPs), low latency, low power consumption, and FPGA compatibility with hardware-constrained digital processors (such as ASICs, FPGAs, or even generic platforms like Raspberry Pi) are essential for efficient, real-time, and portable IFC systems. In all cases, an increased data volume generated by the IFC modality leads to a higher parameter count in the neural network, which in turn results in unavoidably higher power consumption and latency during neural network training.

Based on the above discussion, this work has a twofold objective: first, we enhance classical IFC by incorporating a neuromorphic camera that enables blur-free operation while tracking fast-moving particles, thus reducing sample volume and analysis time. Second, although these event-based modules leverage bio-inspired data sparsity, they still tend to generate vast amounts of data due to their high temporal resolution. This increased data volume places stringent demands on the machine learning hardware responsible for analysis and feature extraction. Typically, the models employed (such as CNNs or RNNs) are large, multi-layer neural networks with a high number of trainable parameters. To address this challenge, we propose replacing the computationally intensive neural network back-end with an analog optical integrated circuit capable of efficiently extracting features from IFC outputs while compressing the data passed to the digital readout, all without sacrificing precision. This photonic pre-processing approach enabled the use of a significantly smaller and less complex machine learning model in the final stage of the setup.

In detail, we realize a low-cost, compact, light emitting diode (LED) based IFC scheme built around a CMOS neuromorphic camera, capable of capturing up to 1 Gevents/sec with μ s-range temporal resolution, delivering an equivalent of 100 kframes/sec¹⁷. The developed IFC module utilizes a generic off-the-shelf microfluidic chip with a single straight channel, without any sheath flow control mechanism. The IFC scenario realized was discriminating aqua solutions of Polymethyl methacrylate (PMMA) spheres of 12, 16 and 20 μ m in diameter. The generated electrical spikes (events) were processed in two distinct ways: the first method involved direct processing using solely digital machine-learning (ML) models, such as fully connected feedforward neural networks (FNNs) and compact RNNs. This direct method yielded high classification accuracy, up to 98.2 and 98.6% for the optimal FNN and RNN configurations, respectively. However, direct processing with FNNs required 30,000 trainable parameters, while RNNs demanded 1 million, significantly increasing hardware requirements and power consumption during training. The second approach adopts an unconventional path by transferring the IFC signals into the optical domain, where they are pre-processed through a photonic neuromorphic scheme utilizing an optical spectrum slicing (OSS) architecture²³. The pre-processed outputs are then also fed into a conventional digital FNN as above. In this scenario, the photonic neuromorphic pre-processor functions as an analogue CNN accelerator, resulting in an improved classification accuracy of 98.6%, surpassing the accuracy achieved when FNNs directly processed the generated data. Additionally, this approach resulted in a significant reduction in the number of trainable parameters for the digital FNN by a factor of > 20 . These experimental results validate that the combination of neuromorphic sensing and processing can enhance accuracy and more importantly generate a massive impact on the power consumption requirements of the overall IFC schemes.

Results

Neuromorphic camera based IFC

The experimental IFC setup is depicted in Fig. 1a–b. The light source is a simple LED emitting at 635 nm, chosen to simplify the experimental system while avoiding the generation of laser-source, diffraction patterns that could obscure the actual physical features of the particles. Two microscope objectives were utilized to focus and collect light into/from a generic microfluidic channel. The three classes of particles tested were PMMA spheres with diameters of 12, 16, and 20 μ m (refer to the Methods section for details). A steady flow was maintained in the channel with the aid of a vacuum pump, allowing the particles to reach mean velocities ranging from 1 to 0.07 m/s, corresponding to an ideal particle flow rate of 500 to 350 particles/s, assuming a sequential, uninterrupted flow of particles. This particle speed is consistent with previous IFC studies using event-based systems, resulting in comparable nominal particle flow². To better assess the experimental system, we computed the actual particle flow by analyzing all captured images and counting unique particle instances, yielding a flow rate of 10–15 particles per second. This lower actual flow rate was necessary to prevent clogging in the generic microfluidic setup, achieved by significantly diluting the solutions. It is important to note that this limitation does not affect the baseline accuracy and performance of the IFC, which is primarily dependent on the camera's temporal resolution relative to particle speed, rather than the dilution. The event-based camera (Prophesee Gen4 sensor) offered a spatial resolution of 640×480 pixels and a temporal resolution of 1 μ s per pixel, providing an effective capture rate of 1 Gevents/sec¹⁷.

The lightweight data processing pipeline is illustrated in Fig. 2. The Prophesee camera generates asynchronous events – tuples that include the coordinates, a timestamp, and a binary polarity signifying whether the pixel contrast is triggered by an intensity increase or decrease (X, Y, t, P). In our case, we sum all events per pixel (regardless of their polarity), over an integration time (T) to generate $I(X, Y)$. The choice of T depends on the speed of the particles and is linked to the number of events per frame; meaning that for a fast-moving sphere, a low value of T would result in too few events per synthetic frame, thus a lower signal-to-noise ratio (SNR). Conversely, a high value of T obscures the recording of the fine spatial features of the particles, similar to motion blur. To compute the optimum T for our setup, we recorded the total number of spiking events over time (see Fig. 2b). This measurement is performed once, and the particles' flow is represented by individual peaks over time. Figure 2b facilitates the extraction of two key observations. The first concerns the duration of each peak, which remains consistent across all recorded peaks, regardless of their amplitude, and corresponds to the time needed for each object to enter and exit the IFC's field of view (Fig. 2c). Therefore, the rise time of these peaks can provide an estimate of the optimum integration window (T) for the specific vacuum pump setting; in our

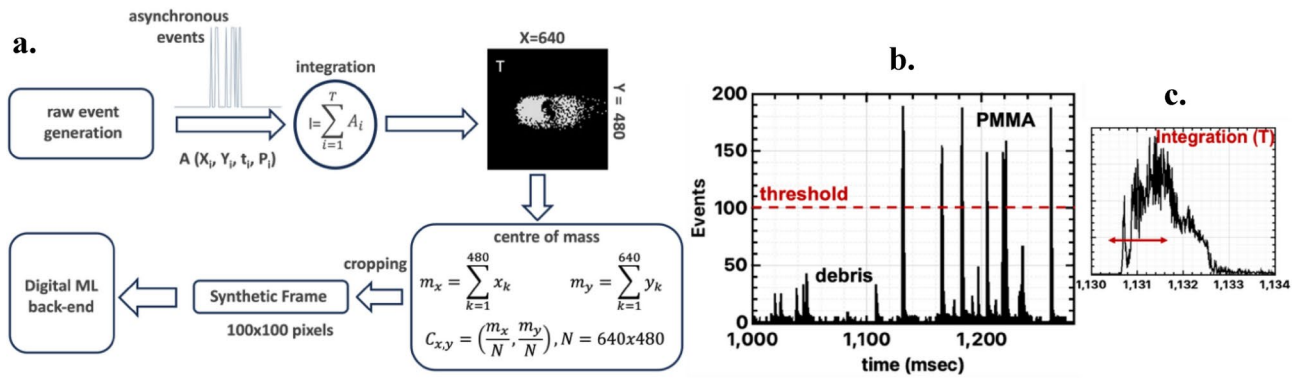


Fig. 2. (a) Data collection and lightweight pre-processing pipeline. (b) Total number of spikes across all 2D coordinates over time for a typical sample of 20 μm PMMA spheres. (c) Zoomed-in version of a single peak from part b.




Fig. 3. (a–c): Images of microfluidic channels from different experimental instances, showing variations in width (in pixels) due to minor misalignments. (d–f): Histograms of the distribution of particle widths, normalized to the width of the corresponding channel for each particle class: (d) PMMA spheres of 12 μm diameter, (e) 16 μm , and (f) 20 μm .

case, T was set to 3 ms (Fig. 2c). The second observation is that the peaks' amplitude (number of events) varies, with several peaks observed in Fig. 2b having a significantly low number of events (< 50). These peaks do not correspond to actual PMMA particles but are debris or air bubbles contaminating the PMMA mixture, evident in all samples. Therefore, a unified criterion of rejecting events with a low event count (< 100) has been used throughout all measurements.

Interestingly, the information in I follows a rate-encoding scheme, where pixels with a high spike count correspond to spatial locations subjected to strong activity (refer to Fig. 2a). Furthermore, considering we did not choose a sophisticated microfluidic channel; the particles do not flow in a single file but exhibit diverse trajectories. This led us to the use of an object tracking algorithm. To preserve the simplicity of the pre-processing stage, we utilized a straightforward centre of mass computation, involving only the summation of the synthetic frame's rows and columns (see Fig. 2a). By locating the centre of each particle, we cropped the dimensions of the synthetic frame to 100×100 pixels. Based on this pipeline, 4378 unique synthetic frames were captured/generated, not equally divided among the three classes: 1216 particles correspond to 20 μm spheres, 1811 to 16 μm , and 1351 to the 12 μm class. A critical issue in IFC data processing, as raised in²⁴, is related to whether experimental bias can influence the accuracy of ML schemes; specifically, if measurements of two particle classes are performed under different experimental conditions, then the ML's accuracy might reflect these experimental biases rather than the inherent differences between the two classes. To avoid this pitfall, we recorded data through multiple experimental instances (on different days, using different microfluidic channels-replaced due to clogging), ensuring multiple classes were recorded in each instance. In Fig. 3a–c, images from three typical microfluidic channels used in different instances are presented. It is evident that minor positioning

Metric	Max accuracy (%)	Mean accuracy (%)	Trainable parameters	FLOPS
Synthetic frames (100,100)	97.5	96.8 ± 0.7	30,003	60,001
Histogram (100)	90	88.6 ± 1.4	303	601
HoG (66,66)	98	97.5 ± 0.5	13,071	26,137

Table 1. Performance of a single layer feedforward neural network for raw events, histograms and HoG.

Metric	Max accuracy (%)	Mean accuracy (%)	Trainable parameters	FLOPS	Thermal design power %	Optimum nodes/function
Synthetic frames (100,100)	97.9/98.2	97.4 ± 0.5 / 97.9 ± 0.3	60,027 / 460,187	120,038/920,278	15.1/19.4	6, relu / 46, relu
HoG (66,66)	98.3/98.7	97.7 ± 0.6 / 98.3 ± 0.4	34,883 / 200,563	69,746 / 401,030	11.3/16.9	8, tanh / 46, tanh

Table 2. Performance of a 2-FCL neural network. The metrics in bold correspond to the most lightweight network in terms of parameters and the metrics in regular font to the optimum network in term of performance. In the last column the number of nodes at the hidden layer and the type of non-linear activation function used are presented.

errors among instances can affect the IFC magnification factor (width of channel in pixels). To address this, we normalized each synthetic frame's spatial features (measured in pixels) to the width of each microfluidic channel used. In Fig. 3d–f, normalized particle size histograms for the three classes are presented; each class exhibits a distinguishable mean value ($m_{d=20\mu m} = 0.25 \pm 0.04$, $m_{d=16\mu m} = 0.17 \pm 0.02$ and $m_{d=12\mu m} = 0.14 \pm 0.016$), while significant overlap between the three distributions is observed. This implies that even after data normalization, the classification task remains challenging, necessitating an ML model.

Digital machine learning IFC

Initially, we focused on lightweight ML that can be implemented on hardware-constrained platforms, such as ASICs or generic systems (e.g., Raspberry Pi, Arduino). A subset of the dataset was randomly chosen to equalize the number of samples per class. Therefore, 1216 unique samples per class were used, of which 70% were used for the training of each model and 30% for testing. All ML models were implemented on the TensorFlow framework²⁵ by a graphics processing unit (GPU) alongside an Adam Optimizer²⁶ with a learning rate of 0.01, and categorical cross-entropy as the loss function. The model training process involved data processed in batches of 250 for a maximum of 500 epochs. To prevent overfitting, an early-stopping callback was implemented to monitor validation loss, maintaining a validation split of 10% during training (see Methods). The architectures used here include a fully connected FNN with 1 and 2 layers. In Table 1, we present the accuracy and the number of trainable parameters for the lightest neural network configuration employed: a single layer consisting of 3 nodes, each equipped with a rectified linear unit (ReLU) activation function. The optimal performance without digital pre-processing was achieved when synthetic frames were used directly as input, yielding an accuracy of 97.5% and requiring 60,000 FLOPs per forward pass. To further enhance performance, we employed a digital pre-processing technique tailored for image processing, specifically the Histogram of Gradients (HoG)²⁷. This approach resulted in a slight accuracy improvement to 98%, while reducing the FLOPs per forward pass to 26,000. However, these benefits come with the subtle cost of additional digital pre-processing of the original data. In particular, HoG requires the computation of both the magnitude and angle of each pixel's contrast relative to its neighboring pixels, row- and column-wise. This process results in a minimum of 105 additional computations during pre-processing²⁷. It is also important to note that the intensity histograms of the synthetic frames (used to generate Fig. 3) were used as input to the same neural network, serving as a basic benchmark for accuracy and complexity. Although histograms require very few parameters (303), their performance is limited to a maximum accuracy of 90%, underscoring the complexity of the IFC classification task. This lower accuracy can be attributed to the fact that the histogram data captures only size (diameter) information, leading to poorer performance compared to the use of synthetic frames. The difference in performance may be due to debris present in the samples, which can be similar in size to the target particles and lead to misclassification when size is the only feature considered. By contrast, when a 2D representation is used, the morphological details of the particles allow the machine learning model to filter out debris more effectively.

Aiming to further boost performance, we use a feedforward neural network with an additional dense layer. In Table 2, we present performance metrics including accuracy, number of trainable parameters and power consumption, using the thermal design power (TDP) for each network. For each metric, we compute two values: one for the ML model with the fewest parameters (highlighted in bold) and another for the model providing the best performance in terms of classification accuracy (indicated in regular font). A compact network with two fully connected layers (FCLs) and only six hidden nodes achieved a precision of 97.9%. However, increasing the number of hidden nodes to 46 yields only a minor improvement of 0.3%, with an accuracy of 98.2%. Similar to the single FCL classification case, the highest accuracy is achieved through HoG pre-processing (98.3%), which halves the number of trainable parameters but requires more complex data pre-processing, as mentioned above. The best overall performance, a 98.7% accuracy, is again achieved with HoG and 2 FCLs. Interestingly, the accuracy boost when using a two-layer FCL network over a single-layer perceptron is marginal. For instance,

the inference score increased from 97.5 to 98.2% for synthetic frames used as direct input, but this came with a substantial rise in FLOPs per forward pass, from 30 to 900 K.

Photonic neuromorphic pre-processor for IFC

Numerical simulations

As shown in the results above, input pre-processing (e.g., HoG) is a potential path towards further increasing accuracy, without massively affecting the number of trainable parameters. On the other hand, digital pre-processing introduces additional computations, thus impacting power consumption during inference. Aiming to circumvent this impediment, we pursue an alternative approach, where input signals are pre-processed directly in the optical/analogue domain. In particular, we utilize a neuromorphic photonic scheme that will alleviate the additional computational workload associated with digital feature extraction techniques^{28,29}, while at the same time, off-loading part of the back-end computational complexity, allowing for slimmer digital back-end ML models. In this context, our group proposed a photonic accelerator, relying on a hardware-friendly optical spectrum slicing (OSS) technique²³. This approach involves the utilization of multiple passive optical filters acting in parallel as convolutional neural nodes. Each OSS node monitors distinct spectral regions of the input optical signal and applies a complex kernel filter directly in the analogue domain without any need for digital processing. This scheme was successfully applied to the field of high-speed image processing, where the OSS has offered performance levels comparable to fully digital sophisticated architectures²³. More importantly, OSS, when targeting datasets such as the MNIST, offered a significant reduction in the number of trainable parameters and consequently to the overall power consumption. The fact that recurrent versions of the OSS have been also tested in transmission impairment mitigation in high baud rate optical communications further solidifies the capabilities of this approach as a ML accelerator³⁰. On the other hand, OSS has been only recently combined with real-life datasets and, in particular, with medical imaging modalities such as IFC²¹. Here, we aim to combine the efficiency of a neuromorphic camera with photonic neuromorphic pre-processing for the first time, so as to enhance the IFC's capabilities overall.

Before implementing the OSS scheme, synthetic frames are converted into 1D vectors, as depicted in Fig. 4. This conversion involves serializing each block of pixel values (patch) in two orientations—row-wise and column-wise. Both orientations are used in the 1D representation of each synthetic frame to enhance their spectro-temporal characteristics (see Methods section). The key components of the OSS scheme are illustrated in Fig. 5. The 1D vectors of each synthetic frame are transferred to the optical domain by modulating the amplitude of a continuous-wave optical carrier through a digital-to-analog converter (DAC) and an electro-optic modulator, such as a Mach–Zehnder Modulator (MZM). The processing core comprises multiple bandpass optical filters. Operation wise, the application of the filter's transfer function in the frequency domain is identical to the convolution of the signal with the filter's impulse response in the temporal domain. Therefore, each OSS node is set at a different central frequency, facilitating the “slicing” of distinct regions of the optical signal. This results in a change in the impulse response of each filter and the application of varying complex weights to the input time-traces²³. Unlike digital kernels, the control over the complex weights is coarsely adjusted through the filter's hyperparameters, which shape the impulse response of the OSS filters, such as central frequency (f_m), bandwidth (f_c) etc. (red and green insets of Fig. 5). Furthermore, the filters can act as tunable optical integrators³¹, where the integration time is governed through tuning the filters' bandwidth and order. The integration time in this scenario is equivalent to the receptive field; meaning the number of spatial pixels' values that are linked during convolution.

Following this purely optical step, the time-traces are transferred back to the electrical domain through a photodiode (PD) and an analog-to-digital converter (ADC) that follows each filter. The PDs introduce an elementwise nonlinear transformation at the filter outputs through their square-law characteristics. Furthermore, by reducing the bandwidth of the PDs or equivalently by integrating the PDs' output, an average-pooling-like operation is performed. The extent of this integration can be controlled through Eq. 1, where the PD's 3 dB bandwidth is set according to the pixel rate (modulation rate at the MZM) of the input signal (PR) and the size of the patch. This is defined by dividing the synthetic frame into square blocks of $M \times M$ pixels.

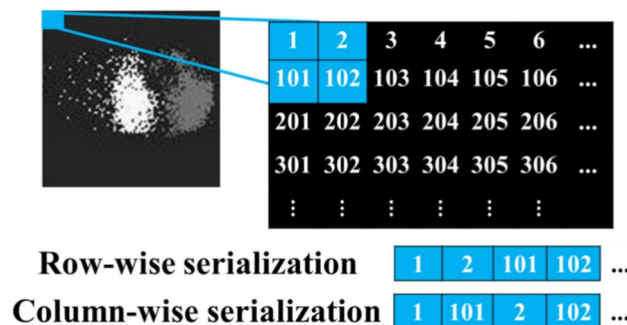


Fig. 4. Conversion of synthetic frames to 1D vectors for OSS processing. Each frame is serialized into a 1D vector by processing blocks of pixel values (patches) in two orientations: row-wise and column-wise (see Methods).

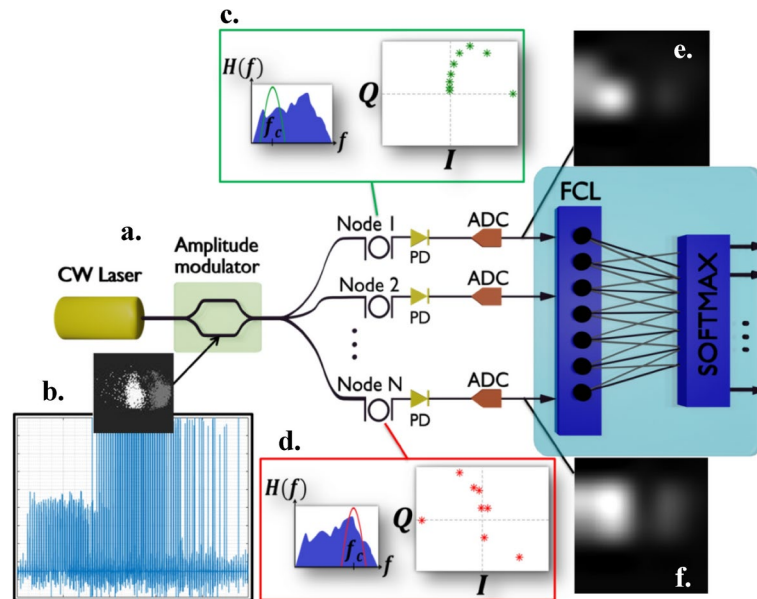


Fig. 5. Schematic diagram of the OSS-CNN for the classification of IFC experimental data. **(a)** The synergy of a continuous-wave laser, a DAC and an MZM, used to imprint the 1D synthetic frame vectors onto the amplitude of the optical carrier. **(b)** A flattened IFC synthetic frame after laser modulation, displayed alongside the original synthetic frame. **(c-d)** Complex-valued coefficients (I/Q) for two identical filters, positioned at different detunings relative to the signal's carrier frequency. **(e-f)** Reconstructed images after processing through two discrete OSS nodes.

$$BW_{PD} = \frac{PR}{M \times M} \quad (1)$$

Towards the same direction, the sampling rate (SR) of the ADCs can be used to further compress data by under-sampling the output. The digitized samples at the output of the OSS are subsequently flattened and fed to a simple digital back-end comprising a lightweight digital FCL with 1 or 2 dense layers, identical to the ones presented above. It is worth mentioning that the number of nodes in the FCL and thus the trainable parameters are directly governed by the number of samples that are generated and fed to the back-end.

Towards validating the concept, the first step consists of numerically simulating the OSS scheme so as to determine the optimum hyperparameters that would maximize classification accuracy. In this case, the OSS nodes are implemented as first-order bandpass Butterworth filters that can be easily realized using conventional photonic structures (e.g. micro-ring resonators-MRRs). The system's hyperparameters under examination included the patch size ($M \times M$), stride, number of OSS nodes and OSS filter characteristics (f_c, f_m). Initially, optical input was assumed to be injected to the filters at a pixel rate of 100 GS/s, with an average input power of 20 mW; a typical optical spectrum and time-trace of a particle are shown in Fig 6 b,a, respectively. Subsequently, the optical signal was split equally to feed the OSS nodes with a $1 \times N$ splitter, where nodes range from 1 to 10. In Fig.6c, three typical filter responses are depicted at different frequencies, "slicing" the spectrum of the input signal. The bandwidth of the PDs was adjusted in accordance with the selected patch size, as defined in Eq. 1, while the sampling rate of the ADCs was modified to extract only two samples over the temporal duration of a single patch.

The generated digital samples were fed to a single FCL identical to the digital scenario described above. The key differentiation in this approach lies in the preliminary optimization phase. Prior to classification, we employed the 'Optuna'³² hyperparameter optimization framework to systematically identify the optimal values for the OSS scheme in this task (bandwidth, central frequency and patch size).

The hyperparameter scan revealed that using a 20×20 pixel patch with a 10-pixel stride, during the conversion of 2D image data to a 1D format for OSS processing, yielded optimal results. Figure 7 graphically illustrates the mean performance of two OSS-CNN configurations, one using a single-layer and the other a two-layer FNN as the classification backend, against the number of OSS nodes utilized. For context, the standalone performances of both single and dual-layer digital FCL models are included for direct comparison, depicted by red and blue dashed lines, respectively. In particular, utilizing a single OSS node, with a filter 3 dB bandwidth (f_c) of 11 GHz and a central frequency detuning (f_m) of 19 GHz, testing accuracy reached as high as 96.4%. Adding just one more OSS node—adjusting f_m to 16 and 37 GHz for each filter, with f_c set to 8 GHz—enhanced maximum classification accuracy to 97.9%, slightly surpassing the standalone single-layer FCL model's performance by 0.4% (see Table 1). A similar accuracy level was also achieved in the cases that node number increased to 3–5. Further exploration with a two-layer FCL classifier, as indicated by the Optuna framework, included 80 nodes in the digital hidden layer with a ReLU activation. This setup showed a notable improvement, where the single-

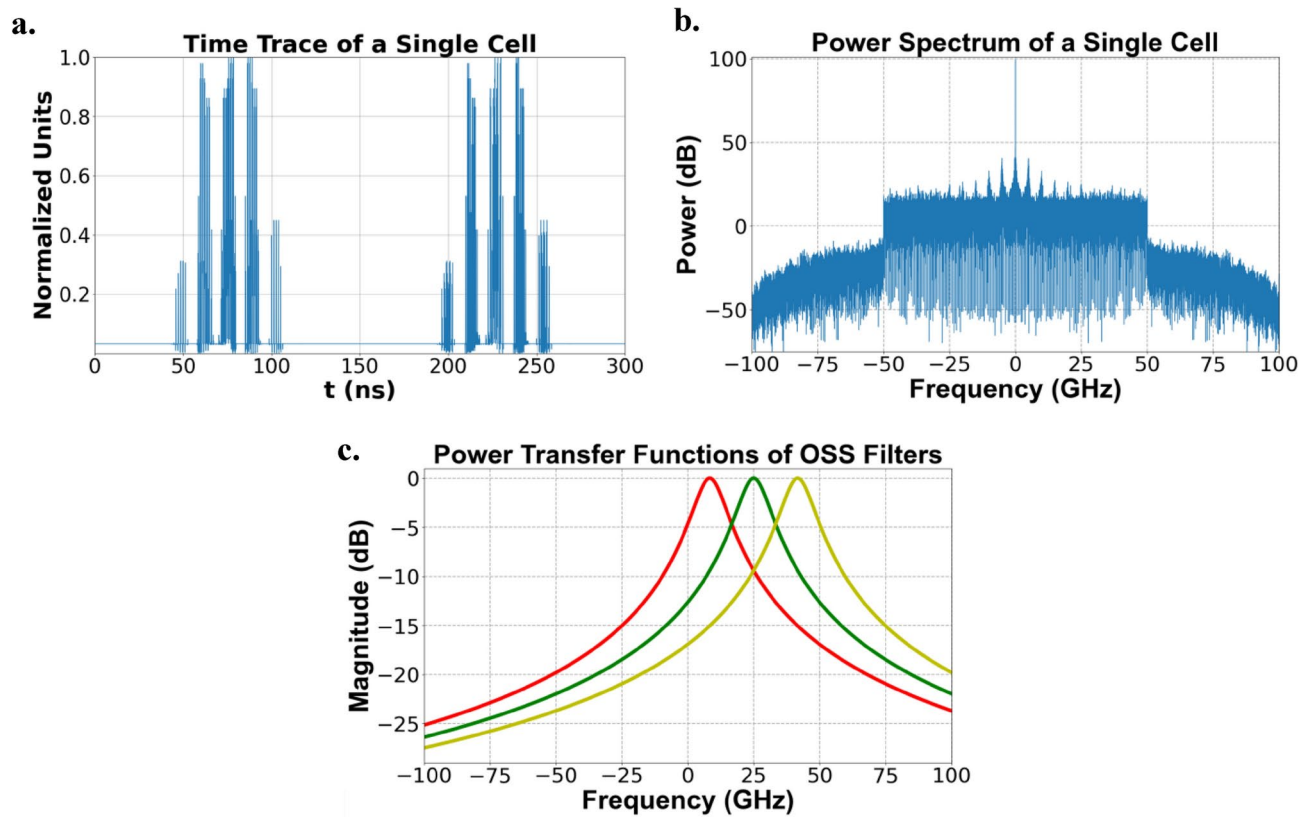


Fig. 6. (a) Time trace at the output of the DAC corresponding to the 1-D vector of a 12 μm cell, (b) Optical power spectrum of the signal modulated by the vector of the 12 μm cell and (c) Power transfer functions of three OSS filters with a 12 GHz 3-dB bandwidth, designed for uniform segmentation of the right-sideband spectrum within the input optical signal.




Fig. 7. Mean classification accuracy of a numerically simulated OSS-CNN with one (red circles) and two FCLs (blue circles) across varying OSS nodes, compared to standalone one-layer (red dashed line) and two-layer FCL models (blue dashed line).

node OSS scheme's accuracy boosted to 98.5%, compared to 98.2% for the purely digital approach. The highest accuracy achieved with the OSS-assisted scheme was 98.6% with two optical nodes, though an increase in the number of OSS nodes beyond three resulted in a slight decline in performance, possibly due to overfitting. Overfitting here arises when more than three OSS nodes are used, due to increase in the model's parameters.

From a first glance, the inclusion of a fully analogue accelerator in this case, does not provide a strong performance boost as in²³ but still improves testing accuracy by 0.5%. However, the key advantage of integrating OSS lies in its impact on reducing the number of trainable parameters, a critical factor in minimizing power consumption during model training. Table 3 provides a detailed comparison of OSS configurations (with one and two FCL layers) against purely digital neural network models, including the standalone single-layer and double-layer feedforward neural networks described above and a sophisticated yet parameter-efficient gated recurrent unit RNN (GRU-RNN). It highlights the differences in classification accuracy, the number of trainable parameters and the number of hidden units for each architecture.

Architectures	Max accuracy (%)	Mean accuracy (%)	Hiddenunits	Parameters
1 FCL	97.5	96.8 ± 0.7	–	30,003
2 FCLs	98.2	97.9 ± 0.3	46	460,187
OSS-1FCL	97.9	97.6 ± 0.3	–	1947
OSS-2FCL	98.6	97.8 ± 0.6	80	52,163
GRU-RNN	98.6	97.9 ± 0.7	38	1,144,677

Table 3. Comparison of OSS-CNN with digital neural network architectures in terms of accuracy and trainable parameters.

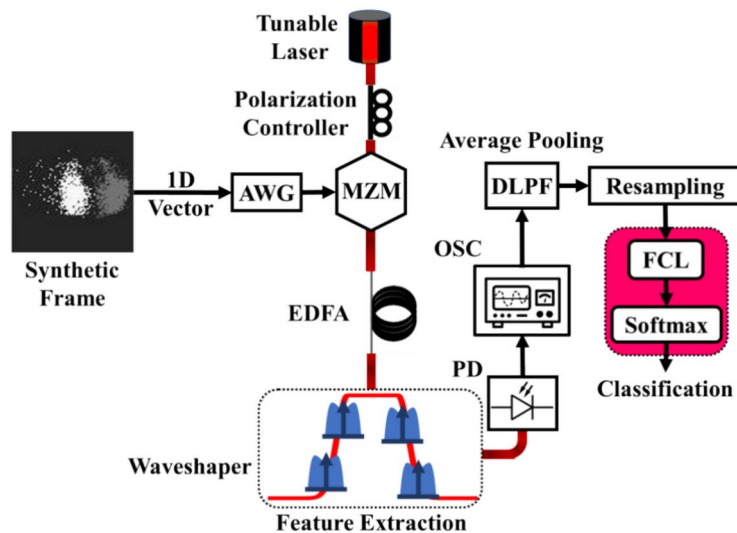


Fig. 8. Schematic diagram of the OSS-CNN experimental setup. The setup includes a tunable laser, a 20 GS/s AWG, and an MZM. The signal is then boosted by an EDFA to enhance SNR and sent to a waveshaper. Detection and analysis are conducted with a 50 GS/s OSC. Post-processing is performed using a DLPF and classification via a FCL.

The results of Table 3 underscore that the OSS was able to achieve a classification precision higher compared to simple FCLs and identical to the more complex GRU-RNN architectures. This is particularly noteworthy since the OSS-2FCL configuration operates with significantly fewer trainable parameters, approximately 8, 15 times lower compared to 1-FCL and 2-FCL respectively, whereas versus the complex RNN, OSS offers a strong parameter reduction by a factor of 22, without any performance degradation. This key property, similar to²³, can be attributed to the CNN-like architecture of the OSS system. Specifically, the slicing process by multiple filters performs a convolutional operation, where the kernel is roughly dictated by the transfer function and position of each filter's response. This process extracts multiple diverse features from the original data, which are also integrated by the filter's integrational property. Additionally, these rich, information-wise features undergo a nonlinear transformation through the PD, which is also set to average these values over the temporal duration of a patch (see Eq. 1). Consequently, the number of samples delivered to the input of the back-end classifier for each synthetic frame is significantly reduced, corresponding to two samples per patch as determined by the sampling rate of the post-nodal ADCs.

Experimental validation of OSS assisted IFC

The aforementioned concept was experimentally validated, as depicted in the schematic diagram of Fig. 8. The setup included a tunable laser (CoBrite-DX Tunable Laser) amplitude-modulated through a 20 GS/s arbitrary waveform generator (AWG) (Tektronix AWG70002B) and a Mach-Zehnder modulator (MZM) (iXblue MXIQ-LN-30) with a 20 GHz 3-dB bandwidth. Subsequently, the modulated signal was directed to an optical waveshaper (Coherent WaveShaper-1000A Programmable Optical Filter), acting as a Butterworth filter with a 3-dB bandwidth around 10 GHz, centred at 1552.6 nm, realizing a single OSS node. It is worth mentioning that this experimental setup aims at a proof of concept, thus near-future implementations could rely on optimized components such as silicon photonic filters and high-speed, miniaturized modulators. The optical output was captured by a 10 GHz photoreceiver (Thorlabs RXM10AF) connected to a 50 GS/s digital signal oscilloscope (OSC) (Tektronix DPO75002SX). The final digital signal was resampled to 20GS/s to match the rate of the AWG and was filtered using a digital low-pass filter to implement an operation akin to average pooling. The final step involves classification using a digital dense feedforward neural network, identical to the one used in the numerical simulations described above.

The data pre-processing pipeline was slightly modified compared to the numerical simulations, aiming to reduce the number of samples processed by the AWG due to time constraints. In this approach, we took the original 100×100 pixel frames and divided them into smaller, 4×4 pixel blocks. By calculating the average value of each 4×4 block, we effectively compressed the frame into a smaller, 25×25 pixel frame, reducing the overall data. Therefore, the compressed frames were then transformed into 1D vectors using 5×5 patches and a stride of 2. To increase the SNR of the optical signal and assess the system's capabilities when not limited by thermal/shot noise during detection, an Erbium-Doped Fiber Amplifier (EDFA) was used prior to the waveshaper to boost the optical power to 8 dBm. To simplify the experimental setup and facilitate the recording of a single trace at a time, the optical convolution process at the waveshaper was performed sequentially. Specifically, the filter response of the waveshaper was maintained constant at 193.22 THz (1552.6 nm), while to implement multiple OSS nodes, the central frequency of the laser source was detuned relative to the central frequency of the filter. The carrier frequency was adjusted in 1 GHz steps, ranging from -9 GHz to $+9$ GHz, allowing for the exploration of 19 distinct filter central frequencies. Figure 9a shows a segment of a normalized time trace corresponding to a $12 \mu\text{m}$ particle as sent to the AWG. In Fig. 9b, the resulting outputs for two different filter-carrier detunings (-6 and $+1$ GHz), as captured by the oscilloscope, are presented. The variation in frequency detuning led to distinct outputs, highlighting the optical filter's differential interaction with the signal's spectral components.

After photodetection, the signal from the oscilloscope was resampled at a rate of 20 GS/s to align with the AWG's output rate. The number of consecutive samples subjected to averaging was then adjusted based on the bandwidth of the following LPF, as per Eq. 1, leading to varying degrees of compression at the output of the OSS system. The compression ratio (CR) serves as a metric that demonstrates the level of compression, compared to the input, applied at the OSS accelerator's output. It is defined as the ratio of the initial data size to the size of the digital outputs of the OSS nodes per dataset sample. For this experimental setup, the compression ratio can be calculated using the formula:

$$CR = \frac{N}{2N_f} \quad (2)$$

where N_f is the number of distinct filter positions contributing to the inputs for the FCL and N is the integer denoting the number of successive samples involved in the averaging process. The averaged outputs from all particle samples were then merged to construct the optical pre-processed version of the original dataset.

The classification process began with analyzing the digitized outputs from each filter position, thereby examining the performance when a single OSS node was used for pre-processing under various averaging scenarios. Subsequently, the back-end FCL was supplied with combinations of outputs from 2, 3, and 5 different OSS nodes, thereby implementing an expanded OSS scheme. Figure 10 illustrates the mean classification accuracy for the optimal combination of two distinct filter positions (optimum found to be -1 and 0 GHz from the carrier) with a single-layer FCL as the digital classifier. This is compared alongside the performance of the optimal simulated 2-node OSS (solid-blue line), in relation to the CR. Additionally, Fig. 10 presents the accuracy of the standalone digital FCL processing the entire synthetic frame dataset (black dashed line) and the accuracy from a time trace derived from an uncompressed synthetic frame without any OSS node intervention (red dashed line). Here, "mean accuracy" refers to the average classification accuracy obtained over ten iterations of the same 3-neuron FCL model, where the only difference across iterations was the initial values of the weights, with all the other training parameters remaining constant. From Fig. 10, it is evident that the performance of the experimental OSS setup is similar to that of the simulated OSS. The highest accuracy recorded in the experimental setup was 98.1% at a CR of 3.3, underscoring the efficacy of OSS by enhancing accuracy by 0.6% over a system without noise, while concurrently reducing the number of trainable parameters by a factor of 3.3. Another significant observation from Fig. 10 is the robustness of OSS pre-processing, which sustains nearly 96% accuracy even under substantial compression (96.1% with a CR of 20).




Fig. 9. (a) Temporal segment of the normalized event-vector time trace from the AWG corresponding to a $12 \mu\text{m}$ cell. (b) Normalized outputs for two distinct filter detunings from the carrier frequency at the OSC, specifically at -6 GHz (red) and $+1$ GHz (green), applied to the input signal depicted in (a).




Fig. 10. Mean classification accuracy for the optimal 2-node experimental OSS configuration, utilizing filters detuned by -1 and 0 GHz, paired with a 1-FCL classifier, in comparison to a simulated 2-node OSS as a function of the CR. For reference, the accuracy of a standalone FCL is shown (black dashed lines), alongside the mean accuracy from the uncompressed output of a wideband (25 GHz) filter aligned with the carrier frequency (red dashed lines) are also presented.

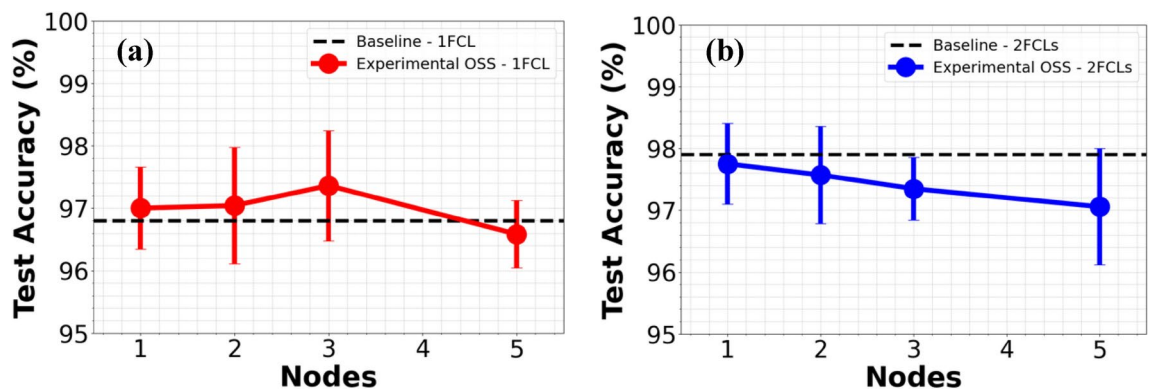


Fig. 11. Mean classification accuracy of the experimental OSS system as a function of the number of OSS nodes, using (a) a single fully connected layer (1-FCL) and (b) a two-layer fully connected network (2-FCL) as the digital back-end. For comparison, the mean performance of the digital FNNs with one and two layers, which directly processed the synthetic frame dataset, are also included.

In Fig. 11, the mean precision of the experimental OSS-CNN is plotted against the number of OSS nodes (filter positions), considering a single-layer (Fig. 11a) and a two-layer FNN back-end (Fig. 11b), with the CR held constant at 3.3. For comparison, the figure also includes the mean classification accuracy of the digital single-layer and two-layer FNNs that directly processed the synthetic frames. In Fig. 11a, the accuracy of the experimental setup is comparable to, or higher than, that of the standalone digital FCL, especially when three OSS nodes are used. Notably, the experimental setup achieves its highest accuracy of 98.1% with three nodes, outperforming the single-layer digital FNN despite the presence of noise from components such as the photodetector and the EDFA. Furthermore, the experimental OSS system demonstrates strong robustness, consistently maintaining a precision above 97% with 1, 2, and 3 OSS nodes. Figure 11b shows that adding a second dense layer in the digital back-end improves accuracy across all cases relative to Fig. 11a. The peak accuracy for the experimental OSS reaches 98.4% with 1–2 OSS nodes, slightly below the 98.6% achieved in the simulated OSS. The overall performance of the experimental setup declines slightly across different node configurations, indicating the influence of experimental factors, particularly in the 2-node case where mean accuracy drops to 97.6% compared to 98.6% in the simulated system. Interestingly, in configurations with 1–2 OSS nodes, the experimental OSS outperforms the dual-layer digital FCL while significantly reducing the number of trainable parameters.

Discussion and conclusion

In this work, experimental data from a neuromorphic event based IFC system have been recorded, regarding the classification of aqueous solution of PMMA particles with different diameter. The generated datasets comprise

rate-encoded synthetic frames subject to a lightweight pre-processing, mainly consisting of object tracking, to compensate for the utilization of a generic micro-fluidic system. The core of this work comprises two different innovative steps, the utilization of an event-based detector for IFC and unconventional data processing. Regarding the first step, event-based sensors have just started infiltrating the IFC landscape, offering rapid visualization, low data-rates and reduced cost compared to conventional FC, even at this early development phase³³.

Regarding the second step, classification of the IFC data has been performed in three ways: lightweight digital feedforward neural networks, offering a classification accuracy of 98.2% comparable to², but requiring more than 460 K trainable parameters. The second approach consisted of using digital pre-processing of the dataset, following the well-established HoG algorithm; topping accuracy at 98.7% with 200 K parameters during training and with the additional digital processing cost during inference. The third route replaces the HoG algorithm with an all-optical convolution-like photonic neuromorphic scheme, based on optical spectrum slicing; this route offloads the pre-processing to the analogue domain while the same feedforward neural networks (FNNs) are used as back-end. Although an all-optical FNN scheme could be considered, we utilize a hybrid electro-optic approach so as to circumvent issues such as the hard-to-scale optical networks and the absence of an efficient optical training approach.

In this work this concept has been validated both experimentally and numerically and it offered an optimum accuracy of 98.6%. The most critical attribute of the proposed scheme is that this performance is achieved with only a fraction of trainable parameters demanded by standalone digital neural networks and without any type of computationally intensive digital preprocessing. Thus, reducing power consumption both during inference and training. These improvements stem from the ability of the OSS to inherently correlate pixel information and apply random, yet complex kernels directly in the analogue domain. This in turn, evokes a coarse projection of image's features to a higher dimensional space, like a reservoir computing scheme, thus simplifies classification by a lightweight digital neural network. This combination of neuromorphic sensing and neuromorphic computing has been presented also in the recent past, where in² a similar IFC generated spiking data that were classified through digital pre-processing (feature extraction) and a software implemented spiking neural network, offering similar accuracy to this work. More recently in¹⁸ the same IFC as in², was merged with a spiking neural network realized in hardware (Intel's Loihi) so as to classify a two-class PMMA particle scenario. The key difference between our work and¹⁸ is the fact that our optimization efforts aim towards applying neuromorphic-analogue pre-processing so as to reduce the trainable parameters of a conventional (floating-point based) back-end, where in¹⁸ the aim is to exploit spiking data sparsity and generate an all-spiking system. The reasoning for our approach is that conventional digital back-end is more mature and efficient compared to experimental spiking processors, whereas independently of the architecture our approach offers a comparable smaller parameter space to train of 1.9 K parameters compared to 394 K parameters derived in¹⁸ with similar accuracy. On the other hand, the proposed IFC entails an additional electro-optic conversion step that is not present in standalone digital solutions or in all-spiking approaches. This conversion step is a necessity in all photonic accelerators^{34,35}, whereas for the explicitly OSS it has been shown that the power consumption of the electro-optics is lower compared to the energy-footprint of the digital back-end. In addition, this step can be omitted if an inherently photonic IFC detector is employed as in¹². Furthermore, in our work we generate 2D synthetic-frames from raw data so as to drive our photonic accelerator, which is more computation-hungry compared to simple down-sampling of raw spiking data¹⁸.

Overall, the results in this work alongside similar works curving complementary strategies confirm that by matching neuromorphic sensing with neuromorphic processing an overall performance enhancement can be achieved outperforming all previous schemes, whereas offering a strong reduction in terms of trainable parameters by a factor > 20 that is of utmost importance for emerging machine learning modalities.

Methods

Experimental setup

The experimental setup features two 40× microscope objective (RMS40X- 40X Olympus Plan Achromat Objective, 0.65 NA, 0.6 mm) lenses with a numerical aperture (NA) of 0.65 and a working distance of 600 μm. These lenses are securely mounted on two 3-Axis MicroBlock. Two planoconvex lenses with a focal length of 10 cm are used before and after the objectives so as to direct/collect light from the microscope system. The light source is a 635 nm emitting LED with average power of 5mW, whereas the recording event-based camera is Gen4 provided by Prophesee with a resolution of 640×480 pixels and a temporal resolution in spike generation of 1μsec. The camera records pixel's contract changes with two polarity values, depending on whether there is an increase or decrease in intensity. In this work, synthetic frames do not include this feature and all events are mapped to the synthetic frame, independently of their polarity. A vacuum pump installed offered a steady liquid flow ranging from 10 μlit/hour to several ml/min thus, regulating the speed of the particles from 0.001 m/sec to well beyond 1 m/sec. In the experiments a particle speed of 0.07 m/sec was used. Finally, the microfluidic channels employed, were straight channels based on TOPAS, offering absorption below <0.5 dB at 635 nm, having a cross section of 100×100 μm. The microfluidic scheme had no seethe control, thus particles propagated at random trajectories within the cross-section.

Sample preparation

The experiments involved three distinct categories of calibrated transparent PMMA spheres (POLYAN) of different diameters 12, 16, and 20 μm. The initial concentration of calibrated transparent PMMA spheres was 5% and it was further diluted to 1:200 by adding purified water. The low particle concentration was used so as to avoid clogging in the microfluidic channel and reduce the probability of particle clustering. All tubes and microfluidic chambers employed during dataset generation were used solely for one class to avoid cross-

contamination that would hinder labeling of the dataset. Prior to any measurement, all the microfluidic channels were thoroughly cleaned by debris by injecting purified water at high flow.

Synthetic frames to 1-D vectors

In order to leverage the OSS optical preprocessing, synthetic frames are serialized using multiple spatial orientations. Although this step does not require any digital preprocessing, it determines the order in which pixels are inserted into the stream, thereby potentially augmenting input data. Specifically, it can generate additional correlations among pixels when combined with a filter. The process consists of the following steps: the values of the synthetic frame are divided into blocks of size $M \times M$ pixels, where M is equal to 20 in numerically simulated OSS and to 5 in experimental OSS (see Fig. 4). Each block is sequentially serialized in two alignments: column-wise and row-wise. Both of these serialized forms are utilized in the representation of each synthetic frame in a one-dimensional vector. The stride, defined as the step by which the patch window shifts across the synthetic frame before serializing the next block, is employed in two orientations. The first orientation is row-wise, serializing pixels within the patch based on horizontal alignment, and the second is column-wise, focusing on vertical alignment of the patch pixels. Utilizing both patch orientations aims to augment the spatio-temporal characteristics of each particle sample within its one-dimensional representation. Consequently, the classification dataset comprises 4378 different one-dimensional vectors, each representing a synthetic frame.

Machine learning training process

The neural network models were formulated using the Keras API³⁶ implemented on the Tensorflow framework²⁵, and their training and evaluation processes were conducted utilizing a graphical processing unit (GPU). In all instances, 70% of the available data particles were designated for training, while 20% out of these were reserved for validation to mitigate the risk of overfitting. The remaining 30% of the data were allocated for evaluation to ensure a comprehensive and unbiased assessment of model performance. The optimization algorithm selected for these models was Adam, with a fixed learning rate of 1×10^{-4} and the training regimen was performed for 300 epochs. Furthermore, in our study, we employed an early stopping mechanism with a patience of 100 epochs after the minimum validation loss was reached to prevent overfitting. This ensures that the model does not continue to train beyond the point where it is effectively learning, thereby mitigating the risk of memorization of the training data. To account for potential variations in performance due to the random initialization of weights, we repeated the training procedure 10 times. Each repetition started with a new random initialization of the model's weights (as the reviewer commented), ensuring the robustness of the reported results. By doing so, we account for any possible variability in the final accuracy and provide a more reliable estimate of the model's performance. Additionally, during each repetition the entire dataset was shuffled, preventing any subset with specific characteristics from being consistently presented to the model during training.

The architectural configurations encompassed a basic digital perceptron, a fully connected feedforward neural network (FNN) with 1–2 layers and lightweight recurrent neural networks (RNN), instantiated as long-short term memory (LSTM), gated-recurrent unit (GRU) and Vanilla-RNN (V-RNN).

To optimize hyperparameters, the 'Optuna' framework³², which utilizes a Tree-structured Parzen Estimator, was employed to maximize the testing accuracy. A comprehensive evaluation of the aforementioned models was undertaken, focusing on several factors: the number of hidden neurons, the learning rate and the batch size within the FCs or the RNN models, as well as the activation function of the hidden nodes in the FC layers.

The data delineated in Tables 2 and 3 encapsulate a robust array of computational characteristics, including the number of parameters, FLOPS, and Thermal Design Power (TDP). The TDP represents the maximum amount of heat a processing component, like a GPU is expected to generate under heavy loads, where a lower TDP generally indicates reduced power consumption. These metrics were derived through the application of TensorFlow libraries in conjunction with NVIDIA tools, ensuring precise and dependable measurements critical for assessing and contrasting computational performance. In this analysis, an NVIDIA 2080Ti GPU with a TDP of 250 Watts was utilized. The determination of the TDP percentage was conducted using the GPU-Z monitoring utility³⁷. The average TDP percentage was calculated over 100 training epochs for each neural network model.

Data availability

All codes generated for the emulation of the optical spectrum slicing concept and the dataset generated in the context of this work are available upon request to the corresponding author.

Received: 26 April 2024; Accepted: 5 October 2024

Published online: 15 October 2024

References

- Barteneva, N. S., Fasler-Kan, E. & Vorobjev, I. A. Imaging flow cytometry: coping with heterogeneity in biological systems. *J. Histochem. Cytochem.* **60**, 723–733 (2012).
- Gouda, M. et al. Improving the classification accuracy in label-free flow cytometry using event-based vision and simple logistic regression. *IEEE J. Sel. Top. Quantum Electron.* **29**, 1–8 (2023).
- Blasi, T. et al. Label-free cell cycle analysis for high-throughput imaging flow cytometry. *Nat. Commun.* **7**, 10256 (2016).
- Rane, A. S., Rutkauskaitė, J., deMello, A. & Stavrikis, S. High-throughput multi-parametric imaging flow cytometry. *Chem* **3**, 588–602 (2017).
- Zhang, F. et al. Intelligent image de-blurring for imaging flow cytometry. *Cytometry A* **95**, 549–554 (2019).
- Isozaki, A. et al. Intelligent image-activated cell sorting 2.0. *Lab. Chip* **20**, 2263–2273 (2020).
- Vembadi, A., Menachery, A. & Qasaimeh, M. A. Cell cytometry: review and perspective on biotechnological advances. *Front. Bioeng. Biotechnol.* **7**, 147 (2019).

8. Rožanc, J., Finšgar, M. & Maver, U. Progressive use of multispectral imaging flow cytometry in various research areas. *The Analyst* **146**, 4985–5007 (2021).
9. Doan, M. et al. Diagnostic potential of imaging flow cytometry. *Trends Biotechnol.* **36**, 649–652 (2018).
10. Rodrigues, M. A. et al. The in vitro micronucleus assay using imaging flow cytometry and deep learning. *Npj Syst. Biol. Appl.* **7**, 20 (2021).
11. Lau, A. K. S., Shum, H. C., Wong, K. K. Y. & Tsia, K. K. Optofluidic time-stretch imaging—an emerging tool for high-throughput imaging flow cytometry. *Lab. Chip* **16**, 1743–1756 (2016).
12. Tsilikas, I. et al. Time-Stretched Imaging Flow Cytometry and Photonic Neuromorphic Processing for Particle Classification. in *2023 Conference on Lasers and Electro-Optics Europe & European Quantum Electronics Conference (CLEO/Europe-EQEC)* 1–1 (IEEE, Munich, Germany, 2023). <https://doi.org/10.1109/CLEO/Europe-EQEC57999.2023.10232590>.
13. Coddington, I., Newbury, N. & Swann, W. Dual-comb spectroscopy. *Optica* **3**, 414 (2016).
14. Mizuno, T. et al. Optical image amplification in dual-comb microscopy. *Sci. Rep.* **10**, 8338 (2020).
15. Lei, C. et al. High-throughput imaging flow cytometry by optofluidic time-stretch microscopy. *Nat. Protoc.* **13**, 1603–1631 (2018).
16. Kiester, A. S., Ibey, B. L., Coker, Z. N., Pakhomov, A. G. & Bixler, J. N. Strobe photography mapping of cell membrane potential with nanosecond resolution. *Bioelectrochemistry* **142**, 107929 (2021).
17. Perot, E., de Tournemire, P., Nitti, D., Masci, J. & Sironi, A. Learning to Detect Objects with a 1 Megapixel Event Camera. in *Advances in Neural Information Processing Systems* vol. 33 16639–16652 (Curran Associates, Inc., 2020).
18. Abreu, S., Gouda, M., Lugnan, A. & Bienstman, P. Flow cytometry with event-based vision and spiking neuromorphic hardware. in *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)* 4139–4147 (IEEE, Vancouver, BC, Canada, 2023). <https://doi.org/10.1109/CVPRW59228.2023.00435>.
19. Tsilikas, I. et al. Neuromorphic Camera Assisted High-Flow Imaging Cytometry for Particle Classification. in *2023 Conference on Lasers and Electro-Optics Europe & European Quantum Electronics Conference (CLEO/Europe-EQEC)* 1–1 (IEEE, Munich, Germany, 2023). <https://doi.org/10.1109/CLEO/Europe-EQEC57999.2023.10232061>.
20. Zhang, Z. et al. Work in Progress: Neuromorphic Cytometry, High-throughput Event-based flow Flow-Imaging. in *2022 8th International Conference on Event-Based Control, Communication, and Signal Processing (EBCCSP)* 1–5 (IEEE, Krakow, Poland, 2022). <https://doi.org/10.1109/EBCCSP56922.2022.9845595>.
21. Tsigotis, A., Tsilikas, I., Sozos, K., Bogris, A. & Mesaritakis, C. Photonic Neuromorphic Accelerator Combined with an Event-Based Neuromorphic Camera for High-Speed Object Classification. in *2023 Conference on Lasers and Electro-Optics Europe & European Quantum Electronics Conference (CLEO/Europe-EQEC)* 1–1 (IEEE, Munich, Germany, 2023). <https://doi.org/10.1109/CLEO/Europe-EQEC57999.2023.10232077>.
22. Haessig, G., Berthelon, X., Ieng, S.-H. & Benosman, R. A spiking neural network model of depth from defocus for event-based neuromorphic vision. *Sci. Rep.* **9**, 3744 (2019).
23. Tsigotis, A. et al. Unconventional integrated photonic accelerators for high-throughput convolutional neural networks. *Intell. Comput.* **2**, 0032 (2023).
24. Lugnan, A., Gooskens, E., Vatin, J., Dambre, J. & Bienstman, P. Machine learning issues and opportunities in ultrafast particle classification for label-free microflow cytometry. *Sci. Rep.* **10**, 20724 (2020).
25. TensorFlow. <https://www.tensorflow.org/>.
26. Kingma, D. P. & Ba, J. Adam: A Method for Stochastic Optimization. Preprint at <https://doi.org/10.48550/arXiv.1412.6980> (2017).
27. Dalal, N. & Triggs, B. Histograms of Oriented Gradients for Human Detection. in *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)* vol. 1 886–893 (IEEE, San Diego, CA, USA, 2005).
28. Li, Y. et al. Deep cytometry: deep learning with real-time inference in cell sorting and flow cytometry. *Sci. Rep.* **9**, 11088 (2019).
29. Yu, H., Yang, L. T., Zhang, Q., Armstrong, D. & Deen, M. J. Convolutional neural networks for medical image analysis: State-of-the-art, comparisons, improvement and perspectives. *Neurocomputing* **444**, 92–110 (2021).
30. Sozos, K. et al. High-speed photonic neuromorphic computing using recurrent optical spectrum slicing neural networks. *Commun. Eng.* **1**, 1–10 (2022).
31. Ferrera, M. et al. On-chip CMOS-compatible all-optical integrator. *Nat. Commun.* **1**, 29 (2010).
32. Akiba, T., Sano, S., Yanase, T., Ohta, T. & Koyama, M. Optuna: A Next-generation Hyperparameter Optimization Framework. in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* 2623–2631 (Association for Computing Machinery, New York, NY, USA, 2019). <https://doi.org/10.1145/3292500.3330701>.
33. Gallego, G. et al. Event-based vision: a survey. *IEEE Trans. Pattern Anal. Mach. Intell.* **44**, 154–180 (2022).
34. Xu, X. et al. 11 TOPS photonic convolutional accelerator for optical neural networks. *Nature* **589**, 44–51 (2021).
35. Feldmann, J. et al. Parallel convolutional processing using an integrated photonic tensor core. *Nature* **589**, 52–58 (2021).
36. Keras: Deep Learning for humans. <https://keras.io/>.
37. TechPowerUp. *TechPowerUp* <https://www.techpowerup.com/gpu/> (2024).

Acknowledgements

This work has received funding from the EU H2020 project NEoteRIC under grant agreement 871330 and by EU Horizon PROMETHEUS project under grant agreement 101070195.

Author contributions

I.T. and C.M. set up and conducted the cytometry experiments. C.M. performed the digital pre-processing on the dataset. A.T. carried out the photonic pre-processor simulations and conducted the photonic machine learning experiments, with assistance from G.S. and A.B. S.D. analyzed the data using digital machine learning for benchmarking purposes. K.P. and G.B. provided expertise and specifications regarding the event-based camera. A.B. and C.M. initiated the project and supervised the entire process. The manuscript was prepared with contributions from all authors.

Declarations

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to C.M.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2024