Training Strategies for Spiking Neural Networks Integrated with Event-Based Vision in Label-Free Flow Cytometry

Muhammed Gouda,[1] *Steven Abreu*,[2] and *Peter Bienstman* [1]

[1] Ghent University, Photonics Research Group, Sint-Pietersnieuwstraat 33, 9000 Ghent
[2] University of Groningen 2, Bernoulli Institute, Broerstraat 5, 9712 CP Groningen

Abstract

Spiking neural networks (SNNs) are computational models inspired by the biological brain, emulating the spiking behavior of neurons. Similarly, event-based vision sensors aim to mimic the functioning of the human eye. In this research, we harness both technologies to facilitate the classification of micro-particles within the domain of label-free flow cytometry. Building upon our prior work employing basic logistic regression with binary labels, which yielded an accuracy exceeding 98%, our objective is to extend the applicability of the system with SNNs to more complex tasks where a broader spectrum of cells could be used. However, this introduces the challenge of training SNNs, as they suffer from the issue of vanishing gradients. In order to address this issue, we employ two distinct training algorithms: EventProp and surrogate gradient learning. Both algorithms demonstrated a performance of over 99% accuracy in a task involving two classes of particles. In the case of a four-class problem, EventProp achieved a 92% accuracy, whereas with the use of surrogate gradients, we were able to maintain a 99% accuracy level.

Introduction

Flow cytometry is a technology concerned with identifying populations of cells/microparticles in a fluid [1]. This could be applied to several domains from medicine to cosmetics and environmental engineering. In almost all those domains a highly accurate device would be required to carry out the classification. In an attempt to achieve this for a wide variety of cells and particles we combine here a novel imaging sensor, known as an event-based camera, and a spiking neural network, both providing sparse dynamics for an energy-efficient system.

Several attempts have been carried out to mimic a biological brain as close as possible. In this regard, a class of neural networks, known as spiking neural networks SNNs, has been designed. These networks differ from artificial neural networks (ANNs) in the sense that they work with short pulses. One of the SNN neuron models known as leaky-integrate-and-fire (LIF) neuron has an internal state often referred to as the membrane potential which gives rise to an output under certain conditions.

To illustrate the difference between ANNs and SNNs, consider figure 1 where both an artificial neuron and a spiking neuron are presented. In both cases a linear combination of input features is applied. However, the spiking neuron will only fire an output at specific time steps when its membrane potential is higher than a certain threshold. On the other hand, an artificial neuron will give an output at all time steps. This makes spiking neural networks attractive for implementation on power-efficient hardware. Moreover, for datasets that are event-based, the sparse dynamics of spiking neurons make them a better fit.

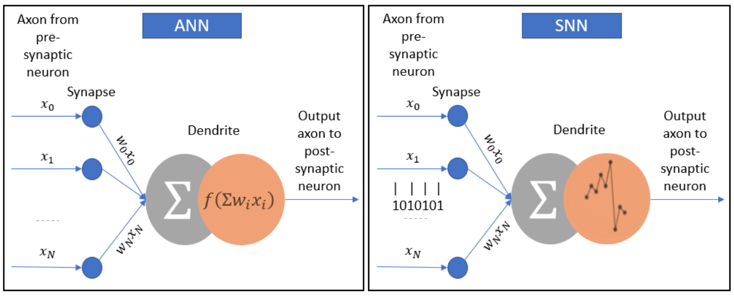

Figure 1 Neuron models in artificial neural networks (left) and spiking neural networks (right). A spiking neuron gives rise to sparse dynamics making it a natural fit for event-based data

The main drawback of working with SNNs is the difficulty of training them. Multiple attempts have been presented in the literature. These were for both supervised and unsupervised learning. Spike time-dependent plasticity (STDP) is an unsupervised learning algorithm thought to be the closest to how actual brains work by making use of the relative delay between spikes. On the other hand, backpropagation has been applied to train SNNs in a supervised manner.

Training with backpropagation requires calculating the gradient of the loss with respect to the network weights. Seeking to do so with SNNs often fails to work due to the dead neuron problem where the gradient vanishes. The problem stems from the fact that the output of the neuron is a non-linear Heaviside function of its membrane potential. Taking the derivative of such output results in a delta function with zeros almost everywhere except at the threshold.

To surmount this problem, we applied two different training algorithms: The surrogate gradient method [2] and the EventProp algorithm [3].

Methodology

- *Experiment*

The setup is shown in Figure 2. It includes a laser source producing light at a wavelength of 632.8 nm. This light traverses a lens and a 25 µm pinhole before converging onto a PMMA microfluidic channel. Within the channel, there is a flow of microparticles pumped by a manual syringe pump linked to the upper port, while a liquid reservoir is connected to the opposite port. The event sensor captures the changes that occur to the diffraction and scattering of light upon the passage of a particle.

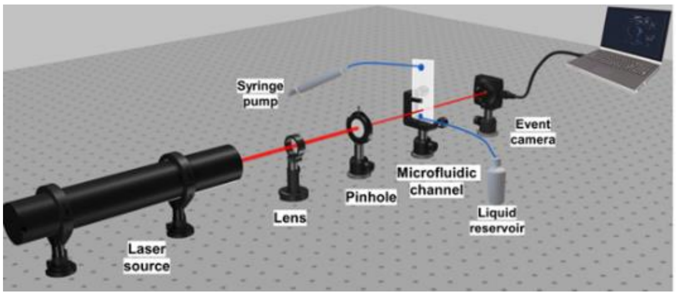

Figure 2 Experimental setup of the optical flow cytometer in this work. Particles flowing inside a microfluidic channel is exposed to coherent light which results in diffraction and scattering being captured by the event sensor.

- *Neural network design*

In this work, a spiking neural network (SNN) consisting of three layers (an input layer of 6,144 neurons, a hidden layer of 100 neurons, and an output layer of 4 neurons) was trained on a two and a four-class dataset to classify particles with diameters of 9 µm,12 µm,16 µm and 20 µm. The data was obtained using Prophesee event-based camera.

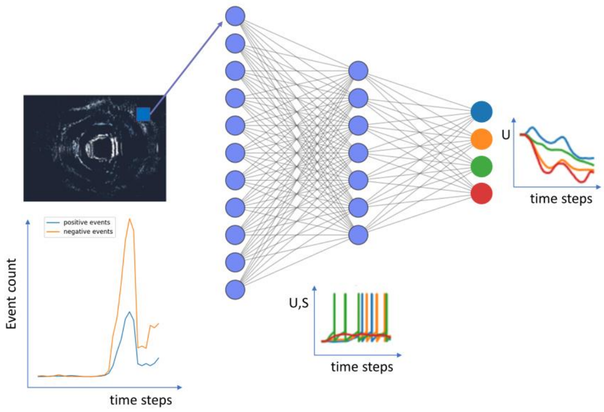

Figure 3 The architecture of the neural network built in this work. There is three layers, input, hidden and output layer. The events corresponding to flowing particles are sent to the network for training and testing.

Results

- Surrogate gradient

By replacing the true gradient of the neural network by an alternative which does not vanish, we could efficiently train the model on different classification problems. Figure 4 shows the train loss when we used different gradients to train the network. We were able to obtain a performance of 99%, both for 2 and 4 class problem.

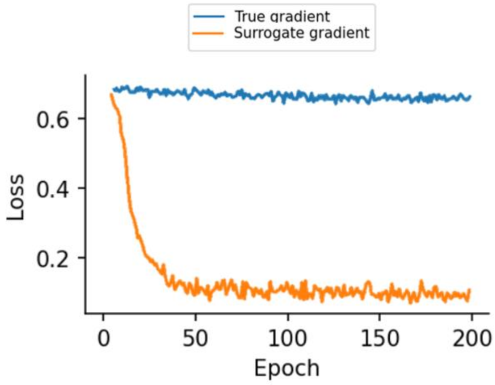

Figure 4 Training loss progression with the number of training epochs both in the case of using the true gradient of the network and the surrogate one.

- Event-Prop

The event-prop algorithm keeps the true gradient of the network but avoids backpropagating it at each time step. It focuses on the individual timesteps where a spike is fired and backpropagates the loss at those times. By doing so we could achieve similar results to those obtained by the first method on the binary classification problem while saving some training time. On the other hand, we found that this algorithm performs poorly on more complex tasks like the four-class problem (92% instead of 99%).

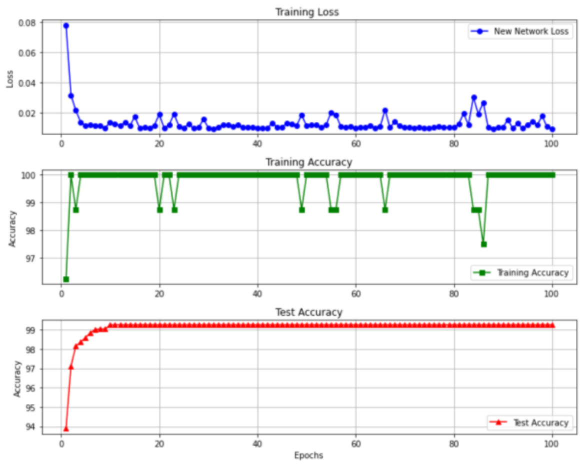

Figure 5 Training loss, training accuracy and test accuracy of the network trained using the EventProp algorithm.

Conclusion

In conclusion, our study showcases the successful utilization of established algorithms, EventProp, and surrogate gradient learning, in the context of micro-particle classification using spiking neural networks and event-based vision sensors. By applying these existing methodologies, we achieved significant improvement in the accuracy of our simple flow cytometer, surpassing 99% accuracy in a two-class task and demonstrating notable adaptability in a four-class scenario with accuracies of 92% and a consistent 99%, respectively.

References

[1] Adan, A., Alizada, G., Kiraz, Y., Baran, Y., & Nalbant, A. (2017). Flow cytometry: basic principles and applications. Critical reviews in biotechnology, 37(2), 163-176.

[2] Neftci, E. O., Mostafa, H., & Zenke, F. (2019). Surrogate gradient learning in spiking neural networks: Bringing the power of gradient-based optimization to spiking neural networks. *IEEE Signal Processing Magazine*, *36*(6), 51-63.

[3] Wunderlich, T.C., Pehle, C. Event-based backpropagation can compute exact gradients for spiking neural networks. *Sci Rep* **11**, 12829 (2021). https://doi.org/10.1038/s41598-021-91786-z

Cooking and baking:

Approaching the thin-film absorption limit with monolayer semiconductor superlattices
S. Elrafey (TU Eindhoven)

On-chip Gain in Neodymium-doped Aluminium Oxide Waveguide Amplifiers
B. Jongebloed (UTwente)

Proximity Effect Correction for E-beam Fabrication of Aluminum Oxide Waveguides
N. Chahir (UTwente)

Pt/Ge/Au ohmic p-contacts with reduced optical loss suitable for integrated opto-electronic devices
J. de Graaf (TU Eindhoven)

Buzz about AI:

How photonic reservoirs improve the performance of deep neural networks
I. Bauwens (VUBrussel)

Training Strategies for Spiking Neural Networks Integrated with Event-Based Vision in Label-Free Flow Cytometry
Muhammed Gouda (UGent)

Improving photonic Ising machines: a guide to reaching the ground state
J. Lamers (VUBrussel)

Feasibility Study of Training Neural Networks Using Photonics
T. Tubbing (TU Eindhoven)

Posters (odd numbers):

1. Thermal design considerations for vertical integration of InP nanophotonic membrane devices with InP electronics

Posters (even numbers):

2. Heterogeneous integration of evanescently-coupled GaAs-based amplifiers for laser systems emitting in the near-infrared