A Spiking Neural Network for Classifying NIR Spectra of Fruits

Anna Wróbel†, Yulia Sandamirskaya†, Thomas Ott†

†Institute of Computational Life Sciences, ZHAW Zurich University of Applied Sciences,
Schloss 1, P.O. Box, 8820, Wädenswil, Switzerland

Email: anna.wrobel@zhaw.ch, yulia.sandamirskaya@zhaw.ch, thomas.ott@zhaw.ch

Abstract— Near-Infrared (NIR) Spectroscopy is widely applied in agriculture and food industry for the determination of fruit ripeness, the content in soluble solids, pH and acidity. In this study, we report on the develeopment of a novel neuromorphic classifier based on Spiking Neural Networks (SNNs) to classify NIR spectra of fruit species. Neuromorphic computing holds the potential for a low-power real-time recognition system based on NIR spectroscopy signals that could be used not only in food industry, but also in pharmaceutical and medical applications. For benchmarking, we compare the performance of the classifier to the performance of non-spiking Artificial Neural Networks (ANN) and Support Vector Machines (SVM). Furthermore, we show how our SNN-based algorithm can be implemented in the mixed-signal analog-digital neuromorphic device DYNAP-SE.

1. Introduction

The use of NIR spectroscopy is widespread in agriculture and the food industry, primarily for rapid, non-destructive analysis of soluble solids content or dry matter and water content [1]. Furthermore, it plays a crucial role in identifying varieties and establishing chemical properties [2]. Near-Infrared (NIR) spectroscopy, which encompasses the electromagnetic spectrum's near-infrared region from 780 nm to 2500 nm, relies heavily on the product's chemical composition and light scattering attributes. The majority of absorption bands found in the near infrared region are the result of the combination or overtone of fundamental absorption bands observed in the infrared region of the electromagnetic spectrum. These fundamental bands arise from transitions in vibrational and rotational states [1].

Deep Neural Networks (DNN), Convolutional Neural Network (CNN) and Residual Network (ResNet) have been successfully applied for the recognition of fruits using NIR spectra [2]. By combining deep learning with robust feature extraction, the authors demonstrated an accuracy of approximately 99% for the task of classifying five types of fruits including Apple, Avocado, Dragon Fruit, Guava, and Mango.

Neuromorphic hardware devices aim to mimic brain functions and promise high energy efficiency and robustness compared to conventional computing devices. In particular mixed-signal VLSI neuromorphic devices come with an interesting potential for applications in the field, e.g. in agriculture, due to their ultra-low power characteristics. The DYNAP-SE is such mixed-signal analogue-digital device. It features 4 chips, each equipped with 4 cores of 256 neurons. These neurons follow Dale's law, meaning they are exclusively either excitatory or inhibitory. The device accommodates four types of synapses: two excitatory (AMPA and NMDA) and two inhibitory (GABA A and GABA B). The synapse weights and neuron biases are shared within a core. This device includes a 64-neuron fan-in [5]. These characteristics necessitate specific considerations during the implementation of a network on the chip.

A spiking neural network (SNN) can be considered a link between conventional artificial neural networks and neuromorphic hardware. An important aspect of SNNs concerns signal encoding. The most common techniques are rate or time encoding. The DYNAP-SE accepts input as either Poisson encoding with a certain frequency or as spike trains generated by the FPGA (Field Programmable Gate Array).

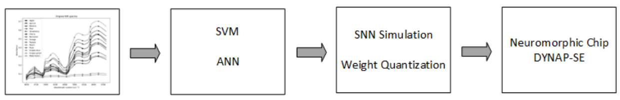

Figure 1: Pipeline overview. We used the NIR-spectral data as direct input for the neural networks. We started with building SVMs as a reference, then built an ANN, which was converted into an SNN and simulated with quantization-aware training. Finally, the chosen architecture with trained weights was transferred onto the neuromorphic chip DYNAP-SE.

In this study, we used NIR spectra from 14 fruit species to develop a neuromorphic classifier based on SNNs. The process involved creating a feed-forward network (ANN), converting it to a spiking network (SNN), and quantizing the weights to transfer them onto a neuromorphic chip (DYNAP-SE), see Fig. 1. For the spike generation from spectra we used rate coding. In our approach, we decided to test the performance of the networks without any feature extraction. The analysis was performed using Scikit-learn, PyTorch, SNNtorch, and Samna software for the DYNAP-SE, and training was carried out on a Tesla P100 GPU.

ORCID iDs First Author: 0009-0000-6545-532X, Second Author: 0000-0003-4684-202X, Third Author: 0000-0003-1635-7816

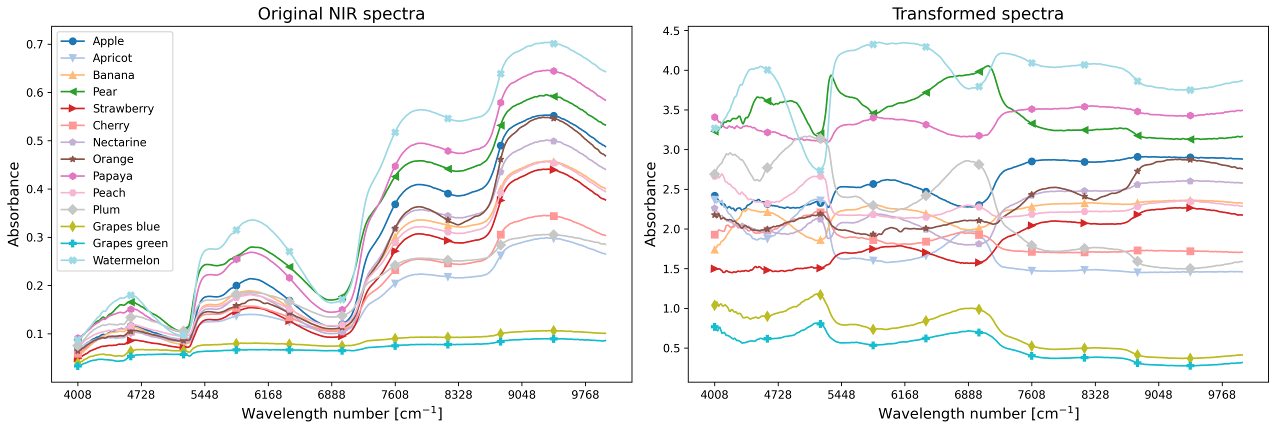

Figure 2: NIR spectra averaged over the fruit types. On the left: original spectra. On the right: z-transformed spectra shifted up by the absolute value of the minimum value.

2. Methods

The NIR spectra were recorded with Bühler NIRVIS in absorption mode, covering a wavelength range from 12500 to 4000 cm^{-1}. It is a Fourier-Transform-NIR-spectrometer, equipped with a polarization interferometer. Measurements were made on the horizontal sample table, with a circular measurement window with a diameter of 1 cm. Prior to each measurement session, white-balance was conducted to ensure consistent conditions. Every product was measured four times at 15 different sites. The four measurements were averaged for the final spectra.

2.1. Dataset

The dataset contains 226 spectra of the selected fruit species: Apple, Apricot, Banana, Pear, Strawberry, Cherry, Nectarine, Orange, Papaya, Peach, Plum, Grapes blue, Grapes green and Watermelon. Fig. 2 shows the spectral data averaged for all the fruit types. Each species contributed between 15 to 23 spectra. Given the size of the dataset, data augmentation was necessary for classifier construction. The augmented dataset was created with randomly scaled contributions from offset, slope and multiplication to simulate baseline offset, slope differences and differences in intensity in the spectral recording described in [3]. For this, the data was divided into a training set containing 10 spectra from each fruit species, each augmented 100 times. The remaining spectra were reserved for the test set, without augmentation. This process resulted in an unbalanced test set, therefore all reported accuracies are balanced over the classes.

In our analysis, we utilized raw spectra and z-transformed input, which involves normalizing the data to a zero mean and variance of one standard deviation. Beyond scaling, we shifted the data upwards by the absolute value of the minimum value. This adjustment was necessary to enable the transformation of the normalized samples into spikes.

2.2. Training

PCA analysis indicated that while most classes could be readily separated, some overlap was observed. For the classification baseline, we built a SVM classifier with a linear kernel and a regularization parameter (C) equal to 0.1. The classifier was tested on raw and z-transformed data, with the scaling significantly improving the separation of the 14 fruit classes. To construct the SNN, we initially tested feed-forward networks with no, one, or two hidden layers, while also reducing the input dimensions through downsampling. The ANNs were trained for 100 epochs using gradient descent and the Adam optimizer, utilizing 90% of the augmented dataset for training and reserving 10% for validation. Subsequently, we built an SNN using snntorch [4] with the selected multilayer perceptron architecture. The network was trained for 30 epochs using back-propagation with a surrogate gradient, Heavyside output function and Adam optimizer.

2.3. Simulation

Before transferring the SNN onto the neuromorphic device DYNAP-SE, we had to address several constraints inherent to the device compared to the ANN. Firstly, the network weights were quantized (integer 8-, 4-, 3-, 2-bit precision), and quantization-aware training was performed using Brevitas library for PyTorch [6], with and without input transformation. To manage the issue of restricted fan-in and simplify parameter tuning, we transferred a reduced network—able to classify only three fruit classes (Apple, Plum, and Watermelon)—onto the chip. These classes were linearly separable, enabling us to simulate a simplified network without hidden layers.

The spiking networks were simulated using snntorch with a membrane potential decay rate (beta) equal to 0.95. The simulation was run for 30 epochs with 50 steps, meaning that each sample was presented 50 times at each iteration.

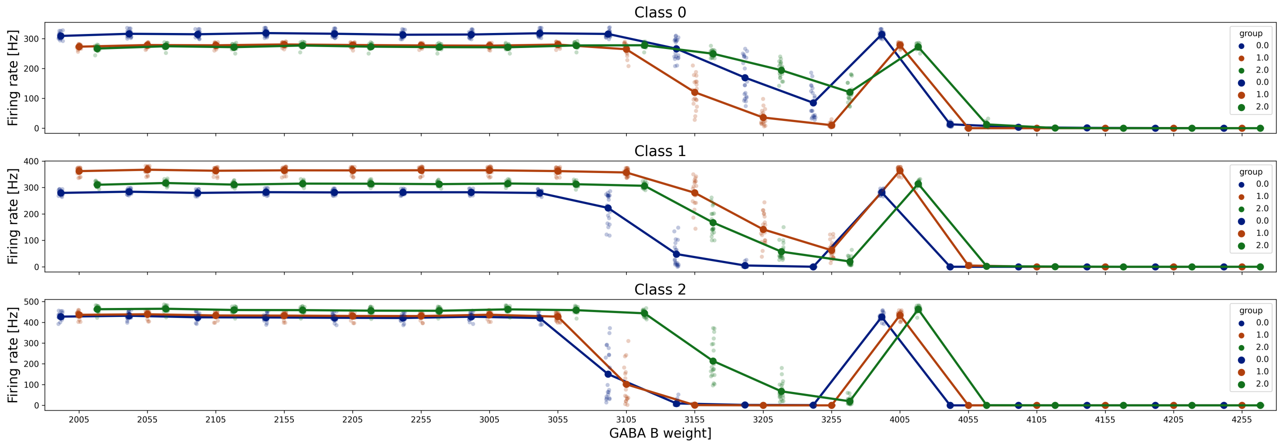

Figure 3: Tuning of GABA B weight. The rows represent samples firing patterns for class 0, 1 and 2 (blue, red, green). The values on the X-axis is a combination of the coarse and fine value of the weight parameter with ranges $(0, 8]$ and $(0, 256]$ respectively. The lower range of values of the weight ensure good separation between the classes.

In the signal encoding for the transfer to the neuromorphic chip the number of steps was increased to 500 to correspond to 500 ms.

2.3.1. Signal encoding

To ensure reproducibility signals were encoded using a rate-encoding and served as inputs on the DYNAP-SE as spike trains via spike generators. For the spectral data, each input feature X_i acted as the probability of an event (spike) occurring at any given time step, yielding a rate-coded value. This can be interpreted as a Bernoulli trial where the number of trials is $n = 1$ and the probability of success (spiking) is $P = X_i$. For z-transform signals, each signal was additionally divided by the maximum value of the spectra, ensuring the signal stayed within the 0 to 1 range. The encoded spikes were used as input for the spike generators by the FPGA on the chip.

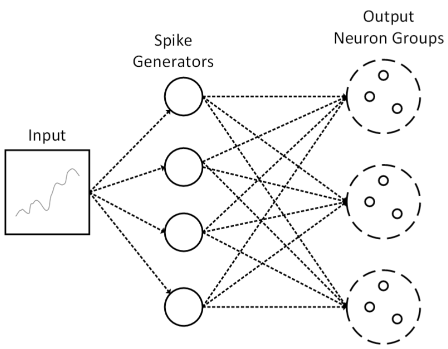

Figure 4: Network topology on the DYNAP-SE. The input is converted into spikes via rate encoding. The spike train is fed into spike generators on the chip. The spike generators are connected via excitatory AMPA synapses or inhibitory GABA B synapses with three populations of neurons, each consisting of 20 neurons.

To account for device mismatch [5] concerning the neurons on the DYNAP-SE, we used population coding on the output layer, meaning that for each class, not just one but a group of 20 neurons were designated, see Fig. 4. Furthermore, the neurons on the chip underwent testing, with the most uniformly spiking neurons selected.

2.3.2. Biases fine-tuning

Due to chip restrictions, primarily the low fan-in, we opted for a small network capable of classifying three fruit types, utilizing a 3-bit weight precision (weight values equal to -3, -2, -1, 0, 1, 2, 3). Negative weights were implemented by GABA B (shunting) synapses, while positive weights were managed by AMPA synapses. The absolute value of the weight corresponded to the quantity of the respective type of synapse. For instance, if the weight equaled 2, then the pre- and post-synaptic neuron were interconnected with two AMPA synapses.

The network on the neuromorphic chip required additional fine-tuning of the biases. This was achieved by sweeping through the synapse and neuron parameters while disabling spike frequency adaptation. An example of tuning curves can be found in Fig. 3. Each row represents firing patterns for one sample from classes 0, 1, 2, respectively. The desired pattern was a clear distinction between the neurons from class populations 0, 1, 2 (colored blue, red, green respectively) within the samples. This pattern is prominently visible within the lower range of GABA B values in the figure.

3. Results

Table 1 summarizes the performance of the chosen models on the test set. As suggested by the SVM results, a shallow feed-forward ANN without any hidden layer delivers good performance, yielding results comparable to those of the SVM. In the classification of 14 fruit types, introducing a single hidden layer results in a slight performance increase (less than 1%), with a comparable increase observed upon the addition of a second hidden layer. Standardizing the input enhances the model performance by ~10%.

Model	Input size	Fruit types	Scaled	Accuracy %
SVM	500	14	yes	92.21
SVM	500	14	no	83.70
SVM	125	14	yes	92.31
SVM	125	14	no	79.53
SVM	25	3	no	100.00
ANN	250	14	no	87.51
ANN	25	3	no	100.00
SNN	250	14	yes	67.31
SNN	250	14	no	57.06
SNN	25	3	no	100.00
SNN	25	3	no	94.44
SNN 3 bit	250	14	yes	55.04
SNN 3 bit	250	14	no	44.48
SNN 3 bit	25	3	yes	100.00
SNN 3 bit	25	3	no	79.63
DYNAP-SE	25	3	yes	100.00

Table 1: The balanced accuracy on the test set for the chosen models indicates a significant drop in accuracy when raw spectra are used for classification. When the input size is reduced (for 14 types of fruits), the performance of the classifier significantly decreases.

However, when the ANN is transformed into an SNN and the weight precision of the quantized weights is subsequently reduced, a significant performance drop is observed, particularly in the 14-class case. For the classification of 3 classes with input standardization, the impact is minimal until a 3-bit weight precision is reached. This network configuration was chosen for transfer onto the neuromorphic chip. For all models, a reduction in weight precision results in a drop in accuracy, as illustrated in Fig. 5. Models that use z-transformed data outperform those using raw data.

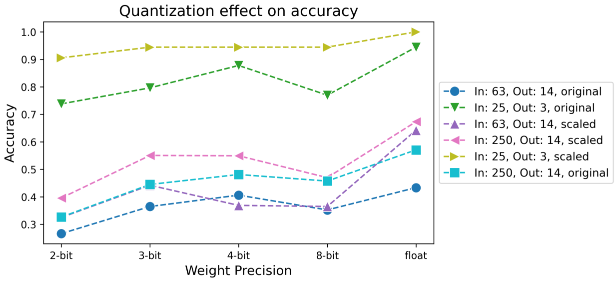

Figure 5: The influence of weight quantization on the model performance for selected models.

4. Discussion

In this work, we demonstrated a pipeline for transferring a feed-forward neural network onto a mixed-signal VLSI neuromorphic chip DYNAP-SE using PyTorch for the simulation, for the classification of fruit NIR spectra. The primary challenge of this approach lies in the weight quantization and the restricted fan-in of the neurons on the chip. Furthermore, there exist substantial differences between the neuron models in the snntorch implementation and those on the DYNAP-SE. A more similar model would be necessary for a smoother transition onto the neuromorphic chip.

Given the substantial differences between the simulation model and the neuromorphic chip, it may be beneficial to conduct on-chip learning to achieve the best results. As demonstrated by the impact of input transformation, signal pre-processing plays a vital role in network performance. For this reason, it would be intriguing to explore the performance of a convolutional neural network (CNN) and a spiking CNN on this dataset. Moving forward, a comparison between the performance of a mixed-signal analog-digital neuromorphic device and a digital neuromorphic chip will be our next step.

References

[1] B. M. Nicolaıi, K. Beullens, E. Bobelyn, A. Peirs, W. Saeys, K. I. Theron, J. Lammertyn, "Nondestructive measurement of fruit and vegetable quality by means of NIR spectroscopy: A review," *Postharvest Biology and Technology*, vol.2, pp.99–118, 2007.

[2] D. Ninh, T. N. C. Doan, C. K. Ninh, T. X. Nguyen-Thi, N. L. Thanh, "Fruit recognition based on near-infrared spectroscopy using deep neural networks," *The 5th International Conference on Machine Learning and Soft Computing*, pp.90–95, 2021.

[3] E. J. Bjerrum, M. Glahder, T. Skov, "Data Augmentation of Spectral Data for Convolutional Neural Network (CNN) Based Deep Chemometrics," *CoRR*, 1710.01927, 2017.

[4] J. K. Eshraghian, M. Ward, E. Neftci, X. Wang, G. Lenz, G. Dwivedi, M. Bennamoun, D. S. Jeong, W. D. Lu, "Training Spiking Neural Networks Using Lessons From Deep Learning," *CoRR*, 2109.12894, 2021.

[5] S. Moradi, N. Qiao, F. Stefanini, G. Indiveri, "A scalable multi-core architecture with heterogeneous memory structures for Dynamic Neuromorphic Asynchronous Processors (DYNAPs)," *IEEE Transactions on Biomedical Circuits and Systems*, vol.12, pp.106–122, 2018.

[6] A. Pappalardo, "Xilinx/brevitas," *Zenodo*, 2023.