

A Lightweight Architecture for Real-Time Neuronal-Spike Classification

Siddiqi, Muhammad Ali; Vrijenhoek, David; Landsmeer, Lennart P.L.; van der Kleij, Job; Gebregiorgis, Anteneh; Romano, Vincenzo; Bishnoi, Rajendra; Hamdioui, Said; Strydis, Christos

DOI

[10.1145/3649153.3649186](https://doi.org/10.1145/3649153.3649186)

Publication date

2024

Document Version

Final published version

Published in

CF '24

Citation (APA)

Siddiqi, M. A., Vrijenhoek, D., Landsmeer, L. P. L., van der Kleij, J., Gebregiorgis, A., Romano, V., Bishnoi, R., Hamdioui, S., & Strydis, C. (2024). A Lightweight Architecture for Real-Time Neuronal-Spike Classification. In *CF '24: Proceedings of the 21st ACM International Conference on Computing Frontiers* (pp. 32-40). ACM. <https://doi.org/10.1145/3649153.3649186>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.



A Lightweight Architecture for Real-Time Neuronal-Spike Classification

Muhammad Ali Siddiqi
Lahore University of Management
Sciences
Pakistan
m.siddiqi@lums.edu.pk

David Vrijenhoek
Delft University of Technology
The Netherlands
davidvrijenhoek@gmail.com

Lennart P. L. Landsmeer
Delft University of Technology
The Netherlands
l.p.l.landsmeer@tudelft.nl

Job van der Kleij
Delft University of Technology
The Netherlands
j.vanderkleij@student.tudelft.nl

Anteneh Gebregiorgis
Delft University of Technology
The Netherlands
a.b.gebregiorgis@tudelft.nl

Vincenzo Romano
Erasmus Medical Center, Rotterdam
The Netherlands
v.romano@erasmusmc.nl

Rajendra Bishnoi
Delft University of Technology
The Netherlands
r.k.bishnoi@tudelft.nl

Said Hamdioui
Delft University of Technology
The Netherlands
s.hamdioui@tudelft.nl

Christos Strydis
Erasmus Medical Center, Rotterdam
The Netherlands
c.strydis@erasmusmc.nl

ABSTRACT

Electrophysiological recordings of neural activity in a mouse's brain are very popular among neuroscientists for understanding brain function. One particular area of interest is acquiring recordings from the Purkinje cells in the cerebellum in order to understand brain injuries and the loss of motor functions. However, current setups for such experiments do not allow the mouse to move freely and, thus, do not capture its natural behaviour since they have a *wired* connection between the animal's head stage and an acquisition device. In this work, we propose a lightweight neuronal-spike detection and classification architecture that leverages on the unique characteristics of the Purkinje cells to discard unneeded information from the sparse neural data in real time. This allows the (condensed) data to be easily stored on a removable storage device on the head stage, alleviating the need for wires. Synthesis results reveal a >95% overall classification accuracy while still resulting in a small-form-factor design, which allows for the *free movement* of mice during experiments. Moreover, the power-efficient nature of the design and the usage of STT-RAM (Spin Transfer Torque Magnetic Random Access Memory) as the removable storage allows the head stage to easily operate on a tiny battery for up to approximately 4 days.

CCS CONCEPTS

• **Applied computing** → **Life and medical sciences**; • **Computer systems organization** → **Embedded and cyber-physical systems**; • **Hardware** → **Emerging technologies**.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CF '24, May 7–9, 2024, Ischia, Italy

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0597-7/24/05

<https://doi.org/10.1145/3649153.3649186>

KEYWORDS

Electrophysiological recordings, spike detection, spike classification, Purkinje cells, cerebellum, low-power computing, STT-RAM

ACM Reference Format:

Muhammad Ali Siddiqi, David Vrijenhoek, Lennart P. L. Landsmeer, Job van der Kleij, Anteneh Gebregiorgis, Vincenzo Romano, Rajendra Bishnoi, Said Hamdioui, and Christos Strydis. 2024. A Lightweight Architecture for Real-Time Neuronal-Spike Classification. In *21st ACM International Conference on Computing Frontiers (CF '24)*, May 7–9, 2024, Ischia, Italy. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3649153.3649186>

1 INTRODUCTION

The cerebellum is crucial for facilitating motor control and hand-eye coordination, among other critical functionalities [23]. In order to unveil the mechanisms underlying its operation, neuroscientists constantly seek to record and understand its activity in living test subjects. The – mostly invasive – nature of these experiments often dictates the use of animal subjects, such as mice [6]. One popular methodology relies on electrophysiological recordings of neural activity in various types of cells, especially the Purkinje cells in the cerebellum because of their critical role in motor coordination [23]. In current setups, a *wired* connection is used between the mouse and an acquisition device, as shown in Figure 1. However, such setups do not mimic natural conditions since the wires do not allow *free movement* of mice. Therefore, in essence, we need to get rid of these wires to enable more *realistic* neuroscientific experiments.

Several *wireless* head stages for different mammals have been proposed [2, 3, 10, 11, 19, 29]. However, they have *at least* one of these two limitations: (1) They are too heavy and, thus, are only suitable for larger animals (in the case of [2, 3, 11, 19]). Simply put, the whole head stage, including the battery, needs to be < 3 grams for mice for the aforementioned cerebellum experiments [6]. (2) The head stages can record for only a short period of time, e.g., 30 and 105 minutes in the case of [29] and [10], respectively, whereas the aforementioned experiments require up to 24 hours of

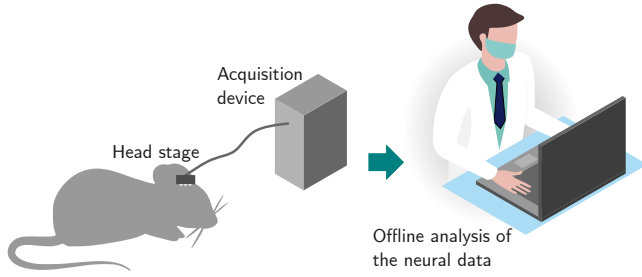


Figure 1: A typical wired neural-signal acquisition setup, which limits the free movement of mice.

neural recordings. This is because the study of learned motor control, information processing, memory consolidation, interactions among distributed brain regions etc., requires long timescales [5], see Table 1. In short, *suitable* wireless head stages for *mice* do not yet exist, which is primarily due to the difficulty in constructing on-site neural-data processing and a wireless transceiver in a small form factor. For instance, algorithms for processing neural data (including that from Purkinje cells), such as [14, 15, 25, 31], are generally designed in software for offline analysis and have not yet reached the required form factor to be incorporated within a mouse’s head stage.

This work addresses both of the aforementioned limitations by introducing a novel approach for real-time detection and classification of neuronal spikes in the sampled neural data from Purkinje cells. The purpose of this scheme is to reduce the dimensionality of this sparse data so that only a small but essential part is stored on a removable storage within the head stage, which alleviates the need for wires. This heavily condensed stored data can then be retrieved after an experiment for offline analysis. In this way, our scheme allows to reach the ultimate *goal* of conducting *long-duration* experiments involving *freely moving* mice. In essence, this work makes the following key contributions:

- A *neuronal-spike classification scheme* that throws out unnecessary information from the stream of data from the Purkinje cells, taking advantage of their unique characteristics, which simplifies the data-storage requirements.
- A *lightweight system architecture* that consists of a controller that orchestrates the data flow from the spike detector to the classification module and ultimately stores the classified data in a non-volatile memory (STT-RAM).
- A *small-form-factor synthesized CMOS design* that enables low-power operation for a prolonged duration of time while staying within the head-stage size requirements.

The rest of the paper is organized as follows. Section 2 provides a brief neuroscientific background of the experiments involving cerebellar recordings. Section 3 explains our proposed scheme followed by the results in Section 4. We draw overall conclusions in Section 5.

2 BACKGROUND

2.1 Neuroscientific Background

Techniques developed during the second half of the nineteenth century allowed neuroscientists to investigate the electrical activity of an animal brain directly using electrophysiological recordings. A particular area of interest is the use of such techniques to acquire and analyze the neuronal activity of the Purkinje cells in the cerebellum. A detailed list of animal behaviours that could be studied using such experiments is provided in Table 1. These experiments shed light on how the activity of a particular brain area, in this case the cerebellum, relates (and possibly give rise) to a particular sensory-motor or cognitive function. This could lay the foundation for a better understanding of the brain and could, potentially, have application in treating brain injuries and loss of motor functions [23]. Mice are usually preferred for these studies since they are easy to manage and very well-suited for genetic manipulations, which significantly increases the number of research questions that can be addressed.

Current experiments involving neural recordings from mice have one major constraint: the animals are head-fixed within the experimental setup that only allows them some body movement, which is unnatural. For instance, when a certain neuronal signal appears to be associated with a particular movement, it is not possible to establish if that neuronal signal is also associated with the movement of other parts of the body that are restrained or immobilized. The size and weight of the head stage (i.e., it should be roughly <3 grams for mice) and the recording duration (i.e., it should support 24-hour-long sessions) constitute the primary challenges in developing technologies that meet the goal of ensuring the natural states of animals during experiments. Significant effort has been devoted over the past decade to addressing these challenges, aiming to enable freely moving animal recordings, as discussed next.

2.2 Related Work

Bilodeau et al. [2] and Gagnon-Turcotte et al. [10] present head stages that utilize Spartan-6 FPGAs for real-time processing of recorded neural signals and wireless transceivers to transmit the reduced data to an external device (base station). These head stages are suitable for mice due to their reasonable weight. However, their maximum recording durations, less than 1.75 hours, are too short for the cerebellar experiments highlighted in Table 1. Grand et al. [11] propose a head stage design where recorded neural data is directly transmitted wirelessly to an external device without any onboard processing. Their system allows for an incredible 72-hour recording duration. However, the large size and weight of the complete system make it suitable only for large animals, such as monkeys. Luan et al. [14] also employ FPGA-based dimensionality reduction, similar to [2, 10]. Their design allows for an impressive 24-hour recording duration. However, like [14], it is only suitable for large animals, in their case, cats.

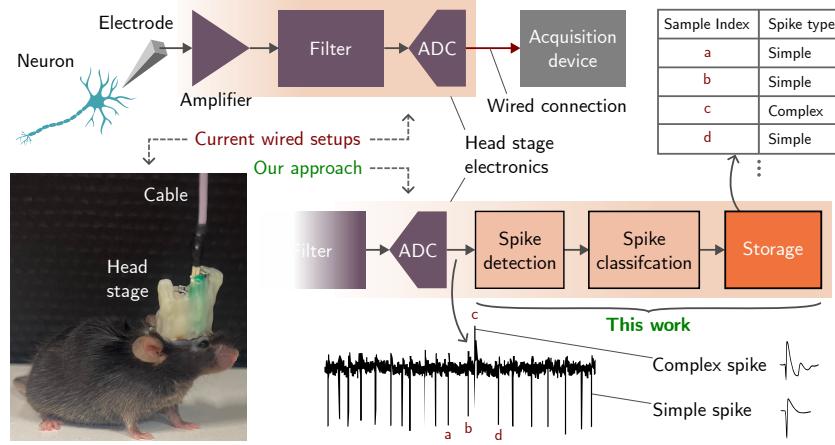
When focusing on the dimensionality reduction of neural data from Purkinje cells, there is a limited body of work available. Notably, all these works are applicable to *offline* software implementations, wherein they process raw digitized data retrieved from a wired setup. However, since these works are intended to be executed on larger computing devices (e.g., PCs, GPUs, etc.), they, in

Table 1: Examples of behaviours and paradigms that could be studied using neuronal recordings involving freely moving mice.

Behaviour	Main topics to be studied	Duration	Examples
Rhythmic movements	Motor control	5-10 minutes	Respiration, whisking, locomotion etc.
Compensatory eye movement	Simple reflex	30-60 min	VOR and OKR
Social interaction	Cognitive functions	30-60 min	Two chambers test
Sleep	Coherent oscillation	12-24 hours	Recording of spontaneous sleep
Grasping task	Learned motor control	Multiple sessions across 2-3 days*	Food pellet or water reaching task
Adaptation of compensatory eye movement	Memory formation	Multiple sessions across 4-5 days*	VOR/OKR adaptation
Eyelid conditioning	Associative learning	Multiple sessions across 6-7 days*	Delay eyelid conditioning
Sensory discrimination task	Cognitive learning	Multiple sessions across 16-20 days*	Directional licking task

*VOR: Vestibulo-ocular Reflex, *OKR: Optokinetic Response

*A recording session can take up to 24 hours.

**Figure 2: Current wired setups vs. our approach, which gets rid of the wire and enables free movement of mice.**

their current form, cannot be executed on a small head stage due to the dearth of computing resources available onboard. These works encompass the research by Sedaghat-Nejad et al. [25], Markanday et al. [15], and Zur et al. [31].

In addition to literature, commercial solutions are also available. For example, the FreeLynx wireless acquisition system from NeuroLynx [19] provides a convenient setup for enabling freely-moving animal recordings. However, its smallest head-stage configuration has a weight of 21 g, making it suitable only for larger animals, such as rats and monkeys. A similar issue is observed with the Cereplex Exilix system from Blackrock Microsystems [3], which features a 9.87 g head stage and a limited recording duration of up to 2.5 hours. The eCube wireless head stage from White Matter [29] shows promise with a 2.5 g head stage, including the battery, making it suitable for mice. Unfortunately, the recording duration is only 30 minutes, which is inadequate for the aforementioned cerebellar experiments. Hence, there is still a pressing need for new head-stage technologies that are small, light, low-power, and wireless.

3 PROPOSED SOLUTION

Our approach deviates from using a wireless transceiver to get rid of the wires during neural-signal acquisition from mice. This is because a wireless transceiver involves a complex design that impacts

system reliability and, of course, power consumption. Instead, we take advantage of the fact that the signal analysis related to the aforementioned cerebellar experiments is usually done *after* the recording phase, i.e., offline. As a result, our proposal to eliminate the wires is to store the neural data on a removable storage on the head stage, which can then be retrieved for offline analysis (see Figure 2). However, this approach, as is, does not scale well for larger number of channels and longer experiments (i.e., up to 24 hours) since it would require a significantly large amount of storage.

We solve this problem by leveraging the unique characteristics of the Purkinje cells and keeping only the information of interest from the very sparse neural recording. More specifically, this crucial information (for the experiments discussed in Section 1) is (a) the time when a spike occurred and, (b) the type of spike, i.e., *simple* or *complex* (since a Purkinje-cell spike can be of one of these two types [15], as shown in Figure 2). This results in storing only a condensed version of the complete recording (see Figure 2, top right), which significantly eases up storage requirements and enables a small-form-factor implementation.

An overview of the proposed scheme is shown in Figure 3. The detector module, which is based on [30], detects the occurrence of a spike in an input stream of digitized neural data. After this detection, the classifier module is enabled, which determines whether the spike is simple or complex. Both these modules are explained next.

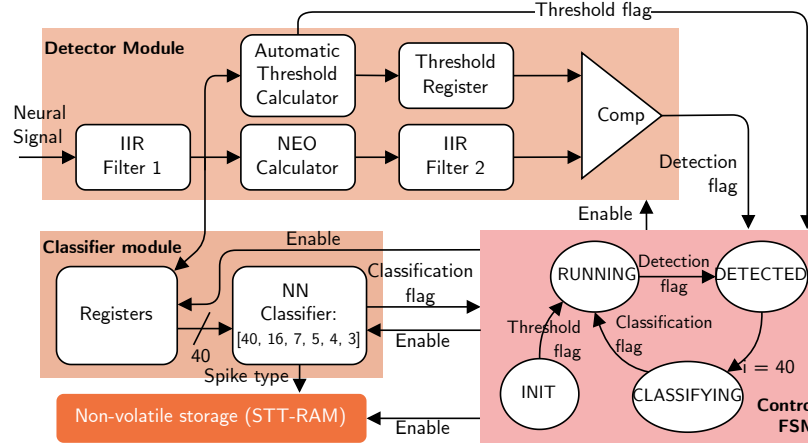


Figure 3: Proposed system architecture.

3.1 Spike Detection

The spike-detection module is based on the work of Yang et al. [30]. This specific design is chosen because it can support the Purkinje-cell firing rate of around 100 Hz while having a small form factor and being extremely power efficient.

The digitized input samples are first filtered using a 1st-order exponential IIR filter, which smooths the input signal to mitigate high-frequency noise components that adversely affect the performance of the spike detector. The filtered signal is then fed to a non-linear energy operator (NEO) [17] and an accompanying 1st-order exponential IIR filter to get the instantaneous estimation of signal energy, which is then compared with a calculated threshold value. If the energy estimate (NEO value) is higher than the threshold then a spike is detected, and vice versa. The threshold calculator is able to dynamically (re)calculate the threshold and minimize the performance loss in case of changes to the input signal, e.g., due to the worsening of the noise level.

3.2 Spike Classification

After the detection of a spike, the spike waveform (i.e., a certain number of samples after the detection event) is saved in order for it to be classified as either being simple or complex. The length of this waveform is set as 40 samples to sufficiently capture the whole spike duration (~ 1.7 ms [15]) considering the sampling frequency of 24.414 kHz. We chose to employ a neural-network (NN)-based design for this classifier. This is because NNs tend to generalize well when there is plenty of training data, as in our case. Such a classifier would, therefore, be robust enough to handle varying signal conditions. For instance, the recorded data set used in this work exhibits significant variations in terms of noise floor, recording-electrode drift, saturated sample points, and data offsets.

3.2.1 Choice of NN Topology. The highly resource-constrained nature of the head stage necessitates the use of a lightweight NN. Thankfully, the classification problem itself is a *ternary* classification problem (i.e., classifying a spike as simple, complex or neither), which means that the NN needs only three output neurons. We

chose to tackle this relatively simple classification problem using a *multilayer perceptron*, i.e., a fully connected feedforward artificial neural network.

Keeping the NN lightweight also implies finding an optimal trade-off between the classification accuracy and area/energy consumption. Thus, an extensive design space exploration (DSE), which will be discussed in Section 4.3, was conducted concerning the choice of NN topology, including the number of hidden layers and neurons per layer, and its impact on classification accuracy. The NN has 40 inputs to correspond with the aforementioned 40-sample spike waveform.

3.2.2 Choice of Activation Functions. For the hidden layers, a multitude of activation functions exist to allow the network to learn the necessary nonlinearities in the classification problem. Popular choices include various sigmoidal-shaped functions and ReLU (rectified linear unit). As hardware implementation of nonlinear functions is very costly, it was decided to choose ReLU for the hidden layers since it is the least computationally-expensive option.

To calculate the final probabilities in the *output* layer of an NN for multi-class classification problems like this spike-classification task, the Softmax function is commonly employed. However, this is a very costly operation in hardware, which would also require IEEE floating point support. Instead, Softmax can be moved to the loss function, effectively making the network output log probabilities (logits) instead of just probabilities. This can be achieved using a simple linear activation function that carries no computational cost during inference. Since the above natural logarithm is a monotonically increasing function of probability, the output class is determined by selecting the index of the neuron with the highest value, just like in the case of using normal probabilities.

3.2.3 NN Quantization. The NN was initially designed and tested in software before implementing it in hardware (see Section 4.1 for more details on the tool flow). Conventionally, software NN implementations use floating-point arithmetic which adds to the computational burden of the classifier. Through the quantization technique presented in [12], we transformed the NN to integer

format to increase the efficiency of the inference. Following this technique, the weights and activation values of the neurons were mapped to signed 8-bit values.

3.3 Control

Given the fact that the occurrence of Purkinje-cell spikes is relatively rare (i.e., roughly 100 Hz on average) compared to the sampling frequency, constantly running all of the modules in the proposed architecture (Figure 3) is unnecessary. In fact, once the threshold has been calculated for the detection, the only block in the system that runs continuously is the NEO calculator and the comparator. The classifier is only enabled once a spike has been detected. Furthermore, there is no need to detect another spike during the classification execution on the initial spike as there is a sufficient time gap (> 2 ms) between two consecutive spikes. To control the information flow and enabling/disabling modules, a finite-state machine (FSM) was designed, whose state diagram is shown in Figure 3.

The system is in the *INIT* state when it is initially started or is reset. During *INIT* a threshold is calculated for the spike detection and all the other subsystems are disabled. When the threshold calculator converges to a threshold, a flag is set. At this point, the FSM transitions to the *RUNNING* state in which the NEO calculator is enabled and spike detection is performed. At this point, there is no need to enable the threshold calculator or classifier. When a calculated NEO value exceeds the threshold, a detection flag is set and the FSM transitions to the *DETECTED* state. During this state, the next 40 filtered input samples are saved. Here, the precision of the samples is reduced to 8 bits as mentioned earlier in Section 3.2.3. After these 40 cycles, the FSM transitions to the *CLASSIFYING* state in which the NN inference is performed and the classified result is stored in the storage. Upon completion, the system transitions back to the *RUNNING* state.

3.4 Storage

One of the important aspects of our approach is the use of on-board storage to store the reduced-size (i.e., classified) data. Since this memory needs to be removed after an experiment to retrieve the recordings, it needs to be non-volatile. Another reason for employing a non-volatile memory is to allow it to stay powered down when the head stage is not classifying, which cuts out its leakage power and improves battery life. For the proposed scheme, we chose the *Spin Transfer Torque Magnetic Random Access Memory* (STT-RAM). This is because apart from non-volatility, STT-RAM has a small form factor, low access latency and energy, high endurance, CMOS compatibility, high maturity and immunity to soft-errors due to radiations [1].

Table 2 illustrates a comprehensive comparison between STT-RAM and various memory technologies, encompassing both conventional and emerging non-volatile memories. Conventional memories like SRAM and DRAM are volatile, necessitating a continuous power supply. SRAM notably faces leakage issues, and DRAM necessitates periodic refresh cycles, rendering it energy-inefficient. Flash memory, another conventional technology, operates at relatively high voltages. Among the emerging non-volatile technologies, RRAM encounters endurance problems, while PCM demands




Figure 4: Relative timing information between complex and simple spikes, which can aid in the offline post-processing

high voltage for switching. Although STT-RAM boasts numerous advantages, it has the drawback of requiring a constant current for writing, resulting in slightly higher energy consumption and latency compared to the traditional SRAM technology. However, this drawback is offset by SRAM's leakage issues. Due to all these benefits, STT-RAM stands out as the only commercially available emerging non-volatile technology in the market due to its overall performance and reliability [9, 18]. Furthermore, since we do not need STT-RAM's prolonged data-retention (of up to 10 years) for the targeted experiments, we take advantage of the tunability of its thermal-stability factor by which we shorten the retention time in return for further improving its energy efficiency and access latency [28]. These features make STT-RAM very suitable for our experiments.

3.5 Post-processing

The calculated NEO value of the spike detector can sometimes exceed the threshold several times in a short duration, leading to one actual spike causing multiple spike detections. To address this limitation of the spike-detection algorithm, a post-processing step is carried out *offline*, i.e., on the reduced data retrieved from the head-stage storage. This step leverages the relative spike timing between the spikes, as illustrated in Figure 4. The time interval between these (false) detection events often falls within the sub-millisecond domain, which is anatomically implausible for Purkinje cells: Neurons require time to accumulate the necessary charge (i.e., ion concentration) to generate a spike, typically exceeding 4 ms [7]. As a result, a dead zone is introduced in post-processing after a detection event. However, since the minimum spike interval pertaining specifically to simple spikes remains relatively consistent, the dead zone is applied *only* after the classification of simple spikes, and any detection events within this period are subsequently discarded.


4 RESULTS

4.1 Experimental Setup

An overview of the employed tool chain is shown in Figure 5. The data used to design, train and verify the system was formatted in Matlab. This data set was acquired using a wired setup (see Figure 2) involving the Mini-Amp-64 head stage from Cambridge NeuroTech [4] with the ADC sampling frequency of 24.414 kHz

Table 2: Comparison of bit-cell design metrics for various memory technologies (data obtained from [20, 24]).

Metrics	SRAM	DRAM	Flash	RRAM	STT-RAM	PCM
Size (F^2)	120–150	10–30	10–30	10–30	10–30	10–30
Volatility	Yes	Yes	No	No	No	No
Write energy	~fJ	~10 fJ	~100 pJ	~1 pJ	~1 pJ	~10 pJ
Write speed	~1 ns	~10 ns	0.1–1 ms	~10 ns	~5 ns	~10 ns
Read speed	~1 ns	~3 ns	~100 ns	~10 ns	~5 ns	~10 ns
Endurance	10^{16}	10^{16}	10^4 – 10^6	10^7	10^{15}	10^{12}
Scalability	medium	medium	medium	high	high	high

**Figure 5: Employed tool flow**

and a resolution of 10 bits. To enable flexibility and multiple DSE iterations, the spike detector and classifier were first implemented in software (Python). Python was employed for this purpose because of its ease of use and well-supported libraries (especially in regards to the NN-based classifier). The hardware was described in VHDL and simulated using Xilinx Vivado. The classifier NN training was done using TensorFlow. The training data itself was sanitized using Uniform Manifold Approximation and Projection for dimension reduction (UMAP) [16]. Subsequently, TFLite was used to optimize (quantize) the TensorFlow models [12]. The Netron software [22] was used to visualize the neural network and extract the weights and biases, which were then used to create a hardware description of the classifier. For our evaluation, we have synthesized the design on Cadence Genus, using the 45-nm NanGate Open-Cell library [26]. The design metrics for the STT-RAM were extracted using the NVSim tool [8].

4.2 Spike-Classifer-NN Training

The acquired neural-recording contains around 10^6 spikes, which were *manually* annotated offline in order to create a reference for training the spike-classifier NN. 40 samples after every NEO-filter-detected sample were saved in order to cover the rough spike duration (as mentioned in Section 3.2). Based on the annotated data, these detected waveforms were labeled as *Complex Spike (CS)*, *Simple Spike (SS)*, or *False positive (F)* (for detected signals that did not correspond to annotated spikes). Since these samples were in a 10-bit format (ADC resolution), the whole waveform was divided by 4 since the NN requires 8-bit inputs.

Furthermore, since the occurrence of complex spikes is far less frequent than that of simple spikes (i.e., there is a large class imbalance), the training data set was re-sampled to obtain a balanced set of simple and complex spikes, and false positives (5953 SS, 5953 CS and 5948 F samples, respectively). This process prevents the NN from being biased towards simple spikes. This reduced data set was further split randomly into the *training set* (80%) and *test set* (20%).

The training was performed over 100 epochs and *early stopping* with a 90%/10% training/validation split on each fold.

The training samples underwent another round of filtering by comparing them against the output of UMAP, which was performed on the aforementioned *balanced* data set (in two dimensions). It was found that a small amount of samples seemed to be outliers or were very similar to other classes, which could be due to possible mislabeling or missing spikes altogether during manual annotation. To avoid confusing the training protocol with these outliers, we discarded samples when the 10 nearest neighbours in the UMAP space contained 9 or more differently-labeled samples at the time of training (Figure 6). It is important to note that the outliers were not discarded during the *testing* of the NN (Section 4.4).

The loss function used for training is the *multi-class cross-entropy*, calculated directly from logits to remove Softmax-computation overhead from the network output layer. As the name suggests, this loss function is well-suited for a ternary-classification problem, such as this one. The weights were updated using the popular Adam optimizer [13].

4.3 Classifier Design Space Exploration

The DSE of the spike-classifier NN was performed to find an optimal trade-off between the classification accuracy and area/energy consumption, which in turn depends on the NN topology (i.e., the number of hidden layers and the number of neurons per layer). To gauge the classification accuracy, we employed the classical formula of $\frac{T_n + T_p}{T_n + T_p + F_n + F_p}$, where T_n and T_p are the true negatives and positives, respectively and F_n and F_p are the false negatives and positives, respectively. To obtain the most ideal NN topology, a 10-fold cross-validated grid search was performed over NN architecture and regularization. The search included all possible hidden layer configurations between 0 and 4 layers, in descending neuron counts for each layer, i.e., 1–40, 1–20, 1–10, and 1–10, respectively. Furthermore, the grid search contained a 0.01 or 0.001 regularization constraint for the weight-matrix orthogonality. It should be noted that all accuracy estimates were performed *after quantization*, i.e., it was incorporated into the DSE.

As CS classification is the hardest task, the most optimal network after the grid search was selected by first choosing those with a CS class accuracy greater than 90% within a 95% confidence interval. This criterion was applied to closely match the performance of the recent software-based technique [15] discussed in Section 1. Subsequently, the network with the lowest computational complexity

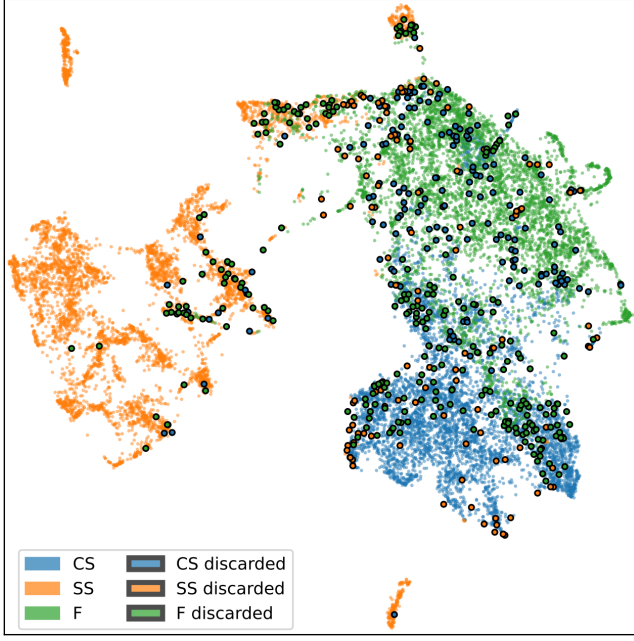


Figure 6: UMAP projection of the *class-balanced* dataset and highlighted outliers to be discarded during training

Table 3: Network-architecture search results

Network Architecture	RF	CS*	SS*	F*	Complexity**
40, 2, 3	0.001	85.4%	93.4%	83.0%	86
40, 2, 3	0.010	86.4%	93.7%	82.2%	86
40, 4, 3, 3	0.010	88.0%	94.3%	84.4%	181
40, 5, 5, 2, 3	0.010	88.9%	93.5%	82.8%	241
40, 7, 7, 4, 3, 3	0.010	90.0%	93.5%	83.0%	378
40, 8, 8, 3, 3, 3	0.001	90.6%	94.2%	84.0%	426
40, 14, 10, 4, 3, 3	0.010	91.3%	94.8%	86.2%	761
40, 16, 7, 5, 4, 3	0.010	91.7%	94.5%	85.1%	819
40, 28, 14, 8, 6, 3	0.010	93.0%	94.5%	89.0%	1690

*RF: Regularization Factor

* Mean post-quantization accuracy calculated across 10-fold cross-validation

** Complexity defined as $\sum_{i \in \text{layers}} n(i)n(i+1)$.

was selected, as estimated by $\sum_{i \in \text{layers}} n(i)n(i+1)$, where $n(i)$ represents the size of the layer. The network-architecture optimization results are shown Table 3. The optimal NN topology consisted of four hidden layers with 16, 7, 5, and 4 neurons, respectively, along with the orthogonal weight regularization set to 0.01.

4.4 Classifier Testing

Finally, the aforementioned optimal network was retrained and quantized on the full training set, and was subsequently tested on the complete *test data set* (which remains unseen until now), leading to the confusion matrix shown in Figure 7 (top left). We can see

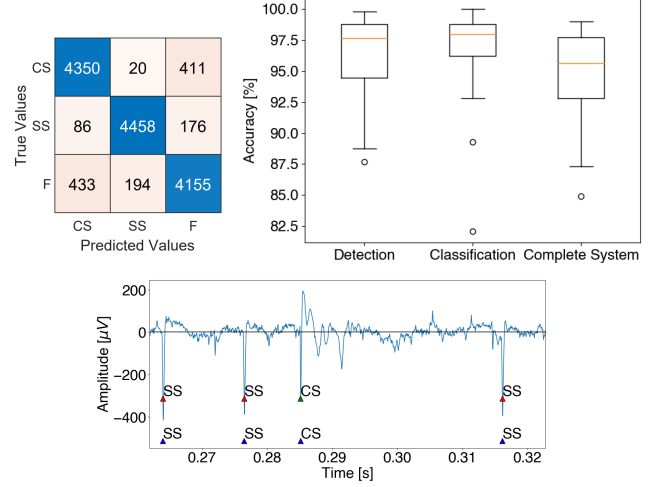


Figure 7: Top left: Final confusion matrix of the classifier NN on the test data set. Top right: Accuracy of our scheme for different neural recordings. Bottom: Manual classification vs. automated classification using our scheme. Blue triangles denote the manually-labeled spikes and the red/green triangles denote the detected simple (SS) and complex (CS) spikes, respectively.

Table 4: Implementation results

	Detector	Classifier	Storage (STT-RAM)
Energy* (nJ)	4.46	311	0.28
Area (mm ²)	0.006	0.081	25.91**

*One complete detection-classification-storage cycle

** For 32-MB capacity to support 24-hour experiments

that both the CS (93.35%) and SS (96.67%) classification accuracies satisfy the target of >90%.

Figure 7 (top right) shows that for the selected NN topology the median accuracy values of the individual blocks and the complete system are greater than 95%. Figure 7 (bottom) shows a snippet of a neural recording in which our scheme detects and classifies the spikes with the same outcome as that of the manual annotation.

4.5 Energy- and Area-Efficiency Analysis

The results of the logic synthesis with the target clock frequency of 24.414 kHz (to match the sampling rate of the input neural data) are summarized in Table 4. It can be seen that most of the energy is spent on the classifier during one complete detection-classification-storage cycle. However, the classification is only invoked roughly a 100 times per second (Purkinje-cell firing rate), which results in a very-low system energy consumption and sufficiently long battery life (as will be discussed shortly). In terms of area, the complete detector and classifier design, and the STT-RAM occupy just under 26 mm², which can be easily accommodated in ~200-mm² PCBs typically used in mouse head stages [6]. Figure 8 (top) shows that the dimensionality reduction performed by our classification scheme significantly reduces the storage requirements.

Table 5: Comparison with the state of the art

	[19]	[2]	[3]	[29]	[10]	[11]	[14]	[25]	[15]	[31]	This work
CS-classification accuracy (%)	N/A ₁							–	–	–	93.35
CS-classification F1 score (%)	N/A ₁							–	92.5	–	90.15
SS-classification accuracy (%)	N/A ₁							–	–	–	96.67
SS-classification F1 score (%)	N/A ₁							–	–	–	94.93
Commercial product	yes	no	yes	yes	no	no	no	no	no	no	no
Purkinje-cell spike classification	no	no	no	no	no	no	no	yes	yes	yes	yes
On-board DR	no	yes	no	no	yes	no	yes	N/A ₂			yes
Weight (g)	>21	4.68	9.87	2.5	4.9	–	–	N/A ₂			<2.5*
Autonomy (hrs)	3	1.6	2.5	0.5	1.75	72	24	N/A ₂			24**
Targeted animal	rats	mice	rats	mice	mice	cats	monkeys	N/A ₂			mice

‘–’: Not provided, ‘N/A₁’: Not applicable since these works do not deal with Purkinje cells, ‘N/A₂’: Not applicable since these works are software based (i.e., not suitable for an animal head stage)

*Worst-case approximation based on the battery choice and design size.

**For the 24-hour-experiment configuration (32-MB STT-RAM).

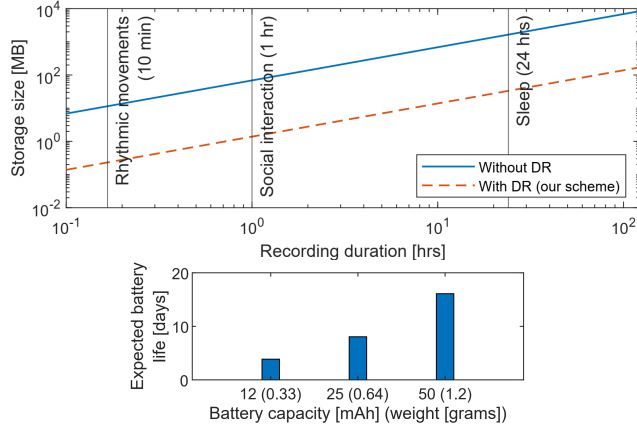


Figure 8: Top: Storage-size savings when employing dimensionality reduction (DR) via the proposed approach for different types of cerebellar experiments. Bottom: Battery life of the head-stage when using the proposed approach for different miniature battery capacities [21].

The next step is to calculate the battery life of the complete head stage in order to prove that it can stay operational throughout the longest of experiments (i.e., 24 hours). Figure 8 (bottom) shows the expected battery life of the head stage for three very-small-sized batteries [21] using an ADC with 0.5 pJ per conversion [27]. It can be seen that our approach easily allows for roughly 4 days of continuous operation for the smallest 0.33-gram battery (12 mAh).

4.6 Discussion

Recall from Section 1 that the goals of our scheme were to achieve (1) *long-duration* experiments, and (2) *freely moving* mice. The first goal can be validated from Figure 8 in which we demonstrated that our scheme easily allows the head stage to operate continuously for the longest-duration experiments. Regarding the second goal, we saw that the use of an on-board storage allows us to get rid

of the wires, which aids in the free movement of mice. Moreover, the design supports the use of very light batteries (down to 0.33 grams). Since the battery is the heaviest component of the head stage, our approach easily meets the requirement of the head-stage weight being less than 3 grams. Table 5 compares our work with the closely-related state of the art. It can be seen that our scheme, to the best of our knowledge, is the only one to date that fulfills the aforementioned goals. Moreover, it is the only one that performs small-form-factor Purkinje-cell spike classification.

5 CONCLUSIONS

In this paper, we proposed a lightweight architecture for classifying Purkinje-cell spikes from a stream of neural data. Our scheme reduces the dimensionality of this very sparse data by classifying spikes with an overall accuracy of >95%, which allows it to be stored on a removable storage on a mouse’s head stage. This results in getting rid of the wires for data acquisition and enabling the free movement of mice during cerebellar-recording experiments. Moreover, our CMOS synthesis results demonstrate that our small-form-factor approach allows for long-duration experiments. The head stage utilizing our scheme can continuously run for approximately 4 days on a small 0.33-gram (12 mAh) battery. Such experiments involving free-moving mice (i.e., more realistic and natural conditions) can significantly help in elucidating the underlying mechanisms behind motor control and loss thereof, in the brain.

ACKNOWLEDGMENTS

This paper is supported by the European Union’s Horizon Europe research and innovation programme under projects SEPTON (Gr. Agr. No. 101094901) and SECURED (Gr. Agr. No. 101095717) and by the Dutch Research Council’s Gravitation programme under project DBI² (No. 024.005.022).

REFERENCES

- [1] Simone Bertolazzi, Paolo Bondavalli, Stephan Roche, Tamer San, Sung-Yool Choi, Luigi Colombo, Francesco Bonaccorso, and Paolo Samori. 2019. Nonvolatile memories based on graphene and related 2D materials. *Advanced materials* 31, 10 (2019), 1806663.

- [2] Guillaume Bilodeau, Gabriel Gagnon-Turcotte, Léonard L Gagnon, Iason Keramidis, Igor Timofeev, Yves De Koninck, Christian Ethier, and Benoit Gosselin. 2021. A wireless electro-optic platform for multimodal electrophysiology and optogenetics in freely moving rodents. *Frontiers in Neuroscience* 15 (2021), 718478.
- [3] Blackrock Microsystems, LLC. 2020. *Cereplex Exilis - Instructions for Use*.
- [4] Cambridge NeuroTech. 2020. *Mini-Amp-64 User Guide - Version 1.0*.
- [5] Jason E Chung, Hannah R Joo, Jiang Lan Fan, Daniel F Liu, Alex H Barnett, Supin Chen, Charlotte Geaghan-Breiner, Mattias P Karlsson, Magnus Karlsson, Kye Y Lee, et al. 2019. High-density, long-lasting, and multi-region electrophysiological recordings using polymer electrode arrays. *Neuron* 101, 1 (2019), 21–31.
- [6] Andres de Groot, Bastijn JG van den Boom, Romano M van Genderen, Joris Coppens, John van Veldhuijzen, Joop Bos, Hugo Hoedemaker, Mario Negrello, Ingo Willuhn, Chris I De Zeeuw, et al. 2020. NINScope, a versatile miniscope for multi-region circuit investigations. *Elife* 9 (2020), e49987.
- [7] Chris I De Zeeuw, Freek E Hoebeek, Laurens WJ Bosman, Martijn Schonewille, Laurens Witter, and Sebastiaan K Koekkoek. 2011. Spatiotemporal firing patterns in the cerebellum. *Nature Reviews Neuroscience* 12, 6 (2011), 327–344.
- [8] Xiangyu Dong, Cong Xu, Yuan Xie, and Norman P Jouppi. 2012. NVSim: A circuit-level performance, energy, and area model for emerging nonvolatile memory. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 31, 7 (2012), 994–1007.
- [9] Everspin Technologies. 2024. *Spin-transfer Torque MRAM Technology*. <https://www.everspin.com/>
- [10] Gabriel Gagnon-Turcotte, Yoan LeChasseur, Cyril Bories, Younès Messaddeq, Yves De Koninck, and Benoit Gosselin. 2017. A wireless headstage for combined optogenetics and multichannel electrophysiological recording. *IEEE transactions on biomedical circuits and systems* 11, 1 (2017), 1–14.
- [11] Laszlo Grand, Sergiu Ftomov, and Igor Timofeev. 2013. Long-term synchronized electrophysiological and behavioral wireless monitoring of freely moving animals. *Journal of neuroscience methods* 212, 2 (2013), 237–241.
- [12] Benoit Jacob, Skirmantas Kligys, Bo Chen, Menglong Zhu, Matthew Tang, Andrew Howard, Hartwig Adam, and Dmitry Kalenichenko. 2018. Quantization and training of neural networks for efficient integer-arithmetic-only inference. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. IEEE, 2704–2713.
- [13] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [14] Song Luan, Ian Williams, Michal Maslik, Yan Liu, Felipe De Carvalho, Andrew Jackson, Rodrigo Quián Quiroga, and Timothy G Constandinou. 2018. Compact standalone platform for neural recording with real-time spike sorting and data logging. *Journal of neural engineering* 15, 4 (2018), 046014.
- [15] Akshay Markanday, Joachim Bellet, Marie E Bellet, Junya Inoue, Ziad M Hafed, and Peter Thier. 2020. Using deep neural networks to detect complex spikes of cerebellar Purkinje cells. *Journal of neurophysiology* 123, 6 (2020), 2217–2234.
- [16] Leland McInnes, John Healy, and James Melville. 2018. UMAP: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426* (2018).
- [17] Sudipta Mukhopadhyay and GC Ray. 1998. A new interpretation of nonlinear energy operator and its efficacy in spike detection. *IEEE Transactions on biomedical engineering* 45, 2 (1998), 180–187.
- [18] Sarath Mohanachandran Nair and othes. 2018. Defect injection, fault modeling and test algorithm generation methodology for STT-MRAM. In *International Test Conference (ITC)*. IEEE, 1–10.
- [19] NeuraLynx, Inc. 2021. *FreeLynx - User Manual*.
- [20] Fabian Oboril, Rajendra Bishnoi, Mojtaba Ebrahimi, and Mehdi B Tahoori. 2015. Evaluation of hybrid memory technologies using SOT-MRAM for on-chip cache hierarchy. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 34, 3 (2015), 367–380.
- [21] PowerStream Technology. 2022. *Ultra low weight lithium ion batteries*. <https://www.powerstream.com/ultra-light.htm>
- [22] Lutz Roeder. 2022. *Netron*. <https://github.com/lutzroeder/netron>
- [23] Vincenzo Romano, Aoibhinn L Reddington, Silvia Cazzanelli, Roberta Mazza, Yang Ma, Christos Strydis, Mario Negrello, Laurens WJ Bosman, and Chris I De Zeeuw. 2020. Functional convergence of autonomic and sensorimotor processing in the lateral cerebellum. *Cell reports* 32, 1 (2020), 107867.
- [24] Sayeef Salahuddin et al. 2018. The era of hyper-scaling in electronics. *Nature Electronics* 1, 8 (2018), 442–450.
- [25] Ehsan Sedaghat-Nejad, Mohammad Amin Fakharian, Jay Pi, Paul Hage, Yoshiko Kojima, Robi Soetedjo, Shogo Ohmae, Javier F Medina, and Reza Shadmehr. 2021. P-sort: an open-source software for cerebellar neurophysiology. *Journal of neurophysiology* 126, 4 (2021), 1055–1075.
- [26] Silicon Integration Initiative, Inc. 2019. *15nm Open-Cell Library and 45nm FreePDK*. <https://si2.org/open-cell-library/>
- [27] Abhairaj Singh, Muath Abu Lebdeh, Anteneh Gebregiorgis, Rajendra Bishnoi, Rajiv V Joshi, and Said Hamdioui. 2021. SRIF: Scalable and reliable integrate and fire circuit adc for memristor-based cim architectures. *IEEE Transactions on Circuits and Systems I: Regular Papers* 68, 5 (2021), 1917–1930.
- [28] Clinton W Smullen, Vidyabhushan Mohan, Anurag Nigam, Sudhanva Gurumurthi, and Mircea R Stan. 2011. Relaxing non-volatility for fast and energy-efficient STT-RAM caches. In *2011 IEEE 17th International Symposium on High Performance Computer Architecture*. IEEE, 50–61.
- [29] White Matter, LLC. 2020. *Details of the eCube standalone and stacking headstages*. <https://docs.white-matter.com/docs/ecube/hardware/headstages/>
- [30] Yuning Yang and Andrew J Mason. 2016. Hardware efficient automatic thresholding for NEO-based neural spike detection. *IEEE Transactions on Biomedical Engineering* 64, 4 (2016), 826–833.
- [31] Gil Zur and Mati Joshua. 2019. Using extracellular low frequency signals to improve the spike sorting of cerebellar complex spikes. *Journal of Neuroscience Methods* 328 (2019), 108423.