

Document Version

Final published version

Licence

Dutch Copyright Act (Article 25fa)

Citation (APA)

Rudge, Z. A., Dold, D., Fieback, M., Izzo, D., & Hamdioui, S. (2026). Memristor-based neural network accelerators for space applications: Enhancing performance with temporal averaging and SIRENs. *Acta Astronautica*, 238, 656-667. <https://doi.org/10.1016/j.actaastro.2025.10.049>

Important note

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

Copyright

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership. Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

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](#)
as part of the Taverne amendment.**

More information about this copyright law amendment
can be found at <https://www.openaccess.nl>.

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.



Research paper

Memristor-based neural network accelerators for space applications: Enhancing performance with temporal averaging and SIRENS[☆]

Zacharia A. Rudge^{a,b}, Dominik Dold^c, Moritz Fieback^a, Dario Izzo^b, Said Hamdioui^a

^a TU Delft, Mekelweg 5, Delft, The Netherlands

^b European Space Agency, Keplerlaan 1, Noordwijk, The Netherlands

^c University of Vienna, Kolingasse 14-16, 1090 Wien, Austria

ARTICLE INFO

Keywords:

Memristors
On-board AI
Guidance
Navigation
And control
Geodesy
Emerging hardware systems

ABSTRACT

Memristors are an emerging technology that enables artificial intelligence (AI) accelerators with high energy efficiency and radiation robustness — properties that are vital for the deployment of AI on-board spacecraft. However, space applications require reliable and precise computations, while memristive devices suffer from non-idealities, such as device variability, conductance drifts, and device faults. Thus, porting neural networks (NNs) to memristive devices often faces the challenge of severe performance degradation. In this work, we show in simulations that memristor-based NNs achieve competitive performance levels on on-board tasks, such as navigation & control and geodesy of asteroids. Through bit-slicing, temporal averaging of NN layers, and periodic activation functions, we improve initial results from around 0.07 to 0.01 and 0.3 to 0.007 for both tasks using RRAM devices, coming close to state-of-the-art levels (0.003 – 0.005 and 0.003, respectively). Our results demonstrate the potential of memristors for on-board space applications, and we are convinced that future technology and NN improvements will further close the performance gap to fully unlock the benefits of memristors.

1. Introduction

Artificial Intelligence (AI) – in particular deep learning – holds the power to revolutionize space missions by vastly increasing the autonomy of deployed spacecraft [1–5]. Recent studies have demonstrated the practical application of deep neural networks in enhancing spacecraft autonomy, including uses such as real-time optimal control for asteroid landing [6], trajectory generation and optimization [7] and minimum-fuel lunar landing [8]. However, deep learning is plagued by rapidly increasing energy requirements [9] incompatible with most spaceborne applications. Operating AI models in space imposes not only tight constraints on on-board energy consumption, but also strict requirements in terms of accuracy, area, temperature extremes, and robustness, e.g., radiation tolerance [2,10].

Existing AI accelerators fail to fully address these needs. Most deep learning implementations are currently based on conventional GPU and CPU design paradigms, which generally do not meet the requirements for on-board application due to – among other reasons – the *Von Neumann bottleneck* [11], i.e., the energy cost overhead of repeatedly moving data and instructions between the memory and the computational units, as well as technology scaling issues and their effect on

radiation resilience [12]. These challenges underscore the necessity to explore beyond conventional architectures and device technologies in order to develop hardware-software solutions that enable us to fully unlock the potential of AI in space — from low Earth orbit to deep space missions where power, communication bandwidths, and robustness requirements are pushed to the extreme. A promising technology that has recently attracted attention in the space community is neuromorphic hardware [5,10,13–17]: transistor-based chip designs that mimic the architecture of the brain to circumvent the Von Neumann bottleneck [18]. However, emerging digital neuromorphic accelerators [19], including Intel's Loihi [20] and FPGA-based implementations, do not meet all space-specific requirements for similar reasons as conventional hardware accelerators [10,14,16,21]. Similarly, analog neuromorphic accelerators [5,14,15], while promising in terms of energy consumption, currently lack the performance guarantees and scalability required for space applications.

To solve the aforementioned issues and enable AI solutions that satisfy all requirements of on-board application, we propose adopting a new chip architecture based on one of the most promising emerging technologies for edge computing: in-memory computing using memristors [22]. By integrating storage and computation into

[☆] This article is part of a Special issue entitled: 'SPAICE2024 AI in and for Space' published in Acta Astronautica.

* Corresponding author at: European Space Agency, Keplerlaan 1, Noordwijk, The Netherlands.

E-mail address: z.a.rudge@tudelft.nl (Z.A. Rudge).

a single element, chip designs based on memristors significantly reduce data movement [23], thereby alleviating the Von Neumann bottleneck. Neural network accelerators utilizing memristors, like PCM (Phase-Change Memory), RRAM (Resistive Random-Access Memory, also known as ReRAM), and MRAM (Magnetic Random-Access Memory) have demonstrated their ability to reduce energy consumption by orders-of-magnitude relative to conventional computing architectures [24–27]. Memristors are non-volatile and do not exhibit leakage currents, but also enable highly parallel computing whilst being resilient to radiation [28]. RRAM devices in particular have been demonstrated to be radiation resistant to radiation events such as single event upsets during space flight [29]. Despite these benefits, the use of memristors also poses challenges, including the introduction of variability and non-ideal behaviors (e.g. read/write noise) that can strongly affect task performance — oftentimes degrading it to levels far below those required in on-board applications [30].

In this work, we close this gap through a simulation-based approach that captures the most crucial characteristics of memristive devices while allowing the evaluation of large-scale, state-of-the-art network architectures and on-board AI applications. We demonstrate for two challenging on-board tasks that memristor-based SIREN neural networks equipped with layerwise temporal averaging reach competitive results — despite the degrading effects of device non-idealities and analog noise. Although a reality gap remains both in terms of modeled device behaviors and the conditions encountered in space, we provide a faithful evaluation of memristive devices for realistic on-board applications. More specifically, the contributions of this paper are as follows:

- **Implementation of memristor-based SIREN neural networks for realistic on-board AI applications.** We deploy and validate guidance and control networks (G&CNETs) [31] and neural networks for asteroid geodesy (geodesyNets) [32] on simulated PCM/RRAM hardware using the IBM Analog Hardware Acceleration Kit [33]. Compared to previous results [30], we show that neural networks using periodic activation functions (SIRENs [34]) perform better on simulated memristive devices than their counterparts using ReLU-like activations, with G&CNETs improving from a test loss of around 0.07 to 0.02.
- **Analysis of device non-idealities in space-relevant contexts.** We quantify the impact of device degradation and non-idealities (conductance drift, read/write noise, faulty devices) on task performance for both applications. In particular, we find that read/write noise strongly affects performance levels, while effects such as faulty devices are automatically compensated for during training in the case of geodesyNets. In contrast, G&CNETs with periodic activation functions showed strong sensitivity to faulty devices and conductance drift despite good initial performance, highlighting the need for identifying neural networks that are robust to parameter perturbations after training.
- **Recovering competitive performance via mitigation strategies.** We show that through bit-slicing and layerwise temporal averaging of neuronal outputs, competitive performance levels can be restored in both on-board applications — with the test loss improving from around 0.02 to 0.01 for G&CNETs and 0.3 to 0.007 for geodesyNets, getting close to competitive results (around 0.003 – 0.005 and 0.003, respectively). To our knowledge, this is the first work that demonstrates the feasibility of memristor-based accelerators for on-board space application tasks, showing that it is possible to reach performance levels satisfying mission requirements with this technology.

The remainder of this paper is structured as follows: Section 2 details the experimental setup, including the architecture of G&CNETs and geodesyNets, the simulation framework for the memristors and its non-idealities (e.g., conductance drift, read/write noise), and the

implementation of non-ideality mitigation strategies such as temporal averaging and linear bit-slicing. Section 3 presents the performance of both G&CNETs and geodesyNets, quantifies the impact of device non-idealities on each network, and evaluates the effectiveness of noise suppression techniques across both architectures. Finally, Section 4 discusses the significance of the results, limitations of the study, and future directions for memristor-based accelerators in on-board AI for spacecraft autonomy.

2. Methods

2.1. On-board applications and models

To evaluate the suitability of a memristor-based neural network accelerator, we first identify two space applications with classical neural network solutions that reach performance levels suitable for deployment on-board an actual spacecraft: guidance and control of a spacecraft (using G&CNETs) and asteroid geodesy (using geodesyNets), i.e., estimating an asteroid’s shape and mass distribution from sensory data. These particular neural networks share similar light-weight architectures (feedforward with 100 s of neurons and several layers), potentially allowing near-term prototyping with actual memristive crossbar arrays. Moreover, both tasks require models producing high precision outputs, making them ideal candidates for testing the capabilities of memristor-based neural networks. Finally, despite their architectural similarities, both tasks are quite different in nature: G&CNETs solve a control task where the network output is applied to change the spacecraft’s state, which in turn changes the input to the neural network. In contrast, geodesyNets solve an inverse problem which involves evaluating an integral of the neural network over a given volume. In the following, we describe both applications and neural network architectures in more detail.

2.1.1. On-board guidance and control: G&CNETs

Guidance and control networks, abbreviated as G&CNETs, are fully-connected feedforward neural networks that take the state of a spacecraft as input and return an action, e.g., the thrust to control the spacecraft. Such networks are trained on a database of optimal state-action pairs, which are pre-computed by solving the two-point boundary value problem from different initial conditions x_0 resulting from the use of Pontryagin’s maximum principle [35]. After training, the neural network represents the optimal state-feedback: A thrust control vector \mathbf{t} , able to steer in (optimal) time t_f^* a system from x_0 to a target state x_f , minimizing the cost function $\mathbf{J}(\mathbf{x}, \mathbf{t}, t_f) = \int_{t_0}^{t_f} \ell(\mathbf{x}, \mathbf{t}) dt$, where $\ell(\mathbf{x}, \mathbf{t})$ denotes the instantaneous cost. The resulting system serves as a suitable on-board substitute for more classical guidance and control in a variety of space applications, such as interplanetary transfer and asteroid landing. In the context of our work, we specifically aim to generate and track an interplanetary trajectory on-board, where the output control vector \mathbf{t} reflects the thrust action of a low-thrust spacecraft (Fig. 1, left).

For the presented experiments, we use a G&CNET network with three hidden layers, each with 128 neurons (Fig. 1, right). Our previous work on this topic evaluated the original architecture of the G&CNETs with softplus and tanh activation functions [36]. However, recent advances have revealed that the performance of G&CNETs is improved by employing the SIREN architecture [31]. Thus, we study here G&CNETs with periodic activation functions, in our case a sine function $\sin(\omega_0 \cdot x)$ with $\omega_0 = 1$ and input x . The weights of each layer are randomly initialized from a uniform distribution: for the first layer, we use $\mathcal{U}(-\frac{1}{n_i}, \frac{1}{n_i})$, while for all other layers we $\mathcal{U}(-\frac{\sqrt{6/n_i}}{\omega_0}, \frac{\sqrt{6/n_i}}{\omega_0})$ is used, with n_i being the number of inputs of layer i . Training is performed for 300 epochs. Alternative methods and deep neural networks exist to perform similar and related tasks, such as Cheng et al.’s work on real-time optimal control for irregular asteroid landing [6], or Li et al.’s work on guidance for powered landing [37]. In this study, G&CNETs are surveyed as a representative example of a guidance and control deep neural network to validate the performance of memristor-based neural network accelerators.




Fig. 1. (Left) A spacecraft transfers between two orbits using a G&CNET. (Right) The architecture of the G&CNET studied. The network is composed of 3 hidden layers with 128 neurons each, using sine activation functions. It receives the spacecraft state (coordinates and velocity) as input and returns the thrust control.

2.1.2. On-board geodesy of small bodies: geodesyNets

Geodesy of small bodies such as asteroids and comets is of great interest in a wide range of fields and industries, from purely scientific interest in the origins of our solar system [38] to the commercial prospect of asteroid mining [39]. Recently, a deep learning-based method for performing geodesy on-board a spacecraft using orbital acceleration measurements has been introduced: geodesyNets [32]. GeodesyNets are inspired by Generative Query Networks [40] and Neural Radiance Fields [41], which introduce novel neural network architectures and training pipelines for three-dimensional scene reconstruction from two-dimensional images. In other words, they solve the inverse problem of going from images of a scene (2D projections from different perspectives) to the scene these images are representing (the 3D object images were taken of). The method introduced in [32] transfers this approach to the realm of geodesy, basically exchanging 2D images by gravitational forces and the 3D scene by the density field of a small body. Thus, geodesyNets approximately solve the inverse problem of finding a gravitational body's mass density field given the gravitational force it exerts on other bodies at different locations, e.g., on a spacecraft along an orbit around it (Fig. 2, left).

GeodesyNets are fully-connected neural networks with periodic activation functions (i.e. SIRENs, with $\omega_0 = 30$). The final layer uses the absolute value as an activation function to ensure densities are positive. The inputs to the network are Cartesian coordinates x, y, z , denoting a point within the three-dimensional space V enclosing the body. In our case, the network consists of 4 fully-connected hidden layers with 300 neurons each. The network's output is the mass density of the small body, $\rho(x, y, z)$, at location (x, y, z) . By using the neural network as a supplement for the actual mass density field of the small body, we calculate the expected gravitational force it enacts on other objects at a given location (e.g. a spacecraft along an orbit around it). The difference between the inferred gravitational acceleration and the ground truth value is evaluated in a custom loss function, which is minimized to improve the neural network's representation of the density field. Both for reference models and the simulated memristor-based geodesyNets, we use 10,000 training epochs and a quadrature of 30,000. Though geodesyNets are generally applicable to any small, irregularly shaped body, training for this work is performed in all cases on the asteroid Eros.

Different from G&CNETs, which are pre-trained before deployment, geodesyNets continuously learn to fully characterize the body they are orbiting [42] — requiring many write operations while being deployed. Thus, they would benefit significantly from a memristive implementation, which provides highly power-efficient write operations.

Alternative methods to perform the same or similar tasks on small bodies include spherical harmonics, using mascon models and polyhedral gravity models or a combination thereof [43–45], which

geodesyNets outperforms in terms of accuracy, needed prior assumptions and on-board applicability. Other neural network-based approaches such as Zhao et al.'s work on using instead Hopfield neural networks to solve a simplified version of the same estimation problem also exist [46]. As such we survey geodesyNets, as a representative method for this particular application.

2.2. Memristor-based neural network accelerators

Originally introduced in 1971 [47] — but only first manufactured in 2011 [48] — memristors are an emerging technology that may enable neural networks to meet essential requirements in terms of accuracy, energy efficiency, radiation tolerance, and latency. The memristor (a portmanteau of “memory resistor”) is a two-terminal nonlinear resistive element, whose resistance R (or conductance G , the inverse of the resistance) depends on the history of voltage across it. The method by which memristors switch their resistance (the so-called “switching mechanism”) differs per technology used to implement the memristor, but is generally performed by one or more electrical pulses forming and rupturing some form of conductive filament, or changing a magnetic orientation [49].

Memristor crossbars (shown in Fig. 3) can store multiple bits of information in each device, and are able to perform computations with very low energy and latency [50]. If an input voltage is applied at the crossbar array's rows, we are able to get the result of a multiplication with the device's resistance as a consequence of Kirchhoff's law, thus performing matrix–vector–multiplication (MVM) in a crossbar array in one computational step. Each element of the crossbar (representing a weight, or element of the vector) is represented by two devices (2R, “Two Resistors”) in a differential configuration (shown schematically in Fig. 4 (left)). In this configuration, each neural network weight w is encoded using a differential pair of memristive devices (commonly referred to as a 2R or “two-resistor” scheme), with associated conductances G^+ and G^- . The effective weight is defined as the difference between these conductances: $w \propto G^+ - G^-$. In this configuration, one device encodes the positive component, and the other the negative component. Typically, either G^+ or G^- is programmed to a conductance value within the supported range $[G_{\min}, G_{\max}]$, while the other is held at G_{\min} (ideally 0). These memristor crossbar arrays can be built using RRAM [51], PCM [52] and other device technologies. As all calculations are performed in the analog domain, DACs and ADCs are required when interfacing with digital inputs our outputs. Although memristors may be implemented in a radiation resistant manner, it is important to note that the peripheral circuits, which include control blocks and the aforementioned DACs and ADCs, may need radiation hardening when a memristive accelerator is to be used in harsh environments [10,28]. Through the experiments performed we wish to

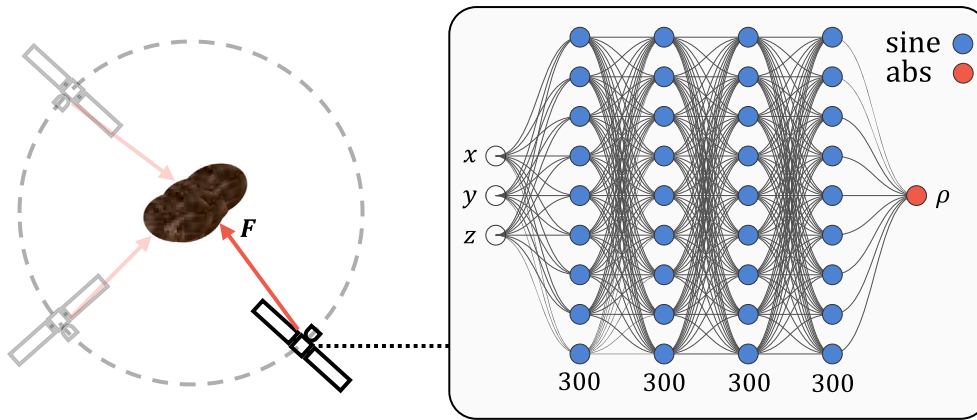


Fig. 2. (Left) A spacecraft orbiting a small, irregularly shaped body. Using the gravitational force (red), an on-board neural network (geodesyNet) learns to reconstruct the density field of the small body. (Right) The architecture of the geodesyNet studied. Similar to the G&CNET, the model is composed of several hidden layers with sine activations. It receives x, y, z coordinates as inputs and returns the density of the asteroid at this coordinate. To calculate forces, an integral of the neural network has to be performed.




Fig. 3. Memristors arranged in crossbar arrays, capable of performing matrix–vector-multiplication (MVM). Whereby a given output current I_i , is the result of a multiplication of input voltage V_j and weight value — represented by conductance $G_{i,j}$.

explore various configurations and parameters for a memristor-based neural network accelerator. For this we use a simulation setup which allows for modification of values such as the device properties and device technologies as necessary.

2.3. Simulation setup

The memristor-based neural network accelerator simulation setup used to perform the experiments is implemented using the *IBM Analog Hardware Acceleration Kit* (IBM AI HW Kit for short) [53], an open-source Python toolkit that is used to explore in-memory computing using analog devices in the context of neural networks in a realistic manner. It is fully integrated within the PyTorch machine learning library for the simulation of training and inference of deep neural networks built with PyTorch on analog crossbar arrays. It also features a variety of analog neural network modules such as fully connected layers and convolutional layers. Both analog training and hardware-aware training are supported. Furthermore, it is highly customizable in every aspect, ranging from mapping algorithms to supported analog layer types. As part of our contribution to the toolkit, bit-slicing [54] (as linear bit sliced layers) has been implemented with the ability to vary the number of slices arbitrarily. Bit-slicing is the use of multiple devices to represent the value of one weight in the neural network, with each device representing a “slice” of a complete weight. Fig. 4 (right)

shows how bit-slicing would be implemented in practice on a circuit level.

In order to achieve realistic simulation of inference accuracy degradation, the IBM AI HW Kit also provides a specific inference configuration that adds carefully calibrated conductance-dependent programming noise, weight read noise and conductance drift based on a 1M PCM device array. The PCM devices exhibit approximately 2% read noise on its programmed conductance on average. The same has been done for RRAM devices, based on data from Wan et al.’s work on RRAM-based in-memory computing chips [25]. The RRAM exhibits 1% read noise on its programmed conductance on average. Conductance drift is a non-ideal behavior of memristors, meaning they may lose their programmed value over time even without any applied voltage. “Refreshing” the devices to restore the weights is possible, but consumes energy [55].

The architecture of the hardware computing the forward pass of the network is as follows: the network consists of multiple layers, each consisting of one or multiple tiles, which contain the memristor crossbar arrays and associated periphery. The weights are mapped to the devices in each tile in a differential configuration, with one device representing a positive synaptic weight and another a negative component, with one of either being set to 0. The peripheral circuitry consists of digital-to-analog converters (DACs, set to a 7-bit resolution) on the input and analog-to-digital converters (ADCs, set to a 9-bit resolution) on the output. It also simulates IR-drop (voltage dropping as

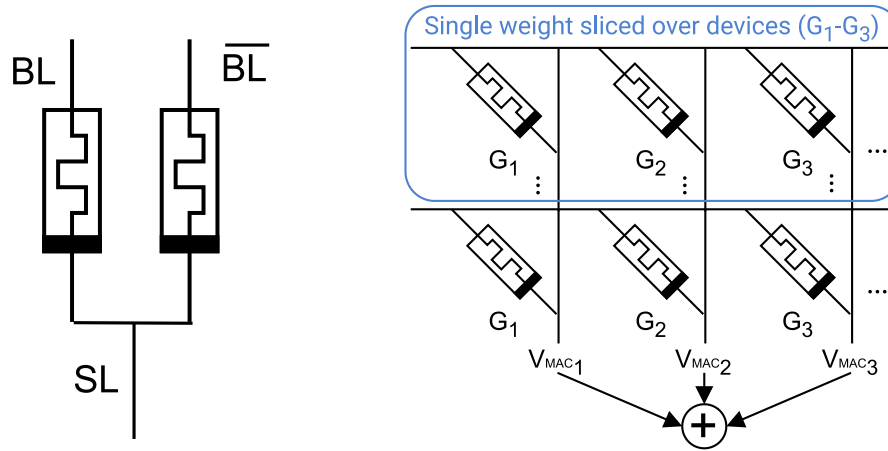


Fig. 4. (Left) shows a schematic representation of 2R scheme commonly used for differential weight encoding, with BL (bit-line) and BLb often connecting to access transistors in a 2T2R (2 Transistor 2 Resistor) configuration. (Right) gives a diagram of bit slicing as being a technique wherein multiple devices are used to represent one weight with higher precision.

current passes through a resistor) and noise from the peripheral circuits (such as from amplifiers used in the ADCs). The result of the computed MVM is then passed to the activation function, which then passes its output to the next layer.

2.4. Device fault and non-ideality mitigation

Within this simulation setup, we propose the concept of “temporal averaging” to deal with hardware-related stochasticity. The core idea is as follows: given a single layer in the neural network, its output is calculated multiple times and averaged before being sent to the next layer. This way, the output of the layer resembles more closely the intended, noise-less and variability free behavior. Moreover, amplification of noise due to propagation through the network is directly inhibited. For instance, in case of noise of a Gaussian distribution on the weights, the standard deviation of the mean output will decrease by a factor $\frac{1}{\sqrt{N}}$, where N is the number of repeats we are performing. Assuming that the means of the weight distributions are the values we intended to program, averaging produces layerwise outputs that are closer to the intended ones without noise.

A large benefit of this mitigation technique is the simplicity of potential implementations in hardware. Combining a shift-register and an adder to capture and combine the result of multiple runs of the layer (or network), which can then be averaged by simply shifting the captured result right by the amount of results captured (thus limiting the system to repeats of powers of 2). It further utilizes the fact that forward passes are cheap both in terms of energy and latency on memristive devices — although excessive averaging will add noticeable costs (both in terms of energy expended performing multiple inferences and the latency this adds) that might detract from the benefits of memristors. Furthermore, benefits will diminish as the number of repeats increases.

As an alternative to layerwise averaging, we explore commonly used bit slicing, with the final weight being the average of the slices. Different from temporal averaging, where averaging was performed for the same device, here we spatially average over different devices, albeit with related read/write behavior. Thus, we expect a similar reduction in noise levels using this technique.

Lastly, to deal with hardware-related non-idealities (e.g. device faults), hardware-aware (HWA) training is used to train the memristor-based neural networks. In HWA training, weight updates are calculated in digitally (i.e., not on-chip) based on inference run on hardware [56–59]. In our work, an ideal backward pass is assumed, with the non-idealities only affecting the forward pass. Weight updates derived

from this process are then applied to the simulated memristive neural network, which is repeated for all training epochs. In a produced chip, the hardware-aware training could be substituted by a digital section of the chip, rather than on a separate machine. The code and data to reproduce the results in this paper is available on Github [60].

3. Results

We perform a thorough set of experiments using the simulation setup outlined in Section 2.2 to assess the current suitability of memristive neural network accelerators implemented with the device technologies introduced in Section 1. Moreover, we demonstrate the feasibility of the proposed noise mitigation techniques, and study the effect of device imperfections (faulty devices and conductance drift) on task performance. In particular, we study:

1. The impact of linear bit-slicing on performance, from 1 slice to 16 slices.
2. The impact of temporal layerwise averaging, with averages (repeats) ranging from no repeats (1) to 64, in intervals matching powers of 2.
3. The performance under various levels of device degradation (as a ratio of devices stuck at any value between G_{min} and G_{max} , randomly distributed), in a realistic range from 1% to 10% [61].
4. Performance recovery through re-training the neural networks post-degradation, i.e., using device-aware training.
5. The effect of conductance drift over time on performance (from $t = 0$ s, to $t = 48 \cdot 3600$ s).

For both the navigation and geodesy task, the performance is defined as the average loss of the test set. All experiments are performed for both the PCM and RRAM devices. Aside from the experiments concerning the number of devices per weight (slices), all experiments are performed with 1 slice per weight for the G&CNETs, and 4 slices per weight for geodesyNets. These values were chosen on the basis of the effectiveness of the slices during the experiments concerning the linear bit-slicing.

3.1. Memristor-based G&CNETs

In Fig. 5 (left), we see that the number of slices has only a minor effect on the accuracy of the network, which is why we default to only 1 slice per weight for the G&CNET. The same is true for temporal averaging (Fig. 5, right). However, by switching from softplus activation functions to sine functions, we see an improvement from 0.07 to almost

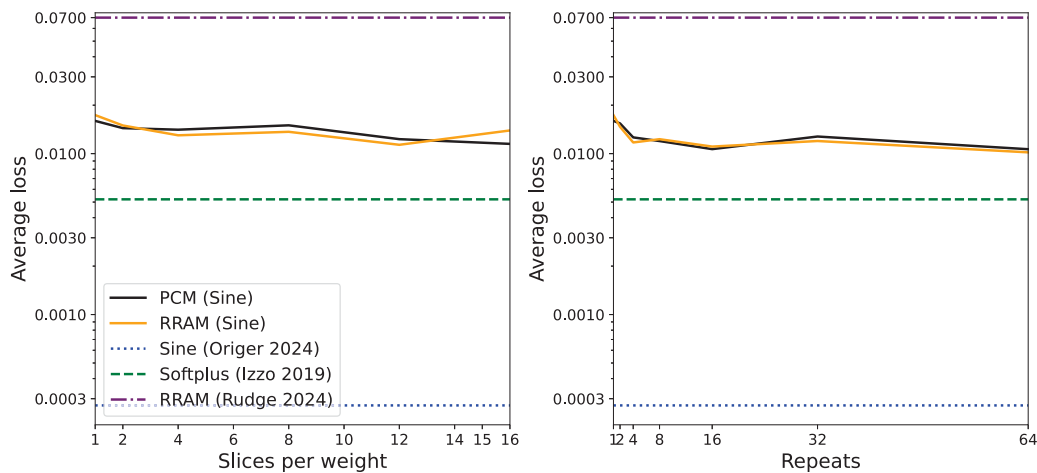


Fig. 5. Average loss over the test dataset of model predictions plotted against (left) slices (with repeats set to 1) and (right) number of temporal averages (with slices set to 1). For comparison, we also show the best result obtained with simulated RRAM for neural networks with softplus activation function (dash-dotted), as well as the digital baseline for softplus (dashed) and sine activation function (dotted).




Fig. 6. (Left) Loss for increasing device fault ratios. The faults are defined as memristors which are “stuck-at” a given value between the maximum and minimum conductance. The two sets of lines depict the PCM and RRAM neural networks as affected by faulty devices with (line) and without (dashed) retraining. (Right) The effect of conductance drift on the performance of the network.

0.01 — although the digital baseline also improves by almost one order of magnitude. Nevertheless, the memristor-based G&CNET gets close to the performance level of the digital baseline using softplus activation functions.

Fig. 6 shows a rapid decline in network performance as device failures increase, with fault ratios above 2%–3% already presenting significant challenges for re-training. In the current model, even a small number of faults severely impacts performance, with the non-retrained network exhibiting substantial degradation. Compared to the softplus-based architecture from [30], this suggests that the G&CNET architecture with periodic activation functions is highly vulnerable to disturbances. For example, when comparing 10% device faults without re-training in Fig. 6 and with re-training, we see that the network recovers significantly from a loss of around 1.1435 to 0.87 — but in the softplus case this was a recovery from around 0.34 to 0.086.

Fig. 6 (right) shows the effect of conductance drift on the loss of the network, which is similar to the effect observed for device faults. Before $t = 0.5$ s the network is mostly unaffected, after which performance deteriorates quickly. In previous network architectures, the softplus

version of G&CNETs [30] was able to hold performance much longer, and as will be shown in the next section, our results for GeodesyNet are more in line with those results (see Section 3.2).

We observed that the G&CNETs with sinusoidal activations exhibit a significantly stronger degradation under conductance drift and device failure compared to their softplus-based counterparts. To investigate this behavior, we analyzed the sensitivity of the networks with respect to perturbations in the weights. In an effort to create a more detailed analysis, we also trained the network with a variety of ω_0 values, ranging from $\omega_0 = 30$ to $\omega_0 = 0.01$ and determined the Lipschitz constant of each of these networks, and also that of the softplus-based network and GeodesyNet. The results of these tests are given in Fig. 7. The Lipschitz constant was estimated numerically by two methods: (1) computing the average gradient norms using automatic differentiation over the dataset, and (2) evaluating $|f(x) - f(x')|/\|x - x'\|$ for a large number (> 7500) of random input pairs. We observe that the Lipschitz constant seems correlated to the sensitivity of a given network to drift. At $\omega_0 = 30$, this value is 850, at $\omega_0 = 1.0$ it is 198, and finally at $\omega_0 = 0.1$ and $\omega_0 = 0.01$ it results in 20 and 1 respectively (all have

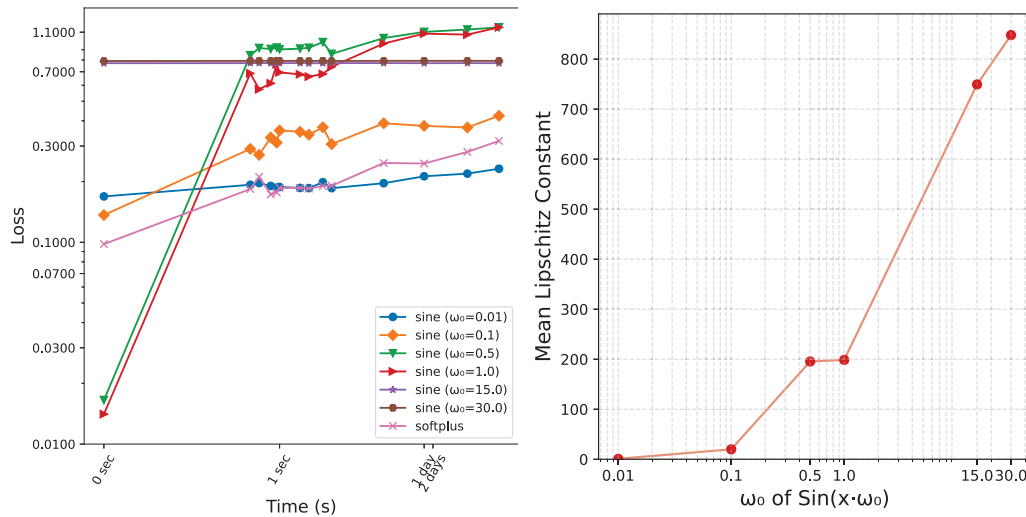


Fig. 7. (Left) A plot showing the performance under drift at various values of ω_0 . Showing that at a high enough ω_0 , the system does not learn at all, and at a small enough ω_0 though robust to drift, the initial performance is very close to the original softplus implementation. (Right) A plot of the computed Lipschitz Constant for a sinusoidal G&CNET under varying ω_0 values.




Fig. 8. The network prediction transformed into spherical coordinates θ and ϕ (lines) compared to the optimal ground truth (points). Both the analog (RRAM, left) and digital (right) models are shown.

been rounded). With $\omega_0 = 0.01$ showing poorer performance at t_0 than $\omega_0 = 1.0$, but much stronger robustness against drift. This phenomenon is discussed further in Section 4.

Finally, the spacecraft control obtained from the current best analog network (1 slice, RRAM, no faults nor drift, 300 epochs at 64 repeats for temporal averaging) for a single interplanetary transfer is shown in Fig. 8, together with the control produced by the digital baseline. The loss of this network is approximately 0.01018. More specifically, we show the target spherical coordinates (θ and ϕ) as predicted by both the digital and analog neural networks — illustrating that the analog model is capable of performing transfers after training, although with noisier control than the digital model.

3.2. Memristor-based geodesyNets

Different to G&CNETs, without bit-slicing or temporal averaging the network completely fails to solve the inverse problem during learning. This is not only seen in the large loss of around 0.36 (compared to the digital baseline of 0.003), but also reflected in the reconstructed density field which fails to capture any structure of the asteroid (Fig. 11, left).

In Fig. 9 we show that increasing the number of slices drastically improves network performance, with diminishing returns for larger number of slices. Similarly, using temporal averaging (repeats) strongly

enhances performance (Fig. 9, right). Of note is that at 4 slices, the loss has reached approximately a value of 0.0655, whereas at 4 repeats (and 4 slices) the best loss reaches 0.0253. This highlights that both techniques can be used complementarily to further enhance performance of the final model. In total, we reach a performance of 0.008 for PCM and 0.007 for RRAM (4 slices, 64 repeats), getting to the same order of magnitude in loss as the digital baseline (0.003).

Fig. 10 (left) shows the performance of the neural network under the effects of device degradation. The network is fairly robust up to 4%–6% device degradation, even without compensatory re-training. After this point, re-training is able to maintain performance within the same range of loss, up to 10% (from the original performance of around 0.0279 to 0.0377 after re-training, with the uncompensated loss at 0.321). In Fig. 10 (right) we see how the network deals with conductance drift. At approximately 1 s, performance is fairly stable, with performance degrading rapidly as we approach 1 day. This loss development as the conductance drift worsens is more in line with expectations, as opposed to the severe impact seen in the G&CNET case, where even the slightest drift or device degradation completely degraded performance. We see also the difference in the onset of the largest effects of the drift between PCM and RRAM, with RRAM’s slope being more gentle, but not tapering off until we reach 2 days of drift.

Fig. 11 compares two neural network models trained to learn the shape and density distribution of the asteroid Eros. The left model,

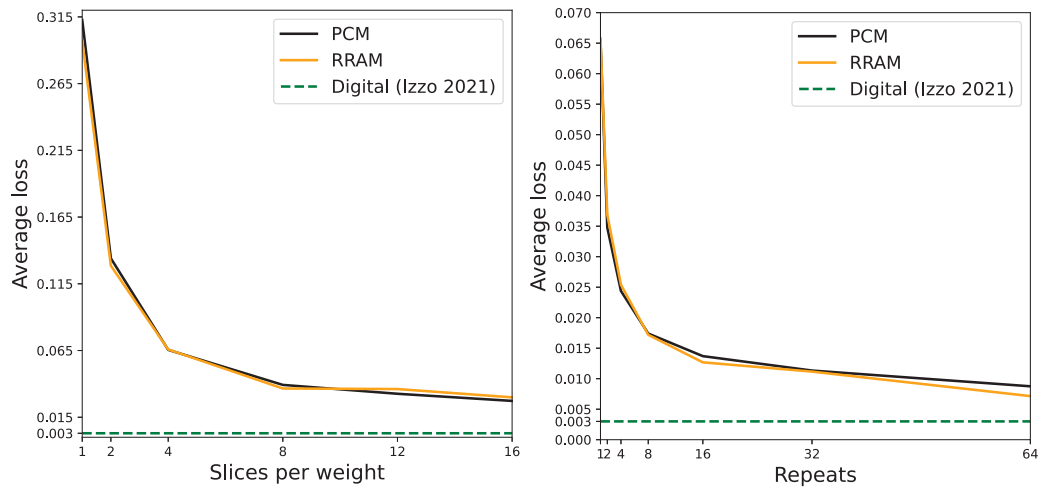


Fig. 9. (Left) Loss of model predictions plotted against slices. Digital baseline is shown as a dashed line. (Right) The effect of using temporal averaging on the network. Each experiment was run with 4 slices per weight. Note that the y-axes have different ranges.

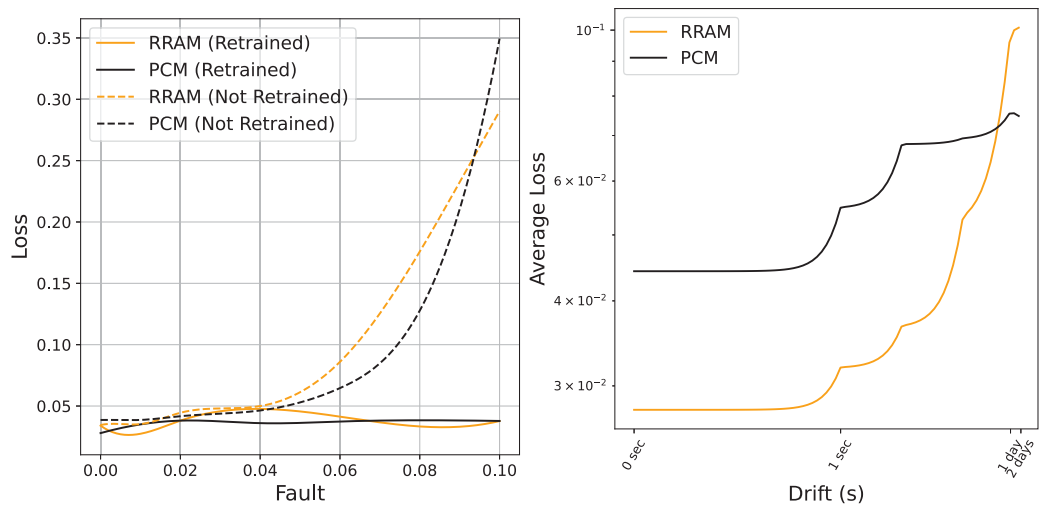


Fig. 10. (Left) Average loss of the model while varying the number of faulty devices in the neural network, up to 10%. The effect of post-degradation re-training is also shown. (Right) The system’s performance under drift, up to 2 days.




Fig. 11. (Left) The predicted asteroid density in a model featuring no temporal averaging and only 1 slice, showing its inability to learn at this level of noise and variability. (Right) With a moderate level of noise mitigation (4 slices and 4 repeats for the temporal averaging), the model learns the density field, but struggles with cleaning up the edges.




Fig. 12. The plot of the best-performing model with 4 slices and 64 repeats. It reaches a final loss of approximately 0.008.

which does not include the mitigation techniques outlined in Section 2, does not learn at all, finishing training with a high loss of 0.36. The learned density field is significantly improved (with a loss of 0.024) when using 4 slices and 4 repeats, as shown in Fig. 11 (right). Finally, in Fig. 12 we show the best performing model, with 64 repeats for temporal averaging and 4 slices per weight. It converges to a final loss of approximately 0.008 and clearly shows the shape and density of Eros, with errant densities at the edges cleaned up. This demonstrates that with temporal averaging and slices, the memristor-based neural network is capable of accurately learning Eros' structure from scratch in under 10,000 epochs.

4. Discussion

Using simulations based on the IBM AI HW KIT, we have demonstrated that memristor-based neural networks are capable of reaching promising performance levels on state-of-the-art onboard AI tasks. In some respects, our simulations may be overly pessimistic, as advances in technology – such as improved RRAM devices [62] – will produce better results than the devices we used in this study (mostly limited by which devices were modeled in the IBM AI HW KIT during the time of this study). In other respects, they may be too optimistic, as a reality gap always remains — especially when modeling microelectronics on an abstract level. This is particularly true regarding the extent to which non-idealities of devices are captured and modeled. An example limitation of the simulation is that the effect of non-ideal transistor matching is not included.

The simulated memristor-based neural networks reach performance levels close to those of their digital counterparts, although a significant gap (around a factor of 2) remains. This performance gap appears to be primarily driven by variability during inference and training, as evidenced by the effectiveness of bit-slicing and temporal averaging. Not only is temporal averaging effective by itself and when combined with bit-slicing, its benefits also extend to the simplicity of its potential implementation in hardware. Since through averaging, variability is only reduced by a factor of $1/\sqrt{N}$ (for N repeats), a future venue is to look into more efficient mitigation techniques compatible with memristor crossbar arrays, e.g., error correcting codes.

Both drift and device degradation have a significant impact on the performance of G&CNETs, with the performance degrading much

more steeply than in the case of the softplus-based neural network architecture presented in [30], despite the initial significantly lower loss. GeodesyNets, which use a similar SIREN architecture, do not suffer from issues with drift and device degradation, possessing sufficient robustness to these non-idealities.

This warranted further research and experimentation, particularly regarding the choice of ω_0 . The selected $\omega_0 = 1.0$ for G&CNETs deviates from the customary 30.0. As such, we investigated the effect of different ω_0 values on the network's response to drift (as outlined in Section 3). This behavior appears linked to the network's Lipschitz constant: the more reactive a network's output is to its input, the more it also amplifies perturbations due to weight drift or noise. We find that G&CNETs with higher ω_0 values exhibit higher Lipschitz constants and are more sensitive to drift, whereas networks with lower ω_0 are more robust but also less expressive. Interestingly, although GeodesyNets also use sinusoidal activations, they exhibit stronger robustness. This could be because the learned density field changes more gradually in space, and the integral over many points averages out local steepness. Furthermore, the Lipschitz constant of the GeodesyNets are in line with those of the softplus-based G&CNETs, which exhibit better robustness to drift noise. Normalization techniques may help address this, but further research is required to confirm this. These findings highlight a crucial consideration in designing networks for memristive applications: loss alone is not a sufficient performance metric. A more holistic approach is required, as networks may train well but still exhibit poor robustness under important conditions such as drift and device degradation. This underscores not only the need for thorough post-training validation, but training procedures and network architectures that exhibit robust parameters.

5. Conclusion

This study has introduced the concept of exploiting the properties of memristors in the design of a neural network accelerator specifically for space applications. There are many avenues to expand this work. One direction of interest is adding energy estimation within the simulation setup to compare the energy-consumption of memristor-based networks with conventional hardware. As at the current stage, quantitative analysis of the energy consumption of the simulated accelerators is unfeasible without having also set the corresponding peripheral circuits, read and write voltages. Without these numbers and decisions made, we can only give what the general direction of the energy consumption will likely be given the technology. This once again reinforces the need for experimenting with fabricated silicon for accurate quantitative results. Further studies on network architectures and hyperparameters are also needed, particularly to explore the trade-offs between temporal averaging and slicing in terms of latency, energy, area, and accuracy in real hardware implementations. For geodesyNets, the current 10,000-epoch training run, chosen due to the prolonged run-times when simulating memristive devices, is relatively short compared to the original 25,000-epoch training, affecting accuracy. Additionally, further experimentation is needed to assess the impact of ω_0 in periodic activation functions, especially in the case of G&CNETs.

To extend the sandbox of models that can inform the design and development of memristive accelerators for space, additional space applications should be explored. With geodesyNets successfully implemented, additional applications using neural implicit representations are suitable candidates for memristors, such as Neural Radiance Methods for Lunar Terrain Modeling [63]. Adaptive compression algorithms [64] are another candidate, optimizing SAR imagery compression and enabling in-field continuous learning. In such cases, memristors' ability to perform low-power inference and in-field training can prove particularly suitable. Another relevant point of interest is also the impact of on-chip training (as opposed to hardware-aware training) on geodesyNets and other networks that learn in the field, as on-chip training may affect the network's ability to learn properly.

Finally, this study is entirely simulation-based and thus inherently constrained by the data, models, and assumptions used. While great effort has been made to ensure realistic conditions, the simulation-to-reality gap remains. In future work we intend to address this by designing, fabricating, and characterizing hardware to assess the feasibility of memristor-based neural network accelerators for on-board AI.

To conclude, coupled with the latency, power, and energy reductions enabled by memristors [27] as well as their radiation hardness [12], our results provide a promising outlook for enabling efficient and reliable on-board AI using memristors.

CRedit authorship contribution statement

Zacharia A. Rudge: Writing – review & editing, Writing – original draft, Visualization, Validation, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization. **Dominik Dold:** Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Methodology, Investigation, Formal analysis, Conceptualization. **Moritz Fieback:** Writing – review & editing, Supervision, Conceptualization. **Dario Izzo:** Supervision, Formal analysis, Conceptualization. **Said Hamdioui:** Supervision, Funding acquisition, Conceptualization.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Zacharia A. Rudge reports financial support was provided by European Space Agency. Dominik Dold reports financial support was provided by European Commission. Co-author Dario Izzo is an editor of the special issue spAIce 2024. Dominik Dold is co-organizer of the scientific committee of spAIce 2024. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

Z.R. acknowledges funding from the European Space Agency's Open Space Innovation Platform (contract number 4000140774). D.D. was supported by the Horizon Europe's Marie Skłodowska-Curie Actions (MSCA) Project 101103062 (BASE). We also acknowledge the use of the open source software IBM Analog Hardware Acceleration Kit, and the open source repositories for G&CNETs and geodesyNets.

References

- [1] G. Furano, G. Meoni, A. Dunne, D. Moloney, V. Ferlet-Cavrois, A. Tavoularis, J. Byrne, L. Buckley, M. Psarakis, K.-O. Voss, L. Fanucci, Towards the use of artificial intelligence on the edge in space systems: Challenges and opportunities, *IEEE Aerosp. Electron. Syst. Mag.* 35 (12) (2020) 44–56, <http://dx.doi.org/10.1109/MAES.2020.3008468>, URL <https://ieeexplore.ieee.org/document/9288809/?arnumber=9288809>.
- [2] V. Kothari, E. Liberis, N.D. Lane, The final frontier: Deep learning in space, in: *Proceedings of the 21st International Workshop on Mobile Computing Systems and Applications, HotMobile '20*, Association for Computing Machinery, New York, NY, USA, 2020, pp. 45–49, <http://dx.doi.org/10.1145/3376897.3377864>.
- [3] D. Izzo, M. Märten, B. Pan, A survey on artificial intelligence trends in spacecraft guidance dynamics and control, *Astrodynamic* 3 (4) (2019) 287–299, <http://dx.doi.org/10.1007/s42064-018-0053-6>.
- [4] D. Izzo, G. Meoni, P. Gómez, D. Dold, A. Zochbauer, *Selected trends in artificial intelligence for space applications*, in: *Artificial Intelligence for Space: AI4SPACE*, CRC Press, 2023, pp. 1–32.
- [5] D. Izzo, A. Hadjiivanov, D. Dold, G. Meoni, E. Blazquez, *Neuromorphic computing and sensing in space*, in: *Artificial Intelligence for Space: AI4SPACE*, CRC Press, 2022, pp. 107–159.
- [6] L. Cheng, Z. Wang, Y. Song, F. Jiang, Real-time optimal control for irregular asteroid landings using deep neural networks, *Acta Astronaut.* 170 (2020) 66–79, <http://dx.doi.org/10.1016/j.actaastro.2019.11.039>, URL <https://www.sciencedirect.com/science/article/pii/S0094576520300151>.

- [7] Y. Ma, B. Pan, Q. Tan, Neural convex optimization for real-time trajectory generation of asteroid landings, *Acta Astronaut.* 229 (2025) 606–615, <http://dx.doi.org/10.1016/j.actaastro.2025.01.009>, URL <https://www.sciencedirect.com/science/article/pii/S0094576525000116>.
- [8] Y. Ni, B. Pan, P.G. Pérez, Accelerating deep neural network training for autonomous landing guidance via homotopy, *Acta Astronaut.* 212 (2023) 654–664, <http://dx.doi.org/10.1016/j.actaastro.2023.08.029>, URL <https://www.sciencedirect.com/science/article/pii/S0094576523004368>.
- [9] S. Luccioni, Y. Jernite, E. Strubell, Power hungry processing: Watts driving the cost of AI deployment? in: *The 2024 ACM Conference on Fairness, Accountability, and Transparency, ACM, Rio de Janeiro Brazil, 2024*, pp. 85–99, <http://dx.doi.org/10.1145/3630106.3658542>, URL <https://dl.acm.org/doi/10.1145/3630106.3658542>.
- [10] J. Schumann, Radiation Tolerance and Mitigation for Neuromorphic Processors, Tech. rep., NASA Ames Research Center, 2022, NTRS Author Affiliations: KBR (United States) NTRS Document ID: 20220013182 NTRS Research Center: Ames Research Center (ARC), URL <https://ntrs.nasa.gov/citations/20220013182>.
- [11] Nature Nanotechnology Editors, Beyond von Neumann, *Nature Nanotechnology* 15 (7) (2020) <http://dx.doi.org/10.1038/s41565-020-0738-x>, 507–507, URL <https://www.nature.com/articles/s41565-020-0738-x>.
- [12] H. Amrouch, N. Du, A. Gebregiorgis, S. Hamdioui, I. Polian, Towards Reliable in-memory computing: From emerging devices to post-von-Neumann architectures, in: *2021 IFIP/IEEE 29th International Conference on Very Large Scale Integration, VLSI-SoC, IEEE, Singapore, 2021*, pp. 1–6, <http://dx.doi.org/10.1109/VLSI-SoC53125.2021.9606966>, URL <https://ieeexplore.ieee.org/document/9606966/>.
- [13] E. Lagunas, F.G. Ortiz, G. Eappen, S. Daoud, W.A. Martins, J. Querol, S. Chatzinotas, N. Skatchkovsky, B. Rajendran, O. Simeone, Performance evaluation of neuromorphic hardware for onboard satellite communication applications, *IEEE Wirel. Commun.* 31 (6) (2024) 78–84.
- [14] P. Lunghi, S. Silvestrini, G. Meoni, D. Dold, A. Hadjiivanov, D. Izzo, et al., Investigation of low-energy spiking neural networks based on temporal coding for scene classification, in: *75th International Astronautical Congress, IAC 2024*, 2024, pp. 1–13.
- [15] E. Arnold, P. Spilger, J.V. Straub, E. Müller, D. Dold, G. Meoni, J. Schemmel, Scalable network emulation on analog neuromorphic hardware, *Front. Neurosci.* 18 (2025) 1523331.
- [16] A.S. Kucik, G. Meoni, Investigating spiking neural networks for energy-efficient on-board AI applications. a case study in land cover and land use classification, in: *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, CVPRW, IEEE, Nashville, TN, USA, 2021*, pp. 2020–2030, <http://dx.doi.org/10.1109/CVPRW53098.2021.00230>, URL <https://ieeexplore.ieee.org/document/9522999/>.
- [17] L. Diana, P. Dini, Review on hardware devices and software techniques enabling neural network inference onboard satellites, *Remote. Sens.* 16 (21) (2024) 3957, <http://dx.doi.org/10.3390/rs16213957>, URL <https://www.mdpi.com/2072-4292/16/21/3957>.
- [18] C. Frenkel, D. Bol, G. Indiveri, Bottom-up and top-down approaches for the design of neuromorphic processing systems: Tradeoffs and synergies between natural and artificial intelligence, *Proc. IEEE* 111 (6) (2023) 623–652.
- [19] C.D. Schuman, T.E. Potok, R.M. Patton, J.D. Birdwell, M.E. Dean, G.S. Rose, J.S. Plank, A survey of neuromorphic computing and neural networks in hardware, 2017, <http://dx.doi.org/10.48550/ARXIV.1705.06963>, URL <https://arxiv.org/abs/1705.06963>.
- [20] M. Davies, N. Srinivasa, T.-H. Lin, G. Chinya, Y. Cao, S.H. Choday, G. Dimou, P. Joshi, N. Imam, S. Jain, Y. Liao, C.-K. Lin, A. Lines, R. Liu, D. Mathaikutty, S. McCoy, A. Paul, J. Tse, G. Venkataraman, Y.-H. Weng, A. Wild, Y. Yang, H. Wang, Loihi: A neuromorphic manycore processor with on-chip learning, *IEEE Micro* 38 (1) (2018) 82–99, <http://dx.doi.org/10.1109/MM.2018.112130359>, URL <https://ieeexplore.ieee.org/document/8259423/>.
- [21] N. Montealegre, D. Merodio, A. Fernández, P. Armbruster, In-flight reconfigurable FPGA-based space systems, in: *2015 NASA/ESA Conference on Adaptive Hardware and Systems, AHS, 2015*, pp. 1–8, <http://dx.doi.org/10.1109/AHS.2015.7231177>, URL https://ieeexplore.ieee.org/abstract/document/7231177?casa_token=nPpHavFgYWMAAAA:8ED_dQ52ItrhoqT92juEf5Ud4xxTUfFaBLTMcbpHVUV-x2keSa7g6RrJhCQ11a0vFT9DINHzhVA.
- [22] A. Singh, S. Diware, A. Gebregiorgis, R. Bishnoi, F. Catthoor, R.V. Joshi, S. Hamdioui, Low-power memristor-based computing for edge-AI applications, in: *2021 IEEE International Symposium on Circuits and Systems, ISCAS, 2021*, pp. 1–5, <http://dx.doi.org/10.1109/ISCAS51556.2021.9401226>, ISSN: 2158-1525, URL https://ieeexplore.ieee.org/abstract/document/9401226?casa_token=Plt6juWj9zAAAAA:CSixmVTSpm_FEDTxySnSng7ikz6spsnTThyOBNLucY5b5zZrdy3-bZWKFKlciVYs-p-zD30.
- [23] S. Diware, A. Singh, A. Gebregiorgis, R.V. Joshi, S. Hamdioui, R. Bishnoi, Accurate and energy-efficient bit-slicing for RRAM-Based neural networks, *IEEE Trans. Emerg. Top. Comput. Intell.* 7 (1) (2023) 164–177, <http://dx.doi.org/10.1109/TETCI.2022.3191397>, URL <https://ieeexplore.ieee.org/document/9840507/>.

- [24] R. Khaddam-Aljameh, M. Stanisavljevic, J. Fornet Mas, G. Karunarathne, M. Braendli, F. Liu, A. Singh, S.M. Muller, U. Egger, A. Petropoulos, T. Antonakopoulos, K. Brew, S. Choi, I. Ok, F.L. Lie, N. Saulnier, V. Chan, I. Ahsan, V. Narayanan, S.R. Nandakumar, M. Le Gallo, P.A. Francese, A. Sebastian, E. Eleftheriou, HERMES core – a 14nm CMOS and PCM-based in-memory compute core using an array of 300ps/LSB linearized CCO-based ADCs and local digital processing, in: 2021 Symposium on VLSI Circuits, IEEE, Kyoto, Japan, 2021, pp. 1–2, <http://dx.doi.org/10.23919/VLSICircuits52068.2021.9492362>, URL <https://ieeexplore.ieee.org/document/9492362/>.
- [25] W. Wan, R. Kubendran, C. Schaefer, S.B. Eryilmaz, W. Zhang, D. Wu, S. Deiss, P. Raina, H. Qian, B. Gao, S. Joshi, H. Wu, H.-S.P. Wong, G. Cauwenberghs, A compute-in-memory chip based on resistive random-access memory, *Nature* 608 (7923) (2022) 504–512, <http://dx.doi.org/10.1038/s41586-022-04992-8>, URL <https://www.nature.com/articles/s41586-022-04992-8>.
- [26] P. Deaville, B. Zhang, N. Verma, A 22nm 128-kb MRAM row/Column-parallel in-memory computing macro with memory-resistance boosting and multi-column ADC readout, in: 2022 IEEE Symposium on VLSI Technology and Circuits (VLSI Technology and Circuits), IEEE, Honolulu, HI, USA, 2022, pp. 268–269, <http://dx.doi.org/10.1109/VLSITechnologyandCir46769.2022.9830153>, URL <https://ieeexplore.ieee.org/document/9830153/>.
- [27] D. Bonnet, T. Hirtzlin, A. Majumdar, T. Dalgaty, E. Esmanhotto, V. Meli, N. Castellani, S. Martin, J.-F. Nodin, G. Bourgeois, J.-M. Portal, D. Querlioz, E. Vianello, Bringing uncertainty quantification to the extreme-edge with memristor-based Bayesian neural networks, *Nat. Commun.* 14 (1) (2023) 7530, <http://dx.doi.org/10.1038/s41467-023-43317-9>, URL <https://www.nature.com/articles/s41467-023-43317-9>.
- [28] D.R. Hughart, A.J. Lohn, P.R. Mickel, S.M. Dalton, P.E. Dodd, M.R. Shaneyfelt, A.I. Silva, E. Bielejec, G. Vizekethy, M.T. Marshall, M.L. McLain, M.J. Marinella, A comparison of the radiation response of TaO_x and TiO₂ Memristors, *IEEE Trans. Nucl. Sci.* 60 (6) (2013) 4512–4519, <http://dx.doi.org/10.1109/TNS.2013.2285516>, URL <https://ieeexplore.ieee.org/document/6678257?arnumber=6678257&tag=1>.
- [29] H. Lyu, H. Zhang, B. Mei, Q. Yu, R. Mo, Y. Sun, W. Gao, Research on single event effect test of a RRAM memory and space flight demonstration, *Microelectron. Reliab.* 126 (2021) 114347, <http://dx.doi.org/10.1016/j.microrel.2021.114347>, URL <https://www.sciencedirect.com/science/article/pii/S0026271421003139>.
- [30] Z.A. Rudge, D. Izzo, M. Fieback, A. Gebregiorgis, S. Hamdioui, D. Dold, Guidance and control neural network acceleration using memristors, in: Proceedings of SPAICE2024: the First Joint European Space Agency/IAA Conference on AI in and for Space, 2024, pp. 162–168, URL <https://ui.adsabs.harvard.edu/abs/2024sais.conf..162R/abstract>.
- [31] S. Origer, D. Izzo, Guidance and control networks with periodic activation functions, 2024, <http://dx.doi.org/10.48550/arXiv.2405.18084>, URL [arXiv:2405.18084 \[cs\]](https://arxiv.org/abs/2405.18084).
- [32] D. Izzo, P. Gómez, Geodesy of irregular small bodies via neural density fields, *Commun. Eng.* 1 (1) (2022) 1–12, <http://dx.doi.org/10.1038/s44172-022-00050-3>, URL <https://www.nature.com/articles/s44172-022-00050-3>.
- [33] M. Le Gallo, C. Lammie, J. Büchel, F. Carta, O. Fagbohngbe, C. Mackin, H. Tsai, V. Narayanan, A. Sebastian, K. El Maghraoui, M.J. Rasch, Using the IBM analog in-memory hardware acceleration kit for neural network training and inference, *APL Mach. Learn.* 1 (4) (2023) 041102, <http://dx.doi.org/10.1063/5.0168089>.
- [34] V. Sitzmann, J. Martel, A. Bergman, D. Lindell, G. Wetzstein, Implicit neural representations with periodic activation functions, in: *Advances in Neural Information Processing Systems*, 33, Curran Associates, Inc., 2020, pp. 7462–7473, URL <https://proceedings.neurips.cc/paper/2020/hash/53c04118df112c13a8c34b38343b9c10-Abstract.html>.
- [35] L.S. Pontryagin, *Mathematical Theory of Optimal Processes*, CRC Press, 1987.
- [36] D. Izzo, E. Öztürk, M. Märten, Interplanetary transfers via deep representations of the optimal policy and/or of the value function, in: *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, ACM, Prague Czech Republic, 2019, pp. 1971–1979, <http://dx.doi.org/10.1145/3319619.3326834>, URL <https://dl.acm.org/doi/10.1145/3319619.3326834>.
- [37] W. Li, Y. Song, L. Cheng, S. Gong, Closed-loop deep neural network optimal control algorithm and error analysis for powered landing under uncertainties, *Astrodynamics* 7 (2) (2023) 211–228, <http://dx.doi.org/10.1007/s42064-022-0153-1>.
- [38] K.-H. Glassmeier, H. Boehnhardt, D. Koschny, E. Kührt, I. Richter, The rosetta mission: Flying towards the origin of the solar system, *Space Sci. Rev.* 128 (1) (2007) 1–21, <http://dx.doi.org/10.1007/s11214-006-9140-8>.
- [39] A.M. Hein, R. Matheson, D. Fries, A techno-economic analysis of asteroid mining, *Acta Astronaut.* 168 (2020) 104–115, <http://dx.doi.org/10.1016/j.actaastro.2019.05.009>, URL <https://www.sciencedirect.com/science/article/pii/S0094576518316357>.
- [40] S.M.A. Eslami, D. Jimenez Rezende, F. Besse, F. Viola, A.S. Morcos, M. Garnelo, A. Ruderman, A.A. Rusu, I. Danihelka, K. Gregor, D.P. Reichert, L. Buesing, T. Weber, O. Vinyals, D. Rosenbaum, N. Rabinowitz, H. King, C. Hillier, M. Botvinick, D. Wierstra, K. Kavukcuoglu, D. Hassabis, Neural scene representation and rendering, *Science* 360 (6394) (2018) 1204–1210, <http://dx.doi.org/10.1126/science.aar6170>, URL <https://www.science.org/doi/full/10.1126/science.aar6170>.
- [41] B. Mildenhall, P.P. Srinivasan, M. Tancik, J.T. Barron, R. Ramamoorthi, R. Ng, NeRF: representing scenes as neural radiance fields for view synthesis, *Commun. ACM* 65 (1) (2021) 99–106, <http://dx.doi.org/10.1145/3503250>, URL <https://dl.acm.org/doi/10.1145/3503250>.
- [42] E. Blazquez, D. Izzo, F. Biscani, R. Walker, F. Perez-Lissi, Small celestial body exploration with CubeSat swarms, 2023, arXiv preprint [arXiv:2308.13333](https://arxiv.org/abs/2308.13333).
- [43] R. Werner, D. Scheeres, Exterior gravitation of a polyhedron derived and compared with harmonic and mascon gravitation representations of asteroid 4769 Castalia, *Celest. Mech. Dyn. Astron.* 65 (3) (1997) <http://dx.doi.org/10.1007/BF00053511>, URL <http://link.springer.com/10.1007/BF00053511>.
- [44] P.T. Wittick, R.P. Russell, Mixed-model gravity representations for small celestial bodies using mascons and spherical harmonics, *Celest. Mech. Dyn. Astron.* 131 (7) (2019) 31, <http://dx.doi.org/10.1007/s10569-019-9904-6>, URL <http://link.springer.com/10.1007/s10569-019-9904-6>.
- [45] J. Sebera, A. Bezděk, I. Pešek, T. Henych, Spheroidal models of the exterior gravitational field of Asteroids Bennu and Castalia, *Icarus* 272 (2016) 70–79, <http://dx.doi.org/10.1016/j.icarus.2016.02.038>, URL <https://linkinghub.elsevier.com/retrieve/pii/S0019103516001160>.
- [46] Y. Zhao, H. Yang, S. Li, Y. Zhou, On-board modeling of gravity fields of elongated asteroids using Hopfield neural networks, *Astrodynamics* 7 (1) (2023) 101–114, <http://dx.doi.org/10.1007/s42064-022-0151-3>.
- [47] L. Chua, Memristor—the missing circuit element, *IEEE Trans. Circuit Theory* 18 (5) (1971) 507–519.
- [48] R.S. Williams, How we found the missing memristor, *IEEE Spectr.* 45 (12) (2008) 28–35.
- [49] F. Zahoor, T.Z. Azni Zulkifli, F.A. Khanday, Resistive random access memory (RRAM): an overview of materials, switching mechanism, performance, multi-level cell (mlc) storage, modeling, and applications, *Nanoscale Res. Lett.* 15 (1) (2020) 90, <http://dx.doi.org/10.1186/s11671-020-03299-9>, URL <https://link.springer.com/10.1186/s11671-020-03299-9>.
- [50] A. Ankit, I.E. Hajj, S.R. Chalamalasetti, G. Ndu, M. Foltin, R.S. Williams, P. Faraboschi, W.-m.W. Hwu, J.P. Strachan, K. Roy, D.S. Milojicic, PUMA: A programmable ultra-efficient memristor-based accelerator for machine learning inference, in: *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*, ACM, Providence RI USA, 2019, pp. 715–731, <http://dx.doi.org/10.1145/3297858.3304049>, URL <https://dl.acm.org/doi/10.1145/3297858.3304049>.
- [51] S. Yu, W. Shim, X. Peng, Y. Luo, RRAM for compute-in-Memory: From inference to training, *IEEE Trans. Circuits Syst. I. Regul. Pap.* 68 (7) (2021) 2753–2765, <http://dx.doi.org/10.1109/TCSI.2021.3072200>, URL https://ieeexplore.ieee.org/abstract/document/9406197?casa_token=d9bwflzC-o4AAAAA:dB6CBZ6Pp789xq1A9xghJqEmX7XjYpyMF4I5C-x4cf1rLbM5GgU1LtEuHM1VRcWpccMalhRrA.
- [52] H.-S.P. Wong, S. Raoux, S. Kim, J. Liang, J.P. Reifenberg, B. Rajendran, M. Asheghi, K.E. Goodson, Phase change memory, *Proc. IEEE* 98 (12) (2010) 2201–2227, <http://dx.doi.org/10.1109/JPROC.2010.2070050>, URL https://ieeexplore.ieee.org/abstract/document/5609179?casa_token=ZNFfgH8G5VIAAAAA:kgTjAeJ-KYvyRKsniS9Gd995JySQs-IrmSpOA5kHqC3aAJYopN79pnNCRwCQmRLCKfUqiRdA.
- [53] M.J. Rasch, D. Moreda, T. Gokmen, M. Le Gallo, F. Carta, C. Goldberg, K. El Maghraoui, A. Sebastian, V. Narayanan, A flexible and fast PyTorch toolkit for simulating training and inference on analog crossbar arrays, in: 2021 IEEE 3rd International Conference on Artificial Intelligence Circuits and Systems, AICAS, 2021, pp. 1–4, <http://dx.doi.org/10.1109/AICAS51828.2021.9458494>, URL <https://ieeexplore.ieee.org/abstract/document/9458494>.
- [54] M.L. Gallo, S.R. Nandakumar, L. Ciric, I. Boybat, R. Khaddam-Aljameh, C. Mackin, A. Sebastian, Precision of bit slicing with in-memory computing based on analog phase-change memory crossbars, *Neuromorphic Comput. Eng.* 2 (1) (2022) 014009, <http://dx.doi.org/10.1088/2634-4386/ac4fb7>.
- [55] A. Baroni, C. Zambelli, P. Olivo, E. Pérez, C. Wenger, D. Ielmini, Tackling the low conductance state drift through incremental reset and verify in RRAM arrays, in: 2021 IEEE International Integrated Reliability Workshop, IIRW, 2021, pp. 1–5, <http://dx.doi.org/10.1109/IIRW53245.2021.9635613>, ISSN: 2374-8036, URL <https://ieeexplore.ieee.org/document/9635613>.
- [56] S.K. Esser, P.A. Merolla, J.V. Arthur, A.S. Cassidy, R. Appuswamy, A. Andreopoulos, D.J. Berg, J.L. McKinstry, T. Melano, D.R. Barch, C. di Nolfo, P. Datta, A. Amir, B. Taba, M.D. Flickner, D.S. Modha, Convolutional networks for fast, energy-efficient neuromorphic computing, *Proc. Natl. Acad. Sci.* 113 (41) (2016) 11441–11446, <http://dx.doi.org/10.1073/pnas.1604850113>.
- [57] S. Schmitt, J. Kl'ahn, G. Bellec, A. Gr'ubl, M. Guettler, A. Hartel, S. Hartmann, D. Husmann, K. Husmann, S. Jeltsch, et al., Neuromorphic hardware in the loop: Training a deep spiking network on the brainscale wafer-scale system, in: *Proceedings International Joint Conference on Neural Networks*, IEEE, 2017, pp. 2227–2234, <http://dx.doi.org/10.1109/IJCNN.2017.7966125>.
- [58] A.F. Kungl, S. Schmitt, J. Kl'ahn, P. M'uller, A. Baumbach, D. Dold, A. Kugele, E. M'uller, C. Koke, M. Kleider, et al., Accelerated physical emulation of bayesian inference in spiking neural networks, *Front. Phys.* 13 (2019) 1201, <http://dx.doi.org/10.3389/fphys.2019.01201>.

- [59] M.J. Rasch, C. Mackin, M. Le Gallo, A. Chen, A. Fasoli, F. Odermatt, N. Li, S.R. Nandakumar, P. Narayanan, H. Tsai, G.W. Burr, A. Sebastian, V. Narayanan, Hardware-aware training for large-scale and diverse deep learning inference workloads using in-memory computing-based accelerators, *Nat. Commun.* 14 (1) (2023) 5282, <http://dx.doi.org/10.1038/s41467-023-40770-4>, URL <https://www.nature.com/articles/s41467-023-40770-4>.
- [60] Guidance and control neural network acceleration using memristors github repo, 2024, URL <https://github.com/HeatPhoenix/memristor-gecnet>. (Accessed May 2024).
- [61] T.-H. Kim, S. Kim, J. Park, S. Youn, H. Kim, Memristor crossbar array with enhanced device yield for in-memory vector–matrix multiplication, *ACS Appl. Electron. Mater.* 6 (6) (2024) 4099–4107, <http://dx.doi.org/10.1021/acsaelm.4c00199>, URL <https://pubs.acs.org/doi/10.1021/acsaelm.4c00199>.
- [62] D.F. Falcone, V. Clerico, W. Choi, T. Stecconi, F. Horst, L. Begon-Lours, M. Galetta, A.L. Porta, N. Garg, F. Alibart, B.J. Offrein, V. Bragaglia, All-in-one analog AI accelerator: On-chip training and inference with conductive-metal-oxide/HfOx ReRAM devices, 2025, <http://dx.doi.org/10.48550/arXiv.2502.04524>, URL [arXiv:2502.04524](http://dx.doi.org/10.48550/arXiv.2502.04524)[cs].
- [63] E.V. Kints, A. Hammond, C. Adams, I.G. Lopez-Francos, Neural radiance methods for lunar terrain modeling, in: 2025 IEEE Aerospace Conference, Big Sky, MT, 2025, pp. 1–17, NTRS Author Affiliations: KBR (United States), Ames Research Center NTRS Document ID: 20250001277 NTRS Research Center: Ames Research Center (ARC), URL <https://ntrs.nasa.gov/citations/20250001277>.
- [64] R.M. Asiyabi, A. Anghel, P. Rizzoli, M. Martone, M. Datcu, Complex-valued autoencoder for multi-polarization SLC SAR data compression with side information, in: IGARSS 2023 - 2023 IEEE International Geoscience and Remote Sensing Symposium, 2023, pp. 1787–1790, <http://dx.doi.org/10.1109/IGARSS52108.2023.10282287>, ISSN: 2153-7003, URL <https://ieeexplore.ieee.org/document/10282287>.