Thermal Anomaly Detection using Physics Aware Neuromorphic Networks: Comparison between Raw and L1C Sentinel-2 Data

Anonymous CVPR submission

Paper ID *****

Abstract

Damage caused by bushfires and volcanic eruptions escalates rapidly when detection is delayed, making fast and reliable early warning capabilities essential. Recent Earth Observation (EO) approaches have shown that thermal anomaly detection can be performed directly on decompressed Level–0 (L0) sensor data, avoiding computationally expensive preprocessing chains. However, direct exploitation of raw data remains challenging due to domain shift, sensor drift, radiometric inconsistencies, and the scarcity of labelled training samples. To address these challenges, this work proposes a Physics-Aware Neuromorphic Network (PANN) framework for onboard thermal anomaly detection. The proposed lightweight architecture, inspired by physical neural network principles and neuromorphic computing paradigms, is evaluated using two Sentinel-2 datasets: decompressed L0 with additional metadata (i.e. raw) and Level–1C (L1C). The PANN achieves a Matthews Correlation Coefficient (MCC) of 0.809 *on raw measurements, compared to* 0.875 *when using ground-processed L1C products. The mean processing latency per L0 granule is* 2.44 ± 0.09 s*, which is below the Sentinel-2 acquisition time of* 3.6 s*, demonstrating the feasibility of real–time, onboard processing. Furthermore, the projected execution time for the corresponding neuromorphic hardware instantiation is substantially lower at* 0.1290 ± 0.0002 s*. Memory usage remains within realistic onboard constraints, with requirements of* 0.673 ± 0.007 Gb *for the software PANN and* 0.393 ± 0.004 Gb *for the estimated hardware realisation. Overall, these results indicate that PANN offers a promising pathway toward low-latency and resource-efficient onboard EO processing for thermal event detection.*

1. Introduction

Thermal hotspots or anomalies, such as bushfires and volcanic activity, can cause severe environmental and societal impacts[1, 2], including adversely impacting air and water quality, ecosystems and human welfare [8, 11, 20]. Thus, it is vital that thermal hotspots and anomalies are detected and reported to emergency services before they become out of control [25], especially as fire seasons become longer, with more intense events [9]. While Earth observation (EO) solutions can provide the necessary coverage, for the information to be most effective, latency must be minimised between data acquisition (with limited revisit times), thermal hotspot detection and arrival of information to decision makers and first responders. An approach to this problem is to move the image processing and thermal event detection to be performed onboard the EO satellite.

In missions such as Sentinel-2 [7] and LandSat-8 [28], multispectral EO images are captured with pushbroom based sensors and require significant processing to enhance image quality. Here, we focus on Sentinel-2 data, but the fundamental processing operations are applied to other data processing chains [10], even though implementation of such processing chains and the corresponding product nomenclature vary between missions. The Sentinel-2 satellites acquire EO images using a multispectral imager [7]. The sensor has a total of 12 detectors that, in a single acquisition, each capture a granule for each spectral band in 3.6 s in a pushbroom motion [10]. These bands are not aligned due to the pushbroom nature of the images and along track parallax error. The images are equalised and compressed before being downlinked to earth, and are defined as the decompressed level-0 (L0) data with additional metadata, often referred as the "raw" data [18]. Once on the ground, the raw data undergoes several processing steps, e.g. spatial coregistration of the spectral bands, radiometric correction, invert onboard equalisation, adding the radiometric corrected geometric model, orthorectification and calculation of top of atmosphere reflectances, among others [18]. At this stage the Sentinel-2 data is defined as level-1C (L1C) data. This is the lowest level data that is routinely made publicly available from Copernicus Data Space Ecosystem and ESA projects, with lower level products such as Level-1B released to expert users only from 2025.

Typically, studies that have involved processing EO im-

ages, particularly for detecting natural disasters, have used the L1C data or higher level processed data as this is what is publicly available. With the growing need for onboard processing, recent studies have begun to shift attention to the raw L0 data [6, 19, 26], to avoid the need to perform the computationally intensive processing required to transform the data from raw to L1C. As such, AI models that are to be deployed onboard EO satellites need to be able to process the raw data, while operating under tight resource constraints present on satellites.

While some models address the need to handle raw data, all models require training on the ground before being deployed on satellites to perform inference. The first challenge is that machine learning models need labelled training data and often do not generalise well to new data from a different sensor, resulting in a significant decrease in performance [12], therefore requiring data and labels to be gathered on an ongoing basis for new sensors and missions. Using simulations is an option, but is often time consuming and expensive [5]. Over the duration of a mission, sensors can drift or the domain problem can shift, requiring models to be retrained with new data to maintain performance [17].

This study investigates a new approach to processing raw data using a Physics Aware Neuromorphic Network (PANN), an AI model based on physical neural networks [4, 13, 27, 29]. These are effectively hardware neural networks with a brain-inspired, recurrent architecture comprised of self-assembled nanowires that form nanoelectronic circuit elements known as memristors (nonlinear resistors with memory) [3]. The physical nanowire network substrate embeds memristive weights that continuously update their states in response to varying input signals according to a set of physical laws (a memristor equation of state and Kirchhoff's circuit laws). As such, the PANN actively extracts dynamical features from each new input without training any network weights. This opens up the possibility of energy-efficient continual learning on-orbit. In contrast, this is not possible with neuromorphic approaches that typically implement spike-based point neuron models, which require significant data preprocessing and model training [15]. The PANN model was previously applied to an EO natural disaster change detection task [23], where the PANN processed a sequence of images from the Sentinel-2 satellites using L1C data. This study extends that work by applying the PANN model to thermal event detection using newly available Sentinel-2 raw data and comparing model performance to corresponding L1C data.

The remainder of this paper is organised as follows: Sec. 2 details the datasets used and presents the PANN model, along with the entire pipeline. Section 3 compares the performance of the PANN model on the raw and L1C datasets used, as well as showing the time and memory requirements of the model and estimates for the hardware equivalent. Section 4 presents a summary the main findings of this study and suggestions for future work.

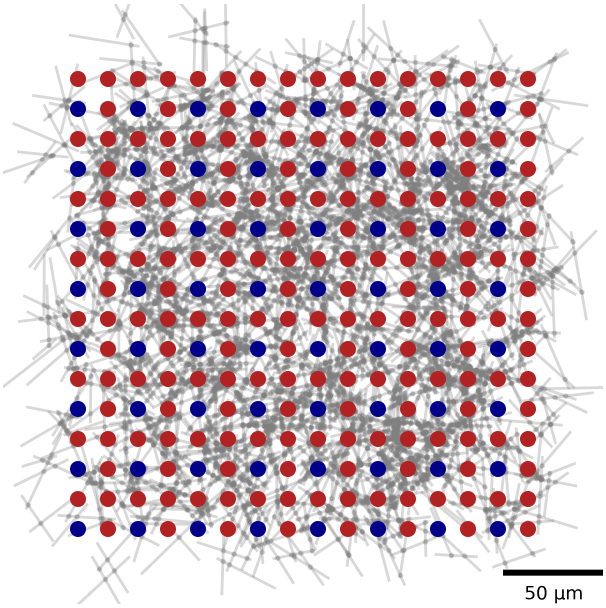

Figure 1. Visualisation of the simulated nanowire network device model used in PANN, showing the input electrodes (large blue circles), readout electrodes (large red circles) and the underlying nanowire network underneath with the nanowires (grey lines) and nanowire–nanowire junctions (small grey dots). Scale bar is 50 μm.

2. Methodology

This study used a modified THRawS dataset [18] containing Sentinel-2 images of thermal hotspots with two levels of processed data, as detailed in Sec. 2.1. Details of the PANN model and parameters used are given in Sec. 2.2, with the full model pipeline given in Sec. 2.3. Section 2.4 describes the method and hardware used to measure the inference time and memory requirement of the model.

2.1. Dataset

The THRawS dataset used in this study contains images of thermal hotspots caused by bushfires and volcanoes, from around the world. In addition, it also contains images of inactive volcanoes to ensure models are learning to detect thermal hotspots, not volcano calderas. For this study, we only used two of the spectral bands, B8A in the near infrared and B12 in the short wave infrared. For each granule location in the dataset there is a raw image and a processed L1C image. Each raw granule had a coarse coregistration [18, 19] applied and adjacent granules were used to fill in the missing data, due to the coarse coregistration of the spectral bands. When no adjacent granule was available, the granule was cropped. This allowed as much of

the dataset as possible to be kept, while ensuring that all inputs into the model had both bands present. To keep the raw and L1C datasets the same, the corresponding L1C granules were also cropped and the SuperGlue neural network [22] was used to align the L1C granules with the raw granules.

Granules, which are 1152×1296 pixels, were tiled into smaller 128×128 pixel patches with no overlap. The patches were automatically labelled using the warm thermal hotspot algorithm by Massimetti et al. [16] on the L1C data. Patches were considered to be an "event", if a hotspot contained at least 9 pixels, to be consistent with [18]. To complete the labelling process, all patches were visually inspected, any patches that contained errors were removed. The final dataset had a total of 21,450 patches of which 442 (2%) were "events" and 21,008 (98%) were "non-events", for both the raw and L1C datasets.

To account for the large class imbalance in the dataset, the Matthew's Correlation Coefficient (MCC) was used to evaluate the performance of the model and is defined as:

$$\text{MCC} = \frac{T_N \cdot T_P - F_N \cdot F_P}{\sqrt{(T_P + F_P)(T_P + F_N)(T_N + F_N)(T_N + F_P)}} \quad (1)$$

where T_P, T_N, F_P and F_N are the total number of True Positives, True Negatives, False Positives and False Negatives, respectively. The precision and recall of the model's performance were also calculated as:

$$\text{precision} = \frac{T_P}{T_P + F_P} \quad , \quad (2)$$

and

$$\text{recall} = \frac{T_P}{T_P + F_N} \quad . \quad (3)$$

2.2. PANN Model Architecture

PANN is based on a physically–motivated model of a nanowire network hardware device comprised of resistive switching memory (memristive) electrical junctions [14, 24]. This is done by first modelling the nanowire network structure to create the topology of the PANN. To create the PANN network used in this study, 1520 nanowires were distributed over a 250×250 μm 2D plane with nanowire centres sampled from a generalised normal distribution with a beta value of 5 and with nanowire orientations sampled on $[0, \pi]$. Nanowire lengths were sampled from a Gaussian distribution with an average of 30 μm and a standard deviation of 6 μm. As memristive nanowire networks are nanoelectronic systems, input and output signals are delivered and read out via electrodes. Thus, a device was modelled by overlaying an evenly spaced 16×16 electrode grid centred over the 2D plane. Each electrode had a 8 μm diameter and a pitch of 8 μm between electrodes. A quarter of the electrodes were selected in an evenly space grid pattern as the input electrodes, thereby forming an input layer. The remaining electrodes were used as readout electrodes forming an output layer. Where the nanowires overlap with other nanowires or electrodes, electrical junctions are formed, resulting in a total of 12,736 memristive junctions for the PANN network used in this study.

Figure 1 depicts the simulated nanowire network used for the PANN model. The nanowires and nanowire–nanowire junctions are displayed as grey lines and small grey dots, respectively. The electrodes are shown as the large circles, with the input electrodes in blue and readout electrodes in red.

Inputs are treated as input voltage signals and Kirchhoff's conservation laws are solved at each time step to calculate the voltage distribution across the network. To do this, the network is converted to its graphical representation, with nanowires and memristors represented as nodes and edges, respectively. As the internal states of the memristive junctions evolve in time, the conductance of each junction can be considered as the weight of an edge in a neural network. Thus, network weights are constrained, rather than trained, allowing the network to continuously extract features from the inputs. The junction conductances evolve in time according to a memristor equation of state [14]:

$$\frac{d\lambda}{dt} = \begin{cases} (|V(t)| - V_{\text{set}})\text{sgn}[V(t)], & |V(t)| > V_{\text{set}} \\ 0, & V_{\text{reset}} < |V(t)| < V_{\text{set}} \\ (|V(t)| - V_{\text{reset}})\text{sgn}[\lambda(t)], & |V(t)| < V_{\text{reset}} \\ 0, & |\lambda| \geq \lambda_{\text{max}} \end{cases} \quad (4)$$

with all junctions set to an initial state $\lambda(t = 0) = 0$. In this study, the following parameters were used: $V_{\text{reset}} = 5 \times 10^{-3}$ V, $V_{\text{set}} = 1 \times 10^{-2}$ V and $\lambda_{\text{max}} = 1.5 \times 10^{-2}$ V s.

2.3. Model Pipeline

Figure 2 shows the full model pipeline. First, each band in the data is normalised according to Eq. (5),

$$x = 1.2 \left[\frac{x}{\max(x)} \right] - 0.4 \quad , \quad (5)$$

and scaled to the interval [-0.4,0.8].

The maximum value used for the different bands was based on values used in other works that used Sentinel-2 data [21, 23]. In the raw data a max value of 3000 was used for both bands and values of 4 and 2 for bands B8A and B12 in the L1C data, respectively. Next, granules were tiled into smaller 128×128 pixel patches with no overlap and a max pooling layer was applied with a pool size of 16×16 and a stride of 16. The two bands were then split and fed into separate PANNs, each of which were interfaced via a 16×16 electrode grid; with the input electrodes functioning as the input layer and the rest of the electrodes used as the readout layer. As such, a feature vector is readout from the PANNs for each tile fed into the PANNs. A skip connection

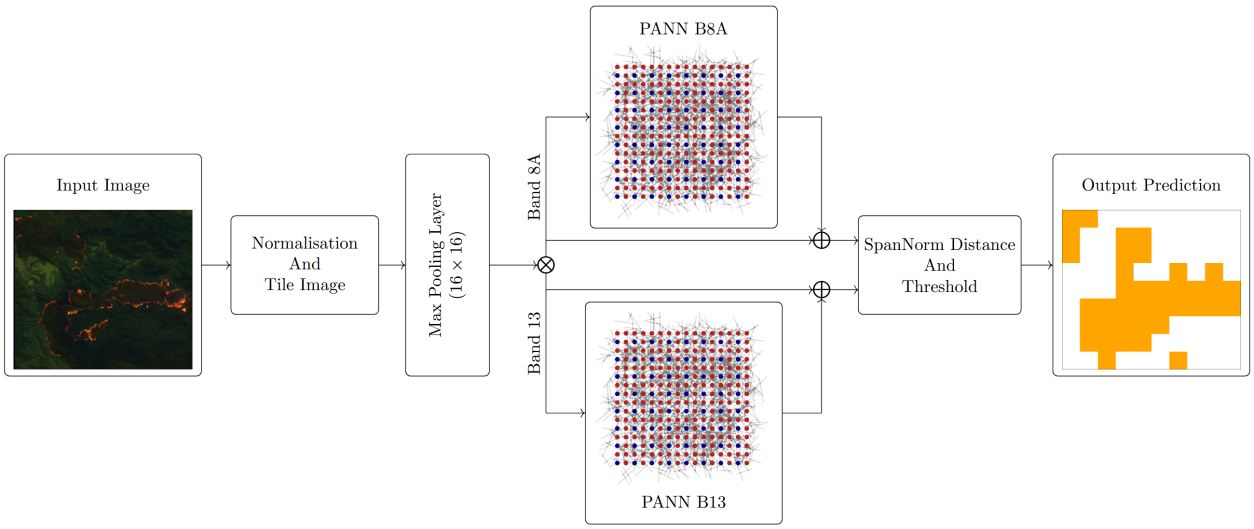

Figure 2. Diagram of the model pipeline. Starting with the input image (left), the image is fist normalised and tiled into smaller 128×128 pixel patches, followed by a 16×16 max pooling layer. The two spectral bands, B8A and B13, are split as denoted by $\otimes$ and fed into identical PANNs. Output feature vectors are concatenated with a skip connection, denoted by $\oplus$. Finally, the spanning normal distances between the two feature vectors is calculated and a threshold value is applied to classify each tile and produce a predicted output map (right).

is used, to combine the inputs into the PANN and with the output feature vector, to create the final feature vector. The PANNs used to process each band were identical. Finally, the spanning normal distance, given by Eq. (6),

$$\text{SpanNorm}(\vec{x}, \vec{y}) = \max(\vec{x} - \vec{y}) - \min(\vec{x} - \vec{y}) \quad (6)$$

between the feature vectors of the two bands is calculated. If the distance is above a predefined threshold, 1.68 and 0.92 for the raw and L1C datasets, respectively, that tile is classified as an "event" tile, otherwise it is classified as a "non-event" tile.

2.4. Compute Metrics

All processing was performed on an Intel® Xeon® Platinum 8268 Processor using only 16 CPU cores. The total time and maximum memory required for the PANN model, including all necessary programs and packages, to process each granule was calculated and then averaged across 5 granules. In addition, a time and memory usage was estimated for the hardware based nanowire network by performing all the steps shown in Fig. 2 except the PANN simulation. Instead, random numbers were generated to replicate the output of the nanowire network. The processing time was also added to estimate the total time and the estimated time and memory were averaged over 5 granules.

Table 1. MCC, precision and recall results for the PANN model for both raw and L1C datasets.

	raw data	L1C data
MCC	0.809	0.875
Precision	0.763	0.887
Recall	0.867	0.869

3. Results and Discussion

3.1. Model Evaluation

The performance of the model measured using the MCC, precision and recall for both the raw and L1C datasets is reported in Table 1. Since the PANN does not need to be trained, the entire dataset was used for performance testing. For MCC, the L1C data performed better with a value of 0.875, however, the raw data is not far behind at 0.809. Thus, the extra cleaning done on the L1C data has helped removed noise to increase performance. The relatively small difference between the L1C and raw MCC scores (0.066) shows that the PANN model is cable of detecting thermal hotspots in the noisier raw data.

Table 1 also reports nearly identical recall performance for both the raw and L1C data at 0.867 and 0.869, respectively. The confusion matrices for the raw and L1C data are shown in Fig. 3, where it can be seen that the difference in recall is due to only 1 more False Negative prediction in the raw data compared to the L1C. The real difference in per-

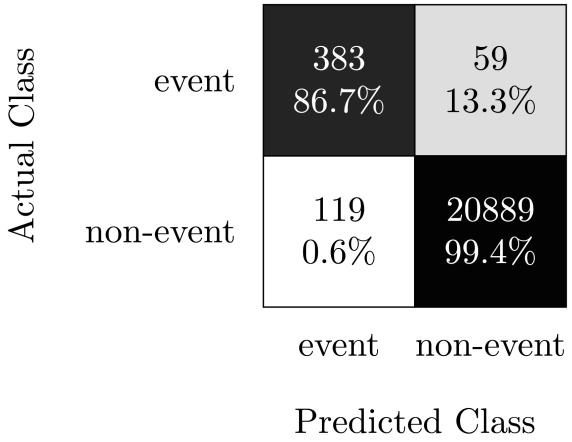

(a) Confusion matrix for the raw data.

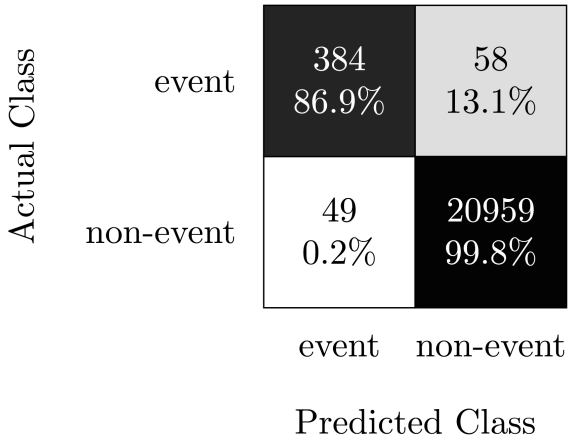

(b) Confusion matrix for the L1C data.

Figure 3. Confusion matrix for the raw (a) and L1C (b) datasets, showing the number of samples and the percentage for that class. The background colour shading represents the class percentage, with white for 0% and black for 100%.

formance between the raw and L1C datasets is due to the number of False Positives as highlighted by the difference in the precision values of 0.763 and 0.887, respectively. The number of False Positives is more than double in raw compared to L1C as seen in Fig. 3, caused by the extra noise in the raw dataset.

Note, however, that the threshold value could be changed to increase either the precision or recall, depending on the mission objectives when deployed. Figure 4 shows how the precision and recall values vary as the threshold value is adjusted and also highlights how the recall between the raw and L1C datasets is closer than the precision. In this study threshold values were chosen to optimise the MCC score.

Figure 5 shows 3 examples of fires and volcanoes each, for both the raw and L1C datasets. Each example, an RGB image using bands B8A, B11 and B12, a target mask and

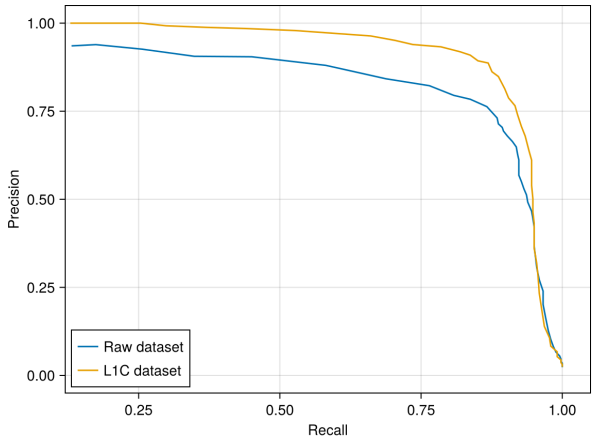

Figure 4. Precision–recall curve for the raw (blue) and L1C (orange) datasets, for a varying threshold value.

a predicted thermal mask, with tiles containing thermal anomalies highlighted in red and tiles predicted to have thermal anomalies highlighted in orange. The examples with significant differences between the raw and L1C predictions were chosen. It can be seen that the largest difference between raw and L1C is the increased number of False Positives as already noted. Although the L1C data gives an overall higher performance, it is not necessarily the case that the L1C data always gives an equivalent or better result than the raw data; an example of such a case is shown in Example 3 for both the fires and volcanoes in Fig. 5, where the model fails to predict thermal anomaly tiles when using the L1C data, but detects them when using the raw data.

3.2. Model Compute Resources

Table 2 reports the time and maximum memory required to process a granule, for the steps shown in Fig. 2, as well as a corresponding estimate for the hardware nanowire network device, which the PANN is based on. The most significant difference between the PANN and the estimated hardware is the time it takes to process a granule, with average times of 2.44 ± 0.09 s and 0.1290 ± 0.0002 s, respectively. Importantly, the time for the sensor to acquire a granule is 3.6 s, which means both the PANN model and hardware should be able to process a single granule within the acquisition time. Additionally, the hardware should be able to operate in real time, as it could process all 12 granules within the acquisition time. While in this study the PANN model was run on hardware that would not be available for deployment in space, we used only 16 CPU cores to limit computing resources and there is the potential to deploy the PANN on GPUs or custom accelerator chips to further speed up processing time. Likewise, the estimated hardware requirement of 0.393 ± 0.004 Gb is substantially less than the mem-

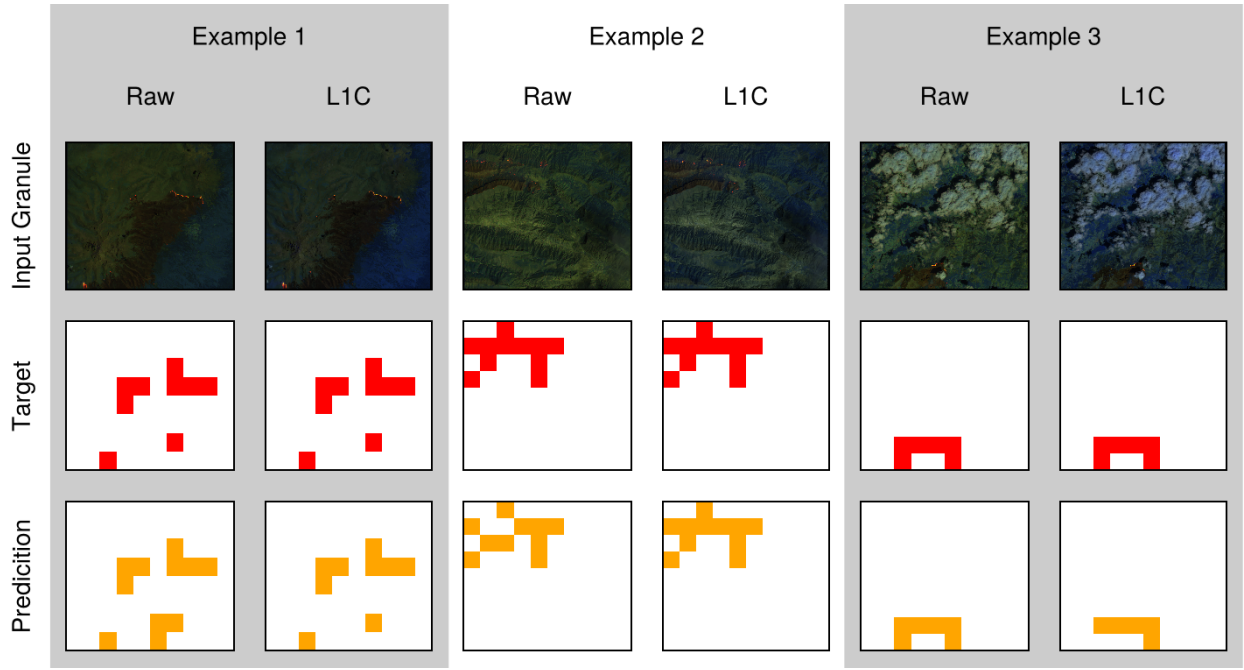

(a) Three example fire thermal hotspots.

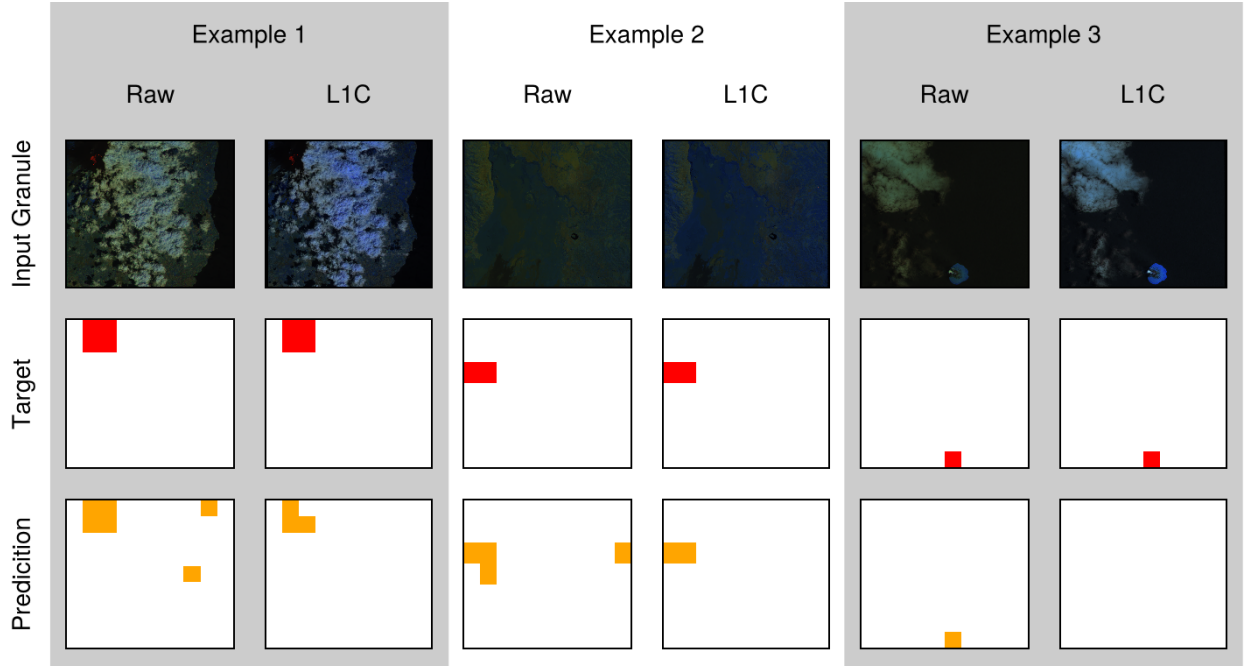

(b) Three example volcano thermal hotspots.

Figure 5. Three examples for fires (a) and volcanoes (b) for both the raw and L1C datasets. Each example shows an RGB image, using bands B8A, B11 and B12 (top row), along with the target output (middle row), where tiles containing thermal anomalies are shown in red, and the predicted output (bottom row), where tiles predicted by the PANN to have thermal anomalies are shown in orange.

Table 2. Compute resources calculated for the PANN model and estimated for the hardware nanowire network device, including time to process a granule in seconds and the total memory requirement of the model and all required software and packages. For all values the ±SEM is over 5 runs.

	PANN model	Estimated Hardware
Time (s / granule)	2.44 ± 0.09	0.1290 ± 0.0002
Memory (GB)	0.673 ± 0.007	0.393 ± 0.004

ory requirement of 0.673 ± 0.007 Gb for the PANN model. We note that the memory required includes loading all necessary packages to run the model, as well as the memory size of the model itself, as this is a more accurate measure of the operational implementation. We also note that both the PANN model and estimated hardware used 32 bit precision and could be further reduced to decrease memory size. Regardless, the memory size of 0.673 ± 0.007 Gb for the PANN model is still within a reasonable size for onboard deployment on modern hardware architectures.

A major advantage of the hardware nanowire network device is its low operational power and energy consumption. Using measured voltage and conductance values from a previous image classification study [30], the average power is $\sim 25\,\mu$W. The estimated time to process a granule, 0.129 s, then implies an energy consumption of just $3.23\,\mu$J per granule. Furthermore, there is no additional energy overhead associated with training time.

4. Conclusion

This study applied a PANN model to detect thermal hotspot anomalies from EO images, as a proof of concept towards onboard processing. The model was evaluated on the THRawS dataset at two different processing levels, raw and L1C from the Sentinel-2 satellites. The results confirmed that the L1C data gave the best model performance with an MCC of 0.875, however, the raw data was comparable at 0.809. This demonstrates the ability of the PANN model to handle the raw data, which is essential for any onboard processing. Additionally, it was found that the time required for the PANN model to process a granule, 2.44 ± 0.09 s using only CPUs, was within the acquisition time of 3.6 s, which would allow for a single granule to be processed, but not all 12 granules acquired by the sensor onboard the Sentinel-2 satellites. For the hardware nanowire network, on which the PANN is based, the estimates processing time is 0.1290 ± 0.0002 s, well below the acquisition threshold and could processes all 12 granules in real time. The memory requirements, for both the PANN and estimated hardware, were within a realistic size for onboard deployment at 0.673 ± 0.007 Gb and 0.393 ± 0.004 Gb, respectively, albeit on the high side. Further memory optimisation, such as quantising the model from 32 bit precision to 16-bit, could make the memory requirements more feasible. Altogether, these results show that the PANN model is a viable candidate for processing data onboard satellites.

Additionally and most importantly, as demonstrated in this thermal hotspot detection task, the PANN model can be implemented in a way that does not require training. This allows the model to be used when labelled data is not available or is too expensive to label for the sensor being used and means it does not need retraining with more new data due to domain shift or sensor drift.

Future work includes implementing the PANN model on a GPU architecture to accelerate processing time and determine the actual time and memory requirements of a hardware nanowire network device to verify the estimates calculated in this study. Testing the PANN model on a variety of tasks and using data from different sensors should also be conducted to determine the generalisability of the model for handling raw sensor data.

References

[1] David MJS Bowman, Grant J Williamson, John T Abatzoglou, Crystal A Kolden, Mark A Cochrane, and Alistair MS Smith. Human exposure and sensitivity to globally extreme wildfire events. *Nature ecology & evolution*, 1(3): 0058, 2017. 1

[2] David MJS Bowman, Crystal A Kolden, John T Abatzoglou, Fay H Johnston, Guido R van der Werf, and Mike Flannigan. Vegetation fires in the anthropocene. *Nature Reviews Earth & Environment*, 1(10):500–515, 2020. 1

[3] Francesco Caravelli and Juan Pablo Carbajal. Memristors for the curious outsiders. *Technologies*, 6(4):118, 2018. 2

[4] Francesco Caravelli, Gianluca Milano, Adam Z Stieg, Carlo Ricciardi, Simon Anthony Brown, and Zdenka Kuncic. Self-organising memristive networks as physical learning systems. *arXiv:2509.00747*, 2025. 2

[5] Bharadwaj Chintalapati, Arthur Precht, Sougata Hanra, Rene Laufer, Marcus Liwicki, and Jens Eickhoff. Opportunities and challenges of on-board ai-based image recognition for small satellite earth observation missions. *Advances in space research*, 75(9):6734–6751, 2025. 2

[6] Roberto Del Prete, Manuel Salvoldi, Domenico Barretta, Nicolas Longépé, Gabriele Meoni, Arnon Karnieli, Maria Daniela Graziano, and Alfredo Renga. Enhancing maritime situational awareness through end-to-end onboard raw data analysis. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 2025. 2

[7] Matthias Drusch, Umberto Del Bello, Sébastien Carlier, Olivier Colin, Veronica Fernandez, Ferran Gascon, Bianca Hoersch, Claudia Isola, Paolo Laberinti, Philippe Martimort, et al. Sentinel-2: Esa's optical high-resolution mission for gmes operational services. *Remote sensing of Environment*, 120:25–36, 2012. 1

[8] Adam J Durant, Costanza Bonadonna, and Claire J Horwell. Atmospheric and environmental impacts of volcanic particulates. *Elements*, 6(4):235–240, 2010. 1

[9] Mike Flannigan, Alan S Cantin, William J De Groot, Mike Wotton, Alison Newbery, and Lynn M Gowman. Global wildland fire season severity in the 21st century. *Forest Ecology and Management*, 294:54–61, 2013. 1

[10] A Gatti and A Bertolini. Sentinel-2 products specification document. *Rapport technique*, pages 4–7, 2015. 1

[11] Daniel A Jaffe, Susan M O'Neill, Narasimhan K Larkin, Amara L Holder, David L Peterson, Jessica E Halofsky, and Ana G Rappold. Wildfire and prescribed burning impacts on air quality in the united states. *Journal of the Air & Waste Management Association*, 70(6):583–615, 2020. 1

[12] Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard Lanas Phillips, Irena Gao, et al. Wilds: A benchmark of in-the-wild distribution shifts. In *International conference on machine learning*, pages 5637–5664. PMLR, 2021. 2

[13] Zdenka Kuncic and Tomonobu Nakayama. Neuromorphic nanowire networks: Principles, progress and future prospects for neuro-inspired information processing. *Advances in Physics: X*, 6(1):1894234, 2021. 2

[14] Zdenka Kuncic, Omid Kavehei, Ruomin Zhu, Alon Loeffler, Kaiwei Fu, Joel Hochstetter, Mike Li, James M. Shine, Adrian Diaz-Alvarez, Adam Stieg, James Gimzewski, and Tomonobu Nakayama. Neuromorphic information processing with nanowire networks. In *2020 IEEE International Symposium on Circuits and Systems (ISCAS)*, pages 1–5, 2020. 3

[15] Paolo Lunghi, Stefano Silvestrini, Dominik Dold, Gabriele Meoni, Alexander Hadjiivanov, and Dario Izzo. Energy efficiency analysis of spiking neural networks for space applications. *arXiv preprint arXiv:2505.11418*, 2025. 2

[16] Francesco Massimetti, Diego Coppola, Marco Laiolo, Sébastien Valade, Corrado Cigolini, and Maurizio Ripepe. Volcanic hot-spot detection using sentinel-2: A comparison with modis–mirova thermal data series. *Remote Sensing*, 12 (5):820, 2020. 3

[17] Gonzalo Mateo-Garcia, Josh Veitch-Michaelis, Cormac Purcell, Nicolas Longepe, Simon Reid, Alice Anlind, Fredrik Bruhn, James Parr, and Pierre Philippe Mathieu. In-orbit demonstration of a re-trainable machine learning payload for processing optical imagery. *Scientific Reports*, 13(1):10391, 2023. 2

[18] Gabriele Meoni, Roberto Del Prete, Federico Serva, Alix De Beusscher, Olivier Colin, and Nicolas Longépé. Unlocking the use of raw multispectral earth observation imagery for onboard artificial intelligence. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 17:12521–12537, 2024. 1, 2, 3

[19] Gabriele Meoni, Roberto Del Prete, Lucia Ancos-Villa, Enrique Albalate-Prieto, David Rijlaarsdam, Jose Luis Espinosa-Aranda, Nicolas Longépé, Maria Daniela Graziano, and Alfredo Renga. E2e: Onboard satellite real-time classification of thermal hotspots events on optical raw data. *Astrodynamics*, pages 1–17, 2025. 2

[20] MJ Paul, SD LeDuc, MG Lassiter, LC Moorhead, PD Noyes, and SG Leibowitz. Wildfire induces changes in receiving waters: A review with considerations for water quality management. *Water Resources Research*, 58(9):e2021WR030699, 2022. 1

[21] Vít Růžička, Anna Vaughan, Daniele De Martini, James Fulton, Valentina Salvatelli, Chris Bridges, Gonzalo Mateo-Garcia, and Valentina Zantedeschi. Ravæn: unsupervised change detection of extreme events using ml on-board satellites. *Scientific reports*, 12(1):16939, 2022. 3

[22] Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superglue: Learning feature matching with graph neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4938–4947, 2020. 3

[23] Stephen Smith, Cormac Purcell, and Zdenka Kuncic. Training-free ai for earth observation change detection using physics aware neuromorphic networks. *Scientific Reports*, 15(1):35292, 2025. 2, 3

[24] Adam Z. Stieg, Audrius V. Avizienis, Henry O. Sillin, Cristina Martin-Olmos, Masakazu Aono, and James K. Gimzewski. Emergent criticality in complex turing b-type atomic switch networks. *Advanced Materials*, 24(2): 286–293, 2011. 3

[25] Nguyen Thanh Toan, Phan Thanh Cong, Nguyen Quoc Viet Hung, and Jun Jo. A deep learning approach for early wildfire detection from hyperspectral satellite images. In *2019 7th International Conference on Robot Intelligence Technology and Applications (RiTA)*, pages 38–45. IEEE, 2019. 1

[26] Cristopher Castro Traba, David Rijlaarsdam, Jian Guo, Roberto Del Prete, and Gabriele Meoni. Towards onboard thermal hotspots segmentation with raw multispectral satellite imagery. *International Journal of Applied Earth Observation and Geoinformation*, 146:105095, 2026. 2

[27] Alexander Vahl, Gianluca Milano, Zdenka Kuncic, Simon A Brown, and Paolo Milani. Brain-inspired computing with self-assembled networks of nano-objects. *Journal of Physics D: Applied Physics*, 57(50):503001, 2024. 2

[28] Michael A Wulder, Thomas R Loveland, David P Roy, Christopher J Crawford, Jeffrey G Masek, Curtis E Woodcock, Richard G Allen, Martha C Anderson, Alan S Belward, Warren B Cohen, et al. Current status of landsat program, science, and applications. *Remote sensing of environment*, 225:127–147, 2019. 1

[29] Ruomin Zhu, Joel Hochstetter, Alon Loeffler, Adrian Diaz-Alvarez, Tomonobu Nakayama, Joseph T Lizier, and Zdenka Kuncic. Information dynamics in neuromorphic nanowire networks. *Scientific Reports*, 11(1):13047, 2021. 2

[30] Ruomin Zhu, Sam Lilak, Alon Loeffler, Joseph Lizier, Adam Stieg, James Gimzewski, and Zdenka Kuncic. Online dynamical learning and sequence memory with neuromorphic nanowire networks. *Nature Communications*, 14(1):6697, 2023. 7