

Event-Driven Neuromorphic AI for Brain Tumor Classification: A CNN–SNN Framework with Hardware Validation

Kalpana Telkar¹, K. Anusudha²

^{1,2}Department of Electronics and Communication Engineering Pondicherry University, India

Emails: telkarkalpana@gmail.com¹

Abstract

Classifying brain tumors using magnetic resonance imaging (MRI) is a crucial medical image processing problem; however, deployment on platforms with limited resources is limited by the computational expense of traditional deep learning models. Convolutional neural network (CNN) feature extraction, spike encoding, spiking neural network (SNN) inference, fixed-point quantization, and hardware validation are all integrated in this paper's Event-Driven Fixed-Point Neuromorphic Inference Engine (ED-FPNIE). Compact features are extracted by a lightweight CNN and transformed into temporal spike sequences for SNN-based classification. While inference employs event-driven conditional accumulation in place of traditional multiplier operations, quantized synaptic weights are mapped to ROM-based memory. A four-class brain MRI dataset consisting of the Glioma, Meningioma, No Tumor, and Pituitary categories was used to assess the framework. With class-wise accuracies of 69.00%, 46.25%, 96.00%, and 82.25%, respectively, RTL validation on 1600 test images yielded an accuracy of 73.38%. Using the Nangate45 library, gate-level synthesis produced a cell area of $7866.68 \mu\text{m}^2$ and an estimated power of 1.38 mW. The results demonstrate the feasibility of a hardware-oriented event-driven neuromorphic inference framework for resource-constrained medical AI applications.

Keywords: Brain Tumor Classification, Magnetic Resonance Imaging (MRI), Convolutional Neural Network (CNN), Spiking Neural Network (SNN), Neuromorphic Computing, Event-Driven Inference, Fixed-Point Quantization, Hardware Acceleration, RTL Validation, Medical Image Processing.

1. Introduction

In computer-aided medical diagnosis, brain tumor categorization using magnetic resonance imaging (MRI) is a crucial task where automated analysis can help precise and effective identification of various tumor classifications. Convolutional neural networks (CNNs) in particular have shown great promise for brain tumor identification and classification from MRI data in recent deep learning approaches (Dorfner et al. 2025; Bouhafra and El Bahi 2025). CNN-based classification utilizing MRI images has been the subject of several recent studies, including improved parallel CNN models, residual and regional CNNs, and customized CNN architectures (Reyes and Sánchez 2024; Albalawi et al. 2024; Zahoor et al. 2024; Li et al. 2024). However, the deployment of traditional deep-learning models on edge devices with limited resources may be limited by their computational requirements, which include frequent memory accesses and dense multiply-accumulate (MAC) operations. By using discrete neuronal events to represent information, neuromorphic computing

offers an alternate method for effective intelligent processing. For energy-efficient information processing, neuromorphic systems' event-driven and memory-centric features have been thoroughly investigated (Indiveri and Liu 2015). Spiking neural networks (SNNs) are appealing for low-power and edge-AI applications because they take advantage of sparse neural activity and temporal spike-based representations. The potential of SNNs for image classification has been shown by converting traditional deep networks into event-driven spiking networks while maintaining the benefits of deep feature extraction (Rueckauer et al. 2017). High-performance SNNs have also been trained thanks to spatiotemporal learning techniques, and more extensive research has shown the advantages and disadvantages of integrating deep learning with spiking neurons (Wu et al. 2018; Pfeiffer and Pfeil 2018). Numerous FPGA- and accelerator-based implementations have been inspired by the hardware efficiency of SNNs. Instead of continually assessing

dense neural computations, Minitaur processed spike events to demonstrate event-driven FPGA acceleration (Neil and Liu 2014). Effective digital neuron and synaptic computation for low-power inference has since been investigated by FPGA-based SNN implementations (Han et al. 2020). In order to minimize needless computation and memory usage in digital SNN hardware, selective input sparsity has been studied (Leigh, Heidarpur, and Mirhassani 2023). In order to increase scalability and energy efficiency, more recent architectures have further utilized multicore processing and structured weight sparsity (Wang et al. 2025). Additionally, trainable and reconfigurable SNN accelerators have been created to offer more versatility. While Yin et al. suggested a reconfigurable FPGA accelerator that takes use of temporal and spatial sparsity (Yin et al. 2024), Mehrabi et al. exhibited an online-trainable FPGA-based SNN architecture (Mehrabi et al. 2024). For distributed neuromorphic systems, event-driven communication, synchronization, routing, and hardware mapping have also been studied at higher scales (Shahsavari et al. 2023). The possibility of specialized hardware for effective SNN inference is further demonstrated by recent event-driven accelerator designs (Cheng et al. 2024). Despite these advances, existing studies generally emphasize either medical-image classification or neuromorphic hardware acceleration. A unified workflow that connects MRI feature extraction, spike generation, fixed-point SNN inference, and synthesizable hardware validation for multi-class brain tumor classification is less commonly addressed. The proposed work targets this integration by establishing a direct software-to-hardware path from MRI features to event-driven digital inference. Additionally, hardware-oriented arithmetic optimization has been studied for edge-AI image processing applications, highlighting the significance of lowering arithmetic complexity in systems with limited resources (Tammineni et al. 2025). Lightweight CNN-based feature extraction, temporal spike encoding, fixed-point SNN inference, ROM-based weight storage, event-driven accumulation, and synthesizable RTL implementation are all integrated into the suggested framework. The suggested architecture conducts conditional fixed-point accumulation and only

processes active spike events, in contrast to traditional multiplier-based inference. 1600 test MRI pictures are used to validate the entire software-to-hardware cycle, which is then further assessed using gate-level synthesis and hardware performance estimation. The goal is to show that an integrated neuromorphic AI system for resource-constrained medical image classification is feasible.

This work's primary contributions are:

- A unified CNN–SNN framework for four-class brain tumor classification from MRI images.
- During SNN inference, an event-driven fixed-point inference architecture avoids traditional multiplier-based operations and processes only active spike events.
- A complete software-to-hardware validation process that includes gate-level hardware analysis, synthesizable RTL implementation, fixed-point weight mapping, and spike creation.
- Gate-level area, timing, and power estimation using the Nangate45 standard-cell library combined with experimental validation on 1600 test images

The remainder of this paper is organized as follows. Section 2 reviews related work in medical image classification, SNNs, and neuromorphic hardware. Section 3 presents the proposed ED-FPNIE architecture. Section 4 describes the hardware-oriented RTL implementation. Section 5 presents gate-level synthesis and hardware analysis. Section 6 discusses the classification and hardware results, and Section 7 concludes the paper.

2. Related Work

An energy-efficient substitute for traditional von Neumann architectures in the implementation of intelligent systems is neuromorphic computing. By proving the potential of Spiking Neural Networks (SNNs) for low-power hardware implementations, Indiveri and Liu (Indiveri and Liu 2015) established the basic concepts of event-driven computation and brain-inspired information processing. A number of FPGA-based SNN accelerators have been put forth to take advantage of spike processing's sparse and event-driven characteristics. By processing only

active spikes, Minitaur (Neil and Liu 2014) presented an event-driven FPGA accelerator that minimizes needless computations. In a similar vein, Han et al. (Han et al. 2020) showed effective FPGA implementation of SNNs for low-power inference applications. Recent research has concentrated on using sparsity exploitation to increase hardware efficiency. To minimize memory accesses and computational overhead in digital SNN implementations, Leigh et al. (Leigh, Heidarpur, and Mirhassani 2023) introduced selected input sparsity approaches. The Cactus accelerator (Wang et al. 2025), developed by Wang et al., uses structured weight sparsity to increase the scalability and energy efficiency of multicore SNN architectures. Reconfigurable and trainable neuromorphic systems have also been investigated to improve flexibility and adaptability. Mehrabi et al. (Mehrabi et al. 2024) presented an FPGA-based online-trainable SNN architecture, while Yin et al. (Yin et al. 2024) proposed a reconfigurable SNN accelerator exploiting spatial and temporal sparsity. For distributed neuromorphic systems, large-scale event-driven communication and synchronization strategies have been investigated (Shahsavari et al. 2023). To increase energy economy in edge-AI systems, low-power hardware optimization strategies, such as approximation arithmetic architectures, have also been investigated. The usefulness of optimized recursive approximate multipliers for image processing applications was shown by Tammineni et al. (Tammineni et al. 2025), who achieved notable energy savings without sacrificing acceptable output quality. Despite these developments, the majority of current research focuses on either hardware acceleration or software-level SNN optimization separately. End-to-end workflows connecting neural-network feature extraction, spike generation, fixed-point SNN inference, and synthesizable RTL validation for medical image classification remain comparatively less explored. A fully synthesizable Event-Driven Fixed-Point Neuromorphic Inference Engine (ED-FPNIE) that combines software training, spike encoding, fixed-point quantization, and RTL-level validation into a single framework is presented in the proposed work to close this gap.

3. Proposed ED-FPNIE Architecture

By combining CNN-based feature extraction, spike encoding, Spiking Neural Network (SNN) inference, fixed-point quantization, and hardware-oriented RTL realization, the suggested Event-Driven Fixed-Point Neuromorphic Inference Engine (ED-FPNIE) offers a comprehensive framework for brain tumor classification. MRI images are first processed using a lightweight CNN to produce compact feature vectors, as illustrated in Fig. 1. For SNN training, these features are transformed into temporal spike sequences. The RTL implementation receives encoded spike sequences as inputs, and the learnt synaptic weights are quantized into fixed-point representation and output to hardware memory.




Figure 1 End-to-end CNN-SNN-hardware workflow for brain tumor classification

4. CNN Feature Extraction

MRI images are first converted into compact feature

representations using a lightweight CNN. The input images are represented as 32×32 grayscale images, and the CNN produces a 256-dimensional feature vector for each image. This compact representation reduces the dimensionality of the input supplied to the subsequent SNN stage while retaining discriminative image information.

5. Spike Encoding

The extracted feature vector is transformed into a temporal spike representation using rate-based encoding. Each feature is represented by a binary spike sequence over $T=30$ time steps, where $S_{(i,t)} \in \{0,1\}$ denotes the spike generated by feature i at time t . This temporal representation enables the subsequent SNN inference stage to operate on discrete events rather than continuous-valued activations [1 – 10].

6. Fixed-Point SNN Inference

The encoded spike sequences are processed by four output class neurons corresponding to Glioma, Meningioma, No Tumor, and Pituitary. The trained synaptic weights are quantized into a fixed-point representation before being transferred to the hardware implementation. With 256 features and $T=30$ temporal steps, each input produces $256 \times 30 = 7680$ binary spike values for hardware processing. The proposed design only performs accumulation when spike events are present, in contrast to traditional neural accelerators that rely on dense multiply-accumulate operations [11–13]. Let the set of active spike events be denoted by Ω . The output neuron c 's accumulated class response is calculated as

$$V_c = \sum_{t=1}^T \sum_{i=1}^N S_{i,t} W_{c,i},$$

where $S_{(i,t)} \in \{0,1\}$ denotes the spike generated by feature i at time t , $W_{(c,i)}$ is the quantized synaptic weight for class neuron c , $N=256$ is the number of extracted features, and $T=30$ is the number of temporal steps. Since the spike value is binary, an active event contributes the corresponding weight, whereas an inactive event contributes zero. The binary nature of the spike representation enables the multiplication term to be implemented as a conditional accumulation. Consequently, the

hardware processes only active spike events and avoids conventional multiplier-based operations during inference.

7. Event-Driven Computation

The proposed ED-FPNIE exploits the binary nature of the spike representation to perform event-driven computation during SNN inference. Instead of continuously evaluating arithmetic operations for all input features and temporal steps, the hardware activates the accumulation operation only when an input spike is present. For each spike-memory entry, the corresponding spike value is first retrieved and used as an enable condition for the accumulation datapath. When $S_{(i,t)}=1$, the corresponding quantized synaptic weight $W_{(c,i)}$ is added to the accumulated response of class neuron c . When $S_{(i,t)}=0$, no weight addition is performed and the previously accumulated value is retained. This conditional operation eliminates the need for a conventional multiplier in the SNN inference datapath because the input spike is binary. The multiplication between $S_{(i,t)}$ and $W_{(c,i)}$ therefore reduces to either adding the stored weight or adding zero. The hardware thus performs fixed-point addition only for active spike events, reducing unnecessary arithmetic activity associated with inactive inputs. This event-driven processing principle is consistent with hardware SNN architectures that exploit sparse spike activity to improve computational efficiency [11–13]. For each input sample, the spike memory contains $N \times T = 256 \times 30 = 7680$ binary spike values. The control and address-generation logic sequentially traverses these entries and provides the corresponding feature address to the class-specific weight memories. Four parallel 32-bit signed accumulators maintain the class responses for Glioma, Meningioma, No Tumor, and Pituitary. Consequently, the same active spike can enable the accumulation of the corresponding weight in all four class paths. After all 7680 spike entries have been processed, the accumulated class responses are provided to the MUX-based argmax decision logic, which selects the class having the maximum response as the final prediction. The resulting computation structure separates spike storage, weight storage, event detection, conditional accumulation, and class selection into dedicated

hardware blocks. This organization provides a simple and synthesizable datapath for fixed-point SNN inference and allows the hardware implementation to exploit the event-driven nature of the input representation [11,13].

8. Hardware - Oriented ED-FPNIE Implementation

The proposed Event-Driven Fixed-Point Neuromorphic Inference Engine (ED-FPNIE) was implemented in Verilog HDL as a synthesizable RTL architecture for hardware-oriented SNN inference. As shown in Fig. 2, the architecture consists of four primary components: (i) Spike ROM, (ii) Weight ROM, (iii) Event-Driven Accumulator, and (iv) MUX-based Decision Logic. This organization follows the event-driven processing principle of SNN hardware, in which computation is activated by incoming spike events rather than continuously processing dense neural operations (Neil and Liu 2014; Han et al. 2020; Leigh, Heidarpur, and Mirhassani 2023). The input MRI features are first converted into temporal spike sequences and stored in spikes.hex, while the quantized synaptic weights obtained from SNN training are stored in weights.hex. For each input sample, the spike memory stores $256 \times 30 = 7680$ binary spike values, whereas four class-specific weight memories store 256 quantized 16-bit weights corresponding to the Glioma, Meningioma, No Tumor, and Pituitary classes. During inference, the control and address-generation logic sequentially retrieves the stored spike events and provides the corresponding feature address to the class-specific weight memories. When an active spike is detected, the associated quantized synaptic weight is conditionally accumulated into the corresponding class response. For an inactive spike, no weight addition is performed and the previous accumulated value is retained. Thus, the binary spike representation eliminates the need for a conventional multiplier in the accumulation datapath, since the product of a binary spike and a weight is equivalent to either adding the stored weight or adding zero. This event-driven conditional accumulation reduces unnecessary arithmetic activity associated with inactive spike events and is consistent with sparsity-aware SNN hardware architectures (Leigh, Heidarpur, and Mirhassani 2023; Wang et al. 2025;

Yin et al. 2024). The hardware uses four parallel 32-bit signed class-response accumulators, one for each output class. Let $V_c^{(k)}$ denote the accumulated response of class neuron c after the k th processed spike-memory entry. The accumulator update is expressed as

$$V_c^{(k+1)} = \begin{cases} V_c^{(k)} + W_{c,i}, & \text{if } S_{i,t} = 1, \\ V_c^{(k)}, & \text{otherwise,} \end{cases}$$

where $S_{i,t} \in \{0,1\}$ denotes the spike generated by feature i at time step t , and $W_{c,i}$ represents the quantized synaptic weight associated with feature i and class c . Since the input spike is binary, the accumulation operation requires only fixed-point addition when an event is present. Conventional multiplier operations are therefore avoided in the SNN inference datapath. After all $256 \times 30 = 7680$ spike-memory entries have been processed, the four accumulated class responses are supplied to the MUX-based decision logic. The predicted class is selected according to the maximum accumulated response as [14 – 19]

$$\hat{y} = \underset{c}{\operatorname{argmax}} V_c,$$

where $\hat{y}$ denotes the predicted tumor class. The MUX-based argmax logic implements the class comparison and generates the corresponding 2-bit class output. This organization separates spike storage, synaptic-weight storage, event detection, conditional accumulation, and class selection into dedicated hardware blocks, providing a simple synthesizable datapath for fixed-point SNN inference.




Figure 2 RTL architecture of the proposed ED-FPNIE showing spike and weight memories, event-driven accumulation, and MUX-based classification




Figure 2 Mux Based Argmax

9. Hardware Validation and Performance Analysis

The proposed ED-FPNIE was implemented in Verilog HDL and evaluated using a gate-level synthesis flow. The RTL was synthesized using the Yosys open-source synthesis framework and mapped to the Nangate45 standard-cell library. The synthesis flow comprised RTL elaboration, logic optimization, technology mapping, and gate-level netlist generation. The resulting technology-mapped netlist was subsequently analyzed using OpenROAD for area, timing, and power estimation. The reported hardware metrics represent ASIC-oriented gate-level pre-layout estimates and are used to assess the hardware feasibility of the proposed architecture. The overall gate-level hardware organization is illustrated in Fig. [fig:gate_flow]. The implementation consists of four principal functional blocks: (i) control and address generation, (ii) spike and weight memories, (iii) event-driven accumulation, and (iv) MUX-based argmax decision logic. The control and address-generation logic sequentially accesses the stored spike events and the corresponding class-specific weights. When an active spike is detected, the associated quantized weight is accumulated into the corresponding class-response accumulator. After all spike events have been processed, the MUX-based

argmax logic selects the class with the maximum accumulated response. The synthesis results are summarized in Table 1. The synthesized design contains 6641 standard cells, comprising 6272 combinational cells and 141 sequential cells, with 228 inverter cells. The resulting cell area is $7866.68 \mu\text{m}^2$. The maximum data arrival time is 8.9789 ns, compared with a required time of 9.9645 ns, resulting in a positive timing slack of 0.9856 ns. The positive slack indicates that the synthesized design satisfies the applied timing constraint under the evaluated synthesis and analysis conditions. As shown in Gate-level synthesis and timing results of the proposed ED-FPNIE.

Table 1 Gate-level synthesis and timing results of the proposed ED-FPNIE

Parameter	Value
Technology Library	Nangate45
Total Standard Cells	6641
Sequential Cells	141
Combinational Cells	6272
Inverters	228
Cell Area	$7866.68 \mu\text{m}^2$
Data Arrival Time	8.9789 ns
Data Required Time	9.9645 ns
Timing Slack	0.9856 ns
Total Power	1.38 mW

The estimated total power consumption of the synthesized design is 1.38 mW. The power distribution comprises 54.8% internal power, 32.2% switching power, and 13.0% leakage power. The event-driven datapath performs conditional accumulation only when an active spike event is present. Consequently, inactive spike events do not trigger weight accumulation, avoiding unnecessary arithmetic activity during inference. The combination of ROM-based spike and weight storage, conditional fixed-point accumulation, and MUX-based decision logic provides a compact digital implementation of the SNN inference stage. Overall, the gate-level results demonstrate the hardware feasibility of the

proposed ED-FPNIE under the evaluated ASIC-oriented synthesis flow. The reported cell area, timing, and power values provide an initial assessment of the implementation cost of the proposed architecture. These results are limited to gate-level pre-layout estimates and do not represent post-layout measurements, fabricated-chip measurements, or direct FPGA implementation results. Further work can investigate physical implementation, memory organization, sparsity-aware processing, and post-layout optimization.

10. Results and Discussion

10.1. Dataset and Experimental Setup

The proposed framework was evaluated using a publicly accessible four-class brain MRI dataset containing Glioma, Meningioma, No Tumor, and Pituitary categories. The input MRI images were represented as 32×32 grayscale images and processed by the CNN feature extraction stage. The CNN generated a compact 256-dimensional feature vector for each input image. These features were subsequently converted into rate-based spike sequences over $T=30$ time steps for SNN-based classification. Thus, each input sample was represented by $256 \times 30 = 7680$ binary spike values, which were supplied to the hardware-oriented SNN inference engine. A total of 1600 test images were evaluated during RTL validation.

11. Classification Performance

Table 2 presents the class-wise classification performance obtained from the RTL validation. The proposed ED-FPNIE correctly classified 1174 out of 1600 test samples, resulting in an overall classification accuracy of 73.38%.

Table 2 Class-wise classification performance

Class	Accuracy (%)
Glioma	69.00
Meningioma	46.25
No Tumor	96.00
Pituitary	82.25
Overall	73.38

Among the four classes, the highest classification accuracy was obtained for the No Tumor category, with an accuracy of 96.00%. The Pituitary class achieved 82.25%, followed by Glioma with 69.00%. The Meningioma class obtained the lowest accuracy of 46.25%. This comparatively lower performance indicates that Meningioma remains a challenging class for the compact feature representation and SNN configuration employed in the present study. The observed classification performance is dependent on the selected CNN feature representation, spike encoding scheme, and SNN configuration. Therefore, the lower Meningioma accuracy represents an important limitation of the current implementation and motivates further investigation of improved feature representations and hardware-aware SNN training strategies.

12. Hardware Validation and Analysis

The gate-level results in Table 1 demonstrate that the proposed ED-FPNIE can be realized using a standard-cell synthesis flow. The positive timing slack of 0.9856 ns indicates that the design satisfies the applied timing constraint under the evaluated conditions. The ROM-based memory organization, conditional accumulation, and MUX-based argmax logic provide a compact implementation of the SNN inference stage. The reported area, timing, and power values are ASIC-oriented gate-level pre-layout estimates obtained using the Nangate45 standard-cell library and the synthesis/analysis flow. They do not represent post-layout measurements, fabricated-chip measurements, or direct FPGA implementation results. Further work can investigate physical implementation, memory optimization, sparsity-aware processing, and post-layout analysis.

13. Overall Discussion

The combined classification and hardware results demonstrate the feasibility of integrating CNN-based MRI feature extraction, rate-based spike encoding, fixed-point SNN inference, and event-driven digital hardware within a unified medical AI workflow. The achieved overall accuracy of 73.38% indicates that the proposed compact CNN-SNN pipeline can perform four-class brain tumor classification while maintaining a hardware-oriented inference structure. From the hardware perspective, the use of binary spike events, ROM-based synaptic weights,

conditional accumulation, and MUX-based decision logic provides a relatively simple digital architecture for implementing the SNN inference stage. The positive timing slack of 0.9856 ns together with the estimated cell area of $7866.68 \mu\text{m}^2$ and total power of 1.38 mW provides an initial assessment of the implementation cost of the proposed architecture. However, the classification results also reveal an important trade-off between model compactness and classification performance. While the No Tumor and Pituitary classes show comparatively higher accuracies, the Meningioma class remains difficult to classify accurately. Further improvements are therefore required in the CNN feature extraction, spike encoding, and SNN training stages. Future hardware optimization can additionally investigate more efficient memory organization, sparsity-aware computation, and post-layout optimization while preserving the classification performance of the complete CNN–SNN framework.

Conclusion

This paper presented an Event-Driven Fixed-Point Neuromorphic Inference Engine (ED-FPNIE) for four-class brain tumor classification from MRI images. The proposed framework integrates CNN-based feature extraction, temporal spike encoding, fixed-point SNN inference, and synthesizable hardware implementation within a unified workflow. RTL validation on 1600 test images achieved an overall classification accuracy of 73.38%, while gate-level synthesis using the Nangate45 library resulted in a cell area of $7866.68 \mu\text{m}^2$ and an estimated power consumption of 1.38 mW. These results demonstrate the feasibility of event-driven neuromorphic inference for resource-constrained medical AI applications. Future work will focus on improved feature representations, hardware-aware SNN training, larger MRI datasets, and post-layout hardware optimization.

References

- [1]. Albalawi, Eid, Arastu Thakur, D. Ramya Dorai, Surbhi Bhatia Khan, T. R. Mahesh, Ahlam Almusharraf, Khursheed Aurangzeb, and Muhammad Shahid Anwar. 2024. "Enhancing Brain Tumor Classification in MRI Scans with a Multi-Layer Customized Convolutional Neural Network Approach." *Frontiers in Computational Neuroscience* 18: 1418546. <https://doi.org/10.3389/fncom.2024.1418546>
- [2]. Bouhafra, Sara, and Hassan El Bahi. 2025. "Deep Learning Approaches for Brain Tumor Detection and Classification Using MRI Images (2020 to 2024): A Systematic Review." *Journal of Digital Imaging* 38 (3): 1403–33. <https://doi.org/10.1007/s10278-024-01283-8>.
- [3]. Cheng, Xi, Shu Cao, Shangmei Wang, Xiaoyang Zeng, et al. 2024. "A 1024-Neuron 1M-Synapse Event-Driven SNN Accelerator for DVS Applications." In *2024 IEEE International Symposium on Circuits and Systems (ISCAS)*. <https://doi.org/10.1109/ISCAS58744.2024.10558416>.
- [4]. Dorfner, Felix J., Jay B. Patel, Jayashree Kalpathy-Cramer, Elizabeth R. Gerstner, and Christopher P. Bridge. 2025. "A Review of Deep Learning for Brain Tumor Analysis in MRI." *Npj Precision Oncology* 9: 2. <https://doi.org/10.1038/s41698-024-00789-2>.
- [5]. Han, J., Z. Li, W. Zheng, and Y. Zhang. 2020. "Hardware Implementation of Spiking Neural Networks on FPGA." *Tsinghua Science and Technology* 25 (4): 479–86. <https://doi.org/10.26599/TST.2019.9010019>.
- [6]. Indiveri, Giacomo, and Shih-Chii Liu. 2015. "Memory and Information Processing in Neuromorphic Systems." *Proceedings of the IEEE* 103 (8): 1379–97.
- [7]. Leigh, A. J., M. Heidarpur, and M. Mirhassani. 2023. "Digital Hardware Implementations of Spiking Neural Networks with Selective Input Sparsity for Edge Inferences in Controlled Image Acquisition Environments." *IEEE Transactions on Circuits and Systems II: Express Briefs* 70 (5): 1724–28. <https://doi.org/10.1109/TCSII.2023.3260704>.
- [8]. Li, Chen, Faxue Zhang, Yongjian Du, and Huachao Li. 2024. "Classification of Brain Tumor Types Through MRIs Using Parallel CNNs and Firefly Optimization." *Scientific Reports* 14: 15057. <https://doi.org/10.1038/s41598-024-65714-w>.
- [9]. Mehrabi, A., Y. Bethi, A. van Schaik, A. Wabnitz, and S. Afshar. 2024. "Efficient

- Hardware Implementation of a Multi-Layer Gradient-Free Online-Trainable Spiking Neural Network on FPGA.” IEEE Access 13: 170980–93.
<https://doi.org/10.1109/ACCESS.2024.3500134>
- [10]. Neil, Daniel, and Shih-Chii Liu. 2014. “Minitaur, an Event-Driven FPGA-Based Spiking Network Accelerator.” IEEE Transactions on Very Large Scale Integration (VLSI) Systems 22 (12): 2621–28.
<https://doi.org/10.1109/TVLSI.2014.2303151>.
- [11]. Pfeiffer, Michael, and Thomas Pfeil. 2018. “Deep Learning with Spiking Neurons: Opportunities and Challenges.” Frontiers in Neuroscience 12: 774.
<https://doi.org/10.3389/fnins.2018.00774>.
- [12]. Reyes, Daniel, and Javier Sánchez. 2024. “Performance of Convolutional Neural Networks for the Classification of Brain Tumors Using Magnetic Resonance Imaging.” Heliyon 10 (3): e25468.
<https://doi.org/10.1016/j.heliyon.2024.e25468>.
- [13]. Rueckauer, Bodo, Iulia-Alexandra Lungu, Yuhuang Hu, Michael Pfeiffer, and Shih-Chii Liu. 2017. “Conversion of Continuous-Valued Deep Networks to Efficient Event-Driven Networks for Image Classification.” Frontiers in Neuroscience 11: 682.
<https://doi.org/10.3389/fnins.2017.00682>.
- [14]. Shahsavari, Mahyar, David Thomas, Marcel van Gerven, Andrew Brown, and Wayne Luk. 2023. “Advancements in Spiking Neural Network Communication and Synchronization Techniques for Event-Driven Neuromorphic Systems.” Array 20: 100323.
<https://doi.org/https://doi.org/10.1016/j.array.2023.100323>.
- [15]. Tammineni, V., S. K. Beura, M. V. H. B. Murthy, S. Majumdar, and P. Saha. 2025. “Optimized Recursive Approximate Multipliers for Edge Detection and Image Smoothing Applications.” Microsystem Technologies 31 (10): 2771–82. <https://doi.org/10.1007/s00542-024-05810-z>.
- [16]. Wang, Z., Z. Ou, Y. Zhong, and Y. Wang. 2025. “Cactus: A Multicore Spiking Neural Network Accelerator with Fine-Grained Structured Weight Sparsity.” IEEE Transactions on Very Large Scale Integration (VLSI) Systems 33 (12): 3535–39.
<https://doi.org/10.1109/TVLSI.2025.3608498>.
- [17]. Wu, Yujie, Lei Deng, Guoqi Li, Jun Zhu, and Luping Shi. 2018. “Spatio-Temporal Backpropagation for Training High-Performance Spiking Neural Networks.” Frontiers in Neuroscience 12: 331.
<https://doi.org/10.3389/fnins.2018.00331>.
- [18]. Yin, Mingqi, Xiaole Cui, Feng Wei, Hanqing Liu, Yuanyuan Jiang, and Xiaoxin Cui. 2024. “A Reconfigurable FPGA-Based Spiking Neural Network Accelerator.” Microelectronics Journal 152: 106377.
<https://doi.org/https://doi.org/10.1016/j.mejo.2024.106377>.
- [19]. Zahoor, Mirza Mumtaz, Saddam Hussain Khan, Tahani Jaser Alahmadi, Tariq Alsahfi, Alanoud S. Al Mazroa, Hesham A. Sakr, Saeed Alqahtani, Abdullah Albanyan, and Bader Khalid Alshemaimri. 2024. “Brain Tumor MRI Classification Using a Novel Deep Residual and Regional CNN.” Biomedicines 12 (7): 1395.
<https://doi.org/10.3390/biomedicines12071395>.