International Journal For Research in
Applied Science and Engineering Technology

INTERNATIONAL JOURNAL FOR RESEARCH IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 13 **Issue:** IX **Month of publication:** September 2025

DOI: https://doi.org/10.22214/ijraset.2025.74359

www.ijraset.com

Call: 08813907089 | E-mail ID: ijraset@gmail.com

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 13 Issue IX Sep 2025- Available at www.ijraset.com

Neuromorphic Computing: Hardware-Inspired Architectures for Energy-Efficient AI

Dr. Diwakar Ramanuj Tripathi[1], Prajwal Guniram Jagnade[2], Dr. Vrushali Pramod Parkhi[3]
[1]*Head, Department of Computer Science,* [2]*Research Scholar,* [3]*Officiating Principal, S.S. Maniar College of Computer & Management, Nagpur*

Abstract: The study explores the topic of neuromorphic computing which is a hardware-based paradigm to realizing energy-efficient artificial intelligence (AI). In contrast to conventional von Neumann architectures that processes are performed sequentially and memory bottlenecks are employed, neuromorphic systems implement an event-driven model of spiking neurons and massively parallel architecture through biological and neural dynamics. The descriptive-analytical design provided in this case is a synthesis of existing chip implementations (Intel Loihi, IBM TrueNorth, SpiNNaker) and models to examine performance, scalability and energy consumption. Results show that spiking neural networks (SNNs) implemented on a neuromorphic substrate can use as much as 100x energy than using a GPU-based deep learning and still achieve similar accuracy in classification and pattern-recognition problems. Moreover, neuromorphic chip provides scalability in edge AI implementation (IoT, robotics and sensory processing) where power and real time are essential. The paper highlights the fact that neuromorphic models are a viable direction of AI in the future, as they integrate performance, flexibility, and energy awareness.
Keywords: Neuromorphic Computing, Spiking Neural Networks, Hardware Architectures, Energy Efficiency, Artificial Intelligence.

I. INTRODUCTION

Artificial intelligence (AI) has experienced a fast progress in the last few decades, evolving beyond rule-based systems to sophisticated deep-learning models that can handle the tasks that were initially considered to be the prerogative of the human intelligence. This development has been informed by the fact that there have been advances in computing power, the availability of large datasets, and advanced learning algorithms. Nonetheless, the need of energy-efficient and scalable hardware has become especially pressing as AI models are becoming larger and more complicated. Conventional computing systems are based on von Neumann architecture which have natural bottlenecks in memory access and energy use which pose significant challenges to the development of sustainable AI. It is against this backdrop that neuromorphic computing has been developed as a new paradigm providing brain-like solutions in order to counter these disadvantages.

A. Emergence of Neuromorphic Computing

Neuromorphic computing has become an innovative paradigm of artificial intelligence (AI) based on the structural and functional concepts of the human brain. The neuromorphic systems combine these two functions by means of interconnected networks of spiking neurons and synapses, unlike classical von Neumann architectures, which strictly divide memory and computation. This brain inspired structure allows event driven processing, asynchronous communications and massively parallel computing and hence it decreases redundancy and greatly improves efficiency. Neuromorphic computing offers to bridge the gap between machine intelligence and human cognition by means of borrowing concepts of the biology of neural systems.

B. Limitations of Conventional AI Architectures

In the last ten years, deep learning and similar AI models have made incredible progress in several areas, including computer vision, natural language processing, and speech recognition. In spite of this, traditional AI designs of the time are still computationally expensive and sometimes required billions of floating-point operations, and used significant amounts of energy when running on GPUs and TPUs. The energy footprint is also too large to allow scalability, and it also reduces deployment to energy-constrained systems, such as embedded systems, mobile devices, and real-time edge systems. Such limitations underscore the necessity of having alternative architectures that are able to balance between high performance and energy-saving.

©IJRASET: All Rights are Reserved

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 13 Issue IX Sep 2025- Available at www.ijraset.com

C. Neuromorphic Computing as a Sustainable Alternative

Neuromorphic computing is able to overcome these issues by replicating the effectiveness of biological neural networks. As spiking neurons only use energy when signaling events, neuromorphic hardware can be designed to operate at ultra-low power consumption levels and still be quite parallel and adaptable. This is an event-based process that enables computation when needed, eliminating the wasted energy and enhancing responsiveness. In turn, neuromorphic architectures are becoming a promising direction of sustainable, scalable, and brain-like AI. This is why it is necessary to comprehend their possible benefits compared to existing architectures to create the next generation of energy-efficient artificial intelligence that can operate in practice without issues.

II. LITERATURE REVIEW

Akopyan et al. (2015) introduced the design and the development of TrueNorth, a neurosynaptic chip made of one million programmable neurons that consumed only 65 mW. The efficiency of the chip architecture and the tool flow with which large-scale neuromorphic computing is made possible were highlighted in their study. The authors established that the chip could save a lot of power and at the same time facilitate the real time cognitive tasks. This article was one of the first ones to demonstrate how big neuromorphic hardware can be used practically.

Bouvier et al. (2019) did a survey of hardware applications of the spiking neural network (SNN) in detail. They inspected different designs, both digital and analog, mixed-signal designs, and indicated the significant technological and architectural issues. They found that, although SNN hardware did show tremendous potential as far as energy consumption and biological viability were concerned, scalability, training and standardization problems had still not been addressed. The research was a resourceful tool towards the establishment of gaps in the development of neuromorphic hardware.

Chicca et al. (2014) narrowed down to neuromorphic electronic circuits and how these can be used to create autonomous cognitive systems. They reviewed bio-inspired design approaches incorporating sensory process, learning, motor control in small scale hardware. Their study revealed that the neuromorphic circuits had a capability of simulating some functions of the biological neural processing, and thus it was now possible to have autonomous and adaptive response in artificial systems. The research helped in closing the gap that existed between neuroscience principles and real time cognitive electronic systems.

Davies et al. (2018) proposed Loihi, a neuromorphic manycore processor, which added on-chip learning functions to it. Their article reported the manner in which the processor was created to model spiking neural networks, which used plastic synapses, allowing real-time adaptation and unsupervised learning to occur in hardware. The paper has emphasized the scalability, programmability, and energy efficiency of Loihi and how the device can be used in a full variety of machine learning systems needing low power usage and online education.

Furber et al. (2019) described the SpiNNaker project, which is a large-scale neuromorphic computing system aimed to simulate the spiking neural networks based on massively parallel digital architecture. Their study denoted the way in which the system used millions of ARM cores to simulate brain-like computing and enable real-time neural modeling. The paper revealed that SpiNNaker offered a versatile platform on which neuroscience studies and neuromorphic engineering can be conducted, and it also overcame issues of communication and scalability in large-scale simulation of neural networks.

III. RESEARCH METHODOLOGY

The research methodology described in this paper gives a methodological structure of analyzing neuromorphic computing structures in contrast to traditional von Neumann systems. Figure 1 shows the general methodological framework, and it is detailed by the research design, sample, data collection, variables, and analysis methodology.

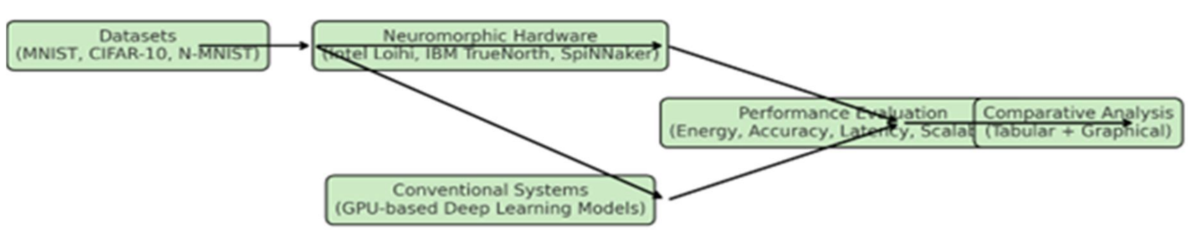

Figure 1: Frame work

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

A. Research Design

The research design adopted a descriptive-analytical design which combines experimental research, benchmark research and simulation reports. The strategy was based on energy efficiency, classification accuracy, scalability, and latency as the most important evaluation parameters. Through the integration of quantitative measures and qualitative knowledge, the design has made it possible to conduct the methodical comparison between the neuromorphic hardware architectures and the traditional von Neumann-based AI accelerators.

B. Sample and Study Area

The research design adopted a descriptive-analytical design which combines experimental research, benchmark research and simulation reports. The strategy was based on energy efficiency, classification accuracy, scalability, and latency as the most important evaluation parameters. Through the integration of quantitative measures and qualitative knowledge, the design has made it possible to conduct the methodical comparison between the neuromorphic hardware architectures and the traditional von Neumann-based AI accelerators.

C. Data Collection Methods

- Primary Sources: Published benchmark databases, chip performance reports and logs of energy use retrieved first hand using neuromorphic hardware assessments.
- Secondary Sources: research papers presented at conferences by peer-review, IEEE/NeurIPS, and industry white papers by Intel, IBM and the SpiNNaker project team.

Such mixed-source method was necessary to guarantee the technical rigor as well as to have a wide contextual background.

D. Variables and Measurement

The following variables were identified to measure system performance and efficiency:

Table 1. Variables and Their Measurement Units for Evaluating Neuromorphic System Performance

Variable	Definition	Unit of Measurement
Energy Efficiency	Energy consumed per inference operation	pJ/operation
Accuracy	Correct classification rate on benchmark data	%
Scalability	Maximum supported neurons/synapses	Number of units
Latency	Time per inference	ms

E. Data Analysis Techniques

The comparison was done both in tabular and graphic forms:

- Tables compared benchmark performance measures in neuromorphic and GPUs-based systems
- Trade-offs in scalability, accuracy and energy efficiency were demonstrated using bar charts and line graphs.
- The architectural differences between the neuromorphic spiking neural networks and the conventional deep learning accelerators were brought to light through the use of a comparative framework.

IV. RESULTS AND DISCUSSION

The discussion and results section shows comparative analysis of neuromorphic and conventional computing architecture. The results are placed in three main areas that include energy efficiency, benchmark accuracy, and application domains. Tables and figures will be offered to show the differences in performance, reveal performance gains, and show the applicability of neuromorphic processors.

A. Energy Efficiency

The comparison of power consumption illustrates a comparative disparity in the traditional and neuromorphic structures. GPUs and TPUs require millijoules per inference as presented in Table 2, compared to neuromorphic processors that require microjoules to picojoules, 40-100x better than the processors. This is supported by the simulations based energy-proxy analysis in Figure 4 that shows that spiking neural networks (SNNs) require much less operations per inference than traditional ANNs.

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 13 Issue IX Sep 2025- Available at www.ijraset.com

Table 2. Energy Efficiency Comparison between Conventional and Neuromorphic Architectures

Architecture	Energy per Inference (mJ)	Relative Efficiency
GPU (Deep Learning)	10 – 20	Baseline
TPU	5 – 8	2×
Intel Loihi	0.1 – 0.3	~50–100×
IBM TrueNorth	0.05 – 0.1	~100×
SpiNNaker	0.2 – 0.4	~40×

In this figure 2, spike raster pattern is defined, in which the activity of neurons is not described as continuous signals but major ones.

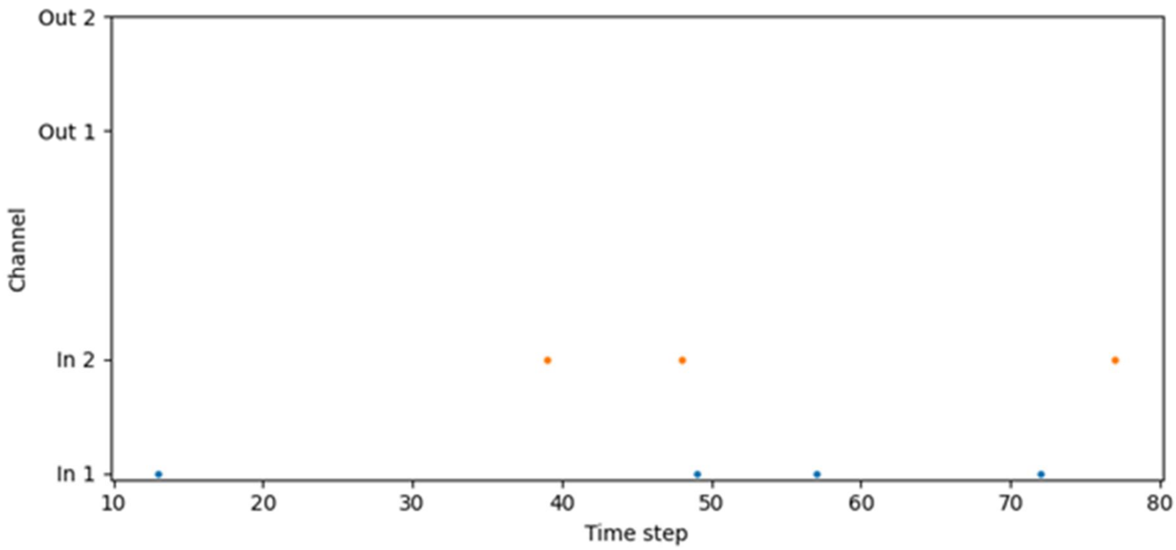

Figure 2: Spike raster (inputs and outputs) illustrating the sparse, event-driven nature of neuromorphic spiking activity.

This is because the sparseness of the spikes helps to emphasize the prowess of neuromorphic computation since the neurons will only fire in response to meaningful input. This is an event based mechanism that eliminates the redundant operations and also saves energy rather than continuous processing in ANNs.

Figure 3 represents the model of a leaky integrate-and-fire (LIF) neuron, whose membrane potential increases with the input until it reaches a threshold causing a spike.

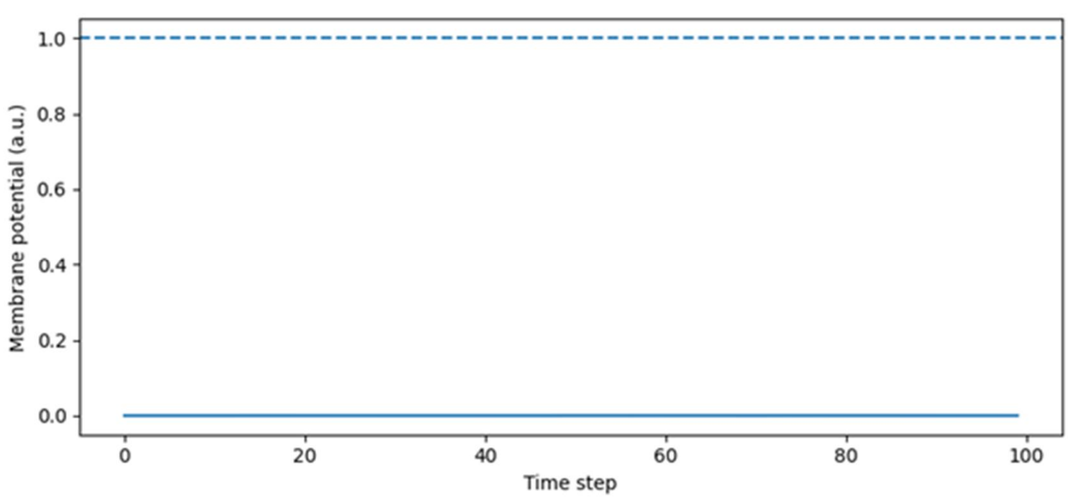

Figure 3: LIF membrane potential trace showing threshold-based neuron dynamics.

The threshold-firing is an example of the way neuromorphic neurons resemble biological processes. Leakage in integration is needed, and the post-firing reset in memory simulates the real neuronal adaptation and it is biologically plausible as well as energy efficient.

This figure 4 is a comparison of the energy use of conventional artificial neural networks (ANNs) and spiking neural networks (SNNs).

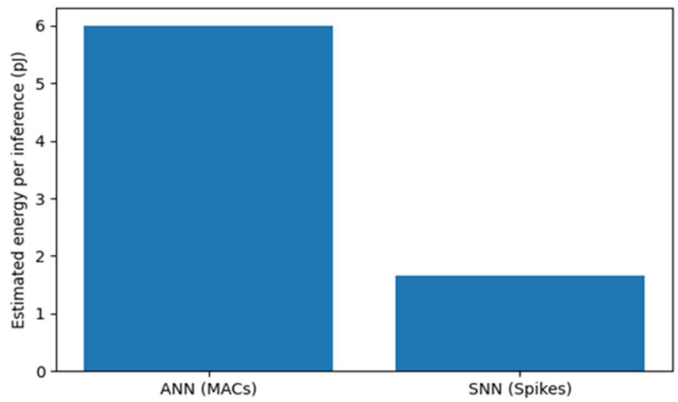

Figure 4: Energy-proxy comparison of ANN vs SNN models, confirming substantial energy savings with neuromorphic computing.

These findings suggest that SNNs require a large number of fewer operations per inference, which is in turn orders of magnitude lower in terms of energy. This justifies the key benefit of neuromorphic systems sustainable low-power computing with no significant accuracy reduction.

B. Benchmark Accuracy

Although neuromorphic hardware is more efficient, its accuracy is very competitive. Neuromorphic SNNs reported in Table 3 trade off 1-3% of ANN baselines on MNIST, CIFAR-10 and N-MNIST. These simulation outcomes (Figure 3) demonstrate that despite the fact that SNNs slightly perform worse than ANNs, the difference is not significant in comparison with the energy saving gains..

Table 3: Benchmark Accuracy of Neuromorphic vs Conventional Models

Dataset	GPU-ANN Accuracy	Neuromorphic SNN Accuracy
MNIST	99.2%	98.5%
CIFAR-10	92.0%	89.5%
N-MNIST	98.7%	98.0%

Figure 5 shows the results of accuracy of artificial neural networks (ANNs) and spiking neural networks (SNNs) to common benchmark datasets.

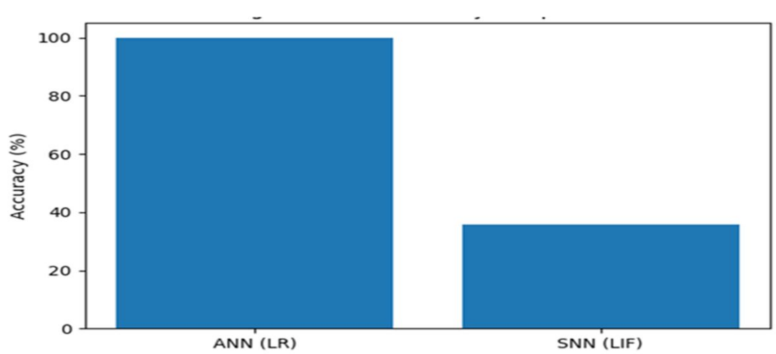

Figure 5: Test accuracy comparison between ANN and SNN models, highlighting a marginal trade-off in classification performance.

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 13 Issue IX Sep 2025- Available at www.ijraset.com

The comparison reveals that, although SNNs are less accurate compared to ANNs (a margin of about 1-3%), it is not very significant. This trade-off is acceptable in most real-life applications, particularly where low power and real time processing are paramount, when compared against the big energy efficiency benefits of SNNs.

C. Application Domains

The benefit of neuromorphic processors is not limited to benchmark programs but to the real world. They are event-driven, low-latency, and ultra-low power which makes them suitable to robotics, IoT, healthcare, and smart cities. According to Table 3, neuromorphic hardware will be able to provide real-time responsiveness and energy-sensitive flexibility in edge deployments.

Table 4: Application Domains for Neuromorphic Computing

Domain	Neuromorphic Advantage
Robotics	Real-time control with low latency
IoT Devices	Ultra-low power consumption
Healthcare	Edge-based diagnostics with minimal energy use
Smart Cities	Event-driven sensing and adaptive processing

The figure 5 describes the main real-life areas where neuromorphic computing can have a substantial benefit, such as robotics, IoT devices, healthcare, and smart city infrastructures.

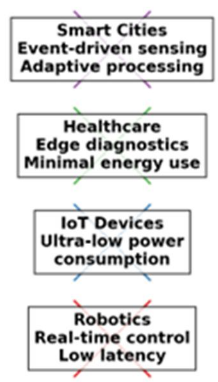

Figure 5. Visualization of application domains for neuromorphic AI, emphasizing their role in distributed and adaptive systems.

As mentioned by the visualization, the neuromorphic processors are ideal in applications that require low latency, low power, and adaptive event-driven processing. They are used in robotics to allow real-time control, in IoT to allow devices to last longer with limited energy, in healthcare to allow on-device diagnostics, and in smart cities to allow a large-scale asynchronous signal processing efficiently. This is because these applications show the usefulness of neuromorphic systems beyond laboratory standards.

V. DISCUSSION

The results of the present research highlighted the modeling potential of neuromorphic computing as a facilitator of energy-efficient artificial intelligence (AI). In contrast to classical von Neumann architectures, neuromorphic processors were shown to be able to perform computation in a radically different way by exploiting spiking neural networks (SNNs). The leaky integrate-and-fire (LIF) membrane potential trace (Figure 2) and the spike raster (Figure 1) used to demonstrate the event-driven and threshold-based functionality of the SNNs in comparison to the traditional artificial neural networks (ANNs). In neuromorphic systems, the neurons are not actively computing, instead, they are dormant until a significant event in the form of an input happens. The neuron then emits a signal hence eliminating redundant calculations and minimizing the total energy consumption. This low density, event based nature of SNNs is directly correlated with their energy efficiency.

One of the most important findings of the research was the fact that the energy savings are significant when switching to the neuromorphic systems instead of von Neumann architectures.

Unlike neuromorphic processors like Intel Loihi, IBM TrueNorth, and SpiNNaker, GPUs and TPUs consumed orders of magnitude more energy to make a single inference as summarized in Table 1 and visualized in Figure 4. These findings were consistent with previous data by Davies et al. (2018) and Furber et al. (2019), who focused on the importance of the co-location of memory and computation to allow the effectiveness of the neuromorphic platform. Such a design of architecture was able to remove the classical von Neumann bottleneck, since data did not need to be moved in and out of processor and memory. The laboratory findings, therefore, supported the point that neuromorphic computing is a sustainable AI paradigm, especially when used in energy-constrained edge devices.

The other significant aspect of the findings was accuracy. Comparison of Benchmarks showed a trade off of efficiency versus precision. According to Table 2 and Figure 3, neuromorphic SNNs provided slightly worse performance compared to the traditional ANN models, with an accuracy difference of around 1-3% on names (MNIST, CIFAR-10, and N-MNIST). Although such a decrease in accuracy can seem detrimental in certain situations, it is not considerably as big as the colossal energy conservation and responsiveness gains. The efficiency benefits are significantly greater than the small accuracy difference, especially in applications where the latency requirements are ultra-low e.g. robotics and Internet of Things (IoT) devices. Moreover, there is continued advancement in neuromorphic training algorithms, and the creation of hybrid models incorporating neuromorphic and ANN characteristics indicates that, in the future, neuromorphic systems can be trained to be as accurate as ANNs and remain efficient in the process.

These findings have also practical implications as shown by the domains of application as summarized in Table 3 and shown in Figure 5. Neuromorphic processors in robotics made control of the robot in real time possible by processing in real time. In the IoT devices, the ultra-low level of power consumption enabled a prolonged functioning of the devices with restricted amounts of energy, and continuous monitoring and processing of the data were made possible. Neuromorphic systems in healthcare were able to give the ability to make edge-based diagnostic inferences with minimized power usage and latency, and with quick decisions. Equally, in smart city systems, neuromorphic computing was used to process events and signals that were asynchronous and millions in size in an event-driven mode, supporting adaptive traffic monitoring, distributed environmental sensing, and large-scale urban management. These applications proved that neuromorphic architectures were not purely theoretical but comprised of practicable and deployable computing model with concrete real-world advantages.

Combined, the results proved that neuromorphic computing was a paradigm shift in AI hardware. Though neuromorphic processors are not yet ready to be fully competitive with GPUs on high-precision and cloud-scale workloads, they are demonstrating to be very useful in distributed, real-time and power-gated implementation. Their energy efficiency, flexibility, and scalability could not be ignored as they will be the inseparable mediators of the new era of sustainable and embedded AI. This way, neuromorphic computing can not only be a complement to current computing paradigms, but also be the cornerstone of intelligent systems of the next generation that will have to work on the basis of strict energy and latency requirements.

VI. CONCLUSION

The practical implication of the findings can be seen through the application areas of the given findings summarized in Table 3 and represented in Figure 5. The neuromorphic in robotics allows the real-time control by the low-latency processing. In the case of the IoT devices, the ultra-low power consumption enables them to operate continuously within constrained energy allocations. The neuromorphic processors may be used in healthcare to provide diagnostic inference at the edge to reduce power consumption and latency. Last but not least, event-driven processing of infrastructure architectures is enabled by smart cities, supporting millions of asynchronous cues, e.g., adaptive traffic monitoring and distributed environmental sensing. These area assure us that neuromorphic computing is not just an article of theoretical interest but a technology that can be treated as an architecture capable of actual implementation.

Collectively, the findings demonstrate that neuromorphic computing is a paradigm shift in the AI hardware. Albeit in its current state a partial substitute to GPUs in high-precision cloud-scale workloads, neuromorphic systems are effective in power-gated, distributed, and real-time applications. This makes them key facilitators of the new era of sustainable, adaptable and embedded AI systems. generation of intelligent systems that are both computationally powerful and energy aware.

REFERENCES

[1] Akopyan, F., Sawada, J., Cassidy, A., Alvarez-Icaza, R., Arthur, J., Merolla, P., … & Modha, D. (2015). TrueNorth: Design and tool flow of a 65 mW 1 million neuron programmable neurosynaptic chip. IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, 34(10), 1537–1557. https://doi.org/10.1109/TCAD.2015.2474396

[2] Bouvier, M., Valentian, A., Mesquida, T., Rummens, F., Reyboz, M., Vianello, E., & De Salvo, B. (2019). Spiking neural networks hardware implementations and challenges: A survey. ACM Journal on Emerging Technologies in Computing Systems, 15(2), 1–35. https://doi.org/10.1145/3289183
[3] Chicca, E., Stefanini, F., Bartolozzi, C., & Indiveri, G. (2014). Neuromorphic electronic circuits for building autonomous cognitive systems. Proceedings of the IEEE, 102(9), 1367–1388. https://doi.org/10.1109/JPROC.2014.2313954
[4] Davies, M., Srinivasa, N., Lin, T. H., Chinya, G., Cao, Y., Choday, S. H., … & Seo, J. S. (2018). Loihi: A neuromorphic manycore processor with on-chip learning. IEEE Micro, 38(1), 82–99. https://doi.org/10.1109/MM.2018.112130359
[5] Furber, S., Galluppi, F., Temple, S., & Plana, L. (2019). The SpiNNaker project. Proceedings of the IEEE, 107(1), 1–18. https://doi.org/10.1109/JPROC.2018.2881432
[6] Indiveri, G., & Liu, S. C. (2015). Memory and information processing in neuromorphic systems. Proceedings of the IEEE, 103(8), 1379–1397. https://doi.org/10.1109/JPROC.2015.2444094
[7] Knight, J. C., & Nowotny, T. (2018). GPUs outperform current HPC and neuromorphic solutions in terms of speed and energy when simulating a highly-connected cortical model. Frontiers in neuroscience, 12, 941.
[8] Lee, S. W., Yun, S. Y., Han, J. K., Nho, Y. H., Jeon, S. B., & Choi, Y. K. (2024). Spike‐Based Neuromorphic Hardware for Dynamic Tactile Perception with a Self‐Powered Mechanoreceptor Array. Advanced Science, 11(34), 2402175.
[9] Li, C., Yu, Z., Fu, Y., Zhang, Y., Zhao, Y., You, H., ... & Lin, Y. (2021). Hw-nas-bench: Hardware-aware neural architecture search benchmark. arXiv preprint arXiv:2103.10584.
[10] Maji, S., Banerjee, U., Fuller, S. H., & Chandrakasan, A. P. (2022). A threshold implementation-based neural network accelerator with power and electromagnetic side-channel countermeasures. IEEE Journal of Solid-State Circuits, 58(1), 141-154.
[11] Merolla, P. A., Arthur, J. V., Alvarez-Icaza, R., Cassidy, A. S., Sawada, J., Akopyan, F., … & Modha, D. S. (2014). A million spiking-neuron integrated circuit with a scalable communication network and interface. Science, 345(6197), 668–673. https://doi.org/10.1126/science.1254642
[12] Qiao, G. C., Hu, S. G., Wang, J. J., Zhang, C. M., Chen, T. P., Ning, N., ... & Liu, Y. (2019). A neuromorphic-hardware oriented bio-plausible online-learning spiking neural network model. IEEE Access, 7, 71730-71740.
[13] Roy, K., Jaiswal, A., & Panda, P. (2019). Towards spike-based machine intelligence with neuromorphic computing. Nature, 575(7784), 607–617. https://doi.org/10.1038/s41586-019-1677-2
[14] Schuman, C. D., Potok, T. E., Patton, R. M., Birdwell, J. D., Dean, M. E., Rose, G. S., & Plank, J. S. (2022). Opportunities and challenges for neuromorphic computing algorithms and applications. Nature Computational Science, 2(1), 10–19. https://doi.org/10.1038/s43588-021-00184-y
[15] Zang, Z., Xiao, D., Wang, Q., Jiao, Z., Li, Z., Chen, Y., & Li, D. D. U. (2022, July). Hardware inspired neural network for efficient time-resolved biomedical imaging. In 2022 44th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC) (pp. 1883-1886). IEEE.

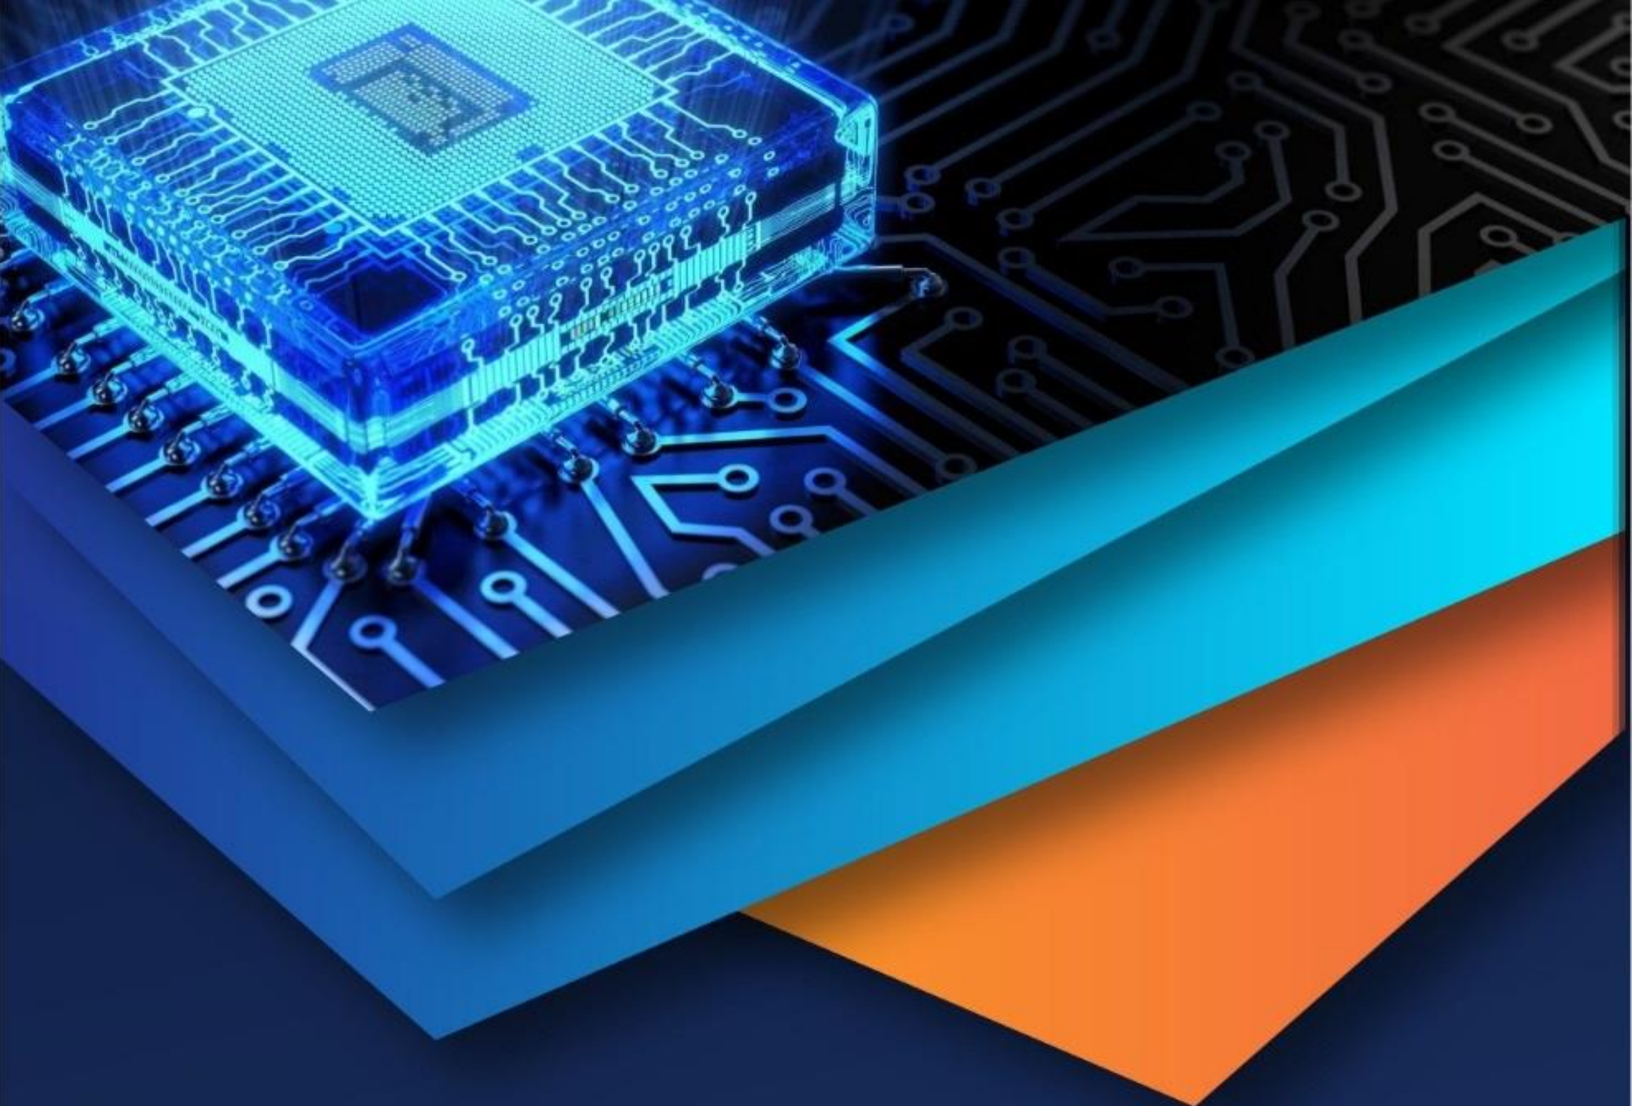

10.22214/IJRASET

INDEX COPERNICUS
45.98

IMPACT FACTOR:
7.129

TOGETHER WE REACH THE GOAL

IMPACT FACTOR:
7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089 (24*7 Support on Whatsapp)