



OPEN ACCESS

EDITED BY

Fakhreddine Zayer,
Khalifa University, United Arab Emirates

REVIEWED BY

Fangwen Yu,
Tsinghua University, China
Thi-Thu-Huong Le,
Pusan National University, Republic of Korea

*CORRESPONDENCE

Rachmad Vidya Wicaksana Putra,
✉ rachmad.putra@nyu.edu

RECEIVED 15 March 2024

ACCEPTED 14 June 2024

PUBLISHED 26 July 2024

CITATION

Putra RVW, Marchisio A and Shafique M (2024), SNN4Agents: a framework for developing energy-efficient embodied spiking neural networks for autonomous agents. *Front. Robot. AI* 11:1401677. doi: 10.3389/frobt.2024.1401677

COPYRIGHT

© 2024 Putra, Marchisio and Shafique. This is an open-access article distributed under the terms of the [Creative Commons Attribution License \(CC BY\)](https://creativecommons.org/licenses/by/4.0/). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

SNN4Agents: a framework for developing energy-efficient embodied spiking neural networks for autonomous agents

Rachmad Vidya Wicaksana Putra*, Alberto Marchisio and Muhammad Shafique

eBrain Lab, Division of Engineering, New York University (NYU) Abu Dhabi, Abu Dhabi, United Arab Emirates

Recent trends have shown that autonomous agents, such as Autonomous Ground Vehicles (AGVs), Unmanned Aerial Vehicles (UAVs), and mobile robots, effectively improve human productivity in solving diverse tasks. However, since these agents are typically powered by portable batteries, they require extremely low power/energy consumption to operate in a long lifespan. To solve this challenge, neuromorphic computing has emerged as a promising solution, where bio-inspired Spiking Neural Networks (SNNs) use spikes from event-based cameras or data conversion pre-processing to perform sparse computations efficiently. However, the studies of SNN deployments for autonomous agents are still at an early stage. Hence, the optimization stages for enabling efficient embodied SNN deployments for autonomous agents have not been defined systematically. Toward this, we propose a novel framework called SNN4Agents that consists of a set of optimization techniques for designing energy-efficient embodied SNNs targeting autonomous agent applications. Our SNN4Agents employs weight quantization, timestep reduction, and attention window reduction to jointly improve the energy efficiency, reduce the memory footprint, optimize the processing latency, while maintaining high accuracy. In the evaluation, we investigate use cases of event-based car recognition, and explore the trade-offs among accuracy, latency, memory, and energy consumption. The experimental results show that our proposed framework can maintain high accuracy (i.e., 84.12% accuracy) with 68.75% memory saving, 3.58x speed-up, and 4.03x energy efficiency improvement as compared to the state-of-the-art work for the NCARS dataset. In this manner, our SNN4Agents framework paves the way toward enabling energy-efficient embodied SNN deployments for autonomous agents.

KEYWORDS

neuromorphic computing, spiking neural networks, autonomous agents, automotive data, neuromorphic processor, energy efficiency

1 Introduction

In recent years, the interest in implementing neuromorphic artificial intelligence based on Spiking Neural Networks (SNNs) for autonomous agents (so-called *SNN-based autonomous agents*) has rapidly increased (Bartolozzi et al., 2022; Putra et al., 2024). The reason is that, SNNs can offer high accuracy due to effective learning mechanism

(Putra and Shafique, 2021b; Putra and Shafique, 2024; Rathi et al., 2023), low computation latency due to efficient neural/spike coding (Guo et al., 2021), and ultra low power/energy consumption due to sparse spike-based operations (Putra and Shafique, 2020; Schuman et al., 2022). To realize such systems in real life, capabilities of solving machine learning (ML) tasks like image classification (Putra et al., 2021b; 2022a; 2023), object detection (Viale et al., 2021; Cordone et al., 2022), or object segmentation (Li et al., 2022) from images/videos are required. Besides such functionalities, SNN-based autonomous agents also require (1) small memory footprint as they typically employ resource-constrained hardware platforms, (2) low power/energy consumption to preserve the battery lifespan as they are typically powered by portable batteries, and (3) real-time output with high accuracy to provide quick decision (Bonnievie et al., 2021; Putra and Shafique, 2022; 2023b). To maximize the benefits of SNN sparse operations, event-based data can be employed as it directly provides a compatible data format for SNN processing and minimizes the pre-processing stage, such as the data-to-spike conversion (e.g., pixel data to spike train) and the spike coding. Therefore, the developments of SNN-based autonomous agents also need to consider event-based data, such as the NCARS dataset (Sironi et al., 2018; Bano et al., 2024).

Motivated by the above-mentioned potentials of SNN-based autonomous agents, *the targeted research problem is how to systematically develop energy-efficient SNN-based autonomous agents considering event-based data workload.* An efficient solution to this problem will enable SNN-based autonomous agents to achieve high accuracy with small memory footprint, low processing latency, and low energy consumption.

1.1 State-of-the-art works and their limitations

The study for developing SNN-based autonomous agents is still at an early stage. The state-of-the-art works are summarized in Table 1. Here, we observe that most of the existing works focus on proposing frameworks and/or techniques for achieving high accuracy (Putra and Shafique, 2022; 2023a; b). However, these works have not considered event-based data workloads, therefore requiring a relatively complex pre-processing stage, including data-to-spike conversion and spike coding. Some other works explore techniques for achieving high accuracy considering event-based automotive data (Viale et al., 2021; 2022; Cordone et al., 2022). However, these works have not considered optimizing the model size to fit into the resource-constrained autonomous agents. All the above-discussed limitations of the state-of-the-art expose that, *the optimization stages for enabling efficient SNN deployments for autonomous agents have not been defined systematically.*

1.2 Motivational case study

To highlight the potentials of further optimizing SNNs for autonomous agents from the current state-of-the-art, we perform an experimental case study that considers applying different levels of weight quantization (i.e., precision) to an SNN model from the work of Viale et al. (2021) with event-based NCARS dataset.

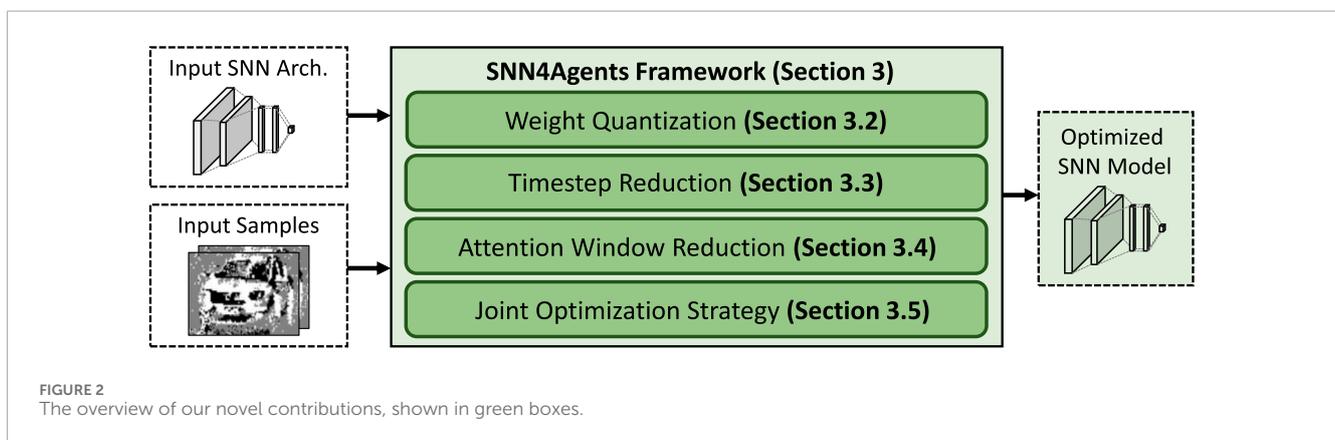
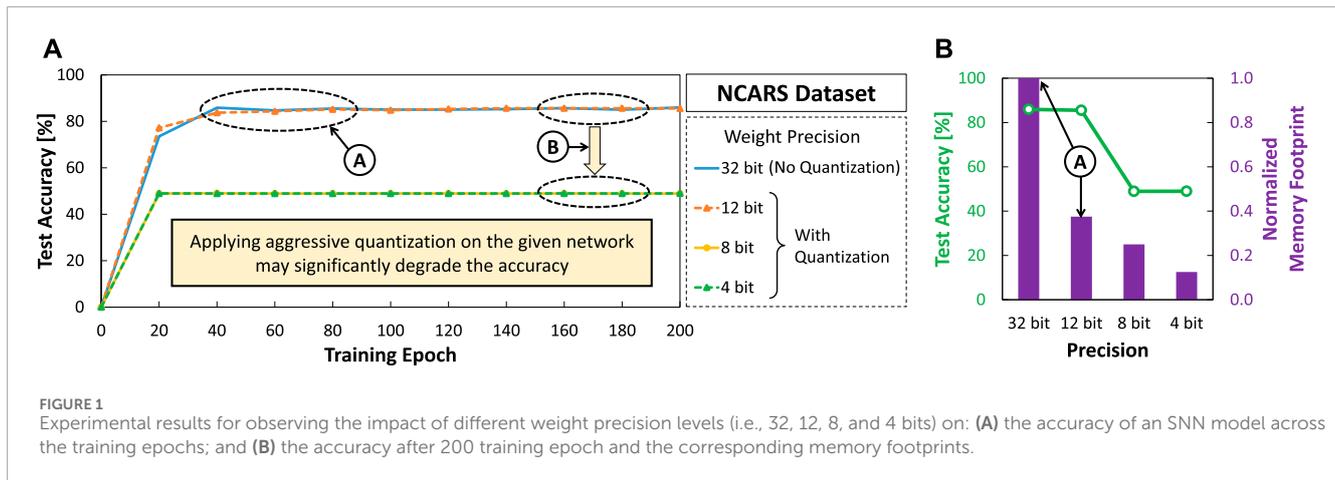
TABLE 1 State-of-the-art for SNN-based autonomous agents.

Work	Key attributes
Mantis (Putra and Shafique, 2023a)	+ It presents a set of techniques for optimizing SNNs for autonomous agents – It considers conventional non-event data, thus its efficiency is sub-optimal
lpSpikeCon (Putra and Shafique, 2022)	+ It enables low-precision SNNs under unsupervised continual learning settings – It considers conventional non-event data, thus its efficiency is sub-optimal
TopSpark (Putra and Shafique, 2023b)	+ It optimizes the computation time for both training and inference – It considers conventional non-event data, thus its efficiency is sub-optimal
CarSNN (Viale et al., 2021)	+ It employs the STBP (Wu et al., 2018) for learning event-based car data – It does not consider optimization for resource-constrained autonomous agents
LaneSNN (Viale et al., 2022)	+ It employs the STBP (Wu et al., 2018) for learning event-based road lane data. – It does not consider optimization for resource-constrained autonomous agents.
Object Detection SNN (Cordone et al., 2022)	+ It employs voxel cube and surrogate gradient for learning event-based data – It does not consider optimization for resource-constrained autonomous agents
SNN4Agents (ours)	+ It presents a set of techniques for optimizing SNNs for autonomous agents + It considers event-based data to enable efficient input activation

Details of the experimental setup will be discussed further in Section 4. The experimental results shown in Figure 1 show several key observations as discussed in the following.

- Ⓐ Quantization with 12-bit precision level achieves comparable accuracy to the original 32-bit precision level (without quantization) across training epochs, while offering 2.7x memory saving.
- Ⓑ Quantization settings with 8-bit and 4-bit precision levels suffer from significant accuracy degradation as they roughly have 50% accuracy across training epochs. Considering that the NCARS dataset has two classes (i.e., “car” or “background”), these results indicate that the network performs random guessing as it is not properly trained.

These observations show that there is an opportunity to optimize further the SNN models to make them fit into the resource-constrained autonomous agents. However, simply performing aggressive quantization on the given network may significantly degrade the accuracy. Therefore, *the research challenge is how to effectively perform different optimization techniques without significantly degrading the accuracy.*



1.3 Our novel contributions

To address the targeted research problem and the related challenges, we propose a novel framework called *SNN4Agents* for developing energy-efficient embodied SNNs for autonomous agents. The overview of our SNN4Agents framework is shown in Figure 2 and its key steps are briefly described in the following.

- **Weight Quantization (Section 3.2):** It aims to find the appropriate quantization settings that meet the memory constraint, we perform design space exploration for different precision levels while observing their impact on the accuracy.
- **Timestep Reduction (Section 3.3):** It aims to find the appropriate processing timesteps that meet the latency constraint, we perform design space exploration for different timestep values while observing their impact on the accuracy.
- **Attention Window Reduction (Section 3.4):** It aims to reduce the compute requirements and hence the processing power/energy, we explore different attention window sizes from the input samples and observe their impact on the accuracy.
- **Joint Optimization Strategy (Section 3.5):** It aims to maximize the benefits from the previous individual optimization steps, we perform a joint optimization strategy by leveraging design space exploration that

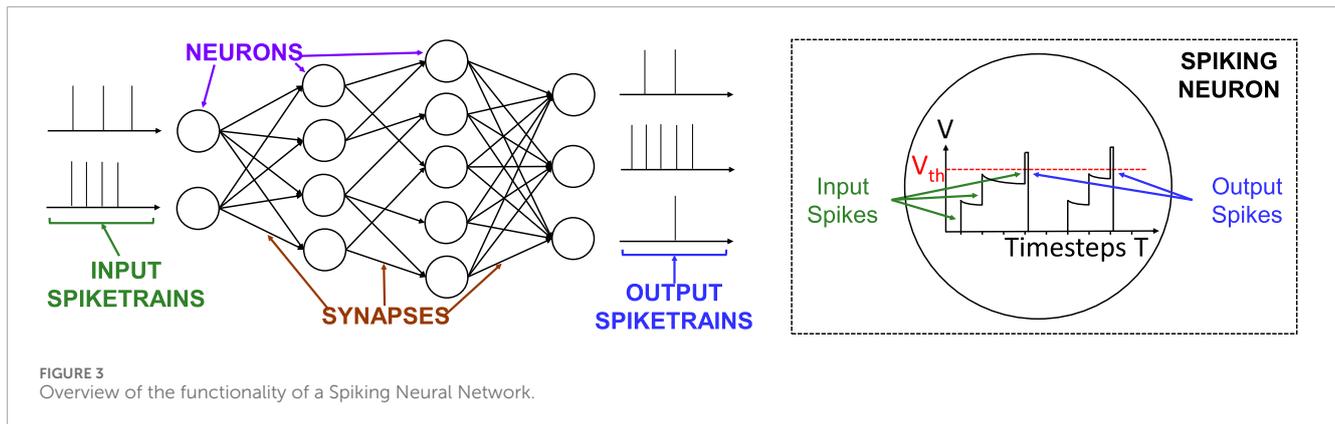
considers the appropriate settings from individual optimization steps.

2 Preliminaries

2.1 Spiking neural networks (SNNs)

2.1.1 Overview

Spiking Neural Networks (SNNs) are considered the third generation neural networks (Maass, 1997). An overview of the SNNs' functionality is shown in Figure 3. The input spike trains are processed by the spiking neurons and the information is propagated to the neurons in the following layers. Among the most popular spiking neuron models, the Leaky-Integrate-and-Fire (LIF) (Izhikevich, 2004; Wang et al., 2014) emerges as an efficient trade-off between complexity and plausibility. The operational time of a neuron to process a spike train from a single input data (e.g., an image pixel) is defined as *timestep* (Putra and Shafique, 2023b). At each timestep, the membrane potential V_m of the neuron increases for each incoming input spike. When V_m overcomes the voltage threshold V_{th} , the neuron emits an output spike (Roy et al., 2019; Putra et al., 2021a; 2022b). In this way, the spikes propagate from input to output and multiple layers of spiking neurons in a network.



While SNNs follow the wave of success of traditional (non-spiking) Deep Neural Networks (DNNs), they offer the following advantages.

- **Biological Plausibility:** The SNNs' functionality is inspired by the behavior of the biological brain, where spikes are propagated across neurons for conveying information. This may open possibilities of cognition and robustness for solving diverse machine learning tasks.
- **Ultra-Low Power/Energy Consumption:** The dynamic power in SNNs is consumed only in the presence of spikes, hence offering ultra low processing power/energy in a long time operational period.
- **Efficient Interface with Event-based Sensors:** The event sequences captured by event-based sensors can directly be utilized as the input of SNNs without complex pre-processing (e.g., data-to-spike conversion).

Besides these advantages, it is actually challenging to efficiently train SNNs due to the non-differentiability of the spiking loss function (Rückauer et al., 2019). Hence, to overcome this challenge, two possible techniques have been proposed in the literature. (1) The *DNN-to-SNN conversion* approach (Bu et al., 2022; Hao et al., 2023) trains a DNN and then converts the model into the equivalent spiking counterpart. (2) The *direct SNN training* approach (Nefci et al., 2019) employs a surrogate gradient function to approximate the spiking loss function, in such a way that it can be differentiated and incorporated into the backpropagation flow. Since approach-(1) requires DNN training, it cannot be directly used when dealing with event-based data (Massa et al., 2020). Moreover, it typically requires a larger number of timesteps than approach-(2) (Chowdhury et al., 2021). Therefore, in this work, we consider employing direct SNN training, i.e., *Spatio-Temporal Back-Propagation (STBP) learning rule*, which leverages both spatial and temporal information within the streaming spikes (Wu et al., 2018; Viale et al., 2021).

2.2 Quantization

Quantization is a prominent optimization technique which can effectively compress SNN models with relatively low overhead, since it only needs to reduce the data precision

(Gupta et al., 2015; Micikevicius et al., 2018). Implementation of quantization requires a specific setting that encompasses *quantization scheme* and *rounding scheme* as discussed in the following.

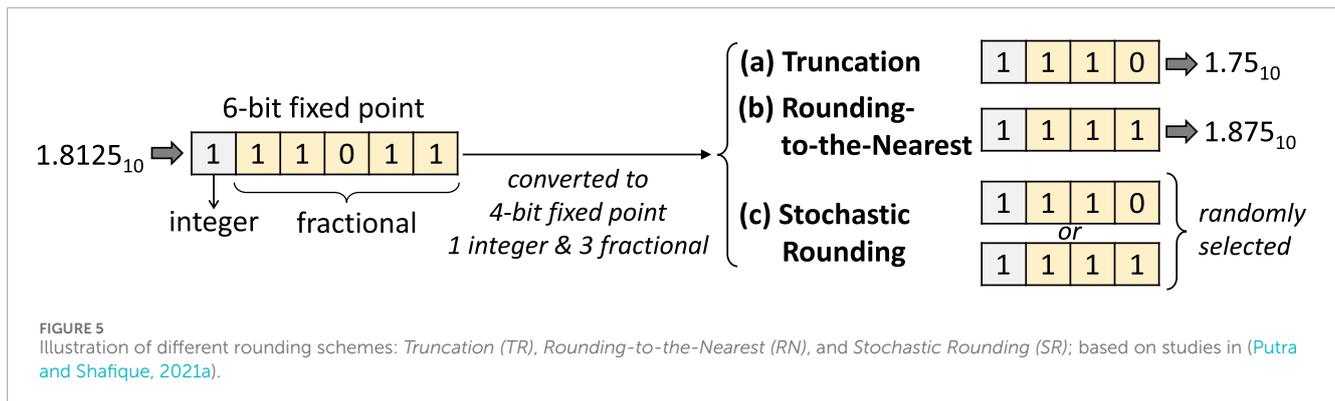
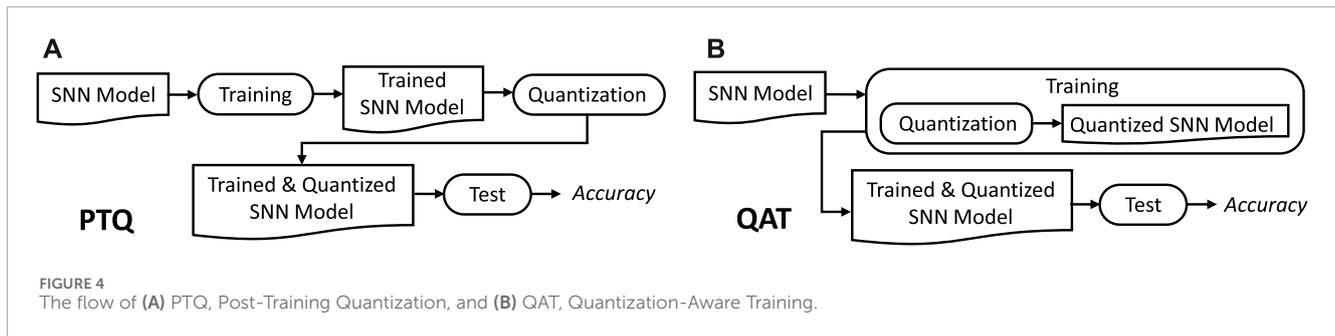
Quantization Schemes: There are two widely used quantization schemes for SNNs: *Post-Training Quantization (PTQ)*, and *Quantization-Aware Training (QAT)* (Putra and Shafique, 2021a); see the illustration of PTQ and QAT flows in Figure 4. PTQ typically trains the given network with a floating-point precision, such as 32-bit floating point (FP32), and then the quantization is applied to the trained SNN model with the given precision level, resulting in a quantized SNN model for the inference phase. Meanwhile, QAT typically performs quantization to the given network with the given precision level during the training phase, resulting in a trained and quantized SNN model which can be directly used for the inference phase. PTQ process is typically simpler and more efficient than QAT as it performs quantization once after the training phase is finished.

Rounding Schemes: The implementation of quantization typically requires a specific rounding scheme for determining how the value will be curtailed. From the literature, there are three widely used rounding schemes for SNN models: *Truncation (TR)*, *Rounding-to-the-Nearest (RN)*, and *Stochastic Rounding (SR)* (Putra and Shafique, 2021a). Illustration of these rounding schemes is shown in Figure 5. Among these rounding schemes, TR has the simplest operation since it simply keeps the defined number of the most significant bits and discards the other remaining bits.

In this work, we employ PTQ scheme with TR rounding since their combination can quickly provide representative results for different quantization settings, thereby enabling fast design space exploration (DSE), which is important for our studies.

2.3 Event-based automotive data

Prophesee NCARS Dataset (Sironi et al., 2018): This event-based dataset contains a collection of 24K samples that have a duration of 100 ms each, recorded with the Asynchronous Time-based Image Sensor (ATIS) camera (Posch et al., 2011). Each sample, labeled as either "car" or "background," is encoded as a sequence of events that contains the following information.



- the timestamp t of when the event occurred;
- the spatial coordinates x and y of the pixel;
- the polarity p of the brightness variation, which can either be positive or negative.

The data is split into 15,422 training and 8,607 testing samples. Each sample has variable sizes and can be cropped. Based on the distribution of events, we can identify an *attention window*, i.e., a region in which the events are more concentrated. Typical sizes of the attention windows used in state-of-the-art works (Viale et al., 2021) can scale down to 100×100 or 50×50 and significantly reduce the computational and memory requirements of the SNN, without compromising the accuracy much.

2.4 SNN architecture

SNN architectures are composed of a sequence of layers of spiking neurons. Their structure and connections define the type of layers. In this work, we employ the CarSNN architectures (Viale et al., 2021) that efficiently execute car recognition with STBP-based learning rule. We implement two SNN architectures built with different sizes of attention window. Both models are composed of an interleaved sequence of two convolutional layers and three average pooling layers, followed by two fully-connected layers. The first SNN architecture, described in Table 2, has a 100×100 attention window, while the second SNN architecture, described in Table 3, has a 50×50 attention window. Note that different attention window sizes affect the feature map sizes of each SNN layer. Hence, these architectures have different numbers of input and output channels in the fully-connected layers than the latter.

3 Our SNN4Agents framework

3.1 Overview

Our SNN4Agents framework employs a set of optimization techniques targeting different design aspects, including *model compression through weight quantization*, *latency optimization through timestep reduction*, *input data optimization through attention window reduction*; and then performs *joint optimization strategy* to maximize benefits from the individual optimization techniques. The flow of our SNN4Agents framework is illustrated in Figure 6 and the detailed discussion for its steps are provided in Section 3.2—Section 3.5.

3.2 Model compression through quantization

To effectively compress the model size, we perform weight quantization through PTQ with TR rounding scheme. To do this, we first train the given network without quantization, while employing baseline settings for time step, attention window, and training epoch. For this scenario, we employ 32-bit precision, 20 time step, 100×100 attention window, and 200 training epoch. Once the training phase is finished, we perform quantization process to the trained network. Afterward, we perform DSE under different weight precision levels (i.e., 32, 16, 12, 10, 8, 6, and 4 bit) to evaluate their impact on the accuracy; see the parameter settings for DSE in Table 4.

Experimental results of DSE are shown in Figure 7, from which we draw the following key observations.

TABLE 2 SNN architecture with 100 × 100 attention window, based on the work of Viale et al. (2021).

Layer type	Input channel	Output channel	Kernel size	Padding	Stride
Average Pooling	2	2	4	-	-
Convolution	2	32	3	1	1
Average Pooling	32	32	2	-	-
Convolution	32	32	3	1	1
Average Pooling	32	32	2	-	-
Fully-Connected	1152	512	-	-	-
Fully-Connected	512	2	-	-	-

TABLE 3 SNN architecture with 50 × 50 attention window, based on the work of Viale et al. (2021).

Layer type	Input channel	Output channel	Kernel size	Padding	Stride
Average Pooling	2	2	4	-	-
Convolution	2	32	3	1	1
Average Pooling	32	32	2	-	-
Convolution	32	32	3	1	1
Average Pooling	32	32	2	-	-
Fully-Connected	288	144	-	-	-
Fully-Connected	144	2	-	-	-

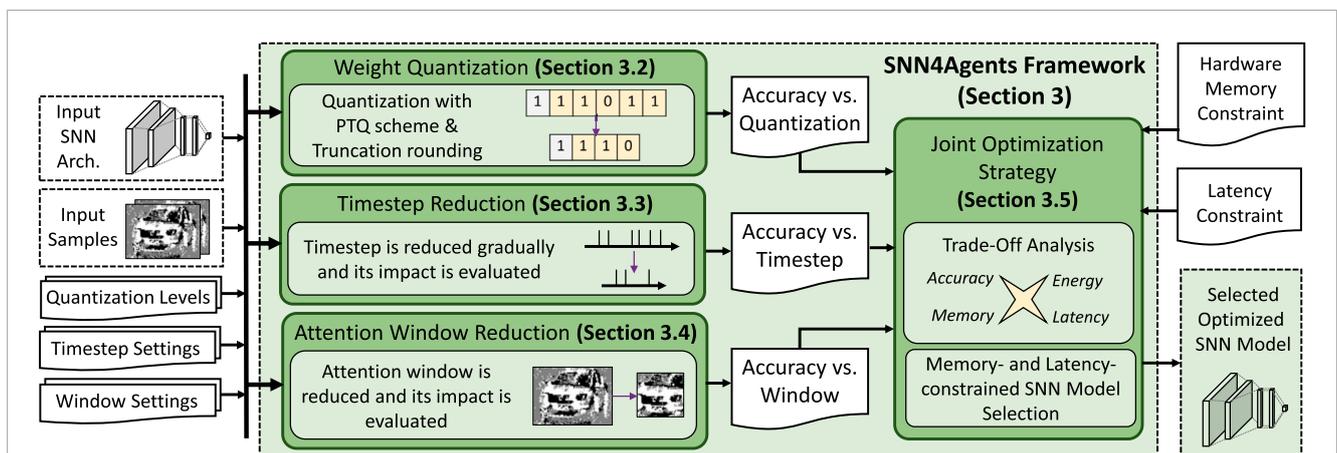


FIGURE 6 The flow of our SNN4Agents framework, with the technical contributions highlighted in green.

- In the early of training phase (e.g., ≤60 training epoch), the network is still learning new information, hence the accuracy curve is increasing for 16-, 12-, and 10-bit precision levels, as shown by ①.
- Employing 16-, 12-, and 10-bit precision levels for SNN weights lead to comparable accuracy to the original SNN model with 32-bit precision (no quantization) after running at least 80 training epoch, as shown by ②.
- Employing 8-, 6-, and 4-bit precision levels for SNN weights lead to significant accuracy degradation, as they can only reach about 50% accuracy across training epochs, as shown by ③. These results indicate that the network is not properly trained.

TABLE 4 Parameter settings for exploring the impact of different precision levels.

Precision (bit)	Timestep	Attention window	Training epoch
32, 16, 12, 10, 8, 6, and 4	20	100 × 100	0–200

These observations expose several key design guides that we should consider when applying quantization. First, *selecting the precision level should be performed carefully*, so that it does not lead to a significant accuracy degradation which diminishes the benefits of quantization. Second, *a 10-bit precision level offers a good trade-off between accuracy and memory footprint* as it can achieve comparable accuracy to that of the larger precision levels after running at least 80 training epoch.

3.3 Latency optimization through timestep reduction

To effectively reduce the processing time, we perform timestep reduction. To do this, we shorten the timestep of SNN processing from the baseline settings, thereby curtailing the time window for presenting the spike trains. Here, we consider different timestep settings (i.e., 20, 15, 10, and 5) for exploring their impact on the accuracy. Once we reduce the timestep, the network is trained under baseline settings of precision level (no quantization), attention window, and training epoch. For this scenario, we employ 32-bit precision, 100 × 100 attention window, and 200 training epoch. If we do not see notable differences from the baseline accuracy, then we may employ a smaller precision level, such as 16- and 4-bit weights, to trigger accuracy variation. The parameter settings for DSE are provided in Table 5.

Experimental results of DSE are shown in Figure 8, from which we draw the following key observations.

- Employing 15, 10, and 5 timestep settings without quantization lead to comparable accuracy to the baseline model (i.e., 20 timestep without quantization) after 60 training epoch, as shown by ④. Similarly, employing 20, 15, 10, and 5 timestep settings with 16-bit precision also lead to comparable accuracy to the baseline after 60 training epoch; see ④.
- Employing 20, 15, 10 and 5 timestep settings with 4-bit precision level lead to significant accuracy degradation, as shown by ⑤. It means that the 4-bit precision is relatively too small for representing temporal and spatial features from the NCARS dataset.

These observations show that *we can apply a relatively aggressive timestep reduction (e.g., 5 timestep) without losing significant accuracy, as long as we also employ an appropriate precision level*, thereby accommodating the spatial and temporal information of the given dataset (e.g., NCARS).

3.4 Attention window reduction for input samples

We also aim at reducing the size of attention window of the input samples to optimize the computational requirements, and hence the latency and energy consumption. To do this, we consider different attention window sizes (i.e., 100 × 100 and 50 × 50), while employing baseline settings for precision level, timestep, and training epoch; see the parameter settings for DSE in Table 6.

Experimental results of our DSE are presented in Figure 9. From these results, we observe that accuracy obtained by a smaller attention window (50 × 50) can saturate faster than a larger one (100 × 100) thus offering a faster training time; see ⑥. Meanwhile, a larger attention window offers better accuracy than a smaller one; see ⑦. The reason is that, a smaller attention window provides less information, hence requiring a shorter time for training the network yet limiting the accuracy that can be achieved. These observations suggest that *the reduction of attention window can be employed as long as the targeted accuracy is met*.

3.5 Joint optimization strategy

Each individual optimization step from previous sub-sections has demonstrated the possibility to reduce memory footprint and latency, while maintaining high accuracy. Therefore, *to maximize these optimization benefits, we propose a strategy to jointly combine the individual optimization steps*. Here, we leverage the key observations and design guides from previous analysis in Section 3.2 and Section 3.4 to devise the following strategy.

- We perform DSE for the following settings: (1) 10–32 bit precision levels of quantization, (2) 5–20 timesteps, and (3) 50 × 50 and 100 × 100 attention window sizes.
- To find the solution candidates, we analyze the experimental results for accuracy, memory footprint, latency, and energy consumption, while considering the memory and latency constraints.
- If there are multiple solution candidates, we can select the most suitable one for the given constraints by trading-off the design metrics, including accuracy, memory footprint, and latency, and energy consumption.

4 Evaluation methodology

To evaluate our SNN4Agents framework, we build and employ the experimental setup shown in Figure 10, while considering the same evaluation conditions as widely used in the SNN community for autonomous agents (Viale et al., 2021; Putra and Shafique, 2023a). We use Python-based implementation and run it on Nvidia RTX 6000 Ada GPU machines for evaluating different performance metrics of our SNN4Agents, including accuracy, processing time, and power consumption. Then, the processing time and power consumption are leveraged to estimate the energy consumption. Meanwhile, memory footprint is estimated by leveraging the precision level and the number of weights in the corresponding SNN architecture. For the case of 100 × 100 attention window, we

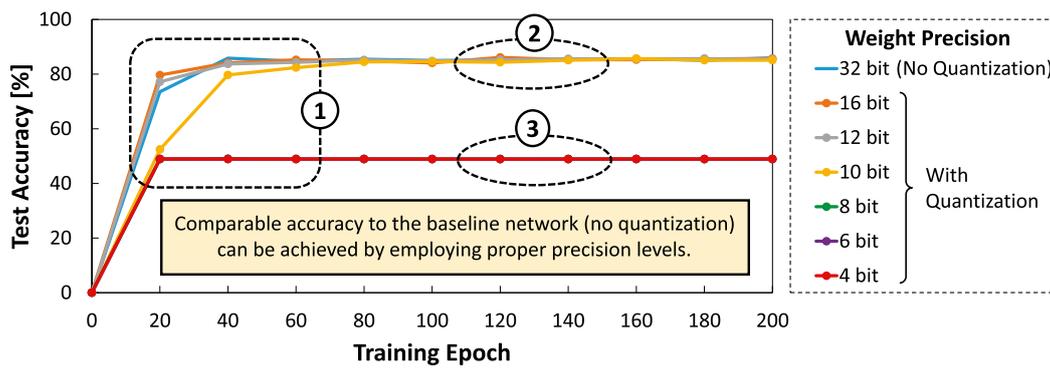


FIGURE 7 Results of accuracy across different precision levels (i.e., 32, 16, 12, 10, 8, and 4 bit).

TABLE 5 Parameter settings for exploring the impact of different timesteps.

Precision (bit)	Timestep	Attention window	Training epoch
32, 16, and 4	20, 15, 10, and 5	100 × 100	0–200

consider the network architecture from Table 2; while for the case of 50×50 , we consider the network architecture from Table 3. For the workload, we consider the NCARS dataset (Sironi et al., 2018). For the DSE purpose, we consider parameter settings shown in Table 7, as well as the state-of-the-art work CarSNN (Viale et al., 2021) as the baseline.

Furthermore, we use several terms to represent the network model for brevity, as the following.

- Bb_Tt : It represents a network with B bit-precision and T timestep. In this case, the attention window will be explained and discussed explicitly to clearly distinguish the network model and its experimental results. This term will be used in Section 5.1, Section 5.3, and Section 5.4.
- Bb_Ww : It represents a network with B bit-precision and $W \times W$ attention window, which encompasses different timestep settings. This term is used when discussing memory footprint in Section 5.2, as networks with different timesteps have the same weight memory footprint.
- Bb_Tt_Ww : It represents a network with B bit-precision, T timestep, and $W \times W$ attention window. This term is used when discussing the trade-off analysis in Section 5.5.
- Tt_Ww : It represents a network with T timestep and $W \times W$ attention window. This term is used when discussing the computational complexity in Section 5.6.

5 Experimental results and discussion

5.1 Maintaining high accuracy

Experimental results for accuracy are presented in Figure 11. These results show that the baseline model (32b_20t) can achieve

85.95% accuracy, while our optimized SNN models with 100×100 attention window achieve 84.12%–85.76% accuracy and our optimized SNN models with 50×50 attention window achieve 76.81%–78.56% accuracy. In the 100×100 attention window case, we observe that 10-bit precision levels generally lead to lower learning quality (i.e., accuracy) across all training epochs as compared to other precision levels; see 1 and 2. These indicate that smaller precision levels have smaller value representation capabilities than the larger ones, and hence they require a longer training period to achieve comparable accuracy. For instance, 10-bit precision levels can achieve close to 80% accuracy in 40 training epoch, while other precision levels can achieve it in 20 training epoch. Due to the same reason, similar patterns are observed in the 50×50 attention window case; see 3 and 4. However, in this case, the accuracy scores from different precision levels are not much different, as they lay within 2% accuracy range; see B4. The reason is that, a smaller attention window means less information to be conveyed and processed by different precision levels, thereby limiting the final accuracy.

5.2 Memory savings

Experimental results for memory footprint are presented in Figure 12. These results show that, in general, our weight quantization step effectively reduces the memory footprint up to 68.75% due to smaller precision levels for representing weight values. For instance, in the 100×100 attention window case, our weight quantization leads to about 50% memory saving for 16b_100w, 62.50% memory saving 12b_100w, and 68.75% memory saving 10b_100w; see 5. Meanwhile, in the 50×50 attention window case, our weight quantization leads to about 91.42% memory saving for 32b_50w, 95.71% memory saving for 16b_50w, 96.78% memory saving 12b_50w, and 97.32% memory saving 10b_50w; see 6. The significant memory savings obtained in the 50×50 attention window case come from the quantization as well as the reduction on the number of weights due to architectural differences in the network as shown in Table 2 and Table 3.

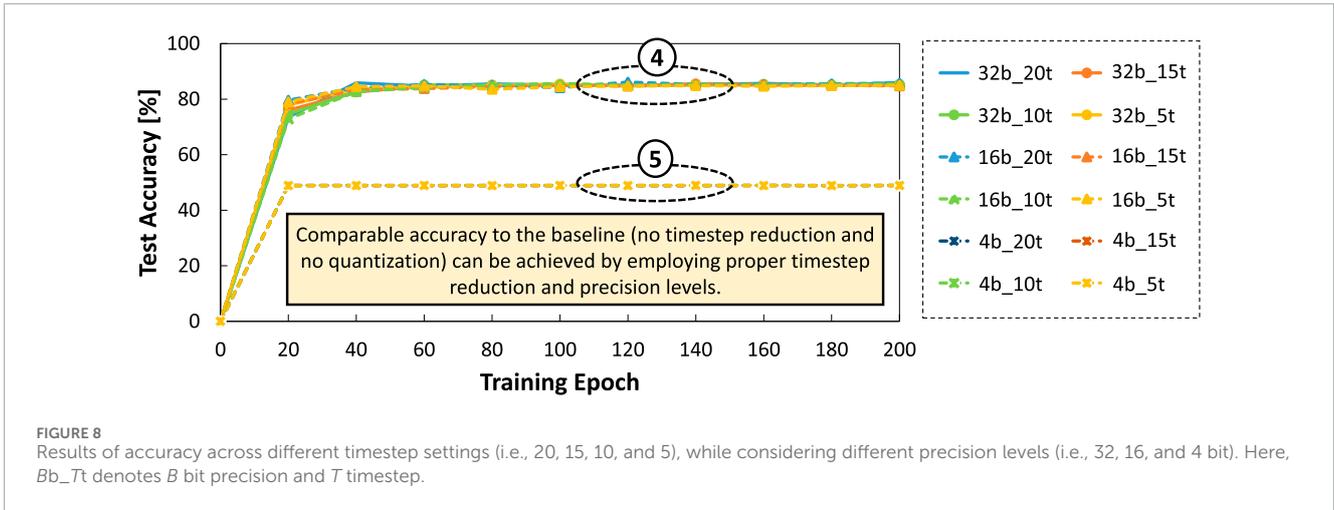


FIGURE 8 Results of accuracy across different timestep settings (i.e., 20, 15, 10, and 5), while considering different precision levels (i.e., 32, 16, and 4 bit). Here, Bb_Tt denotes B bit precision and T timestep.

TABLE 6 Parameter settings for exploring impact of different attention window sizes.

Precision (bit)	Timestep	Attention Window	Training epoch
32	20	100 × 100 and 50 × 50	0–200

as shown by ⑨. Meanwhile, in the 50 × 50 attention window case, our optimizations lead to energy improvements by 1.35x–1.54x for settings 16b_15t, 12b_15t, and 10b_15t; by 2.24x–2.48x for settings 16b_10t, 12b_10t and 10b_10t; as well as by 4.32x–4.66x for settings 16b_5t, 12b_5t, and 10b_5t; as shown by ⑩. These significant energy efficiency improvements come from the joint benefits from weight quantization, timestep reduction, and attention window reduction that decrease the processing power and processing latency, and hence the energy consumption.

5.3 Processing time speed-up

Experimental results for processing latency are presented in Figure 13. These results show that, in general, our timestep reduction step effectively reduce the processing latency as compared to the baseline model (32b_20t). For instance, in the 100 × 100 attention window case, our timestep reduction leads to speed-ups by 1.27x–1.31x for timestep = 15 (i.e., 32b_15t, 16b_15t, 12b_15t, and 10b_15t); by 1.86x–2.02x for timestep = 10 (i.e., 32b_10t, 16b_10t, 12b_10t, and 10b_10t); as well as by 3.58x–3.69x for timestep = 5 (i.e., 32b_5t, 16b_5t, 12b_5t, and 10b_5t); as shown by ⑦. Meanwhile, in the 50 × 50 attention window case, our timestep reduction leads to speed-ups by 1.28x - 1.32x for timestep = 15 (i.e., 32b_15t, 16b_15t, 12b_15t, and 10b_15t); by 2.04x - 2.11x for timestep = 10 (i.e., 32b_10t, 16b_10t, 12b_10t, and 10b_10t); as well as by 3.85x–3.95x for timestep = 5 (i.e., 32b_5t, 16b_5t, 12b_5t, and 10b_5t); as shown by ⑧. These significant speed-ups come from the reduced timesteps which curtail the processing time for spike trains.

5.5 Trade-off analysis

To properly select the appropriate SNN model for the given memory and latency constraints, we perform a trade-off analysis. To do this, we analyze the correlation between accuracy and memory footprint, accuracy and latency, as well as accuracy and energy consumption. For the accuracy-memory analysis, we observe that accuracy decreases as the memory footprint decreases, which is indicated by lower bit precision levels with smaller attention window sizes; see Figure 15a.1 It means that we need to select the appropriate network model whose memory footprint meets the given memory constraint. For instance, if the memory constraint is 8MB, then we can select network models with 12- and 10-bit precision under 100 × 100 attention window, as well as all network models under 50 × 50 attention window. For the accuracy-latency analysis, we observe that accuracy decreases as the latency decreases, which is indicated by shorter timesteps; see Figure 15a.2 It means that we need to select the appropriate network model whose latency meets the given memory constraint. For instance, if the latency constraint is 0.25x from the baseline latency, then we can select network models with 5 timestep under both 100 × 100 and 50 × 50 attention window sizes. For the accuracy-energy analysis, we observe that accuracy also decreases as the energy consumption decreases, which is indicated by lower bit precision levels with shorter timesteps and smaller attention window sizes; see Figure 15a.3 For instance, if the given constraints are 8 MB memory with 0.25x latency from the baseline latency, then we can select the network model with 10b_5t_100w setting, which achieves 84.12% with 68.75% memory saving and 3.58x speed-up, and 4.03x energy efficiency improvement from the baseline model;

5.4 Energy efficiency improvements

Experimental results for energy consumption are presented in Figure 14. These results show that, in general, our optimization techniques effectively reduce the energy consumption as compared to the baseline model (32b_20t). For instance, in the 100 × 100 attention window case, our optimizations lead to energy improvements by 1.36x–1.41x for settings 16b_15t, 12b_15t, and 10b_15t; by 1.94x–2.08x for settings 16b_10t, 12b_10t and 10b_10t; as well as by 3.82x–4.03x for settings 16b_5t, 12b_5t, and 10b_5t;

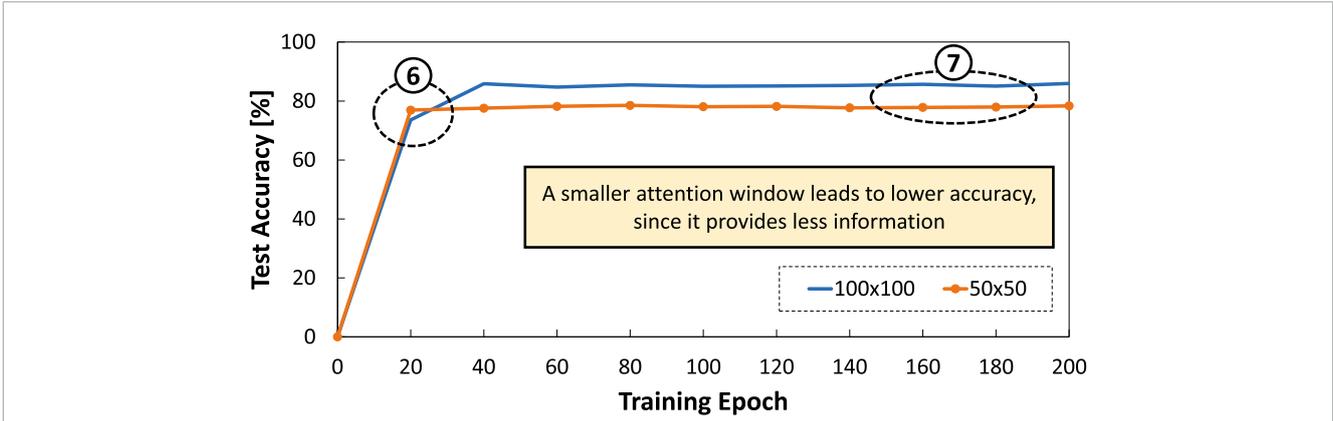


FIGURE 9 Results of accuracy across different attention window sizes (i.e., 100 × 100 and 50 × 50).

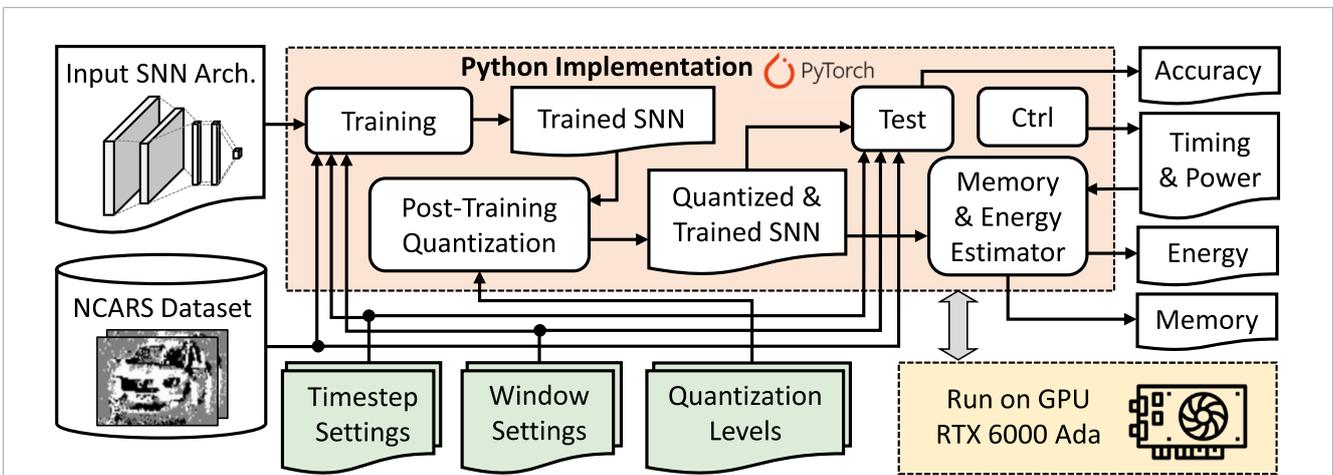


FIGURE 10 Experimental setup for evaluating our SNN4Agents framework. The proposed settings from our SNN4Agents are incorporated into the experimental setup as highlighted in green.

TABLE 7 Parameter settings for our DSE in evaluating our SNN4Agent framework.

Precision (bit)	Timestep	Attention window	Training epoch
32, 16, 12, and 10	20, 15, 10, and 5	100 × 100 and 50 × 50	0–200

see label-11. Furthermore, if the memory constraint is significantly smaller (e.g., 1 MB memory) with 0.25x latency from the baseline latency, then we can select the network model with 10b_5t_50w setting, which achieves 77.10% with 97.32% memory saving and 4x speed-up, and 5.85x energy efficiency improvement from the baseline model; see label-12.

In summary, all these experimental results show that our SNN4Agents framework effectively improves energy efficiency of SNN models for autonomous agent applications. Furthermore, our framework also enables the users to find and select

the suitable SNN model to meet the given memory and latency constraints, i.e., by tuning the optimization settings for weight quantization, timestep setting, and attention window size.

5.6 Computational complexity

Besides performance and efficiency benefits, a set of optimization techniques in our SNN4Agents framework also leads to different levels of computational complexity, which can be quantified through the number of synaptic and neuronal operations across different layers and the given timestep setting. Figure 15B shows the comparison of computational complexity from different designs across different attention window sizes (i.e., 100 × 100 and 50 × 50) and timestep settings (i.e., 20, 15, 10, and 5). In general, a smaller attention window incurs lower computational complexity than the bigger ones. The reason is that, a smaller attention window leads to smaller model size as indicated in

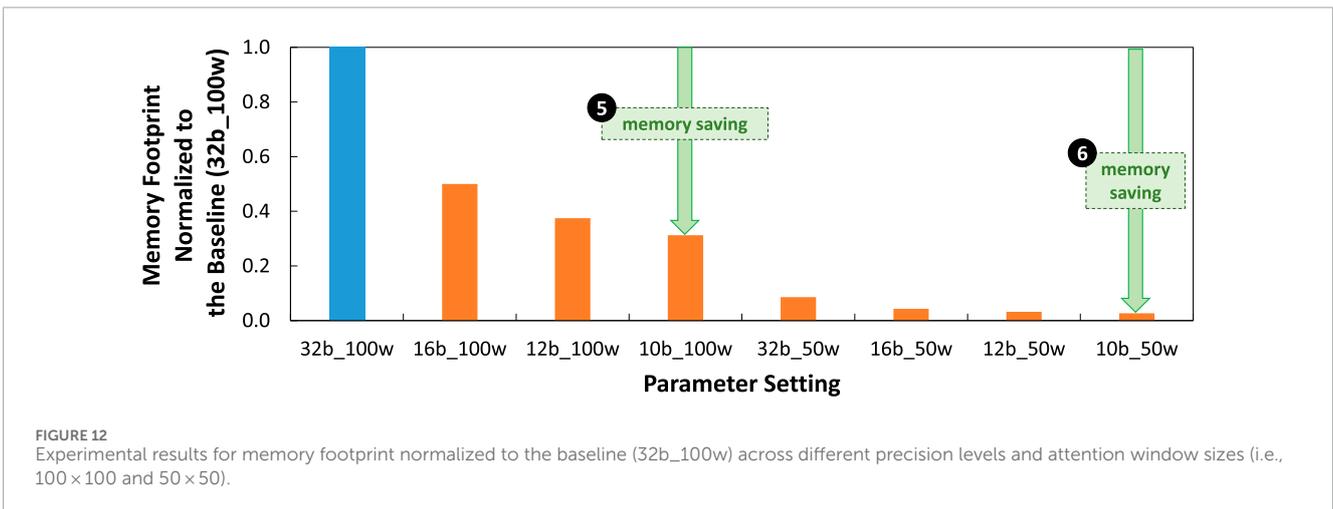
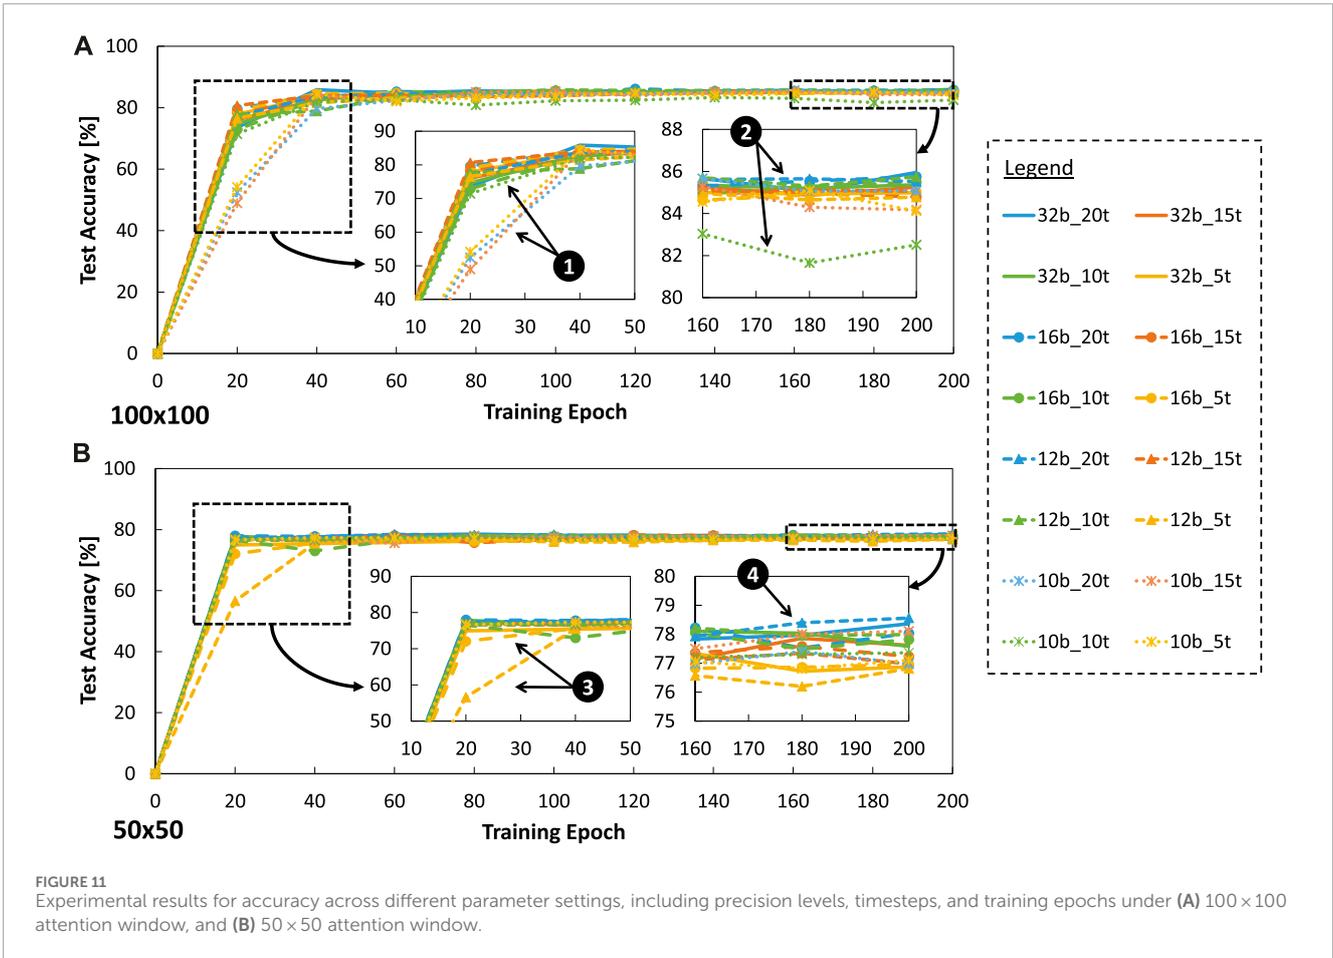


Table 2 and Table 3, and hence smaller number of synaptic and neuronal operations. Specifically, if the larger attention window size is $W_0 \times W_0$ and the smaller attention window size is $W_1 \times W_1$, then the reduction factor of the computational complexity achieved by employing the smaller attention window can be estimated with $\approx (W_0/W_1)^2$. For instance, the 20t_50w model has around 80% lower computation than the state-of-the-art (i.e.,

20t_100w). Furthermore, a smaller number of timestep setting also incurs lower computational complexity than the bigger ones, as the synaptic and neuronal operations are performed in a timestep unit. Specifically, if the larger timestep is T_0 and the smaller timestep is T_1 , then the reduction factor of the computational complexity achieved by employing the smaller timestep can be computed with (T_0/T_1) . For instance, the 15t_100w, 10t_100w,

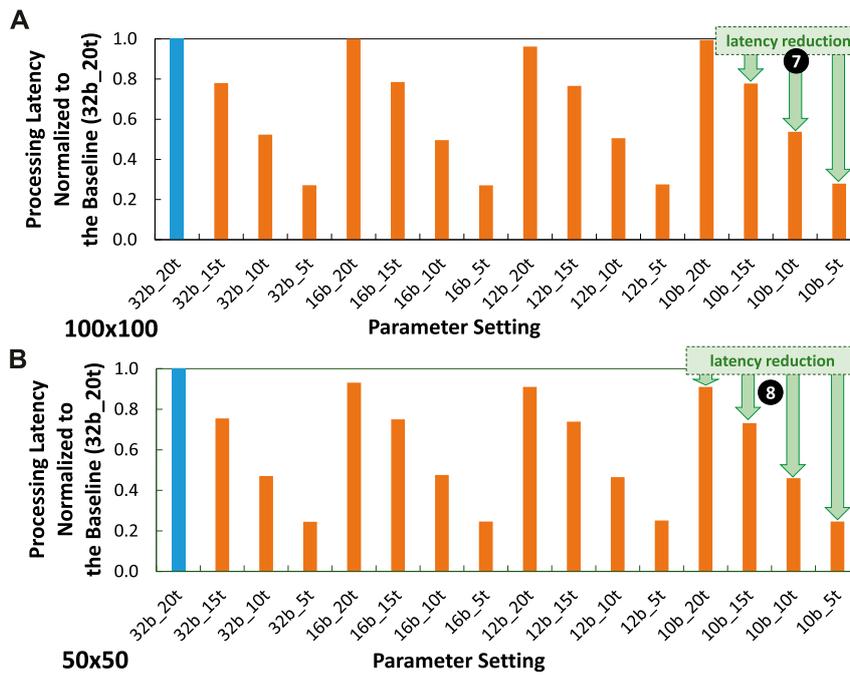


FIGURE 13 Experimental results for latency normalized to the baseline (32b_20t) across different parameter settings, including precision levels, timesteps, and training epochs under (A) 100 × 100 attention window, and (B) 50 × 50 attention window.

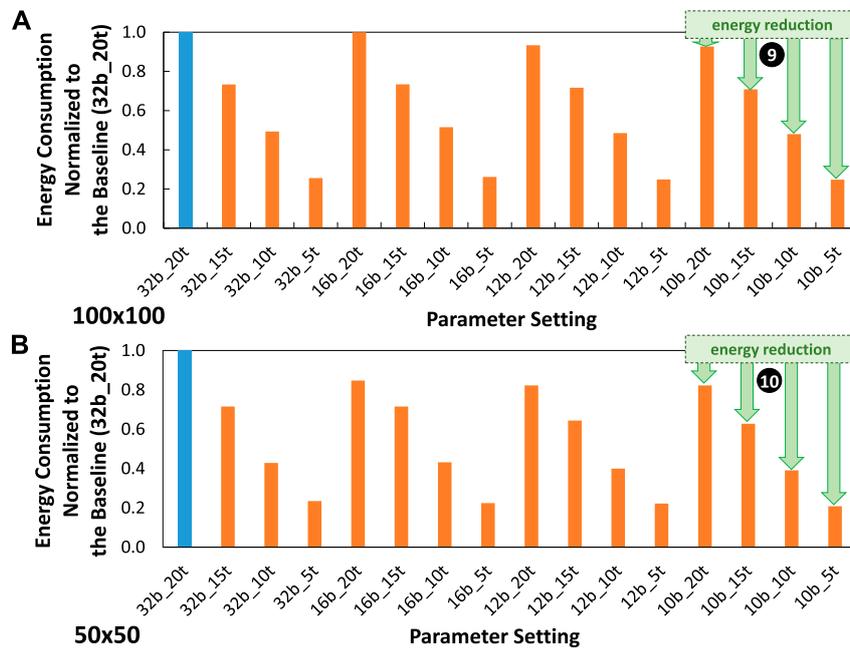


FIGURE 14 Experimental results for energy consumption normalized to the baseline (32b_20t) across different parameter settings, including precision levels, timesteps, and training epochs under (A) 100 × 100 attention window, and (B) 50 × 50 attention window.

and 5t_100w models incur lower computation than the 20t_100w by 25%, 50%, and 75%, respectively. Consequently, jointly employing smaller attention window size and smaller number of timestep will further decrease the computational complexity, whose

reduction factor can be estimated with $\approx (W_0/W_1)^2 \cdot (T_0/T_1)$. For instance, the 15t_50w, 10t_50w, and 5t_50w models incur lower computation than the 20t_100w by about 85%, 90%, and 95%, respectively.

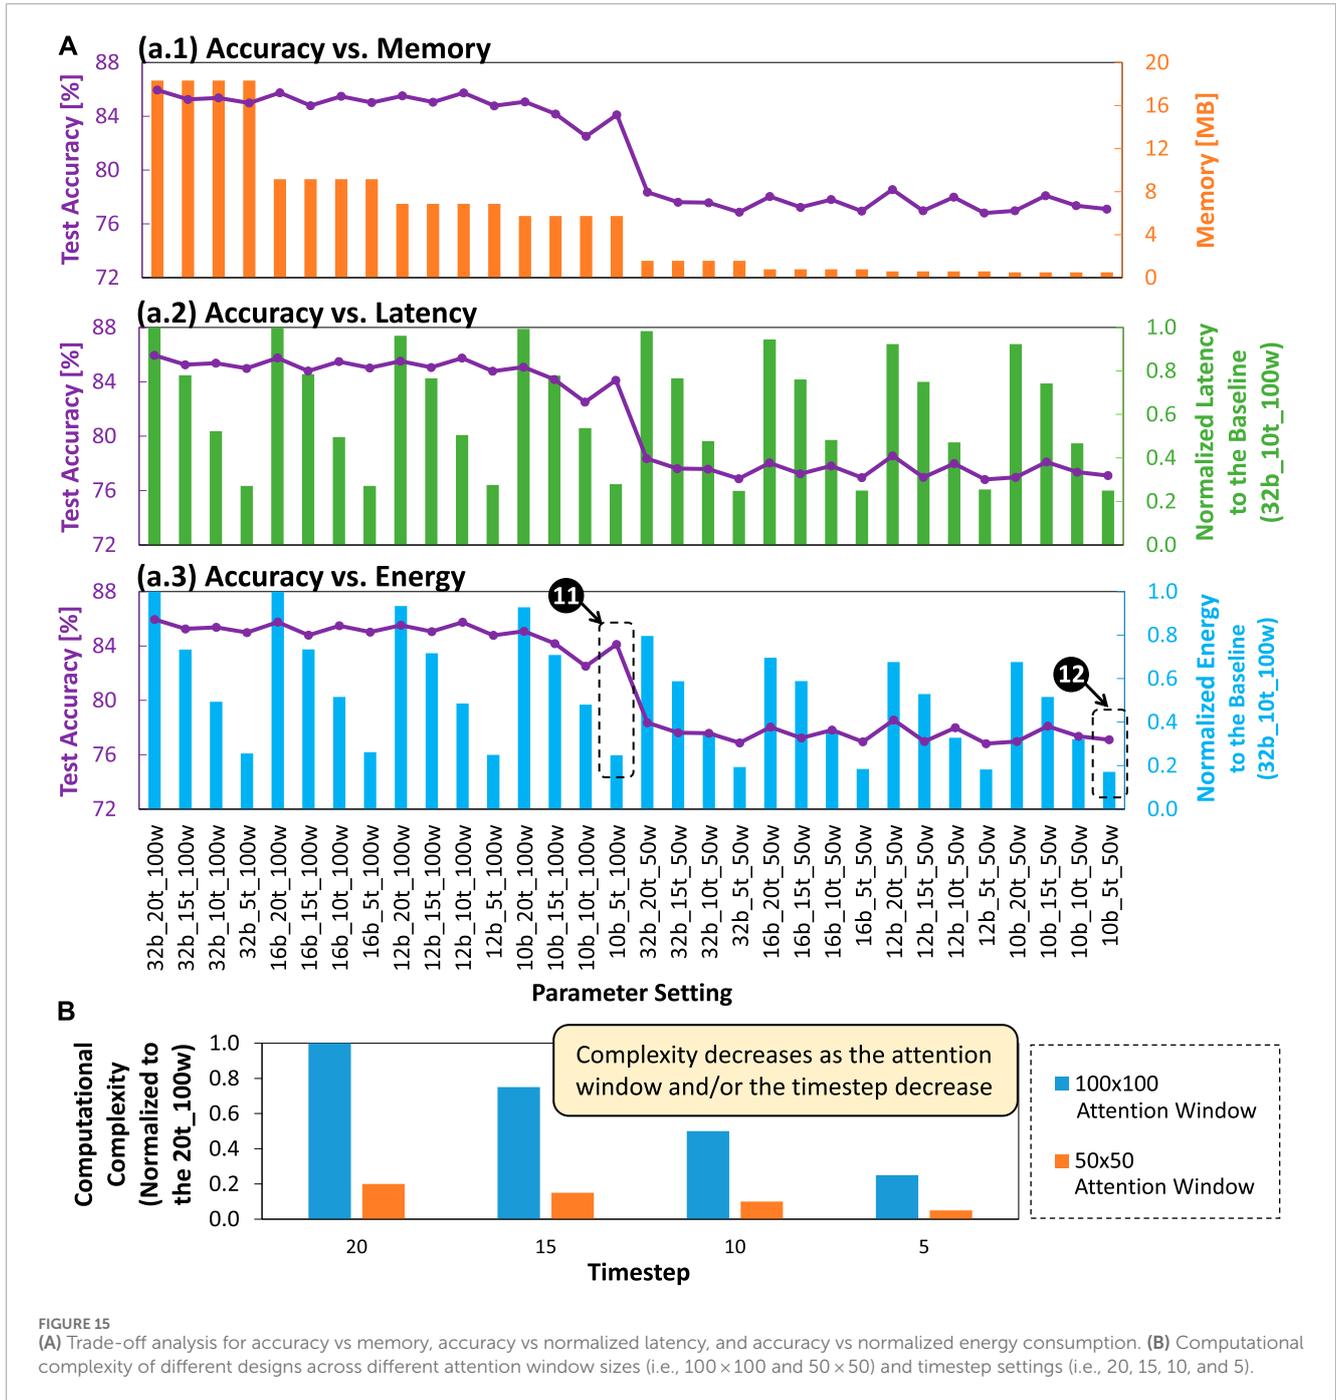


FIGURE 15 (A) Trade-off analysis for accuracy vs memory, accuracy vs normalized latency, and accuracy vs normalized energy consumption. (B) Computational complexity of different designs across different attention window sizes (i.e., 100 × 100 and 50 × 50) and timestep settings (i.e., 20, 15, 10, and 5).

5.7 Further discussion

Our SNN4Agents framework is the first work that incorporates a set of optimization techniques while considering the event-based automotive dataset for enabling the efficient development of SNN-based autonomous agents. Therefore, it has several advantages as the following.

- Our framework incorporates each optimization technique in a modular form. Therefore, the existing optimization techniques in the framework can be activated or deactivated as per

the design requirements. Furthermore, new optimization techniques can also be incorporated feasibly into the framework.

- Our framework can be used for design space exploration to investigate and understand the role of different SNN parameters. For instance, this framework can be utilized to observe the impact of specific parameters (e.g., membrane threshold potential) on the accuracy (Bano et al., 2024).
- The generated SNN model has direct interfacing with event-based input data stream, thereby enabling efficient integration with event-based sensors (e.g., DVS Camera).

Besides these advantages, our framework in the current form can still be improved further to enhance the generated SNN model. Our SNN4Agents currently supports the STBP learning rule with offline-based training scenarios, which may be insufficient for some application use-cases. For instance, some autonomous agents may need to continuously adapt to dynamic operational environments, hence our SNN4Agents framework needs to be enhanced with advanced neural architectures and/or continual learning algorithms, while considering both offline and online training scenarios (Putra et al., 2024). Toward this, in the future, we plan to continue extending the work by incorporating more complex datasets in our framework, such as EventKITTI (Liang et al., 2022), then evaluating the performance as well as testing it for a real-world robotic application use-case (e.g., UGV rover).

6 Conclusion

In this work, we propose an SNN4Agents framework that employs a set of optimization techniques for developing energy-efficient SNNs targeting autonomous agent applications. Here, our SNN4Agents compresses the SNN model through weight quantization, optimizes processing time through timestep reduction, and optimizes input samples through attention window reduction. The experimental results show that our proposed framework effectively improves energy efficiency, reduces memory footprint and latency, while maintaining high accuracy. If we consider a tolerance range of 2% lower accuracy from the baseline, we can achieve 84.12% accuracy with 68.75% memory saving, 3.58x speed-up, and 4.03x energy efficiency improvement. In this manner, our SNN4Agents framework paves the way for further research and studies toward enabling the efficient development of SNN-based autonomous agents, and can be enhanced by incorporating other optimization techniques.

Data availability statement

Publicly available datasets were analyzed in this study. NCARS dataset can be accessed at <https://www.prophesee.ai/2018/03/13/dataset-n-cars/>.

References

- Bano, I., Putra, R. V. W., Marchisio, A., and Shafique, M. (2024). A methodology to study the impact of spiking neural network parameters considering event-based automotive data. *arXiv Prepr. arXiv:2404.03493*.
- Bartolozzi, C., Indiveri, G., and Donati, E. (2022). Embodied neuromorphic intelligence. *Nat. Commun.* 13, 1024. doi:10.1038/s41467-022-28487-2
- Bonnevie, R., Duberg, D., and Jensfelt, P. (2021). "Long-term exploration in unknown dynamic environments," in 2021 7th International Conference on Automation, Robotics and Applications (ICARA), 32–37. doi:10.1109/ICARA51699.2021.9376367
- Bu, T., Fang, W., Ding, J., Dai, P., Yu, Z., and Huang, T. (2022). "Optimal ANN-SNN conversion for high-accuracy and ultra-low-latency spiking neural networks," in International Conference on Learning Representations.
- Chowdhury, S. S., Rathi, N., and Roy, K. (2021). One timestep is all you need: training spiking neural networks with ultra low latency. *Corr. abs/2110.05929*.
- Cordone, L., Miramond, B., and Thierion, P. (2022). "Object detection with spiking neural networks on automotive event data," in 2022 International Joint Conference on Neural Networks (IJCNN), 1–8. doi:10.1109/IJCNN55064.2022.9892618
- Guo, W., Fouda, M. E., Eltawil, A. M., and Salama, K. N. (2021). Neural coding in spiking neural networks: a comparative study for robust neuromorphic systems. *Front. Neurosci. (FNINS)* 15, 638474. doi:10.3389/fnins.2021.638474
- Gupta, S., Agrawal, A., Gopalakrishnan, K., and Narayanan, P. (2015). "Deep learning with limited numerical precision," in International Conference on Machine Learning (ICML). Editors F. Bach, and D. Blei, 1737–1746.
- Hao, Z., Ding, J., Bu, T., Huang, T., and Yu, Z. (2023). "Bridging the gap between ANNs and SNNs by calibrating offset spikes," in The Eleventh International Conference on Learning Representations.
- Izhikevich, E. M. (2004). Which model to use for cortical spiking neurons? *IEEE Trans. Neural Netw. (TNN)* 15, 1063–1070. doi:10.1109/tnn.2004.832719
- Li, Y., He, X., Dong, Y., Kong, Q., and Zeng, Y. (2022). Spike calibration: fast and accurate conversion of spiking neural network for object detection and segmentation. *arXiv Prepr. arXiv:2207.02702*.
- Liang, Z., Cao, H., Yang, C., Zhang, Z., and Chen, G. (2022). "Global-local feature aggregation for event-based object detection on eventkitti," in 2022 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI), 1–7. doi:10.1109/MFI5806.2022.9913852

Codes of SNN4Agents are available at <https://github.com/rachmadvwp/SNN4Agents>.

Author contributions

RP: Conceptualization, Data curation, Investigation, Methodology, Software, Validation, Visualization, Writing—original draft, Writing—review and editing. AM: Conceptualization, Methodology, Writing—original draft, Writing—review and editing, Visualization. MS: Conceptualization, Supervision, Writing—original draft, Writing—review and editing.

Funding

The author(s) declare that financial support was received for the research, authorship, and/or publication of this article. This work was partially supported by the NYUAD Center for Artificial Intelligence and Robotics (CAIR), funded by Tamkeen under the NYUAD Research Institute Award CG010.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

- Maass, W. (1997). Networks of spiking neurons: the third generation of neural network models. *Neural Netw.* 10, 1659–1671. doi:10.1016/s0893-6080(97)00011-7
- Massa, R., Marchisio, A., Martina, M., and Shafique, M. (2020). “An efficient spiking neural network for recognizing gestures with a DVS camera on the loihi neuromorphic processor,” in International Joint Conference on Neural Networks (IJCNN) (IEEE), 1–9. doi:10.1109/IJCNN48605.2020.9207109
- Micikevicius, P., Narang, S., Alben, J., Diamos, G., Elsen, E., Garcia, D., et al. (2018). “Mixed precision training,” in International Conference on Learning Representation (ICLR).
- Neftci, E. O., Mostafa, H., and Zenke, F. (2019). Surrogate gradient learning in spiking neural networks: bringing the power of gradient-based optimization to spiking neural networks. *IEEE Signal Process. Mag.* 36, 51–63. doi:10.1109/msp.2019.2931595
- Posch, C., Matolin, D., and Wohlgenannt, R. (2011). A QVGA 143 db dynamic range frame-free PWM image sensor with lossless pixel-level video compression and time-domain CDS. *IEEE J. Solid State Circuits* 46, 259–275. doi:10.1109/JSSC.2010.2085952
- Putra, R. V. W., Hanif, M. A., and Shafique, M. (2021a). “Respawn: energy-efficient fault-tolerance for spiking neural networks considering unreliable memories,” in 2021 IEEE/ACM International Conference On Computer Aided Design (ICCAD), 1–9. doi:10.1109/ICCAD51958.2021.9643524
- Putra, R. V. W., Hanif, M. A., and Shafique, M. (2021b). “Sparkxd: a framework for resilient and energy-efficient spiking neural network inference considering approximate dram,” in 2021 58th ACM/IEEE Design Automation Conference, 379–384. doi:10.1109/DAC18074.2021.9586332
- Putra, R. V. W., Hanif, M. A., and Shafique, M. (2022a). Enforcesnn: enabling resilient and energy-efficient spiking neural network inference considering approximate drams for embedded systems. *Front. Neurosci. (FNINS)* 16, 937782. doi:10.3389/fnins.2022.937782
- Putra, R. V. W., Hanif, M. A., and Shafique, M. (2022b). “Softsnn: low-cost fault tolerance for spiking neural network accelerators under soft errors,” in 59th ACM/IEEE Design Automation Conference (DAC), 151–156.
- Putra, R. V. W., Hanif, M. A., and Shafique, M. (2023). Rescuesnn: enabling reliable executions on spiking neural network accelerators under permanent faults. *Front. Neurosci. (FNINS)* 17, 1159440. doi:10.3389/fnins.2023.1159440
- Putra, R. V. W., Marchisio, A., Zayer, F., Dias, J., and Shafique, M. (2024). Embodied neuromorphic artificial intelligence for robotics: perspectives, challenges, and research development stack. *arXiv Prepr. arXiv:2404.03325*. doi:10.48550/arXiv.2404.03325
- Putra, R. V. W., and Shafique, M. (2020). Fspinn: an optimization framework for memory-efficient and energy-efficient spiking neural networks. *IEEE Trans. Computer-Aided Des. Integr. Circuits Syst. (TCAD)* 39, 3601–3613. doi:10.1109/TCAD.2020.3013049
- Putra, R. V. W., and Shafique, M. (2021a). “Q-spinn: a framework for quantizing spiking neural networks,” in 2021 International Joint Conference on Neural Networks (IJCNN), 1–8. doi:10.1109/IJCNN52387.2021.9534087
- Putra, R. V. W., and Shafique, M. (2021b). “Spikedyn: a framework for energy-efficient spiking neural networks with continual and unsupervised learning capabilities in dynamic environments,” in 58th ACM/IEEE Design Automation Conference (DAC), 1057–1062.
- Putra, R. V. W., and Shafique, M. (2022). “Ipspikecon: enabling low-precision spiking neural network processing for efficient unsupervised continual learning on autonomous agents,” in 2022 International Joint Conference on Neural Networks (IJCNN), 1–8. doi:10.1109/IJCNN55064.2022.9892948
- Putra, R. V. W., and Shafique, M. (2023a). “Mantis: enabling energy-efficient autonomous mobile agents with spiking neural networks,” in 2023 9th International Conference on Automation, Robotics and Applications (ICARA), 197–201. doi:10.1109/ICARA56516.2023.10125781
- Putra, R. V. W., and Shafique, M. (2023b). “Topspark: a timestep optimization methodology for energy-efficient spiking neural networks on autonomous mobile agents,” in IEEE RSJ International Conference on Intelligent Robots and Systems (IROS), 3561–3567. doi:10.1109/IROS55552.2023.10342499
- Putra, R. V. W., and Shafique, M. (2024). Spikenas: a fast memory-aware neural architecture search framework for spiking neural network-based autonomous agents. *arXiv Prepr. arXiv:2402.11322*. doi:10.48550/arXiv.2402.11322
- Rathi, N., Chakraborty, I., Kosta, A., Sengupta, A., Ankit, A., Panda, P., et al. (2023). Exploring neuromorphic computing based on spiking neural networks: algorithms to hardware. *ACM Comput. Surv.* 55, 1–49. doi:10.1145/3571155
- Roy, K., Jaiswal, A., and Panda, P. (2019). Towards spike-based machine intelligence with neuromorphic computing. *Nature* 575, 607–617. doi:10.1038/s41586-019-1677-2
- Rückauer, B., Känzig, N., Liu, S., Delbrück, T., and Sandamirskaya, Y. (2019). Closing the accuracy gap in an event-based visual recognition task. *Corr. abs/1906.08859*.
- Schuman, C. D., Kulkarni, S. R., Parsa, M., Mitchell, J. P., Date, P., and Kay, B. (2022). Opportunities for neuromorphic computing algorithms and applications. *Nat. Comput. Sci.* 2, 10–19. doi:10.1038/s43588-021-00184-y
- Sironi, A., Brambilla, M., Bourdis, N., Lagorce, X., and Benosman, R. (2018). “Hats: histograms of averaged time surfaces for robust event-based object classification,” in 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 1731–1740. doi:10.1109/CVPR.2018.00186
- Viale, A., Marchisio, A., Martina, M., Masera, G., and Shafique, M. (2021). “Carsnn: an efficient spiking neural network for event-based autonomous cars on the loihi neuromorphic research processor,” in 2021 International Joint Conference on Neural Networks (IJCNN), 1–10. doi:10.1109/IJCNN52387.2021.9533738
- Viale, A., Marchisio, A., Martina, M., Masera, G., and Shafique, M. (2022). “Lanesnns: spiking neural networks for lane detection on the loihi neuromorphic processor,” in IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (IEEE), 79–86.
- Wang, Z., Guo, L., and Adjouadi, M. (2014). “A biological plausible generalized leaky integrate-and-fire neuron model,” in 36th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBC 2014, Chicago, IL, USA, August 26–30, 2014 (IEEE), 6810–6813. doi:10.1109/EMBC.2014.6945192
- Wu, Y., Deng, L., Li, G., Zhu, J., and Shi, L. (2018). Spatio-temporal backpropagation for training high-performance spiking neural networks. *Front. Neurosci. (FNINS)* 12, 331. doi:10.3389/fnins.2018.00331