

MD-SNN: Membrane Potential-aware Distillation on Quantized Spiking Neural Network

Donghyun Lee¹, Abhishek Moitra², Youngeun Kim³†, Ruokai Yin⁴, Priyadarshini Panda¹
¹University of Southern California, ²Washington State University, ³Amazon, ⁴Yale University
 donghyun.lee.1@usc.edu

Abstract—Spiking Neural Networks (SNNs) offer a promising and energy-efficient alternative to conventional neural networks, thanks to their sparse binary activation. However, they face challenges regarding memory and computation overhead due to complex spatio-temporal dynamics and the necessity for multiple backpropagation computations across timesteps during training. To mitigate this overhead, compression techniques such as quantization are applied to SNNs. Yet, naively applying quantization to SNNs introduces a mismatch in membrane potential, a crucial factor for the firing of spikes, resulting in accuracy degradation. In this paper, we introduce Membrane-aware Distillation on quantized Spiking Neural Network (MD-SNN), which leverages membrane potential to mitigate discrepancies after weight, membrane potential, and batch normalization quantization. To our knowledge, this study represents the first application of membrane potential knowledge distillation in SNNs. We validate our approach on various datasets, including CIFAR10, CIFAR100, N-Caltech101, and TinyImageNet, demonstrating its effectiveness for both static and dynamic data scenarios. Furthermore, for hardware efficiency, we evaluate the MD-SNN with SpikeSim platform, finding that MD-SNNs achieve $14.85\times$ lower energy-delay-area product (EDAP), $2.64\times$ higher TOPS/W, and $6.19\times$ higher TOPS/mm² compared to floating point SNNs at iso-accuracy on N-Caltech101 dataset. Code is available at Github.

Index Terms—Spiking Neural Network, Quantization, Knowledge Distillation

I. INTRODUCTION

In recent years, Spiking Neural Networks (SNNs) have emerged as a compelling energy-efficient alternative to conventional Artificial Neural Networks (ANNs) [1], [2]. Distinguished by their event-driven computation and sparse binary spike communication across temporal dimension, SNNs offer a biologically-inspired paradigm for neural information processing. This unique computational model not only promises significant energy savings but also enables deployment across diverse hardware platforms, from neuromorphic chips to edge devices [3]–[7].

The pursuit of energy-efficient SNNs has spawned two primary training methodologies: ANN-to-SNN conversion [8]–[12] and direct training via backpropagation with surrogate gradients [13]–[16]. While backpropagation-based approaches have achieved state-of-the-art accuracy with fewer timesteps, they impose substantial memory overhead due to the storage requirement for temporal activations during training [5]. To address these computational challenges, researchers have developed various compression techniques, including quantization [17]–[19], pruning [20], Knowledge Distillation (KD) [21]–[23], and tensor decomposition [24].

†Work done prior to joining Amazon.




Fig. 1: Impact of knowledge distillation on membrane potential distributions. (a) Before distillation: significant mismatch between Floating-Point (FP) and 4-bit quantized models. (b) After distillation: aligned distributions through MD-SNN. (c) Performance comparison on CIFAR100, where MD-SNN (green star) surpasses both the quantized student and FP teacher models.

Among these techniques, quantization has proven particularly effective in reducing both memory footprint and computational complexity. By reducing the bit-precision of weights and membrane potentials, quantization methods have successfully compressed SNNs while maintaining competitive accuracy [17], [18]. For instance, recent work has demonstrated that SNNs can operate with weights and membrane potentials quantized to as low as 2-bit integer precision with minimal accuracy loss [19]. However, aggressive quantization introduces a critical challenge: *significant distortion in membrane potential distributions across timesteps, which directly impacts spike generation and network performance*.

Fig. 1(a) illustrates the critical challenge: membrane potential distributions between Floating-Point (FP) and 4-bit quantized models exhibit substantial divergence, disrupting the precise spike timing essential for SNN information encoding. This divergence is particularly problematic because membrane potentials, the continuous state variables that integrate temporal inputs and govern spike generation, encode rich dynamical information absent from binary spike outputs. Our KD approach successfully aligns these distributions (Fig. 1(b)), achieving performance that surpasses even the full-precision teacher (Fig. 1(c)). While KD has been explored for SNNs, existing methods exhibit fundamental limitations. Approaches using ANN teachers [22], [23] risk losing spike-specific temporal dynamics, while SNN-based methods focusing solely on output spikes [21] overlook the informative membrane potential trajectories. By directly transferring membrane distributions, our work preserves the complete spatiotemporal dynamics critical

for accurate spike generation in quantized networks.

In this paper, we introduce Membrane-aware Distillation for quantized Spiking Neural Networks (MD-SNN), a novel knowledge distillation framework that explicitly leverages membrane potential dynamics to guide quantized student networks. As illustrated in Fig. 2(a), MD-SNN employs dual distillation pathways, i.e., logit and membrane potential, between a full-precision SNN teacher and a quantized student. The key innovation lies in recognizing membrane potentials as feature representations analogous to hidden activations in ANNs, but with additional temporal dynamics unique to SNNs. Our approach makes several critical design choices. First, we employ Kullback-Leibler (KL) divergence rather than the conventional Mean Squared Error (MSE) for membrane distillation. Our empirical analysis reveals that MSE leads to excessive spike generation, creating redundant information flow that degrades performance. KL divergence, by treating membrane distributions probabilistically, better preserves the relative importance of different membrane states. Second, we introduce a versatile teacher framework where a single teacher trained with T timesteps can effectively guide multiple students operating at timestep $t \leq T$, diminishing training costs across different temporal configurations. The main contributions of this work are as:

- We present the first knowledge distillation method that explicitly transfers membrane potential dynamics between SNN models, from full-precision teachers to quantized students. Our temporal and spatial (group-wise) distillation design achieves an optimal balance between effectiveness and computational efficiency.
- Our framework enables a single teacher model with timestep T to guide students with varying timesteps ($t \leq T$), reducing the training overhead by 30% for multi-configuration deployments.
- Extensive experiments on static (CIFAR-10/100, TinyImageNet) and neuromorphic (N-Caltech101) datasets demonstrate that MD-SNN consistently outperforms existing quantization methods, achieving hardware efficiency with up to $14.85\times$ EDAP reduction while maintaining competitive accuracy.

II. RELATED WORK

A. Spiking Neural Network

The advancement of SNNs has spurred an extensive investigation into methods for effectively training and deploying SNNs. Particularly, the training method of SNNs can be divided into two categories, i.e., ANN-to-SNN conversion and backpropagation with surrogate gradients. Within the domain of SNN research, one prominent avenue is the conversion of conventional ANNs into their spiking counterparts. This approach involves the transformation of pre-trained ANNs, typically designed for continuous value operations, into SNN configurations. The primary objective is to capitalize on the wealth of knowledge cultivated in the ANN domain to expedite and enhance the training of SNNs. Eminent researchers in this field, including Diehl et al. [8]–[11], have made substantial

contributions by elucidating the conversion process and demonstrating the potential of ANNs as instructive mentors for SNNs.

Backpropagation, a well-established method for training conventional ANNs, has found application in the context of SNNs, albeit with requisite modifications to accommodate the discrete nature of spike events. An influential facet of this research realm is the utilization of surrogate gradients, which streamline the integration of backpropagation techniques for SNN training. [13]–[16] have significantly contributed to this domain by offering valuable insights into the training of SNNs.

B. Efficient methods for compressing SNNs

The training of SNNs involves computationally intensive backward computations, particularly in surrogate gradient calculations. To mitigate this computational burden, a range of efficient methods have been introduced, all focused on reducing memory and computation costs such as KD [21]–[23], quantization [17]–[19], pruning [20], and tensor decomposition [24]. In terms of KD, Xu et al. [22] propose an ANN-based teacher model that imparts its knowledge to an SNN-based student model, resulting in accelerated convergence. Kushawaha et al. [21] take a different route, employing an SNN-based teacher model, which shares its spike distribution knowledge with the student model to stabilize the training process. Guo et al. [23] introduce a novel KD approach by transferring knowledge from an ANN teacher model to an SNN student model using both logit and feature distillation. Additionally, they enforce constraints to ensure that the ANN and SNN models share the same singular vectors during the training process, resulting in improved accuracy. These KD strategies have demonstrated their effectiveness in enabling small SNN models to achieve performance levels comparable to those of the teacher model. Other approaches, such as pruning and tensor decomposition, have been introduced to make SNN architectures more hardware-friendly. Unstructured pruning, guided by the Lottery Ticket Hypothesis (LTH) [25], has been effectively employed in SNNs [20]. Moreover, tensor decomposition techniques, particularly the Tensor Train (TT) decomposition, have been harnessed to optimize SNNs, leading to reductions in weight count, training energy expenditure, and training time [24].

III. PRELIMINARY

A. Leaky Integrate-and-Fire Neuron

As a fundamental building block of SNNs, the Leaky Integrate-and-Fire (LIF) neuron has emerged as an important component for energy-efficient computation. The LIF neuron is a non-linear activation function that determines whether neurons fire spikes as follows:

$$\mathbf{u}_l^{t+1} = \tau_m \mathbf{u}_l^t + \mathbf{W}_l f(\mathbf{u}_{l-1}^t) \quad (1)$$

$$f(\mathbf{u}_l^t) = \mathbb{1}_{\mathbf{u}_l^t > V_{th}} \quad (2)$$

where, \mathbf{u}_l^t is the membrane potential in l -th layer at timestep t , $\tau_m \in (0, 1]$ is the leaky factor for membrane potential leakage, \mathbf{W} is the weight, and $f(\cdot)$ is the LIF function with firing threshold V_{th} . Therefore, when the membrane \mathbf{u}_l^t is higher than V_{th} , the LIF function fires a spike and the membrane potential is reset to 0.

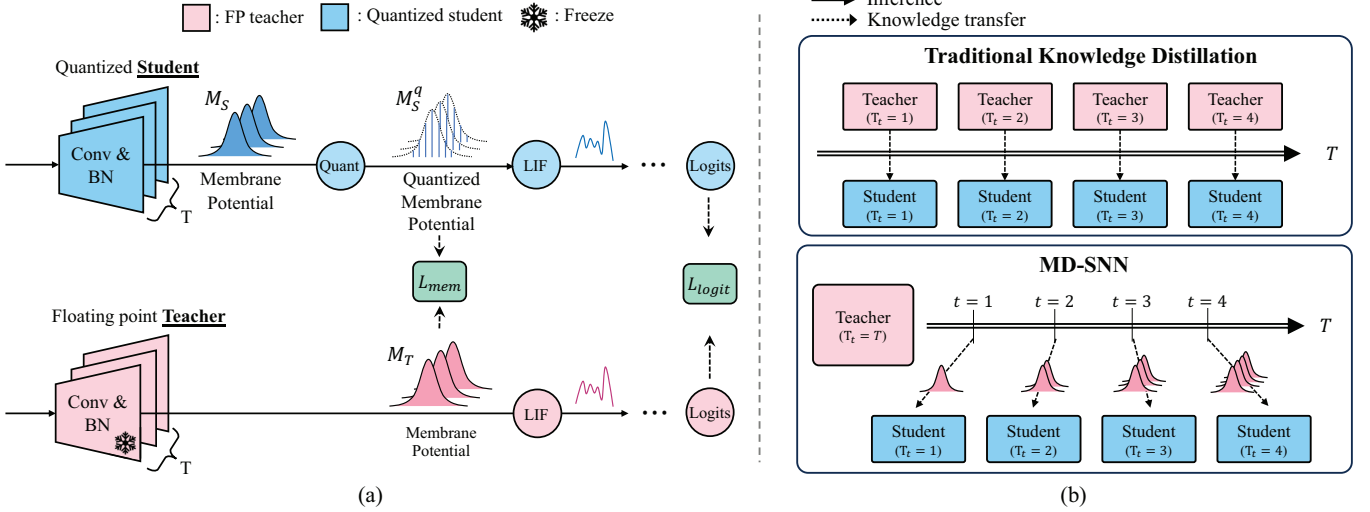


Fig. 2: MD-SNN architecture and training paradigm. (a) Dual distillation paths using membrane potentials and logits between FP teacher and quantized student. (b) Versatile teacher framework: single $T = 4$ teacher guides multiple students via temporal membrane alignment, contrasted with traditional one-to-one knowledge distillation requiring separate teachers per timestep.

B. Knowledge Distillation

Knowledge distillation is a potent technique designed to bridge the gap between larger, well-built teacher models and more compact student models, enabling the students to quickly attain optimal solutions. There are two primary methods in KD: logit distillation and feature distillation. The logit distillation method employs softened logits along with KL divergence, effectively minimizing the distance between the logits of the teacher and student models [26]. This approach encourages the student model to inherit the knowledge captured within the teacher model's predictions. On the other hand, feature distillation takes a different route. It aims to align the hidden features of both teacher and student models, facilitating rapid convergence and enhancing the student model's performance [27]. Our MD-SNN is a branch of feature distillation because we align the membrane potential of the student and the teacher. Critically, membrane potentials (\mathbf{u}_i^t in Eqn.(1)) represent pre-activation states before the LIF function, capturing rich temporal dynamics that directly influence spike generation. Furthermore, we adopt the KL divergence function for membrane distillation instead of the commonly used MSE employed when measuring the distance between hidden features in the teacher and student architectures.

IV. PROPOSED METHOD

A. Quantization Schemes

MD-SNN framework employs distinct quantization strategies for different network components to maintain computational efficiency while preserving information flow. We apply uniform quantization to weights and membrane potentials, and adopt a different approach for Batch Normalization (BN) parameters.

1) *Weight and Membrane Potential Quantization*: Inspired by MINT [19], we employ symmetric uniform quantization with adaptive scaling for both weights and membrane potentials. Given a tensor w and target bit-width b , our quantization

applies a hyperbolic tangent transformation followed by uniform quantization:

$$Q(w, b) = \frac{\alpha \cdot \text{Round}(\bar{w} \cdot \Delta(b))}{\Delta(b)}, \quad (3)$$

where $\bar{w} = \text{clamp}(\tanh(w)/\alpha, -1, 1)$, $\alpha = \max(|\tanh(w)|)$ is the dynamic scale factor, and $\Delta(b) = 2^{b-1} - 1$ represents the quantization levels. Here, $\text{clamp}(x, a, b) = \min(\max(x, a), b)$ constrains values to the range $[a, b]$, and $\max(|\cdot|)$ denotes the maximum absolute value across all tensor elements. We employ the Straight-Through Estimator (STE) for gradient backpropagation during training. The tanh transformation bounds weight values while maintaining smooth gradients, and the adaptive scaling factor α ensures efficient utilization of quantization levels. This approach is particularly effective for membrane potentials, as the adaptive scaling and symmetric quantization provide effective representation of the temporal dynamics in spiking neurons.

2) *Batch Normalization Quantization*: Batch normalization layers require special consideration as they contain multiple learnable parameters (scale γ , shift β) and running statistics (mean μ , variance σ^2). Direct quantization of these parameters can lead to significant accuracy degradation due to their different value ranges and sensitivities. Following the WAGEUBN approach [28], we quantize the weights and bias of BN layers as follows:

$$Q_{BN}(w, b) = \text{clamp}\left[Q(w, b), -1 + \frac{1}{\Delta(b) + 1}, 1 - \frac{1}{\Delta(b) + 1}\right], \quad (4)$$

where $Q(w, b)$ is the quantization function defined in Eq. (3) with $\alpha = 1$. The clipping range is slightly reduced from $[-1, 1]$ to ensure proper gradient flow during backpropagation.

B. Membrane-aware Distillation

In this section, we present the details of membrane-aware distillation and training strategies. The membrane potential, computed by weights and input spikes shown as Eqn. (1),

TABLE I: Accuracy comparison of MD-SNN against floating-point baselines on CIFAR-100, N-Caltech101, and TinyImageNet datasets. FP32 refers to full-precision (32-bit) trained without quantization. MS-N denotes our MD-SNN with N-bit quantization for both weights and membrane potentials.

Dataset	Model	Timestep	Method	Accuracy
CIFAR-100	ResNet18	4	FP32	74.34%
			MS-8	74.91%
			MS-4	74.41%
N-Caltech101	ResNet18	10	FP32	80.31%
			MS-8	81.73%
			MS-4	81.10%
TinyImageNet	ResNet34	4	FP32	60.34%
			MS-8	60.59%
			MS-4	59.90%

TABLE II: Comparisons between existing quantization methods on SNNs. W and U represent the bit-widths of weight and membrane potential. † represents our implementation.

Data	Method	Model	Precision (W / U)	Timestep	Accuracy
CIFAR10	STBP-Quant [29]	AlexNet	8 / 14	8	86.65 %
	STBP-Quant [29]	AlexNet	4 / 10	8	84.09 %
	ST-Quant [30]	VGG9	5 / 32	25	88.60 %
	ADMM-Quant [31]	7Conv+3FC	4 / 32	8	89.40 %
	SpikeSim [32]	VGG9	8 / 32	5	88.11 %
	MINT [19]	ResNet19	8 / 8	4	91.36 %
	MINT [19]	ResNet19	4 / 4	4	91.45 %
	MD-SNN (ours)	ResNet18	8 / 8	4	94.27 %
	MD-SNN (ours)	ResNet18	4 / 4	4	94.05 %
CIFAR100	ST-Quant [30]	VGG11	5 / 32	30	66.20 %
	SpikeSim [32]	VGG16	8 / 32	10	65.23 %
	MINT [19]†	ResNet19	8 / 8	4	68.53 %
	MINT [19]†	ResNet19	4 / 4	4	68.79 %
	MD-SNN (ours)	ResNet18	8 / 8	4	74.91 %
	MD-SNN (ours)	ResNet18	4 / 4	4	74.41 %
TinyImageNet	SpikeSim [32]	VGG16	8 / 32	10	54.03 %
	MINT [19]†	ResNet34	8 / 8	4	42.94 %
	MINT [19]†	ResNet34	4 / 4	4	44.94 %
	MD-SNN (ours)	ResNet34	8 / 8	4	60.59 %
	MD-SNN (ours)	ResNet34	4 / 4	4	59.90 %

(2), is an important intrinsic component as it determines the output spikes of the specific neuron. Therefore, the membrane distribution of the teacher model can be a good guide to help the student achieve optimal solutions. The holistic scheme of MD-SNN is shown in Fig. 2(a). We group-wisely and time-wisely extract the membrane distributions and compute the distance of membrane distribution between the teacher and student model as follows:

$$L_{mem} = \sum_{i=0}^{group} \sum_{j=0}^{time} D(M_{T,(i,j)}, M_{S,(i,j)}^q) \quad (5)$$

$$D(a, b) = \tau_a \tau_b \times KL(\log\text{-softmax}(a/\tau_a), \text{softmax}(b/\tau_b)), \quad (6)$$

where *group*, *time* are the total number of groups in a ResNet model, timesteps respectively, M_T , M_S^q are extracted membrane potential from teacher and quantized student model, KL is KL divergence, and τ_a , τ_b are temperature parameters of a , b in KL divergence. The decision for group-wise distillation is discussed in Section VI-A. The grand goal of L_{mem} is to minimize the distribution similarities between the teacher and student model. The final loss L for training the student model is shown below.

$$L = \alpha L_{CE} + \beta L_{logit} + \gamma L_{mem} \quad (7)$$

$$L_{CE} = \text{Cross-Entropy}(z_s^q) \quad (8)$$

$$L_{logit} = D(z_t, z_s^q). \quad (9)$$




Fig. 3: Hardware efficiency comparison of floating-point and MD-SNN approaches on the N-Caltech101 dataset. Lower EDAP values indicate better overall hardware efficiency, while higher TOPS/W and TOPS/mm² values indicate better energy and area efficiency, respectively. FP32 refers to full-precision, MS-N denotes MD-SNN with N-bit quantization for both weights and membrane potentials.

Here, L_{CE} , L_{logit} represent the Cross-Entropy loss of the student model and logit distillation loss respectively, z_t, z_s^q are the logits of the teacher model and quantized student model respectively. The α , β , and γ are the hyperparameters of each component of the loss function.

C. Teacher Versatility

A distinctive advantage of MD-SNN is exploiting SNN’s temporal computation for efficient knowledge transfer. Since a teacher model trained with T timesteps inherently computes membrane potentials at all intermediate timesteps $t \in 1, 2, \dots, T$, a single teacher can guide multiple students with different temporal configurations. As shown in Fig. 2(b), we extract membrane potentials at timestep t from a $T = 4$ teacher to supervise students configured for $t \leq T$ timesteps. This one-to-many paradigm eliminates training separate teachers for each configuration. The efficiency gain is substantial: instead of $2N$ model trainings (N teachers + N students), our approach requires only $N+1$ trainings (1 teacher + N students), significantly reducing training costs.

V. EXPERIMENTS

A. Experimental Setup

Software Evaluation Setup: We verify our proposed method on both static and dynamic datasets, e.g., CIFAR10/100, N-Caltech101, and TinyImageNet with ResNet architectures. Our baseline of SNN-based ResNet architecture is MS-ResNet [33] and we use direct coding to convert a float pixel value into binary spikes [34]. While the teacher model is trained with FP values, the weights and membrane potentials of students are quantized to 4 or 8-bit. We also use 8-bit BN parameters in the student model. We set α, β, γ in Eqn. (9) as 1, 3, 1 respectively. We set the number of epochs as 200 for all datasets. The CIFAR10/100, N-Caltech experiments are conducted on RTX2080ti GPUs and the TinyImageNet results are obtained from A100 GPU. For inference evaluation, we use the following hardware evaluation method.

Hardware Evaluation Setup: All the baselines and MD-SNN, optimized SNNs are evaluated on SpikeSim, an end-to-end inference evaluation platform based on in-memory computing architectures for SNNs. We use the Spikesim code available




Fig. 4: Teacher versatility on CIFAR-10. (a) Accuracy improvements of MD-SNN over quantized baselines using a single $T = 4$ teacher for all timestep configurations. Quantized baselines (dashed line) are models trained with MINT [19] for $t = 1, 2, 3, 4$ timesteps individually without any distillation. MD-SNN results (solid line) undergo distillation. (b) Training FLOPs comparison showing 30% reduction with the versatile teacher approach. Traditional refers to individual timestep teacher-student distillation (see Fig. 2(b)).

in [32] to benchmark the baseline as well as the MD-SNN model’s energy, delay, area and performance metrics. Tab. III shows the detailed hardware parameters used to initialize the SpikeSim framework.

Note that the results shown in the tables and figures in the below sections for MD-SNN correspond to the final student model’s accuracy or energy obtained with membrane potential distillation.

TABLE III: Hardware Implementation Parameters.

Technology	32nm CMOS
Crossbar Size	64
Device	4-bit FeFET ($\sigma/\mu=20\%$)
R_{off}/R_{on}	10 at $R_{on}=20k\Omega$
V_{DD} & V_{read}	0.9V & 0.1V

B. Accuracy and Energy-efficiency Results

Software Results: Overall performance results of CIFAR100, N-Caltech101, and TinyImageNet are shown in Tab. I. We use ResNet-18 for CIFAR100 and N-Caltech101, and ResNet-34 for TinyImageNet. The conventional FP models are our comparison baseline. On CIFAR100 and N-Caltech101 datasets, our proposed distilled model’s accuracies with 4 and 8 bits are higher than the results of the FP model. In the large dataset, e.g., TinyImageNet, our distilled accuracy achieves competitive performance with the accuracy of the FP architecture.

Hardware Results: Fig. 3 provides the SpikeSim evaluation, demonstrating significant efficiency gains on the N-Caltech101 dataset. MD-SNN with 4-bit quantization achieves $14.85\times$ lower energy-delay-area product (EDAP), $2.64\times$ higher TOPS/W and $6.19\times$ higher TOPS/mm² compared to 32-bit baselines at iso-accuracy. The 8-bit variants show intermediate improvements. Notably, our 4-bit models significantly surpass FP models’ efficiency while maintaining higher accuracy, validating that membrane distillation effectively compensates for aggressive quantization.

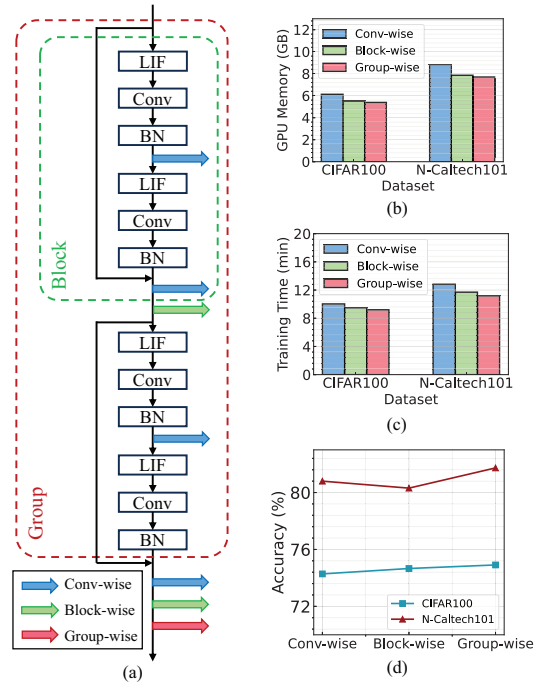


Fig. 5: Ablation study on membrane distillation granularity. (a) Three extraction strategies for membrane potential: Conv-wise (after each Conv-BN layer), Block-wise (after each residual block), and Group-wise (after each stage). (b) GPU memory consumption and (c) training time comparison across datasets. (d) Accuracy comparison on CIFAR-100 and N-Caltech101.

C. Comparison to Previous Works

To validate the performance of MD-SNN, we compare our work to previous quantization methods on SNNs, including STBP-Quant [29], ST-Quant [30], ADMM-Quant [31], SpikeSim [32], and MINT [19] shown in Tab. II. We compare the results of quantized SNN architecture on CIFAR10, CIFAR100, and TinyImageNet datasets with different bit-width of weight and membrane potential. Across all datasets, our MD-SNN achieves higher or comparable accuracy than other prior quantized SNN works.

D. Effectiveness of Versatile Teacher

Fig. 4 demonstrates the versatile teacher framework on CIFAR-10. A single $T = 4$ teacher successfully guides students at all timesteps ($t = 1, 2, 3, 4$), achieving 0.48-1.06% accuracy improvement over quantized models while reducing the total training FLOPs by 30%. The efficiency gain comes from reusing intermediate membrane potentials. The teacher transfers the knowledge for all $t \leq 4$ during forward propagation, eliminating the need for separate teachers per timestep configuration.

VI. ABLATION STUDY

A. Group-wise Membrane Distillation

Given the presence of membrane potential before every LIF neuron, selecting the extraction point for membrane potential distillation is crucial. We propose three distinct approaches: Conv-wise, Block-wise, and Group-wise, as depicted

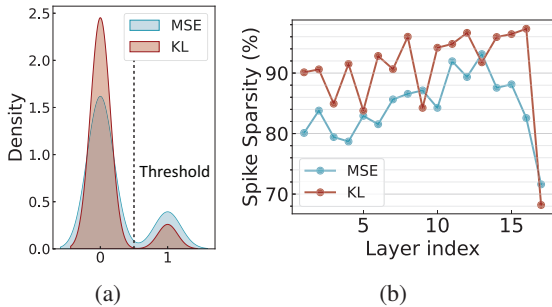


Fig. 6: Spike Sparsity and distribution of membrane potential of a student model trained on N-Caltech101 dataset on ResNet18 architecture. (a) Membrane potential distribution of a student model trained with KL-divergence and MSE distillation loss. The dashed line represents the threshold V_{th} . (b) The sparsity according to the layer index.

TABLE IV: The accuracy and sparsity comparison for varying loss methods on membrane distillation. X/Y represent the results when the weights and membrane potential are quantized to 8-bit/4-bit respectively.

Dataset	Loss	Accuracy (%)	Sparsity (%)
CIFAR100	CE+Logit+KL	74.91 / 74.41	83.08 / 81.51
	CE+Logit+MSE	73.83 / 73.09	80.00 / 77.78
N-Caltech101	CE+Logit+KL	81.73 / 81.18	89.64 / 90.33
	CE+Logit+MSE	80.74 / 78.01	83.09 / 82.61

in Fig. 5(a). In the Conv-wise approach, we extract the membrane potential after every Convolutional (Conv) and BN layer, transferring this knowledge to the membrane potential of the student model. Alternatively, the block-wise approach involves using the membrane potential immediately after each block, which comprises two Conv layers, two BN layers, and LIF neurons. Lastly, we introduce the group-wise approach, where a group consists of multiple blocks. For example, in ResNet34, a group may consist of 3, 4, or 6 blocks, and the membrane potential knowledge is transferred after each group.

To select the optimal membrane potential distillation method, we analyze GPU memory, training time, and accuracy on CIFAR-100 and N-Caltech101 datasets, as depicted in Fig. 5(b)-(d). Group-wise distillation proves most efficient in both GPU memory and training time. Moreover, it achieves the highest accuracy compared to Conv-wise and Block-wise methods. This highlights the importance of balanced knowledge transfer from the teacher model, as excessive teaching can degrade student model performance [35].

B. KL Divergence vs. MSE for Membrane Distillation

We investigate the choice of loss function for membrane potential distillation between teacher and student models. Unlike conventional ANN feature distillation that employs MSE [27], [36], we adopt KL divergence for membrane distribution alignment. Table IV compares the effectiveness of both approaches on CIFAR-100 and N-Caltech101 datasets, demonstrating superior accuracy when using KL divergence over MSE.

Beyond accuracy improvements, KL divergence also enhances spike efficiency. We measure spike sparsity as the ratio

TABLE V: Impact of distillation components on teacher versatility.

Method	T=1	T=2	T=3	T=4
Quantized (no distillation)	91.80	93.07	93.47	93.51
Logit-only	92.15	93.21	93.55	93.78
Logit+Membrane (Ours)	92.86	93.73	93.91	94.05

of total spikes across T timesteps to the total number of neurons in the network. Models trained with KL divergence consistently achieve higher sparsity than those using MSE. This difference is evident in the membrane potential distributions shown in Fig. 6(a). MSE produces a larger density area above the firing threshold compared to KL divergence, indicating excessive spike generation. Consequently, as illustrated in Fig. 6(b), KL divergence maintains higher spike sparsity across most network layers. This distinction is crucial: MSE loss generates redundant spikes that degrade both computational efficiency and model performance, while KL divergence preserves the sparse, event-driven nature essential for efficient SNN computation.

C. Distillation Components

To validate the importance of membrane distillation for teacher versatility, we compare different distillation strategies in Tab. V. Using a single $T = 4$ teacher to guide students at various timesteps on CIFAR-10 with 4-bit quantized weights and membrane potentials, we evaluate three configurations: quantized model without distillation, logit-only distillation, and our complete MD-SNN (logit+membrane distillation). Full usage of logit and membrane distillation achieves higher improvements across all timesteps compared to logit-only distillation. This demonstrates that membrane potentials, not just logits, enable effective knowledge transfer from a versatile teacher to temporally-diverse students.

VII. CONCLUSION

We present MD-SNN, the first knowledge distillation framework that leverages membrane potential dynamics to guide quantized SNN training. By transferring membrane distributions from full-precision teachers to quantized students using KL divergence, our method addresses the critical accuracy degradation problem in SNN quantization. MD-SNN achieves competitive performance over FP baselines across CIFAR-10/100, N-Caltech101, and TinyImageNet datasets, while delivering better efficiency on hardware. Furthermore, the versatile teacher framework, where a single teacher trained at $T = 4$ timesteps can effectively guide students at any $t \leq 4$, reduces training costs by 30% while maintaining superior performance at every operating point. This enables flexible deployment across different latency-accuracy requirements without retraining multiple models.

ACKNOWLEDGMENT

This work was supported in part by CoCoSys, a JUMP2.0 center sponsored by DARPA and SRC, the National Science Foundation (CAREER Award, Grant #2312366, Grant #2318152), the DARPA Young Faculty Award, the DoE MMICC center SEA-CROGS (Award #DE-SC0023198) and the Global Industrial Technology Cooperation Center (GITCC) program.

REFERENCES

- [1] Wolfgang Maass. Networks of spiking neurons: the third generation of neural network models. *Neural networks*, 10(9):1659–1671, 1997.
- [2] Kaushik Roy, Akhilesh Jaiswal, and Priyadarshini Panda. Towards spike-based machine intelligence with neuromorphic computing. *Nature*, 575(7784):607–617, 2019.
- [3] Filipp Akopyan, Jun Sawada, Andrew Cassidy, Rodrigo Alvarez-Icaza, John Arthur, Paul Merolla, Nabil Imam, Yutaka Nakamura, Pallab Datta, Gi-Joon Nam, et al. 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, 2015.
- [4] Mike Davies, Narayan Srinivasa, Tsung-Han Lin, Gautham Chinya, Yongqiang Cao, Sri Harsha Choday, Georgios Dimou, Prasad Joshi, Nabil Imam, Shweta Jain, et al. Loihi: A neuromorphic manycore processor with on-chip learning. *Ieee Micro*, 38(1):82–99, 2018.
- [5] Ruokai Yin, Abhishek Moitra, Abhiroop Bhattacharjee, Youngeun Kim, and Priyadarshini Panda. Sata: Sparsity-aware training accelerator for spiking neural networks. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 2022.
- [6] Ruokai Yin, Youngeun Kim, Di Wu, and Priyadarshini Panda. Loas: Fully temporal-parallel dataflow for dual-sparse spiking neural networks. In *2024 57th IEEE/ACM International Symposium on Microarchitecture (MICRO)*, pages 1107–1121. IEEE, 2024.
- [7] Congyang Li, Nabil Imam, and Rajit Manohar. A deterministic neuromorphic architecture with scalable time synchronization. *Nature Communications*, 16(1):10329, 2025.
- [8] Peter U Diehl et al. Fast-classifying, high-accuracy spiking deep networks through weight and threshold balancing. In *2015 International joint conference on neural networks (IJCNN)*, pages 1–8. iee, 2015.
- [9] Bing Han and Kaushik Roy. Deep spiking neural network: Energy efficiency through time based coding. In *European Conference on Computer Vision*, pages 388–404. Springer, 2020.
- [10] Yuhang Li, Shikuang Deng, Xin Dong, Ruihao Gong, and Shi Gu. A free lunch from ann: Towards efficient, accurate spiking neural networks calibration. In *International conference on machine learning*, pages 6316–6325. PMLR, 2021.
- [11] Shikuang Deng and Shi Gu. Optimal conversion of conventional artificial neural networks to spiking neural networks. *arXiv preprint arXiv:2103.00476*, 2021.
- [12] Tong Bu, Wei Fang, Jianhao Ding, PengLin Dai, Zhaofei Yu, and Tiejun Huang. Optimal ann-snn conversion for high-accuracy and ultra-low-latency spiking neural networks. *arXiv preprint arXiv:2303.04347*, 2023.
- [13] Yujie Wu, Lei Deng, Guoqi Li, Jun Zhu, and Luping Shi. Spatio-temporal backpropagation for training high-performance spiking neural networks. *Frontiers in neuroscience*, 12:331, 2018.
- [14] Sumit B Shrestha and Garrick Orchard. Slayer: Spike layer error reassignment in time. *Advances in neural information processing systems*, 31, 2018.
- [15] Hanle Zheng, Yujie Wu, Lei Deng, Yifan Hu, and Guoqi Li. Going deeper with directly-trained larger spiking neural networks. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pages 11062–11070, 2021.
- [16] Donghyun Lee, Yuhang Li, Youngeun Kim, Shiting Xiao, and Priyadarshini Panda. Spiking transformer with spatial-temporal attention. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 13948–13958, 2025.
- [17] Rachmad Vidya Wicaksana Putra and Muhammad Shafique. Q-spinn: A framework for quantizing spiking neural networks. In *2021 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2021.
- [18] Chen Li, Lei Ma, and Steve Furber. Quantization framework for fast spiking neural networks. *Frontiers in Neuroscience*, 16:918793, 2022.
- [19] Ruokai Yin, Yuhang Li, Abhishek Moitra, and Priyadarshini Panda. Mint: Multiplier-less integer quantization for energy efficient spiking neural networks. In *2024 29th Asia and South Pacific Design Automation Conference (ASP-DAC)*, pages 830–835. IEEE, 2024.
- [20] Ruokai Yin, Youngeun Kim, Yuhang Li, Abhishek Moitra, Nitin Satpute, Anna Hambitzer, and Priyadarshini Panda. Workload-balanced pruning for sparse spiking neural networks. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 8(4):2897–2907, 2024.
- [21] Ravi Kumar Kushawaha, Saurabh Kumar, Biplab Banerjee, and Rajbabu Velmurugan. Distilling spikes: Knowledge distillation in spiking neural networks. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 4536–4543. IEEE, 2021.
- [22] Qi Xu et al. Constructing deep spiking neural networks from artificial neural networks with knowledge distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7886–7895, 2023.
- [23] Yufei Guo, Weihang Peng, Yuanpei Chen, Liwen Zhang, Xiaode Liu, Xuhui Huang, and Zhe Ma. Joint a-snn: Joint training of artificial and spiking neural networks via self-distillation and weight factorization. *Pattern Recognition*, 142:109639, 2023.
- [24] Donghyun Lee, Ruokai Yin, Youngeun Kim, Abhishek Moitra, Yuhang Li, and Priyadarshini Panda. Tt-snn: Tensor train decomposition for efficient spiking neural network training. *arXiv preprint arXiv:2401.08001*, 2024.
- [25] Youngeun Kim, Yuhang Li, Hyoungseob Park, Yeshwanth Venkatesha, Ruokai Yin, and Priyadarshini Panda. Exploring lottery ticket hypothesis in spiking neural networks. In *European Conference on Computer Vision*, pages 102–120. Springer, 2022.
- [26] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.
- [27] Byeongho Heo, Jeessoo Kim, Sangdoon Yun, Hyojin Park, Nojun Kwak, and Jin Young Choi. A comprehensive overhaul of feature distillation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1921–1930, 2019.
- [28] Yukuan Yang, Lei Deng, Shuang Wu, Tianyi Yan, Yuan Xie, and Guoqi Li. Training high-performance and large-scale deep neural networks with full 8-bit integers. *Neural Networks*, 125:70–82, 2020.
- [29] Pai-Yu Tan and Cheng-Wen Wu. A low-bitwidth integer-stbp algorithm for efficient training and inference of spiking neural networks. In *Proceedings of the 28th Asia and South Pacific Design Automation Conference*, pages 651–656, 2023.
- [30] Sayeed Shafayet Chowdhury, Isha Garg, and Kaushik Roy. Spatio-temporal pruning and quantization for low-latency spiking neural networks. In *2021 International Joint Conference on Neural Networks (IJCNN)*, pages 1–9. IEEE, 2021.
- [31] Lei Deng, Yujie Wu, Yifan Hu, Ling Liang, Guoqi Li, Xing Hu, Yufei Ding, Peng Li, and Yuan Xie. Comprehensive snn compression using admm optimization and activity regularization. *IEEE transactions on neural networks and learning systems*, 2021.
- [32] Abhishek Moitra, Abhiroop Bhattacharjee, Runcong Kuang, Gokul Krishnan, Yu Cao, and Priyadarshini Panda. Spikesim: An end-to-end compute-in-memory hardware evaluation tool for benchmarking spiking neural networks. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 2023.
- [33] Yifan Hu, Lei Deng, Yujie Wu, Man Yao, and Guoqi Li. Advancing spiking neural networks towards deep residual learning. *arXiv preprint arXiv:2112.08954*, 2021.
- [34] Yujie Wu, Lei Deng, Guoqi Li, Jun Zhu, Yuan Xie, and Luping Shi. Direct training for spiking neural networks: Faster, larger, better. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 1311–1318, 2019.
- [35] Seyed Iman Mirzadeh, Mehrdad Farajtabar, Ang Li, Nir Levine, Akihiro Matsukawa, and Hassan Ghasemzadeh. Improved knowledge distillation via teacher assistant. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 5191–5198, 2020.
- [36] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chas-sang, Carlo Gatta, and Yoshua Bengio. Fitnets: Hints for thin deep nets. *arXiv preprint arXiv:1412.6550*, 2014.