

TSkips: Efficiency Through Explicit Temporal Delay Connections in Spiking Neural Networks

Anonymous authors

Paper under double-blind review

Abstract

Spiking Neural Networks (SNNs) with their bio-inspired Leaky Integrate-and-Fire (LIF) neurons inherently capture temporal information. This makes them well-suited for sequential tasks like processing event-based data from Dynamic Vision Sensors (DVS) and event-based speech tasks. Harnessing the temporal capabilities of SNNs requires mitigating vanishing spikes during training, capturing spatio-temporal patterns and enhancing precise spike timing. To address these challenges, we propose *TSkips*, augmenting SNN architectures with forward and backward skip connections that incorporate explicit temporal delays. These connections capture long-term spatio-temporal dependencies and facilitate better spike flow over long sequences. The introduction of *TSkips* creates a vast search space of possible configurations, encompassing skip positions and time delay values. To efficiently navigate this search space, this work leverages training-free Neural Architecture Search (NAS) to identify optimal network structures and corresponding delays. We demonstrate the effectiveness of our approach on four event-based datasets: DSEC-flow for optical flow estimation, DVS128 Gesture for hand gesture recognition and Spiking Heidelberg Digits (SHD) and Spiking Speech Commands (SSC) for speech recognition. Our method achieves significant improvements across these datasets: up to 18% reduction in Average Endpoint Error (AEE) on DSEC-flow, 8% increase in classification accuracy on DVS128 Gesture, and up to $\sim 8\%$ and $\sim 16\%$ higher classification accuracy on SHD and SSC, respectively.

1 Introduction

Spiking Neural Networks (SNNs) are a class of bio-physically realistic models with Leaky Integrate-and-Fire (LIF) (Abbott, 1999) neurons that offer a promising alternative for processing sequential data. By encoding information in the precise timing of spikes, SNNs capture spatio-temporal dependencies (Rathi & Roy, 2020; Chowdhury et al., 2021) through the membrane potential of their LIF neurons. This intrinsic ability to process temporal information allows SNNs to excel at capturing dynamic patterns, effectively functioning as specialized RNNs (Rathi & Roy, 2024) with reduced complexity.

The ability of SNNs to leverage the precise timing of spikes offers a distinct advantage over traditional Artificial Neural Networks (ANNs) for sequential tasks. Note, ANNs are “stateless” in nature (Hagenaars et al., 2021) and their reliance on specially curated input encoding schemes (Zhu et al., 2018b; Paredes-Vallés & De Croon, 2021) hinder their ability to accurately capture precise timing information crucial in event streams, as generated by Dynamic Vision Sensors (DVS) (Lichtsteiner et al., 2008; Brandli et al., 2014). Standard Recurrent Neural Networks (vRNNs) (Rumelhart et al., 1986) and Long Short-Term Memory (LSTM) networks (Hochreiter, 1997; Ponghiran & Roy, 2021) attempt to address this, but remain difficult to train (Pascanu et al., 2013) and computationally expensive. Furthermore, transformers and spiking transformers (Li et al., 2022; Zhou et al., 2023) excel at sequential processing but come at the cost of significantly larger models and computational overhead.

Deep SNNs offer the appealing combination of rich temporal information processing and computational efficiency, but training them effectively presents unique challenges (Rathi et al., 2021). The combination of multiple time steps with Back Propagation through Time (BPTT), non-differentiable spikes and sparse event

data can worsen the vanishing spike problem (Cao et al., 2015). To circumvent this, previous works have proposed adaptive LIF neurons with trainable threshold and leak (Kosta & Roy, 2023) and approximate surrogate gradients (Fang et al., 2021; Neftci et al., 2019). Other works (Shrestha & Orchard, 2018; Sun et al., 2023b; Hammouamri et al., 2023) have considered synaptic delays as an additional parameter alongside weights, thus offering a richer representation of the spike patterns present in event data. We posit that these learnable synaptic delays only capture spatio-temporal patterns over a short window of time. Note, however, event data for speech recognition and optical flow also have long-term temporal dependencies that need to be captured through better spike representation.

In contrast, this work introduces explicit temporal delays in forward and backward skip connections, named temporal skips (*TSkips*), for SNNs and hybrid ANN-SNN (Negi et al., 2024) models. *TSkips* offer finer control over spike timing, enhance responsiveness to temporal patterns, mitigate vanishing spikes, all while capturing long-term spatio-temporal dependencies for event-based data. In addition, our method utilizes adaptive LIF neurons (Kosta & Roy, 2023) that have learnable leaks and thresholds that capture local temporal patterns.

The next question to address would be, what are the optimal *TSkips* configurations within the network architecture and what are the corresponding temporal delays associated with them? To efficiently identify optimal *TSkips* configurations across different architectures and sequence lengths, we leverage training-free Neural Architecture Search (NAS) tailored for SNNs (Kim et al., 2022). *TSkips* introduces minimal increase in model size with very few additional trainable parameters, minimal overhead and improves inference on event-based tasks. In addition, we have observed that training time is faster for the proposed networks compared to standard SNNs, RNNs and LSTMs. The contributions of this work can be summarized as follows:

- We introduce *TSkips*, a novel mechanism that enhances SNNs and hybrid ANN-SNN architectures by enabling direct transmission of spike information between non-adjacent layers with temporal delays, effectively capturing long-term spatio-temporal patterns. (Section 3)
- We analyzed various *TSkips* configurations, revealing the complexities of the search space, particularly the interplay between the temporal delay (Δt), *TSkips* position and network depth. This complexity motivated the use of NAS to efficiently identify optimal *TSkips* architectures. (Section 4.2)
- We show that *TSkips* achieves strong performance across sequential tasks. On the DSEC-flow dataset, they reduce average endpoint error (AEE) and inference energy, up to 18% and 56%, respectively, as shown in Fig. 1. Furthermore, *TSkips* achieves significant accuracy improvements on DVS128 Gesture, SHD and SSC datasets — up to 8.3%, 8.6% and 16.04%, respectively. This highlights the scalability and effectiveness of our approach across diverse sequential tasks and architectures. (Section 4)

2 Related Work

This section explores the challenges of training deep SNNs and incorporating delays to capture richer temporal dynamics.




Figure 1: Model size vs. average endpoint error (AEE) for fully spiking EV-FlowNet architectures (Kosta & Roy, 2023) modified with Forward and Backward *TSkips* on the DSEC-flow dataset (lower AEE is better). Each tick on the horizontal axis represents a $3\times$ increase in model size. *TSkips* achieve an average reduction of 15% in AEE and 40.75% in inference energy (E) while maintaining comparable AEE to larger baselines.

2.1 Training Deep SNNs

Early methods to train SNNs relied on ANN-to-SNN conversion (Cao et al., 2015; Panda et al., 2020). However, this approach often struggled to match ANN performance on complex sequential tasks (Deng et al., 2020), particularly with event data. A key obstacle in training deep SNNs is the non-differentiable nature of spike activations (Fang et al., 2021), which prevents the direct application of gradient-based optimization techniques, leading to subpar performance.

To overcome this, approximate surrogate gradients (Fang et al., 2021; Neftci et al., 2019) have been proposed, that replace non-differentiable spike activations with smooth gradient approximations, enabling the use of BPTT (Werbos, 1990) to train deep SNNs. However, this approach introduces other challenges like vanishing gradients, especially with long sequences and sparse event data (Cao et al., 2015), and reduce accuracy compared to ANNs (due to approximate gradients). To address these challenges, previous work has investigated two main avenues: adaptive LIF neurons with learnable thresholds and leaks (Kosta & Roy, 2023), and traditional skip connections (without delays) in SNNs (Benmeziane et al., 2023). While adaptive LIF neurons allow for dynamic adjustment and better capture of spatio-temporal patterns, and skip connections can improve accuracy and efficiency, neither approach fully addresses the challenges of representing and capturing complex temporal patterns. To achieve this, we propose *TSkips*, incorporating explicit temporal delays into the network architecture with forward and backward skip connections.

2.2 Delays in SNNs

The importance of precise timing information in event-based data, such as that from dynamic vision sensors, has motivated research into effectively capturing temporal dynamics in SNNs (Rathi & Roy, 2020; Chowdhury et al., 2021). Some works have introduced delays alongside the weights in SNNs to offer a richer representation of the spike patterns. For instance, Shrestha & Orchard (2018) introduced a fixed delay in SNNs to improve learning of temporal complexities, Sun et al. (2023b) extended this by making these delays learnable and introducing a layer-specific maximum delay. Hammouamri et al. (2023) further refined this by making the per-layer maximum delay learnable and utilizing Dilated Convolutions with Learnable Spacings (DCLS) (Khalfaoui-Hassani et al., 2023) to capture spatio-temporal patterns.

However, learning these delays, as in Sun et al. (2023b); Hammouamri et al. (2023), introduces a new set of parameters with the same dimension as the weights, essentially doubling the number of trainable parameters. Additionally, these methods only capture local temporal patterns over a short window of time, due to their reliance on the dynamics of LIF neurons.

In contrast to these methods that rely on learning delays, our proposed method introduces explicit temporal delays in the form of forward and backward skip connections (*TSkips*) within the SNN architecture. The use of *TSkips* with adaptive LIF neurons effectively captures both local and long-term spatio-temporal patterns with minimal additional computational overhead.

3 Methodology

This section details our methodology, starting with the input representation and the neuron model employed in our SNNs, followed by the proposed method, our NAS search space, and the training process.

3.1 Input Representation and Sensors

For efficient computation, SNNs require data in the form of discrete spikes, mirroring communication in biological neurons. DVS (Lichtsteiner et al., 2008) cameras output asynchronous spikes representing real-time changes in pixel intensity (I_t). This eliminates redundant data transmission and achieves higher temporal resolution than traditional frame-based cameras or rate coding methods (Cao et al., 2015). These cameras trigger events when the log intensity change at a pixel crosses a threshold (θ) (Gallego et al., 2022), represented as $||\log(I_t) - \log(I_{t-1})|| \geq \theta$. Events are encoded in the Address Event Representation (AER) format as tuples (x, y, t, p) — representing pixel coordinates, timestamp, and polarity. Similarly, event-based audio

datasets (Cramer et al., 2022) mimic human auditory spiking activity, with spikes represented as (x, t) — denoting spike unit and timestamp.

3.2 Neuron Model

The fundamental computational unit of SNNs is the spiking neuron. We employ Leaky Integrate-and-Fire (LIF) (Abbott, 1999) neurons for their biological plausibility and computational efficiency. The LIF neuron integrates input spike information over time, accumulating it in its membrane potential. This potential gradually decays, allowing controlled forgetting of less relevant information. The dynamics of the LIF neuron can be described by

$$U_l^t = \lambda U_l^{t-1} + W_l O_{l-1}^t - V_l^{th} O_l^{t-1} \quad (1)$$

where U_l^t represents the membrane potential of the neurons in layer l at time step t , λ is the leak factor, controlling the decay rate of the membrane potential. W_l is the weight matrix for neurons in layer l , V_l^{th} is the voltage threshold of layer l and O_l^t is the output generated by layer l at time step t .

When the membrane potential surpasses the voltage threshold, the neuron emits a spike. The first term in Eq. (1) denotes the leakage in the membrane potential, the second term computes the weighted summation of output spikes from layer $(l - 1)$ and the third term denotes the reduction in membrane potential when an output spike is generated at layer l . The membrane potential is then reset, either to be zero (hard reset) or by subtracting V_l^{th} (soft reset). Equation (2) describes the spike generation.




Figure 2: LIF Neuron

$$Z_l^t = \frac{U_l^t}{V_l^{th}} - 1, \quad O_l^t = \begin{cases} 1, & \text{if } Z_l^t > 0 \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

We use a hard reset for optical flow estimation and a soft reset for gesture and speech classification. Fig. 2 shows the LIF neuron dynamics.

3.3 Proposed Method

Efficiently training deep SNNs is often hindered by vanishing gradients (Cao et al., 2015), which impede accurate spike propagation across time steps. To address this, we propose a novel method that augments skip connections in SNNs by introducing explicit temporal delays (Δt) within the connections. We refer to these augmented skip connections, which can be forward, backward, or a combination of both, as *TSkips*. The duration of these explicit delays is constrained by $0 < t - \Delta t < T$, where t is the current time step and T is the sequence length of the data. Specifically, t represents the set of all discrete time steps within the sequence, denoted as $t = \{t_i, \forall i = 0, \dots, T\}$. This constraint on Δt ensures that the network does not receive future information.

To illustrate the information flow in a neural network with *TSkips*, let us consider how they modify the output of a layer. If h_l^t denotes the input to a layer l at time step t , then a forward *TSkips* provides input h_{l-k}^t and a backward *TSkips* provides input h_{l+k}^t . For brevity, we denote these skip inputs as $h_{l\pm k}^t$, where k indicates the layer separation between layer l and the source of the skip connection. *TSkips* can be represented as:

$$h_l^t = f_l(h_{l-1}^t \oplus W_s h_{l\pm k}^{(t-\Delta t)}) \quad (3)$$

where f_l is the affine function of layer l , \oplus represents either concatenation (Huang et al., 2017) or element-wise addition (He et al., 2015) and W_s is a fixed shortcut path matrix (He et al., 2016) that randomly selects channels in $h_{l\pm k}^{(t-\Delta t)}$ to match the dimensions of h_l^t . Eq. 3 demonstrates that the output of layer l at time

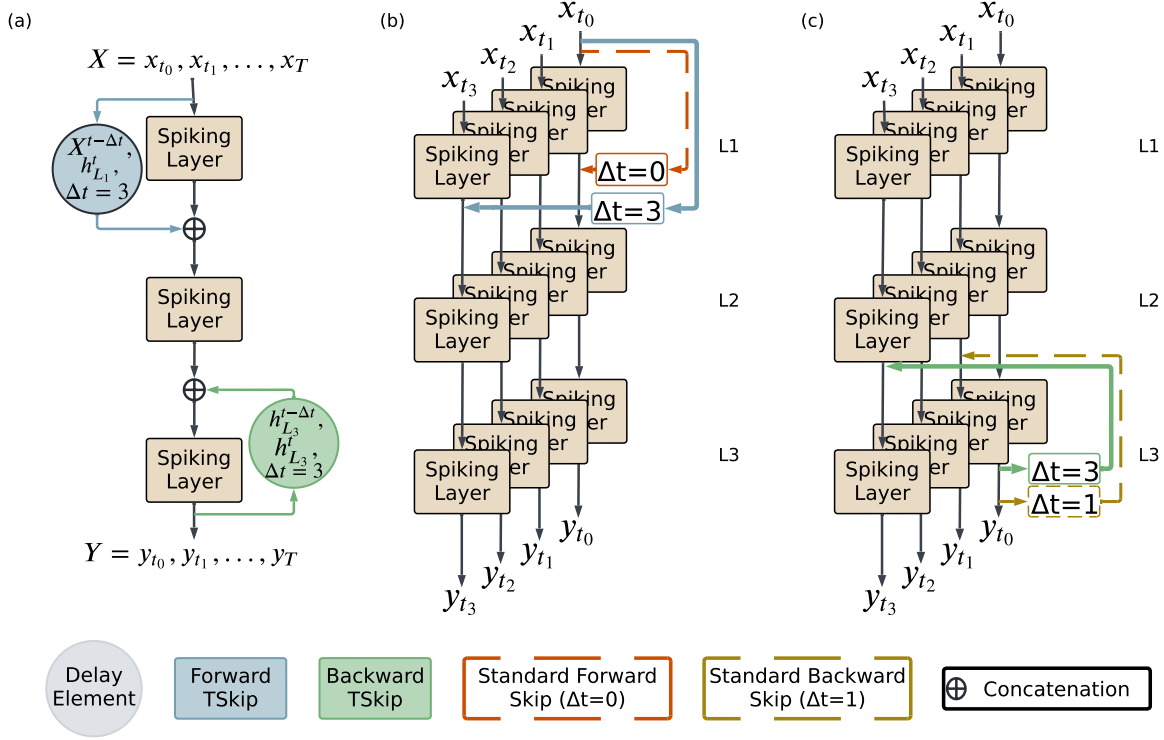


Figure 3: Illustration of $TSkip$ s in an SNN. (a) A 3-layer SNN with forward and backward $TSkip$ s, annotated with their origin and destination layers and the delay (Δt). (b) Unrolled SNN with $T = 4$ with a forward $TSkip$ s and $\Delta t = 3$ compared with a standard skip connection (no delay). (c) Unrolled SNN with $T = 4$ with a backward $TSkip$ s and $\Delta t = 3$ compared with a standard backward connection ($\Delta t = 1$) (Rumelhart et al., 1986)

step t is now influenced by the output of layer $l \pm k$ at time step $(t - \Delta t)$. Here, k represents the difference in layer indices between connected layers. For example, if $l = 4$ and $k = 2$, a forward $TSkip$ s connects layers 2 and 4 with delay Δt .

Our method, incorporating $TSkip$ s, is illustrated in Fig. 3. Fig. 3(a) shows a simple 3-layer SNN architecture, illustrating both forward and backward $TSkip$ s between adjacent layers. Each connection is annotated with the origin and destination layers and (Δt). Fig. 3(b) depicts the temporal unrolling of the network over $T = 4$ time steps. It contrasts a standard skip connection ($\Delta t = 0$) and a forward $TSkip$ s with $\Delta t = 3$. Fig. 3(c) similarly illustrates a backward $TSkip$ s with $\Delta t = 3$ and contrasts it with a standard backward connection that has a delay of $\Delta t = 1$, typically seen in vRNNs (Rumelhart et al., 1986). This comparison emphasizes that $TSkip$ s allow for more flexibility and can thus capture longer temporal dependencies compared to standard RNNs. This flexibility comes with minimal overhead, as $TSkip$ s are identity mappings that introduces very few additional parameters and no computational complexity to the network.

$TSkip$ s enhance the ability of SNNs and hybrid models to process sparse event streams by facilitating the propagation of temporally relevant information across non-adjacent layers. By capturing these long-term dependencies, $TSkip$ s can improve the network’s ability to learn complex temporal patterns. Furthermore, the explicit temporal delays in $TSkip$ s offer finer control over spike timing within the network. This finer control enhances the network’s responsiveness to temporal patterns by allowing it to selectively integrate information from different time steps. Additionally, by enabling the network to access relevant information from the past, $TSkip$ s can mitigate the issue of vanishing spikes. To further enhance temporal representation, we introduce a learnable scaling factor, α , that controls the ratio of data at the current time-step (t) and data from the delayed time step ($t - \Delta t$) within the $TSkip$ s. This is incorporated into the skip as shown in Eq. 4:

$$h_l^t = f_l(h_{l-1}^t \oplus W_s(\alpha h_{l \pm k}^t + (1 - \alpha)h_{l \pm k}^{(t-\Delta t)})) \quad (4)$$

3.4 Exploring the search space

Identifying optimal models with *TSkips* is challenging due to the vast number of potential architectures and associated *TSkips* configurations. To manage this complexity, we initialize our search with a backbone network (Cai et al., 2020). This backbone provides a starting point for exploration, reducing the number of possible configurations and allowing the search to focus on optimizing the specific parameters of our proposed method. However, significant complexity remains due to the numerous choices for Δt , connection placement, and network depth.

To illustrate this complexity, consider a 3-layer network with the possibility of forward and backward connections between any two layers (Fig. 4). This simple example results in 12 possible combinations of skip connections. With a sequence length of $T = 6$, this grows to 120 possible delay combinations and 2^{120} total possible *TSkips* configurations. This combinatorial problem, grows exponentially with the sequence lengths, as each additional time step introduces new *TSkips* configurations.

This large search space necessitates an efficient exploration method to identify optimal architectures. To address this, we leverage NASWOT-SAH (Kim et al., 2022), a training-free NAS method tailored for SNNs. NASWOT-SAH scores networks at initialization using the Sparsity Aware Hamming Distance (SAHD), which quantifies the dissimilarity in spike patterns between different layers. By favoring networks with diverse spiking activity, SAHD acts as a proxy for identifying optimal sub networks of the backbone architecture augmented with *TSkips*. Specifically, (Kim et al., 2022) performs a random search for architectures with optimal Δt , *TSkips* positions (origin and destination) and network depth.

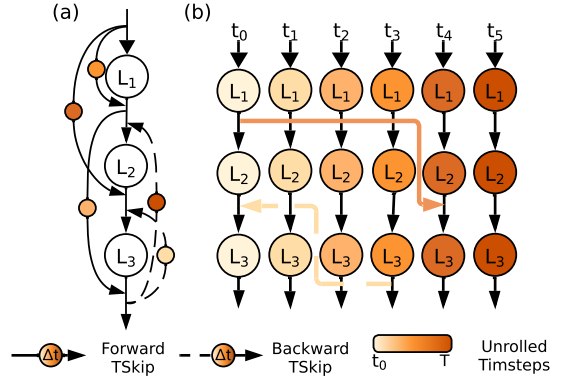


Figure 4: Illustration of the NAS Search Space. (a) Example of possible *TSkips* combinations in a 3-layer network with $\Delta t = 6$. (b) Top-ranked forward and backward *TSkips* with $\Delta t = 4$ and $\Delta t = 3$, respectively, found by (Kim et al., 2022). For visual clarity, the figure depicts the *TSkips* spanning a single delay window (Δt). Here t_0 refers to the first time step and T is the sequence length.

3.5 Training with Explicit Temporal Delays

Training SNNs and hybrid models with gradient descent requires addressing the non-differentiable nature of spike activations (Fang et al., 2021) in LIF neurons. We employ adaptive LIF neurons (Kosta & Roy, 2023) with learnable leaks and thresholds, and use the ArcTangent (Fang et al., 2021) surrogate gradient to approximate the derivative of the spike activation function. To capture precise spike timing and spatio-temporal dependencies, we unroll the network in time. This, combined with adaptive LIF neurons and *TSkips*, allows for explicit representation of the network’s temporal evolution, capturing both long-term and local temporal patterns. We then use back propagation through time (BPTT) (Werbos, 1990) for gradient propagation and Batch Normalization Through Time (BNTT) (Kim & Panda, 2021) for improved training stability.

4 Experiments

4.1 Experimental Setup

DSEC-Flow The DSEC-Flow dataset (Gehrig et al., 2021a) consists of 24 challenging driving sequences with varying lighting conditions, fast motion, and occlusions. To assess our models on unseen data, we created a custom test set by splitting the training set, following the approach in Ponghiran et al. (2023). To evaluate the accuracy of the predicted optical flow, we use average end point error (AEE) (Zhu et al., 2018a),

which measures the average Euclidean distance between the predicted and ground truth flow vectors.

$$AEE = \frac{1}{n} \sum_n \|(u, v)_{pred} - (u, v)_{gt}\|_2 \quad (5)$$

For experiments on DSEC-flow (Gehrig et al., 2021a), we use a fully spiking (Kosta & Roy, 2023) and hybrid (Negi et al., 2024) multi-scale encoder-decoder network inspired by EV-FlowNet (Zhu et al., 2018a) as our backbone architecture. We augment this architecture by integrating *TSkips* into the existing skip connections between the encoder and decoder blocks, as depicted in Fig. 5. Due to the inherent structure of our backbone, the search space for optimal *TSkips* configurations is relatively small, focusing primarily on selecting which existing skip connection to modify and determining the appropriate temporal delay (Δt).




Figure 5: Fully spiking EV-FlowNet (Zhu et al., 2018a) architecture with forward and backward *TSkips*. Each *TSkips* replaces a pre-existing skip connection between the encoder and decoder. Although shown together for conciseness, forward and backward connections are evaluated independently. In practice, when used together, they connect different encoder-decoder layers. This architecture can be modified to be a hybrid ANN-SNN model, by replacing the LIF neurons with ReLU activations.

DVS128-Gesture The IBM DVS128 Gesture dataset (Hu et al., 2022) contains 1342 instances of 11 hand gestures, captured with a DVS128 (Lichtsteiner et al., 2008) camera. For experiments on DVS128 Gesture, we use a ResNet18 (He et al., 2015) backbone architecture (Cai et al., 2020) and explore various configurations by modifying its basic blocks. This search includes modifying the input and output channels of each layer, kernel sizes, stride, and network depth. To integrate *TSkips* in the backbone, we remove the original ResNet18 skip connections and search for optimal *TSkips* placement and corresponding Δt .

For optical flow estimation and gesture recognition, we introduce a learnable scaling factor, α , within the *TSkips* to further enhance temporal processing (detailed in Section. 3.3).

SHD and SSC The SHD (Cramer et al., 2022) dataset comprises 10,000 recordings of spoken digits (zero to nine in English and German), while the larger SSC (Cramer et al., 2022) dataset contains 100,000 recordings of spoken words from various speakers. These datasets present challenges for audio classification, requiring the network to capture and process subtle temporal patterns within spike trains that encode spoken sounds. We focus on the top-one classification task for all 35 classes in SSC and all 20 classes in SHD. For our experiments on speech recognition, we use a fully connected Multi Layer Perceptron (MLP) (Rumelhart et al., 1986) as our backbone architecture. We explore various MLP configurations by adjusting the input and output features of each layer, the overall network depth, and the placement and Δt of *TSkips*. In these experiments, we omit α used in the EV-FlowNet and DVS128 Gesture *TSkips* architectures.

To improve generalization, we augment the training sets for DVS128 Gesture, SHD and SSC with channel jitter and random noise (Shen et al., 2023). We use sequence lengths (T) of 10, 30, and 99 for DSEC-flow, DVS128 Gesture, and SHD/SSC, respectively. All models are trained on an NVIDIA A40 GPU using the Adam optimizer (Kingma & Ba, 2014). For DSEC-flow, we use a multi-step learning rate scheduler (initial

rate: 10^{-3} , scaled by 0.7 every 10 epochs) and the supervised mean squared error (MSE) loss (Kosta & Roy, 2023; Negi et al., 2024) for 200 epochs. For all other datasets, we use a Cosine Annealing Learning Rate Scheduler (Loshchilov & Hutter, 2017) (initial rate: 0.001, minimum rate: 5×10^{-6} , updated every 10 iterations) for 100 epochs.

Details on the constraints on the search parameters and *TSkips* configurations for DSEC-flow, DVS128 Gesture, SHD, and SSC can be found in Appendix D.

4.2 Ablation Study




Figure 6: Ablation study results on the SHD dataset demonstrating the impact of varying the (a) temporal delay (Δt), (b) *TSkips* position and (c) network depth on classification accuracy.

As detailed in Section 3.4, the search space of possible *TSkips* configurations is exponentially large. To effectively navigate this complex search space, we performed an ablation study on an 8-layer MLP baseline found using Kim et al. (2022) on the SHD dataset. We varied the temporal delay (Δt), *TSkips* position and network depth using concatenation-based (Huang et al., 2017) *TSkips* for our analysis (see Appendix C.5 for detailed baseline network configuration).

Impact of Temporal Delay (Δt): We investigated the impact of varying Δt for both forward and backward *TSkips* (Fig. 6(a)). Forward connections connected the input data to the second layer, while backward connections connected the final output to the seventh layer, with a delay of Δt . Varying Δt significantly influenced accuracy, with larger delays generally leading to better classification. This highlights the importance of optimizing Δt relative to the sequence length (T), as both excessively small and large delays can hinder performance.

Impact of TSkips Position: Next, we investigated the impact of *TSkips* position on classification accuracy (Fig. 6(b)). Fixing $\Delta t = 16$, we varied the destination layer of both forward and backward connections. Forward connections originated from the input data, while backward connections originated from the final output. Our analysis revealed that connecting *TSkips* to deeper layers consistently improved accuracy. This suggests that propagating temporal information to deeper layers is crucial for enhancing model performance. Note, the optimal placement varies based on forward or backward *TSkips*.

Impact of Network Depth: We analyzed the impact of *TSkips* on models of varying depth (Fig. 6(c)), with forward connections originating from the input and backward connections originating from the final output with $\Delta t = 16$. Our analysis revealed that deeper networks benefited more from forward connections, while shallower architectures favored backward connections. This suggests the optimal configuration of *TSkips* and network depth are intertwined.

These findings underscore the complex interplay between Δt , *TSkips* position (origin and destination layers), and network depth, which we set as our search parameters. For the results presented in the following sections, we leveraged Kim et al. (2022) to identify optimal baseline and concatenation-based (Huang et al., 2017) *TSkips* architectures (see Appendix C.4 for comparisons with addition-based *TSkips*), constraining the architecture search by the number of parameters.




Figure 7: Qualitative results on the DSEC-flow dataset. (a) Ground Truth (GT) mask, (b) Mini (3.4M) SNN baseline model mask, (c) Mini SNN (3.4M) with forward (F) TSkips, (d) Mini SNN (3.4M) with backward (B) TSkips, and (e) Base (13M) SNN baseline model mask. As observed, the Mini model that is $3.8\times$ smaller than Base model performs just as well qualitatively when *TSkips* are incorporated.

4.3 DSEC-Flow Results

Incorporating *TSkips* into the fully spiking and hybrid EV-FlowNet architectures (Kosta & Roy, 2023; Negi et al., 2024), as detailed in Section 4.1, significantly lowers AEE on the DSEC-flow dataset (Table 1). Across both fully spiking and hybrid architectures, *TSkips* consistently yielded an average reduction in AEE of 14% and 9.5%, respectively. Notably, these improvements are achieved with models $3\times$ smaller than baselines with comparable AEE. While SOTA methods like E-RAFT (Gehrig et al., 2021b) and E-FlowFormer (Li et al., 2023) achieve 6% and 10% lower AEE, respectively, these reductions come at the cost of increased complexity. E-RAFT, for example, while being $3\times$ smaller than our best model, relies on complex mechanisms such as gated recurrent units and iterative flow refinement. Similarly, E-FlowFormer employs a computationally expensive transformer architecture and is trained on both a custom dataset and DSEC-flow. These complexities likely lead to higher inference energy compared to our lightweight *TSkips* architectures. Importantly, our approach achieves comparable qualitative performance to baselines $3\times$ larger (Fig. 7), highlighting its efficiency in enabling accurate temporal processing without increasing model complexity. Hyper parameters and a detailed analysis of convergence and inference energy for *TSkips* on DSEC-flow are in Appendix B.

4.4 DVS128-Gesture, SHD and SSC Results

Using *TSkips* in ResNet18 and MLP backbone architectures, as detailed in Section 4.1, consistently improved the classification accuracy of SNNs across the DVS Gesture, SHD, and SSC datasets (Tables 2 and 3).

On the DVS Gesture dataset, incorporating *TSkips* resulted in a significant 8.37% improvement in accuracy over the baseline model, a convolutional SNN found by Kim et al. (2022). Importantly, our approach achieved this with significantly smaller models compared to SOTA, including a spiking transformer (Qin & Liu, 2024), demonstrating the efficiency of our proposed method. Furthermore, our method outperformed several other approaches, including spiking RNNs (Xing et al., 2020), methods incorporating learnable delays (Shrestha & Orchard, 2018), and methods using specialized LIF neurons or spatio-temporal feature extraction (Jiang & Zhang, 2024; Samadzadeh et al., 2023). This highlights the effectiveness of our approach in capturing and leveraging temporal dependencies for improved gesture recognition.

On the SHD dataset, *TSkips* achieved an average accuracy gain of 8.6% over the 4-layer baseline (Baseline - 1) and 8.31% over the 8-layer baseline (Baseline - 2). Our approach also outperformed various recurrent architectures (vRNNs, spiking RNNs, and LSTMs) by substantial margins. Similarly, on the SSC dataset, *TSkips* yielded an average accuracy gain of 14% over Baseline - 1 and 16.04% over the 8-layer Baseline - 2, again surpassing the performance of recurrent architectures.

While *TSkips* achieves accuracy within 0.34% of SOTA model DCLS-Delays (Hammouamri et al., 2023) on SHD and 0.46% on SSC, it offers several advantages: applicability across various network architectures, as shown in Section 4.3, efficiency due to very few new trainable parameters, and strong performance. Notably *TSkips* achieves comparable accuracy to Hammouamri et al. (2023) on the SSC dataset with a 44% smaller model.

Table 1: AEE on DSEC-flow (lower is better). Best performing SNN and hybrid baselines and forward/backward *TSkips* models are highlighted in bold.

Architecture	#Params (M)	AEE	
		SNN	Hybrid
EV-FlowNet (Gehrig et al., 2021b)	13.04	2.32	-
LSTM-FlowNet (Ponghiran et al., 2023)	N/A ¹	1.28	-
E-RAFT (Gehrig et al., 2021b)	5.3	0.79	-
E-FlowFormer (Li et al., 2023)	N/A	0.76	-
Our Models			
Base - Baseline	13.04	1.35	0.96
Base + Forward <i>TSkips</i>		1.12	0.86
Base + Backward <i>TSkips</i>		1.13	0.84
Mini - Baseline	3.41	1.65	1.11
Mini + Forward <i>TSkips</i>		1.44	1.02
Mini + Backward <i>TSkips</i>		1.46	1.01
Micro - Baseline	0.93	1.80	1.22
Micro + Forward <i>TSkips</i>		1.57	1.12
Micro + Backward <i>TSkips</i>		1.56	1.09
Nano - Baseline	0.27	2.17	1.47
Nano + Forward <i>TSkips</i>		1.86	1.37
Nano + Backward <i>TSkips</i>		1.77	1.35
Pico - Baseline	0.092	2.57	2.02
Pico + Forward <i>TSkips</i>		2.19	1.78
Pico + Backward <i>TSkips</i>		2.28	1.81

Table 2: Classification accuracy on DVS-Gesture. Comparisons between forward/backward *TSkips* and SOTA models. Best models in bold.

Method	Base Model	#Params (M)	Accuracy (%)
SLAYER (Shrestha & Orchard, 2018)	SNN (8 layer)	N/A	93.6
KLIF (Jiang & Zhang, 2024)	VGG-11	N/A	93.75
DECOLE (Kaiser et al., 2020)	Custom	N/A	95.2
SNN-Skip (Benmeziene et al., 2023)	ResNet18	N/A	95.43
SCRNN (Xing et al., 2020)	Custom	N/A	96.59
STS-ResNet (Samadzadeh et al., 2023)	ResNet18	N/A	96.7
Mamba-Spike (Qin & Liu, 2024)	Mambda	6.1	97.8
Our Models			
Baseline		0.30	88.75
vRNN ²	ResNet18	0.30	89.12
Forward <i>TSkips</i>	Backbone	0.42	94.82
Backward <i>TSkips</i>		0.38	95.97
Forward + Backward <i>TSkips</i>		0.55	97.52

These improvements can be attributed to the ability of *TSkips* to capture long-term temporal dependencies, facilitated by the unrolled network structure, detailed in Section 3. Our method improves both the accuracy and convergence rate of the models by mitigating vanishing spikes and enabling more efficient gradient propagation during training. The analysis presented in Appendices A and C demonstrates lower energy

¹N/A - not reported #Params²Backward skips are added to all baseline layers with $\Delta t = 1$.

consumption and faster convergence compared to standard SNNs and recurrent architectures. Additionally, we show that incorporating *TSkips* in ANNs and hybrid models significantly improves speech classification (Appendix C.3).

Table 3: Classification accuracy on SHD and SSC datasets. Comparisons between forward, backward and a combination of both *TSkips* against SOTA models and two baseline models: a shallow 4-layer MLP network (Baseline - 1) and a deep 8-layer MLP (Baseline - 2). Best models with a single and two *TSkips* in bold.

Method	#Params (M)		Accuracy (%)	
	SHD	SSC	SHD	SSC
SOTA Models				
LSTM-LIF (Zhang et al., 2023)	0.14	0.11	88.91	63.46
SRNN (Yin et al., 2021)	N/A	N/A	90.4	74.2
DL256-SNN-DLoss (Sun et al., 2023a)	0.14	-	92.56	-
SpikGRU (Dampfhofer et al., 2022)	-	0.28	-	77
radLIF (Bittar & Garner, 2022)	3.9	3.9	94.62	77.40
Spikingformer (Zhou et al., 2023) ³	1.98	1.99	82.68	72.43
DCLS-Delays (Sun et al., 2023b)	0.21	2.5	95.07	80.69
Our Baseline MLPs				
Baseline - 1	0.16	0.12	84.32	64.19
Baseline - 2	1.04	1.04	86.42	67.54
Our MLP Models with <i>TSkips</i>				
vRNN ² - 1	0.16	0.12	68.5	71.2
Forward <i>TSkips</i> - 1	0.24	0.42	92.32	76.5
Backward <i>TSkips</i> - 1	0.19	0.24	93.64	79.87
Forward + Backward <i>TSkips</i> - 1	0.20	0.54	93.01	78.64
vRNN ² - 2	1.04	1.04	71.35	72.24
Forward <i>TSkips</i> - 2	1.12	1.08	94.15	78.98
Backward <i>TSkips</i> - 2	1.16	1.14	93.86	79.65
Forward + Backward <i>TSkips</i> - 2	1.28	1.42	94.73	80.23

5 Conclusion

In conclusion, this work demonstrates the potential of incorporating temporal delays in forward and backward skip connections across SNNs and hybrid architectures. Our proposed method, optimized through training-free NAS, enhances accuracy, temporal representation, and efficiency while maintaining small model sizes, as shown by our results on the DSEC-flow, DVS128 Gesture, SHD, and SSC datasets. Furthermore, our ablation studies reveal the importance of understanding the interplay between network depth, temporal delays, and connection placement in optimizing SNN performance. Overall, this work demonstrates a promising direction for the development of efficient and accurate SNN architectures, achieving competitive performance with reduced complexity, thus offering a compelling alternative for real-world applications requiring fast and accurate processing of temporal information.

³Zhou et al. (2023) was evaluated on SHD and SSC with minimal updates to the architecture and tuned learning rate.

References

- Larry F Abbott. Lapicque’s introduction of the integrate-and-fire model neuron (1907). *Brain research bulletin*, 50(5-6):303–304, 1999.
- Hadjer Benmeziiane, Amine Ziad Ounnoughene, Imane Hamzaoui, and Younes Bouhadjar. Skip connections in spiking neural networks: An analysis of their effect on network training, 2023. URL <https://arxiv.org/abs/2303.13563>.
- Alexandre Bittar and Philip N Garner. A surrogate gradient spiking baseline for speech command recognition. *Frontiers in Neuroscience*, 16:865897, 2022.
- Christian Brandli, Raphael Berner, Minhao Yang, Shih-Chii Liu, and Tobi Delbruck. A 240×180 130 db $3 \mu\text{s}$ latency global shutter spatiotemporal vision sensor. *IEEE Journal of Solid-State Circuits*, 49(10):2333–2341, 2014.
- Han Cai, Chuang Gan, Tianzhe Wang, Zhekai Zhang, and Song Han. Once-for-all: Train one network and specialize it for efficient deployment, 2020. URL <https://arxiv.org/abs/1908.09791>.
- Yongqiang Cao, Yang Chen, and Deepak Khosla. Spiking deep convolutional neural networks for energy-efficient object recognition. *International Journal of Computer Vision*, 113:54–66, 2015.
- Sayed Shafayet Chowdhury, Nitin Rath, and Kaushik Roy. One timestep is all you need: Training spiking neural networks with ultra low latency, 2021.
- Benjamin Cramer, Yannik Stradmann, Johannes Schemmel, and Friedemann Zenke. The heidelberg spiking data sets for the systematic evaluation of spiking neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 33(7):2744–2757, 2022. doi: 10.1109/TNNLS.2020.3044364.
- Manon Dampfhofer, Thomas Mesquida, Alexandre Valentian, and Lorena Anghel. Investigating current-based and gating approaches for accurate and energy-efficient spiking recurrent neural networks. In *International Conference on Artificial Neural Networks*, pp. 359–370. Springer, 2022.
- Lei Deng, Yujie Wu, Xing Hu, Ling Liang, Yufei Ding, Guoqi Li, Guangshe Zhao, Peng Li, and Yuan Xie. Rethinking the performance comparison between snns and anns. *Neural networks*, 121:294–307, 2020.
- Wei Fang, Zhaofer Yu, Yanqi Chen, Timothée Masquelier, Tiejun Huang, and Yonghong Tian. Incorporating learnable membrane time constant to enhance learning of spiking neural networks. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 2661–2671, 2021.
- Guillermo Gallego, Tobi Delbrück, Garrick Orchard, Chiara Bartolozzi, Brian Tabar, Andrea Censi, Stefan Leutenegger, Andrew J. Davison, Jörg Conradt, Kostas Daniilidis, and Davide Scaramuzza. Event-based vision: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(1):154–180, 2022. doi: 10.1109/TPAMI.2020.3008413.
- Mathias Gehrig, Willem Aarents, Daniel Gehrig, and Davide Scaramuzza. Dsec: A stereo event camera dataset for driving scenarios. *IEEE Robotics and Automation Letters*, 2021a. doi: 10.1109/LRA.2021.3068942.
- Mathias Gehrig, Mario Millhäusler, Daniel Gehrig, and Davide Scaramuzza. E-raft: Dense optical flow from event cameras. In *2021 International Conference on 3D Vision (3DV)*, pp. 197–206, 2021b. doi: 10.1109/3DV53792.2021.00030.
- Jesse Hagenaars, Federico Paredes-Vallés, and Guido de Croon. Self-supervised learning of event-based optical flow with spiking neural networks, 2021.
- Ilyass Hammouamri, Ismail Khalfaoui-Hassani, and Timothée Masquelier. Learning delays in spiking neural networks using dilated convolutions with learnable spacings, 2023.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition, 2015.

- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- S Hochreiter. Long short-term memory. *Neural Computation MIT-Press*, 1997.
- Mark Horowitz. 1.1 computing’s energy problem (and what we can do about it). In *2014 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*, pp. 10–14, 2014. doi: 10.1109/ISSCC.2014.6757323.
- Yu Hu, Ziteng Li, Xinpeng Li, Jianfeng Li, Xintong Yu, Xiaofan Chen, and Lei Wang. Hand gesture recognition system using the dynamic vision sensor. In *2022 5th International Conference on Circuits, Systems and Simulation (ICSSS)*, pp. 102–110, 2022. doi: 10.1109/ICSSS55260.2022.9802196.
- Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4700–4708, 2017.
- Chunming Jiang and Yilei Zhang. Klif: An optimized spiking neuron unit for tuning surrogate gradient function. *Neural Computation*, pp. 1–15, 2024.
- Jacques Kaiser, Hesham Mostafa, and Emre Neftci. Synaptic plasticity dynamics for deep continuous local learning (decolle). *Frontiers in Neuroscience*, 14:424, 2020.
- Ismail Khalfaoui-Hassani, Thomas Pellegrini, and Timothée Masquelier. Dilated convolution with learnable spacings, 2023. URL <https://arxiv.org/abs/2112.03740>.
- Youngeun Kim and Priyadarshini Panda. Revisiting batch normalization for training low-latency deep spiking neural networks from scratch. *Frontiers in neuroscience*, 15:773954, 2021.
- Youngeun Kim, Yuhang Li, Hyoungseob Park, Yeshwanth Venkatesha, and Priyadarshini Panda. Neural architecture search for spiking neural networks, 2022.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Adarsh Kumar Kosta and Kaushik Roy. Adaptive-spikenet: Event-based optical flow estimation using spiking neural networks with learnable neuronal dynamics, 2023.
- Chankyu Lee, Adarsh Kumar Kosta, and Kaushik Roy. Fusion-flownet: Energy-efficient optical flow estimation using sensor fusion and deep fused spiking-analog network architectures, 2021. URL <https://arxiv.org/abs/2103.10592>.
- Yijin Li, Zhaoyang Huang, Shuo Chen, Xiaoyu Shi, Hongsheng Li, Hujun Bao, Zhaopeng Cui, and Guofeng Zhang. Blinkflow: A dataset to push the limits of event-based optical flow estimation. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 3881–3888, 2023. doi: 10.1109/IROS55552.2023.10341802.
- Yudong Li, Yunlin Lei, and Xu Yang. Spikeformer: A novel architecture for training high-performance low-latency spiking neural network, 2022.
- Patrick Lichtsteiner, Christoph Posch, and Tobi Delbruck. A 128×128 120 db 15 μ s latency asynchronous temporal contrast vision sensor. *IEEE Journal of Solid-State Circuits*, 43(2):566–576, 2008. doi: 10.1109/JSSC.2007.914337.
- Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts, 2017. URL <https://arxiv.org/abs/1608.03983>.
- Emre O Neftci, Hesham Mostafa, and Friedemann Zenke. Surrogate gradient learning in spiking neural networks: Bringing the power of gradient-based optimization to spiking neural networks. *IEEE Signal Processing Magazine*, 36(6):51–63, 2019.

- Shubham Negi, Deepika Sharma, Adarsh Kumar Kosta, and Kaushik Roy. Best of both worlds: Hybrid snn-ann architecture for event-based optical flow estimation, 2024.
- Priyadarshini Panda, Sai Aparna Aketi, and Kaushik Roy. Toward scalable, efficient, and accurate deep spiking neural networks with backward residual connections, stochastic softmax, and hybridization. *Frontiers in Neuroscience*, 14:535502, 2020.
- Federico Paredes-Vallés and Guido CHE De Croon. Back to event basics: Self-supervised learning of image reconstruction for event cameras via photometric constancy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3446–3455, 2021.
- Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. On the difficulty of training recurrent neural networks, 2013.
- Wachirawit Ponghiran and Kaushik Roy. Hybrid analog-spiking long short-term memory for energy efficient computing on edge devices. In *2021 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pp. 581–586. IEEE, 2021.
- Wachirawit Ponghiran, Chamika Mihiranga Liyanagedera, and Kaushik Roy. Event-based temporally dense optical flow estimation with sequential learning, 2023.
- Jiahao Qin and Feng Liu. Mamba-spike: Enhancing the mamba architecture with a spiking front-end for efficient temporal data processing, 2024. URL <https://arxiv.org/abs/2408.11823>.
- Nitin Rathi and Kaushik Roy. Diet-snn: Direct input encoding with leakage and threshold optimization in deep spiking neural networks, 2020.
- Nitin Rathi and Kaushik Roy. Lite-snn: Leveraging inherent dynamics to train energy-efficient spiking neural networks for sequential learning. *IEEE Transactions on Cognitive and Developmental Systems*, pp. 1–10, 2024. doi: 10.1109/TCDS.2024.3396431.
- Nitin Rathi, Amogh Agrawal, Chankyu Lee, Adarsh Kumar Kosta, and Kaushik Roy. Exploring spike-based learning for neuromorphic computing: Prospects and perspectives. In *2021 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pp. 902–907. IEEE, 2021.
- B Rueckauer, IA Lungu, Y Hu, M Pfeiffer, and SC Liu. Conversion of continuous-valued deep networks to efficient event-driven networks for image classification. *front. neurosci.* 11, 682 (2017), 2017.
- David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning internal representations by error propagation, parallel distributed processing, explorations in the microstructure of cognition, ed. de rumelhart and j. mcclelland. vol. 1. 1986. *Biometrika*, 71:599–607, 1986.
- Ali Samadzadeh, Fatemeh Sadat Tabatabaei Far, Ali Javadi, Ahmad Nickabadi, and Morteza Haghiri Chehreghani. Convolutional spiking neural networks for spatio-temporal feature extraction. *Neural Processing Letters*, 55(6):6979–6995, 2023.
- Guobin Shen, Dongcheng Zhao, and Yi Zeng. Eventmix: An efficient data augmentation strategy for event-based learning. *Information Sciences*, 644:119170, 2023.
- Sumit Bam Shrestha and Garrick Orchard. Slayer: Spike layer error reassignment in time, 2018.
- Pengfei Sun, Yansong Chua, Paul Devos, and Dick Botteldooren. Learnable axonal delay in spiking neural networks improves spoken word recognition. *Frontiers in Neuroscience*, 17:1275944, 2023a.
- Pengfei Sun, Ehsan Eqlimi, Yansong Chua, Paul Devos, and Dick Botteldooren. Adaptive axonal delays in feedforward spiking neural networks for accurate spoken word recognition. In *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1–5, 2023b. doi: 10.1109/ICASSP49357.2023.10094768.

- P.J. Werbos. Backpropagation through time: what it does and how to do it. *Proceedings of the IEEE*, 78(10):1550–1560, 1990. doi: 10.1109/5.58337.
- Yannan Xing, Gaetano Di Caterina, and John Soraghan. A new spiking convolutional recurrent neural network (scrnn) with applications to event-based hand gesture recognition. *Frontiers in neuroscience*, 14:590164, 2020.
- Bojian Yin, Federico Corradi, and Sander M Bohté. Accurate and efficient time-domain classification with adaptive spiking recurrent neural networks. *Nature Machine Intelligence*, 3(10):905–913, 2021.
- Shimin Zhang, Qu Yang, Chenxiang Ma, Jibin Wu, Haizhou Li, and Kay Chen Tan. Long short-term memory with two-compartment spiking neuron, 2023.
- Chenlin Zhou, Liutao Yu, Zhaokun Zhou, Zhengyu Ma, Han Zhang, Huihui Zhou, and Yonghong Tian. Spikingformer: Spike-driven residual learning for transformer-based spiking neural network, 2023.
- Alex Zhu, Liangzhe Yuan, Kenneth Chaney, and Kostas Daniilidis. Ev-flownet: Self-supervised optical flow estimation for event-based cameras. In *Robotics: Science and Systems XIV*, RSS2018. Robotics: Science and Systems Foundation, June 2018a. doi: 10.15607/rss.2018.xiv.062. URL <http://dx.doi.org/10.15607/RSS.2018.XIV.062>.
- Alex Zihao Zhu, Liangzhe Yuan, Kenneth Chaney, and Kostas Daniilidis. Unsupervised event-based learning of optical flow, depth, and egomotion, 2018b.

A Inference Energy Estimation

To validate the observation that *TSkips* exhibit reduced energy consumption compared to baseline SNNs and standard recurrent methods (vRNNs (Rumelhart et al., 1986) and LSTMs (Hochreiter, 1997)) for various tasks (Sections 4.3 and 4.4), we estimate inference energy. These estimations are made using the method from Rueckauer et al. (2017); Lee et al. (2021), which considers the energy required for sparse accumulate (AC) operations in SNNs ($E_{AC} = 0.9pJ$) and dense multiply-and-accumulate (MAC) operations in ANNs ($E_{MAC} = 4.6pJ$) based on a 45nm CMOS technology (Horowitz, 2014). We calculated the total inference energy (E_{total}) as:

$$E_{total} = \begin{cases} \#OPS_{SNN} \times E_{AC}, & \text{for SNNs} \\ \#OPS_{ANN} \times E_{MAC}, & \text{for ANNs} \end{cases} \quad (6)$$

where $\#OPS_{SNN}$ and $\#OPS_{ANN}$ represent the layer-wise synaptic connections in SNNs and ANNs, respectively, calculated as:

$$\begin{aligned} \#OPS_{SNN} &= T \times \mathcal{N} \times \mathcal{C} \times \mathcal{M}, \\ \#OPS_{ANN} &= \mathcal{N} \times \mathcal{C}. \end{aligned}$$

Here, T is the sequence length, \mathcal{N} is the number of neurons in the layer, \mathcal{C} is the number of synaptic connections per neuron, and \mathcal{M} is the average spike rate over T .

A.1 Inference Energy Ablation Study

Building upon our ablation study in Section 4.2, we further investigated the impact of different *TSkips* configurations on inference energy, estimated using Eq. 6. Specifically, we explored the impact of varying temporal delay (Δt), *TSkips* position, and network depth on an 8-layer multilayer perceptron (MLP) baseline model with concatenation-based *TSkips*, found using the Neural Architecture Search (NAS) framework (Kim et al., 2022). This analysis is crucial because *TSkips* increases the average spike rate within the network, potentially leading to higher inference energy consumption.




Figure 8: Ablation study on the SHD dataset demonstrating the impact of varying the (a) temporal delay (Δt), (b) *TSkips* position and (c) network depth on inference energy.

Impact of Temporal Delay (Δt): We investigated the impact of varying Δt for both forward and backward *TSkips* on inference energy (Fig. 8(a)). Forward connections originated from the input data and terminated at the second layer, while backward connections originated from the final output and terminated at the seventh layer. Our analysis revealed that varying Δt significantly influenced inference energy. Larger delays resulted in lower energy consumption, likely because fewer (but more temporally relevant) spikes were propagated through the connections.

Impact of TSkip Position: Next, we investigated the impact of varying *TSkips* placement on inference energy (Fig. 8(b)). With a fixed $\Delta t = 16$, we varied the destination layer of both forward (originating from the input data) and backward (originating from the final output) connections. While no clear trend emerged, terminating both forward and backward *TSkips* at deeper layers generally resulted in lower inference energy consumption.

Impact of Network Depth: Finally, we analyzed the impact of *TSkips* on inference energy across models of varying depth (Fig. 8(c)). Forward connections originated from the input, and backward connections originated from the final output, with $\Delta t = 16$. Our analysis revealed that increasing the number of layers led to a significant increase in inference energy.

This analysis of inference energy with varying *TSkips* parameters further emphasizes the complex interplay between Δt , *TSkips* position, and network depth. It highlights the importance of using methods like Kim et al. (2022) to find optimal network architectures and *TSkips* configurations. While Kim et al. (2022) does not explicitly optimize for minimal energy consumption, it prioritizes networks with diverse spiking activity by quantifying the dissimilarity between spiking patterns. As a result, the identified configurations often utilize *TSkips* with larger Δt values, enabling our networks to efficiently capture long-term patterns while maintaining low energy consumption.

B Additional Results - DSEC-flow

B.1 Inference Energy Analysis

As demonstrated in Section 4.3, *TSkips* reduce average endpoint error (AEE) without increasing model complexity or energy consumption for optical flow estimation on the DSEC-flow dataset. To further analyze this, we evaluated the inference energy of our fully spiking EV-FlowNet architectures augmented with *TSkips* using Eq. 6.

Table 4 demonstrates the effectiveness of incorporating *TSkips* into fully spiking EV-FlowNet architectures (Kosta & Roy, 2023). By efficiently utilizing temporal information, our approach achieves comparable AEE to baseline models $3\times$ larger, while significantly reducing inference energy. This improvement stems from *TSkips* ability to enhance spike propagation with minimal added parameters. While adding *TSkips* increases the average spike rate, the resulting models achieve comparable AEE to larger baselines (Table 1) with substantially lower energy consumption. This highlights the potential of our method for accurate, robust, and efficient optical flow prediction.

Table 4: Comparison of inference energy (E_{total}) using Eq. 6 for fully spiking DSEC-flow models with and without forward/backward $TSkips$ across all scales.

Model	#Params ($\times 10^6$) (\mathcal{N})	Variant	#OPS _{SNN} ($\times 10^9$)	Spike Rate (\mathcal{M})	E_{total} (mJ)
Base	13.04	Baseline	25.9	45.83	23.3
		Forward $TSkips$	32.7	57.8	29.4
		Backward $TSkips$	30.7	54.32	27.6
Mini	3.41	Baseline	9.71	45.15	8.75
		Forward $TSkips$	11.7	54.23	10.5
		Backward $TSkips$	11.4	53.21	10.3
Micro	0.93	Baseline	5.81	61.9	5.25
		Forward $TSkips$	6.14	65.4	5.52
		Backward $TSkips$	5.86	62.4	5.27
Nano	0.27	Baseline	2.67	55.55	2.40
		Forward $TSkips$	2.81	58.47	2.53
		Backward $TSkips$	2.71	56.3	2.44
Pico	0.092	Baseline	2.11	74.12	1.90
		Forward $TSkips$	2.29	80.6	2.07
		Backward $TSkips$	2.25	78.9	2.02

B.2 Convergence Analysis

As discussed in Section 3, incorporating $TSkips$ can lead to easier training and faster convergence for both fully spiking (Kosta & Roy, 2023) and hybrid (Negi et al., 2024) EV-FlowNet architectures. To validate this, we analyzed the convergence behavior of our DSEC-flow models across different scales. Fig. 9 shows that both our largest (Base vs. Mini with $TSkips$) and smallest (Nano vs. Pico with $TSkips$) models exhibit faster convergence when incorporating $TSkips$, for both fully spiking and hybrid variants. This faster convergence can be attributed to $TSkips$ ability to mitigate vanishing spikes and enhance gradient flow throughout the network, enabling quicker optimization and improved learning dynamics. Additionally, by facilitating the propagation of temporally relevant information, $TSkips$ can help the network learn important temporal dependencies more efficiently, leading to faster convergence. This consistent trend across different model sizes highlights the benefits of $TSkips$ in facilitating fast, efficient and effective training.

B.3 Hyper parameters for DSEC-flow

Table 5 specifies the $TSkips$ configurations used to replace existing skip connections in the fully spiking (Kosta & Roy, 2023) and hybrid (Negi et al., 2024) EV-FlowNet architectures. Here, “ $TSkips$ Position” indicates the targeted connection, with ‘1’ representing the longest skip (between the first encoder and last decoder blocks) and ‘3’ the shortest (between the last encoder and first decoder blocks), see Fig. 5.

C Additional Results - DVS128 Gesture, SHD and SSC

C.1 Inference Energy Analysis

As demonstrated in Section 4.4, $TSkips$ improves classification accuracy for speech recognition compared to standard recurrent models (vRNNs and LSTMs). Additionally, we state that $TSkips$ also reduce inference energy compared to these models. This energy efficiency is further confirmed by the estimations presented in Table 6 for both shallow (4-layer) and deep (8-layer) MLP (Rumelhart et al., 1986) networks on the SHD (Cramer et al., 2022) dataset. These estimations were calculated using Eq. 6, which considers the energy consumed by both sparse accumulate (AC) operations in SNNs and dense multiply-and-accumulate (MAC) operations in ANNs.

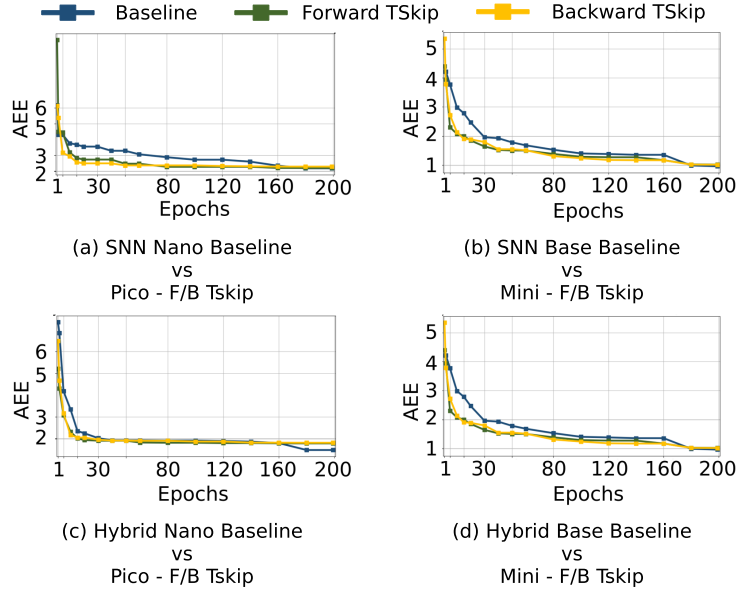


Figure 9: Convergence analysis of fully spiking (Kosta & Roy, 2023) and hybrid (Negi et al., 2024) EV-FlowNet architectures on the DSEC-flow dataset. The plots compare the convergence behavior of baseline models (Base and Nano) with their smaller forward (F) and backward (B) *TSkips*-augmented counterparts (Mini and Pico, respectively).

Table 5: *TSkip* parameters used for fully spiking and hybrid DSEC-flow models.

Architecture	SNN		Hybrid	
	Δt	<i>TSkip</i> Pos	Δt	<i>TSkip</i> Pos
Base + Forward <i>TSkips</i>	3	1	3	1
Base + Backward <i>TSkips</i>	4	1	5	1
Mini + Forward <i>TSkips</i>	3	1	4	2
Mini + Backward <i>TSkips</i>	3	2	3	1
Micro + Forward <i>TSkips</i>	4	2	3	2
Micro + Backward <i>TSkips</i>	3	3	3	3
Nano + Forward <i>TSkips</i>	3	1	3	2
Nano + Backward <i>TSkips</i>	4	2	3	2
Pico + Forward <i>TSkips</i>	4	1	3	1
Pico + Backward <i>TSkips</i>	3	2	4	2

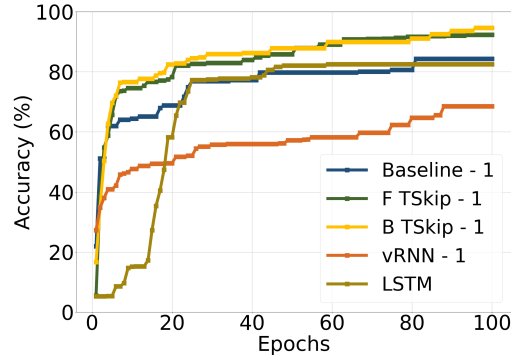
For comparison, the vRNNs in Table 6 are our baseline SNNs with backward skip connections added to every layer with $\Delta t = 1$. The LSTMs are implemented as ANNs, with the 0.2×10^6 and 1.1×10^6 parameter models having 700 input channels and hidden sizes of 75 and 400, respectively.

The observed reduction in energy consumption can be attributed to several key factors. Firstly, *TSkips* introduce a limited number of additional synaptic connections (\mathcal{C}) compared to the densely connected nature of RNNs and LSTMs. Secondly, they promote sparse spiking activity by selectively propagating information across time steps, leading to lower energy consumption than the dense computations inherent in recurrent layers. Finally, the explicit temporal delays in *TSkips* enable the network to focus on the most relevant information from previous time steps, thereby reducing unnecessary computations and further improving energy efficiency.

Table 6: Inference energy (E_{total}) best-performing models on the SHD dataset. Energy is calculated using Eq. 6.

Model	#Params ($\times 10^6$) (\mathcal{N})	#OPS ($\times 10^{12}$)	Spike Rate ($\times 10^4$) (\mathcal{M})	E_{total} (mJ)
Baseline - 1	0.16	0.0083	3.16	7.52
vRNN - 1	0.16	0.213	38.08	192.41
LSTM - 1	0.2	0.016	-	74.63
Forward <i>TSkips</i> - 1	0.24	0.038	5.62	34.79
Backward <i>TSkips</i> - 1	0.19	0.022	4.95	20.37
Forward + Backward <i>TSkips</i> - 1	0.20	0.032	7.42	29.06
Baseline - 2	1.04	0.815	7.63	734.26
vRNN - 2	1.04	10.854	58.23	9768.88
LSTM - 2	1.1	1.196	-	5505.28
Forward <i>TSkips</i> - 2	1.12	1.204	8.21	1083.90
Backward <i>TSkips</i> - 2	1.16	1.213	8.23	1091.87
Forward + Backward <i>TSkips</i> - 2	1.28	1.761	9.13	1585.70

C.2 Convergence Analysis

Figure 10: Test accuracy convergence on the SHD dataset, comparing the best-performing models with *TSkips* (from Table 3) to a vRNN and an LSTM with 0.2×10^6 parameters.

As discussed in Sections 3 and 4.4, incorporating *TSkips* can lead to easier training and faster convergence compared to recurrent models like vRNNs and LSTMs, as demonstrated in Fig. 10 for a 4-layer MLP network on the SHD dataset. This accelerated convergence can be attributed to several factors: improved gradient flow by mitigating the vanishing gradient problem, efficient information propagation by facilitating the access to temporally relevant information, and reduced computational complexity compared to the dense connectivity of recurrent networks. These factors, combined with the ability of *TSkips* to capture long-term temporal dependencies, contribute to the improved training efficiency observed in our experiments.

C.3 TSkips in ANNs and Hybrid Models

As discussed in Section 4.4, incorporating *TSkips* into ANNs and hybrid ANN-SNN models can lead to improved classification accuracy. Table 7 presents these results, demonstrating the effectiveness of *TSkips* in enhancing both ANN and hybrid model performance, surpassing the performance of vRNNs and LSTMs, respectively.

This improvement can be attributed to the ability of *TSkips* to facilitate more effective temporal processing. By incorporating explicit temporal delays, these connections allow the network to access and integrate information from past time steps, which can be crucial for understanding complex temporal patterns. This

Table 7: Classification accuracy of ANN and Hybrid ANN-SNN models on the SHD and SSC datasets. We compare the 4-layer MLP baseline (Baseline-1) with forward/backward *TSkips*.

Method	#Params ($\times 10^6$)	Accuracy (%)	
		ANN	Hybrid
SHD			
Baseline	0.16	49.95	74.11
vRNN	0.16	60.51	67.45
Forward <i>TSkips</i>	0.24	69.08	84.54
Backward <i>TSkips</i>	0.19	73.45	86.02
Forward + Backward <i>TSkips</i>	0.20	73.25	85.74
LSTM	0.20	80.54	-
SSC			
Baseline	0.12	55.41	59.61
vRNN	0.12	68.29	71.43
Forward <i>TSkips</i>	0.42	71.19	73.48
Backward <i>TSkips</i>	0.24	70.23	72.87
Forward + Backward <i>TSkips</i>	0.54	71.09	74.96
LSTM	0.20	72.3	-

enhanced temporal awareness enables the network to make more accurate predictions, leading to improved classification accuracy. Furthermore, *TSkips* can mitigate the vanishing spikes problem often encountered in deep networks, by providing more direct pathways for gradient flow during training. This improved gradient propagation can lead to more stable and efficient training, further contributing to the enhanced performance observed in ANNs and hybrid models.

C.4 Comparison between TSkip Operators

To further analyze the impact of different *TSkips* connection types, we compared the performance of concatenation-based and addition-based *TSkips* on the 4-layer and 8-layer baselines (Table 3).

Table 8 demonstrates the superior performance of concatenation-based *TSkips* over addition-based on the baseline models. Concatenation-based *TSkips* consistently outperformed their addition-based counterparts. This can be attributed to the increased representational capacity of concatenation, which expands the feature space by combining information from the skip connection and the destination layer. This richer representation allows the network to learn more complex temporal patterns, particularly crucial in deep SNNs where information is encoded in sparse spike trains and membrane potentials. Furthermore, concatenation preserves information from both sources, preventing potential information loss that can occur with addition.

C.5 Hyper parameters for DVS128 Gesture, SHD and SSC

Table 9 presents the top-ranked network architectures and *TSkips* configurations identified by NASWOT-SAHD (Kim et al., 2022), with corresponding results shown in Tables 2 and 3.

The table specifies the following for each configuration:

1. Network architecture:
 - (a) DVS128 Gesture: $2 \times 64 \times 64$ represents the input dimensions (polarity 2, 64×64 spike resolution) for all networks. Convolutional layers are represented concisely, for example, 3c80s1 denotes a layer with a 3×3 kernel, 80 output channels, and a stride of 1.
 - (b) SHD and SSC: Specifies the number of input channels for each MLP layer, including the final fully connected layer. Channels represented as $n * 2$ indicate that concatenation-based *TSkips* were used, doubling the number of channels at that layer. This representation is used to highlight

Table 8: Comparison of classification accuracy on SHD and SSC datasets with concatenation-based (C) and addition-based (A) forward/backward *TSkips* on a 4-layer (Baseline -1) and 8-layer (Baseline - 2) MLP network.

Method	#Params ($\times 10^6$)		Accuracy (%)	
	SHD	SSC	SHD	SSC
Baseline - 1				
Baseline	0.16	0.12	84.32	64.19
Forward <i>TSkips</i> - C	0.24	0.42	92.32	76.5
Backward <i>TSkips</i> - C	0.19	0.24	93.64	79.87
Forward + Backward <i>TSkips</i> - C	0.20	0.54	93.01	78.64
Forward <i>TSkips</i> - A	0.19	0.37	87.65	72.13
Backward <i>TSkips</i> - A	0.11	0.21	87.54	73.15
Forward + Backward <i>TSkips</i> - A	0.11	0.40	88.13	75.12
Baseline - 2				
Baseline	1.04	1.04	86.42	67.54
Forward <i>TSkips</i> - C	1.12	1.08	94.15	78.98
Backward <i>TSkips</i> - C	1.16	1.14	93.86	79.65
Forward + Backward <i>TSkips</i> - C	1.28	1.42	94.73	80.23
Forward <i>TSkips</i> - A	0.99	1.05	86.39	72.15
Backward <i>TSkips</i> - A	0.97	1.01	87.32	75.53
Forward + Backward <i>TSkips</i> - A	0.98	1.14	89.45	77.12

that the results shown in Table 8, which uses addition-based *TSkips*, are based on the same underlying network architectures.

2. *TSkips* parameters: The temporal delay (Δt) and the origin and destination layers (*from* \rightarrow *to*) for Forward (F), Backward (B), and combined (F+B) *TSkips*.
3. Other hyper parameters:
 - (a) Leak and threshold initialization values for the adaptive LIF neurons: 0.6 and 15, respectively.
 - (b) Dataset-specific dropout rates: 0.4 for SHD and 0.2 for SSC.

D NASWOT-SAHD Correlation on *TSkips*

To validate the effectiveness of NASWOT-SAHD (Kim et al., 2022) in identifying optimal network configurations, we analyzed the correlation between the accuracy and score of the identified networks, as shown in Fig. 11. The scores are based on the Sparsity Aware Hamming Distance (SAHD) (Kim et al., 2022), and the correlation is measured using Kendall’s τ . Our baseline models achieved a τ correlation of 0.63, consistent with the results reported in Kim et al. (2022). However, the forward and back-

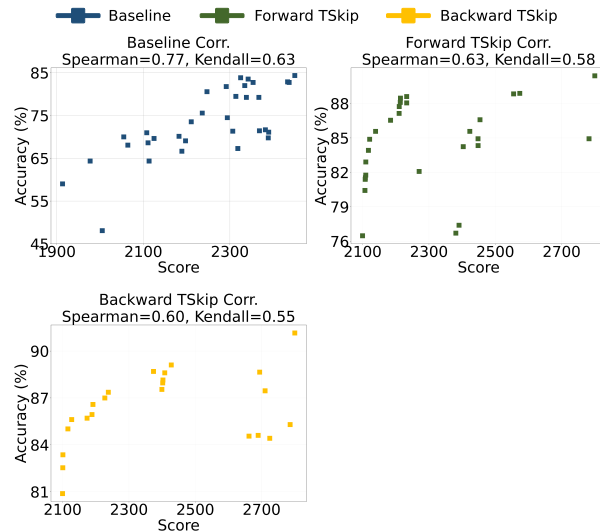


Figure 11: Correlation between accuracy and SAHD score for different network configurations (Baseline, forward *TSkips*, and backward *TSkips*).

ward *TSkips* network searches resulted in slightly lower correlations of 0.58 and 0.55, respectively. All results reported are for the top-ranked networks identified by NASWOT-SAHD (Kim et al., 2022).

To ensure a focused and efficient search, we introduced several constraints:

1. Temporal Delay (Δt):
 - (a) DSEC-flow ($T = 10$): $2 \leq \Delta t \leq 6$
 - (b) DVS Gesture ($T = 30$): $5 \leq \Delta t \leq 14$
 - (c) SHD and SSC ($T = 99$): $10 \leq \Delta t \leq 45$
2. *TSkips* Position: We constrained connections to not be within the same layer, ensuring that all connections originated from a different layer than the destination.
3. Model Size:
 - (a) DSEC-Flow: No constraint was applied as we directly replaced existing skip connections in the fully spiking (Kosta & Roy, 2023) and hybrid (Negi et al., 2024) EV-FlowNet architecture with *TSkips*, maintaining the original model size.
 - (b) DVS128 Gesture: 0.6×10^6 parameters
 - (c) SHD and SSC (Baseline-1): 0.3×10^6 parameters
 - (d) SHD and SSC (Baseline-2): 1.3×10^6 parameters

Table 9: Optimal Parameters for Training MLP Networks with Temporal Skip Connections on SHD and SSC Datasets.

Variant	Network	Delay (Δt)	TSkip Position (from \rightarrow to)
DVS128 Gesture			
Baseline	$2 \times 64 \times 64$ -3c80s1-3c80s1-5c86s1-5c64s1-1c64s1-1c32s11	-	-
F <i>TSkips</i>	$2 \times 64 \times 64$ -3c83s1-3c83s1-5c110s1-3c94s1-1c94s1-1c32s11	5	1 \rightarrow 2
B <i>TSkips</i>	$2 \times 64 \times 64$ -3c64s1-5c64s1-5c64s1-5c76s1-1c76s1-1c32s11	5	4 \rightarrow 3
F + B	$2 \times 64 \times 64$ -3c122s1-5c122s1-	F: 6	F: 1 \rightarrow 3
<i>TSkips</i>	3c122s1-5c64s1-1c64s1+1c32s11	B: 8	B: 5 \rightarrow 4
SHD			
Baseline - 1	700-124-288-144-20	-	-
F <i>TSkips</i> - 1	700-124*2-432-115-20	24	1 \rightarrow 2
B <i>TSkips</i> - 1	700*2-124-115-96-20	14	4 \rightarrow 1
F + B	700*2-124-144-72*2-20	F: 18	F: 2 \rightarrow 4
<i>TSkips</i> - 1		B: 14	B: 2 \rightarrow 1
Baseline - 2	700-512-288-224-192-896-288-20	-	-
F <i>TSkips</i> - 2	700-532-288-524-192-448*2-288-20	17	2 \rightarrow 6
B <i>TSkips</i> - 2	700-321-558*2-334-292-452-328-20)	12	3 \rightarrow 5
F + B	700-321-228*2-334-	F: 10	F: 2 \rightarrow 5
<i>TSkips</i> - 2	492*2-462-448-20	B: 13	B: 7 \rightarrow 3
SSC			
Baseline - 1	700-124-148-130-35	-	-
F <i>TSkips</i> - 1	700-324-298*2-145-35	15	1 \rightarrow 3
B <i>TSkips</i> - 1	700-184*2-179-230-35	12	3 \rightarrow 2
F + B	700-364*2-257*2-195-35	F: 14	F: 1 \rightarrow 3
<i>TSkips</i> - 1		B: 8	B: 4 \rightarrow 2
Baseline - 2	700-731-410-340-150-184-58-35	-	-
F <i>TSkips</i> - 2	700-646-410-400-250-184*2-135-35	17	2 \rightarrow 6
B <i>TSkips</i> - 2	700-663-373*2-337-292-182-135-35	15	7 \rightarrow 3
F + B	700-674-493*2-347*2-	F: 9	F: 1 \rightarrow 2
<i>TSkips</i> - 2	278-192-89-35	B: 11	B: 7 \rightarrow 3