

Pruning of Deep Spiking Neural Networks through Gradient Rewiring

Yanqi Chen^{1,3}, Zhaofei Yu^{1,2,3*}, Wei Fang^{1,3}, Tiejun Huang^{1,2,3} and Yonghong Tian^{1,3*}

¹Department of Computer Science and Technology, Peking University

²Institute for Artificial Intelligence, Peking University

³Peng Cheng Laboratory

{chyq, yuzf12, fwei, tjhuang, yhtian}@pku.edu.cn

Abstract

Spiking Neural Networks (SNNs) have been attached great importance due to their biological plausibility and high energy-efficiency on neuromorphic chips. As these chips are usually resource-constrained, the compression of SNNs is thus crucial along the road of practical use of SNNs. Most existing methods directly apply pruning approaches in artificial neural networks (ANNs) to SNNs, which ignore the difference between ANNs and SNNs, thus limiting the performance of the pruned SNNs. Besides, these methods are only suitable for shallow SNNs. In this paper, inspired by synaptogenesis and synapse elimination in the neural system, we propose gradient rewiring (Grad R), a joint learning algorithm of connectivity and weight for SNNs, that enables us to seamlessly optimize network structure without retraining. Our key innovation is to redefine the gradient to a new synaptic parameter, allowing better exploration of network structures by taking full advantage of the competition between pruning and regrowth of connections. The experimental results show that the proposed method achieves minimal loss of SNNs' performance on MNIST and CIFAR-10 datasets so far. Moreover, it reaches a $\sim 3.5\%$ accuracy loss under unprecedented 0.73% connectivity, which reveals remarkable structure refining capability in SNNs. Our work suggests that there exists extremely high redundancy in deep SNNs. Our codes are available at <https://github.com/Yanqi-Chen/Gradient-Rewiring>.

1 Introduction

Spiking Neural Networks (SNNs), as the third generation of Artificial Neural Networks (ANN) [Maass, 1997], have increasingly aroused researchers' great interest in recent years due to their high biological plausibility, temporal information processing capability, and low power consumption on neuromorphic chips [Gerstner and Kistler, 2002; Roy *et al.*, 2019].

As these chips are usually resource-constrained, the compression of SNNs is thus of great importance to lower computation cost to make the deployment on neuromorphic chips more feasible, and realize the practical application of SNNs.

Researchers have made a lot of efforts on the pruning of SNNs. Generally, most methods learn from the existing compression techniques in the field of ANNs and directly apply them to SNNs, such as directly truncating weak weights below the threshold [Neftci *et al.*, 2016; Rathi *et al.*, 2019; Liu *et al.*, 2020] and using Lottery Ticket Hypothesis [Frankle and Carbin, 2019] in SNNs [Martinelli *et al.*, 2020]. Some works use other metrics, such as levels of weight development [Shi *et al.*, 2019] and similarity between spike trains [Wu *et al.*, 2019a], as guidance for removing insignificant or redundant weights. Notwithstanding some spike-based metrics for pruning are introduced in these works, the results on shallow SNNs are not persuasive enough.

There are some attempts to mimic the way the human brain learns to enhance the pruning of SNNs. Qi *et al.* [2018] enabled structure learning of SNNs by adding connection gates for each synapse. They adopted the biological plausible local learning rules, Spike-Timing-Dependent Plasticity (STDP), to learn gated connection while using Tempotron rules [Gütig and Sompolinsky, 2006] for regular weight optimization. However, the method only applies to the SNN with a two-layer structure. Deng *et al.* [2019] exploited the ADMM optimization tool to prune the connections and optimize network weights alternately. Kappel *et al.* [2015] proposed synaptic sampling to optimize network construction and weights by modeling the spine motility of spiking neurons as Bayesian learning. Bellec *et al.* went a further step and proposed Deep Rewiring (Deep R) as a pruning algorithm for ANNs [2018a] and Long Short-term Memory Spiking Neural Networks (LSNNs) [2018b], which was then deployed on SpiNNaker 2 prototype chips [Liu *et al.*, 2018]. Deep R is an on-the-fly pruning algorithm, which indicates its ability to reach target connectivity without retraining. Despite these methods attempting to combine learning and pruning in a unified framework, they focus more on pruning, rather than the rewiring process. Actually, the pruned connection is hard to regrow, limiting the diversity of network structures and thus the performance of the pruned SNNs.

In this paper, inspired by synaptogenesis and synapse elimination in the neural system [Huttenlocher and Nelson, 1994;

*Corresponding author

Petzoldt and Sigrüst, 2014], we propose gradient rewiring (Grad R), a joint learning algorithm of connectivity and weight for SNNs, by introducing a new synaptic parameter to unify the representation of synaptic connectivity and synaptic weights. By redefining the gradient to the synaptic parameter, Grad R takes full advantage of the competition between pruning and regrowth of connections to explore the optimal network structure. To show Grad R can optimize synaptic parameters of SNNs, we theoretically prove that the angle between the accurate gradient and redefined gradient is less than 90° . We further derive the learning rule and realize the pruning of deep SNNs. We evaluate our algorithm on the MNIST and CIFAR-10 dataset and achieves state-of-the-art performance. On CIFAR-10, Grad R reaches a classification accuracy of 92.03% (merely $<1\%$ accuracy loss) when constrained to $\sim 5\%$ connectivity. Further, it achieves a classification accuracy of 89.32% ($\sim 3.5\%$ accuracy loss) under unprecedented 0.73% connectivity, which reveals remarkable structure refining capability in SNNs. Our work suggests that there exists extremely high redundancy in deep SNNs.

2 Methods

In this section, we present the pruning algorithm for SNNs in detail. We first introduce the spiking neural model and its discrete computation mechanism. Then we propose the network structure of the SNNs used in this work. Finally, we present the gradient rewiring framework to learn the connectivity and weights of SNNs jointly and derive the detailed learning rule.

2.1 Spiking Neural Model

The spiking neuron is the fundamental computing unit of SNNs. Here we consider the Leaky Integrate-and-Fire (LIF) model [Gerstner and Kistler, 2002], which is one of the simplest spiking neuron models widely used in SNNs. The membrane potential $u(t)$ of the LIF neuron is defined as:

$$\begin{aligned} \tau_m \frac{du(t)}{dt} &= -(u(t) - u_{\text{rest}}) + \sum_i w_i I_i(t), \\ \lim_{\Delta t \rightarrow 0^+} u(t^f + \Delta t) &= u_{\text{rest}}, \text{ if } u(t^f) = u_{\text{th}}, \end{aligned} \quad (1)$$

where $I_i(t)$ denotes the input current from the i -th presynaptic neuron at time t , w_i denotes the corresponding *synaptic weight* (connection weight). τ_m is the membrane time constant, and u_{rest} is the resting potential. When the membrane potential exceeds the firing threshold u_{th} at time t^f , a spike is generated and transmitted to postsynaptic neurons. At the same time, the membrane potential $u(t)$ of the LIF neuron goes back to the reset value $u_{\text{rest}} < u_{\text{th}}$ (see Eq. (1)).

For numerical simulations of LIF neurons in SNNs, we need to consider a version of the parameters dynamics that is discrete in time. Specifically, Eq. (1) is approximated by:

$$\begin{aligned} m_{t+1} &= u_t + \frac{1}{\tau_m} \left(-(u_t - u_{\text{rest}}) + \sum_i w_i I_{i,t} \right), \\ S_{t+1} &= H(m_{t+1} - u_{\text{th}}), \\ u_{t+1} &= S_{t+1} u_{\text{rest}} + (1 - S_{t+1}) m_{t+1}. \end{aligned} \quad (2)$$

Here m_t and u_t denote the membrane potential after neural dynamics and after the trigger of a spike at time-step t , respectively. S_t is the binary spike at time-step t , which equals

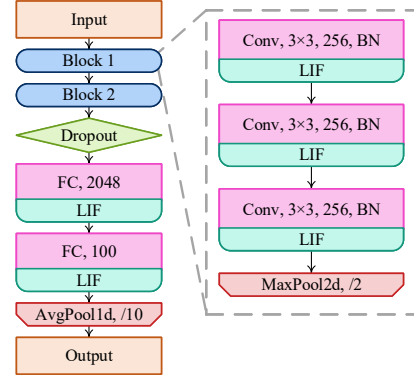


Figure 1: The architecture of the SNN consisted of LIF neurons.

1 if there is a spike and 0 otherwise. $H(\cdot)$ is the Heaviside step function, which is defined by $H(x) = \begin{cases} 1 & x \geq 0 \\ 0 & x < 0 \end{cases}$. Note that $H(0) = 1$ is specially defined.

2.2 Deep Spiking Neural Network

A typical architecture of deep SNNs is exemplified as Fig. 1, which consists of neuronal layers (LIF layers) generating spikes and synaptic layers including convolutional layers, pooling layers, and fully-connected (FC) layers. The synaptic layers define connective patterns of synapses between spiking neurons, which are leveraged for spatial correlative feature extraction or classification. A deep SNN such as CIFARNet proposed in [Wu *et al.*, 2019b] usually has multiple large-scale synaptic layers and requires pruning for reducing computation cost.

2.3 Gradient Rewiring Framework

Inspired by synaptogenesis and synapse elimination of the brain, we present a gradient-based rewiring framework for deep SNNs pruning and call it *Gradient Rewiring*. Rewiring denotes the structural change of the brain’s wiring diagram. Gradient Rewiring combines pruning and learning to explore both connective patterns and weight spaces. The combination of pruning and learning requires continuous parameterized expression of wiring change (synaptic connection change) and weight optimization. Previous neuroscience research has revealed that the wiring plasticity of neurons can be viewed as a special case of weight plasticity [Chklovskii *et al.*, 2004]. Thus we can unify the representation of synaptic connectivity and synaptic weights by introducing a new hidden synaptic parameter θ in addition to the synaptic weight w .

The synaptic parameter θ is related to the volume of a dendritic spine, which defines the synaptic weights and connectivity for each synapse. A synapse is considered connected if $\theta > 0$ and disconnected otherwise. Similar to [Bellec *et al.*, 2018a], the synaptic parameter θ is mapped to synaptic weight through the relation $w = s\text{ReLU}(\theta)$, where $s \in \{+1, -1\}$ is *fixed* and used to distinguish excitatory and inhibitory synapses. In this way, the sign change of

synaptic parameter θ can describe both formation and elimination of synaptic connection.

Optimizing Synaptic Parameters θ

With the above definition, the gradient descent algorithm and its variants can be directly applied to optimize all synaptic parameters θ of the SNNs. Specifically, we assume that the loss function has the form $\mathcal{L}(x_i, y_i, \mathbf{w})$, where x_i, y_i denotes the i -th input sample and the corresponding target output, and \mathbf{w} is the vector of all synaptic weights. The corresponding vectors of all synaptic parameters and signs are $\boldsymbol{\theta}, \mathbf{s}$ respectively. Suppose w_k, θ_k, s_k stand for the k -th element of $\mathbf{w}, \boldsymbol{\theta}, \mathbf{s}$. Then we have $\mathbf{w} = \mathbf{s} \odot \text{ReLU}(\boldsymbol{\theta})$, where \odot is the Hadamard product, and ReLU is an element-wise function. The gradient descent algorithm would adjust all the synaptic parameters to lower the loss function \mathcal{L} , that is:

$$\Delta\theta_k \propto \frac{\partial \mathcal{L}}{\partial \theta_k} = \frac{\partial \mathcal{L}}{\partial w_k} \frac{\partial w_k}{\partial \theta_k} = s_k \frac{\partial \mathcal{L}}{\partial w_k} H(\theta_k), \quad (3)$$

where $H(\cdot)$ represents the Heaviside step function. Ideally, if $\theta_k < 0$ during learning, the corresponding synaptic connection will eliminate. It will regrow when $\theta_k > 0$. Nonetheless, the pruned connection will never regrow. To see this, supposing that $\theta_k = \theta < 0$, we get $H(\theta_k) = 0$, and $w_k = s_k \text{ReLU}(\theta_k) = 0$. Then:

$$\left. \frac{\partial \mathcal{L}}{\partial \theta_k} \right|_{\theta_k = \theta} = H(\theta) s_k \left. \frac{\partial \mathcal{L}}{\partial w_k} \right|_{w_k = 0} = 0, \quad \text{for } \theta < 0. \quad (4)$$

As the gradient of the loss function \mathcal{L} with respect to θ_k is 0, θ_k will keep less than 0, and the corresponding synaptic connection will remain in the state of elimination. This might result in a disastrous loss of expression capabilities for SNNs. So far, the current theories and analyses can also systematically explain the results of deep rewiring [Bellec *et al.*, 2018a], where the pruned connection is hard to regrow even noise is added to the gradient. There is no preference for minimizing loss when adding noise to gradient.

Gradient Rewiring

Our new algorithm adjusts θ in the same way as gradient descent, but for each θ_k , it uses a new formula to redefine the gradient when $\theta_k < 0$.

$$(\Delta\theta_k)_{\text{GradR}} \propto \frac{\partial \bar{\mathcal{L}}}{\partial \theta_k} := s_k \left. \frac{\partial \bar{\mathcal{L}}}{\partial w_k} \right|_{w_k = s_k \text{ReLU}(\theta_k)}, \quad \text{for } \theta_k \in \mathbb{R}. \quad (5)$$

For an active connection $\theta_k = \theta > 0$, the definition makes no difference as $H(\theta_k) = 1$. For an inactive connection $\theta_k = \theta < 0$, the gradients are different as $\Delta\theta_k = 0$ and $(\Delta\theta_k)_{\text{GradR}} \propto s_k \left. \frac{\partial \bar{\mathcal{L}}}{\partial w_k} \right|_{w_k = 0}$. The change of gradient for inactive connections results in a different loss function $\bar{\mathcal{L}}$. As illustrated in Fig. 2(a), when $s_k \left. \frac{\partial \bar{\mathcal{L}}}{\partial w_k} \right|_{w_k = 0} < 0$, the gradient of $\bar{\mathcal{L}}$ guides a way to jump out the plateau in \mathcal{L} (red arrow). With the appropriate step size chosen, a depleted connection will regrow without increasing the real loss \mathcal{L} on this occasion. It should be noted that the situation in Fig. 2(b) can also occur and push the parameter θ_k even more negative, which

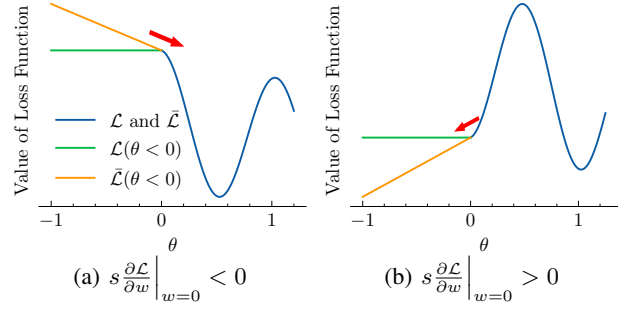


Figure 2: The comparison of gradient descent on $\bar{\mathcal{L}}$ and \mathcal{L} . The red arrows specify the direction of gradient descent on $\bar{\mathcal{L}}$ for θ .

can be overcome by adding a constraint scope to the parameter θ (discussed in the subsequent part).

Similar to the gradient descent algorithm (Eq. (3)), Grad R can also optimize synaptic parameters θ of the SNNs. An intuitive explanation is the angle between the accurate gradient $\nabla_{\theta} \mathcal{L}$ and the redefined gradient $\nabla_{\theta} \bar{\mathcal{L}}$ is below 90° , i.e., $\langle \nabla_{\theta} \mathcal{L}, \nabla_{\theta} \bar{\mathcal{L}} \rangle \geq 0$, which can be derived from Lemma 1. Further, a rigorous proof is illustrated in Theorem 1.

Theorem 1 (Convergence). *For a spiking neural network with synaptic parameters $\boldsymbol{\theta}$, synaptic weights \mathbf{w} , signs \mathbf{s} , and the SoftPlus mapping, a smooth approximation to the ReLU, $w_k = \frac{s_k}{\beta} \log(1 + \exp(\beta\theta_k))$, $\beta \gg 1$, where w_k, θ_k, s_k stand for the k -th element of $\mathbf{w}, \boldsymbol{\theta}, \mathbf{s}$, the loss function $\mathcal{L}(\boldsymbol{\theta})$ with lowerbound, if there exists a constant C such that $\forall k, \theta_k \geq C$ and a Lipschitz constant L such that $\|\nabla_{\theta} \mathcal{L}(\boldsymbol{\theta}_1) - \nabla_{\theta} \mathcal{L}(\boldsymbol{\theta}_2)\| \leq L \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|$ for any differentiable points $\boldsymbol{\theta}_1$ and $\boldsymbol{\theta}_2$, then \mathcal{L} will decrease until convergence using gradient $\nabla_{\theta} \bar{\mathcal{L}}$, if step size $\eta > 0$ is small enough.*

Proof. Let $\boldsymbol{\theta}^t$ be the parameter $\boldsymbol{\theta}$ after t -th update steps using $\nabla_{\theta} \bar{\mathcal{L}}$ as direction for each step, that is, $\boldsymbol{\theta}^{t+1} = \boldsymbol{\theta}^t - \eta \nabla_{\theta} \bar{\mathcal{L}}(\boldsymbol{\theta}^t)$. With Lipschitz continuous gradient, we have

$$\begin{aligned} & \mathcal{L}(\boldsymbol{\theta}^{t+1}) - \mathcal{L}(\boldsymbol{\theta}^t) \\ &= \int_0^1 \langle \nabla_{\theta} \mathcal{L}(\boldsymbol{\theta}^t(\tau)), \boldsymbol{\theta}^{t+1} - \boldsymbol{\theta}^t \rangle d\tau \\ &= \langle \nabla_{\theta} \mathcal{L}(\boldsymbol{\theta}^t), \boldsymbol{\theta}^{t+1} - \boldsymbol{\theta}^t \rangle \\ & \quad + \int_0^1 \langle \nabla_{\theta} \mathcal{L}(\boldsymbol{\theta}^t(\tau)) - \nabla_{\theta} \mathcal{L}(\boldsymbol{\theta}^t), \boldsymbol{\theta}^{t+1} - \boldsymbol{\theta}^t \rangle d\tau \\ &\leq \langle \nabla_{\theta} \mathcal{L}(\boldsymbol{\theta}^t), \boldsymbol{\theta}^{t+1} - \boldsymbol{\theta}^t \rangle \\ & \quad + \int_0^1 \|\nabla_{\theta} \mathcal{L}(\boldsymbol{\theta}^t(\tau)) - \nabla_{\theta} \mathcal{L}(\boldsymbol{\theta}^t)\| \cdot \|\boldsymbol{\theta}^{t+1} - \boldsymbol{\theta}^t\| d\tau \\ &\leq \langle \nabla_{\theta} \mathcal{L}(\boldsymbol{\theta}^t), \boldsymbol{\theta}^{t+1} - \boldsymbol{\theta}^t \rangle + L \|\boldsymbol{\theta}^{t+1} - \boldsymbol{\theta}^t\|^2 \int_0^1 \tau d\tau \\ &= -\eta \langle \nabla_{\theta} \mathcal{L}(\boldsymbol{\theta}^t), \nabla_{\theta} \bar{\mathcal{L}}(\boldsymbol{\theta}^t) \rangle + \frac{L\eta^2}{2} \|\nabla_{\theta} \bar{\mathcal{L}}(\boldsymbol{\theta}^t)\|^2 \\ &\leq -\eta(1 - AL\eta/2) \langle \nabla_{\theta} \mathcal{L}(\boldsymbol{\theta}^t), \nabla_{\theta} \bar{\mathcal{L}}(\boldsymbol{\theta}^t) \rangle, \end{aligned} \quad (6)$$

where $\boldsymbol{\theta}^t(\tau)$ is defined by $\boldsymbol{\theta}^t(\tau) = \boldsymbol{\theta}^t + \tau(\boldsymbol{\theta}^{t+1} - \boldsymbol{\theta}^t)$, $\tau \in [0, 1]$. The first two inequality holds by Cauchy-Schwarz

inequality and Lipschitz continuous gradient, respectively. The last inequality holds with constant $A > 0$ according to Lemma 1.

When $\|\nabla_{\theta} \bar{\mathcal{L}}(\theta^t)\| > 0$ and $\eta < 2/AL$, the last term in Eq. (6) is strictly negative. Hence, we have a strictly decreasing sequence $\{\mathcal{L}(\theta^n)\}_{n \in \mathbb{N}}$. Since the loss function has lower-bound, this sequence must converges. \square

Prior-based Connectivity Control

There is no explicit control of connectivity in the above algorithm. We need to explore speed control of pruning targeting diverse levels of sparsity. Accordingly, we introduce a prior distribution to impose constraints on the range of parameters. To reduce the number of extra hyperparameters, we adopt the same Laplacian distribution for each synaptic parameter θ_k :

$$p(\theta_k) = \frac{\alpha}{2} \exp(-\alpha|\theta_k - \mu|), \quad (7)$$

where α and μ are scale parameter and location parameter respectively. The corresponding penalty term is:

$$\frac{\partial \log p(\theta_k)}{\partial \theta_k} = -\alpha \text{sign}(\theta_k - \mu), \quad (8)$$

which is similar to ℓ^1 -penalty with $\text{sign}(\cdot)$ denoting the sign function.

Here we illustrate how the prior influences network connectivity. We assume that all the synaptic parameters θ follow the same prior distribution in Eq. (7). Since the connection with the negative synaptic parameter ($\theta_k < 0$) is pruned, the probability that a connection is pruned is:

$$\int_{-\infty}^0 p(\theta_k) d\theta_k = \int_{-\infty}^0 \frac{\alpha}{2} \exp(-\alpha|\theta_k - \mu|) d\theta_k \quad (9)$$

$$= \begin{cases} \frac{1}{2} \exp(-\alpha\mu), & \text{if } \mu > 0 \\ 1 - \frac{1}{2} \exp(\alpha\mu), & \text{if } \mu \leq 0 \end{cases}$$

As all the synaptic parameters follow the same prior distribution, Eq. (9) also represents the network’s sparsity. Conversely, given target sparsity p of the SNNs, one can get the appropriate choice of location parameter μ with Eq. (9), which is related to the scale parameter α

$$\mu = \begin{cases} -\frac{1}{\alpha} \ln(2p), & \text{if } p < \frac{1}{2} \\ \frac{1}{\alpha} \ln(2 - 2p), & \text{if } p \geq \frac{1}{2} \end{cases}. \quad (10)$$

Note that when $p < 0.5$, we have $\mu > 0$, which indicates that the peak of the density function $p(\theta_k)$ will locate at positive θ_k . As the sign s_k is chosen randomly and set fixed, the corresponding prior distribution $p(w_k)$ will be bimodal (explained in Appendix C). In view of that the distribution of weight (bias excluded) in convolutional and FC layers of a directly trained SNN is usually zero-meaned unimodal (See Fig. 3), we assume $p \geq 0.5$ in the following discussion.

The Grad R algorithm is given in Algorithm 1. It is important to notice that the target sparsity p is not bound to be reached. Since the gradient of the loss function, i.e., $s_k \frac{\partial \mathcal{L}(x_i, y_i, \theta)}{\partial w_k}$, also contributes to the posterior distribution (line 7). Consequently, the choice of these hyperparameters is

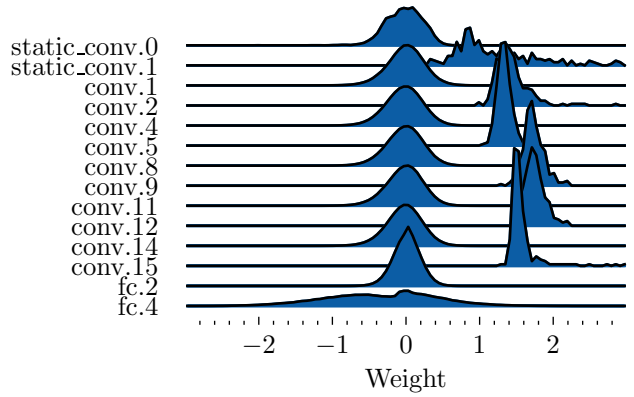


Figure 3: The distribution of the parameters in a trained SNN without pruning. Note that *static_conv.1*, *conv.2*, *conv.5*, *conv.9*, *conv.12* and *conv.15* are all BN layers.

Algorithm 1: Gradient Rewiring with SGD

Input: Input data X , output labels Y , learning rate η , target sparsity p , penalty term α

- 1 Initialize weight w ;
 - 2 $s \leftarrow \text{sign}(w)$, $\theta \leftarrow w \odot s$;
 - 3 Calculate μ according to p , α ;
 - 4 **for** t in $[1, N]$ **do**
 - 5 **for all synaptic parameters** θ_k **and sign** $s_k \in s$ **do**
 - 6 **for input** $x_i \in X$ **and label** $y_i \in Y$ **do**
 - 7 $\theta_k \leftarrow \theta_k - \eta \left(s_k \frac{\partial \mathcal{L}(x_i, y_i, \theta)}{\partial w_k} + \alpha \text{sign}(\theta_k - \mu) \right)$;
 - 8 **end**
 - 9 $w_k \leftarrow s_k \text{ReLU}(\theta_k)$;
 - 10 **end**
 - 11 **end**
-

more a reference rather than a guarantee on the final sparsity. The prior encourages synaptic parameters to gather around μ , wherein the weights are more likely to be zero. Meanwhile, α controls the speed of pruning. The penalty term prevent weights from growing too large, alleviating the concern in Fig. 2(b).

2.4 Detailed Learning Rule in SNNs

Here we apply Grad R to the SNN in Fig. 1 and derive the detailed learning rule. Specifically, the problem is to compute the gradient of the loss function $s_k \frac{\partial \mathcal{L}(x_i, y_i, \theta)}{\partial w_k}$ (line 7 in Algorithm 1). We take a single FC layer without bias parameters as an example. Suppose that m_t and u_t represent the membrane potential of all the LIF neurons after neural dynamics and after the trigger of a spike at time-step t , S_t denotes the vector of the binary spike of all the LIF neurons at time-step t , W denotes the weight matrix of FC layer, $\mathbf{1}$ denotes a vector composed entirely of 1 with the same size as S_t . With these definitions, we can compute the gradient of output spikes S_t with respect to W by unfolding the network over the simulation time-steps, which is similar to the idea of

backpropagation through time (BPTT). We obtain

$$\begin{aligned} \frac{\partial \mathbf{S}_t}{\partial \mathbf{W}} &= \frac{\partial \mathbf{S}_t}{\partial \mathbf{m}_t} \frac{\partial \mathbf{m}_t}{\partial \mathbf{W}}, \\ \frac{\partial \mathbf{u}_t}{\partial \mathbf{W}} &= \frac{\partial \mathbf{u}_t}{\partial \mathbf{S}_t} \frac{\partial \mathbf{S}_t}{\partial \mathbf{W}} + \frac{\partial \mathbf{u}_t}{\partial \mathbf{m}_t} \frac{\partial \mathbf{m}_t}{\partial \mathbf{W}}, \\ \frac{\partial \mathbf{m}_{t+1}}{\partial \mathbf{W}} &= \frac{\partial \mathbf{m}_{t+1}}{\partial \mathbf{u}_t} \frac{\partial \mathbf{u}_t}{\partial \mathbf{W}} + \frac{1}{\tau_m} \mathbf{I}_t^\top, \end{aligned} \quad (11)$$

where

$$\begin{aligned} \frac{\partial \mathbf{u}_t}{\partial \mathbf{S}_t} &= \text{Diag}(u_{\text{rest}} \mathbf{1} - \mathbf{m}_t), \\ \frac{\partial \mathbf{m}_{t+1}}{\partial \mathbf{u}_t} &= 1 - \frac{1}{\tau_m}, \\ \frac{\partial \mathbf{u}_t}{\partial \mathbf{m}_t} &= \frac{\partial \mathbf{u}_t}{\partial \mathbf{S}_t} \frac{\partial \mathbf{S}_t}{\partial \mathbf{m}_t} + \text{Diag}(\mathbf{1} - \mathbf{S}_t). \end{aligned} \quad (12)$$

When calculating $\frac{\partial \mathbf{S}_t}{\partial \mathbf{m}_t}$, the derivative of spike with respect to the membrane potential after charging, we adopt the surrogate gradient method [Wu *et al.*, 2018; Neftci *et al.*, 2019]. Specifically, we utilize shifted ArcTan function $H_1(x) = \frac{1}{\pi} \arctan(\pi x) + \frac{1}{2}$ as a surrogate function of the Heaviside step function $H(\cdot)$, and we have $\frac{\partial \mathbf{S}_t}{\partial \mathbf{m}_t} \approx H_1'(\mathbf{m}_t - u_{\text{th}})$.

3 Experiments

We evaluate the performance of the Grad R algorithm on image recognition benchmarks, namely the MNIST and CIFAR-10 datasets. Our implementation of deep SNN are based on our open-source SNN framework *SpikingJelly* [Fang *et al.*, 2020a].

3.1 Experimental Settings

For MNIST, we consider a shallow fully-connected network (Input-Flatten-800-LIF-10-LIF) similar to previous works. On the CIFAR-10 dataset, we use a convolutional SNN with 6 convolutional layers and 2 FC layers (Input-[[256C3S1-BN-LIF]×3-MaxPool2d2S2]×2-Flatten-Dropout-2048-LIF-100-LIF-AvgPool1d10S10, illustrated in Fig. 1) proposed in our previous work [Fang *et al.*, 2020b]. The first convolutional layer acts as a learnable spike encoder to transform the input static images to spikes. Through this way, we elude from the Poisson encoder that requires a long simulation time-step to restore inputs accurately [Wu *et al.*, 2019b]. The dropout mask is initiated at the first time-step and remains constant within an iteration of mini-batch data, guaranteeing a fixed connective pattern in each iteration [Lee *et al.*, 2020]. All bias parameters are omitted except for BatchNorm (BN) layers. Note that pruning weights in BN layers is equivalent to pruning all the corresponding input channels, which will lead to an unnecessary decline in the expressiveness of SNNs. Moreover, the peak of the probability density function of weights in BN layers are far from zero (see Fig. 3), which suggest the pruning is improper. Based on the above reasons, we exclude all BN layers from pruning in our algorithm.

To facilitate the learning process, we use the Adam optimizer along with our algorithm for joint optimization of connectivity and weight. Furthermore, the gradients of reset $\frac{\partial u_t}{\partial S_t}$




Figure 4: Top-1 accuracy (solid line) on the test set and connectivity (dashed line) during training on CIFAR-10. α is the penalty term in Algorithm 1.

in LIF neurons’ forward steps (Eq. 2) are dropped, i.e. manually set to zero, according to [Zenke and Vogels, 2020]. The choice of all hyperparameters is shown in Table 3 and Table 4.

3.2 Performance

Fig. 4 shows the top-1 accuracy on the test set and connectivity during training on the CIFAR-10 datasets. The improvement of accuracy and network sparsity is simultaneous as the Grad R algorithm can learn the connectivity and weights of the SNN jointly. As the training goes into later stages, the growth of sparsity has little impact on the performance. The sparsity will likely continue to increase, and the accuracy will still maintain, even if we increase the training epochs. The effect of prior is also clearly shown in Fig. 4. Overall, a larger penalty term α will bring higher sparsity and more apparent performance degradation. Aggressive pruning setting like huge α , e.g. $\alpha = 0.005$, will soon prune most of the connections and let the network leaves few parameters to be updated (see the brown line).

As most of the previous works do not conduct experiments on datasets containing complicated patterns, such as CIFAR-10, we also test the algorithm on a similar shallow fully-connected network for the classification task of the MNIST dataset. We further implement Deep R as a baseline, which also can be viewed as a pruning-while-learning algorithm. The highest performance on the test set against connectivity is presented in Fig. 5. The performance of Grad R is comparable to Deep R on MNIST and much better than Deep R on CIFAR-10. As illustrated in Fig. 5(b), the performance of Grad R is close to the performance of the fully connected network when the connectivity is above 10%. Besides, it reaches a classification accuracy of 92.03% (merely $<1\%$ accuracy loss) when constrained to $\sim 5\%$ connectivity. Further, Grad R achieves a classification accuracy of 89.32% ($\sim 3.5\%$ accuracy loss) under unprecedented 0.73% connectivity. These results suggest that there exists extremely high redundancy

Pruning Methods	Training Methods	Arch.	Acc. (%)	Acc. Loss (%)	Connectivity (%)
Threshold-based [Neftci <i>et al.</i> , 2016]	Event-driven CD	2 FC	95.60	-0.60	26.00
Deep R [Bellec <i>et al.</i> , 2018b] ¹	Surrogate Gradient	LSNN ²	93.70	+2.70	12.00
Threshold-based [Shi <i>et al.</i> , 2019]	STDP	1 FC	94.05	≈-19.05 ³	≈30.00
Threshold-based [Rathi <i>et al.</i> , 2019]	STDP	2 layers ⁴	93.20	-1.70	8.00
				+0.03	50.00
ADMM-based [Deng <i>et al.</i> , 2019]	Surrogate Gradient	LeNet-5 like	99.07	-0.43	40.00
				-2.23	25.00
Online APTN [Guo <i>et al.</i> , 2020]	STDP	2 FC	90.40	-3.87	10.00
				-0.36	37.14
				-0.56	13.30
Deep R ⁵	Surrogate Gradient	2 FC	98.92	-2.47	4.18
				-9.67	1.10
				-0.33	25.71
				-0.43	17.94
Grad R	Surrogate Gradient	2 FC	98.92	-2.02	5.63
				-3.55	3.06
				-8.08	1.38

¹ Using sequential MNIST.

² 100 adaptive and 120 regular neurons.

³ A data point of conventional pruning. Since the soft-pruning will only freeze weight during training, bringing no real sparsity.

⁴ The 1st layer is FC and the 2nd layer is a one-to-one layer with lateral inhibition.

⁵ Our implementation of Deep R.

Table 1: Performance comparison between the proposed method and previous work on MNIST dataset.

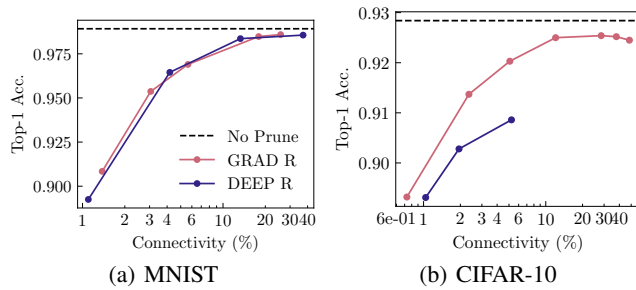


Figure 5: Top-1 accuracies on test set. Different penalty term α is chosen for different sparsity (connectivity). The accuracy of unpruned model is shown as a black dashed line.

in deep SNNs. We compare our algorithm to other state-of-the-art pruning methods of SNNs, and the results are listed in Table 1 and Table 2, from which we can find that the proposed method outperforms previous works.

3.3 Behaviors of Algorithm

Ablation study of prior. We test the effects of prior. The Grad R algorithm without prior ($\alpha = 0$) achieves almost the highest accuracy after pruning, but the connectivity drops the most slowly (the red curve in Fig. 4). The performance gap between unpruned SNN and pruned SNN without prior suggests the minimum accuracy loss using Grad R. It also clarifies that prior is needed if higher sparsity is required.

Distribution drift of parameters. Fig. 6(a) illustrates the distribution of the parameter θ during training ($\alpha = 0.005$).

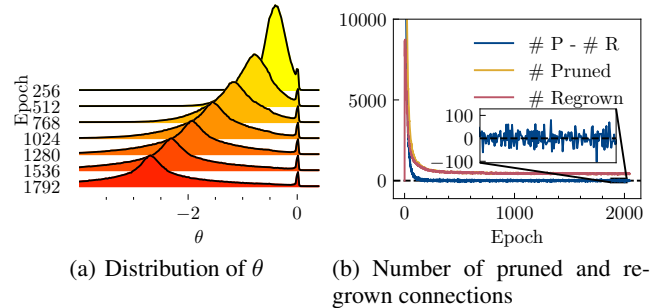


Figure 6: The statistics of the 2nd convolutional layer during training.

Since we use a Laplacian distribution with $\mu < 0$ as prior, the posterior distribution will be dragged to fit the prior.

Competition between pruning and regrowth. A cursory glance of the increasing sparsity over time may mislead one that the pruning process always gains huge ascendancy over the regrowth process. However, Fig. 6(b) sheds light on the fierce competition between these two processes. The pruning takes full advantage of the regrowth merely in the early stages (the first 200 epochs). Afterward, the pruning process only wins by a narrow margin. The competition encourages SNNs to explore different network structures during training.

Firing rate and connectivity. Reduction of firing rate is essential for energy-efficient neuromorphic chips (event-driven hardware) [Oster *et al.*, 2007]. Therefore, we explore the impact of synaptic pruning on the firing rate. As shown in Fig. 7, there is a relatively weak correlation between the distribution

Pruning Methods	Training Methods	Arch.	Acc. (%)	Acc. Loss (%)	Connectivity (%)
ADMM-based [Deng <i>et al.</i> , 2019]	Surrogate Gradient	7 Conv, 2 FC	89.53	-0.38	50.00
				-2.16	25.00
				-3.85	10.00
Deep R ¹	Surrogate Gradient	6 Conv, 2 FC	92.84	-1.98	5.24
				-2.56	1.95
				-3.53	1.04
Grad R	Surrogate Gradient	6 Conv, 2 FC	92.84	-0.30	28.41
				-0.34	12.04
				-0.81	5.08
				-1.47	2.35
				-3.52	0.73

¹ Experiment on CIFAR-10 is not involved in SNNs' pruning work of Deep R [Bellec *et al.*, 2018b]. Therefore, we implement and generalize Deep R to deep SNNs.

Table 2: Performance comparison between the proposed method and previous work on CIFAR-10 dataset.

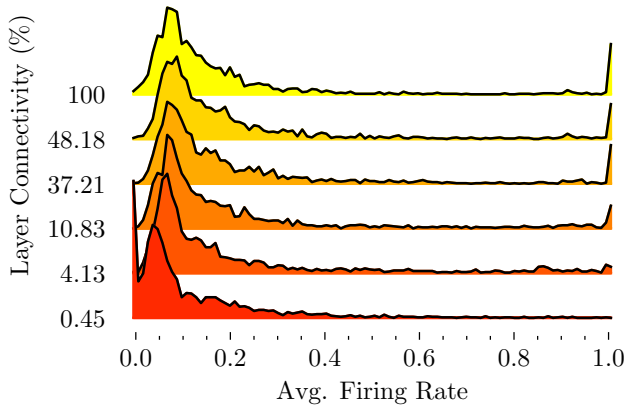


Figure 7: The distribution of the average firing rate of post-synaptic neurons in the 1st FC layer under different connectivity. The firing rate of each neuron is averaged on samples of the whole test dataset and on time domain.

of neurons' average firing rate and connectivity. However, when the connectivity drastically drops extremely low under 1%, the level of neural activities is severely suppressed. The peak on the left side has a significant shift, causing by silent neurons. Meanwhile, the right peak representing saturated neurons vanishes.

4 Conclusion

In this paper, we present Gradient Rewiring, a pruning method for deep SNNs, enabling efficient joint learning of connectivity and network weight. It utilizes the competition between pruning and regrowth to strengthen the exploration of network structures and achieves the minimal loss of SNN's performance on the MNIST and CIFAR-10 dataset so far. Our work reveals that there exists high redundancy in deep SNNs. An interesting future direction is to further investigate and exploit such redundancy to achieve low-power computation on neuromorphic hardware.

Acknowledgements

This work is supported by grants from the National Natural Science Foundation of China under contract No. 62027804, No. 61825101 and No. 62088102.

A Supporting Lemma

Lemma 1. *If there exists a constant C , such that $\forall k, \theta_k \geq C$, then there exists a constant $A > 0$ such that*

$$\|\nabla_{\theta} \bar{\mathcal{L}}\|^2 \leq A \langle \nabla_{\theta} \mathcal{L}, \nabla_{\theta} \bar{\mathcal{L}} \rangle. \quad (13)$$

Proof. Since $w_k = \frac{s_k}{\beta} \log(1 + \exp(\beta\theta_k))$, we have $\frac{\partial w_k}{\partial \theta_k} = s_k \sigma(\beta\theta_k)$, where $\sigma(x) = \frac{1}{1 + \exp(-x)}$ is sigmoid function.

Note that $\sigma(x) > 0$ is monotonically increasing, we have

$$\begin{aligned} \frac{\|\nabla_{\theta} \bar{\mathcal{L}}\|^2}{\langle \nabla_{\theta} \mathcal{L}, \nabla_{\theta} \bar{\mathcal{L}} \rangle} &= \frac{\sum_k \left(s_k \frac{\partial \mathcal{L}}{\partial w_k} \right)^2}{\sum_k \left[s_k^2 \sigma(\beta\theta_k) \left(\frac{\partial \mathcal{L}}{\partial w_k} \right)^2 \right]} \\ &= \frac{\sum_k \left(\frac{\partial \mathcal{L}}{\partial w_k} \right)^2}{\sum_k \left[\sigma(\beta\theta_k) \left(\frac{\partial \mathcal{L}}{\partial w_k} \right)^2 \right]} \\ &\leq \frac{\sum_k \left(\frac{\partial \mathcal{L}}{\partial w_k} \right)^2}{\sum_k \left[\sigma(\beta C) \left(\frac{\partial \mathcal{L}}{\partial w_k} \right)^2 \right]} \\ &= \frac{1}{\sigma(\beta C)}. \end{aligned} \quad (14)$$

Thus there exists a constant $A = \frac{1}{\sigma(\beta C)}$ such that $\|\nabla_{\theta} \bar{\mathcal{L}}\|^2 \leq A \langle \nabla_{\theta} \mathcal{L}, \nabla_{\theta} \bar{\mathcal{L}} \rangle$. \square

Lemma 1 also suggests that $\langle \nabla_{\theta} \mathcal{L}, \nabla_{\theta} \bar{\mathcal{L}} \rangle \geq 0$. When $\|\nabla_{\theta} \bar{\mathcal{L}}(\theta^t)\| > 0$, the inner product $\langle \nabla_{\theta} \mathcal{L}, \nabla_{\theta} \bar{\mathcal{L}} \rangle$ is strictly positive, thus the angle between the accurate gradient $\nabla_{\theta} \mathcal{L}$ and the redefined gradient $\nabla_{\theta} \bar{\mathcal{L}}$ is below 90° .

Parameters	Descriptions	CIFAR-10	MNIST
N	# Epoch	2048	512
Batch Size	-	16	128
T	# Timestep	8	
τ_m	Membrane Constant	2.0	
u_{th}	Firing Threshold	1.0	
u_{rest}	Resting Potential	0.0	
p	Target Sparsity	95%	
η	Learning Rate	1e-4	
β_1, β_2	Adam Decay	0.9, 0.999	

Table 3: Choice of hyperparameters for both pruning algorithm.

Dataset	Algorithm	Penalty	Connectivity (%)
CIFAR-10	Grad R	0	48.35
		10^{-4}	37.79
		2×10^{-4}	28.41
		5×10^{-4}	12.04
		10^{-3}	5.08
		2×10^{-3}	2.35
	Deep R	0	5.24
		5×10^{-4}	1.95
		10^{-3}	1.04
		5×10^{-3}	0.73
MNIST	Grad R	10^{-2}	25.71
		10^{-1}	17.94
		5×10^{-2}	5.63
		2×10^{-1}	3.06
	Deep R	0	1.38
		10^{-2}	37.14
		5×10^{-2}	13.30
		2×10^{-1}	4.18
			1.10

 Table 4: Choice of penalty term for Grad R (target sparsity $p = 0.95$) and Deep R (maximum sparsity $p = 1$ for CIFAR-10, $p = 0.99$ for MNIST) with corresponding final connectivity.

B Detailed Setting of Hyperparameters

The hyperparameters of Grad R and our implementation of Deep R are illustrated in Table 3 and Table 4 for reproducing results in this paper. Note that for Deep R, zero penalty still introduces low connectivity.

C Range of Target Sparsity

Recall that the network weights w are determined by corresponding signs s and synaptic parameters θ , if the signs are randomly assigned and the distribution of θ_k is given, we are able to deduce the distribution of w_k . For this algorithm, we denote the cumulative distribution function (CDF) of θ_k is $F_{\theta_k}(\theta) := P[\theta_k < \theta]$. So we obtain

$$\begin{aligned}
 P[\text{ReLU}(\theta_k) < \theta] &= P[\theta > \max\{\theta_k, 0\}] \\
 &= \begin{cases} 0, & \text{if } \theta \leq 0 \\ F_{\theta_k}(\theta), & \text{if } \theta > 0 \end{cases} \quad (15)
 \end{aligned}$$

Since the distribution of sign is similar to the Bernoulli distribution, the CDF of w_k can also be derived as

$$\begin{aligned}
 F_{w_k}(w) &:= P[w_k < w] \\
 &= P[s_k \text{ReLU}(\theta_k) < w] \\
 &= \sum_{s \in \{1, -1\}} P[s_k \text{ReLU}(\theta_k) < w \mid s_k = s] P[s_k = s] \\
 &= \frac{1}{2} (P[\text{ReLU}(\theta_k) < w] + P[-\text{ReLU}(\theta_k) < w]) \\
 &= \frac{1}{2} (P[\text{ReLU}(\theta_k) < w] + 1 - P[\text{ReLU}(\theta_k) < -w]) \\
 &\quad - P[\text{ReLU}(\theta_k) = -w]) \\
 &= \begin{cases} \frac{1}{2}(1 - F_{\theta_k}(-w)), & \text{if } w < 0 \\ \frac{1}{2}(1 - P(\theta_k \leq 0)), & \text{if } w = 0. \\ \frac{1}{2}(1 + F_{\theta_k}(w)), & \text{if } w > 0 \end{cases} \quad (16)
 \end{aligned}$$

Note that $F'_{w_k}(w) = \frac{1}{2} F'_{\theta_k}(|w|)$, which indicates the probability density function (PDF) of w_k is an even function. Therefore, the distribution of w is symmetric with respect to the w -axis. We hope it to be unimodal, which matches our observation of weight distribution in SNNs (See Fig. 3). So there should not be any peak in $F'_{w_k}(w)$ when $w > 0$, otherwise we will get the other peak symmetric to it.

Note that the requirement that $F'_{w_k}(w) = \frac{1}{2} F'_{\theta_k}(|w|)(w > 0)$ has no peak is equivalent to that $F'_{\theta_k}(w)(w > 0)$ has no peak. As $F'_{\theta_k}(\cdot)$ is the prior of θ_k , which is a Laplacian distribution where the PDF has only one peak at μ , we must constrain $\mu \leq 0$. This explains why we require target sparsity $p \geq 0.5$.

References

- [Bellec *et al.*, 2018a] Guillaume Bellec, David Kappel, Wolfgang Maass, and Robert Legenstein. Deep rewiring: Training very sparse deep networks. In *ICLR*, 2018.
- [Bellec *et al.*, 2018b] Guillaume Bellec, Darjan Salaj, Anand Subramoney, Robert Legenstein, and Wolfgang Maass. Long short-term memory and learning-to-learn in networks of spiking neurons. In *NeurIPS*, page 795–805. Curran Associates Inc., 2018.
- [Chklovskii *et al.*, 2004] D. B. Chklovskii, B. W. Mel, and K. Svoboda. Cortical rewiring and information storage. *Nature*, 431(7010):782–788, 2004.
- [Deng *et al.*, 2019] 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. *arXiv preprint arXiv:1911.00822*, 2019.
- [Fang *et al.*, 2020a] Wei Fang, Yanqi Chen, Jianhao Ding, Ding Chen, Zhaofei Yu, Huihui Zhou, Yonghong Tian, and other contributors. Spikingjelly. <https://github.com/fangwei123456/spikingjelly>, 2020. Accessed: 2020-11-09.
- [Fang *et al.*, 2020b] Wei Fang, Zhaofei Yu, Yanqi Chen, Timothée Masquelier, Tiejun Huang, and Yonghong Tian.

- Incorporating learnable membrane time constant to enhance learning of spiking neural networks. *arXiv preprint arXiv:2007.05785*, 2020.
- [Frankle and Carbin, 2019] Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. In *ICLR*, 2019.
- [Gerstner and Kistler, 2002] Wulfram Gerstner and Werner M. Kistler. *Spiking Neuron Models: Single Neurons, Populations, Plasticity*. Cambridge University Press, 2002.
- [Guo *et al.*, 2020] Wenzhe Guo, Mohammed E. Fouda, Hasan Erdem Yantir, Ahmed M. Eltawil, and Khaled Nabil Salama. Unsupervised adaptive weight pruning for energy-efficient neuromorphic systems. *Frontiers in Neuroscience*, 14:1189, 2020.
- [Gütig and Sompolinsky, 2006] Robert Gütig and Haim Sompolinsky. The tempotron: a neuron that learns spike timing-based decisions. *Nature Neuroscience*, 9(3):420–428, 2006.
- [Huttenlocher and Nelson, 1994] Peter R Huttenlocher and Ch A Nelson. Synaptogenesis, synapse elimination, and neural plasticity in human cerebral cortex. In *Threats to Optimal Development: Integrating Biological, Psychosocial, and Social Risk Factors: The Minnesota Symposia on Child Psychology*, volume 27, pages 35–54, 1994.
- [Kappel *et al.*, 2015] David Kappel, Stefan Habenschuss, Robert Legenstein, and Wolfgang Maass. Network plasticity as Bayesian inference. *PLOS Computational Biology*, 11(11):1–31, 11 2015.
- [Lee *et al.*, 2020] Chanky Lee, Syed Shakib Sarwar, Priyadarshini Panda, Gopalakrishnan Srinivasan, and Kaushik Roy. Enabling spike-based backpropagation for training deep neural network architectures. *Frontiers in Neuroscience*, 14:119, 2020.
- [Liu *et al.*, 2018] Chen Liu, Guillaume Bellec, Bernhard Vogginger, David Kappel, Johannes Partzsch, Felix Neumärker, Sebastian Höppner, Wolfgang Maass, Steve B. Furber, Robert Legenstein, and Christian G. Mayr. Memory-efficient deep learning on a SpiNNaker 2 prototype. *Frontiers in Neuroscience*, 12:840, 2018.
- [Liu *et al.*, 2020] Yanchen Liu, Kun Qian, Shaogang Hu, Kun An, Sheng Xu, Xitong Zhan, JJ Wang, Rui Guo, Yuancong Wu, Tu-Pei Chen, et al. Application of deep compression technique in spiking neural network chip. *IEEE Transactions on Biomedical Circuits and Systems*, 14(2):274–282, 2020.
- [Maass, 1997] Wolfgang Maass. Networks of spiking neurons: the third generation of neural network models. *Neural Networks*, 10(9):1659–1671, 1997.
- [Martinelli *et al.*, 2020] Flavio Martinelli, Giorgia Dellaferera, Pablo Mainar, and Milos Cernak. Spiking neural networks trained with backpropagation for low power neuromorphic implementation of voice activity detection. In *ICASSP*, pages 8544–8548, 2020.
- [Neftci *et al.*, 2016] Emre O. Neftci, Bruno U. Pedroni, Siddharth Joshi, Maruan Al-Shedivat, and Gert Cauwenberghs. Stochastic synapses enable efficient brain-inspired learning machines. *Frontiers in Neuroscience*, 10:241, 2016.
- [Neftci *et al.*, 2019] 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.
- [Oster *et al.*, 2007] Matthias Oster, Rodney Douglas, and Shih-Chii Liu. Quantifying input and output spike statistics of a winner-take-all network in a vision system. In *IEEE International Symposium on Circuits and Systems*, pages 853–856, 2007.
- [Petzoldt and Sigrist, 2014] Astrid G. Petzoldt and Stephan J. Sigrist. Synaptogenesis. *Current Biology*, 24(22):1076–1080, 2014.
- [Qi *et al.*, 2018] Yu Qi, Jiangrong Shen, Yueming Wang, Huajin Tang, Hang Yu, Zhaohui Wu, and Gang Pan. Jointly learning network connections and link weights in spiking neural networks. In *IJCAI*, pages 1597–1603, 7 2018.
- [Rathi *et al.*, 2019] Nitin Rathi, Priyadarshini Panda, and Kaushik Roy. STDP-based pruning of connections and weight quantization in spiking neural networks for energy-efficient recognition. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 38(4):668–677, 2019.
- [Roy *et al.*, 2019] Kaushik Roy, Akhilesh Jaiswal, and Priyadarshini Panda. Towards spike-based machine intelligence with neuromorphic computing. *Nature*, 575(7784):607–617, 2019.
- [Shi *et al.*, 2019] Yuhan Shi, Leon Nguyen, Sangheon Oh, Xin Liu, and Duygu Kuzum. A soft-pruning method applied during training of spiking neural networks for in-memory computing applications. *Frontiers in Neuroscience*, 13:405, 2019.
- [Wu *et al.*, 2018] 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.
- [Wu *et al.*, 2019a] Doudou Wu, Xianghong Lin, and Pangao Du. An adaptive structure learning algorithm for multi-layer spiking neural networks. In *15th International Conference on Computational Intelligence and Security (CIS)*, pages 98–102, 2019.
- [Wu *et al.*, 2019b] Yujie Wu, Lei Deng, Guoqi Li, Jun Zhu, Yuan Xie, and Luping Shi. Direct training for spiking neural networks: Faster, larger, better. *AAAI*, 33(01):1311–1318, July 2019.
- [Zenke and Vogels, 2020] Friedemann Zenke and Tim P. Vogels. The remarkable robustness of surrogate gradient learning for instilling complex function in spiking neural networks. *bioRxiv*, 2020.