

Deep activity propagation via weight initialization in spiking neural networks

Aurora Micheli, Olaf Booij, Jan van Gemert, Nergis Tomen

• What is a good weight initialization for a spiking neural network (SNN)?



• What is a good weight initialization for a spiking neural network (SNN)?





- What is a good weight initialization for a spiking neural network (SNN)? ٠
- SNNs are often initialized following standard strategies designed for conventional networks (ANNs) ٠

Xavier Glorot [1] \rightarrow Sigmoid Kaiming He [2] \rightarrow ReLu

[1] X. Glorot and Y. Bengio. Understanding the difficulty of training deep feedforward neural networks [2] K. He, X. Zhang, S. Ren, J. Sun. Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification



- What is a good weight initialization for a spiking neural network (SNN)?
- SNNs are often initialized following standard strategies designed for conventional networks (ANNs)



[1] X. Glorot and Y. Bengio. Understanding the difficulty of training deep feedforward neural networks
 [2] K. He, X. Zhang, S. Ren, J. Sun. Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification

TUDelft

- What is a good weight initialization for a spiking neural network (SNN)?
- SNNs are often initialized following standard strategies designed for conventional networks (ANNs)



[1] X. Glorot and Y. Bengio. Understanding the difficulty of training deep feedforward neural networks
 [2] K. He, X. Zhang, S. Ren, J. Sun. Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification

TUDelft

In this paper

1. Analytically derive a weight initialization method tailored for SNNs

 Empirically validate on simulations that our method ensures stable spike propagation in deep SNNs, preventing dissipation or amplification

3. Show that our weight initialization enhances training speed and improves classification accuracy on 4 different datasets



• Feed-forward fully-connected SNN initialized at t=0. For a generic layer *l* :

$$egin{aligned} oldsymbol{u}_l &= oldsymbol{w}_l oldsymbol{x}_l \ oldsymbol{x}_l &= f(oldsymbol{u}_{l-1}) \end{aligned}$$

$$f(u_{l-1}) = egin{cases} 1, & ext{if} \;\; u_{l-1} > heta \ 0, & ext{if} \;\; u_{l-1} < heta \end{cases}$$

$oldsymbol{u}_l~\in~\mathbb{R}^m$: membrane potentials
$oldsymbol{w}_l ~\in~ \mathbb{R}^{m imes m}$: weight matrix
$oldsymbol{x}_l \in \{0,1\}^n$: input spikes
f	: activation function
heta	: firing threshold



• Feed-forward fully-connected SNN initialized at t=0. For a generic layer *l* :

$$egin{aligned} oldsymbol{u}_l &= oldsymbol{w}_l oldsymbol{x}_l \ oldsymbol{x}_l &= f(oldsymbol{u}_{l-1}) \end{aligned}$$

$f(u_{l-1}) = \langle$	∫ 1,	if	$u_{l-1} > \theta$
	0,	if	$u_{l-1} < \theta$

$oldsymbol{u}_l~\in$	\mathbb{R}^m : membrane potentials
$oldsymbol{w}_l~\in$	$\mathbb{R}^{m imes n}$: weight matrix
$oldsymbol{x}_l \in \{$	$0,1\}^n$: input spikes
f	: activation function
θ	: firing threshold

[2] Good initialization method: avoid reducing or amplifying the magnitudes of the input signals ...

[2] K. He, X. Zhang, S. Ren, J. Sun. Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification



• Feed-forward fully-connected SNN initialized at t=0. For a generic layer *l* :

$$u_{l} = w_{l}x_{l}$$

$$x_{l} = f(u_{l-1})$$

$$f(u_{l-1}) = \begin{cases} 1, & \text{if } u_{l-1} > \theta \\ 0, & \text{if } u_{l-1} < \theta \end{cases}$$

$$f(u_{l-1}) = \begin{cases} 1, & \text{if } u_{l-1} > \theta \\ 0, & \text{if } u_{l-1} < \theta \end{cases}$$

$$f(u_{l-1}) = \begin{cases} 1, & \text{if } u_{l-1} > \theta \\ 0, & \text{if } u_{l-1} < \theta \end{cases}$$

$$f(u_{l-1}) = \begin{cases} 1, & \text{if } u_{l-1} > \theta \\ 0, & \text{if } u_{l-1} < \theta \end{cases}$$

$$f(u_{l-1}) = \begin{cases} 1, & \text{if } u_{l-1} > \theta \\ 0, & \text{if } u_{l-1} < \theta \end{cases}$$

$$f(u_{l-1}) = \begin{cases} 1, & \text{if } u_{l-1} > \theta \\ 0, & \text{if } u_{l-1} < \theta \end{cases}$$

$$f(u_{l-1}) = \begin{cases} 1, & \text{if } u_{l-1} > \theta \\ 0, & \text{if } u_{l-1} < \theta \end{cases}$$

$$f(u_{l-1}) = \begin{cases} 1, & \text{if } u_{l-1} > \theta \\ 0, & \text{if } u_{l-1} < \theta \end{cases}$$

$$f(u_{l-1}) = \begin{cases} 1, & \text{if } u_{l-1} > \theta \\ 0, & \text{if } u_{l-1} < \theta \end{cases}$$

$$f(u_{l-1}) = \begin{cases} 1, & \text{if } u_{l-1} > \theta \\ 0, & \text{if } u_{l-1} < \theta \end{cases}$$

$$f(u_{l-1}) = \begin{cases} 1, & \text{if } u_{l-1} > \theta \\ 0, & \text{if } u_{l-1} < \theta \end{cases}$$

$$f(u_{l-1}) = \begin{cases} 1, & \text{if } u_{l-1} > \theta \\ 0, & \text{if } u_{l-1} < \theta \end{cases}$$

$$f(u_{l-1}) = \begin{cases} 1, & \text{if } u_{l-1} > \theta \\ 0, & \text{if } u_{l-1} < \theta \end{cases}$$

$$f(u_{l-1}) = \begin{cases} 1, & \text{if } u_{l-1} > \theta \\ 0, & \text{if } u_{l-1} < \theta \end{cases}$$

$$f(u_{l-1}) = \begin{cases} 1, & \text{if } u_{l-1} > \theta \\ 0, & \text{if } u_{l-1} < \theta \end{cases}$$

$$f(u_{l-1}) = \begin{cases} 1, & \text{if } u_{l-1} < \theta \\ 0, & \text{if } u_{l-1} < \theta \end{cases}$$

$$f(u_{l-1}) = \begin{cases} 1, & \text{if } u_{l-1} < \theta \\ 0, & \text{if } u_{l-1} < \theta \end{cases}$$

$$g(u_{l-1}) = \begin{cases} 1, & u_{l-1} < \theta \\ 0, & u_{l-1} < \theta \\ 0, & u_{l-1} < \theta \end{cases}$$

$$g(u_{l-1}) = \begin{cases} 1, & u_{l-1} < \theta \\ 0, & u_{l-1} < \theta \\$$

... by keeping the variance of the input to each layer constant

[2] K. He, X. Zhang, S. Ren, J. Sun. Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification

• Feed-forward fully-connected SNN initialized at t=0. For a generic layer *l* :

$$u_{l} = w_{l}x_{l}$$

$$x_{l} = f(u_{l-1})$$

$$f(u_{l-1}) = \begin{cases} 1, & \text{if } u_{l-1} > \theta \\ 0, & \text{if } u_{l-1} < \theta \end{cases}$$

$$I_{l} \in \mathbb{R}^{m \times n} : \text{weight matrix}$$

$$x_{l} \in \{0, 1\}^{n} : \text{input spikes}$$

$$f : \text{activation function}$$

$$\theta : \text{firing threshold}$$

$$[2] \text{ Good initialization method: avoid reducing or amplifying the magnitudes of the input signals ...}$$

$$Var[u_{l}] = n_{l}Var[w_{l}]E[x_{l}^{2}] = \text{constant}$$

$$n_{l} : \text{input size}$$

... by keeping the variance of the input to each layer constant

[2] K. He, X. Zhang, S. Ren, J. Sun. Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification

$$\operatorname{Var}[u_l] = n_l \operatorname{Var}[w_l] E[x_l^2]$$
 [2



[2] K. He, X. Zhang, S. Ren, J. Sun. Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification







[2] K. He, X. Zhang, S. Ren, J. Sun. Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification

″UDelft





[2] K. He, X. Zhang, S. Ren, J. Sun. Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification

TUDelft



[2] K. He, X. Zhang, S. Ren, J. Sun. Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification

TUDelft

- 100 layers fully-connected SNN with Leaky Integrate-and-Fire (LIF) neurons
- Input I_0 drawn from $\mathcal{N}(\mu = 0, \sigma^2 = 1), \theta \in [0, 1]$



With the proposed method, unlike Kaiming, the variance of the membrane potentials stays constant across layers

- 100 layers fully-connected SNN with Leaky Integrate-and-Fire (LIF) neurons
- Input I_0 drawn from $\mathcal{N}(\mu = 0, \sigma^2 = 1), \theta \in [0, 1]$



With the proposed method, unlike Kaiming, the variance of the membrane potentials stays constant across layers

Delft

- 100 layers fully-connected SNN with Leaky Integrate-and-Fire (LIF) neurons
- Input I_0 drawn from $\mathcal{N}(\mu = 0, \sigma^2 = 1)$



Unlike other commonly used schemes, our initialization enables spiking activity to propagate from input to output

- 100 layers fully-connected SNN with Leaky Integrate-and-Fire (LIF) neurons
- Input I_0 drawn from $\mathcal{N}(\mu = 0, \sigma^2 = 1)$



Unlike other commonly used schemes, our initialization enables spiking activity to propagate from input to output





- 100 layers fully-connected SNN of Leaky Integrate-and-fire (LIF) neurons with soft reset
- Input I_0 drawn from $\mathcal{N}(\mu=0,\sigma^2=1)$
- Time steps T=20

While our method doesn't explicitly account for time, it approximately conserves variance and spike count





- 100 layers fully-connected SNN of Leaky Integrate-and-fire (LIF) neurons with soft reset
- Input I_0 drawn from $\mathcal{N}(\mu=0,\sigma^2=1)$
- Time steps T=20

While our method doesn't explicitly account for time, it approximately conserves variance and spike count





3. Evaluation on classification datasets

- Image classification on 4 datasets: MNIST, Fashion-MNIST, Neuromorphic-MNIST, CIFAR10
- Fully-connected (MLP) and Convolutional (CNN) SNNs trained with backpropagation



Our weight initialization can improve the classification accuracy with only T = 3 time steps



Conclusions

- Derive a new variance-conserving weight initialization method tailored to the activation function of SNNs
- Numerically validate that our method ensures stable spike propagation in deep SNNs, unlike baseline methods
- Show effectiveness of proper weight initialization on training across four image classification datasets



Conclusions

- Derive a new variance-conserving weight initialization method tailored to the activation function of SNNs
- Numerically validate that our method ensures stable spike propagation in deep SNNs, unlike baseline methods
- Show effectiveness of proper weight initialization on training across four image classification datasets

Limitations & future work

- Our method does not explicitly account for temporal dynamics \rightarrow extend theory
- Validation on more complex architectures and datasets



Conclusions

- Derive a new variance-conserving weight initialization method tailored to the activation function of SNNs
- Numerically validate that our method ensures stable spike propagation in deep SNNs, unlike baseline methods
- Show effectiveness of proper weight initialization on training across four image classification datasets

Limitations & future work

- Our method does not explicitly account for temporal dynamics \rightarrow extend theory
- Validation on more complex architectures and datasets



Thank you!

- 100 layers fully-connected SNN of Leaky Integrate-and-fire (LIF) neurons with soft reset
- Input I_0 drawn from $\mathcal{N}(\mu=0,\sigma^2=1)$



While our method doesn't explicitly account for time, it approximately conserves variance and spike count