

Fully Spiking Variational Autoencoder

Hiromichi Kamata¹, Yusuke Mukuta^{1,2}, Tatsuya Harada^{1,2}

¹ The University of Tokyo

² RIKEN

{kamata, mukuta, harada}@mi.t.u-tokyo.ac.jp

Abstract

Spiking neural networks (SNNs) can be run on neuromorphic devices with ultra-high speed and ultra-low energy consumption because of their binary and event-driven nature. Therefore, SNNs are expected to have various applications, including as generative models being running on edge devices to create high-quality images. In this study, we build a variational autoencoder (VAE) with SNN to enable image generation. VAE is known for its stability among generative models; recently, its quality advanced. In vanilla VAE, the latent space is represented as a normal distribution, and floating-point calculations are required in sampling. However, this is not possible in SNNs because all features must be binary time series data. Therefore, we constructed the latent space with an autoregressive SNN model, and randomly selected samples from its output to sample the latent variables. This allows the latent variables to follow the Bernoulli process and allows variational learning. Thus, we build the Fully Spiking Variational Autoencoder where all modules are constructed with SNN. To the best of our knowledge, we are the first to build a VAE only with SNN layers. We experimented with several datasets, and confirmed that it can generate images with the same or better quality compared to conventional ANNs. The code is available at <https://github.com/kamata1729/FullySpikingVAE>.

Introduction

Recently, artificial neural networks (ANNs) have been evolving rapidly, and have achieved considerable success in computer vision and NLP. However, ANNs often require significant computational resources, which is a challenge in situations where computational resources are limited, such as on edge devices.

Spiking neural networks (SNNs) are neural networks that more accurately mimic the structure of a biological brain than ANNs; notably, SNNs are referred to as the third generation of artificial intelligence (Maass 1997). In a SNN, all information is represented as binary time series data, and is driven by event-based processing. Therefore, SNNs can run with ultra-high speed and ultra-low energy consumption on neuromorphic devices, such as Loihi (Davies et al. 2018),

Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

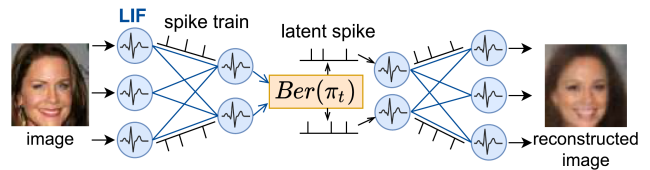


Figure 1: Illustration of our FSVAE. Entire model is constructed with SNNs. All features are represented as spike trains, and the latent spike trains follow Bernoulli processes.

TrueNorth (Akopyan et al. 2015), and Neurogrid (Benjamin et al. 2014). For example, on TrueNorth, the computational time is approximately 1/100 lower and the energy consumption is approximately 1/100,000 times lower than on conventional ANNs (Cassidy et al. 2014).

With the recent breakthroughs on ANNs, research on SNNs has been progressing rapidly. Additionally, SNNs are outperforming ANNs in accuracy in MNIST, CIFAR10, and ImageNet classification tasks (Zheng et al. 2021; Zhang and Li 2020). Moreover, SNNs are used for object detection (Kim et al. 2020), sound classification (Wu et al. 2018), optical flow estimation (Lee et al. 2020); however, their applications are still limited.

In particular, image generation models based on SNNs have not been studied sufficiently. Spiking GAN (Kotariya and Ganguly 2021) built a generator and discriminator with shallow SNNs, and generated images of handwritten digits by adversarial learning. However, its generation quality was low, and some undesired images were generated that could not be interpreted as numbers. In (Skatchkovsky, Simeone, and Jang 2021), SNN was used as the encoder and ANN as the decoder to build a VAE (Kingma and Welling 2014); however, the main focus of their research was efficient spike encoding, and not the image generation task.

In ANNs, image generation models have been extensively studied and can generate high-quality images (Razavi, van den Oord, and Vinyals 2019; Karras et al. 2020). However, in general, image generation models are computationally expensive, and some problems must be solved for edge devices, or for real-time generation. If SNNs can generate images comparably to ANNs, their high speed and low energy consumption can solve these problems.

Therefore, we propose Fully Spiking Variational Autoen-

coder (FSVAE), which can generate images with the same or better quality than ANN. VAEs are known for their stability among generative models and are related to the learning mechanism of the biological brain (Han et al. 2018). Hence, it is compatible to build VAE with SNN. In our FSVAE, we built the entire model in SNN, so that it can be implemented in a neuromorphic device in the future. We conducted experiments using MNIST (Deng 2012), FashionMNIST (Xiao, Rasul, and Vollgraf 2017), CIFAR10 (Krizhevsky and Hinton 2009), and CelebA (Liu et al. 2015), and confirmed that **FSVAE can generate images of equal or better quality than ANN VAE of the same structure**. FSVAE can be implemented on neuromorphic devices in the future, and is expected to improve in terms of speed and energy consumption.

The most difficult aspect of creating VAEs in SNNs is how to create the latent space. In ANN VAEs, the latent space is often represented as a normal distribution. However, within the framework of SNNs, sampling from a normal distribution is not possible because all features must be binary time series data. Therefore, we propose the autoregressive Bernoulli spike sampling. First, we incorporated the idea of VRNN (Chung et al. 2015) into SNN, and built prior and posterior models with autoregressive SNNs. The latent variables are randomly selected from the output of the autoregressive SNNs, which enables sampling from the Bernoulli processes. This can be realized on neuromorphic devices because it does not require floating-point calculations during sampling as in ANNs, and sampling using a random number generator is possible on actual neuromorphic devices (Wen et al. 2016; Davies et al. 2018). In addition, the latent variables can be sampled sequentially; thus, they can be input to the decoder incrementally, which saves time.

The main contributions of this study are summarized as follows.

- We propose the autoregressive Bernoulli spike sampling, which uses autoregressive SNNs and constructs the latent space as Bernoulli processes. This sampling method is feasible within the framework of SNN.
- We propose Fully Spiking Variational Autoencoder (FSVAE), where all modules are constructed in the SNN.
- We experimented with multiple datasets; FSVAE could generate images of equal or better quality than ANN VAE of the same architecture.

Related Work

Development of SNNs

SNNs are neural networks that accurately mimic the structure of the biological brain. In the biological brain, information is transmitted as spike trains (binary time series data with only on/off). This information is transmitted between neurons via synapses, and subsequently, the neuron’s membrane potential changes. When it exceeds a threshold, it fires and becomes a spike train to the next neuron.

SNNs mimic these characteristics of the biological brain, modeling biological neurons using differential equations and representing all features as spike trains. This allows

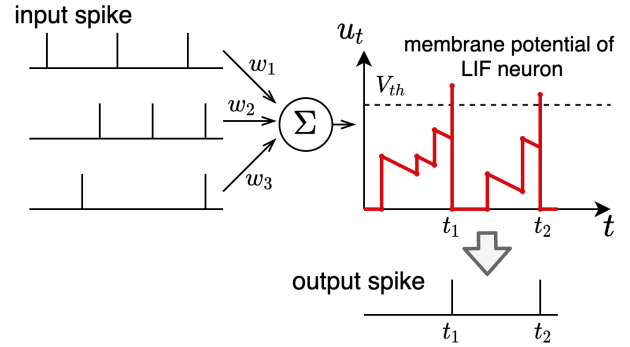


Figure 2: LIF neuron. When spike trains of the previous layer neurons enter, the internal membrane potential u_t changes by Eq. (2). If u_t exceeds V_{th} , it fires a spike $o_t = 1$; otherwise, $o_t = 0$.

SNNs to run faster and asynchronously, because they require fewer floating-point computations and only need computations when the input spike arrives. SNNs can be considered as recurrent neural networks (RNNs) with the membrane potential as its internal state.

Learning algorithms for SNNs have been studied extensively recently. (Diehl and Cook 2015) used a two-layer SNN to recognize MNIST with STDP, an unsupervised learning rule, and achieved 95% accuracy. Later, (Wu et al. 2019) made it possible to train a deep SNN with backpropagation. Recently, (Zhang and Li 2020) exceeded the ANN’s accuracy for MNIST and CIFAR10 (Krizhevsky and Hinton 2009) in only 5 timesteps (length of spike trains). (Zheng et al. 2021) has even higher accuracy in 2 timesteps for CIFAR10 and ImageNet (Deng et al. 2009).

Spike Neuron Model

Although there are several learning algorithms for SNNs, in this study, we follow (Zheng et al. 2021), which currently has the highest recognition accuracy. First, as a neuron model, we use the iterative leaky integrate-and-fire (LIF) model (Wu et al. 2019), which is a LIF model (Stein and Hodgkin 1967) solved using the Euler method.

$$u_t = \tau_{decay} u_{t-1} + x_t \quad (1)$$

where u_t is a membrane potential, x_t is a presynaptic input, and τ_{decay} is a fixed decay factor.

When u_t exceeds a certain threshold V_{th} , the neuron fires and outputs $o_t = 1$. Then, u_t is reset to $u_{rest} = 0$. This can be written as follows:

$$u_{t,n} = \tau_{decay} u_{t-1,n} (1 - o_{t-1,n}) + x_{t,n-1} \quad (2)$$

$$o_{t,n} = H(u_{t,n} - V_{th}) \quad (3)$$

Here, $u_{t,n}$ is the membrane potential of the n th layer, and $o_{t,n}$ is its binary output. H is the heaviside step function. Input $x_{t,n}$ is described as a weighted sum of spikes from neurons in the previous layer, $x_{t,n-1} = \sum_j w^j o_{t,n-1}^j$. By changing the connection way of w^j , we can implement convolution layers, FC layers, etc.

The next step is to enable learning with backpropagation. As Eq. (3) is non-differentiable, we approximate it as follows:

$$\frac{\partial o_{t,n}}{\partial u_{t,n}} = \frac{1}{a} \text{sign} \left(|u_{t,n} - V_{\text{th}}| < \frac{a}{2} \right) \quad (4)$$

Variational Autoencoder

Variational Autoencoder (VAE) (Kingma and Welling 2014) is a generative model that explicitly assumes a distribution of the latent variable z over input x . Typically, the distribution $p(x|z)$ is represented by deep neural networks, so its inverse transformation is approximated by a simple approximate posterior $q(z|x)$. This allows us to calculate the evidence lower bound (ELBO) of the log likelihood.

$$\log p(x) \geq \mathbb{E}_{q(z|x)} [\log p(x|z)] - \text{KL}[q(z|x)||p(z)] \quad (5)$$

where $\text{KL}[Q||P]$ is the Kullback–Leibler (KL) divergence for distributions Q and P . In $q(z|x)$, reparameterization trick is used to sample from $\mathcal{N}(\mu(x), \text{diag}(\sigma(x)^2))$.

VAEs have stable learning among generative models, and can be applied to various tasks, such as anomaly detection (An and Cho 2015). As VAE could generate high-quality images (Razavi, van den Oord, and Vinyals 2019; Vahdat and Kautz 2020), we aimed to build a VAE using SNN.

Variational Recurrent Neural Network

Variational Recurrent Neural Network (VRNN) (Chung et al. 2015) is a VAE for time series data. Its posterior and prior distributions are set as follows:

$$q(z_{1:T}|\mathbf{x}_{1:T}) = \prod_{t=1}^T q(z_t|\mathbf{x}_{\leq t}, \mathbf{z}_{<t}) \quad (6)$$

$$p(z_{1:T}|\mathbf{x}_{1:T-1}) = \prod_{t=1}^T p(z_t|\mathbf{x}_{<t}, \mathbf{z}_{<t}) \quad (7)$$

Here, $q(z_t|\mathbf{x}_{\leq t}, \mathbf{z}_{<t})$ and $p(z_t|\mathbf{x}_{<t}, \mathbf{z}_{<t})$ are defined using LSTM (Hochreiter and Schmidhuber 1997). When sampling, prior inputs z_{t-1} to the decoder to reconstruct x_{t-1} , which is used to generate z_t repeatedly.

As SNN is a type of RNNs, we use VRNN to build FS-VAE.

Generative Models in SNN

Spiking GAN (Kotariya and Ganguly 2021) uses two-layer SNNs to construct a generator and discriminator to train a GAN; however, the quality of the generated image is low. One reason for this is that the time-to-first spike encoding cannot grasp the entire image in the middle of spike trains. In addition, as the learning of SNN is unstable, it would be difficult to perform adversarial learning without regularization.

Applying VAE on SNN

Some studies partially used SNN to create VAEs. In (Stewart et al. 2021), human gesture videos captured with a DVS camera were input to an SNN encoder, and the latent variables were generated from the membrane potential of the output neuron; the ANN decoder reconstructed the input from that. Their research aimed to generate pseudo labels for new gestures, not generate images. Moreover, as the decoder was built with ANN, the entire model could not be implemented on neuromorphic devices. Similarly, (Skatchkovsky, Simeone, and Jang 2021) had the same problem because their decoder is also built with ANN.

In SVAE (Talaflha et al. 2020), after training a VAE built on ANN, they converted it to SNN to perform unsupervised learning using STDP, and then, converted it again to ANN to improve the quality of the generated images. Thus, image generation with SNN was not studied.

Some studies used probabilistic neurons to perform variational learning (Rezende, Wierstra, and Gerstner 2011; Bagheri 2019). However, as probabilistic neuron uses the sampling from a Poisson process, it requires additional sampling modules. Therefore, we created an entire image generation model in deterministic SNN, which is more commonly used.

Autoregressive Model for Spike Train Modeling

When constructing VAEs with SNNs, prior and posterior distributions are required to generate spike trains based on a stochastic process. In a related study, MMD-GLM (Arribas, Zhao, and Park 2020) modeled a real biological spike train as a Poisson process using an autoregressive ANN. The Maximum Mean Discrepancy (MMD) was used to measure the consistency with actual spike trains. This is because using KL divergence may cause runaway self-excitation.

In this study, we propose a method for modeling spike trains using an autoregressive SNN in prior and posterior distributions. The details are described in the next section.

Proposed Method

Overview of FSVAE

The detailed network architecture can be found in **Supplementary Material A**.

Figure 3 shows the model overview. The input image x is transformed into spike trains $\mathbf{x}_{1:T}$ using direct input encoding (Rueckauer et al. 2017), and is subsequently input to the SNN encoder. From the output spike trains of Encoder $\mathbf{x}_{1:T}^E$, the posterior outputs the latent spike trains $\mathbf{z}_{1:T}$ incrementally. Then, the SNN decoder generates output spike trains $\hat{\mathbf{x}}_{1:T}$, which is decoded to obtain the reconstructed image \hat{x} . When sampling, prior generates $\mathbf{z}_{1:T}$ incrementally and inputs it into the SNN decoder to generate image \hat{x} .

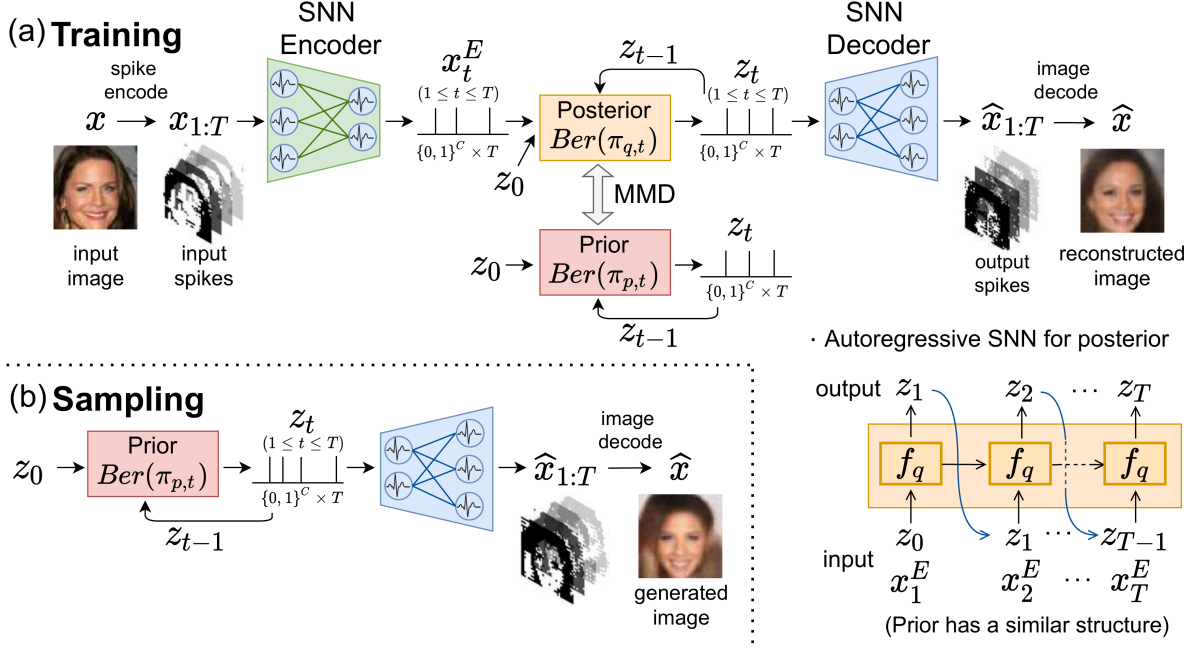


Figure 3: Overview of FSVAE. (a) During the training, the input image x is spike encoded to $x_{1:T}$, which is passed through the SNN encoder to obtain $x_{1:T}^E$. In addition to $x_{1:T}^E$, posterior takes the previously generated latent variables z_{t-1} as input, and sequentially outputs z_t . The lower right figure shows this process in detail. Here, f_q is the SNN model described in Figure 4. In prior, only z_{t-1} is used to generate z_t . Next, z_t is sequentially input to the SNN decoder, which outputs $\hat{x}_{1:T}$ and decodes it to obtain the reconstructed image \hat{x} . For the loss, we take the reconstruction error of x and \hat{x} and the MMD of the posterior and prior. (b) During the sampling, the image is generated from $z_{1:T}$ sampled in prior.

Autoregressive Bernoulli Spike Sampling

We define the posterior and prior probability distributions of posterior and prior as follows:

$$q(z_{1:T} | x_{1:T}) = \prod_{t=1}^T q(z_t | x_{\leq t}, z_{< t}) \quad (8)$$

$$p(z_{1:T}) = \prod_{t=1}^T p(z_t | z_{< t}) \quad (9)$$

We need to model $q(z_t | x_{\leq t}, z_{< t})$ and $p(z_t | z_{< t})$ with SNNs. Notably, all SNN features must be binary time series data. Therefore, we cannot use the reparameterization trick to sample from the normal distribution as in the conventional VAE.

Consequently, we define $q(z_t | x_{\leq t}, z_{< t})$ and $p(z_t | z_{< t})$ as Bernoulli distributions which take binary values. First, we need to generate z_t sequentially from $z_{< t}$; thus, we use the autoregressive SNN model. By randomly selecting one by one from its output, we can sample from a Bernoulli distribution. The overview of this sampling method is shown in Figure 4.

When the autoregressive SNN model of $q(z_t | x_{\leq t}, z_{< t})$ is f_q and that of $p(z_t | z_{< t})$ is f_p , their outputs are below.

$$\zeta_{q,t} := f_q(z_{q,t-1}, x_t^E; \Theta_{q,t}) \in \{0, 1\}^{kC} \quad (10)$$

$$\zeta_{p,t} := f_p(z_{p,t-1}; \Theta_{p,t}) \in \{0, 1\}^{kC} \quad (11)$$

where C is the dimension of z_t , k is a natural number starting from 2. x_t^E is the output of the encoder. $\Theta_{q,t}$ and $\Theta_{p,t}$ are the sets of membrane potentials of the neurons in f_q and f_p , which are updated by the input.

Sampling is performed by randomly selecting one for every k value:

$$z_{q,t,c} = \text{random_select}(\zeta_{q,t}[k(c-1) : kc]) \quad (12)$$

$$z_{p,t,c} = \text{random_select}(\zeta_{p,t}[k(c-1) : kc]) \quad (13)$$

By doing this in $1 \leq c \leq C$, we can sample $z_{q,t}, z_{p,t} \in \{0, 1\}^C$.

In TrueNorth, one of the neuromorphic devices, there is a pseudo random number generator mechanism that can replace the synaptic weight randomly (Wen et al. 2016), which makes this sampling method feasible.

This is equivalent to sampling from the following Bernoulli distribution.

$$z_{q,t} | x_t, z_{q,t-1} \sim \text{Ber}(\pi_{q,t}) \quad (14)$$

$$z_{p,t} | z_{p,t-1} \sim \text{Ber}(\pi_{p,t}) \quad (15)$$

$$\text{where } \begin{cases} \pi_{q,t,c} = \text{mean}(\zeta_{q,t}[k(c-1) : kc]) \\ \pi_{p,t,c} = \text{mean}(\zeta_{p,t}[k(c-1) : kc]) \end{cases} \quad (16)$$

Therefore,

$$q(z_t | x_{\leq t}, z_{< t}) = \text{Ber}(\pi_{q,t}) \quad (17)$$

$$p(z_t | z_{< t}) = \text{Ber}(\pi_{p,t}) \quad (18)$$

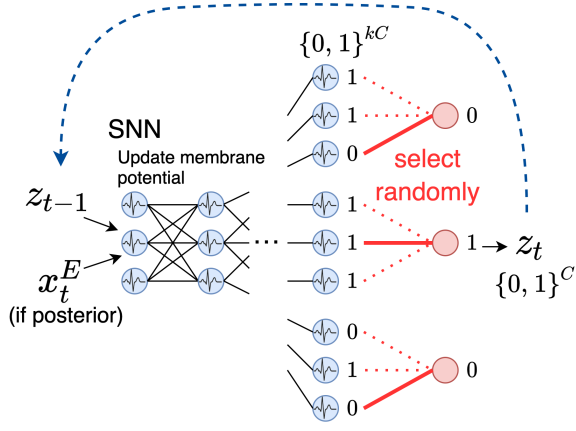


Figure 4: Autoregressive Bernoulli spike sampling for prior and posterior distributions. The input is z_{t-1} for prior, and (z_{t-1}, x_t^E) for posterior. We sample from the Bernoulli distribution by randomly selecting C channels from kC channels of the output.

Spike to Image Decoding Using Membrane Potential

Sampled z_t is sequentially input to the SNN decoder, which outputs the spike trains $\hat{x}_{1:T} \in \{0, 1\}^{C_{\text{out}} \times H \times W \times T}$. We need to decode this into the reconstructed image $\hat{x} \in \mathbb{R}^{C_{\text{out}} \times H \times W}$. To utilize the framework of SNN, we use the membrane potential of the output neuron for spike-to-image decoding. We use non-firing neurons in the output layer, and convert $\hat{x}_{1:T}$ into a real value by measuring its membrane potential at the last time T . It is given by Eq. (2) as follows:

$$\begin{aligned} u_T^{\text{out}} &= \tau_{\text{out}} u_{T-1}^{\text{out}} + \hat{x}_T \\ &= \tau_{\text{out}}^2 u_{T-2}^{\text{out}} + \tau_{\text{out}} \hat{x}_{T-1} + \hat{x}_T = \sum_{t=1}^T \tau_{\text{out}}^{T-t} \hat{x}_t \end{aligned} \quad (19)$$

We set $\tau_{\text{out}} = 0.8$, and obtain the real-valued reconstructed image as $\hat{x} = \tanh(u_T^{\text{out}})$.

Loss Function

The ELBO is as follows:

$$\begin{aligned} ELBO &= \mathbb{E}_{q(z_{1:T}|\mathbf{x}_{1:T})} [\log p(\mathbf{x}_{1:T}|z_{1:T})] \\ &\quad - \text{KL}[q(z_{1:T}|\mathbf{x}_{1:T})||p(z_{1:T})] \end{aligned} \quad (20)$$

The first term is the reconstruction loss, which is $\text{MSE}(x, \hat{x})$ as in the usual VAE. The second term, KL divergence, represents the closeness of the prior and posterior probability distributions. Traditional VAEs use KL divergence, but we use MMD, which has been shown to be a more suitable distance metric for spike trains in MMD-GLM (Arribas, Zhao, and Park 2020). MMD can be written using the kernel function k as follows:

$$\begin{aligned} &\text{MMD}^2[q(z_{1:T}|\mathbf{x}_{1:T}), p(z_{1:T})] \\ &= \mathbb{E}_{z, z' \sim q} [k(z_{1:T}, z'_{1:T})] + \mathbb{E}_{z, z' \sim p} [k(z_{1:T}, z'_{1:T})] \\ &\quad - 2 \mathbb{E}_{z \sim q, z' \sim p} [k(z_{1:T}, z'_{1:T})] \end{aligned} \quad (21)$$

We set $k(z_{1:T}, z'_{1:T}) = \sum_t \text{PSP}(z_{\leq t}) \text{PSP}(z'_{\leq t})$, as the model based kernel in MMD-GLM. The PSP stands for a postsynaptic potential function that can capture the time-series nature of spike trains (Zenke and Ganguli 2018). We use the first-order synaptic model (Zhang and Li 2020) as the PSP. The following update formula is used to calculate the $\text{PSP}(z_{\leq t})$.

$$\text{PSP}(z_{\leq t}) = \left(1 - \frac{1}{\tau_{\text{syn}}}\right) \text{PSP}(z_{\leq t-1}) + \frac{1}{\tau_{\text{syn}}} z_t \quad (22)$$

where τ_{syn} is the synaptic time constant. We set $\text{PSP}(z_{\leq 0}) = 0$.

This gives us Eq. (21), as follows. The detailed derivation can be found in **Supplementary Material B**.

$$\begin{aligned} &\text{MMD}^2[q(z_{1:T}|\mathbf{x}_{1:T}), p(z_{1:T})] \\ &= \sum_{t=1}^T \|\text{PSP}(\mathbb{E}_{z \sim q}[z_{\leq t}]) - \text{PSP}(\mathbb{E}_{z \sim p}[z_{\leq t}])\|^2 \end{aligned} \quad (23)$$

$$= \sum_{t=1}^T \|\text{PSP}(\pi_{q, \leq t}) - \text{PSP}(\pi_{p, \leq t})\|^2 \quad (24)$$

The loss function is calculated as follows:

$$\mathcal{L} = \text{MSE}(x, \hat{x}) + \sum_{t=1}^T \|\text{PSP}(\pi_{q, \leq t}) - \text{PSP}(\pi_{p, \leq t})\|^2 \quad (25)$$

Experiments

We implemented FSAE in PyTorch (Paszke et al. 2019), and evaluated it using MNIST, Fashion MNIST, CIFAR10, and CelebA. The results are summarized in Table 1.

Datasets

For MNIST and Fashion MNIST, we used 60,000 images for training and 10,000 images for evaluation. The input images were resized to 32×32 . For CIFAR10, we used 50,000 images for training and 10,000 images for evaluation. For CelebA, we used 162,770 images for training and 19,962 images for evaluation. The input images were resized to 64×64 .

Network Architecture

The SNN encoder comprises several convolutional layers, each with `kernel_size=3` and `stride=2`. The number of layers is 4 for MNIST, Fashion MNIST and CIFAR10, and 5 for CelebA. After each layer, we set a tBN (Zheng et al. 2021), and then, input the feature to the LIF neuron to obtain the output spike trains. The encoder's output is $x_t^E \in \{0, 1\}^C$ with latent dimension $C = 128$. We combine it with $z_{t-1} \in \{0, 1\}^C$, and input it to the posterior model; thus, its input dimension is $128 + 128 = 256$. Additionally, we set $z_0 = \mathbf{0}$. The posterior model comprises three FC layers, and increases the number of channels by a factor of $k = 20$. z_t is sampled from it by the proposed autoregressive Bernoulli spike sampling. We repeat this process $T = 16$ times. The

Dataset	Model	Reconstruction Loss ↘	Inception Score ↗	Fréchet Distance ↘	
				Inception (FID)	Autoencoder
MNIST	ANN	0.048	5.947	112.5	17.09
	FSVAE (Ours)	0.031	6.209	97.06	35.54
Fashion MNIST	ANN	0.050	4.252	123.7	18.08
	FSVAE (Ours)	0.031	4.551	90.12	15.75
CIFAR10	ANN	0.105	2.591	229.6	196.9
	FSVAE (Ours)	0.066	2.945	175.5	133.9
CelebA	ANN	0.059	3.231	92.53	156.9
	FSVAE (Ours)	0.051	3.697	101.6	112.9

Table 1: Results for each dataset. In all datasets, our model outperforms ANN in the inception score. Moreover, our model outperforms MNIST and Fashion MNIST in FID, CIFAR10 in all metrics, and CelebA in Autoencoder’s Fréchet distance. Reconstruction losses are better for our model in all datasets.

prior model has the same model architecture; however, as the input is only z_{t-1} , its input dimension is 128.

Sampled z_t is input to the SNN decoder, which contains the same number of deconvolution layers as the encoder. Decoder’s output is spike trains of the same size as the input. Finally, we perform spike-to-image decoding to obtain the reconstructed image, according to Eq. (19).

Training Settings

We use AdamW optimizer (Loshchilov and Hutter 2019), which trains 150 epochs with a learning rate of 0.001 and a weight decay of 0.001. The batch size is 250. In prior model, teacher forcing (Williams and Zipser 1989) is used to stabilize training, so that the prior’s input is $z_{q,t}$, which is sampled from the posterior model. In addition, to prevent posterior collapse, scheduled sampling (Bengio et al. 2015) is performed. With a certain probability, we input $z_{p,t}$ to the prior instead of $z_{q,t}$. This probability varies linearly from 0 to 0.3 during training.

Evaluation Metrics

The quality of the sampled images is measured by the inception score (Salimans et al. 2016) and FID (Heusel et al. 2017; Parmar, Zhang, and Zhu 2021). However, as FID considers the output of the ImageNet pretrained inception model, it may not work well on datasets such as MNIST, which have a different domain from ImageNet. Consequently, we trained an autoencoder on each dataset beforehand, and used it to measure the Fréchet distance of the autoencoder’s latent variables between sampled and real images. We sampled 5,000 images to measure the distance. As a comparison method, we prepared vanilla VAEs of the same network architecture built with ANN, and trained on the same settings.

Table 1 shows that our FSVAE outperforms ANN in the inception score for all datasets. For MNIST and Fashion MNIST, our model outperforms in terms of FID, for CIFAR10 in all metrics, and for CelebA in the Fréchet distance of the pretrained autoencoder. As SNNs can only use binary values, it is difficult to perform complex tasks; in contrast,

our FSVAE achieves equal or higher scores than ANNs in image generation.

Computational Cost We measured how many floating-point additions and multiplications were required for inferring a single image. We summarized the results in Table 2. The number of additions is 6.8 times less for ANN, but the number of multiplications is 14.8 times less for SNN. In general, multiplication is more expensive than addition, so fewer multiplications are preferable. Moreover, as SNNs are expected to be about 100 times faster than ANNs when implemented in neuromorphic device, FSVAE can significantly outperform ANN VAE in terms of speed.

Model	Computational complexity	
	Addition	Multiplication
ANN	7.4×10^9	7.4×10^9
FSVAE (Ours)	5.0×10^{10}	5.6×10^8

Table 2: Comparison of the amount of computation required to infer a single image in MNIST.

Figure 6 shows the change in FID depending on the timestep and k , the number of output choices. The best FID was obtained when the timestep was 16. A small number of timesteps does not have enough expressive power, whereas a large number of timesteps makes the latent space too large; thus, the best timestep was determined to be 16. Moreover, FID does not change so much with k , but it is best when $k = 20$.

Ablation Study The results are summarized in Table 3. When calculating KLD, $\epsilon = 0.01$ was added to $\pi_{q,t}$ and $\pi_{p,t}$ to avoid divergence. The best FID score was obtained using MMD as the loss function and applying PSP to its kernel.

Qualitative Evaluation

Figure 5 shows examples of the generated images. In CIFAR10 and Fashion MNIST, the ANN VAE generated blurry

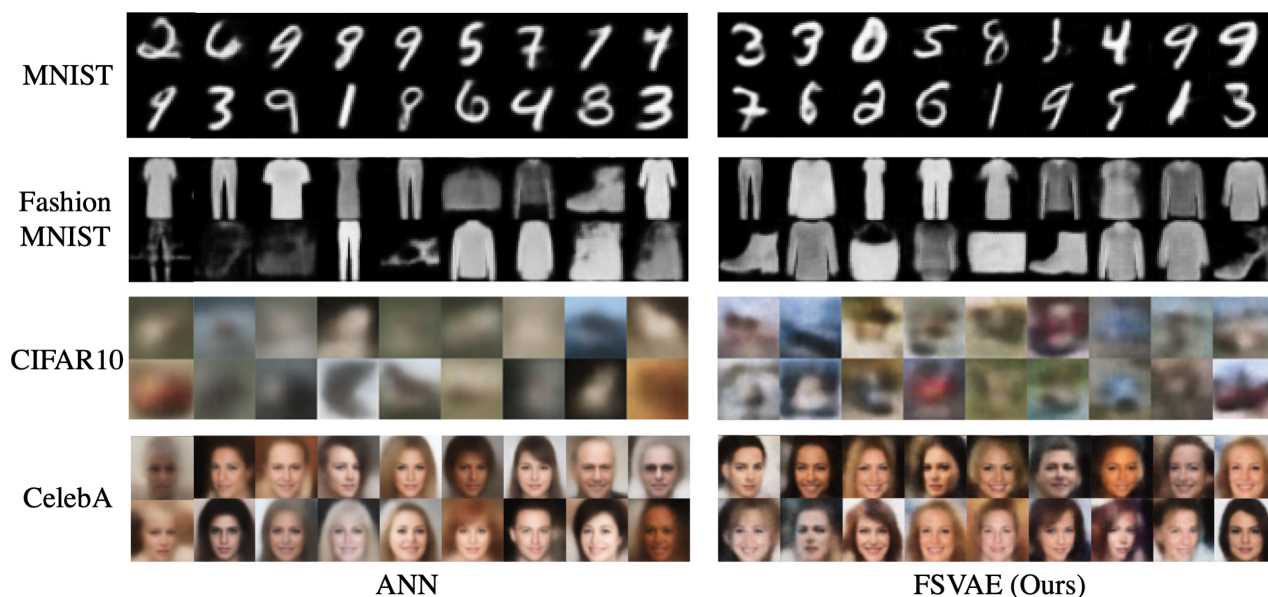


Figure 5: Generated images of ANN VAE and our FSVAE (SNN).

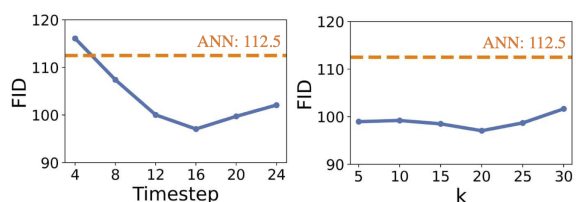


Figure 6: Left: relationship between the timestep and FID score of generated images in MNIST. Right: relationship between k (multiplier for the channel) and FID in MNIST.

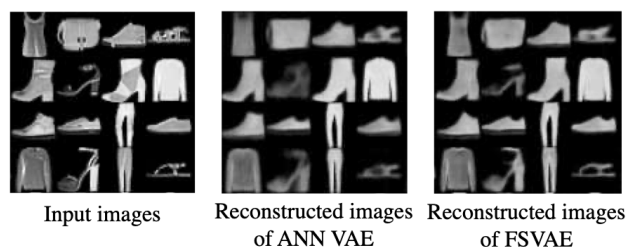


Figure 7: Reconstructed images of Fashion MNIST.

KLD	MMD	apply PSP	FID \searrow
✓			114.3
	✓		106.0
	✓	✓	101.6

Table 3: Ablation study on CelebA. KLD and MMD indicate the distance metric used in the loss function. We also measured whether to apply PSP to the MMD kernel.

and hazy images, whereas our FSVAE generated clearer images. In CelebA, FSVAE generated images with more distinct background areas than ANN. This is because the latent variables of FSVAE are discrete spike trains, thus it can avoid posterior collapse, as VQ-VAE (van den Oord, Vinyals, and Kavukcuoglu 2017).

Figure 7 shows the reconstructed images of Fashion MNIST. FSVAE reconstruct images more clearly than ANN, especially the details. This is because that the latent variable of FSVAE is discrete. Posterior collapse is caused by the latent variables being ignored by the decoder. With discrete

latent variables, every latent variable becomes meaningful, which can prevent posterior collapse.

Conclusion

In this study, we proposed FSVAE, which allows image generation on SNNs with equal or higher quality than ANNs. We proposed autoregressive Bernoulli spike sampling, which is a spike sampling strategy that can be implemented on neuromorphic devices. This sampling method is used in prior and posterior distribution, and we model the latent spike trains as a Bernoulli process. Experiments on multiple datasets show that FSVAE can generate images with equal or higher quality than ANN VAE of the same architecture. As SNNs can be significantly faster on neuromorphic devices, FSVAE can significantly outperform ANN VAE in terms of speed. In the future, when we incorporate the recent VAE researches, we will soon be able to achieve high-resolution image generation with SNNs.

Acknowledgments

This work was partially supported by JST AIP Acceleration Research JPMJCR20U3, Moonshot R&D Grant Number JPMJPS2011, CREST Grant Number JPMJCR2015, and Basic Research Grant (Super AI) of Institute for AI and Beyond of the University of Tokyo. We would like to thank Yang Li and Muyuan Xu for some helpful advice of SNNs. We also thank Tomoyuki Takahata, Yusuke Mori, and Atsuhiko Noguchi for helpful discussions.

References

- Akopyan, F.; Sawada, J.; Cassidy, A.; Alvarez-Icaza, R.; Arthur, J. V.; Merolla, P.; Imam, N.; Nakamura, Y. Y.; Datta, P.; Nam, G.; Taba, B.; Beakes, M. P.; Brezzo, B.; Kuang, J. B.; Manohar, R.; Risk, W. P.; Jackson, B. L.; and Modha, D. S. 2015. TrueNorth: Design and Tool Flow of a 65 mW 1 Million Neuron Programmable Neurosynaptic Chip. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 34(10): 1537–1557.
- An, J.; and Cho, S. 2015. Variational Autoencoder based Anomaly Detection using Reconstruction Probability. *Special Lecture on IE*, 2(1): 1–18.
- Arribas, D. M.; Zhao, Y.; and Park, I. M. 2020. Rescuing neural spike train models from bad MLE. In *NeurIPS*.
- Bagheri, A. 2019. *Probabilistic Spiking Neural Networks : Supervised, Unsupervised and Adversarial Trainings*. Ph.D. thesis, New Jersey Institute of Technology.
- Bengio, S.; Vinyals, O.; Jaitly, N.; and Shazeer, N. 2015. Scheduled Sampling for Sequence Prediction with Recurrent Neural Networks. In *NeurIPS*.
- Benjamin, B. V.; Gao, P.; McQuinn, E.; Choudhary, S.; Chandrasekaran, A.; Bussat, J.; Alvarez-Icaza, R.; Arthur, J. V.; Merolla, P.; and Boahen, K. 2014. Neurogrid: A Mixed-Analog-Digital Multichip System for Large-Scale Neural Simulations. *Proceedings of the IEEE*, 102(5): 699–716.
- Cassidy, A. S.; Alvarez-Icaza, R.; Akopyan, F.; Sawada, J.; Arthur, J. V.; Merolla, P.; Datta, P.; Tallada, M. G.; Taba, B.; Andreopoulos, A.; Amir, A.; Esser, S. K.; Kusnitz, J.; Appuswamy, R.; Haymes, C.; Brezzo, B.; Moussalli, R.; Bellofatto, R.; Baks, C. W.; Mastro, M.; Schleupen, K.; Cox, C. E.; Inoue, K.; Millman, S. E.; Imam, N.; McQuinn, E.; Nakamura, Y. Y.; Vo, I.; Guok, C.; Nguyen, D.; Lekuch, S.; Asaad, S. W.; Friedman, D. J.; Jackson, B. L.; Flickner, M.; Risk, W. P.; Manohar, R.; and Modha, D. S. 2014. Real-Time Scalable Cortical Computing at 46 Giga-Synaptic OP-S/Watt with $\sim 100\times$ Speedup in Time-to-Solution and $\sim 100,000\times$ Reduction in Energy-to-Solution. In *Supercomputing Conference*.
- Chung, J.; Kastner, K.; Dinh, L.; Goel, K.; Courville, A. C.; and Bengio, Y. 2015. A Recurrent Latent Variable Model for Sequential Data. In *NeurIPS*.
- Davies, M.; Srinivasa, N.; Lin, T.; China, G. N.; Cao, Y.; Choday, S. H.; Dimou, G. D.; Joshi, P.; Imam, N.; Jain, S.; Liao, Y.; Lin, C.; Lines, A.; Liu, R.; Mathaikutty, D.; McCoy, S.; Paul, A.; Tse, J.; Venkataramanan, G.; Weng, Y.; Wild, A.; Yang, Y.; and Wang, H. 2018. Loihi: A Neuro-morphic Manycore Processor with On-Chip Learning. *IEEE Micro*, 38(1): 82–99.
- Deng, J.; Dong, W.; Socher, R.; Li, L.; Li, K.; and Li, F. 2009. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR*.
- Deng, L. 2012. The MNIST Database of Handwritten Digit Images for Machine Learning Research. *IEEE Signal Processing Magazine*, 29(6): 141–142.
- Diehl, P. U.; and Cook, M. 2015. Unsupervised Learning of Digit Recognition using Spike-Timing-Dependent Plasticity. *Frontiers in Computational Neuroscience*, 9: 99.
- Han, K.; Wen, H.; Shi, J.; Lu, K.-H.; Zhang, Y.; and Liu, Z. 2018. Variational Autoencoder: An Unsupervised Model for Modeling and Decoding fMRI Activity in Visual Cortex. *bioRxiv*.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium. In *NeurIPS*.
- Hochreiter, S.; and Schmidhuber, J. 1997. Long Short-Term Memory. *Neural Computation*, 9(8): 1735–1780.
- Karras, T.; Laine, S.; Aittala, M.; Hellsten, J.; Lehtinen, J.; and Aila, T. 2020. Analyzing and Improving the Image Quality of StyleGAN. In *CVPR*.
- Kim, S. J.; Park, S.; Na, B.; and Yoon, S. 2020. Spiking-YOLO: Spiking Neural Network for Energy-Efficient Object Detection. In *AAAI*.
- Kingma, D. P.; and Welling, M. 2014. Auto-Encoding Variational Bayes. In *ICLR*.
- Kotariya, V.; and Ganguly, U. 2021. Spiking-GAN: A Spiking Generative Adversarial Network Using Time-To-First-Spike Coding. arXiv:2106.15420.
- Krizhevsky, A.; and Hinton, G. 2009. Learning Multiple Layers of Features from Tiny Images. Technical report, University of Toronto.
- Lee, C.; Kosta, A.; Zhu, A. Z.; Chaney, K.; Daniilidis, K.; and Roy, K. 2020. Spike-FlowNet: Event-Based Optical Flow Estimation with Energy-Efficient Hybrid Neural Networks. In *ECCV*.
- Liu, Z.; Luo, P.; Wang, X.; and Tang, X. 2015. Deep Learning Face Attributes in the Wild. In *ICCV*.
- Loshchilov, I.; and Hutter, F. 2019. Decoupled Weight Decay Regularization. In *ICLR*.
- Maass, W. 1997. Networks of spiking neurons: The third generation of neural network models. *Neural Networks*, 10(9): 1659–1671.
- Parmar, G.; Zhang, R.; and Zhu, J. 2021. On Buggy Resizing Libraries and Surprising Subtleties in FID Calculation. arXiv:2104.11222.
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; Desmaison, A.; Köpf, A.; Yang, E.; DeVito, Z.; Raison, M.; Tejani, A.; Chilamkurthy, S.; Steiner, B.; Fang, L.; Bai, J.; and Chintala, S. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *NeurIPS*.

- Razavi, A.; van den Oord, A.; and Vinyals, O. 2019. Generating Diverse High-Fidelity Images with VQ-VAE-2. In *NeurIPS*.
- Rezende, D. J.; Wierstra, D.; and Gerstner, W. 2011. Variational Learning for Recurrent Spiking Networks. In *NeurIPS*.
- Rueckauer, B.; Lungu, I.-A.; Hu, Y.; Pfeiffer, M.; and Liu, S.-C. 2017. Conversion of Continuous-Valued Deep Networks to Efficient Event-Driven Networks for Image Classification. *Frontiers in Neuroscience*, 11: 682.
- Salimans, T.; Goodfellow, I. J.; Zaremba, W.; Cheung, V.; Radford, A.; and Chen, X. 2016. Improved Techniques for Training GANs. In *NeurIPS*.
- Skatchkovsky, N.; Simeone, O.; and Jang, H. 2021. Learning to Time-Decode in Spiking Neural Networks Through the Information Bottleneck. In *NeurIPS*.
- Stein, R. B.; and Hodgkin, A. L. 1967. The Frequency of Nerve Action Potentials Generated by Applied Currents. *Proceedings of the Royal Society of London. Series B. Biological Sciences*, 167(1006): 64–86.
- Stewart, K.; Danieleescu, A.; Supic, L.; Shea, T. M.; and Neftci, E. 2021. Gesture Similarity Analysis on Event Data Using a Hybrid Guided Variational Auto Encoder. arXiv:2104.00165.
- Talafha, S.; Rekabdar, B.; Mousas, C.; and Ekenna, C. 2020. Biologically Inspired Sleep Algorithm for Variational Auto-Encoders. In *ISVC*.
- Vahdat, A.; and Kautz, J. 2020. NVAE: A Deep Hierarchical Variational Autoencoder. In *NeurIPS*.
- van den Oord, A.; Vinyals, O.; and Kavukcuoglu, K. 2017. Neural Discrete Representation Learning. In *NeurIPS*.
- Wen, W.; Wu, C.; Wang, Y.; Nixon, K. W.; Wu, Q.; Barnell, M.; Li, H.; and Chen, Y. 2016. A New Learning Method for Inference Accuracy, Core Occupation, and Performance Co-optimization on TrueNorth Chip. In *DAC*.
- Williams, R. J.; and Zipser, D. 1989. A Learning Algorithm for Continually Running Fully Recurrent Neural Networks. *Neural Computation*, 1(2): 270–280.
- Wu, J.; Chua, Y.; Zhang, M.; Li, H.; and Tan, K. C. 2018. A Spiking Neural Network Framework for Robust Sound Classification. *Frontiers in Neuroscience*, 12: 836.
- Wu, Y.; Deng, L.; Li, G.; Zhu, J.; Xie, Y.; and Shi, L. 2019. Direct Training for Spiking Neural Networks: Faster, Larger, Better. In *AAAI*.
- Xiao, H.; Rasul, K.; and Vollgraf, R. 2017. Fashion-MNIST: a Novel Image Dataset for Benchmarking Machine Learning Algorithms. <https://github.com/zalandoresearch/fashion-mnist>. Accessed: 2021-09-08, arXiv:1708.07747.
- Zenke, F.; and Ganguli, S. 2018. SuperSpike: Supervised Learning in Multilayer Spiking Neural Networks. *Neural Computation*, 30(6).
- Zhang, W.; and Li, P. 2020. Temporal Spike Sequence Learning via Backpropagation for Deep Spiking Neural Networks. In *NeurIPS*.
- Zheng, H.; Wu, Y.; Deng, L.; Hu, Y.; and Li, G. 2021. Going Deeper With Directly-Trained Larger Spiking Neural Networks. In *AAAI*.