

Spiking NeRF: Representing the Real-World Geometry by a Discontinuous Representation

Zhanfeng Liao, Yan Liu, Qian Zheng*, Gang Pan*

Zhejiang University
{zhanfengliao, qianzheng, rliuyan2, gpan}@zju.edu.cn

Abstract

A crucial reason for the success of existing NeRF-based methods is to build a neural density field for the geometry representation via multiple perceptron layers (MLPs). MLPs are continuous functions, however, real geometry or density field is frequently discontinuous at the interface between the air and the surface. Such a contrary brings the problem of unfaithful geometry representation. To this end, this paper proposes spiking NeRF, which leverages spiking neurons and a hybrid Artificial Neural Network (ANN)-Spiking Neural Network (SNN) framework to build a discontinuous density field for faithful geometry representation. Specifically, we first demonstrate the reason why continuous density fields will bring inaccuracy. Then, we propose to use the spiking neurons to build a discontinuous density field. We conduct a comprehensive analysis for the problem of existing spiking neuron models and then provide the numerical relationship between the parameter of the spiking neuron and the theoretical accuracy of geometry. Based on this, we propose a bounded spiking neuron to build the discontinuous density field. Our method achieves SOTA performance. The source code and the supplementary material are available at <https://github.com/liaozhanfeng/Spiking-NeRF>.

Introduction

3D reconstruction from RGB images is a challenging and complex task in computer vision (Mildenhall et al. 2021; Wang et al. 2021). Neural Radiance Fields (NeRF) (Mildenhall et al. 2021), a recently promising solution for novel view synthesis in an implicit manner, has also achieved very competitive results in 3D reconstruction (Wang et al. 2021; Wang, Skorokhodov, and Wonka 2022, 2023).

One of the reasons for the success of NeRF is its ability to implicitly represent geometric information using a neural network based on conventional MLPs, which are continuous functions. However, geometric information in the real world is discontinuous at the interface between the air and the surface, which is inconsistent with the computational representation in conventional MLPs. This inconsistency poses three problems. **Post-processing.** Existing methods require a post-processing approach (i.e., filtering

with a threshold (Peng et al. 2021; Pumarola et al. 2021; Boss et al. 2021)) to extract discontinuous geometric representations from learned continuous fields. However, determining the optimal threshold requires empirical and tedious tuning, and many methods rely on manual selection of the threshold for different scenarios (Mildenhall et al. 2021; Pumarola et al. 2021; Boss et al. 2021) (see the left of Fig. 1). **Optimal threshold perturbation.** NeRF-based methods optimize the network view by view, resulting in the optimal threshold of the surface to perturbate (see supplementary material for more details). However, existing methods use the same threshold to filter output values (Mildenhall et al. 2021; Peng et al. 2021; Pumarola et al. 2021; Boss et al. 2021), which cannot obtain accurate geometric information because the optimal filtering thresholds under different views vary (see the middle of Fig. 1). **Light density scenarios**¹. The inconsistency can result in even greater errors in light density scenarios (see the right of Fig. 1). Because it is difficult for continuous functions like conventional MLPs to generate quite different densities (i.e., significantly different outputs) for points close to each other (i.e., similar inputs), resulting in the disappearance of the object.

Although there are some attempts to address the aforementioned problems, these methods are still based on continuous representation and do not provide a fundamental solution to those problems. For example, some works replace the density field with other fields to avoid threshold selection (e.g., (Oechsle, Peng, and Geiger 2021; Wang et al. 2021; Yariv et al. 2021)). However, these methods cannot use the filter value to get an accurate geometric representation (see the left of Fig. 1). Moreover, they fail on light density scenarios (see the right of Fig. 1). Some works modify the input to improve its frequency or discretization, thereby mitigating the inconsistency between continuous and discontinuous representations. However, these methods still fail on light density scenarios (Wang, Skorokhodov, and Wonka 2022, 2023) (see the right of Fig. 1). Some works are designed for specific scenarios (e.g., (Guo et al. 2022; Levy et al. 2023)). However, these methods have poorer performance in general

¹These scenarios contain several rays, along which the density with non-zero values distribute within a very narrow range (e.g., thin objects) or with small values (e.g., semi-transparent objects). In real life, semi-transparent objects (e.g., windows, glasses) and thin objects (e.g., branches, axles, nets) can be seen everywhere.

*Corresponding authors.
Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: Left: The extracted surfaces from NeRF. Each row in the first big red box represents a surface extracted by a trained NeRF using different thresholds, indicating that the optimal thresholds corresponding to different scenarios are different. The tick represents that the threshold is optimal. The cross represents that the threshold is not optimal. Middle: The error maps from different views in the same scene. Each row in the second big red box represents the depth error map of a trained NeRF’s surface extracted with different thresholds from different views. It can be seen that the optimal thresholds corresponding to different views are different. Right: The extracted surfaces from NeRF. These figures show that the inconsistency can result in even greater errors in light density scenarios. The image in the bottom right corner of each part represents the original image from the corresponding view. The number displayed in the bottom left corner of each image represents either the Chamfer distance (left and right) or the depth error (middle).

scenarios. Moreover, they mainly focus on novel view synthesis, and the quality of the 3D reconstruction is low (see the right of Fig. 1).

Compared to the ANN transmitting continuous values, the SNN transmits discontinuous spikes. The discontinuity of the SNN is expected to represent the discontinuous geometric information in the real world and fundamentally resolve the inconsistency between real-world representation and computational representation in conventional MLPs. Moreover, the threshold of spiking neurons in SNNs can serve as a filtering threshold through a parameter learning scheme (e.g., (Li et al. 2022; Zhang et al. 2022)), which avoids the issue of manually specifying the threshold. Different from traditional cognitive applications (Li et al. 2022; Zhang et al. 2022; Han et al. 2023; Zou, Huang, and Wu 2022), using the spiking representation to model 3D geometry from a numerical perspective is a non-cognitive application (Aimone et al. 2022; Ren et al. 2023).

In this paper, we address the problems of continuous geometric representation in the conventional NeRF by proposing spiking NeRF, which is based on a hybrid ANN-SNN framework to model the real-world 3D geometry in a discontinuous representation. First, we build the relationship of the spiking threshold, maximum activation, and depth error for our discontinuous computational representation. Second, based on this relationship, we observe that when the spiking threshold is sufficiently large, the depth error is sufficiently small. However, the spiking threshold cannot be set sufficiently high in semi-transparent scenarios and cannot be set to infinity in practical implementation. We further discover a way to maintain a small error under a finite spiking threshold, which is to control the maximum activation. Last, based on our analysis, we propose the Bounded Full-precision Integrate and Fire (B-FIF) spiking neuron to build a hybrid ANN-SNN framework. Meanwhile, we design a corresponding training pipeline for the hybrid ANN-SNN framework and verify the effectiveness on mainstream datasets and light density scenarios. Our contributions can be summarized as follows:

- We build the relationship of the spiking threshold, maximum activation, and depth error. Moreover, we constrain the bound of depth error by the spiking threshold and maximum activation. This bound can be leveraged to facilitate the real-world application for geometric estimation without knowing the ground truth.
- We propose the B-FIF spiking neuron based on the aforementioned analysis and build a hybrid ANN-SNN framework. Meanwhile, we design the corresponding training pipeline and training strategy for the hybrid ANN-SNN framework. We verify the effectiveness on mainstream datasets and light density scenarios.

Related Work

Neural Implicit Representations

Recently, neural implicit functions have emerged as an effective representation of 3D geometry (Atzmon et al. 2019; Peng et al. 2020; Saito et al. 2019; Xu et al. 2019) and appearance (Liu et al. 2020a,b; Schwarz et al. 2020; Sitzmann, Zollhöfer, and Wetzstein 2019) as they represent 3D content continuously and without discretization while simultaneously having a small memory footprint. Most of these methods require 3D supervision. However, several recent works (Niemeyer et al. 2020; Yariv et al. 2020) demonstrated differentiable rendering for training directly from images. Some works use estimated depth information for surface rendering without pixel-accurate object masks (Verbin et al. 2022). Some works enhance the accuracy of geometric information by incorporating point cloud information and warp operations (Fu et al. 2022; Darmon et al. 2022). Some works do not model the density field and use other fields to avoid threshold selection (Oechsle, Peng, and Geiger 2021; Wang et al. 2021; Fu et al. 2022; Darmon et al. 2022; Yariv et al. 2021). However, these works are still based on a continuous representation, the same threshold cannot accurately divide the surface. Different from previous methods, we model the 3D geometric information in a discontinuous representation by proposing a hybrid ANN-SNN framework.

Spiking Neural Networks in Computer Vision

Over the past few years, brain-inspired (Li et al. 2021a) deep SNNs using spike trains (Marchisio et al. 2020) have gradually replaced ANNs in various tasks (Zhang, Huang, and He 2023; Zhang et al. 2023; Ororbia 2023), and it is extensively utilized to develop energy-efficient solutions for various tasks. Some works use SNNs to process video or event streams (Zhu et al. 2022; Gehrig et al. 2020). Some works exploit the discontinuity of the network to enhance its robustness (Bagheri, Simeone, and Rajendran 2018; Sharmin et al. 2020). Some works leverage the neural parameter of SNNs to further improve the utilization of available information (Li et al. 2022; Zhang et al. 2022). Although previous works have demonstrated SNN applications on a wide range of tasks, they are still limited in their performance (Li et al. 2021b; Fang et al. 2021). Meanwhile, there has been a growing interest in exploring the potential benefits of combining ANNs and SNNs (Kugele et al. 2021; Liu and Zhao 2022). By combining ANNs and SNNs, better performance can be achieved while reducing the time step of SNNs. Different combination strategies have been explored for a variety of tasks. A group of work employs the strategy of processing the accumulated spike train of SNNs with ANNs (Kugele et al. 2021; Lee et al. 2020). In these works, the SNN is used as an efficient encoder of spatio-temporal data. The output of the SNN is accumulated to summarize the temporal dimension before the ANN processes the accumulated features (Kugele et al. 2021; Lee et al. 2020). A second line of work uses a strategy where the output of the independently operating SNN and ANN is fused (Lele et al. 2022; Zhao et al. 2022). In these works, the outputs of both networks are fused based on heuristics (Lele et al. 2022) or accumulation based on the output of the ANN (Zhao et al. 2022). Different from the previous methods, we apply the hybrid ANN-SNN framework to a non-cognitive application.

Preliminary

Neural Radiance Fields

NeRF (Mildenhall et al. 2021) represents a scene as a continuous volumetric field, where the density $\sigma \in R$ and radiance $c \in R^3$ at any 3D position $\mathbf{x} \in R^3$ under viewing direction $\mathbf{d} \in R^2$ are modeled by a MLP $f_\theta : (\mathbf{x}, \mathbf{d}) \rightarrow (c, \sigma)$, with θ as learnable parameters. To render a pixel, the MLP first evaluates points sampled from the camera ray $r = \mathbf{o} + t\mathbf{d}$ to get their densities and radiance. The light starts at 0 and stops at a hypothetical large density T . Define C_p as the estimated pixel color and \hat{C}_p as the ground truth. NeRF is optimized by minimizing the photometric loss:

$$L_{\text{rgb}} = \sum_p \left\| C_p - \hat{C}_p \right\|. \quad (1)$$

Here, p refers to each pixel. There are currently two main methods for depth estimation: one based on the integration (Zhang et al. 2021; Deng et al. 2022) and the other based on the threshold to get global shape (Wang et al. 2021; Ichnowski et al. 2021). Since we focus on a global geometric shape extracted from the network by threshold rather than a

depth map dependent on perspective, this paper mainly follows the second method to define the depth.

Spiking Neuron

Integrate-and-fire model. We introduce the well-known Integrate-and-Fire (IF) model (Dayan and Abbott 2005). Given a membrane potential u_t at time step t , the membrane potential u_{t+1}^{pre} before firing at time step $t + 1$ is updated as:

$$u_{t+1}^{\text{pre}} = u_t + Wx_{t+1}. \quad (2)$$

Here, W and x_{t+1} represent respectively the weight and the output from the previous layer at time step $t + 1$. V_{th} is the firing threshold. The spiking neuron will fire a spike when u_{t+1}^{pre} exceeds V_{th} (see Fig. 2), and then u_{t+1} resets to 0. The spike output is given by:

$$o_{t+1} = \begin{cases} 0 & u_{t+1}^{\text{pre}} < V_{\text{th}} \\ 1 & \text{otherwise,} \end{cases} \quad (3)$$

$$y_{t+1} = o_{t+1} \cdot V_{\text{th}}. \quad (4)$$

After firing, the spike output y_{t+1} will propagate to the next layer and become the input x_{t+1} of the next layer. Note that we omit the layer index for simplicity.

Full-precision integrate-and-fire model. However, the spikes convey less information than the floating-point number, and networks using IF do not perform well in regression tasks (Rançon et al. 2022; Lu et al. 2021; Shrestha and Orchard 2018). It's difficult to ensure accuracy by directly using IF for modeling. To obtain full-precision information, (Li et al. 2022) changes the Eqn. 4 to:

$$y_{t+1} = o_{t+1} \cdot u_{t+1}^{\text{pre}}. \quad (5)$$

By replacing the output V_{th} with u_{t+1}^{pre} , the full-precision information is reserved. Due to the full-precision information meeting our requirements, we will only consider one time step and omit t in the following formulations.

Spiking NeRF

To make the density field discontinuous, a trivial way is to introduce spiking neurons in NeRF, which replaces the last activation layer of the density network with spiking neurons (e.g., IF, FIF). However, on the one hand, as previously mentioned, using IF is difficult to ensure accuracy, and on the other hand, directly replacing the activation layer with FIF will result in a significant depth error when the spiking threshold is small (see ablation study for more details).

In this section, we first analyze what kind of spiking neurons can solve problems of post-processing and optimal threshold perturbation, and alleviate the problem of light density scenarios by building the relationship of the spiking threshold, maximum activation, and depth error. Then, based on this relationship, we further propose our method.

Relationship between Parameters of Spiking Neuron and Depth Error

We define that d_v is the estimated depth and d is the accurate depth. T refers to the length of the sampling range and Δt

refers to the sampling interval. V_{th} is the spiking threshold and V_{max} is the maximum density value. For a well-trained NeRF, the V_{th} is also the non-zero minimum value of its density field, and the density value of the first point that a ray encounters with a non-zero density is V_{th} . We build the relationship based on the following Proposition 1 (see supplementary material for proof).

Proposition 1. For a well-trained NeRF, we have:

$$|d - d_v| < \max((\Delta t - T e^{-V_{max}\Delta t})e^{-V_{th}\Delta t}, T(1 - e^{-V_{max}T})e^{-V_{th}\Delta t}). \quad (1)$$

Based on the relationship, when V_{th} is sufficiently large $e^{-V_{th}\Delta t}$ is sufficiently small, resulting in a sufficiently small error. Meanwhile, for a fixed V_{th} , a small V_{max} can decrease the error (see supplementary material for more details).

However, V_{th} cannot be set sufficiently high in semi-transparent scenarios and cannot be set to infinity in practical implementation. Meanwhile, FIF does not have V_{max} constraint (see the right of Fig. 2), resulting in a large positive $(\Delta t - T e^{-V_{max}\Delta t})$. In this case, significant errors may occur in semi-transparent scenarios and practical implementation. Direct use of FIF cannot completely solve the above-mentioned problems. Therefore, V_{max} should have a constraint to decrease the error and we need a spiking neuron with a relatively small maximum activation.

B-FIF: Bounded full-precision integrate-and-fire spiking neuron. Based on previous analysis, the maximum activation close to the spiking threshold can ensure lower error. Therefore, we need to constrain the maximum activation of spiking neurons to decrease the error. In this paper, we use $\tanh()$ to constrain the maximum activation (see supplementary material for other bounded functions). The u^{pre} in Eq. 5 is reformulated as:

$$u^{pre} = k \cdot \tanh(u + Wx), \quad (7)$$

$$\sigma = y = o \cdot u^{pre}. \quad (8)$$

Here, k is a learnable parameter. By increasing k , it is ensured that the spiking neuron can have a larger spiking threshold. The proposed neuron constrains the maximum output value while ensuring accuracy as it is also a full-precision neuron. So the spiking neuron can ensure that the network is trained accurately and also decrease errors when the spiking threshold is relatively small as its maximum activation approaches its spiking threshold.

Hybrid ANN-SNN Framework

B-FIF implementation. Directly using the previously mentioned spiking neuron can lead to slow training. It is potentially caused by the derivative of the initial network output approaching 0 (see Fig. 2), leading to a smaller learning step size. Therefore, to better train the network, we introduce a learnable parameter r for proposed neurons to increase the derivative of the initial network output (see Fig. 2), which can improve training. The u^{pre} in Eq. 7 is reformulated as:

$$u^{pre} = k \cdot r \cdot \tanh\left(\frac{u + Wx}{r}\right). \quad (9)$$

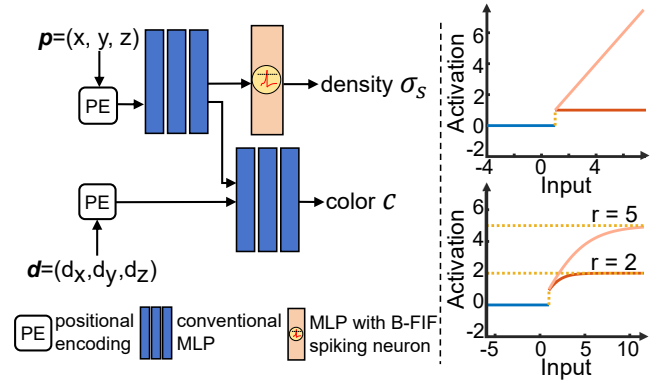


Figure 2: Framework overview of spiking NeRF and an illustration of different existing spiking neuron models and the proposed one. Left: The network structure of our approach. We use a NeRF model following (Mildenhall et al. 2021) but excluding the last activation layer of the density network. Instead of using ReLU, we use B-FIF spiking neurons to make the density field discontinuous. Right top: the IF and FIF. Right bottom: B-FIF with different r ($r = 2$ and 5). These curves show that B-FIF becomes more similar to FIF as the parameter r increases. And when the r is sufficiently large, B-FIF degenerates to the FIF.

We set the initial value of r and k to 100 and 1 to maintain similarity to the original NeRF output (i.e., ReLU (Mildenhall et al. 2021; Azinović et al. 2022; Wang et al. 2021)) near point 0 (see Fig. 2). This setting ensures that the network maintains performance similar to NeRF during the initial training stage, which can ensure that the range of the derivative approaching 1 is expanded. In addition to the above advantage, this initial strategy and the proposed neuron in Eq. 9 have another advantage the network is trained with an increased penalty for small density regions that should not appear, and it can limit the density value that should not appear above the spiking threshold more easily during training.

When r is sufficiently large, the B-FIF degenerates into the FIF, and when the maximum activation approaches the threshold, the B-FIF degenerates into the IF (please find Fig. 2 for the comparison of different spiking neuron models).

Loss function. Based on Proposition 1, it is necessary to have a large spiking threshold to obtain accurate geometric information. The network can learn the spiking threshold during the training process. However, due to the difficulty of the network spontaneously pushing the spiking threshold to a larger value, the spiking threshold at the end of the training process may not be relatively large. Therefore, we propose a regularization term L_v to increase the spiking threshold and ensure that the spiking threshold does not remain a small value after the training is completed. The L_v is as follows:

$$L_v = \frac{1}{V_{th}}. \quad (10)$$

Following (Verbin et al. 2022), we use a regularization term L_g to maintain the smoothness of the geometric representation in the initial stage, improving the network conver-

gence. The L_g is as follows:

$$L_g = \sum_p \sum_i w_i \max(-\mathbf{d}_p \nabla \sigma_i, 0). \quad (11)$$

Here, p refers to each pixel as defined in (Guo et al. 2022). i refers to each sampling point and w_i refers to the weight at each sampling point as defined in (Verbin et al. 2022). However, in (Verbin et al. 2022), they specifically designed a MLP to predict the normal vector, which generated more parameters that needed to be learned. We directly use the normal vector calculated by the gradient of density on the input coordinates to calculate L_g , reducing the number of parameters while reducing network complexity.

Finally, We optimize the following loss function:

$$L = L_{\text{rgb}} + \lambda_1 L_v + \lambda_2 L_g. \quad (12)$$

Surrogate gradient. The non-differentiability of the firing function remains one of the most significant challenges in training SNNs (Neftci, Mostafa, and Zenke 2019). Direct training requires the use of the surrogate gradient (Pfeiffer and Pfeil 2018; Li et al. 2022, 2021c; Bagheri, Simeone, and Rajendran 2018). There are 2 kinds of surrogate gradient, the surrogate gradient for conventional spiking neurons and the surrogate gradient for full-precision spiking neurons. The proposed neuron is based on the full-precision spiking neuron, so we use the surrogate gradient for full-precision spiking neurons. However, there are multiple time steps in (Li et al. 2022), and the surrogate gradient of the spiking threshold is based on the number of spikes, while our method only uses one time step. Directly using the surrogate gradient in (Li et al. 2022) can lead to unstable training (see supplement material for more details). In this paper, we use the piecewise linear function similar to (Li et al. 2021c) as the surrogate gradient for V_{th} to avoid the gradient that is 0 almost everywhere. The surrogate gradient of B-FIF is as follows:

$$\frac{\partial y}{\partial u^{\text{pre}}} = o, \quad (13)$$

$$\frac{\partial y}{\partial V_{\text{th}}} = \lambda \max(0, \frac{k - |u^{\text{pre}} - V_{\text{th}}|}{k^2}). \quad (14)$$

Meanwhile, due to the use of full-precision spiking neurons, we only need one time step to meet requirements. So theoretically, our training complexity is similar to that of ANN.

Training strategy. During the initial stage of network training, the density field exhibits significant deviations. A larger spiking threshold will render the network untrainable because most of the membrane potential cannot reach the spiking threshold and is not in the range of surrogate gradient non-zero. So we initially set the spiking threshold to 0. Furthermore, we set λ_1 to 0.15 at first and continue to increase λ_1 as the training process progresses. To preserve true high-frequency geometric information, we initially set the λ_2 to 0.0001 and gradually decrease the proportion of L_g in L as the training progresses by decreasing λ_2 .

Experiments

Experimental Settings

Datasets. We evaluate our methods in 8 scenes from the Blender dataset (Mildenhall et al. 2021) and 6 scenes from

Dex-NeRF (Ichnowski et al. 2021). Following most 3D reconstruction methods (Wang et al. 2021; Wang, Skorokhodov, and Wonka 2022, 2023; Oechsle, Peng, and Geiger 2021), we also evaluate our methods in 15 scenes from the DTU dataset (Jensen et al. 2014). These datasets cover different types of scenes or objects and are benchmarks for NeRF-based methods. Furthermore, we additionally use 2 self-created scenes to show that our method is better than previous methods in special scenarios (see supplementary material for more details).

Implementation details. We sample 1024 rays per batch and train our model for 400k iterations for 10 hours on a single NVIDIA RTX3090 GPU. Our network architecture and initialization scheme are similar to those of NeRF (Mildenhall et al. 2021) and we model the background using NeRF++ (Zhang et al. 2020) with the same settings as NeuS in the real scene dataset.

Metrics. Following most 3D reconstruction methods, we measure the reconstruction quality with the Chamfer distances (Wang et al. 2021; Wang, Skorokhodov, and Wonka 2022, 2023; Fu et al. 2022; Darmon et al. 2022).

Overall Performance

Comparison methods. We compared our method with 2 kinds of NeRF-based methods, conventional NeRF-based methods and SDF-based methods. For conventional NeRF-based methods, we compared with NeRF (Mildenhall et al. 2021) and MipNeRF (Barron et al. 2021) (referred to as ‘‘Mip’’ in the table). For SDF-based methods, we compare with NeuS (Wang et al. 2021), HF-NeuS (Wang, Skorokhodov, and Wonka 2022) (referred to as ‘‘HFS’’ in the table), and PET-NeuS (Wang, Skorokhodov, and Wonka 2023) (referred to as ‘‘PET’’ in the table). PET-NeuS and HF-NeuS are follow-up works on NeuS and achieve more detailed geometry reconstruction. We did not include IDR (Yariv et al. 2020), UNISURF (Oechsle, Peng, and Geiger 2021), or VolSDF (Yariv et al. 2021) as NeuS had shown superior performance on Chamfer distances. We report the Chamfer distances in Tab. 1, Tab. 2, and Tab. 3, and conduct qualitative comparisons in Fig. 3.

Quantitative comparison. These SDF-based methods failed in semi-transparent scenes (shown in Tab. 1), resulting in significant Chamfer distances. This indicates that the SDF-based method cannot reconstruct semi-transparent scenes. Meanwhile, as shown in Tab. 2, our average Chamfer distances are the lowest in the Blender dataset (Mildenhall et al. 2021). This indicates that our method can obtain more accurate geometric information. It is rare (Ren et al. 2023) that the result of SNN is better than that of ANN. As shown in Tab. 3, we have a similar average Chamfer distance to SDF-based methods in the DTU dataset. While our method achieves leading results, some of our results are less competitive compared to some mainstream methods (e.g. PET-NeuS, HF-NeuS). Because SDF-based methods adopt lower frequency position encoding and constrain the variation of output with input, they have good performance

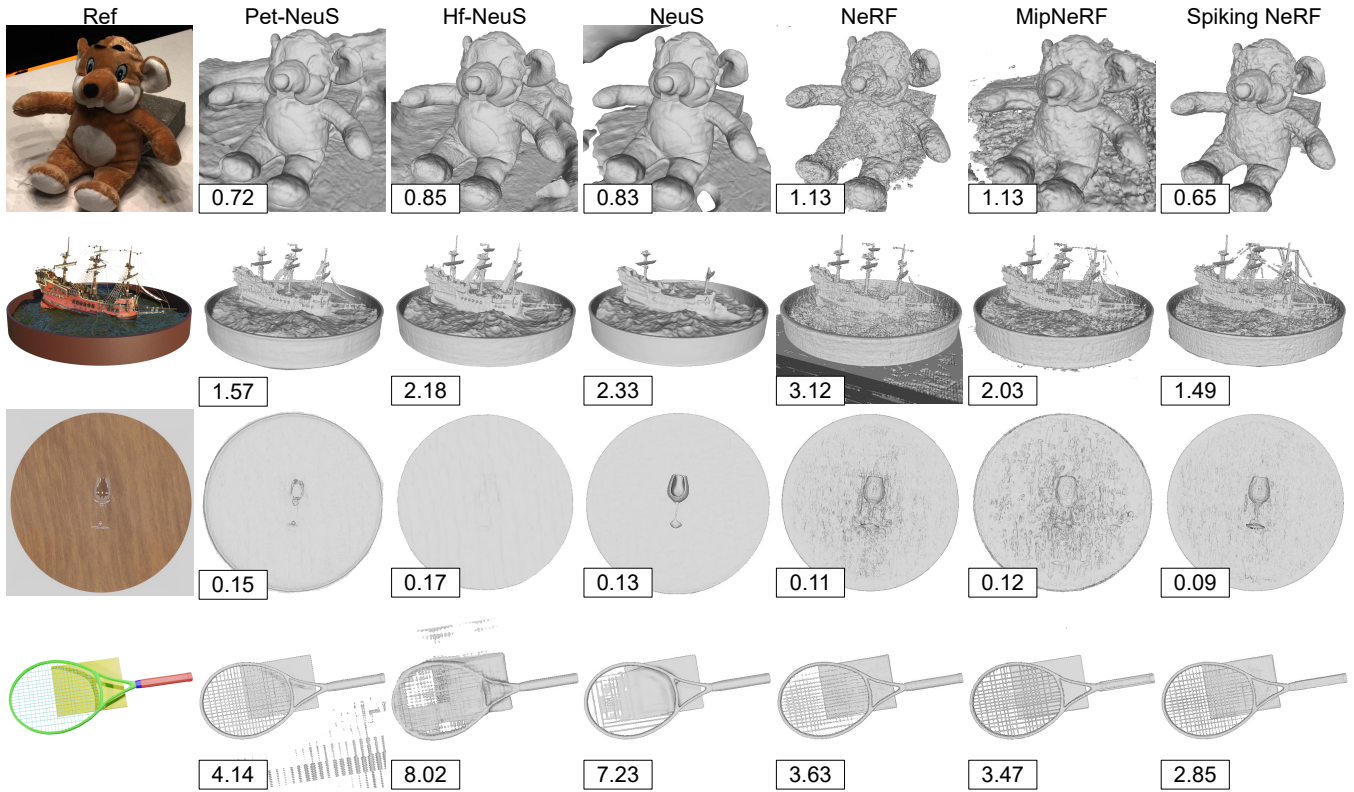


Figure 3: Visual Quality Comparisons on surface reconstruction on Blender dataset (Mildenhall et al. 2021), DTU dataset (Jensen et al. 2014), semi-transparent dataset (Ichnowski et al. 2021), and thin object dataset. We show the Chamfer distance in the bottom left corner of the image. The results of the 2nd and 4th rows are multiplied by 10^2 .

	wheel	tennis	mount	glass	turbine	clutter	pawn	pipe	Avg.
PET	13.23	4.14	15.94	15.32	15.34	23.28	14.31	14.23	14.47
HFS	30.41	8.02	15.35	17.24	11.53	17.23	15.78	12.57	16.01
NeuS	23.85	7.23	12.75	13.14	11.70	25.32	14.30	11.53	14.97
NeRF	7.47	3.63	10.22	11.53	10.06	13.15	12.04	10.12	9.77
Mip	11.31	3.47	10.30	12.74	10.25	12.36	11.67	9.23	10.16
Ours	6.42	2.85	9.87	9.45	9.82	10.44	10.87	10.32	8.75

Table 1: Quantitative Comparison on light scenarios (Ichnowski et al. 2021). We show the Chamfer distance $\times 10^{-2}$ for the reconstructed geometry on 6 scenes from the Dex-NeRF dataset and 2 scenes from our Blender dataset.

for low-frequency and smooth scenes. However, when these premises are not met, their performance will decrease.

Since our method is based on the conventional NeRF and is similar to NeRF, the robustness of our method closely resembles that of NeRF. The actual performance can be seen in Tab. 2 and Tab. 3 which show results on the general real scene (i.e. DTU dataset (Jensen et al. 2014)) and the general synthetic scene (i.e. Blender dataset (Mildenhall et al. 2021)). For conventional NeRF-based methods, the results obtained by manually selecting the optimal threshold are worse than our methods. Because our method uses the hybrid ANN-SNN framework to model the 3D geometric information in a discontinuous representation, the recon-

	Lego	Chair	Mic	Ficus	Hotdog	Drums	Mats	Ship	Avg.
PET	0.58	0.65	0.59	0.71	1.02	2.53	1.05	1.57	1.09
HFS	0.96	0.65	0.72	0.87	1.35	3.82	1.08	2.18	1.45
NeuS	1.52	0.70	0.85	1.67	1.40	4.27	1.08	2.33	1.73
NeRF	2.06	0.75	0.95	0.56	1.83	3.38	1.12	3.12	1.72
Mip	1.92	0.90	1.13	0.55	1.98	3.34	1.30	2.03	1.64
Ours	0.70	0.66	0.72	0.54	0.94	2.43	1.10	1.49	1.07

Table 2: Quantitative Comparison on Blender. We show the Chamfer distance $\times 10^{-2}$ for the reconstructed geometry on 8 scenes from the Blender dataset.

structed geometric information is more accurate.

Visual quality comparison. As shown in the 5th and 6th columns of Fig. 3, with conventional NeRF-based methods, due to the lack of surface constraints, the reconstruction results have many erroneous high-frequency information. Additionally, there is a loss of reconstruction accuracy for thin objects and semi-transparent scenes. As shown in the 2nd, 3rd, and 4th columns of Fig. 3, with SDF-based methods, the reconstruction resulted in significant errors in high-frequency changes in the image, such as failing to reconstruct holes that should have been present and surfaces that should have been separate. Moreover, NeuS demonstrated excessive smoothness in the 2nd row of Fig. 3.

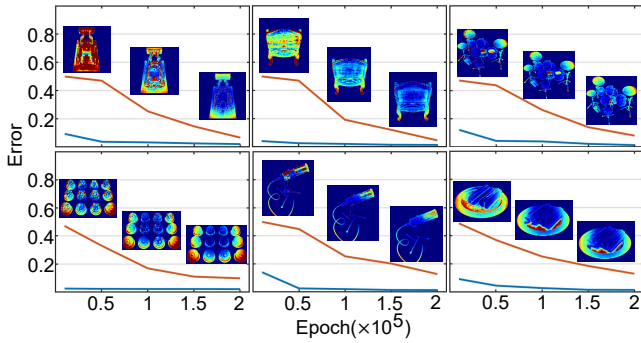


Figure 4: The relationship between the upper bound and the average depth error during training. We show 6 scenes from Blender dataset (Mildenhall et al. 2021). We randomly choose a view for displaying from each scene, and compute error and upper bound in epoch = 10K, 50K, 100K, 150K and 200K. The red curve represents the upper bound while the blue curve represents the average depth error during training. It can be seen that the average depth error decreases with the upper bound and the average depth error keeps being less than the upper bound during training.

HF-NeuS, and PET-NeuS fail in the reconstruction of semi-transparent scenes. Even though HF-NeuS uses some methods to reconstruct the high-frequency information, it still fails in some thin high-frequency positions. Because our method uses the hybrid ANN-SNN framework to model the 3D geometric information in a discontinuous representation, the reconstructed geometric information is more accurate.

The Practicality of Upper Bound of Proposition 1

To explore whether the bound is practical, we verify whether there is a correlation between the bound and the average depth error and whether the bound has similar patterns across different scenarios. Because the Blender dataset provides depth information, we conducted our experiments on it. As shown in Fig. 4, the average depth error is under the bound. Moreover, there is indeed a correlation between the bound and the average depth error, and the bound has similar patterns across different scenarios. It means that the bound can be leveraged to facilitate real-world applications for geometric estimation without knowing the ground truth. However, note that there is still a relatively large gap between the bound and the average depth error. While we made a new attempt in the field of NeRF and this bound can be used to a certain extent, it is still not enough to be practical.

Ablation Study

Validation for proposition 1. When L_v is minimized, the spiking threshold will continue to increase, and when L_v is removed, the spiking threshold of the final network will remain at a lower value. To validate proposition 1, we design an experiment without L_v and compare results with our method. It can be seen that there are many areas that should not exist above the reconstructed surface (see the 3rd column of Fig. 5), indicating that a relatively large spiking threshold is necessary, which is consistent with our analysis.




Figure 5: Ablation studies. We show qualitative results and report the quantitative metrics in Chamfer distance.

	37	40	55	63	65	83	97	105	110	114	118	Avg.
PET	0.84	0.72	0.40	0.95	0.84	1.42	1.14	0.72	1.05	0.44	0.59	0.83
HFS	1.37	0.78	0.47	1.11	0.68	1.20	1.17	0.85	1.27	0.38	0.54	0.90
NeuS	1.37	0.93	0.43	1.10	0.70	1.48	1.16	0.83	1.69	0.35	0.49	0.96
NeRF	1.63	1.75	0.60	2.03	1.07	1.70	1.95	1.13	2.33	0.87	1.05	1.46
Mip	1.83	1.65	1.61	2.90	1.79	1.93	2.19	1.32	2.52	1.38	1.49	1.87
Ours	1.20	1.02	0.38	1.15	0.72	1.10	1.19	0.65	1.60	0.49	0.55	0.90

Table 3: Quantitative Comparison on DTU (Jensen et al. 2014). We show a quantitative comparison for the reconstructed geometry on 11 scans from the DTU dataset. See supplementary material for more scenes.

Effectiveness of B-FIF. To assess the effectiveness of B-FIF, we design an experiment that replaced B-FIF with FIF which does not have a maximum activation constraint, and compare the result with our method. It can be seen that the reconstructed surface has significant deviation (see the 2nd column of Fig. 5), indicating that a bound is necessary, which is consistent with our analysis.

Effectiveness of smoothness. To assess the effectiveness of the initial smoothing, we design an experiment without L_g and compare the result with our method. It can be seen that although the Chamfer distances are similar, the reconstructed surface contains wrong high-frequency information (see the 4th column of Fig. 5), indicating that training from the wrong high-frequency surface is difficult.

Conclusion

We present a novel neural surface reconstruction method for reconstructing objects. NeRF and its variants use volume rendering to produce a neural scene representation. However, extracting high-quality surfaces from this learned implicit representation is difficult because there is not a definite interface density. In our work, we propose to use a hybrid ANN-SNN framework to reconstruct density fields. We observe that the conventional volume rendering method causes inherent geometric errors for surface reconstruction, and propose a new spiking neuron to get more accurate surface reconstruction with an adaptive spiking threshold.

Limitation. Since NeRF does not specifically consider scenes with high light and low brightness, it will struggle to accurately represent the geometry on these scenes (Verbin et al. 2022). Our method also did not specifically consider these issues, so it is highly likely to struggle to accurately represent the geometry on these scenes.

Acknowledgements

This work is supported by the National Natural Science Foundation of China under Grand No. 62376247. This work is supported by the State Key Lab of Brain-Machine Intelligence.

References

- Aimone, J.; Date, P.; Fonseca-Guerra, G.; Hamilton, K.; Henke, K.; Kay, B.; Kenyon, G.; Kulkarni, S.; Mniszewski, S.; Parsa, M.; et al. 2022. A Review of Non-Cognitive Applications for Neuromorphic Computing. *Neuromorphic Computing and Engineering*.
- Atzman, M.; Haim, N.; Yariv, L.; Israelov, O.; Maron, H.; and Lipman, Y. 2019. Controlling Neural Level Sets. *NIPS*.
- Azinović, D.; Martin-Brualla, R.; Goldman, D. B.; Nießner, M.; and Thies, J. 2022. Neural RGB-D Surface Reconstruction. In *CVPR*.
- Bagheri, A.; Simeone, O.; and Rajendran, B. 2018. Adversarial Training for Probabilistic Spiking Neural Networks. In *2018 IEEE 19th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*.
- Barron, J. T.; Mildenhall, B.; Tancik, M.; Hedman, P.; Martin-Brualla, R.; and Srinivasan, P. P. 2021. Mip-NeRF: A Multiscale Representation for Anti-Aliasing Neural Radiance Fields. In *ICCV*.
- Boss, M.; Braun, R.; Jampani, V.; Barron, J. T.; Liu, C.; and Lensch, H. 2021. NeRD: Neural Reflectance Decomposition From Image Collections. In *ICCV*.
- Darmon, F.; Bascle, B.; Devaux, J.-C.; Monasse, P.; and Aubry, M. 2022. Improving Neural Implicit Surfaces Geometry with Patch Warping. In *CVPR*.
- Dayan, P.; and Abbott, L. F. 2005. *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*. MIT Press.
- Deng, K.; Liu, A.; Zhu, J.-Y.; and Ramanan, D. 2022. Depth-Supervised NeRF: Fewer Views and Faster Training for Free. In *CVPR*.
- Fang, W.; Yu, Z.; Chen, Y.; Huang, T.; Masquelier, T.; and Tian, Y. 2021. Deep Residual Learning in Spiking Neural Networks. *NIPS*.
- Fu, Q.; Xu, Q.; Ong, Y. S.; and Tao, W. 2022. Geo-Neus: Geometry-Consistent Neural Implicit Surfaces Learning for Multi-view Reconstruction. *NIPS*.
- Gehrig, M.; Shrestha, S. B.; Mouritzen, D.; and Scaramuzza, D. 2020. Event-Based Angular Velocity Regression with Spiking Networks. In *ICRA*.
- Guo, Y.-C.; Kang, D.; Bao, L.; He, Y.; and Zhang, S.-H. 2022. NeRFReN: Neural Radiance Fields with Reflections. In *CVPR*.
- Han, J.; Wang, Z.; Shen, J.; and Tang, H. 2023. Symmetric-Threshold ReLU for Fast and Nearly Lossless ANN-SNN Conversion. *Machine Intelligence Research*.
- Ichnowski, J.; Avigal, Y.; Kerr, J.; and Goldberg, K. 2021. Dex-NeRF: Using a Neural Radiance Field to Grasp Transparent Objects. In *5th Annual Conference on Robot Learning*.
- Jensen, R.; Dahl, A.; Vogiatzis, G.; Tola, E.; and Aanaes, H. 2014. Large Scale Multi-View Stereopsis Evaluation. In *CVPR*.
- Kugele, A.; Pfeil, T.; Pfeiffer, M.; and Chicca, E. 2021. Hybrid SNN-ANN: Energy-Efficient Classification and Object Detection for Event-Based Vision. In *DAGM German Conference on Pattern Recognition*.
- Lee, C.; Kosta, A. K.; 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*.
- Lele, A. S.; Fang, Y.; Anwar, A.; and Raychowdhury, A. 2022. Bio-Mimetic High-Speed Target Localization with Fused Frame and Event Vision for Edge Application. *Frontiers in Neuroscience*.
- Levy, D.; Peleg, A.; Pearl, N.; Rosenbaum, D.; Akkaynak, D.; Korman, S.; and Treibitz, T. 2023. SeaThru-NeRF: Neural Radiance Fields in Scattering Media. In *CVPR*.
- Li, S.; Zhang, Z.; Mao, R.; Xiao, J.; Chang, L.; and Zhou, J. 2021a. A Fast and Energy-Efficient SNN Processor With Adaptive Clock/Event-Driven Computation Scheme and Online Learning. *IEEE Transactions on Circuits and Systems I: Regular Papers*.
- Li, W.; Chen, H.; Guo, J.; Zhang, Z.; and Wang, Y. 2022. Brain-Inspired Multilayer Perceptron with Spiking Neurons. In *CVPR*.
- Li, Y.; Deng, S.; Dong, X.; Gong, R.; and Gu, S. 2021b. A Free Lunch from ANN: Towards Efficient, Accurate Spiking Neural Networks Calibration. In *ICML*.
- Li, Y.; Guo, Y.; Zhang, S.; Deng, S.; Hai, Y.; and Gu, S. 2021c. Differentiable Spike: Rethinking Gradient-Descent for Training Spiking Neural Networks. *NIPS*.
- Liu, F.; and Zhao, R. 2022. Enhancing Spiking Neural Networks with Hybrid Top-Down Attention. *Frontiers in Neuroscience*.
- Liu, L.; Gu, J.; Zaw Lin, K.; Chua, T.-S.; and Theobalt, C. 2020a. Neural Sparse Voxel Fields. *NIPS*.
- Liu, S.; Zhang, Y.; Peng, S.; Shi, B.; Pollefeys, M.; and Cui, Z. 2020b. DIST: Rendering Deep Implicit Signed Distance Function With Differentiable Sphere Tracing. In *CVPR*.
- Lu, L.; Jin, P.; Pang, G.; Zhang, Z.; and Karniadakis, G. E. 2021. Learning Nonlinear Operators via DeepONet Based on the Universal Approximation Theorem of Operators. *Nature Machine Intelligence*.
- Marchisio, A.; Nanfa, G.; Khalid, F.; Hanif, M. A.; Martina, M.; and Shafique, M. 2020. Is Spiking Secure? A Comparative Study on the Security Vulnerabilities of Spiking and Deep Neural Networks. In *IJCNN*.
- Mildenhall, B.; Srinivasan, P. P.; Tancik, M.; Barron, J. T.; Ramamoorthi, R.; and Ng, R. 2021. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. *Communications of the ACM*.
- Neftci, E. O.; Mostafa, H.; and Zenke, F. 2019. Surrogate Gradient Learning in Spiking Neural Networks: Bringing the Power of Gradient-Based Optimization to Spiking Neural Networks. *IEEE Signal Processing Magazine*.

- Niemeyer, M.; Mescheder, L.; Oechsle, M.; and Geiger, A. 2020. Differentiable Volumetric Rendering: Learning Implicit 3D Representations Without 3D Supervision. In *CVPR*.
- Oechsle, M.; Peng, S.; and Geiger, A. 2021. UNISURF: Unifying Neural Implicit Surfaces and Radiance Fields for Multi-View Reconstruction. In *ICCV*.
- Ororbia, A. 2023. Spiking Neural Predictive Coding for Continually Learning from Data Streams. *Neurocomputing*.
- Peng, S.; Niemeyer, M.; Mescheder, L.; Pollefeys, M.; and Geiger, A. 2020. Convolutional Occupancy Networks. In *ECCV*.
- Peng, S.; Zhang, Y.; Xu, Y.; Wang, Q.; Shuai, Q.; Bao, H.; and Zhou, X. 2021. Neural Body: Implicit Neural Representations With Structured Latent Codes for Novel View Synthesis of Dynamic Humans. In *CVPR*.
- Pfeiffer, M.; and Pfeil, T. 2018. Deep Learning with Spiking Neurons: Opportunities and Challenges. *Frontiers in Neuroscience*.
- Pumarola, A.; Corona, E.; Pons-Moll, G.; and Moreno-Noguer, F. 2021. D-NeRF: Neural Radiance Fields for Dynamic Scenes. In *CVPR*.
- Rañon, U.; Cuadrado-Anibarro, J.; Cottureau, B. R.; and Masquelier, T. 2022. StereoSpike: Depth Learning With a Spiking Neural Network. *IEEE Access*.
- Ren, D.; Ma, Z.; Chen, Y.; Peng, W.; Liu, X.; Zhang, Y.; and Guo, Y. 2023. Spiking PointNet: Spiking Neural Networks for Point Clouds. In *NIPS*.
- Saito, S.; Huang, Z.; Natsume, R.; Morishima, S.; Kanazawa, A.; and Li, H. 2019. PIFu: Pixel-Aligned Implicit Function for High-Resolution Clothed Human Digitization. In *ICCV*.
- Schwarz, K.; Liao, Y.; Niemeyer, M.; and Geiger, A. 2020. GRAF: Generative Radiance Fields for 3D-Aware Image Synthesis. *NIPS*.
- Sharmin, S.; Rathi, N.; Panda, P.; and Roy, K. 2020. Inherent Adversarial Robustness of Deep Spiking Neural Networks: Effects of Discrete Input Encoding and Non-linear Activations. In *ECCV*.
- Shrestha, S. B.; and Orchard, G. 2018. SLAYER: Spike Layer Error Reassignment in Time. *NIPS*.
- Sitzmann, V.; Zollhöfer, M.; and Wetzstein, G. 2019. Scene Representation Networks: Continuous 3D-Structure-Aware Neural Scene Representations. *NIPS*.
- Verbin, D.; Hedman, P.; Mildenhall, B.; Zickler, T.; Barron, J. T.; and Srinivasan, P. P. 2022. Ref-NeRF: Structured View-Dependent Appearance for Neural Radiance Fields. In *CVPR*.
- Wang, P.; Liu, L.; Liu, Y.; Theobalt, C.; Komura, T.; and Wang, W. 2021. NeuS: Learning Neural Implicit Surfaces by Volume Rendering for Multi-view Reconstruction. In *NIPS*.
- Wang, Y.; Skorokhodov, I.; and Wonka, P. 2022. HF-NeuS: Improved Surface Reconstruction Using High-Frequency Details. *NIPS*.
- Wang, Y.; Skorokhodov, I.; and Wonka, P. 2023. PET-NeuS: Positional Encoding Tri-Planes for Neural Surfaces. In *CVPR*.
- Xu, Q.; Wang, W.; Ceylan, D.; Mech, R.; and Neumann, U. 2019. DISN: Deep Implicit Surface Network for High-quality Single-view 3D Reconstruction. *NIPS*.
- Yariv, L.; Gu, J.; Kasten, Y.; and Lipman, Y. 2021. Volume Rendering of Neural Implicit Surfaces. *NIPS*.
- Yariv, L.; Kasten, Y.; Moran, D.; Galun, M.; Atzmon, M.; Ronen, B.; and Lipman, Y. 2020. Multiview Neural Surface Reconstruction by Disentangling Geometry and Appearance. *NIPS*.
- Zhang, C.; Huang, C.; and He, J. 2023. Defects Recognition of Train Wheelset Tread Based on Improved Spiking Neural Network. *Chinese Journal of Electronics*.
- Zhang, J.; Dong, B.; Zhang, H.; Ding, J.; Heide, F.; Yin, B.; and Yang, X. 2022. Spiking Transformers for Event-Based Single Object Tracking. In *CVPR*.
- Zhang, J.; Huang, L.; Ma, Z.; and Zhou, H. 2023. Predicting the Temporal-Dynamic Trajectories of Cortical Neuronal Responses in Non-Human Primates Based on Deep Spiking Neural Network. *Cognitive Neurodynamics*.
- Zhang, K.; Riegler, G.; Snavely, N.; and Koltun, V. 2020. NeRF++: Analyzing and Improving Neural Radiance Fields. *arXiv preprint arXiv:2010.07492*.
- Zhang, X.; Srinivasan, P. P.; Deng, B.; Debevec, P.; Freeman, W. T.; and Barron, J. T. 2021. NeRFactor: Neural Factorization of Shape and Reflectance Under an Unknown Illumination. *TOG*.
- Zhao, R.; Yang, Z.; Zheng, H.; Wu, Y.; Liu, F.; Wu, Z.; Li, L.; Chen, F.; Song, S.; Zhu, J.; et al. 2022. A Framework for the General Design and Computation of Hybrid Neural Networks. *Nature Communications*.
- Zhu, L.; Wang, X.; Chang, Y.; Li, J.; Huang, T.; and Tian, Y. 2022. Event-Based Video Reconstruction via Potential-Assisted Spiking Neural Network. In *CVPR*.
- Zou, X.-L.; Huang, T.-J.; and Wu, S. 2022. Towards a New Paradigm for Brain-Inspired Computer Vision. *Machine Intelligence Research*.