

# Reconstruction of natural visual scenes from neural spikes with deep neural networks

Yichen Zhang<sup>a,b</sup>, Shanshan Jia<sup>a,b</sup>, Yajing Zheng<sup>a,b</sup>, Zhaofei Yu<sup>a,b,\*</sup>, Yonghong Tian<sup>a,b</sup>, Siwei Ma<sup>a,b</sup>, Tiejun Huang<sup>a,b</sup>, Jian K. Liu<sup>c,\*</sup>

<sup>a</sup> National Engineering Laboratory for Video Technology, School of Electronics Engineering and Computer Science, Peking University, Beijing, China

<sup>b</sup> Peng Cheng Laboratory, Shenzhen, China

<sup>c</sup> Centre for Systems Neuroscience, Department of Neuroscience, Psychology and Behavior, University of Leicester, Leicester, UK

## ARTICLE INFO

### Article history:

Received 1 July 2019

Received in revised form 16 January 2020

Accepted 28 January 2020

Available online 8 February 2020

### Keywords:

Vision

Natural scenes

Neural decoding

Neural spikes

Deep learning

Artificial retina

## ABSTRACT

Neural coding is one of the central questions in systems neuroscience for understanding how the brain processes stimulus from the environment, moreover, it is also a cornerstone for designing algorithms of brain-machine interface, where decoding incoming stimulus is highly demanded for better performance of physical devices. Traditionally researchers have focused on functional magnetic resonance imaging (fMRI) data as the neural signals of interest for decoding visual scenes. However, our visual perception operates in a fast time scale of millisecond in terms of an event termed neural spike. There are few studies of decoding by using spikes. Here we fulfill this aim by developing a novel decoding framework based on deep neural networks, named spike-image decoder (SID), for reconstructing natural visual scenes, including static images and dynamic videos, from experimentally recorded spikes of a population of retinal ganglion cells. The SID is an end-to-end decoder with one end as neural spikes and the other end as images, which can be trained directly such that visual scenes are reconstructed from spikes in a highly accurate fashion. Our SID also outperforms on the reconstruction of visual stimulus compared to existing fMRI decoding models. In addition, with the aid of a spike encoder, we show that SID can be generalized to arbitrary visual scenes by using the image datasets of MNIST, CIFAR10, and CIFAR100. Furthermore, with a pre-trained SID, one can decode any dynamic videos to achieve real-time encoding and decoding of visual scenes by spikes. Altogether, our results shed new light on neuromorphic computing for artificial visual systems, such as event-based visual cameras and visual neuroprostheses.

© 2020 Elsevier Ltd. All rights reserved.

## 1. Introduction

In everyday life, various types of sensory information are processed by our brain through different sensory modalities, which are then processed to generate a series of behavior reactions. Neuronal networks in the brain, then, play an essential role in efficient and powerful computation where a sequence of input-output mappings is carried out by single and networked neurons. At the level of the system, single neurons receive and respond to input stimuli by changing their membrane potential to generate a sequence of fast events, termed neural spikes. Thus, spikes have been suggested as a fundamental element to represent input-output neural computation of visual scenes (Chichilnisky, 2001; Gollisch & Meister, 2008). Such observations lead to a central question for systems neuroscience on the sensory system: how

neurons represent the input-output relationship between stimuli and their spikes (Rieke, Warland, de Ruyter van Steveninck, & Bialek, 1996; Simoncelli & Olshausen, 2001; Wu, David, & Gallant, 2006). This question, formulated as neural coding, consists of two essential parts, encoding and decoding. For the visual system, it is to understand how visual scenes are represented by neural spiking activity, and how to decode them to reconstruct the input visual scenes.

The retina serves as a useful system to study these questions. Visual scenes are projected into the eyes, where the retinal neuronal network processes the input by a few types of neurons, but the only output neurons are the retinal ganglion cells (RGCs). Therefore, all information of visual scenes is encoded by RGCs that produce a sequence of action potentials or spikes, which are transmitted via the optic nerve to the downstream brain regions. Essentially, visual scenes are represented by spikes of a population of RGCs. One expects that encoding and decoding of visual scenes from RGC spikes can be formulated into a closed form where a computational model can be developed to set up

\* Corresponding authors.

E-mail addresses: [yuzf12@pku.edu.cn](mailto:yuzf12@pku.edu.cn) (Z. Yu), [jian.liu@leicester.ac.uk](mailto:jian.liu@leicester.ac.uk) (J.K. Liu).

a mapping between visual scenes and RGC spikes. Indeed, much effort has been given to study the encoding part of RGCs. Various neuroscience mechanisms have been identified in understanding the retinal computation of visual scenes by its neurons and neural circuitry (Demb & Singer, 2015; Gollisch & Meister, 2010; Grimes, Songco-Aguas, & Rieke, 2018; Jadzinsky & Baccus, 2013; O'Brien & Bloomfield, 2018; Rivlin-Etzion, Grimes, & Rieke, 2018). In addition, for understanding the encoding principles of the retina, a number of models are developed based on different properties of neurons and neural circuits in the retina (Meyer, Williamson, Linden, & Sahani, 2017; Pillow et al., 2008; Yan, Yu, Chen, & Liu, 2017; Yan et al., 2018; Yu et al., 2020).

Decoding of RGC spikes to obtain the essential information of visual scenes is processed by the downstream neurons from the lateral geniculate nucleus to the visual cortex, which involves feedback from the higher part of the cortex for cognition. However, as the retina does not receive the feedback, the RGC spikes can be thought as a minimal computational device to represent visual information as a whole. Therefore, it is suitable to develop a decoding model that can reconstruct visual scenes from the RGC spikes directly (Botella-Soler, Deny, Martius, Marre, & Tkačik, 2018; Gollisch & Meister, 2008). This is in particular important for neuroengineering, such as neuroprosthesis, where efficient algorithms of decoding neuronal signals for controlling physical devices are highly demanded (Gilja et al., 2012). In general, an ideal neuronal decoder should be able to read out and reconstruct stimulus from neural responses. The existing methods for visual scenes reconstruction are mainly divided into two types. The first one is to decode the stimulus category: choosing a stimulus according to neural responses, then classifying neural signals to find the corresponding stimulus in a candidate set (Wen et al., 2018). The second type is to reconstruct the original stimulus directly by using neural responses, i.e., obtaining every pixel of visual scenes from the neural signal. It is obvious that the second category is more challenging (Nishimoto et al., 2011).

Reconstruction of visual scenes has been studied over many years. The neural signals of interest can be functional magnetic resonance imaging (fMRI) activity (Naselaris, Prenger, Kay, Oliver, & Gallant, 2009; Nishimoto et al., 2011; Qiao et al., 2018; Thirion et al., 2006; Wen et al., 2018), neural spikes in the retina (Botella-Soler et al., 2018; Gollisch & Meister, 2008; Marre et al., 2015; Parthasarathy et al., 2017) and lateral geniculate nucleus (Stanley, Li, & Dan, 1999), neural calcium imaging signals in the visual cortex (Garasto, Bharath, & Schultz, 2018). However, the decoding performance of current methods is rather low for natural scenes, in particular for dynamical videos, which can be seen from some typical examples of the videos reconstructed from fMRI signals (Nishimoto et al., 2011; Wen et al., 2018). For the retina, one would expect to decode visual scenes by using the spiking responses of a population of RGCs in a complete way as these spikes are the only output of the eyes. Decoding visual scenes, at least for static natural images, is possible when sampling the whole image with a number of RGCs (Gollisch & Meister, 2008). Yet, it remains unclear how to deal with the dynamic natural scenes, where the temporal complexity of stimulus information has a strong coupling with the temporal adaption of neurons (Liu & Gollisch, 2015).

In this study, we propose such a general-purpose decoding framework, termed spike-image decoder (SID), that performs an end-to-end decoding process from neural spikes to visual scenes, based on deep learning neural networks. The proposed SID can achieve state-of-the-art performance, compared to previous studies, for reconstructing natural visual scenes, including both static images and dynamic videos, from spikes of a population of RGCs recorded simultaneously in the isolated animal retina. The workflow of the SID is illustrated in Fig. 1. When a large population of

RGCs is recorded simultaneously, and their spikes are extracted, a spike-image converter based on a neural network is used to map the spikes of every RGC to intermediate images at the pixel level. After that, an autoencoder-type deep learning neural network is applied to spike-based intermediate images to match original stimulus images for every pixel.

Essentially, the SID has two stages with one as spike-image converter and another one as image-image autoencoder. Most of the previous studies focused on the first stage, such that the spikes are directly mapped to obtain the image pixels, where, traditionally, a decoder can be optimized by some statistical models in a linear or nonlinear fashion (Botella-Soler et al., 2018; Gollisch & Meister, 2008; Marre et al., 2015; Naselaris et al., 2009; Nishimoto et al., 2011; Stanley et al., 1999; Thirion et al., 2006). A recent study used a separate autoencoder network as the second stage to enhance the quality of images (Parthasarathy et al., 2017), which essentially employed a similar network model dedicated for image denoising and enhancement (Yue et al., 2018). Here, we systematically tested different training approaches of neural networks from spikes to images, and found the SID can achieve similar performance for static images across three types of loss functions, which is consistent with those previous studies, however for dynamic videos, the direct training from spikes to images shows an outstanding performance.

We tested the performance of SID on experimentally recorded RGCs spikes, and found it can obtain state-of-the-art performance for the reconstruction of natural visual scenes of static image and dynamic video. Our SID can also be used to decode the fMRI signal, and has a better performance on the reconstruction of visual stimulus compared to the existing state-of-the-art fMRI decoding models. Furthermore, we found that the computational ability of SID can be generalized to arbitrary visual scenes with the aid of an encoding model that can simulate spikes for any given images. By using the image datasets of MNIST, CIFAR10, and CIFAR100, the SID can reconstruct natural images with high precision. In addition, with a pre-trained SID model, one can reconstruct any given video scenes to realize real-time encoding and decoding of dynamic visual scenes as an artificial vision system.

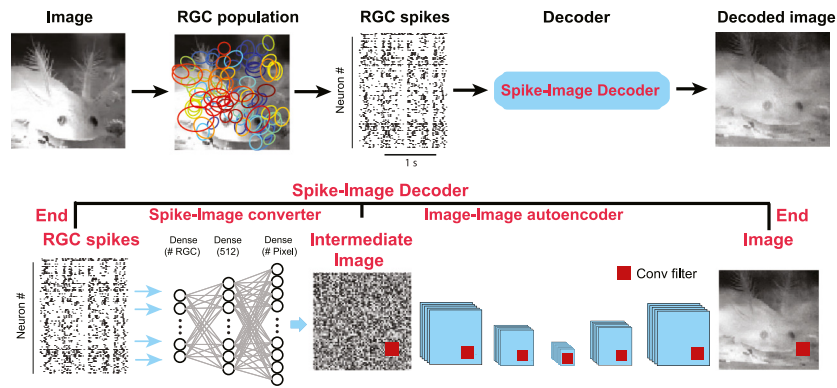
## 2. Methods

### 2.1. Spike-image decoder

Traditionally, it has been shown that deep neural networks can extract low-level features in earlier layers, and more semantic features in later layers, which is similar to the information processing of the visual cortex in the brain (Yamins & Dicarlo, 2016; Zeiler & Fergus, 2014). Here we used a similar deep neural network to reconstruct natural scenes from the retinal spikes. Our decoding model consists of two parts: (1) spike-to-image converter and (2) image-to-image autoencoder. The spike-to-image converter is able to map every spike to every pixel to get intermediate images, which plays a role of nonlinear re-sampling, and the image-to-image autoencoder is used to match all pixels to those of the target images of visual scenes.

#### 2.1.1. Spike-to-image converter

The spike-to-image converter is a three-layer fully-connected neural network, similar to a multilayer perceptron (MLP). The first layer receives the spikes of all RGCs as input such that the number of neurons of the first layer is matched to the number of RGCs used. The second layer is a hidden layer, which consists of 512 neurons. ReLU is used as an activation function of the hidden layer. The output layer has the same number of neurons as the pixel number of the input image, in our case, 4096 neurons were used since all stimulus images are in  $64 \times 64$  pixels. ReLU as activation function is used for the third layer, and we can get all pixel values of the intermediate image from this output layer.



**Fig. 1.** Illustration of the spike-image decoder model. (Top) Workflow of decoding visual scenes. Here a salamander swimming video was presented to a salamander retina to stimulate a firing of spike trains in a population of RGCs. The spikes are used to train the decoder for the given video to reconstruct the visual scene. Receptive fields of ganglion cells are mapped onto the image. Each colored circle is an outline of receptive field. (Bottom) Spike-image decoder is an end-to-end decoder with two stages: spike-image converter used to map a population of spikes to a pixel-level intermediate image, and image-autoencoder for mapping every pixel to the target pixels in the desired images.

### 2.1.2. Image-to-image autoencoder

The image-to-image autoencoder is a typical deep autoencoder based on convolutional neural network (CNN), where the whole information processing procedure can be split into two phases. In the first phase, convolution and down-sampling are used to process and decrease the size of the input image. This phase contains four convolutional layers, the kernel sizes of these four layers are (64,7,7), (128,5,5), (256,3,3), (256,3,3) and the stride sizes are (2,2) for all these four layers. After this phase, the most important components of the input image are kept, and noise and redundant contents are filtered. In the second phase, convolution and up-sampling are used to process the image, which recovers the texture of the down-sampled image while increasing the size of the down-sampled image. The upsampling phase also consists of four convolutional layers, the kernel sizes are (256,3,3), (128,3,3), (64,5,5), (3,7,7) and the stride sizes are (1,1) for all these layers. Up-sampling layers with size (2,2) are added before these four convolutional layers. ReLU is used as an activation function in the whole Image-to-image autoencoder.

### 2.1.3. SID model architecture

Given an input image  $\mathbf{I}$ , it will trigger a response  $\mathbf{s} = \{s_1, s_2, \dots, s_n\}$  on the retinal ganglion cells, here we use rate coding such that  $s_i$  represents the spike count of each RGC within a time bin depending on the sampling rate of visual scenes. As in experimental recordings, a typical 30 Hz refreshing rate is used, so the time bin is about 33 ms. Then the triggered responses are first fed into the spike-to-image MLP converter, which outputs an intermediate image  $\mathbf{O}_1 = f_1(\mathbf{s})$  from RGC responses. Then the image-to-image CNN autoencoder takes the intermediate image as input, and is able to map it to match the target image, so one can get a clear and refining reconstruction result  $\mathbf{O}_2 = f_2(\mathbf{O}_1)$ . These two parts are all implemented with deep neural networks, so end-to-end training can be used to train both networks.

Three different ways of defining loss function were explored in this study. In the first one, one can use a multi-task loss function to make sure both stages have the capability to decode natural scenes from spikes, so that we can get a high quality reconstruction result from our decoding model. Here,  $\lambda_1$  and  $\lambda_2$  are weights of these two parts mean square error loss,  $\mathbf{L}_1 : Loss = \lambda_1 \|\mathbf{O}_1 - \mathbf{I}\| + \lambda_2 \|\mathbf{O}_2 - \mathbf{I}\|$ . In the second one, we optimize each stage of decoding separately such that  $\mathbf{L}_2 : (Loss_1 = \lambda_1 \|\mathbf{O}_1 - \mathbf{I}\|, Loss_2 = \lambda_2 \|\mathbf{O}_2 - \mathbf{I}\|)$ . The third one uses one loss for the whole process of two stages as follows:  $\mathbf{L}_3 : Loss = \lambda \|\mathbf{O} - \mathbf{I}\|$ , where there is only one image output  $\mathbf{O}$  as the final reconstruction to compare the target image such that one error signal stays in both

stages. In this way, the intermediate image is like noise without visible structures as in Fig. 1. Three loss functions were compared in detail.

The SID was implemented with Keras deep learning library, and Tensorflow as backend was used. During training for both spike-to-image converter and image-to-image autoencoder, batch normalization using 256 as batch size for both converter MLP and autoencoder CNN was used to achieve better performance (Ioffe & Szegedy, 2015) and dropout (Hinton, Srivastava, Krizhevsky, Sutskever, & Salakhutdinov, 2012) was used to deal with overfitting. Adam method was used to train the model. The code is available online (<https://sites.google.com/site/jiankliu>).

## 2.2. Experimental spike data

Our decoder was used to reconstruct natural visual scenes, including both static image and dynamic video, from spikes of a population of RGCs recorded simultaneously in isolated retinas of salamanders. Datasets of spikes recorded and stimulus images and movies used in this study can be found publicly (Onken et al., 2016). Experimental details can be found in previous studies (Liu et al., 2017; Onken et al., 2016). Briefly, neural data had been collected in isolated retinas obtained from axolotl salamanders. The retina was placed onto a 60- or 252-channel multielectrode array in a recording chamber. Visual stimuli were displayed on an OLED monitor that was projected onto the photoreceptor layer through a telecentric lens above the retina.

### 2.2.1. Image stimulus

The static images were taken from a dataset used previously (Liu et al., 2017; Onken et al., 2016), which is part of the McGill Calibrated Color Image Database (Olmos & Kingdom, 2004), including 300 natural images with a size of  $64 * 64$  pixels. Briefly, each image covers  $1920 * 1920 \mu\text{m}$  area on the retina, in which there are 80 RGCs recorded. Each individual image was presented for 200 ms and then followed by an 800 ms empty display in a pseudo-random sequence. For each image to each RGC, the number of spikes during 300 ms in one trial was collected. In addition, we averaged the spike counts among 13 trials for each image. Finally, we got the spike counts of a population of 80 RGCs for one image as the input data for reconstruction. Among all 300 images, 270 images were used to train the decoder model and the remaining 30 images were used as the test set.



### 2.2.2. Video stimulus

The video data consists of salamander swimming which includes 1800 frames as the same dataset in Onken et al. (2016). Each frame covered a total area of  $2700 \times 2700 \mu\text{m}$  on the retina with a spatial resolution of  $360 \times 360$  pixels, which was down-sampled into  $90 \times 90$  pixels, and cropped into  $64 \times 64$  pixels around the central region of images, in which 90 RGCs responses were collected. A segment of the video was presented at a frame rate of 30 Hz. For the recorded responses of each RGC, one can bin them into 33 ms that is the same with the presentation time of each frame. As a result, a spike train was collected with 1800 counts for each RGC in response to the video. Eventually for the population of RGCs, a  $90 \times 1800$  spike count matrix, averaged over 31 trials of repeated presentation, was used to reconstruct the video. To train the decoder model, we randomly selected 1620 frames from the video as a training set and the remaining frames as a test set. In this way, the strong temporal correlation within the video can be washed out.

### 2.3. Simulated spike data

To test our SID model for any arbitrary visual scenes, we used a simple encoder based on the typical linear–nonlinear model (Chichilnisky, 2001; Liu et al., 2017) to simulate neural responses of a population of RGCs. For comparison, the encoding model is based on the same experimental retinal data used for the analysis of the video. The linear filters are based on the receptive fields of a population of 90 RGCs obtained with white noise analysis (Chichilnisky, 2001; Liu et al., 2017; Onken et al., 2016) fitted with a 2D Gaussian for each cell. With the given natural scene stimulus (a static image or a frame of dynamic video), each cell implements a linear computation on pixels within its receptive field, and generates a spike count for this stimulus after a rectification nonlinearity. All these simulated spike counts (modeled responses of 90 RGCs) were then fed into our SID model to reconstruct natural scene stimulus.

With this retinal encoding model, one can simulate RGC responses under any given stimuli. In this study, simulations were done on three popular image datasets, MNIST, CIFAR10, and CIFAR100 (Krizhevsky, 2012; Lecun, Bottou, Bengio, & Haffner, 1998). MNIST is a dataset of handwritten digits that consists of a training set of 60,000 examples and a test set of 10,000 examples. CIFAR10 and CIFAR100 are datasets used for image classification, consisting of 60,000 color images in 10 and 100 classes respectively. There are 50,000 training images and 10,000 test images in both of CIFAR10 and CIFAR100. All images were first resized to the size of  $64 \times 64$  pixels, which is comparable to our retinal experimental data. Three channels of each color image were first encoded into spikes by the encoding model separately. The encoded (90,3) spikes are fed into the spike-to-image converter to get a color intermediate image, here the weights are shared for all three channels. Then image-to-image autoencoder takes the intermediate image as input to get the final reconstructed color image.

Similarly, for comparison, dynamic videos were also used for testing the SID. With the same encoding model, we used the same movie chips from the fMRI decoding experiments (Nishimoto et al., 2011), and obtained a set of RGC spiking responses similar to the biological data. Then the SID model pre-trained with CAFAR10 dataset was used to decode these dynamic videos directly.

Finally, we implemented real-time coding of arbitrary visual scenes by embedding our models into the power-efficient computing device, here NVIDIA Jetson TX2, which is built around a 256-core NVIDIA Pascal GPU. In order to better take advantage of the GPU to accelerate the computation, we combined the image-spike encoding process into the decoding neural networks. In the

real-time coding system, we can choose the input files such as image or video from the embedded device, or directly read the inputs, i.e. real-time videos, from a camera. However, as our SID was pre-trained with  $64 \times 64$  pixels as input stimulus and target output, the resolution is not very high such that it misses the fine detailed structure of visual scenes. Therefore, based on the pre-trained SID, we cropped the inputs stimulus into several patches with  $64 \times 64$  pixels, and then joined the decoding results after sending these data as a batch into the model simultaneously. Considering the GPU processing speed as well as the resolution, we cropped the input frames as  $5 \times 5$  patches, such that the resolution of input images and decoding results can be increased to  $320 \times 320$  in a real-time fashion. Higher resolution with real-time speed can be further achieved depending on the hardware used for this artificial vision system. The reconstruction results of the same videos used by the previous study (Nishimoto et al., 2011) and our real-time video decoding are available online (<https://sites.google.com/site/jiankliu>).

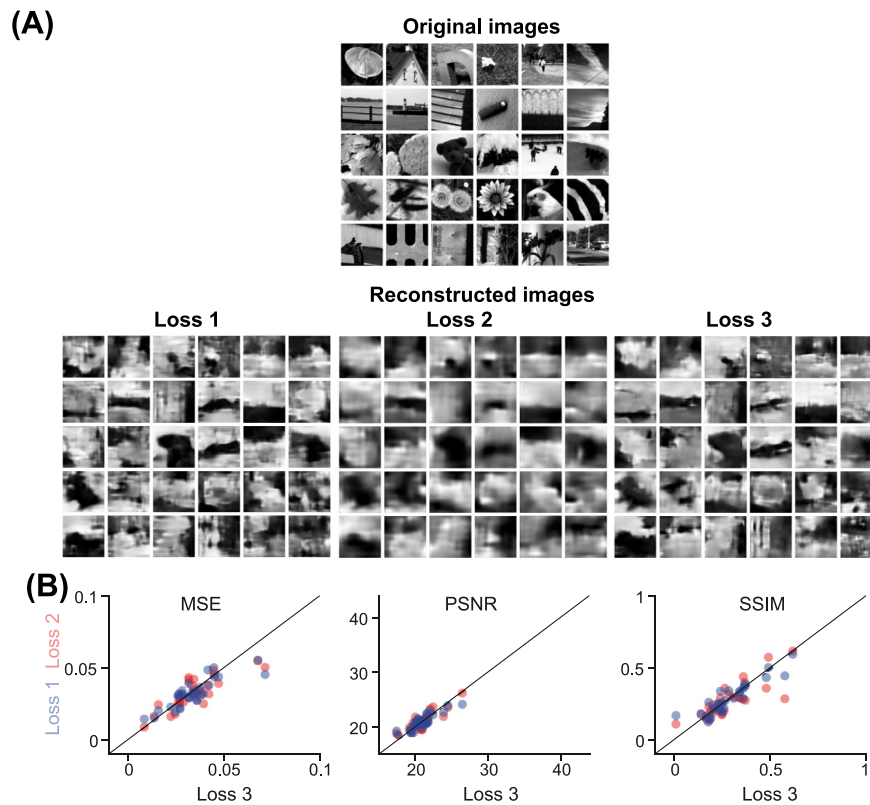
### 2.4. Experimental fMRI data

For comparison, we consider experimental fMRI data recorded in humans. The stimuli contain a set of 100 gray-scale handwritten digit images (equal number of 6's and 9's) extracted from the MNIST dataset, and the corresponding fMRI recordings were from V1-V3 areas of a single participant as detailed in previous studies (Van Gerven, Cseke, De Lange, & Heskes, 2010; van Gerven, de Lange, & Heskes, 2010). We evaluated this fMRI dataset by our SID directly with the same network structure. First, we convolved the stimulus images, here, 6's and 9's, with the same set of the receptive fields of RGCs used previously to simulated spikes, which is a population of 90 simulated cells, similar to the simulation of MNIST and CIFAR image datasets.

For fMRI data, each stimulus image triggered a response represented by a 3092-dimension vector as activity patterns. For a detailed comparison, we conducted two-fold cross-validation: decoding two sets of neural signals, fMRI and spikes, with two different models. We choose a recent state-of-the-art method termed deep generative multi-view model (DGMM) (Du, Du, Huang, & He, 2019), which is developed in the context of fMRI decoding. Thus, both models, DGMM and SID, are used for decoding with both fMRI data and our simulated spikes to see if there is an advantage of neural data type, as well as an advancement of decoding method. Within such a validation, one can observe decoding results in four conditions: decoding fMRI by DGMM (DGMM-fMRI), decoding spikes by DGMM (DGMM-Spike), decoding fMRI by SID (SID-fMRI), and decoding spike by SID (SID-Spike). In all cases, we used 90 images for training and 10 images for test.

## 3. Results

To test the capability of our SID model for the reconstruction of visual scenes, we first use our decoding model to reconstruct natural stimulus from biological experimental data that consists of a population of RGCs recorded in salamander retina triggered by natural image and video. To further test the generalization capability of our model, a set of simulated experiments with the encoding model were conducted, in which we reconstructed stimulus images of MNIST, CIFAR10, and CIFAR100 from simulated spikes. In addition, dynamic videos were also tested with simulated RGC spikes and with the SID model trained from CAFAR10. The results show that our SID can decode natural scenes from both biological and simulated spikes with very good precision.



**Fig. 2.** Reconstructed natural images from experimental RGC spikes with different loss functions. (A) Example original stimulus images in the test data (top) and final reconstructed images (bottom) from RGC spikes with three loss functions. (B) Reconstruction errors measured by MSE, PSNR, and SSIM are compared between Loss 3 and Loss 1 (blue) and Loss 2 (red). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

### 3.1. Decoding visual scenes from experimental data

Here the experimental data are a population of RGCs recorded from salamander retina (see Methods), where the stimuli include static natural images and a dynamic video, and the spike counts of the recorded RGCs are used as the input of our decoding model to reconstruct natural scenes.

#### 3.1.1. Decoding natural image stimulus

For natural images, there are 300 images as stimulus and 80 RGCs showing spike responses. For the decoder, the spike counts of 80 RGCs in response to each image are collected at one end, the other end is a training set of 270 images of  $64 \times 64$  pixels each. For each image, spike-image converter up-sample a population of 80 spike counts into an image with  $64 \times 64$  pixels to get an intermediate image, which then is mapped to the target training image with the image-image autoencoder.

Fig. 2(A) shows the decoding results on some example test images, where the global contents of images are reconstructed well while some details are missed. Depending on the loss functions used for training, intermediate images can be very different (Supplemental Figure 1), however, final reconstructions are similar to each other with some fine differences in the detailed textures. When there are two losses for each stage of the decoder as in  $L_1$  and  $L_2$ , the intermediate image (denoted as  $O_1$ ) is an preliminary image close to the target image. When there is one loss for both stages of the model as in  $L_3$ , the intermediate image is like noise without any visual structures.

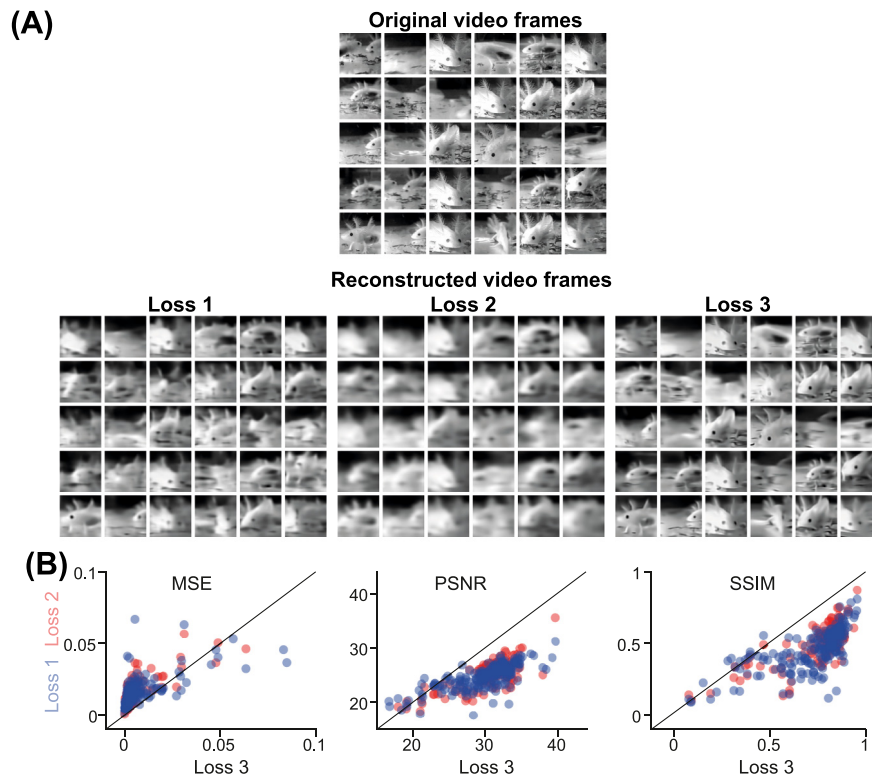
Fig. 2(B) shows the performance of our model characterized by three typical measures of reconstructed images: the mean square error (MSE) that describes the absolute difference of every pixel, the peak-signal-to-noise-ratio (PSNR) that characterizes the global quality, and the structural similarity index measure

(SSIM) that captures the details or image distortion, for evaluating the reconstruction results. Whereas there is a variation between different test images, in general, all three loss functions give similar performances in terms of three measures. However, these measures are not perfect characteristics. For example, the reconstructed images with  $L_2$  look as not good as the other two, however, it gives the similar measure errors. Therefore, these measures are a general guideline rather than the precise characteristic, as it seems that there is no perfect measure (Hore & Ziou, 2010; Wang, Simoncelli, & Bovik, 2003).

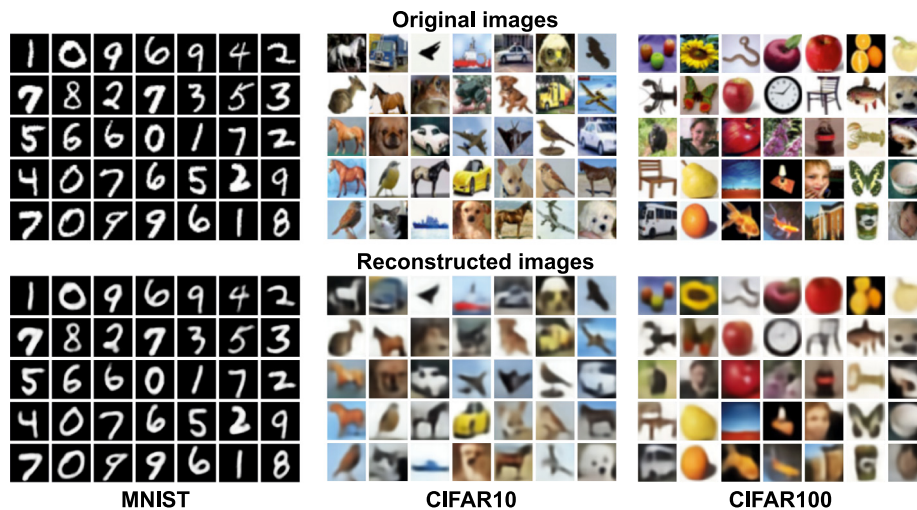
#### 3.1.2. Decoding dynamic video stimulus

So far, most of the decoding studies focus on static natural images, or some simple artificial dynamic scenes (Botella-Soler et al., 2018; Gollisch & Meister, 2008). To study dynamical visual scenes, such as videos with highly complex dynamics in both spatial and temporal domains, one has to, traditionally, understand how to deal with strong temporal adaptation observed in single neurons and neural circuits in the retina (Liu & Gollisch, 2015). Here with the video stimulus containing 1800 frames, and a population of 90 RGCs spike trains, our SID can overcome this difficulty by training of randomly selected frames out of the whole video, and testing on a subset untrained frames (see Methods).

From the reconstructed sample frames of test dataset shown in Fig. 3(A), one can see that our decoding model obtains high precise results, which can not only reconstruct the global content of each frame, but also some details of images that are missed in previous decoding studies (Botella-Soler et al., 2018; Nishimoto et al., 2011; Wen et al., 2018). The quantified performance is shown in Fig. 3(B). In contrast to the image results, all three measures, MSE, PSNR, and SSIM, indicate that the best results are obtained with  $L_3$ . Such a good performance can also be seen from the sample images where both global contents and fine details are obtained.



**Fig. 3.** Reconstructed video frames from experimental RGC spikes with different loss functions. Similar to Fig. 2, except that here are 30 random test video frames of a continuous video in (A), and a set of 50 random test frame images in (B), where each point is one of the test images.



**Fig. 4.** Reconstructed images from simulated RGC spikes in datasets of MNIST, CIFAR10, and CIFAR100. (Top) Example original images in test dataset. (Bottom) Reconstructed images from the spikes of 90 simulated RGCs.

### 3.2. Decoding visual scenes from simulated spikes

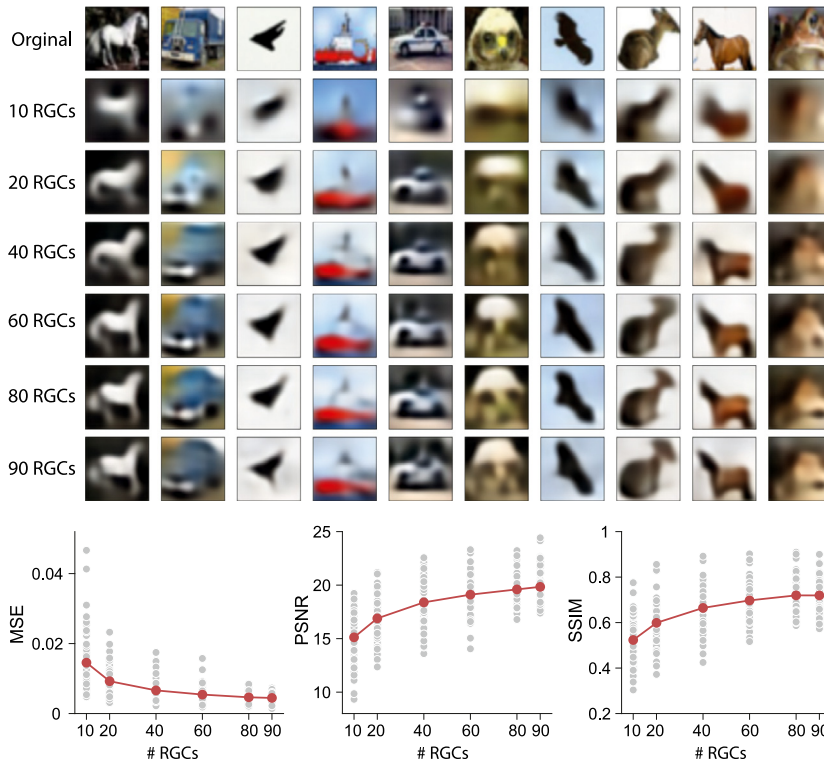
To further test the generalization capability of our decoding model, we did numerical experiments on simulated RGC data by creating a retinal encoding model based on the typical linear–nonlinear model and biological experimental data to simulate RGC responses (see Methods). In this way, one can simulate RGC spiking responses for any given stimuli. For comparison with experimental RGC data, we simulated the same number of 90 RGCs to obtain spikes.

We first did experiments on three popular image datasets, MNIST, CIFAR10, and CIFAR100. A population of neural spikes was obtained for each image with the encoding model. Then the

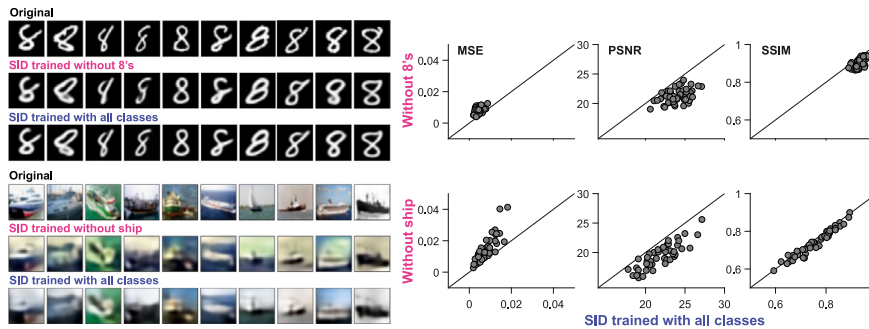
SID was trained for the training set, and tested on separate test data. The reconstructed example images are shown in Fig. 4 with good quality. In particular, the decoding performance for MNIST is almost perfect.

With the good reconstruction from simulated spikes with a population of only 90 RGCs, one can ask the question that how the decoding performance can be changed by a reduced number of RGCs. Thus, we re-trained the SID with a randomly selected subset of 90 RGCs as the input for the model, as shown in Fig. 5. Surprisingly, with only 20 RGC input, the SID can still recover a large part of the information from the original stimulus images. Indeed, the performance measured by MSE, PSNR, and SSIM is systematically decreasing when fewer RGCs are used for input.





**Fig. 5.** Reconstruction with a reduced number of RGC inputs. (Top) Reconstructed example images from CIFAR10 with different subsets of RGCs, from 10 to 90 RGCs. (Bottom) The performance of SID increased with more RGCs. Each gray point represents one of 50 randomly selected test images. Red points are the averaged measures over all test images.



**Fig. 6.** SID generalization to the untrained images of one category of MNIST and CIFAR. (Left, top) Original test 8's images from MNIST data, reconstructed example 8's images by the SID trained on other images but without the 8's category, and reconstructed example 8's images by the SID trained on images with all the categories. (Left, bottom) Similar results from CIFAR data, where the ship category is not used for the SID training. (Right) Comparison of generalization ability between the SID with and without one category of MNIST (top) and CIFAR (bottom). Each gray point represents one of 50 randomly selected test images.

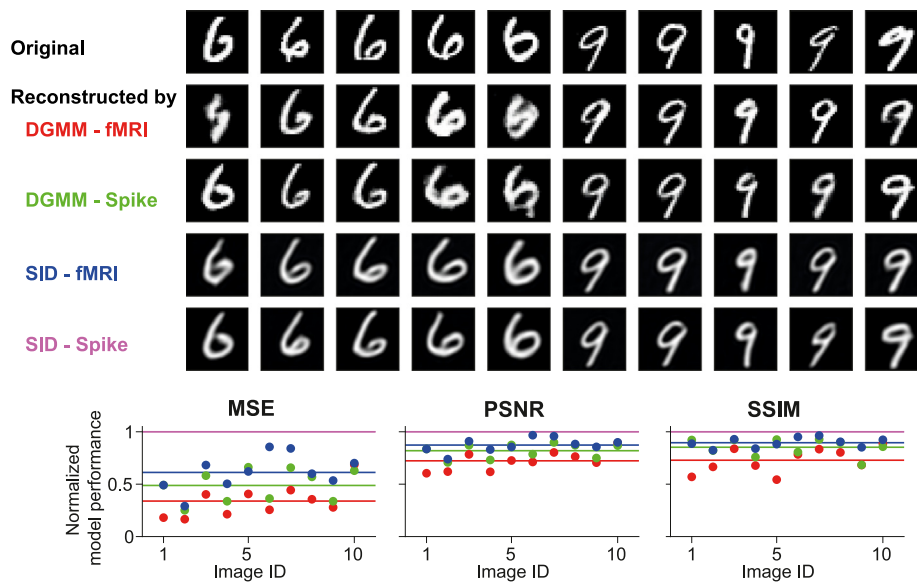
However, the performance tends to be convergent with 90 RGCs, which is the maximum number of RGCs recorded in our experimental data. It might be possible that the performance can be a little higher when more than 90 RGCs are used. We leave this possibility for the future study when more RGCs can be recorded (See Discussion).

As our decoder reconstructs the image per pixel directly, we expect that the training of SID can be done without any information about the category. Both MNIST and CIFAR datasets are originally used for image classification with labels of category. Thus, we re-trained our SID on these datasets but removing one of the known categories. Then we tested the reconstruction of the missing images from this untrained category. As shown in Fig. 6, the 8's digital images are reconstructed from the SID that is trained without 8's of digital images. In both training conditions, the sample size of the training set is the same, so that the only difference between the two is the missing category. We found

that the 8's can be decoded with a high accuracy, which means that the training of SID does not use any information about the category of the dataset. Similarly, the various ship images from the missing category during the training can also be reconstructed well on the more complex CIFAR dataset as in Fig. 6. However, the performance characterized by three measures of MSE, PSNR, and SSIM shows that there is a subtle difference between two versions of the test results: the SID trained with all classes has better performance, compared to the SID training without one category. Taken together, these results suggest that our SID has a good generalization ability, and the rich image content of the dataset can improve the decoding ability of the SID.

### 3.3. Cross-validation with existing methods

As there is no direct experimental data of neuronal spikes recorded for MNIST so far, we compared our results with the



**Fig. 7.** Reconstructed stimulus images from fMRI responses and simulated spikes. (Top) All of the 10 original images in the test dataset used for fMRI recordings, and the reconstructed images from fMRI data and the simulated spikes with the decoder of DGMM and SID cross-validated by a pair of model-data test, e.g. DGMM-fMRI stands for fMRI decoding by the DGMM model. For comparison, cross-validation is done by using DGMM for simulated spikes and our SID for fMRI data. (Bottom) Normalized model performance measured by MSE, PSNR, and SSIM. The performance of each pair of model-data is normalized by the SID-Spike. Each colored data point corresponds to each test image in the same order of image ID as in the original stimulus images on the top row. Colored solid lines represent the mean values of each error measure: DGMM-fMRI (red), DGMM-Spike (green), SID-fMRI (blue), and SID-Spike (purple). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

decoding results based on the models developed for fMRI signals. Previous studies evaluated the performance of quite a few decoders based on one particular dataset where a small part of MNIST images, 6's and 9's, were used as stimulus and the fMRI data was collected for human subject (Du et al., 2019; Fujiwara, Miyawaki, & Kamitani, 2013; Miyawaki et al., 2008; Qiao et al., 2018; Van Gerven et al., 2010; van Gerven et al., 2010; Wang, Arora, Livescu, & Bilmes, 2015; Wen et al., 2018). We compared our SID with a recent state-of-the-art method, termed Deep Generative Multi-view Model (DGMM), developed for decoding of fMRI data (Du et al., 2019). For detailed comparison, we used the same fMRI dataset but evaluated the decoding performance based on both fMRI responses and our simulated spikes, as well as two decoding models: DGMM and our SID, such that there are four pairs of model-data test: DGMM-fMRI (e.g. fMRI decoding by DGMM), DGMM-Spike, SID-fMRI, and SID-Spike.

The reconstructed example images are shown in Fig. 7, which can be characterized by MSE, PSNR and SSIM from a single run and a set of 10 runs with different initial conditions of the decoding models (Supplemental Figure 2). Interestingly, the SID gives a better performance for decoding of fMRI as well. Compared to DGMM, our SID outperforms on decoding with both fMRI signals and simulated spikes. Fig. 7 shows the normalized model performance for each individual test images, where all model-data pairs are normalized by the SID-Spike. In this way, SID-Spike gives a ratio as 1, and other model-data pairs have a performance smaller than 1. In particular, the performance increases in ascending order: DGMM-fMRI < DGMM-Spike < SID-fMRI < SID-Spike. These results suggest that our SID, as a general decoding method, shows a good capability for fMRI decoding as well.

### 3.4. Model generalization to real visual scenes

Next, we used the encoding model to simulated RGCs in response to dynamic videos. For comparison, the same set of videos used in Nishimoto et al. (2011) were encoded into RGC spikes. Then the SID model trained with CIFAR10 was used to reconstruct

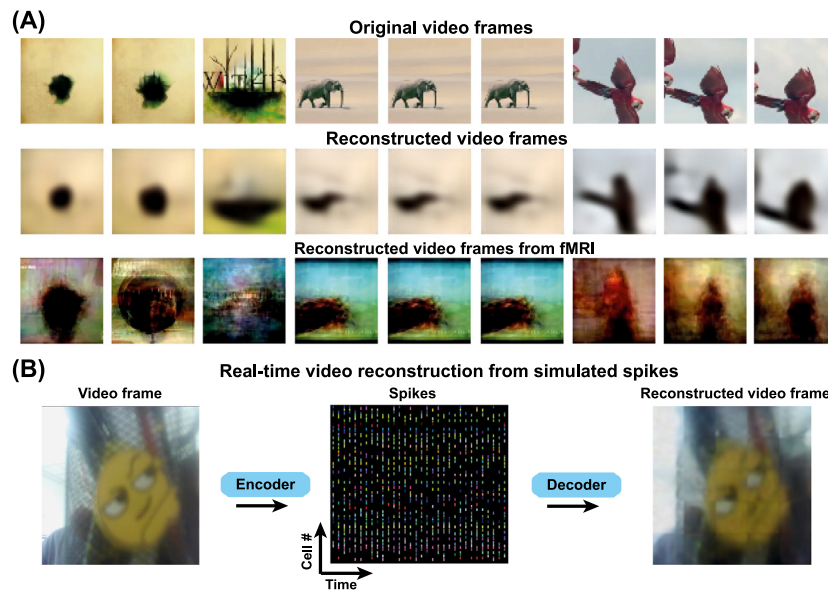
the videos from simulated RGC spikes. As shown in Fig. 8(A), the SID can reconstruct the dynamic videos from spikes very well. Compared with the fMRI decoded results (Nishimoto et al., 2011), our reconstructed videos contain more semantic contents and are much more refined. What is more, the SID model used here is trained with CIFAR10, which means that our SID model has a great generalization capability for arbitrary visual scenes. The complete videos reconstructed are available online (see Methods).

Finally, based on the results above, we implemented a real-time artificial vision system (see Methods), where real-time videos are captured by a standard camera, then encoded into a population of neural spikes, which are fed into the SID to reconstruct the visual scenes. Fig. 8(B) shows a snapshot of a real-time recorded video. Certainly, the speed of reconstruction and resolution of decoded video can be greatly improved with more advanced hardware. For this real-time system, the SID is also trained with the same CIFAR-10 dataset. These results suggest that our model can serve as a general framework to analyze those data recorded by artificial event-based vision systems (Gallego et al., 2019).

## 4. Discussion

In this study, we proposed an end-to-end spike-image decoder to reconstruct stimulus images from neural spikes based on the retinal ganglion cells. The capability of SID was tested on experimental RGC data for both static and dynamic natural scenes to obtain state-of-the-art reconstruction performance. Furthermore, by using an additional encoding model, the great performance of reconstruction of arbitrary visual scenes was demonstrated on popular image datasets. In addition, our SID can also be used for decoding fMRI data and show very good performance compared to other methods developed in the context of fMRI decoding. Finally, we examined the generalization ability of SID by using a pre-trained SID to decode various types of dynamic videos to achieve real-time encoding and decoding visual scenes by neural spikes.





**Fig. 8.** SID can be generalized to reconstruct arbitrary videos with a pre-trained model on CIFAR10 without any fine tuning. (A) Three segments of original videos (top), reconstructed videos from simulated spikes with a SID model pre-trained on CIFAR10 data (middle), and reconstructed videos based on fMRI signals from Nishimoto et al. (2011) (bottom). (B) Decoding of real-time videos. (Left) One example frame of a real-time video. (Middle) A population of spike trains obtained by the encoder. Each line represents one RGC's spike train. (Right) The same frame reconstructed in a real-time fashion by the SID. Real-time videos are available to view online.

#### 4.1. Rate code vs. temporal code

In our current study, the firing rate averaged over a number of trials for experimental data was used for our decoder and reconstruction. Both firing rate code and temporal code have been used in the field of neuronal coding (Gollisch & Meister, 2008; Jacobs et al., 2009; Onken et al., 2016). In practice, for analyzing experimental data, the difference between these two depends on many factors. In our case presented here, experimental data were collected under a condition that a 30 Hz frame rate was used for video stimulation, therefore, to match the sample rate of video frames, we binned the spike trains into 30 Hz as well. For image stimulation, each individual image was presented for 200 ms and then followed by an 800 ms empty display, therefore, for each image, we used the number of spikes during 300 ms after the onset of image presentation. We also conducted the reconstruction by binarizing the firing rate of RGCs into a single spike train with 0's and 1's only, and the results are similar to the current presentation.

From the experimental side, one could conduct new experiments with a high frame rate of video sampling, for instance, 120 Hz can be used, so that one only needs to consider a smaller bin,  $1000/120 \approx 8.3$  ms, in which more precise spike times can be recovered given that the sparse firing of the retinal ganglion cells in general cannot generate 2 or more spikes within 8.3 ms in salamander, the animal species used here. On the other hand, these neurons in this animal species have a very low firing rate such that within a bin of 33 ms used in our study, there are 93.12% of zero spike, 5.37% of single spike, 1.14% of double spikes, and 0.37% of triple or more spikes. Therefore, although the firing rate was used for our decoding, the low firing rate in neurons is an approximation of temporal coding of spike times.

Given the experimental issues listed above, that is the reason why we use simulated spikes to demonstrate our method. In our simulations of image datasets and real-time videos, the decoding is based on the temporal code with the exact spikes in a single trial. Thus, our results presented here are provided as a proof-of-concept for using a population of single-trial spikes to reconstruct visual scenes.

#### 4.2. Implication for other neuronal systems

Our method could serve as a general model of neural coding, similar to other neural coding models, such as the general linear model (Pillow et al., 2008), to study the relationship between neuronal signal and stimulus. In our case, the SID is a non-linear decoder to reconstruct image stimulus from neural spikes, which can be seen as an inverse function of spike-image mapping. Thus our method could be generalized to other parts of the brain as long as the neurons deliver different spike patterns for different stimuli.

The retina has been thought of as a relatively simple neuronal circuit, however, it has been suggested that many complex computations can be realized by their neurons, in particular, the RGCs (Gollisch & Meister, 2010). It has been studied intensively to develop various types of models for the encoding of visual scenes in the retina (Chichilnisky, 2001; Pillow et al., 2008; Yu et al., 2020). Similar to our current study, other recent work also focuses on neuronal spikes in the retina to develop the novel decoding methods (Botella-Soler et al., 2018; Marre et al., 2015; Parthasarathy et al., 2017). All the methods used in these studies take a similar approach by using the advantage of the recently advanced machine learning technique. Furthermore, the similar CNN decoding approach has been used for other parts of the visual system and different types of neural signals, such as calcium imaging in animal visual cortex (Garasto et al., 2018), and classical fMRI signal in human visual cortex (Wen et al., 2018). Therefore, our method could be applied to these types of neural signals from other parts of the visual system.

In the current study, we provide a case for using the SID to the fMRI signal. It is easy to expect that the decoding performance is higher for neural spikes compared to the fMRI signal due to the nature of neural signal, where spikes are much more refined than fMRI. The two-fold cross-validation shown in Fig. 7 confirm this expectation. With the same model, either SID or DGMM, the decoding results by using neural spikes are better than those by fMRI. However, interestingly, we also observe a methodological advantage by comparing two decoding methods: SID and DGMM.

It should be noted that DGMM is a recent and state-of-the-art method developed in the context of fMRI decoding (Du et al., 2019). Even in this condition, our SID can achieve a better decoding performance than DGMM by using the same fMRI data. Therefore, our method can be seen as a more general model to cope with different types of neural signals beyond neural spikes.

Recent experimental advancements in neuroscience can collect a large population of neurons simultaneously. In particular, in the retina, a population of spike trains from hundreds of retinal ganglion cells can be obtained with well-controlled visual scenes, such as images and movies (Botella-Soler et al., 2018; Onken et al., 2016). New techniques may record several thousands of neurons simultaneously (Hilgen et al., 2017; Maccione et al., 2014; Portelli et al., 2016). Together with other recent studies (Botella-Soler et al., 2018; Parthasarathy et al., 2017), our results in this work show that a better decoding performance can be archived with neural spikes, which is beyond the scope of fMRI signal (Du et al., 2019; Nishimoto et al., 2011; Wen et al., 2018).

#### 4.3. Model limitation and further study

Our decoder can be considered as a general framework, in which the detailed network structures are flexible. Here we only explored the simple and perhaps the most basic structure, e.g., a simple dense connected and layered network for the spike-image converter, similar to a multilayer perceptron, and a typical network for the image-image autoencoder. One of the advantages of the current framework is that one can use any other types of spike-image converter, which is similar to existing models of decoding by transferring spikes into image pixels (see (Botella-Soler et al., 2018)), and any other types of image-image encoder, which is similar to existing models of image analysis (such as generative adversarial network or GAN (Goodfellow et al., 2014)) in the deep learning field. Given the rapid development of architecture design for artificial, spiking, and a mixture of both (Xu et al., 2018), neural networks, it is possible to use other network architectures to further improve the decoding ability.

Our SID does not include the temporal filter, but only the spatial filter, i.e., the spatial receptive field of RGCs. RGCs have both spatial and temporal filters to take into account of spatial structure (Liu et al., 2017) and temporal dynamics of stimulus (Liu & Gollisch, 2015). Nonlinear structure in space has been intensively studied in the retina (Liu et al., 2017; Schwartz et al., 2012; Schwartz & Rieke, 2011). In contrast, the temporal dynamics are more complex and entangled with the stimulus itself. Even for the simple artificial stimulus, such as white noise, there is a strong neuronal adaption although there is no correlation within white noise stimulus (Liu & Gollisch, 2015). For video stimulus with natural scenes, neuronal adaption entangled with correlation in the video in the temporal domain makes the investigation of neuronal coding of natural scene extremely difficult (Heitman et al., 2016).

Here we overcome this difficulty by using a training set of randomly selected frames from the video in experimental data. In this way, one can break the temporal correlation between video frames. Essentially, this manipulation resembles presenting a large set of natural images sequentially, termed as naturalistic video, such that the temporal correlation between frames can be reduced to a minimal level (Heitman et al., 2016). The limited recording time during neuroscience experiments restricts the stimulus to a short segment of video used for RGCs. Thus, we tested the decoding capability of our method by using the stimulation of a large set of images, CIFAR, to mimic the full pixel space encountered by the model. To our surprise, our model trained by CIFAR can be generalized to any dynamic natural scenes without any re-training or fine-tuning.

Dynamic visual scenes are highly complex with the information presented in a spatiotemporal fashion and high-order correlation (Simoncelli & Olshausen, 2001). Recent advancements of computer vision make some breakthroughs for analyzing these complex natural scenes (Lecun, Bengio, & Hinton, 2015), for instance, the PredNet model can decode the next frame based on the previous frame of video (Lotter, Kreiman, & Cox, 2017). However, the efficiency, generalization ability, and adaption or transfer learning between different tasks, of well-trained models are still far from human performance (Marblestone, Wayne, & Kording, 2016). Instead, our modeling framework shows a great generalization ability for the decoding of visual scenes with a pre-trained model. In future work, we expect to develop a time-based SID and use temporal filters for decoding.

#### 4.4. Implication for artificial visual systems

The proposed model framework could also be used for visual neuroprostheses and other bioengineering devices of brain-machine interface. For the retinal neuroprostheses, the benefit of a decoding model is to justify the spiking patterns produced by the targeted downstream neurons, such that electrical stimulation should be able to close to those desired patterns of retinal neural activity in a prosthesis (Nirenberg & Pandarinath, 2012). Such control is important and beyond the traditional way of computing the distance between two or more spike trains in general (Rossum, 2001; Shah, Madugula, Chichilnisky, Shlens, & Singer, 2017; Victor, 2005), as recent studies suggest that the performance of neuroprosthesis for restoring vision can be improved by the encoding/decoding algorithms besides of hardware design (Nirenberg & Pandarinath, 2012; Yan & Nirenberg, 2018; Yue, Weiland, Roska, & Humayun, 2016). Our study here proposes a novel decoding framework, as well as provides a proof-of-principle of encoding and decoding visual scenes with only a small population of neurons. Given the limitation of physical devices, such as only a small patch of brain area can be stimulated at one time, our framework suggests a feasible approach by using a small set of neurons for high precision decoding, which may systematically improve the capability of brain-machine interface.

The model framework proposed in this study could be used for other artificial visual systems. The main feature of our framework is to make use of neural spikes. Advancements of recent artificial intelligence computing align with the development of the next generation of neuromorphic chips and devices, where the new data format is processed as spikes or events (Davies et al., 2018; Esser et al., 2016; Gallego et al., 2019; Ma et al., 2017; Pei et al., 2019). Therefore, our method can be applied to neuromorphic retina cameras with spiking or events, such as the dynamic vision sensor (Lichtsteiner, Posch, & Delbruck, 2008). Taken together with neuromorphic hardware and event/spiking computing algorithm, the next generation of computational vision can develop a better system for artificial vision.

#### Acknowledgments

We thank W. Fang and Q. Xu for discussions, and the NVIDIA NVAIL project for providing hardware. This work is supported by the National Basic Research Program of China [2015CB351806], the National Natural Science Foundation of China [61806011, 61825101, 61425025, 61961130392 and U1611461], the National Postdoctoral Program for Innovative Talents, China [BX20180005], the China Postdoctoral Science Foundation [2018M630036], the Zhejiang Lab, China [2019KCOAB03 and 2019KCOAD02], and the Royal Society Newton Advanced Fellowship, UK [NAF-R1-191082].

## Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.neunet.2020.01.033>.

## References

- Botella-Soler, V., Deny, S., Martius, G., Marre, O., & Tkačik, G. (2018). Non-linear decoding of a complex movie from the mammalian retina. *PLoS Computational Biology*, *14*(5), e1006057.
- Chichilnisky, E. J. (2001). A simple white noise analysis of neuronal light responses. *Network*, *12*(2), 199–213.
- Davies, M., Srinivasa, N., Lin, T.-H., Chinya, G., Cao, Y., Choday, S. H., et al. (2018). Loihi: A neuromorphic manycore processor with on-chip learning. *IEEE Micro*, *38*(1), 82–99.
- Demb, J. B., & Singer, J. H. (2015). Functional circuitry of the retina. *Annual Review of Vision Science*, *1*, 263–289.
- Du, C., Du, C., Huang, L., & He, H. (2019). Reconstructing perceived images from human brain activities with bayesian deep multiview learning. *IEEE Transactions on Neural Networks and Learning Systems*, *30*(8), 2310–2323.
- Esser, S., Merolla, P., Arthur, J., Cassidy, A., Appuswamy, R., Andreopoulos, A., et al. (2016). Convolutional networks for fast, energy-efficient neuromorphic computing. *Proceedings of the National Academy of Sciences*, *113*(41), 11441–11446.
- Fujiwara, Y., Miyawaki, Y., & Kamitani, Y. (2013). Modular encoding and decoding models derived from bayesian canonical correlation analysis. *Neural Computation*, *25*(4), 979–1005.
- Gallego, G., Delbruck, T., Orchard, G., Bartolozzi, C., Taba, B., Censi, A., et al. (2019). Event-based vision: A survey. arXiv preprint arXiv:1904.08405.
- Garasto, S., Bharath, A. A., & Schultz, S. R. (2018). Visual reconstruction from 2-photon calcium imaging suggests linear readout properties of neurons in mouse primary visual cortex. (p. 300392). bioRxiv.
- Gilja, V., Nuyujukian, P., Chestek, C. A., Cunningham, J. P., Byron, M. Y., Fan, J. M., et al. (2012). A high-performance neural prosthesis enabled by control algorithm design. *Nature Neuroscience*, *15*(12), 1752.
- Gollisch, T., & Meister, M. (2008). Rapid neural coding in the retina with relative spike latencies. *Science*, *319*(5866), 1108–1111.
- Gollisch, T., & Meister, M. (2010). Eye smarter than scientists believed: Neural computations in circuits of the retina. *Neuron*, *65*(2), 150–164.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., et al. (2014). Generative adversarial networks. *Advances in Neural Information Processing Systems*, *3*.
- Grimes, W. N., Songco-Aguas, A., & Rieke, F. (2018). Parallel processing of rod and cone signals: Retinal function and human perception. *Annual Review of Vision Science*, *4*, 123–141.
- Heitman, A., Brackbill, N., Greschner, M., Sher, A., Litke, A. M., & Chichilnisky, E. (2016). Testing pseudo-linear models of responses to natural scenes in primate retina. bioRxiv.
- Hilgen, G., Pirmoradian, S., Pamplona, D., Kornprobst, P., Cessac, B., Hennig, M. H., et al. (2017). Pan-retinal characterisation of light responses from ganglion cells in the developing mouse retina. *Scientific Reports*, *7*, 42330.
- Hinton, G. E., Srivastava, N., Krizhevsky, A., Sutskever, I., & Salakhutdinov, R. R. (2012). Improving neural networks by preventing co-adaptation of feature detectors. arXiv preprint arXiv:1207.0580.
- Hore, A., & Ziou, D. (2010). Image quality metrics: PSNR vs. SSIM. In *2010 20th international conference on pattern recognition*. IEEE.
- Ioffe, S., & Szegedy, C. (2015). Batch normalization: Accelerating deep network training by reducing internal covariate shift. arXiv preprint arXiv:1502.03167.
- Jacobs, A. L., Fridman, G., Douglas, R. M., Alam, N. M., Latham, P. E., Prusky, G. T., et al. (2009). Ruling out and ruling in neural codes. *Proceedings of the National Academy of Sciences of the United States of America*.
- Jadzinsky, P. D., & Baccus, S. A. (2013). Transformation of visual signals by inhibitory interneurons in retinal circuits. *Annual Review of Neuroscience*, *36*, 403–428.
- Krizhevsky, A. (2012). *Learning multiple layers of features from tiny images*. University of Toronto.
- Lecun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, *521*(7553), 436–444.
- Lecun, Y., Bottou, L., Bengio, Y., & Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, *86*(11), 2278–2324.
- Lichtsteiner, P., Posch, C., & Delbruck, T. (2008). A 128\*128 120 db 15 us latency asynchronous temporal contrast vision sensor. *IEEE Journal of Solid-State Circuits*, *43*(2), 566–576.
- Liu, J. K., & Gollisch, T. (2015). Spike-triggered covariance analysis reveals phenomenological diversity of contrast adaptation in the retina. *PLoS Computational Biology*, *11*(7), e1004425.
- Liu, J. K., Schreyer, H. M., Onken, A., Rozenblit, F., Khani, M. H., Krishnamoorthy, V., et al. (2017). Inference of neuronal functional circuitry with spike-triggered non-negative matrix factorization. *Nature Communications*, *8*(1), 149.
- Lotter, W., Kreiman, G., & Cox, D. (2017). Deep predictive coding networks for video prediction and unsupervised learning. arXiv:1605.08104v5.
- Ma, D., Shen, J., Gu, Z., Zhang, M., Zhu, X., Xu, X., et al. (2017). Darwin: A neuromorphic hardware co-processor based on spiking neural networks. *Journal of Systems Architecture*, *77*, 43–51.
- Maccione, A., Hennig, M. H., Gandolfo, M., Muthmann, O., Copenhagen, J., Eglén, S. J., et al. (2014). Following the ontogeny of retinal waves: pan-retinal recordings of population dynamics in the neonatal mouse. *The Journal of Physiology*, *592*(7), 1545–1563.
- Marblestone, A. H., Wayne, G., & Kording, K. P. (2016). Toward an integration of deep learning and neuroscience. *Frontiers in Computational Neuroscience*, *10*, 94.
- Marre, O., Botella-Soler, V., Simmons, K. D., Mora, T., Tkačik, G., & Berry II, M. J. (2015). High accuracy decoding of dynamical motion from a large retinal population. *PLoS Computational Biology*, *11*(7), e1004304.
- Meyer, A. F., Williamson, R. S., Linden, J. F., & Sahani, M. (2017). Models of neuronal stimulus-response functions: elaboration, estimation, and evaluation. *Frontiers in Systems Neuroscience*, *10*, 109.
- Miyawaki, Y., Uchida, H., Yamashita, O., aki Sato, M., Morito, Y., Tanabe, H. C., et al. (2008). Visual image reconstruction from human brain activity using a combination of multiscale local image decoders. *Neuron*, *60*, 915–929.
- Naselaris, T., Prenger, R. J., Kay, K. N., Oliver, M., & Gallant, J. L. (2009). Bayesian reconstruction of natural images from human brain activity. *Neuron*, *63*(6), 902–915.
- Nirenberg, S., & Pandarinath, C. (2012). Retinal prosthetic strategy with the capacity to restore normal vision. *Proceedings of the National Academy of Sciences*, *109*(37), 15012–15017.
- Nishimoto, S., Vu, A. T., Naselaris, T., Benjamini, Y., Yu, B., & Gallant, J. L. (2011). Reconstructing visual experiences from brain activity evoked by natural movies. *Current Biology*, *21*(19), 1641–1646.
- O'Brien, J., & Bloomfield, S. A. (2018). Plasticity of retinal gap junctions: Roles in synaptic physiology and disease. *Annual Review of Vision Science*, *4*, 79–100.
- Olmos, A., & Kingdom, F. A. A. (2004). A biologically inspired algorithm for the recovery of shading and reflectance images. *Perception*, *33*(12), 1463–1473.
- Onken, A., Liu, J. K., Karunasekara, P. C. R., Delis, I., Gollisch, T., & Panzeri, S. (2016). Using matrix and tensor factorizations for the single-trial analysis of population spike trains. *PLoS Computational Biology*, *12*(11), e1005189.
- Parthasarathy, N., Batty, E., Falcon, W., Rutten, T., Rajpal, M., Chichilnisky, E., et al. (2017). Neural networks for efficient bayesian decoding of natural images from retinal neurons. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, R. Garnett (Eds.), *Advances in neural information processing systems*, Vol. 30 (pp. 6437–6448). Curran Associates, Inc.
- Pei, J., Deng, L., Song, S., Zhao, M., Zhang, Y., Wu, S., et al. (2019). Towards artificial general intelligence with hybrid tianjic chip architecture. *Nature*, *572*(7767), 106–111.
- Pillow, J. W., Shlens, J., Paninski, L., Sher, A., Litke, A. M., Chichilnisky, E. J., et al. (2008). Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature*, *454*(7207), 995.
- Portelli, G., Barrett, J. M., Hilgen, G., Masquelier, T., Maccione, A., Di Marco, S., et al. (2016). Rank order coding: a retinal information decoding strategy revealed by large-scale multielectrode array retinal recordings. *Eneuro*, *ENEURO-0134*.
- Qiao, K., Zhang, C., Wang, L., Chen, J., Zeng, L., Tong, L., et al. (2018). Accurate reconstruction of image stimuli from human functional magnetic resonance imaging based on the decoding model with capsule network architecture. *Frontiers in Neuroinformatics*, *12*, 62.
- Rieke, F., Warland, D., de Ruyter van Steveninck, R., & Bialek, W. (1996). *Spikes: Exploring the neural code (computational neuroscience)*. The MIT Press.
- Rivlin-Etzion, M., Grimes, W. N., & Rieke, F. (2018). Flexible neural hardware supports dynamic computations in retina. *Trends in Neurosciences*, *41*, 224–237.
- Rossum, M. v. (2001). A novel spike distance. *Neural Computation*, *13*(4), 751–763.
- Schwartz, G. W., Okawa, H., Dunn, F. A., Morgan, J. L., Kerschensteiner, D., Wong, R. O., et al. (2012). The spatial structure of a nonlinear receptive field. *Nature Neuroscience*.
- Schwartz, G., & Rieke, F. (2011). Perspectives on: information and coding in mammalian sensory physiology: nonlinear spatial encoding by retinal ganglion cells: when 1 + 1 not equal 2. *The Journal of General Physiology*, *138*(3), 283–290.
- Shah, N. P., Madugula, S., Chichilnisky, E., Shlens, J., & Singer, Y. (2017). Learning a neural response metric for retinal prosthesis. (p. 226530). bioRxiv.
- Simoncelli, E. P., & Olshausen, B. A. (2001). Natural image statistics and neural representation. *Annual Review of Neuroscience*, *24*(24), 1193.
- Stanley, G. B., Li, F. F., & Dan, Y. (1999). Reconstruction of natural scenes from ensemble responses in the lateral geniculate nucleus. *Journal of Neurophysiology*, *19*(18), 8036–8042.



- Thirion, B., Duchesnay, E., Hubbard, E., Dubois, J., Poline, J.-B., LeBihan, D., et al. (2006). Inverse retinotopy: inferring the visual content of images from brain activation patterns. *Neuroimage*, 33(4), 1104–1116.
- Van Gerven, M. A., Cseke, B., De Lange, F. P., & Heskes, T. (2010). Efficient bayesian multivariate fmri analysis using a sparsifying spatio-temporal prior. *NeuroImage*, 50(1), 150–161.
- van Gerven, M. A., de Lange, F. P., & Heskes, T. (2010). Neural decoding with hierarchical generative models. *Neural Computation*, 22(12), 3127–3142.
- Victor, J. D. (2005). Spike train metrics. *Current Opinion in Neurobiology*, 15(5), 585–592.
- Wang, W., Arora, R., Livescu, K., & Bilmes, J. (2015). On deep multi-view representation learning. In *International conference on machine learning* (pp. 1083–1092).
- Wang, Z., Simoncelli, E., & Bovik, A. (2003). Multiscale structural similarity for image quality assessment. In *The thirty-seventh asilomar conference on signals, systems & computers*. IEEE.
- Wen, H., Shi, J., Zhang, Y., Lu, K.-H., Cao, J., & Liu, Z. (2018). Neural encoding and decoding with deep learning for dynamic natural vision. *Cerebral Cortex*, 28(12), 4136–4160.
- Wu, M. C.-K., David, S. V., & Gallant, J. L. (2006). Complete functional characterization of sensory neurons by system identification. *Annual Review of Neuroscience*, 29, 477–505.
- Xu, Q., Qi, Y., Yu, H., Shen, J., Tang, H., & Pan, G. (2018). Csnns: An augmented spiking based framework with perceptron-inception. In *IJCAI* (pp. 1646–1652).
- Yamins, D. L. K., & Dicarlo, J. J. (2016). Using goal-driven deep learning models to understand sensory cortex. *Nature Neuroscience*, 19(3), 356.
- Yan, B., & Nirenberg, S. (2018). An embedded real-time processing platform for optogenetic neuroprosthetic applications. *IEEE Transactions on Neural Systems and Rehabilitation Engineering: a Publication of the IEEE Engineering in Medicine and Biology Society*, 26(1), 233–243.
- Yan, Q., Yu, Z., Chen, F., & Liu, J. K. (2017). Revealing structure components of the retina by deep learning networks. In *NIPS symposium on interpretable machine learning*. arXiv:1711.02837.
- Yan, Q., Zheng, Y., Jia, S., Zhang, Y., Yu, Z., Chen, F., et al. (2018). Revealing fine structures of the retinal receptive field by deep learning networks. arXiv preprint arXiv:1811.0229.
- Yu, Z., Liu, J. K., Jia, S., Zhang, Y., Zheng, Y., Tian, Y., et al. (2020). Towards the next generation of retinal neuroprosthesis: Visual computation with spikes. arXiv preprint arXiv:2001.04064.
- Yue, Y., He, L., He, G., Liu, J., Du, K., Tian, Y., et al. (2018). A simple blind-denoising filter inspired by electrically coupled photoreceptors in the retina. arXiv preprint arXiv:1806.05882.
- Yue, L., Weiland, J. D., Roska, B., & Humayun, M. S. (2016). Retinal stimulation strategies to restore vision: Fundamentals and systems. *Progress in Retinal and Eye Research*, 53, 21–47.
- Zeiler, M. D., & Fergus, R. (2014). Visualizing and understanding convolutional networks. In *European conference on computer vision* (pp. 818–833).