



Learning to Recognize Actions From Limited Training Examples Using a Recurrent Spiking Neural Model

Priyadarshini Panda^{1*} and Narayan Srinivasa^{2†}

¹ Nanoelectronics Research Laboratory, School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN, United States, ² Intel Labs, Hillsboro, OR, United States

OPEN ACCESS

Edited by:

Arindam Basu,
Nanyang Technological University,
Singapore

Reviewed by:

Subhrajit Roy,
IBM Research, Australia
Chetan Singh Thakur,
Indian Institute of Science, India
Yulia Sandamirskaya,
University of Zurich, Switzerland

*Correspondence:

Priyadarshini Panda
pandap@purdue.edu
Narayan Srinivasa
nara@etacompute.com

† Present Address:

Narayan Srinivasa,
Independent Researcher, Westlake
Village, CA, United States

Specialty section:

This article was submitted to
Neuromorphic Engineering,
a section of the journal
Frontiers in Neuroscience

Received: 31 October 2017

Accepted: 16 February 2018

Published: 02 March 2018

Citation:

Panda P and Srinivasa N (2018)
Learning to Recognize Actions From
Limited Training Examples Using a
Recurrent Spiking Neural Model.
Front. Neurosci. 12:126.
doi: 10.3389/fnins.2018.00126

A fundamental challenge in machine learning today is to build a model that can learn from few examples. Here, we describe a reservoir based spiking neural model for learning to recognize actions with a limited number of labeled videos. First, we propose a novel encoding, inspired by how microsaccades influence visual perception, to extract spike information from raw video data while preserving the temporal correlation across different frames. Using this encoding, we show that the reservoir generalizes its rich dynamical activity toward signature action/movements enabling it to learn from few training examples. We evaluate our approach on the UCF-101 dataset. Our experiments demonstrate that our proposed reservoir achieves 81.3/87% Top-1/Top-5 accuracy, respectively, on the 101-class data while requiring just 8 video examples per class for training. Our results establish a new benchmark for action recognition from limited video examples for spiking neural models while yielding competitive accuracy with respect to state-of-the-art non-spiking neural models.

Keywords: reservoir model, driven-autonomous construction, action recognition, limited training example, supervised plasticity, micro-saccade spike encoding, eigenvalue spectra

1. INTRODUCTION

The exponential increase in digital data with online media, surveillance cameras among others, creates a growing need to develop intelligent models for complex spatio-temporal processing. Recent efforts in deep learning and computer vision have focused on building systems that learn and think like humans (Mnih et al., 2013; LeCun et al., 2015). Despite the biological inspiration and remarkable performance of such models, even beating humans in certain cognitive tasks (Silver et al., 2016), the gap between humans and artificially engineered intelligent systems is still great. One important divide is the size of the required training datasets. Studies on mammalian concept understanding have shown that humans and animals can rapidly learn complex visual concepts from single training examples (Lake et al., 2011). In contrast, the state-of-the-art intelligent models require vast quantities of labeled data with extensive, iterative training to learn suitably and yield high performance. While the proliferation of digital media has led to the availability of massive raw and unstructured data, it is often impractical and expensive to gather annotated training datasets for all of them. This motivates our work on learning to recognize from few labeled data. Specifically, we propose a reservoir based spiking neural model that reliably learns to recognize actions in video data by extracting long-range structure from a small number of training examples.

Reservoir or Liquid Computing approaches have shown surprising success in recent years on a variety of temporal data based recognition problems (though effective vision applications are still scarce) (Lukoševičius and Jaeger, 2009; Maass, 2011; Srinivasa and Cho, 2014). This can be attributed to populations of recurrently connected spiking neurons, similar to the mammalian neocortex anatomy (Wehr and Zador, 2003), which create a high-dimensional dynamic representation of an input stream. The advantage with such an approach is that the reservoir (with sparse/random internal connectivity) implicitly encodes the temporal information in the input and provides a unique non-linear description for each input sequence, that can be trained for read out by a set of linear output neurons. In fact, in our view, the non-linear integration of input data by the high-dimensional dynamical reservoir results in generic stable internal states that generalize over common aspects of the input thereby allowing rapid learning from limited examples. While the random recurrent connectivity generates favorable complex dynamic activity within a reservoir, it poses unique challenges that are generally not encountered in constructing feedforward/deep learning networks.

A pressing problem with recurrently connected networks is to determine the connection matrix that would be suitable for making the reservoir perform well on a particular task, because it is not entirely obvious how individual neurons should spike while interacting with other neurons. To address this, Abbott et al. (2016) proposed an elegant method for constructing recurrent spiking networks, termed as Autonomous (A) models, from continuous variable (rate) networks, called Driven (D) models. The basic idea in the D/A approach is to construct a network (Auto) that performs a particular task by copying another network (Driven) statistics that does the same task. The copying provides an estimate of the internal dynamics for the autonomous network including currents at individual neurons in the reservoir to produce the same outputs. Here, we adopt the D/A based reservoir construction approach and modify it further by introducing approximate delay apportioning to develop a recurrent spiking model for practical action recognition from limited video data. Please note, we use Auto/Autonomous interchangeably in the remainder of the text.

Further, in order to facilitate limited example training in a reservoir spiking framework, we propose a novel spike based encoding/pre-processing method to convert the raw pixel valued videos in the dataset into spiking information that preserves the temporal statistics and correlation across different frames. This approach is inspired by how MicroSaccades (MS) cause retinal ganglion cells to fire synchronously corresponding to contrast edges after the MS. The emitted synchronous spike volley thus rapidly transmits the most salient edges of the stimulus, which often constitute the most crucial information (Masquelier et al., 2016). In addition, our encoding eliminates irrelevant spiking information due to ambiguity such as, noisy background activity or jitter (due to unsteady camera movement), further making our approach more robust and reliable. Hence, our encoding, in general, captures the signature action or movement of a subject across different videos as spiking information. This, in turn, enables the reservoir to recognize/generalize over motion

cues from spiking data, to enable learning various types of actions from few video samples per class. Our proposed spiking model is much more adept at learning object/activity types, due in part to our ability to analyze activity dynamically as it takes place over time. Furthermore, besides the advantage of learning with limited training examples, the spiking foundation of the proposed model provides an energy-efficient alternative for neuromorphic hardware implementation than conventional machine learning models (Merolla et al., 2014; Han et al., 2015).

It is worth mentioning that there has been substantial effort toward implementing a neural model with human-level concept learning capability (Fei-Fei et al., 2007; Lake et al., 2015; Santoro et al., 2016; Vinyals et al., 2016; Burgess et al., 2017). For instance, Burgess et al. (2017), Lake et al. (2015), and Fei-Fei et al. (2007) use Bayesian updates on a pre-trained network to classify a new image category with limited instances. Santoro et al. (2016) use relatively complex architectures such as memory-augmented neural Turing machines to perform one-shot learning. Recently, Google proposed a combined non-parametric/parametric matching network for one-shot learning demonstrating state-of-the-art accuracy on ImageNet dataset (Vinyals et al., 2016). While our paper complements these works, our spiking model approach to few-shot learning is first of its kind. While most of the previous works have shown the efficacy of one-shot learning on static image data, the inherent time-based processing capability of reservoir models allows us to demonstrate the effectiveness of our model on complex temporal video data.

2. MATERIALS AND METHODS

2.1. Reservoir Model: Framework and Implementation

2.1.1. Model Architecture

Let us first discuss briefly the Driven/Autonomous model approach (Abbott et al., 2016). The general architecture of the Driven/Autonomous reservoir model is shown in **Figure 1A**. Both driven/auto models are 3-layered networks. The input layer is connected to a reservoir that consists of nonlinear (Leaky-Integrate-and-Fire, LIF), spiking neurons recurrently connected to each other in a random sparse manner (generally with a connection probability of 10%), and the spiking output produced by the reservoir is then projected to the output/readout neurons. In case of driven, the output neurons are rate neurons that simply integrate the spiking activity of the reservoir neurons. On the other hand, the auto model is an entirely spiking model with LIF neurons at the output as well as within the reservoir. Please note, each reservoir neuron here is excitatory in nature and the reservoir does not have any inhibitory neuronal dynamics in a manner consistent with conventional Liquid or Reservoir Computing.

Essentially, the role of the driven network is to provide targets (f_{out}) for the auto model, which is the spiking network we are trying to construct for a given task. To achieve this, the driven network is trained on an input (f_D) that is a high pass filtered

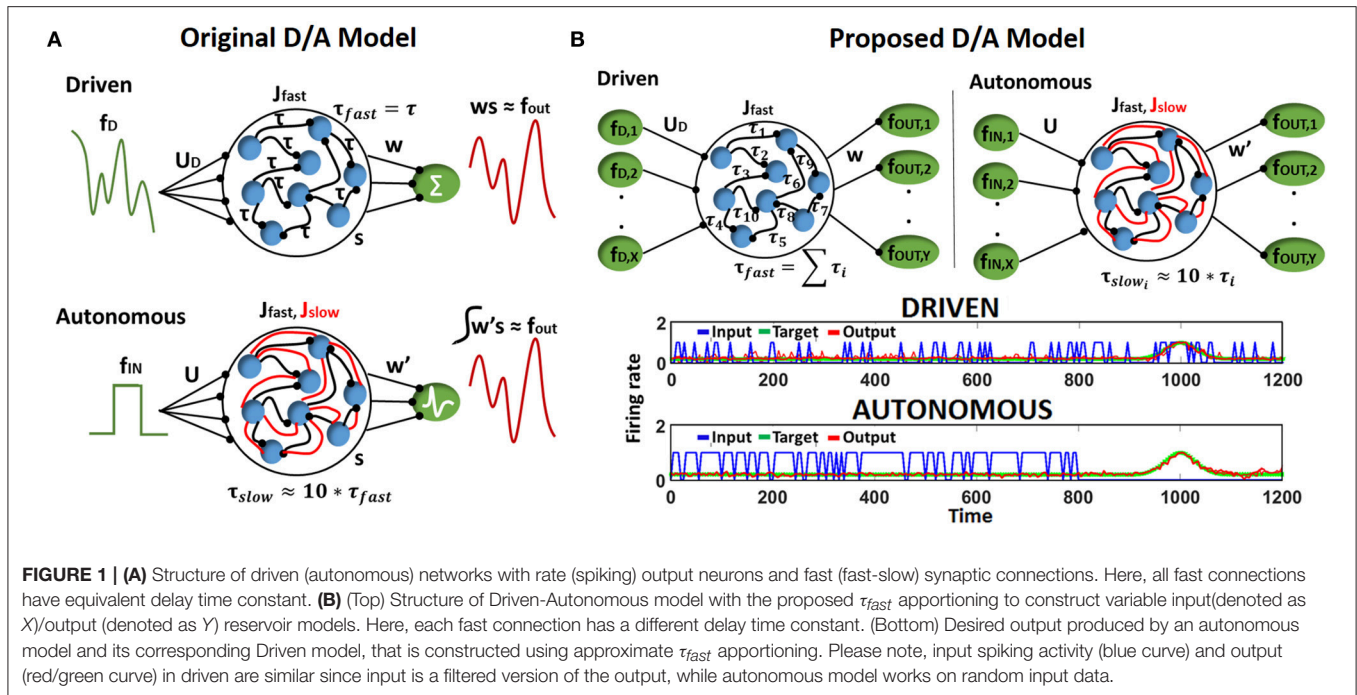


FIGURE 1 | (A) Structure of driven (autonomous) networks with rate (spiking) output neurons and fast (fast-slow) synaptic connections. Here, all fast connections have equivalent delay time constant. **(B)** (Top) Structure of Driven-Autonomous model with the proposed τ_{fast} apportioning to construct variable input (denoted as X)/output (denoted as Y) reservoir models. Here, each fast connection has a different delay time constant. (Bottom) Desired output produced by an autonomous model and its corresponding Driven model, that is constructed using approximate τ_{fast} apportioning. Please note, input spiking activity (blue curve) and output (red/green curve) in driven are similar since input is a filtered version of the output, while autonomous model works on random input data.

version of the desired output (f_{out}),

$$f_D = f_{out} + \tau_{fast} \frac{df_{out}}{dt} \tag{1}$$

The connections from the reservoir to the readout neurons (w) of the driven model are then trained to produce desired activity using the Recursive Least Square (RLS) rule (Haykin, 2008; Sussillo and Abbott, 2009). Once trained, the driven network connections (J_{fast}) are imported into the autonomous model and a new set of network connections (J_{slow}) are added for each fast connection. The combination of the two sets of synapses, J_{fast} and J_{slow} , allows the auto model to match its internal dynamics to that of the driven network that is already tuned to producing desired f_{out} . In fact, the connections u and J_{slow} in the auto model are derived from the driven network as, $u = u_D u_R$ and $J_{slow} = u_D B w$, where u_R and B are randomly chosen. Note, the time constants of the two sets of synapses is different i.e., $\tau_{slow} \sim 10 * \tau_{fast}$ that allows the autonomous network to slowly adjust its activity to produce f_{out} during training. For further clarification, please refer to Abbott et al. (2016) to gain more details and insights on D/A approach. After the autonomous model is constructed, we train its output layer (that consist of LIF spiking neurons), w' , using a supervised Spike Timing Dependent Plasticity (STDP) rule (Ponulak and Kasiński, 2010) based on the input f_{in} .

The motif behind choosing f_D as a filtered version of f_{out} is to compensate for the synaptic filtering at the output characterized by time constant, τ_{fast} (Eliasmith, 2005; Abbott et al., 2016). However, it is evident that the dependence of f_D on f_{out} from Equation (1) imposes a restriction on the reservoir construction, that is, the number of input and output neurons must be same.

To address this limitation and extend the D/A construction approach to variable input-output reservoir models, we propose an approximate definition of the driven input using a delay apportioning method as,

$$(f_{Dk})_j = (f_{out})_j + \tau_{fast_k} \frac{(df_{out})_j}{dt} \tag{2}$$

Here, k varies from 1 to X , j varies from 1 to Y , where X/Y are the number of input/output neurons in our reservoir model respectively and $\tau_{fast} = \sum_{i=1}^N \tau_{fast_i}$, where N is the number of neurons in the reservoir. **Figure 1B** (Top) shows the architecture of the D/A model with the proposed delay apportioning. For each output function ($(f_{out})_j$), we derive the driven inputs ($(f_{Dk} = 1..X)_j$) by distributing the time constant across different inputs. We are essentially compensating the phase delay due to τ_{fast} in an approximate manner. Nevertheless, it turns out this approximate version of f_D still drives the driven network to produce the desired output and derive the connection matrix of the autonomous model that is then trained on the real input (video data in our case). **Figure 1B** (Bottom) illustrates the desired output generated by a D/A model of 2 (input) \times 400 (reservoir neurons) \times 1 (output) configuration constructed using Equation (2). Note, for the sake of convenience in representation and comparative purposes, the output of both driven/auto models are shown as rate instead of spiking activity. It is clearly seen in the driven case that the input spiking activity for a particular input neuron f_{Dk} , even with τ_{fast} apportioning, approximates the target quite well. Also, the autonomous model derived from the driven network in **Figure 1B** (Bottom), produces the desired spiking activity in response to random input upon adequate training of w' connections.

2.1.2. Supervised STDP and One-Hot Encoding

Reservoir framework relaxes the burden of training by fixing the connectivity within the reservoir and that from input to the reservoir. Merely, the output layer of readout neurons (w, w' in **Figure 1A**) is trained (generally in a supervised manner) to match the reservoir activity with the target pattern. In our D/A based reservoir construction, readout of the driven network is trained using standard RLS learning (Sussillo and Abbott, 2009) while we use a modified version of the supervised STDP rule proposed in Ponulak and Kasiński (2010) to conduct autonomous model training. An illustration of our STDP model is shown in **Figure 2A**. For a given target pattern, the synaptic weights are potentiated or depressed as

$$\Delta w = x_{trace}(t_{target} - t_{actual}) \quad (3)$$

Where t_{actual}/t_{target} denotes the time of occurrence of actual/desired spiking activity at the post-synaptic neuron during the simulation and x_{trace} , namely the presynaptic trace, models the pre-neuronal spiking history. Every time a pre-synaptic neuron fires, x_{trace} increases by 1, otherwise it decays exponentially with τ_{pre} that is of the same order as τ_{fast} . Fundamentally, as illustrated in **Figure 2A**, Equation (3) depresses (or potentiates) the weights at those time instants when actual (or target) spike activity occurs. This ensures that actual spiking converges toward the desired activity as training progresses. Once the desired/actual activity become similar, the learning stops, since at time instants where both desired and actual spike occur simultaneously, the weight update value as per Equation (3) becomes 0.

Since the main aim of our work is to develop a spiking model for action recognition in videos, we model the desired spiking activity at the output layer based on one-hot encoding. As with standard machine learning applications, the output neuron assigned to a particular class is trained to produce high spiking

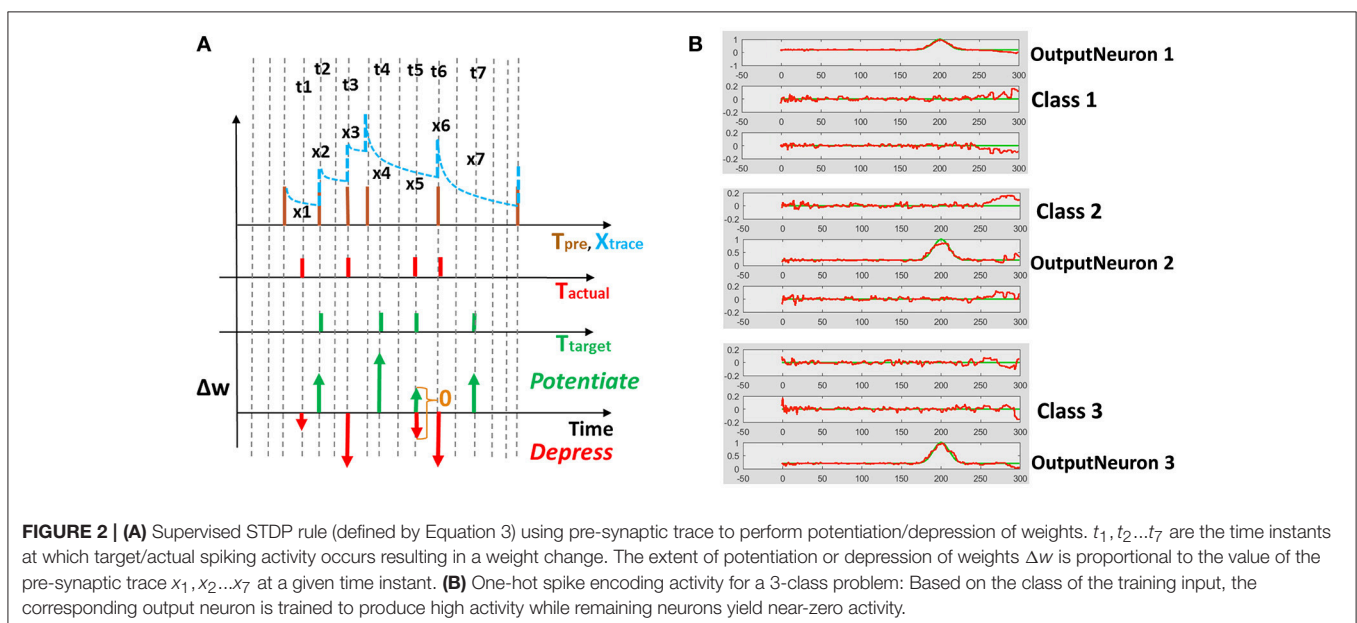
activity while the remaining neurons are trained to generate zero activity. **Figure 2B** shows the one-hot spike encoding for a 3-class problem, wherein the reservoir's readout has 3 output neurons and based on the training input's class, the desired spiking activity (sort of "bump") for each output neuron varies. Again, we use the supervised STDP rule described above to train the output layer of the autonomous model. Now, it is obvious that since the desired spiking activity of the auto model (f_{out} referring to **Figure 1**) follows such one-hot spike encoding, the driven network (from which we derive our auto model) should be constructed to produce equivalent target output activity. To reiterate, the auto model (entirely spiking) works on the real-world input data and is trained to perform required tasks such as classification. The connection matrix (J_{slow}) of the auto model is derived from a driven network (with continuous rate-based output neurons) that, in turn, provides the target activity for the output neurons of the auto model.

2.2. Input Processing

In this work, we use the UCF101 dataset (Soomro et al., 2012) to perform action recognition with our reservoir approach. UCF101 consists of 13,320 realistic videos (collected from YouTube) that fall into 101 different categories or classes, 2 of which are shown in **Figure 3**. Since the reservoir computes using spiking inputs, we convert the pixel valued action information to spike data in steps outlined below.

2.2.1. Pixel Motion Detection

Unlike static images, video data are a natural target in event-based learning algorithms due to their implicit time-based processing capability. First, we track the moving pixels across different frames so that the basic action/signature is captured. This is done by monitoring the difference of the pixel intensity values (P_{diff}) between consecutive frames and comparing the difference against some threshold (ϵ , that is user-defined) to



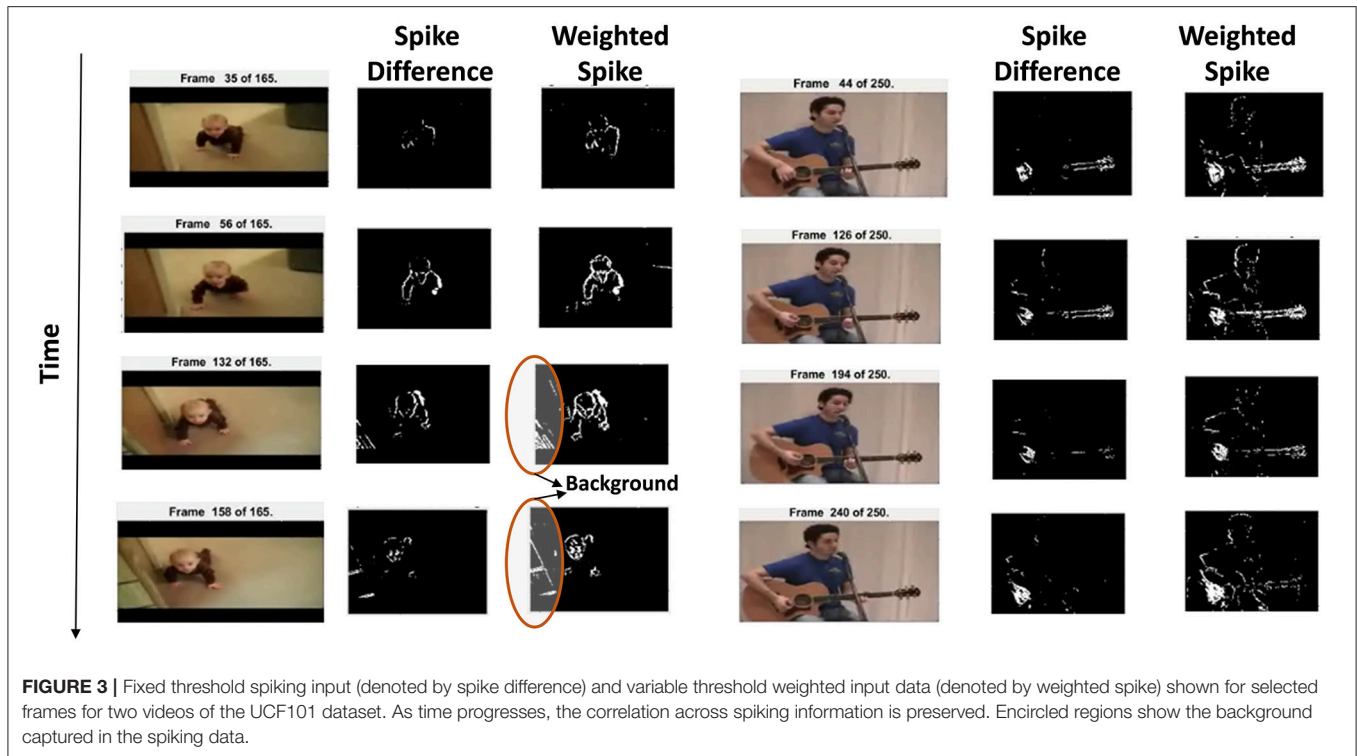


FIGURE 3 | Fixed threshold spiking input (denoted by spike difference) and variable threshold weighted input data (denoted by weighted spike) shown for selected frames for two videos of the UCF101 dataset. As time progresses, the correlation across spiking information is preserved. Encircled regions show the background captured in the spiking data.

detect the moving edges (for each location (x,y) within a frame), as

$$P_{diff}(x,y) = Frame_j - Frame_{j-1};$$

(where $j = 2$ to Number Of Frames)

$$\text{If } P_{diff}(x,y) \geq \epsilon, Spike_{diff}(x,y) = 1 \tag{4}$$

$$WSpike_{diff}(x,y) = \frac{\sum_i \epsilon_i Spike_{diff}(x,y)}{\sum_i \epsilon_i}; \text{ (where } i=1 \text{ to } N)$$

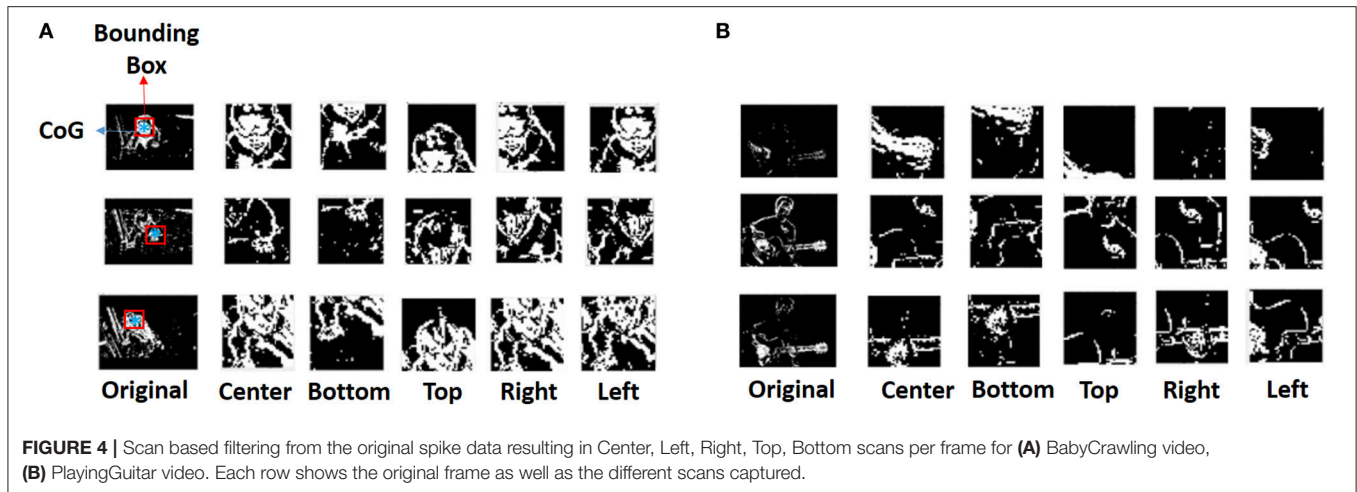
The thresholding detects the moving pixels and yields an equivalent spike pattern ($Spike_{diff}$) for each frame. It is evident that the conversion of the moving edges into spikes preserves the temporal correlation across different frames. While $Spike_{diff}$ encodes the signature action adequately, we further make the representation robust by varying the threshold ($\epsilon_i = \{1,2,4,8,16,32\}$; $N = 6$ in our experiments) and computing the weighted summation of the spike differences for each threshold as shown in the last part of Equation (4). The resultant weighted spike difference $WSpike_{diff}(x,y)$ is then binarized to yield the final spike pattern for each frame. This kind of pixel-change based motion detection to generate spike output is now possible with a new breed of energy efficient sensors inspired by the retina (Hu et al., 2016).

Figure 3 shows the spiking inputs captured from the moving pixels based on fixed threshold spike input and varying threshold weighted input for selected frames. It is clearly seen that the weighted input captures more relevant edges and even subtle movements, for instance, in the *PlayingGuitar* video, the delicate body movement of the person is also captured

along with the significant hand gesture. Due to the realistic nature of the videos in UCF101, there exists a considerable amount of intra class variation along with ambiguous noise. Hence, a detailed input pattern that captures subtle gestures facilitates the reservoir in differentiating the signature motion from noise. It should be noted that we convert the RGB pixel video data into grayscale before performing motion detection.

2.2.2. Scan Based Filtering

We observe from **Figure 3**, for the *BabyCrawling* case, that along with the baby’s movement, some background activity is also captured. This background spiking occurs on account of the camera movement or jitter. Additionally, we see that majority of the pixels across different frames in **Figure 3** are black since the background is largely static and hence does not yield any spiking activity. To avoid the wasteful computation due to non-spiking pixels as well as circumvent insignificant activity due to camera motion, we create a Bounding Box (BB) around the Centre of Gravity (CoG) of spiking activity for each frame and scan across 5 directions as shown in **Figure 4**. The CoG (x_g, y_g) is calculated as the average of the active pixel locations within the BB. In general, the dimensionality of the video frames in the UCF101 dataset is $\sim 200 \times 300$. We create a BB of size 41×41 centered at the CoG, capture the spiking activity enclosed within the BB, in what we refer to as the Center (C) scan, and then stride the BB along different directions (by number of pixels equal to \sim half the distance between the CoG and the edge of the BB), namely, Left (L), Right (R), Top (T), Bottom (B), to capture the activity in the corresponding directions.



These five windows of spiking activity can be interpreted as capturing the salient edge pixels (derived from pixel motion detection above) after each microsaccade (Masquelier et al., 2016). In our model, the first fixation is at the center of the image. Then, the fixation is shifted from the center of the image to the CoG using the first microsaccade. The rest of the four microsaccades are assumed to occur in four directions (i.e., top, bottom, left, and right) with respect to CoG of the image by the stride amount. It should be noted that this sequence of four microsaccades is assumed to be always fixed in its direction. The stride amount can, however, vary based on the size of the BB.

Figure 4A clearly shows that the CRLTB scans only retain the relevant region of interest across different frames of the BabyCrawling video while filtering out the irrelevant spiking activity. Additionally, we also observe that in the PlayingGuitar video shown in **Figure 4B**, different scans capture diverse aspects of the relevant spiking activity, for instance, hand gesture in C scan while guitar shape in L scan shown in the 3rd row. This ensures that all significant movements for a particular action in a given frame are acquired in a holistic manner.

In addition to the scan based filtering, we also delete certain frames to further eliminate the effect of jitter or camera movement on the spike input data. We delete the frames based on two conditions:

- If the overall number of active pixels for a given spiking frame are greater/lesser than a certain maximum/minimum threshold, then, that frame is deleted. This eliminates all the frames that have flashing activity due to unsteady camera movement. In our experiments, we use a min/max threshold value of 8/75% of the total number of pixels in a given frame. Note, we perform this step before we obtain the CRLTB scans.
- Once the CRLTB scans are obtained, we monitor the activity change for a given frame in RLTB as compared to C (i.e., for $Frame_i$ check $C_i - \{R_i, L_i, T_i, B_i\}$). If the activity change (say, $C_i - R_i$) is greater than the mean activity across C_i , then, we delete that particular frame (R_i) corresponding to the scan that yields high variance.

It is noteworthy to mention that the above methods of filtering eliminate the ambiguous motion or noise in videos with minimal background activity or jitter motion (such as PlayingGuitar) to a large extent. However, the videos where the background has significant motion (such as HorseRacing) or those where multiple human subjects are interacting (such as Fencing, IceDancing), the extracted spike pattern does not encode a robust and reliable signature motion due to extreme variation in movement. For future implementations, we are studying using a DVS camera (Hu et al., 2016; Li et al., 2017) to extract reliable spike data despite ego motion. Nonetheless, our proposed input processing method yields reasonable accuracy (shown later in Results section) with the D/A model architecture for action recognition with limited video examples. Supplementary Videos 1, 2 give better visualization of the spiking data obtained with the pixel motion detection and scan based filtering scheme.

2.3. Network Parameters

The proposed reservoir model and learning was implemented in MATLAB. The dynamics of the reservoir neurons in the Driven/Autonomous models are described by

$$\begin{aligned} \tau_{mem} \frac{dV}{dt} &= V_{rest} - V + (J_{fast}f(t) + u_D f_D); \quad \text{Driven} \\ \tau_{mem} \frac{dV}{dt} &= V_{rest} - V + (J_{fast}f(t) + J_{slow}s(t) + u_{fin}); \quad \text{Auto} \end{aligned} \quad (5)$$

where $\tau_{mem} = 20$ ms, f_{in} denotes the real input spike data, f_D is the derived version of the target output (refer to Equation 2), u_D/u are the weights connecting the input to the reservoir neurons in the driven/auto model respectively. Each neuron fires when its membrane potential, V , reaches a threshold $V_{th} = -50$ mV and is then reset to $V_{reset} = V_{rest} = -70$ mV. Following a spike event, the membrane potential is held at the reset value V_{reset} for a refractory period of 5 ms. The parameters $f(t), s(t)$ correspond to fast/slow synaptic currents observed in the reservoir due to the inherent recurrent activity. In the driven model, only fast synaptic connections are present while auto model has both slow/fast connections. The synaptic currents are

encoded as traces, that is, when a neuron in the reservoir spikes, $f(t)$ (and $s(t)$ in case of auto model) is increased by 1, otherwise it decays exponentially as

$$\tau_{slow} \frac{ds(t)}{dt} = -s(t); \quad \tau_{fast} \frac{df(t)}{dt} = -f(t) \quad (6)$$

where $\tau_{slow} = 100$ ms, $\tau_{fast} = 5$ ms. Both J_{slow} and J_{fast} recurrent connections within the reservoir of D/A model are fixed during the course of training. J_{fast} of the driven model are randomly drawn from a normal distribution. J_{slow} and u (refer to **Figure 1A**) of the auto network are derived from the driven model using $u_R = 1$, $B = \text{random number drawn from a uniform distribution in the range } [0, 0.15]$. We discuss the impact of variation of B on the decoding capability of the autonomous reservoir in a later section. If the reservoir outputs a spike pattern, say $x(t)$, the weights from the reservoir to the readout (w) are trained to produce the desired activity, f_{out} . The continuous rate output neurons in driven network are trained using RLS such that $w * x(t) \sim f_{out}(t)$. The spiking output neurons of the auto model (with similar neuronal parameters as Equation 5) integrate the net current $w' * x(t)$ to fire action potentials, that converge toward the desired spiking activity with supervised STDP learning.

2.4. D/A Model for Video Classification

The input processing technique yields 5 different CRLTB scan spike patterns per frame for each video. In the D/A based reservoir construction, each of the 5 scans for every video is processed by 5 different autonomous models. **Figure 5A** shows an overview of our proposed D/A construction approach for video classification. Each scan is processed separately by the corresponding autonomous model. An interesting observation here is that since the internal topology across all auto models is equivalent, they can all be derived from a single driven model as shown in **Figure 5A**. In fact, there is no limit to the number of autonomous models that can be extracted from a driven network, which is the principal advantage of our approach. Thus, in case, more scans based on depth filtering or ego motion monitoring are obtained during input processing, our model can be easily extended to incorporate more autonomous models corresponding to the new scans.

During training, the driven network is first trained to produce the target pattern (similar to the one-hot spike encoding described in **Figure 2B**). Then, each autonomous model is trained separately on the corresponding input scan to produce the target pattern. In all our experiments, we train each auto model for 30 epochs. In each training epoch, we present the training input patterns corresponding to all classes sequentially (for instance, Class1→Class 2→Class 3) to produce target patterns similar to **Figure 2B**. Each input pattern is presented for a time period of 300 ms (or 300 time steps). In each time step, a particular spike frame for a given scan is processed by the reservoir. For patterns where total number of frames is < 300 , the network still continues to train (even in absence of input activity) on the inherent reservoir activity generated due to the recurrent dynamics. In fact, this enables the reservoir to generalize its' dynamical activity over similar input patterns belonging to the

same class while discriminating patterns from other classes. Before presenting a new input pattern, the membrane potential of all neurons in the reservoir are reset to their resting values. It is worth mentioning that the training of the driven model accounts for a slight portion ($\sim 10\%$) of the total training complexity (training 1 Driven + 5 Auto models) in the overall D/A based reservoir construction. This can be attributed to the fact that the driven model essentially receives a slightly modified version of the target output to produce the same target output and hence requires lesser training epochs (15 epochs). In contrast, the autonomous models are trained on the real world input-data to produce the desired targets and eventually require more training epochs (30 epochs per auto model, Total $30 \times 5 = 150$ epochs) to reach an optimal point. Generally, since numerous autonomous models can be derived from the same driven model, the overall training overhead of the driven model (that is trained only once to produce the required targets) tends to be negligible.

After training is done, we pass the test instance CRLTB scans to each auto model simultaneously and observe the output activity. The test instance is assigned a particular class label based on the output neuron class that generated the highest spiking response. This assignment is done for each auto model. Finally, we take a majority vote across all the assignments (or all auto models) for each test input to predict its class. For instance, if majority of the auto models (i.e., ≥ 3) predict a class (say Class 1), then predicted output class from the model is Class 1. Now, a single auto model might not have sufficient information to recognize the test input correctly. The remaining scans or auto models, in that case, can compensate for the insufficient or inaccurate output prediction from the individual models. While we reduce the computational complexity of training by using the scan based filtering approach that partitions the input space into smaller scans, the parallel voting during inference enhances the collective decision making of the proposed architecture, thereby improving the overall performance.

Figure 5B summarizes the parameter values of the D/A model used for the action recognition experiments described below. The reservoir (both driven/auto) topology in case of the video classification experiments consist of $X(\text{Input Neurons}) - N(\text{Reservoir Neurons}) - Y(\text{Output Neurons})$ with $p\%$ connectivity within the reservoir. X is equivalent to the size of the bounding box used during scan based filtering (in our case $X = 41 \times 41 = 1,681$). The number of output neurons vary based on the classification problem, that is, $Y = 5$ for a 5-class problem. In most experiments below, we use $p = 10\%$ connectivity to maintain the sparseness in the reservoir. In the Results section below, we further elaborate on the role of connectivity on the performance of the reservoir. The number of neurons N are varied based upon the complexity of the problem in order to achieve maximum classification accuracy.

3. RESULTS

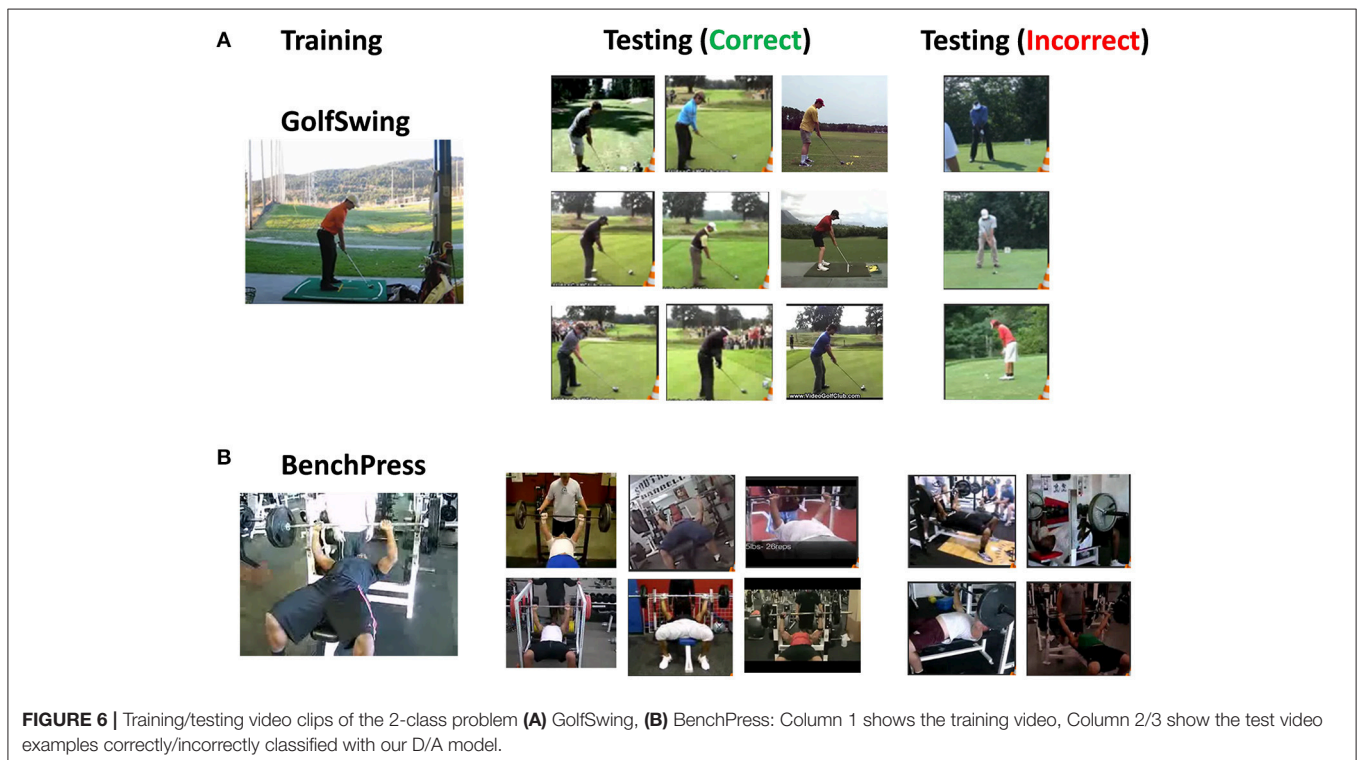
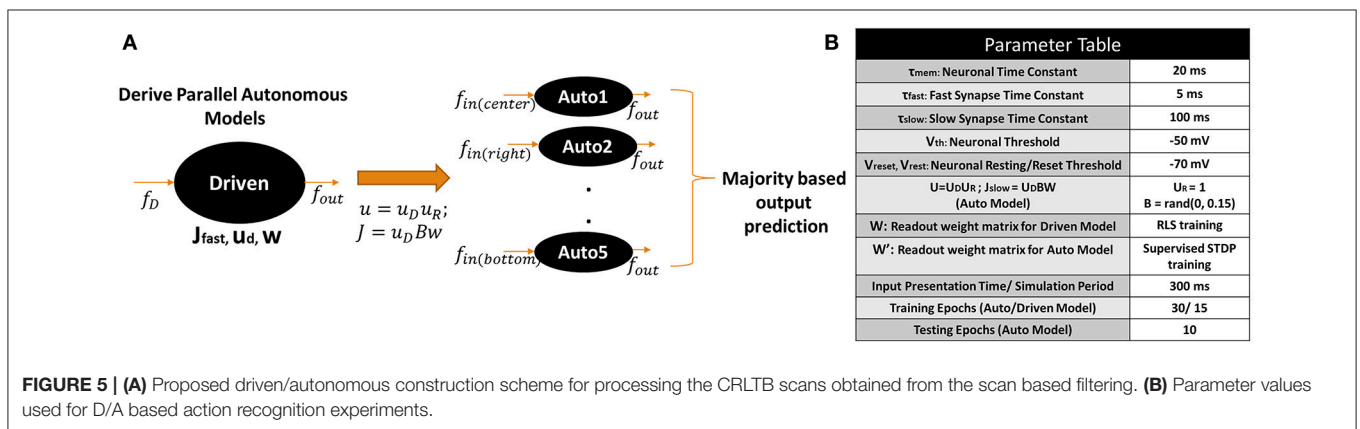
3.1. Classification Accuracy on UCF101

First, we considered a 2-class problem of recognizing BenchPress/GolfSwing class in UCF101 to test the effectiveness of our overall approach. We specifically chose these 2 classes

in our preliminary experiment since majority of videos in these categories had a static background and the variation in action across different videos was minimal. We simulated an 800N-2Output, 10% connectivity D/A reservoir (N:number of reservoir neurons) and trained the autonomous models with a single training video from each class as shown in **Figure 6**. It is clearly seen that even with single video training, the reservoir correctly classifies a host of other examples irrespective of the diverse subjects performing the same action. For instance, in the golf examples in **Figure 6A**, the reservoir does not discriminate between a person wearing shorts or trousers while recognizing the action. Hence, we can gather that the complex transient dynamics within a reservoir latches onto a particular action signature (swing gesture in this case) that enables it to classify several test examples even under diverse conditions.

Investigating the incorrect predictions from the reservoir for both classes, we found that the reservoir is sensitive toward variation in views and inconsistencies in action, which results in inaccurate prediction. **Figure 6** shows that the incorrect results are observed when the view angle of the training/test video changes from frontal to adjacent views. In fact, in case of BenchPress, even videos with similar view as training are incorrectly classified due to disparity in motion signature (for instance, number of times the lifting action is performed). Note, we only show a selected number of videos classified correctly/incorrectly in **Figure 6**.

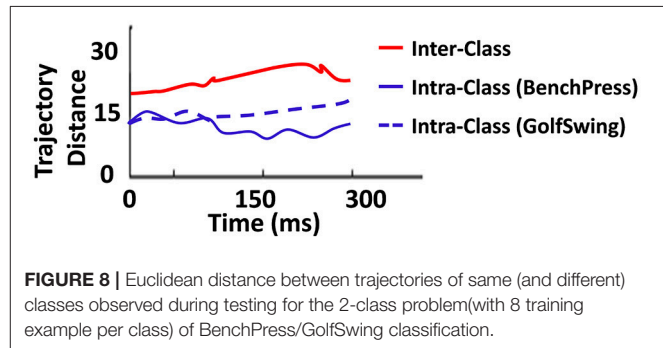
To account for such variations, we increased the number of training examples from 1 to 8. In the extended training set, we tried to incorporate all kinds of view/signature based variations for each action class as shown in **Figure 7A**. For the



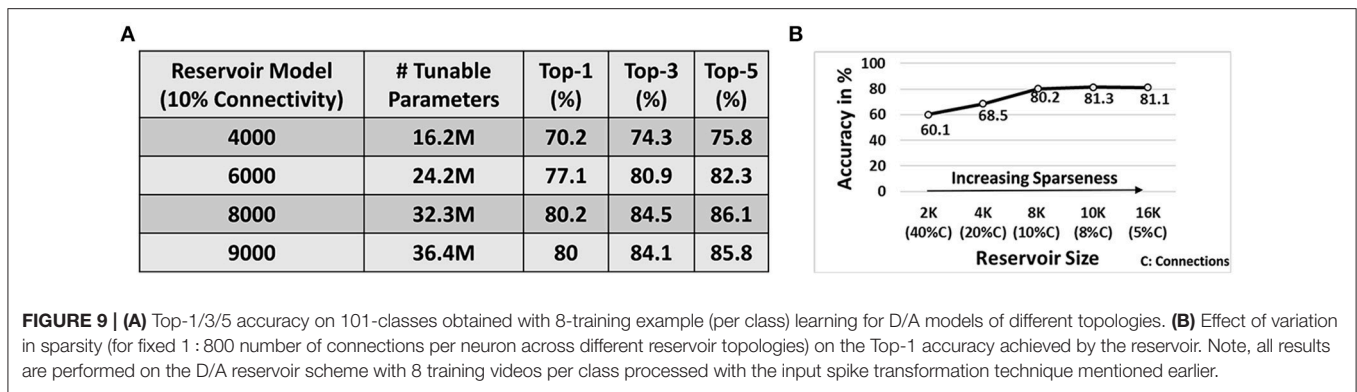


2-class problem, with 8 training examples for each class, we simulated a 2,000N-10% connectivity reservoir that yields an accuracy of 93.25%. **Figure 7B** shows the confusion matrix for the 2-class problem. BenchPress has more misses than GolfSwing as the former has more variation in action signature across different videos. Now, it is well-known that the inherent chaotic or spontaneous activity in a reservoir enables it to produce a wide variety of complex output patterns in response to an input stimulus. In our opinion, this behavior enables data efficient learning as the reservoir learns to generalize over a small set of training patterns and produce diverse yet converging internal dynamics that transform different test inputs toward the same output. Note, testing is done on all the remaining videos in the dataset that were not selected for training.

To further characterize the discriminative performance of the D/A based reservoir network, we plotted the intra-/inter-trajectory distance of the states of the reservoir neurons in the Auto model for same class (for instance, within-BenchPress) vs. those of different classes (between-BenchPress and GolfSwing). The trajectory distance is measured from the Euclidean distance (as $\sqrt{1/N \sum_{i=1}^N (r_{i,1}(t) - r_{i,2}(t))^2}$, where $r_1(t)$ ($r_2(t)$) is the firing rate activity of the reservoir neurons corresponding to input instances I_1 (I_2)), at each time step. Here I_1, I_2 correspond to the different test instances presented to the trained Auto network (learnt with 8 examples per class) for the 2-class problem. **Figure 8** shows the inter-/intra-class trajectory distance averaged across the Euclidean measured for all the testing instances presented to the network. It is desirable to have larger inter-trajectory distance and small intra-trajectory distance such that the network easily distinguishes between two different class of inputs while being able to reproduce the required output response for varying instances belonging to the same class. We observe a similar behavior in **Figure 8**. Now, it is clear that inter-trajectory distance quantifies discriminative capability while generalization ability can be inferred from the intra-trajectory. This ascertains the ability of our reservoir model to classify even with limited training. An interesting observation here is that the intra-trajectory distance measured across test instances belonging to GolfSwing class is higher than that of BenchPress. This implies that the network finds GolfSwing instances to be more confusing/difficult than BenchPress. This further corroborates the declined prediction accuracy observed for GolfSwing in the Confusion Matrix in **Figure 7B**.



Next, we simulated a reservoir of 4K/6K/8K/9K neurons (with 10% connectivity) to learn the entire 101 classes in the UCF101 dataset. Note, for each of the 101 classes, we used 8 different training videos that were randomly chosen from the dataset. To ensure that more variation is captured for training, we imposed a simple constraint during selection, that each training example must have a different subject/person performing the action. **Figure 9A** illustrates the Top-1, Top-3, Top-5 accuracy observed for each topology. Note, in all our experiments with D/A model, during testing, we present the test patterns for 10 epochs (wherein the entire test data corresponding to all classes is presented in each epoch). The prediction accuracy across 10 epochs is then averaged and then reported as the final accuracy. The standard deviation for accuracy across the 10 epochs, during testing, ranges between 0.8 and 1.3% across all experiments. We also report the average number of tunable parameters for each topology. The reservoir has fixed connectivity from the input to reservoir as well as within the reservoir that are not affected during training. Hence, we only compare the total number of trainable parameters (comprising the weights from reservoir to output layer) as the reservoir size varies in **Figure 9A**. We observe that the reservoir with 8K neurons yields maximum accuracy (*Top-1*: 80.2%, *Top-5*: 86.1%). Also, the classification accuracy improves as the reservoir size increases from 4K \rightarrow 8K neurons. This is apparent since larger reservoirs with sparse random connectivity generate richer and heterogeneous dynamics that overall boosts the decoding capability of the network. Thus, for the same task difficulty (in this case 101-class problem), increasing the reservoir size improves the performance. However, beyond a certain point, the accuracy saturates. Here, the reservoir with 9K neurons achieves similar accuracy as 8K. In fact, there is



a slight drop in accuracy from 8K \rightarrow 9K. This might be a result of some sort of overfitting phenomenon as the reservoir (with more parameters than necessary for a given task) might begin to assign, instead of generalizing, its dynamic activity toward a certain input. Note, all reservoir models discussed so far have 10% connectivity. Hence, the reservoirs in **Figure 9A** have increasing number of connections per neuron (1 : 400 \rightarrow 1 : 900) as we increase the number of neurons in the reservoir from 4,000 to 9,000. Please refer to the supplementary information (see Figures S1, S2) that details out the confusion matrix for the 101-class problem with 8K neurons and 10% connectivity.

3.2. Impact of Variation of Reservoir Sparsity on Accuracy

Since sparsity plays an important role in determining the reservoir dynamics, we simulated several topologies in an iso-connectivity scenario to see the effect of increasing sparseness on the reservoir performance. **Figure 9B** shows the variation in Top-1 accuracy for 101-classes as the reservoir sparseness changes. From **Figure 9A**, we observe that the 8,000N-10% connectivity reservoir (with 1 : 800 number of connections per neuron) gave us the best results. Hence, in the sparsity analysis, we fixed the number of connections per neuron to 1 : 800 across different reservoir topologies. Initially, for a 2,000N-40% connectivity reservoir, the accuracy observed is pretty low (60.1%). This is evident as lesser sparsity limits the richness in the complex dynamics that can be produced by the reservoir that further reduces its' learning capability. Increasing the sparseness improves the accuracy as expected. We would like to note that the 4,000N-20% connectivity reservoir yields lesser accuracy than that of the sparser 4,000N-10% connectivity reservoir of **Figure 9A**. Another interesting observation here is that the 10,000N-8% reservoir yields higher accuracy (\sim 1% more) than that of the 8,000N-10% reservoir. This result supports the fact that in an iso-connectivity scenario, more sparseness yields better results. However, the accuracy again saturates with a risk of overfitting or losing generalization capability beyond a certain level of sparseness as observed in **Figure 9B**. Hence, there is no fixed rule to determine the exact connectivity/sparsity ratio for a given task. We empirically arrived at the results by simulating a varying set of reservoir topologies. However, in most cases, we observed that for an N-reservoir model 10%

connectivity generally yields reasonable results as seen from **Figure 9A**. It is worth mentioning that while the 10K neuron reservoir yields the highest accuracy of 81.3%, it consists of larger number of tunable parameters (total connections/weights from reservoir to output) 40.4M as compared to 32.3M observed with the 8K-10% connectivity reservoir (**Figure 9A**). In fact, the best accuracy achieved with our approach corresponds to the 10K-8% connectivity reservoir with Top-1/Top-3/Top-5 accuracy as 81.3/85.2/87%, respectively. Hence, depending upon the requirements, the connectivity/size of the reservoir can be varied to achieve a favorable efficiency-accuracy tradeoff. Please note that the total number of tunable parameters in the reservoir model include all 5-autonomous models. To highlight the universality of our approach, we verified the D/A model for speech classification on TI46 speech corpus dataset. Please refer to the supplementary to get further details and insights (see **Figure S4**).

3.3. EigenValue Spectra Analysis of Reservoir Connections

We analyzed the EigenValue (EV) spectra of the synaptic connections (both J_{slow} and J_{fast} of the autonomous models) to study the complex dynamics of the reservoir. EV spectra has been shown to characterize the memory/learning ability of a dynamical system (Rajan, 2009). Rajan et al. (Rajan and Abbott, 2006; Rajan, 2009) have shown that stimulating a reservoir with random recurrent connections activates multiple modes where neurons start oscillating with individual frequencies. These modes then superpose in a highly non-linear manner resulting in the complex chaotic/persistent activity observed from a reservoir. The EVs of the synaptic matrix typically encode the oscillatory behavior of each neuron in the reservoir. The real part of an EV corresponds to the decay rate of the associated neuron, while frequency is given by the imaginary part. If the real part of a complex EV exceeds 1 (or $\text{Re}(\text{EV}) > 1$), the activated mode leads to long lasting oscillatory behavior. If there are a few modes with $\text{Re}(\text{EV}) > 1$, the network is at a fixed point, meaning the reservoir exhibits the same pattern of activity for a given input. It is desirable to construct a reservoir with multiple fixed points so that each one can be used to retain a different memory. In a recognition scenario, each of the fixed points latch onto a particular input pattern or video in our case, thereby yielding a

good memory model. As the number of modes with $\text{Re}(\text{EV}) > 1$ increases, the reservoir starts generating chaotic patterns, that are complex and non-repeating. In order to construct a reservoir model with good memory, we must operate in a region between singular fixed point activity ($\text{Re}(\text{EV}) \ll 1$) and complete chaos ($\text{Re}(\text{EV}) \gg 1$).

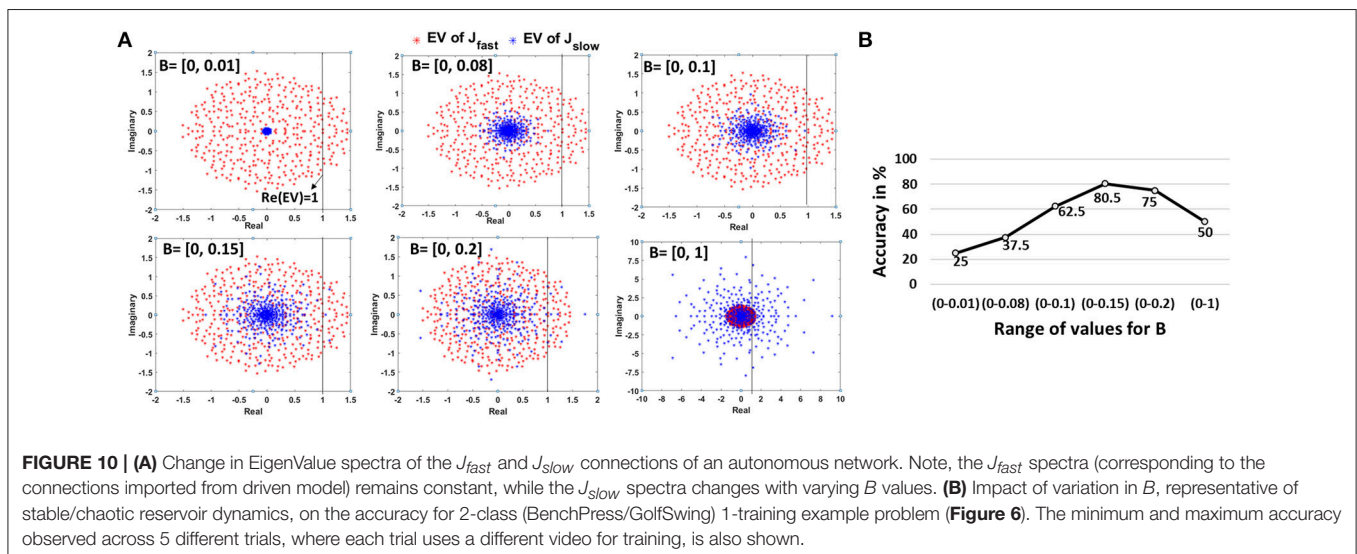
Since our autonomous models have two kinds of synapses with different time constants, we observe two different kinds of EV spectra that contribute concurrently to the dynamical state of the auto network. While the overall reservoir dynamics are dictated by the effective contribution of both connections, J_{slow} (due to larger time constant of synaptic decay) will have a slightly more dominating effect. **Figure 10A** illustrates the EV spectra corresponding to the bottom scan autonomous 800N-reservoir model simulated earlier for the 2-class problem (BenchPress/GolfSwing) with just 1 training video (refer **Figure 6**). Now, the EV of J_{fast} have considerable number of modes with $\text{Re}(\text{EV}) > 1$ that is characteristic of chaotic dynamics. Hence, as discussed above, we need to compensate for the overwhelming chaotic activity with modes that exhibit fixed points to obtain a reliable memory model. This can be attained by adjusting the J_{slow} connections. From **Figure 10A**, we see that as the range of B (random constant used to derive the J_{slow} connections from the driven model) increases, the EV spectral radius of the J_{slow} connections grows. For $B = [0 - 0.01]$, the EVs are concentrated at the center. Such networks have almost no learning capability as the dynamics converge to a single point that produces the same time-independent pattern of activity for any input stimulus. As the spectral circle enlarges, the number of modes with $\text{Re}(\text{EV}) > 1$ also increases. For $B = [0 - 0.15]$, the number of fixed point modes reaches a reasonable number that in coherence with the chaotic activity due to J_{fast} results in generalized reservoir learning for variable inputs. Increasing B further expands the spectral radius for J_{slow} (with more $\text{Re}(\text{EV}) > 1$) that adds onto the chaotic activity generated due to J_{fast} . In such scenarios, the reservoir will not be able to discern between inputs that look similar due to persistent chaos. This is akin to the

observations about the need for the reservoir to strike a balance between pattern approximation and separability (Rajan, 2009).

Figure 10B demonstrates the accuracy variation for the same 2-class 1-training example learning model with different values of B . We repeated the experiment 5 times wherein we chose a different training video per class in each case. As expected, the accuracy for a fixed state reservoir with $B = [0 - 0.01]$ is very low. As the number of modes increase, the accuracy improves justifying the relevance of stable and chaotic activity for better learning. As the chaotic activity begins to dominate with increasing B , the accuracy starts declining. We also observe that the overall accuracy trend (for the 2-class 1-training example model) represented by the min-max accuracy bar is consistent even when the training videos are varied. This indicates that the reservoir accuracy is sensitive to the B -initialization and is not biased toward any particular training video in our limited training example learning scheme. While the empirical result corroborates our theoretical analysis and justifies the effectiveness of using the D/A approach in reservoir model construction, further investigation needs to be done to gauge the other advantages/disadvantages of the proposed approach. Please refer to Rajan and Abbott (2006) and Rajan (2009) for more clarification on random matrices and EV spectra implication on dynamical reservoir models.

3.4. Comparison With State-of-the-Art Recognition Models

Deep Learning Networks (DLNs) are the current state-of-the-art learning models that have made major advances in several recognition tasks (Mnih et al., 2013; Cox and Dean, 2014). While these large-scale networks are very powerful, they inevitably require large training data and enormous computational resources. Spiking networks offer an alternative solution by exploiting event-based data-driven computations that makes them attractive for deployment on real-time neuromorphic hardware where power consumption and speed become vital



constraints (Merolla et al., 2014; Han et al., 2015). While research efforts in spiking models have soared in the recent past (Masquelier and Thorpe, 2007; Srinivasa and Cho, 2012; Maass, 2016; Panda and Roy, 2016), the performance of such models are not as accurate as compared to their artificial counterparts (or DLNs). From our perspective, this work provides a new standard establishing the effectiveness of spiking models and their inherent timing-based processing for action recognition in videos.

Figure 11 compares our reservoir model to state-of-the-art recognition results (Simonyan and Zisserman, 2014a; Srivastava et al., 2015; Feichtenhofer et al., 2016). We simulated a VGG-16 (Simonyan and Zisserman, 2014b) DLN architecture that only processes the spatial aspects of the video input. That is, we ignore the temporal features of the video data and attempt to classify each clip by looking at a single frame. We observe that the spatial VGG-16 model yields reduced accuracy of 66.2% (42.9%) when presented with the full training dataset (8-training examples per class), implying the detrimental effect of disregarding temporal statistics in the data. Please refer to the supplementary information (see Figure S3) that details out the efficiency vs. accuracy tradeoff between our proposed reservoir and the spatial VGG-16 model. The static hierarchical/feedforward nature of DLNs serves as a major drawback for sequential (such as video) data processing. While deep networks extract good flexible representations from static images, addition of temporal dimension in the input necessitates the model to incorporate recurrent connections that can capture the underlying temporal correlation from the data.

Recent efforts have integrated Recurrent Neural Networks (RNNs) and Long Short-Term Memory (LSTMs) with DLNs to process temporal information (Donahue et al., 2015; Srivastava et al., 2015). In fact, the state-of-the-art accuracy on UCF-101 is reported as 92.5% (Feichtenhofer et al., 2016) as illustrated in **Figure 11**, where a two-stream convolutional architecture with two VGG-16 models, one for spatial and another for temporal information processing is used. While the spatial VGG model is trained on the real-valued RGB data, the temporal VGG model uses optical flow stacking to capture the correlated motion. Here, both the models are trained on the entire UCF-101 dataset that significantly increases the overall training complexity. All models noted in **Figure 11** (including the LSTM model) use the temporal information in the input data by computing the flow features and subsequently training the deep learning/recurrent models on such inputs. The overall accuracy shown in **Figure 11** is the combined prediction obtained from the spatial RGB and temporal flow models. The incorporation of temporal information drastically improves a DLN's accuracy.

It is evident that the main idea of using the D/A based reservoir construction is to determine a suitable recurrent connection matrix such that the reservoir (Auto model in our case) performs well on a given task. In order to quantify the effectiveness of the D/A approach, we compared the performance of our D/A based reservoir with that of a vanilla Liquid State Machine (LSM) (Maass et al., 2002; Maass, 2011). The LSM was constructed with randomly initialized recurrent connections. Moreover, since our input consists of 5 different scanned spike patterns, namely CRLTB, each of the 5 scans were processed

by 5 different LSMs (each with a different set of recurrent connectivity). Considering same connectivity for all 5 LSMs yielded lower accuracy. Hence, we conducted the analysis on LSMs with diverse connectivity. As with our D/A model, a majority based output prediction yielded the final prediction class for a given instance from the 5-LSM system. While both D/A model and vanilla LSM were able to give a lower error and better convergence during training on a small 10-class problem, the prediction capability of the vanilla LSM severely decreased when we increase the number of classes to 20-/30-classes. Note, we trained both the networks (D/A, vanilla LSM) with 8 training examples per class. **Figure 11** shows the accuracy obtained from a vanilla LSM with similar topology as our D/A reservoir model for 101-class problem. We believe that the combination of fast/slow synapses in our reservoir model renders the network more flexibility to learn and hence discriminate better than that of a vanilla LSM. However, in spite of the low accuracy, the vanilla LSM still performs better than the spatial VGG-16 model in the iso-training scenario. Thus, we can deduce that the recurrent and inherent spike-time based processing capability of a spiking network enables training with fewer examples. However, the D/A based approach provides an additional avenue of constructing a better reservoir with improved discriminative capability. Note, the simulation parameters (neuronal constants, training/test epochs) of the vanilla LSM were similar to that of the parameters of the autonomous model. As there are only one set of synapses in the vanilla LSM, the synaptic time constant in this case was set to $\tau = 40$ ms to calculate the synaptic current (refer to Equation 6).

Figure 11 also quantifies the benefits with our reservoir approach as compared to the state-of-the-art models from a computational efficiency perspective. We quantify efficiency in terms of total number of resources utilized during training a particular model. The overall resource utilization is defined as the product of the Number of (tunable) Parameters \times Number of training Data Examples. The number of data examples are calculated as Number of Frames per video \times Number of Training videos ($\equiv 180 \times 9,537$ for full training data). Referring to the results for spatial VGG-16 model, we observe that the reservoir (10K-10% connectivity) is $2.4\times$ more efficient, while yielding significantly higher accuracy, in the iso-data scenario where both the models are shown 8 video samples per class. In contrast, the efficiency improves extensively to $28.5\times$ in the iso-accuracy scenario wherein the reservoir still uses 8 video samples per class (808 total videos), while the VGG-16 model requires the entire training dataset (9,537 total videos) during training. Comparing with the models that include temporal information, our reservoir (10K-10% connectivity), although with lower accuracy, continues to yield significantly higher benefits of approximately $54\times$ and $7\times$ as compared to Two-stream architecture and Temporal Convolutional network, respectively, on account of being trained on limited data. However, in limited training example scenario, the reservoir (81.3% accuracy) continues to outperform an LSTM based machine learning model (67% accuracy) while requiring lesser training data. It is worth mentioning that our reservoir due to the inherent temporal processing capability requires only 30 epochs for training. On the other hand, the spatial

Model	# Training Videos	Accuracy	# Parameters	# Data Examples	Resource Utilization (B)	Model Efficiency
Spatial VGG-16 Model	Full Training	66.2 %	97.5M	1716660	167374.4	0.035
Temporal ConvNet [Simoyan et al. 2014a]	Full Training	83.7 %	23.5M	1716660	40341.5	0.146
Composite LSTM Model [Srivastava et al. 2015]	Full Training	84.3 %	-	1716660	-	-
Two-Stream ConvNet [Feichtenhofer et al. 2016]	Full Training	92.5 %	185M	1716660	317582.1	0.019
Spatial VGG-16 Model	8 videos per class	42.9 %	97.5M	145440	14180.4	0.414
Composite LSTM Model [Srivastava et al. 2015]	10 videos per class	67 %	-	181800	-	-
Vanilla LSM 8K-10%Connectivity	8 videos per class	44.1%	32.3M	145440	4697.7	1.251
Vanilla LSM 10K-10%Connectivity	8 videos per class	47.3%	40.4M	145440	5875.8	1
Ours 8K-10%Connectivity	8 videos per class	80.2%	32.3M	145440	4697.7	1.251
Ours 10K-10%Connectivity	8 videos per class	81.3%	40.4M	145440	5875.8	1

FIGURE 11 | Accuracy and Efficiency comparison with state-of-the-art recognition models. Resource Utilization (B) is the product of Number of Parameters \times Number of Data Examples, where B denotes billion. The model efficiency is calculated by normalizing the net resource utilization of each model against our reservoir model (10K-10% connectivity). Full Training (in column corresponding to # Training Videos) denotes the entire training dataset that comprises of 9,537 total videos across 101 classes.

VGG-16 model needs 250 epochs. As the number of epochs required for training increase, the overall training cost (can be defined as, Number of Epochs \times Resource Utilization) would consequently increase. This further ascertains the potential of reservoir computing for temporal data processing. Note, in terms of inference time, spiking networks (such as our reservoir) unlike deep learning models will have higher latency as the reservoir has to process each test video input for the entire simulation time period (300 ms in our case) before we monitor the output at the readout neurons. However, spiking networks owing to the inherent sparse event driven activity will be more power-efficient than their non-spiking (or standard DLN) counterparts. Essentially, a large part of the spiking network (in terms of activity) is idle at any given time instant that makes it a low power option in comparison to non-spiking models where each and every unit is active throughout the testing/inference period.

4. DISCUSSION

In this paper, we demonstrated the effectiveness of a spiking-based reservoir computing architecture for learning from limited video examples on UCF101 action recognition task. We adopted the Driven/Autonomous reservoir construction approach originally proposed for neuroscientific use and gave it a new dimension to perform multi input-output transformation for practical recognition tasks. In fact, the D/A approach opens up novel possibilities for multi-modal recognition. For instance,

consider a 5-class face recognition scenario. Here, the driven network is trained to produce 5 different desired activity. The auto models derived from the driven network will process the facial data to perform classification. Now, the same driven model can also be used to derive autonomous models that will work on, say voice data. Combining the output results across an ensemble of autonomous models will boost the overall confidence of classification. Apart from performance, the D/A model also has several hardware oriented advantages. Besides sparse power-efficient spike communication, the reservoir internal topology across all autonomous models for a particular problem remains equivalent. This provides a huge advantage of reusability of network structures while scaling from one task to another.

Our analysis using eigenvalue spectra results on the reservoir provide a key insight about D/A based reservoir construction such that the complex dynamics can converge to fixed memory states. We believe that this internal stabilization enables it to generalize its memory from a few action signatures. We also compared our reservoir framework with state-of-the-art models and reported the cost/accuracy benefits gained from our model. We would like to note that our work delivers a new benchmark for action recognition from limited training videos in spiking scenario. We believe that we can further improve the accuracy with our proposed approach by employing better input encoding techniques. The scan-based filtering approach discussed in this paper does not account for depth variations and is susceptible to extreme spatial/rotational variation. In fact,

we observe that our reservoir model yields higher accuracy for classes with more consistent action signatures and finds it difficult to recognize classes with more variations in movement (refer to Figure S1 for action specific accuracies). Using an adequate spike processing technique (encompassing depth based filtering, motion based tracking or a fusion of these) will be key to achieving improved accuracy. Another intriguing possibility is to combine our spike based model that can rapidly recognize simpler and repeatable actions (e.g., GolfSwing) with minimal samples for training with more other DLN models for those actions that have several variations (e.g., Knitting) or the actions that appear different dependent on the view point or both. We are also studying the extension of our SNN model to a hierarchical model that could possibly capture more complex and subtle aspects of the spatio-temporal signatures in these actions. Further, we would like to mention the state-of-the-art reservoir based techniques (Norton and Ventura, 2010; Roy and Basu, 2016; Panda and Roy, 2017) that rearrange/learn the connections within the reservoir, in addition to training the reservoir to readout connections, to yield a network that performs better than a randomly generated reservoir. Our D/A based reservoir construction approach provides a simpler avenue (without synaptic plasticity within the reservoir) to build functional spiking models that yield competitive accuracy with lesser training complexity. However, combining the separation driven/structural plasticity mechanisms from Roy and Basu (2016), Norton and Ventura (2010), and Panda and Roy (2017) with the D/A approach is a promising direction of future research. The nature and mode of action of such mechanisms might help us replace the least-squares adjustment of synaptic weights in the driven model with more biophysically realistic learning mechanisms.

Finally, our reservoir model, although biologically inspired, only has excitatory components. The brain's recurrent circuitry consists of both excitatory and inhibitory neurons that operate in a balanced regime giving it the staggering ability to make sense of a complex and ever-changing world in the most energy-efficient manner (Herculano-Houzel, 2012). In the future, we will concentrate on incorporating the inhibitory scheme that will further augment the sparseness in reservoir activity

thereby improving the overall performance. An interesting line of research in this regard can be found here (Srinivasa and Cho, 2014; Brendel et al., 2017). Although there have been multiple efforts in the spiking domain exploring hierarchical feedforward and reservoir architectures with biologically plausible notions, most of them have been focused on static images (and some voice recognition) that use the entire training dataset to yield reasonable results (Masquelier and Thorpe, 2007; Norton and Ventura, 2010; Diehl et al., 2015; Kheradpisheh et al., 2016; Panda and Roy, 2016; Roy and Basu, 2016). To that effect, this work provides a new perspective on the power of temporal processing for rapid recognition in cognitive applications using spiking networks. In fact, we are currently exploring the use of this model for learning at the edge (such as cell phones and sensors where energy is of paramount importance) using a recently developed neural chip called the Loihi (Morris, 2017) that enables energy efficient on-chip spike timing based learning. We believe that results from this work could be leveraged to enable Loihi based applications such as video analytics, video annotation and rapid searches for actions in video databases.

AUTHOR CONTRIBUTIONS

NS conceived the study. PP designed and conducted the experiment(s). Both PP and NS analyzed the results. Both the authors contributed to writing the manuscript.

ACKNOWLEDGMENTS

This work was conducted in Intel Labs, Hillsboro. We would like to acknowledge the support of PP through a summer internship in the Microarchitecture Research Labs at Intel Corporation. We would also like to thank Michael Mayberry and Daniel Hammerstrom for their helpful comments and suggestions.

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fnins.2018.00126/full#supplementary-material>

REFERENCES

- Abbott, L. F., DePasquale, B., and Memmesheimer, R. M. (2016). Building functional networks of spiking model neurons. *Nat. Neurosci.* 19:350. doi: 10.1038/nn.4241
- Brendel, W., Bourdoukan, R., Verterchi, P., Machens, C. K., and Denève, S. (2017). Learning to represent signals spike by spike. arXiv 1703.03777.
- Burgess, J., Lloyd, J. R., and Ghahramani, Z. (2017). One-shot learning in discriminative neural networks. arXiv:1707.05562.
- Cox, D. D., and Dean, T. (2014). Neural networks and neuroscience-inspired computer vision. *Curr. Biol.* 24, R921–R929. doi: 10.1016/j.cub.2014.08.026
- Diehl, P. U., Neil, D., Binas, J., Cook, M., Liu, S.-C., and Pfeiffer, M. (2015). “Fast-classifying, high-accuracy spiking deep networks through weight and threshold balancing,” in *International Joint Conference on Neural Networks (IJCNN)* (Killarney: IEEE), 1–8.
- Donahue, J., Anne Hendricks, L., Guadarrama, S., Rohrbach, M., Venugopalan, S., Saenko, K., et al. (2015). “Long-term recurrent convolutional networks for visual recognition and description,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2625–2634.
- Eliasmith, C. (2005). A unified approach to building and controlling spiking attractor networks. *Neural Comput.* 17, 1276–1314. doi: 10.1162/0899766053630332
- Fei-Fei, L., Fergus, R., and Perona, P. (2007). Learning generative visual models from few training examples: an incremental bayesian approach tested on 101 object categories. *Comput. Vis. Image Unders.* 106, 59–70. doi: 10.1016/j.cviu.2005.09.012
- Feichtenhofer, C., Pinz, A., and Zisserman, A. (2016). Convolutional two-stream network fusion for video action recognition. arXiv:1604.06573.
- Han, S., Pool, J., Tran, J., and Dally, W. J. (2015). “Learning both weights and connections for efficient neural network,” in *NIPS'15 Proceedings of the 28th*

- International Conference on Neural Information Processing Systems*, Vol. 1 (Montreal, QC), 1135–1143.
- Haykin, S. S. (2008). *Adaptive Filter Theory*. Pearson Education India.
- Herculano-Houzel, S. (2012). The remarkable, yet not extraordinary, human brain as a scaled-up primate brain and its associated cost. *Proc. Natl. Acad. Sci. U.S.A.* 109(Suppl. 1), 10661–10668. doi: 10.1073/pnas.1201895109
- Hu, Y., Liu, H., Pfeiffer, M., and Delbruck, T. (2016). Dvs benchmark datasets for object tracking, action recognition, and object recognition. *Front. Neurosci.* 10:405. doi: 10.3389/fnins.2016.00405
- Kheradpisheh, S. R., Ganjtabesh, M., and Masquelier, T. (2016). Bio-inspired unsupervised learning of visual features leads to robust invariant object recognition. *Neurocomputing* 205, 382–392. doi: 10.1016/j.neucom.2016.04.029
- Lake, B. M., Salakhutdinov, R., Gross, J., and Tenenbaum, J. B. (2011). “One shot learning of simple visual concepts,” in *Proceedings of the Annual Meeting of the Cognitive Science Society*, Vol. 33.
- Lake, B. M., Salakhutdinov, R., and Tenenbaum, J. B. (2015). Human-level concept learning through probabilistic program induction. *Science* 350, 1332–1338. doi: 10.1126/science.aab3050
- LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *Nature* 521, 436–444. doi: 10.1038/nature14539
- Li, H., Liu, H., Ji, X., Li, G., and Shi, L. (2017). CIFAR10-DVS: an event-stream dataset for object classification. *Front. Neurosci.* 11:309. doi: 10.3389/fnins.2017.00309
- Lukoševičius, M., and Jaeger, H. (2009). Reservoir computing approaches to recurrent neural network training. *Comput. Sci. Rev.* 3, 127–149. doi: 10.1016/j.cosrev.2009.03.005
- Maass, W. (2011). “Liquid state machines: motivation, theory, and applications,” in *Computability in Context: Computation and Logic in the Real World* (World Scientific), 275–296.
- Maass, W. (2016). Energy-efficient neural network chips approach human recognition capabilities. *Proc. Natl. Acad. Sci. U.S.A.* 113, 11387–11389. doi: 10.1073/pnas.1614109113
- Maass, W., Natschläger, T., and Markram, H. (2002). Real-time computing without stable states: a new framework for neural computation based on perturbations. *Neural Comput.* 14, 2531–2560. doi: 10.1162/089976602760407955
- Masquelier, T., Portelli, G., and Kornprobst, P. (2016). Microsaccades enable efficient synchrony-based coding in the retina: a simulation study. *Sci. Rep.* 6:24086. doi: 10.1038/srep24086
- Masquelier, T., and Thorpe, S. J. (2007). Unsupervised learning of visual features through spike timing dependent plasticity. *PLoS Comput. Biol.* 3:e31. doi: 10.1371/journal.pcbi.0030031
- Merolla, P. A., Arthur, J. V., Alvarez-Icaza, R., Cassidy, A. S., Sawada, J., Akopyan, F., et al. (2014). A million spiking-neuron integrated circuit with a scalable communication network and interface. *Science* 345, 668–673. doi: 10.1126/science.1254642
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., et al. (2013). Playing atari with deep reinforcement learning. arXiv:1312.5602.
- Morris, J. (2017). Available online at: <http://www.zdnet.com/article/why-intel-built-a-neuromorphic-chip/>
- Norton, D., and Ventura, D. (2010). Improving liquid state machines through iterative refinement of the reservoir. *Neurocomputing* 73, 2893–2904. doi: 10.1016/j.neucom.2010.08.005
- Panda, P., and Roy, K. (2016). “Unsupervised regenerative learning of hierarchical features in spiking deep networks for object recognition,” in *Neural Networks (IJCNN), 2016 International Joint Conference on (IEEE)*, 299–306.
- Panda, P., and Roy, K. (2017). Learning to generate sequences with combination of Hebbian and non-hebbian plasticity in recurrent spiking neural networks. *Front. Neurosci.* 11:693. doi: 10.3389/fnins.2017.00693
- Ponulak, F., and Kasinski, A. (2010). Supervised learning in spiking neural networks with resume: sequence learning, classification, and spike shifting. *Neural Comput.* 22, 467–510. doi: 10.1162/neco.2009.11-08-901
- Rajan, K. (2009). *Spontaneous and Stimulus-Driven Network Dynamics*. Columbia University.
- Rajan, K., and Abbott, L. (2006). Eigenvalue spectra of random matrices for neural networks. *Phys. Rev. Lett.* 97:188104. doi: 10.1103/PhysRevLett.97.188104
- Roy, S., and Basu, A. (2016). An online structural plasticity rule for generating better reservoirs. *Neural Comput.* 28, 2557–2584. doi: 10.1162/NECO_a_00886
- Santoro, A., Bartunov, S., Botvinick, M., Wierstra, D., and Lillicrap, T. (2016). One-shot learning with memory-augmented neural networks. arXiv:1605.06065.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., van den Driessche, G., et al. (2016). Mastering the game of go with deep neural networks and tree search. *Nature* 529, 484–489. doi: 10.1038/nature16961
- Simonyan, K., and Zisserman, A. (2014a). “Two-stream convolutional networks for action recognition in videos,” in *Advances in Neural Information Processing Systems*, 568–576.
- Simonyan, K., and Zisserman, A. (2014b). Very deep convolutional networks for large-scale image recognition. arXiv:1409.1556.
- Soomro, K., Zamir, A. R., and Shah, M. (2012). Ucf101: a dataset of 101 human actions classes from videos in the wild. arXiv:1212.0402.
- Srinivasa, N., and Cho, Y. (2012). Self-organizing spiking neural model for learning fault-tolerant spatio-motor transformations. *IEEE Trans. Neural Netw. Learn. Syst.* 23, 1526–1538. doi: 10.1109/TNNLS.2012.2207738
- Srinivasa, N., and Cho, Y. (2014). Unsupervised discrimination of patterns in spiking neural networks with excitatory and inhibitory synaptic plasticity. *Front. Comput. Neurosci.* 8:159. doi: 10.3389/fncom.2014.00159
- Srivastava, N., Mansimov, E., and Salakhutdinov, R. (2015). “Unsupervised learning of video representations using LSTMs,” in *International Conference on Machine Learning (Sweden)*, 843–852.
- Sussillo, D., and Abbott, L. F. (2009). Generating coherent patterns of activity from chaotic neural networks. *Neuron* 63, 544–557. doi: 10.1016/j.neuron.2009.07.018
- Vinyals, O., Blundell, C., Lillicrap, T., Kavukcuoglu, K., and Wierstra, D. (2016). “Matching networks for one shot learning,” in *Advances in Neural Information Processing Systems*, 3630–3638.
- Wehr, M., and Zador, A. M. (2003). Balanced inhibition underlies tuning and sharpens spike timing in auditory cortex. *Nature* 426, 442–446. doi: 10.1038/nature02116

Conflict of Interest Statement: PP and NS were employed by company Intel Corporation.

Copyright © 2018 Panda and Srinivasa. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.