

SpykeTorch: Efficient Simulation of Convolutional Spiking Neural Networks with at most one Spike per Neuron

Milad Mozafari^{1,2}, Mohammad Ganjtabesh¹,
Abbas Nowzari-Dalini¹, Timothée Masquelier²

¹U of Tehran, Iran. ²CERCO, CNRS - U Toulouse 3, France.



1. Introduction

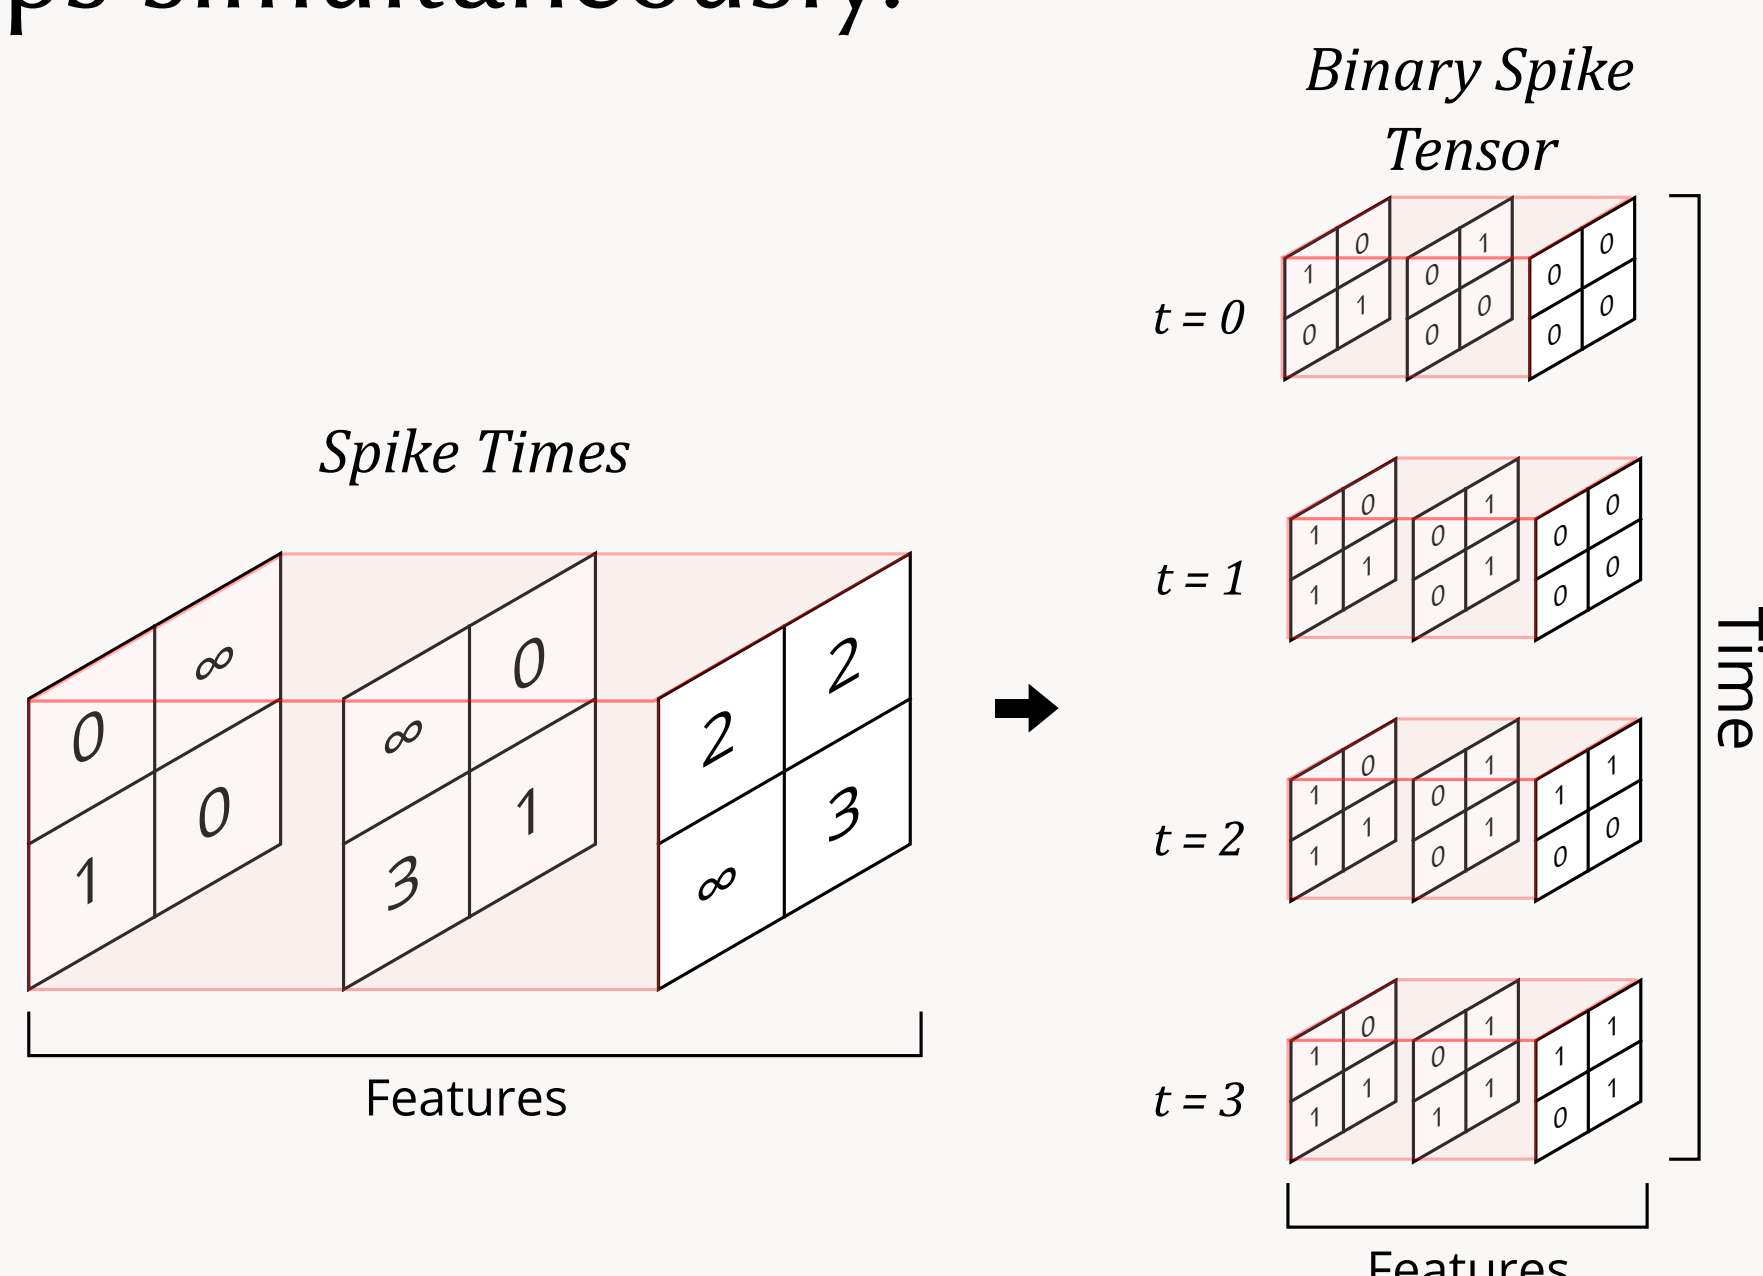
- ▶ Spiking neural networks (SNNs) are energy-efficient and hardware-friendly.
- ▶ SNNs with time-to-first-spike coding and spike-timing-dependent plasticity (STDP) learning rules have done well in various artificial intelligence (AI) tasks.
- ▶ **SpykeTorch** is a simulator for such SNNs that
 - makes use of **PyTorch** highly optimized **tensor operations**, on CPUs and (multi-) GPUs.
 - is fully **compatible** with and **integrated** to **PyTorch**,
 - and it is **user-friendly** and easy to learn specially for deep learning community.

2. Time Dimension

- ▶ SpykeTorch considers an extra dimension in tensors for representing time.
- ▶ SpykeTorch divides all of the spikes of a particular stimulus into a pre-defined number of spike bins, where each bin corresponds to a single time-step.
- ▶ If $T_{f,r,c}$ spike time of the neuron placed at position (r, c) of the feature map f , then the SpykeTorch's corresponding tensor S will be:

$$S[t, f, r, c] = \begin{cases} 0 & t < T_{f,r,c} \\ 1 & \text{otherwise.} \end{cases} \quad (1)$$

- ▶ Keeping spikes in accumulative format increases the memory usage, but let us perform network operations on all time steps simultaneously.



3. Resources

Source Code:

<https://github.com/miladmozafari/SpykeTorch>

Documentation:

<http://cnr1.ut.ac.ir/SpykeTorch/doc/>

Method Paper:

<https://arxiv.org/abs/1903.02440>

4. Tutorial

Network Structure

We implement the network proposed in [1] with seven layers: DoG, Conv1, Pool1, Conv2, Pool2, Conv3, Decision.

Forward Pass

- Override the forward function in `nn.Module`.
- Apply convolutions.
- Apply pooling by functional module.
- Saving data for plasticity (input, output, winners).
- Testing pass is the same but without saving plasticity data.

Plasticity

STDP for the first two Conv layers, and Reward-Modulated STDP (R-STDP) for the last Conv layer.

Input Transform

- Prepare data (MNIST here).
- Use `CacheDataset` to cache the pre-processed data.
- Read image.
- Apply DoG filters.
- Apply intensity-to-latency.

Execution

- Create network instance.
- Run unsupervised learning on Conv1 and Conv2.
- Run reinforcement learning on Conv3.
- After each epoch, evaluate network on testing set.
- test function is similar to `train_r1`, but without calling plasticity functions.

```
1 import torch.nn as nn
2 import SpykeTorch.snn as snn
3 import SpykeTorch.functional as sf
4 class DCSNN(nn.Module):
5     def __init__(self):
6         super(DCSNN, self).__init__()
7
8     # (in_channels, out_channels, kernel_size, weight_mean=0.8, weight_std=0.02)
9     self.conv1 = snn.Convolution(6, 30, 5, 0.8, 0.05)
10    self.conv2 = snn.Convolution(30, 250, 3, 0.8, 0.05)
11    self.conv3 = snn.Convolution(250, 200, 5, 0.8, 0.05)
12
13    # (conv_layer, learning_rate, use_stabilizer=True, lower_bound=0, upper_bound=1)
14    self.stdp1 = snn.STDP(self.conv1, (0.004, -0.003))
15    self.stdp2 = snn.STDP(self.conv2, (0.004, -0.003))
16    self.stdp3 = snn.STDP(self.conv3, (0.004, -0.003), False, 0.2, 0.8)
17    self.anti_stdp3 = snn.STDP(self.conv3, (-0.004, 0.0005), False, 0.2, 0.8)
```

```
1 def forward(self, input, max_layer):
2     input = sf.pad(input, (2,2,2,2))
3     if self.training: #forward pass for train
4         pot = self.conv1(input)
5         spk, pot = sf.fire(pot, 15, True)
6         if max_layer == 1:
7             winners = sf.get_k_winners(pot, 5, 3)
8             self.save_data(input, pot, spk, winners)
9             return spk, pot
10        spk_in = sf.pad(sf.pooling(spk, 2, 2), (1,1,1,1))
11        pot = self.conv2(spk_in)
12        spk, pot = sf.fire(pot, 10, True)
13        if max_layer == 2:
14            winners = sf.get_k_winners(pot, 8, 2)
15            self.save_data(spk_in, pot, spk, winners)
16            return spk, pot
17        spk_in = sf.pad(sf.pooling(spk, 3, 3), (2,2,2,2))
18        pot = self.conv3(spk_in)
19        spk = sf.fire_(pot)
20        winners = sf.get_k_winners(pot, 1)
21        self.save_data(spk_in, pot, spk, winners)
22        output = -1
23        if len(winners) != 0:
24            output = self.decision_map[winners][0][0]
25        return output
26    else:
27        # forward pass for testing process
```

```
1 def stdp(self, layer_idx):
2     if layer_idx == 1:
3         self.stdp1(self.ctx["input_spikes"], self.ctx["potentials"],
4                   -> self.ctx["output_spikes"], self.ctx["winners"])
5     if layer_idx == 2:
6         self.stdp2(self.ctx["input_spikes"], self.ctx["potentials"],
7                   -> self.ctx["output_spikes"], self.ctx["winners"])
8
9     def reward(self):
10        self.stdp3(self.ctx["input_spikes"], self.ctx["potentials"], self.ctx["output_spikes"],
11                  -> self.ctx["winners"])
12
13    def punish(self):
14        self.anti_stdp3(self.ctx["input_spikes"], self.ctx["potentials"],
15                       -> self.ctx["output_spikes"], self.ctx["winners"])
```

```
1 import SpykeTorch.utils as utils
2 import torchvision.transforms as transforms
3 class InputTransform:
4     def __init__(self, filter):
5         self.to_tensor = transforms.ToTensor()
6         self.filter = filter
7         self.temporal_transform = utils.Intensity2Latency(15, to_spike=True)
8     def __call__(self, image):
9         image = self.to_tensor(image) * 255
10        image.unsqueeze_(0)
11        image = self.filter(image)
12        image = sf.local_normalization(image, 8)
13        return self.temporal_transform(image)
14
15    kernels = [ utils.DoGKernel(3,3/9,6/9), utils.DoGKernel(3,6/9,3/9),
16              utils.DoGKernel(7,7/9,14/9), utils.DoGKernel(7,14/9,7/9),
17              utils.DoGKernel(13,13/9,26/9), utils.DoGKernel(13,26/9,13/9)]
18    filter = utils.Filter(kernels, padding = 6, thresholds = 50)
19    transform = InputTransform(filter)
```

```
1 def train_unsupervised(network, data, layer_idx):
2     network.train()
3     for i in range(len(data)):
4         data_in = data[i].cuda() if use_cuda else data[i]
5         network(data_in, layer_idx)
6         network.stdp(layer_idx)
```

```
1 import numpy as np
2 def train_r1(network, data, target):
3     network.train()
4     perf = np.array([0,0,0]) # correct, wrong, silent
5     for i in range(len(data)):
6         data_in = data[i].cuda() if use_cuda else data[i]
7         target_in = target[i].cuda() if use_cuda else target[i]
8         d = network(data_in, 3)
9         if d != -1:
10            if d == target_in:
11                perf[0]+=1
12                network.reward()
13            else:
14                perf[1]+=1
15                network.punish()
16        else:
17            perf[2]+=1
18    return perf/len(data)
```

5. Results

- ▶ Almost the same results as original implementations.
- ▶ 97% [1], and 98.4% [2] on MNIST.

References

1. Mozafari et al 2019, Pattern Recogn, (in press)
2. Kheradpisheh et al 2018, Neural Networks