

## BTP-I PRESENTATION


# REINFORCED LIQUID STATE MACHINE

by Devashri Deulkar (21CS10021)  
under the guidance of Prof. Saumik Bhattacharya

# Table of Contents


- 1. Introduction and Motivation**
  - a. Understanding SNN
    - i. Inspiration -Why
    - ii. what is SNN
- 2. Literature Review**
- 3. Problem Statement**
- 4. LSM - Liquid state machine**
  - a. Inspiration - Why
  - b. What is LSM
- 5. Methodology**
  - a. Implementation of LSM (Maass et al)
  - b. Dataset
- 6. Results**
- 7. Conclusion and Future Works:**
  - a. Understanding RLSM
    - i. DLSM
    - ii. WTA and R-STDP
    - iii. What is RLSM

# Introduction and Motivation:



## A Novel Spiking Neural Network method inspired from neuroscience theory and analytical investigation


- Liquid state Machine is still unexplored potential area of interest



# Inspiration


## 1. Neuroscience:

SNNs are what the brain does. If we want to fully understand the brain we need to understand SNNs.




## 2. Energy Consumption:

Neuromorphic hardware. Low power consumption



## 3. Different Level of Bio plausibility:


### Perceptron: 1st generation



Simple classification-  
Threshold non-linearity

LOW


### SNN: 3rd generation



Neuromorphic computing  
Integrate and fire Non-linearity

BETTER

### ANNs: 2nd generation



Multiclass classification-  
sigmoid, Tanh, Relu

> LOW

# Spiking Neural Networks

What is SNNs?



LIF neuron Model

Synapse Model

---

# Leaky-Integrate and Fire Model


A single spiking neuron is modeled as LIF (Leaky-Integrate and Fire) Model:



leaky  
Integrate  
Fire

We model membrane potential is modeled as capacitor, with Ion channels as pathway with resistance  $R_i$  and its equilibrium potential  $E_i$ .

When given spike  $\Rightarrow I_{\text{ext}} \Rightarrow \text{LIF}$



```

def simulate(self, Iinj, stop=False):
    # Initialize voltage
    v = np.zeros(self.Lt)
    v[0] = self.V_init

    # Set current time course
    Iinj = Iinj * np.ones(self.Lt)

    # If current pulse, set beginning and end to 0
    if stop:
        Iinj[:int(len(Iinj) / 2) - 1000] = 0
        Iinj[int(len(Iinj) / 2) + 1000:] = 0

    # Loop over time
    rec_spikes = [] # record spike times
    spike_train = np.zeros(self.Lt) # spike train
    tr = 0. # the count for refractory duration
    count = 0

    for it in range(self.Lt - 1):

        if tr > 0: # check if in refractory period
            v[it] = self.V_reset # set voltage to reset
            tr = tr - 1 # reduce running counter of refractory period

        elif v[it] >= self.V_th: # if voltage over threshold
            rec_spikes.append(it) # record spike event
            spike_train[it] = 1 # record spike
            count+=1
            v[it] = self.V_reset # reset voltage
            tr = self.tref / self.dt # set refractory time

        # Calculate the increment of the membrane potential
        dv = (-(v[it]-self.E_L) + Iinj[it] / self.g_L)*self.dt/self.tau_m


        # Update the membrane potential
        v[it + 1] = v[it] + dv

    # Get spike times in ms
    rec_spikes = np.array(rec_spikes) * self.dt
    #print("rec_spikes", rec_spikes)
    spike_train = spike_train.astype(int)

    return v, rec_spikes, spike_train

```


## Neuron Model: LIF



(My results)

# synapse model

## MSSM (Modified Stochastic Synaptic Model)



Simulate liquid synapses when given spike trains as input. Working as shown below:

`synapse = MSSM(pars)`

`C, V, E_final, W, Iinj, P = synapse(spike_train, is_inh)`

`synapse` is an instance, `C, V, E_final, W, Iinj` are Ca<sup>2+</sup>, Vesicle concentration, Potential, Weights, current

$$w = C \cdot V \cdot N_t \quad (0)$$


$$P(t) = 1 - \exp(-w), \quad (1)$$

$$\frac{dC}{dt} = \frac{(C_0 - C)}{\tau_C} + \alpha \cdot \sum_i \delta(t - t_i), \quad (2)$$

$$\frac{dV}{dt} = \frac{(V_0 - V)}{\tau_V} - P(t) \cdot \sum_i \delta(t - t_i), \quad (3)$$

$$\frac{dN_t}{dt} = \max\left(0, -\frac{dV}{dt}\right) + \frac{(N_{t0} - N_t)}{\tau_{N_t}}, \quad (4)$$

$$\tau_{E_{psp}} \frac{dE_{psp}}{dt} = -E_{psp} + k_{epsp} \cdot N_t. \quad (5)$$



(My results)

# Spiking Neural Network





Fig. Simple Spiking Neural Network

# Literature review:



# Problem statement

- Applying Reinforcement learning(RL) in Recurrent Spiking Neural Network (SNN)
  - Predictive coding is prevalent in neural microcircuits, minimizing the prediction error through hierarchical feedback loops.
  - R-STDP reward modulated Spike time-dependant plasticity for adaptive, reward-based synaptic updates.
  - Our goal is to make model that performs well with complex tasks such as activity detection real-time videos, real-time fmri dataset etc.




Fig. 1: The architecture of the Reinforced Liquid State Machine.

[Ref. BLSTM](#)

# Liquid State Machine (LSM)

Inspiration

Components


What's LSM?

---


# Inspiration

1. Almost 80% synapses are used in multiple **recurrent loops**, which is a challenge for traditional ANN.
2. Neocortex works well with rapidly changing spatiotemporal inputs and it is **high dimensional dynamical** system.
3. Issue with Turing Machine


Offline computation:



Online computation:




[More detail](#)



computation theory and algorithm design have focused for several decades on offline computations


But our cortical microcircuit does real-time computation (online computation)

Cortical Microcircuits:




## 4. Properties of Cortical microcircuits vs ANNs:

Cortical microcircuit:



- Fix Architecture
- Multiplexing
- Continuously changing spatiotemporal inputs
- Energy Efficient



VS/




- Different Architecture
- Single task
- Turing Machine
- Energy demanding

LSM tries to Mediate the differences...

# Liquid state machine: Overview



## Liquid/reservoir:



- Recurrent SNN
- Encodes input into rich high dimension
$$x^M(t) = (L_M u)(t)$$
$$x^M(t) = \{ x_1^M(t), x_2^M(t), \dots, x_n^M(t) \}$$
- Untrained
  - Already provides High dimensional temporal space
  - Training readout layer is suffice
- Follows Separation Property (SP)
$$u_1 \neq u_2 \Rightarrow (L_M u_1)(t) \neq (L_M u_2)(t)$$
- Follows Fading memory
$$\| (L_M u)(t) - (L_M v)(t) \| \rightarrow 0 \quad \text{if} \quad u(s) \approx v(s) \text{ for } s < t$$


Empirically/theoretically random reservoir/liquid follows SP and Fading memory.

## Liquid Architecture:

```
194 # setting seed
195 np.random.seed(1821)
196 rng = np.random.default_rng(seed=1821)
197
198 # parameters
199 pars = default_pars()
200
201 # Liquid
202 rsnn = SRN(pars)
203
204 # temp synthetic data
205 pre_spike_train_ex = Poisson_generator(pars, rate=1500, n=27, myseed=2020) #1, 3, 4 have spikes
206
207 m, e, i, inpt = rsnn.architecture()
208 liquid_state = rsnn.simulate(pre_spike_train_ex)
209
210 # plot
211 my_neuronNet_plot(m, e, i, inpt)
212
```


```
Number of neurons: 135
Number of excitatory neurons: 108
Number of Input neurons: 16
Index of Input neurons: [ 44  43  65  33  35 108  20 102  18  59  51   9 109 126  80  45]
Neurons in position: 41th neuron is in [6 2 1] position
```

```
Created 1454 MSSM synapses.
synapse[0]: (0, 15)
```



## Readout Layer:

The readout layer is where all task-specific learning happens



A linear readout is enough to extract the desired pattern.

- Normal Layer (AN or SN)
- Linear mapping of high dimension liquid state to output
- Trained
  - W matrix is updated by supervised learning
  - No backprop through time
- Follows Approximation Property (AP)
  - If the liquid satisfies SP and FM, then a linear readout can approximate any causal function of the input (Maass et al., 2002).

$$y(t) = f_M(x^M(t)) = f_M((L_M u)(t))$$





```
1 #@title Readout layer
2 class Net(nn.Module):
3     def __init__(self):
4         super(Net, self).__init__()
5         self.fc1 = nn.Linear(135,10)
6
7     def forward(self, x):
8         x = self.fc1(x)
9         x = F.relu(x)
10    return x
11
12
13
14 class ReadoutLayer:
15     def init(self, pars):
16         self.pars = pars
17         self.net = Net()
18         self.P = 10
19
20     def simulate(self, liquid_state):
21         # liquid state N neurons getting added to 10 readout neurons
22         input = np.mean(liquid_state, axis=1)
23         input = torch.from_numpy(input).float().unsqueeze(0) # (1, 135)
24         output = self.net(input) # (1, 10)
25
26         return output
27
28 pars = default_pars()
29 readout = ReadoutLayer()
30 readout.init(pars)
31 output = readout.simulate(liquid_state)
32 print("train:", output.shape)
33
```

Readout Layer

## Input: Spike Encoding + Dataset

Way to converting input signals to spike trains ([MNIST](#))


- Neuromorphic-MNIST (NMNIST) is a spiking version of the classic MNIST dataset
- Event-based encoding : Generates a spike only when the pixel intensity changes
- It mimics biological eye-movements (Neurons emit spikes only when input signals change)



Dataset:


- N-MNIST converts the static 28×28 MNIST digits into spike trains by **physically moving a camera (ATIS)** in three small **microsaccades** while viewing the images on a monitor
- Each pixel emits an ON/OFF spike on intensity change due to motion  
ON = brightness ↑, OFF = brightness ↓

Raster graph from the N-MNIST dataset:



From Local code, 2312 neuron spike trains  
for digit 3

# Putting all together: LSM



Input: encoded spike trains

Liquid: It is recurrent SNN, with few special properties.

Readout: trained output layer

# Next Thing to do: RLSM pipeline




Fig. 1: The architecture of the Reinforced Liquid State Machine.


## Features:

- We use multiple liquid layer, Deep Liquid State Machine
- How do we connect liquid layer?
  - Winner-Takes-All layer, with STDP (spike-time dependant plasticity) synapse model.
  - Trainable synapses
- Reward-system: reward modulated stdp, again mimicking neurobiology of neuromodulators

Predictive coding: Through hierarchical feedback loop, minimising predictive error.

# Upgraded Synaptic Model: STDP

STDP is governed by the relative timing of these spikes, leading to either long-term potentiation (LTP) or long-term depression (LTD) of the synapse




O1 (Postsynaptic) spike just after I3 (Presynaptic) correlated

Weight update happens according to above rule, unsupervised

There is also concept Homeostasis, which I will not cover here...

Strength of synapse (W) depends on ISI between spikes




I2 (Presynaptic) spike after O1 (Postsynaptic) uncorrelated

Pre-synaptic trace:


$$\frac{dx_{\text{pre}}}{dt} = -\frac{x_{\text{pre}}}{\tau_{\text{pre}}} + \sum_f \delta(t - t^f)$$

$$\Delta w_{\text{local}} = \eta(x_{\text{pre}} - x_{\text{tar}})(w_{\text{max}} - w)^{\mu}$$
$$w_{\text{local}} = \frac{w_{i,j}}{\sum_{j=1}^N w_{i,j}} \cdot \alpha$$

Biphasic STDP: (My work)



# WTA- Layer (R-STDP)



```

28 # Initialization
29 N, M = W.shape
30 dt = pars['dt']
31 V_th = pars['V_vth']
32 V_rst = pars['V_reset']
33 spikes = np.zeros(M)
34
35
36 # 1. Excitatory drive from liquid -> WTA
37 I_exc = ls @ W # (M_r)
38
39 # 2. Top-k selection
40 winners = np.argpartition(I_exc, -k)[-k:]
41
42 # 3. Global inhibition current (shared)
43 I_inh = g_inh * np.sum(I_exc[winners])
44
45 # 4. Apply inhibition: winners keep full drive, others get clamped down
46 I_exc_post = I_exc - I_inh
47 I_exc_post[winners] = I_exc[winners]
48 I_exc_post = np.clip(I_exc_post, 0.0, None) # no negative currents
49
50 # 5. Single non-leaky LIF integration step: (C=1) V(0)=0
51 V = I_exc_post * dt # integrate over one dt
52
53 widx = V >= V_th # which neurons cross threshold?
54 #print("widx:",widx)
55
56 spikes[widx] = 1.0 # spike if above threshold
57 # 6. Reset all spiking neurons
58 V[widx] = V_rst # reset to V_reset
59
60 return winners, spikes
61

```



$$\frac{de_{\text{trace}}}{dt} = -\frac{e_{\text{trace}}}{\tau_{\text{trace}}} + \Delta w_{\text{local}}$$

$$\Delta w_{\text{feedback}} = r \cdot e_{\text{trace}}$$

1. **Eligibility trace** acts as a short-term memory — it decays over time but is refreshed whenever a local synaptic change occurs.
2. This trace stores **when and how much** a synapse was active, waiting for a later reward signal.
3. When a **reward** (  $r$  ) arrives, the actual weight change is computed as linking past activity to current reward — the essence of **reward-modulated plasticity**.

## Section 5: WTA and R-STDP

1. Presynaptic trace created by incoming presynaptic spikes:  $\frac{dx_{\text{pre}}}{dt} = \frac{x_{\text{pre}}}{\tau_{\text{pre}}} + \sum_f \delta(t - t^f)$

2. When the postsynaptic neuron fires, the weight of the synapse is updated based on  $x_{\text{pre}}$  and  $x_{\text{tar}}$ :

$$\Delta w_{\text{local}} = \eta(x_{\text{pre}} - x_{\text{tar}})(w_{\text{max}} \cdot w)^\mu$$

$$3. \text{Synaptic scaling: } w_{\text{local}} = \frac{w_{ij}}{\sum_j^N w_{ij}} \cdot \alpha$$




4. To enable faster learning by crediting actions that contributed to future rewards, an eligibility trace  $e_{\text{trace}}$  is introduced. This trace acts as a bridge linking the local synaptic updates:  $\frac{de_{\text{trace}}}{dt} = \frac{e_{\text{trace}}}{\tau_{\text{trace}}} + \Delta w_{\text{local}}$

5. Overall weight update:  $\Delta w_{\text{feedback}} = r \cdot e_{\text{trace}}$

```

180
181 pre_spike_train_ex = Poisson_generator(pars, rate=10, n=5, myseed=2020)
182 t_pre = generate_presynaptic_trace(pars, pre_spike_train_ex)
183 my_example_trace(pre_spike_train_ex, pars, t_pre)
184
185 dW,W = Delta_W(pars, t_pre)
186

```



## New Features Summary:

### 1. Local temporal learning:


STDP captures fine-grained spike-timing correlations between liquid (pre-) and WTA (post-) neurons — enabling each WTA neuron to become selective to specific temporal input patterns.

### 2. Global performance feedback:


The reward signal adds a third factor that reinforces learning— linking local plasticity to task outcome.

### 3. Competition and specialization:

The WTA mechanism ensures only the K fastest (strongest) neurons fire; R-STDP helps to specialize on distinct input features.




# RLSM:




## references:

1. Krenzer, D., and Bogdan, M. (2025). Reinforced Liquid State Machines—New Training Strategies for Spiking Neural Networks Based on Reinforcements. *Frontiers in Computational Neuroscience*, 19, 1569374. <https://doi.org/10.3389/fncom.2025.1569374>
2. Krenzer, D., and Bogdan, M. (2025). The Reinforced Liquid State Machine: A New Training Architecture for Spiking Neural Networks. *ESANN 2025 Proceedings*, 699–704. Bruges, Belgium.
3. Maass, W., Natschläger, T., and Markram, H. (2002). Real-Time Computing Without Stable States: A New Framework for Neural Computation Based on Perturbations. *Neural Computation*, 14(11), 2531–2560. <https://doi.org/10.1162/089976602760407955>
4. Soures, N., and Kudithipudi, D. (2019). Deep Liquid State Machines with Neural Plasticity for Video Activity Recognition. *Frontiers in Neuroscience*, 13, 686.
5. Frémaux, N., and Gerstner, W. (2016). Neuromodulated Spike-Timing-Dependent Plasticity, and Theory of Three-Factor Learning Rules. *Frontiers in Neural Circuits*, 9, 85.
6. El-Laithy, K., and Bogdan, M. (2010). Predicting Spike Timing of a Thalamic Neuron Using a Stochastic Synapse. *European Symposium on Artificial Neural Networks (ESANN)*, 358–363.
7. Ellatihi, K., and Bogdan, M. (2017). Enhancements on the Modified Stochastic Synaptic Model: The Functional Heterogeneity. *ICANN 2017, LNCS 10613*, 389–396.
8. Oster, M., Douglas, R., and Liu, S.-C. (2009). Computation with spikes in a winner-take-all network. *Neural Comput.* 21, 2437–2465. <https://doi.org/10.1162/neco.2009.07-08-829>
9. Maass, W. (1997). Networks of spiking neurons: the third generation of neural network models. *Neural Netw.* 10, 1659–1671. doi: 10.1016/S0893-6080(97)00011-7
10. Maass, W. (2011). Liquid state machines: motivation, theory, and applications. *Comput. Context* 2011, 275–296. doi: 10.1142/9781848162778\_0008
11. Spike times make sense VanRullen, Rufin et al. *Trends in Neurosciences*, Volume 28, Issue 1, 1 – 4 <https://doi.org/10.1016/j.tins.2004.10.010>



**Thank You!**

---



[back](#)

## NMNIST

Generates a spike only when the pixel intensity changes i.e., when there is new information.

NMNIST is created from event-based spike encoding

### **Method:**

#### **1. Event based Camera**

ATIS (Asynchronous Time-based Image Sensor)

-The camera/sensor was mounted on a **pan-tilt** unit facing a monitor displaying static MNIST digits.

#### **2. Monitor**

-screen displaying one MNIST digit at given time

#### **3. Motion/ saccades**

- performs three tiny saccades while looking at a static MNIST digit on a monitor

#### **4. Duration**

-each digit recorded for  $\approx 300$  ms (three saccades of 100 ms each)



FIGURE 2 | (A) A picture of the ATIS mounted on the pan tilt unit used in the conversion system. (B) The ATIS placed viewing the LCD monitor.

