

Biologically Plausible Variational Policy Gradient with Spiking Recurrent Winner-Take-All Networks

Zhile Yang¹

sczy@leeds.ac.uk

Shangqi Guo^{†2}

shangqi_guo@foxmail.com

Ying Fang³

fy20@fjnu.edu.cn

Jian Liu^{†1}

j.liu9@leeds.ac.uk

¹ School of Computing,
University of Leeds,
Leeds, UK

² Department of Precision Instrument
and Department of Automation,
Tsinghua University,
Beijing, China

³ College of Computer and Cyber
Security,
Fujian Normal University,
Fujian, China

Abstract

One stream of reinforcement learning research is exploring biologically plausible models and algorithms to simulate biological intelligence and fit neuromorphic hardware. Among them, reward-modulated spike-timing-dependent plasticity (R-STDP) is a recent branch with good potential in energy efficiency. However, current R-STDP methods rely on heuristic designs of local learning rules, thus requiring task-specific expert knowledge. In this paper, we consider a spiking recurrent winner-take-all network, and propose a new R-STDP method, spiking variational policy gradient (SVPG), whose local learning rules are derived from the global policy gradient and thus eliminate the need for heuristic designs. In experiments of MNIST classification and Gym InvertedPendulum, our SVPG achieves good training performance, and also presents better robustness to various kinds of noises than conventional methods. Code can be found at https://github.com/yz1c080733/BMVC2022_SVPG.

1 Introduction

Reinforcement learning (RL) with artificial neural networks (ANNs) has been applied to various scenarios [4, 7]. However, the adaptability, robustness, learning speed, etc. are still not satisfying. Recently some research build *biologically plausible* learning systems similar to humans' brains to achieve human-level performances [1, 5, 31]. In these brain-inspired methods, neural spiking signals are the key elements, which have demonstrated a great advantage for visual coding [9, 34]. The enriched analysis of neural signals using RL also provide insights for explaining the formation of capabilities of brains in the field of computational neuroscience [28].

One mainstream biologically plausible method is spiking neural networks (SNNs). Current SNN RL methods can be divided into three categories: ANN-to-SNN conversion ANN2SNN [22], heuristically approximated backpropagation [3, 32], and modulated STDP (e.g., R-STDP) [8, 35]. Compared with the first two, R-STDP has the advantages of being biologically plausible in both training and testing, and being energy efficient. Recent R-STDP methods consider simulating the activity patterns of biological neurons, which are characterized by types of interactions between neurons, strengths of modulation, and modulation timings [8], and have outperformed other state-of-the-art algorithms in many tasks in terms of control capabilities and energy efficiency [1, 5, 31]. However, one drawback of them is their reliance on heuristic designs of local learning rules, since there is a gap between the local learning rules and the overall RL optimization target. This makes it inconvenient to apply R-STDP to RL tasks.

To bridge the gap, we adopt the variational inference method to study spiking neural networks, and propose a novel method, named spiking variational policy gradient (SVPG). It has been shown that variational inference can decompose the global target of a pattern generation or classification task into local learning rules [11, 13, 23]. Besides, there is evidence that variational inference can be made biologically plausible using specific network structures [2, 13, 23]. As far as we know, this paper is the first to investigate variational inference for R-STDP.

To establish the theoretical relationship between local R-STDP rules and the global target, we propose a spiking neural network recurrently constructed of winner-take-all (WTA) circuits [11, 33] (named *RWTA network*), together with an energy-based policy formulation. Under these settings, we prove that the fixed point of the RWTA network equals the inference result of the RL policy distribution, and that the R-STDP mechanism is a first-order approximation of the policy gradient. These bridge the R-STDP learning rule with the overall target. We apply SVPG to two typical RL tasks, including reward-based MNIST classification and Gym InvertedPendulum. SVPG successfully solved the tasks, achieving similar performance as a conventional ANN-based method. We also tested the robustness of the trained policies against various kinds of input noise [22], network parameter noise [16], and environment variation [19]. Results show SVPG has better robustness to the tested noises than popular biologically plausible RL methods.

2 Preliminaries and Notations

We begin with notations and formulations of the policy gradient algorithm, the SNN model, and the R-STDP framework.

2.1 Policy Gradient

In RL, Markov decision processes (MDPs) are commonly used for task formulation [29]. Here we adopt the notations from [29] and use tuple $\langle \mathcal{S}, \mathcal{A}, P, R, \gamma \rangle$ to denote an MDP. The training objective is to find a policy $\pi : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ that maximizes the expected return $J = \mathbb{E}_{\tau \sim \pi} \sum_{t=0}^{T-1} \gamma^t r_t$, where τ is the sampled trajectory (denoted as $\langle s_0, a_0, r_0, s_1, a_1, \dots, r_{T-1}, s_T \rangle$) and T is the total length of an episode.

Policy gradient is a widely used branch of algorithms in RL and has many popular derivatives [17, 20, 26]. A fundamental form of policy gradient is the REINFORCE algorithm,

which optimizes the expected return through ascending the following gradient [29]:

$$\nabla_{\theta} J(\pi_{\theta}) = \mathbb{E}_{\tau \sim \pi_{\theta}} \left[\sum_{t=0}^{T-1} \gamma^t r_t \sum_{k=0}^{T-1} \nabla_{\theta} \log \pi_{\theta}(a_k | s_k) \right], \quad (1)$$

where θ is the parameters of the policy π . In this paper we propose a new model for π_{θ} , so the key point is to derive the gradient $\nabla_{\theta} \log \pi_{\theta}(a_k | s_k)$.

2.2 Spiking Neurons

Various neuron models have been proposed in the field of SNN. Here we consider the leaky integrate-and-fire (LIF) spike response model [10]. We consider that all spikes are binary and happen at discrete time steps. At each spike time step l , the firing probability ρ of a neuron is determined by its membrane potential $u(l)$:

$$\rho(l) = \exp\{u(l) - u_{\text{th}}\}, \quad (2)$$

where u_{th} is a threshold voltage. Note that l is different from the RL step t . We consider that each t -step corresponds to a certain number of spike time steps. The membrane potential $u(l)$ is determined by the spike train S from the presynaptic neurons [10]:

$$u(l) = \sum_{j \in N(\cdot)} w_j \int_0^{\infty} \kappa(y) S_j(l-y) dy + b, \quad (3)$$

where $N(\cdot)$ denotes the set of presynaptic neurons of the considered neuron, w_j is the synapse weight, κ is the excitatory postsynaptic potential, S_j is the spike train from neuron j , and b is the self-excitation voltage of the considered neuron.

R-STDP models the learning of synapse weights given an external reward and local firing activities. We use w_{ij} to denote the weight of synapse between presynaptic neuron i and postsynaptic neuron j . The R-STDP learning takes the form of $\Delta w_{ij} = R(l) \cdot \text{STDP}(l)$, where $R(l)$ is the external reward signal and $\text{STDP}(l)$ is a coefficient determined by the STDP learning rule in the following form:

$$\text{STDP}(l) = S_j(l) \left[W_{\text{pre}} + \int_0^{\infty} A_+ W_+(y) S_i(l-y) dy \right] + S_i(l) \left[W_{\text{post}} + \int_0^{\infty} A_- W_-(y) S_j(l-y) dy \right], \quad (4)$$

where W_{pre} and W_{post} are respectively the constants of presynaptic and postsynaptic activity, A_+ and A_- characterize the extent to which synaptic changes depend on the current synaptic weights. W_+ and W_- are respectively the time windows of the long-term potentiation (LTP) and the long-term depression (LTD) processes, and they satisfy $\int_0^{\infty} W(l) dl = 1$. With these notations, each tuple $\langle W_{\text{pre}}, W_{\text{post}}, A_+(w_{ij}), A_-(w_{ij}) \rangle$ defines an STDP learning rule.

3 Method

In this paper, we develop a spiking-based variational policy gradient (SVPG) method with R-STDP and an RWTA network. In this section, we first introduce the formulation of the policy and the network. Then we derive the parts of the method respectively for inference and optimization.

3.1 Network Structure

The key demand in network design is the capability for probabilistic inference so that the relationship between local optimization and global target can be built. It has been shown that WTA circuits with the STDP learning rule can perform probabilistic inference and learning [11]. Here we consider rate coding – each input (state) neuron represents a normalized state element; each output (action) state represents the probability of corresponding action being selected; each hidden WTA circuit encodes the probability distribution of a hidden variable.

The RWTA network structure is shown in Figure 1. Red rectangles represent WTA circuits. State neurons are denoted as s_i ($i = 1, \dots, d_s$); action neurons are denoted with a_i ($i = 1, \dots, d_a$); the j -th neuron in the i -th hidden circuit is denoted with h_{ij} ($i = 1, \dots, n_h$, $j = 1, \dots, d_h$). Here d_s , d_h , and d_a are the corresponding dimensions, and n_h is the number of hidden groups. “FC” means “fully connected”.

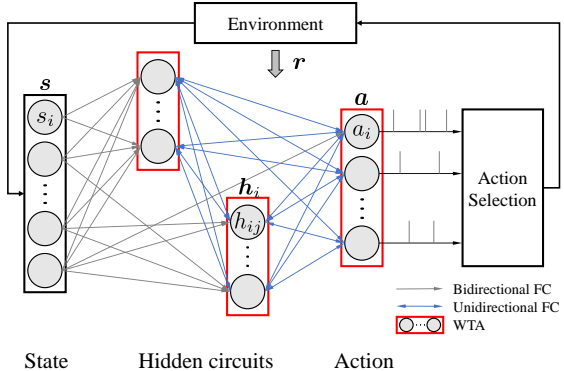


Figure 1: RWTA Network structure.

Note that similar circuits have been investigated in [11]. However, their structure is constrained to trees of circuits, while ours allows arbitrary connections and so covers any recurrent structures (regardless of layer number and layer size).

At each spike time step, each neuron has two properties: firing probability $q \in [0, 1]$ and a binary firing status $v \in \{0, 1\}$. For example of notations, the properties of neuron h_{ij} are $q_{h_{ij}}$ and $v_{h_{ij}}$. We use bold symbols to denote vectors. For example, the firing probability of WTA circuit h_i is $\mathbf{q}_{h_i} := [q_{h_{i1}}, \dots, q_{h_{id_h}}]^T$, where $[\cdot]^T$ is the transpose operation. For conciseness, we denote $\mathbf{q}_h := [\mathbf{q}_{h_1}^T, \dots, \mathbf{q}_{h_{n_h}}^T]^T$, $\mathbf{v}_h := [v_{h_1}^T, \dots, v_{h_{n_h}}^T]^T$, and then $\mathbf{q} = [\mathbf{q}_h^T, \mathbf{q}_a^T, \mathbf{q}_s^T]^T$, $\mathbf{v} = [v_h^T, v_a^T, v_s^T]^T$. The total number of neurons is $N = n_h d_h + d_a + d_s$. The synapse weights between all the neurons are altogether parameterized as $\mathbf{W} \in \mathbb{R}^{N \times N}$. The self-activation parameters are $\mathbf{b} \in \mathbb{R}^N$. The aforementioned parameters of a policy, θ , is thus the combination of $\theta = \langle \mathbf{W}, \mathbf{b} \rangle$. Note that different network structures can be represented by constraining certain parts of \mathbf{W} to zero.

3.2 Policy Inference

We consider the RL policy to be a probability distribution over the action space, and define it using an energy function $E(\mathbf{v})$:

$$\pi(\mathbf{v}_a | s) = \sum_{\mathbf{v}_h} p(\mathbf{v}_a, \mathbf{v}_h | s), \quad (5)$$

$$p(\mathbf{v}_a, \mathbf{v}_h | s) := \frac{1}{Z(s)} \exp\{E(\mathbf{v})\}, \quad E(\mathbf{v}) := \mathbf{v}^T \mathbf{W} \mathbf{v} + \mathbf{b}^T \mathbf{v}, \quad (6)$$

where $Z(s)$ is the normalization $Z(s) = \sum_{\mathbf{v}'_h, \mathbf{v}'_a} \exp\{E(\mathbf{v}')\}$. Given a state s , first the encoding \mathbf{v}_s is generated (see experiment section for details), and then the action distribution

is calculated as the marginal distribution of \mathbf{v}_a . Although the energy function is linear, the normalization operation enables representation of complex distributions [12].

This policy representation is computable in principle. However, when there is a large number of hidden neurons, the calculation can be intractable. To address this problem, we adopt mean-field inference to derive a tractable method. We will then show that the mean-field inference can be realized with the RWTA network.

Policy Mean Field Inference

First, we use a variational distribution $\hat{p}(\mathbf{v}_a, \mathbf{v}_h | s)$ to approximate $p(\mathbf{v}_a, \mathbf{v}_h | s)$, and assume that the states of all circuits are independent from each other. This allows decomposition of \hat{p} : $\hat{p}(\mathbf{v}_a, \mathbf{v}_h | s) := \hat{p}(\mathbf{v}_a | s) \hat{p}(\mathbf{v}_{h_1} | s) \cdots \hat{p}(\mathbf{v}_{h_{n_h}} | s)$, where $\hat{p}(\mathbf{v}_{h_1} | s) := \mathbf{q}_{h_1}^T \mathbf{v}_{h_1}, \dots, \hat{p}(\mathbf{v}_a | s) = \mathbf{q}_a^T \mathbf{v}_a$.

Second, we use the KL divergence to measure the difference between \hat{p} and p : $D_{\text{KL}}(s) \doteq D_{\text{KL}}[\hat{p}(\mathbf{v}_a, \mathbf{v}_h | s) \| p(\mathbf{v}_a, \mathbf{v}_h | s)]$. By letting $\partial D_{\text{KL}}(s) / \partial q_i = 0$ for each q_i (the i -th element of vector \mathbf{q}), we get the mean-field inference function [30]:

$$q_i = \frac{1}{Z(q_{G(i)})} \exp\{\mathbf{w}_{\text{row},i}^T \mathbf{q} + \mathbf{w}_{\text{col},i}^T \mathbf{q} + b_i\}, \quad (7)$$

where $Z(q_{G(i)}) = \sum_{j \in G(i)} \exp\{\mathbf{w}_{\text{row},j}^T \mathbf{q} + \mathbf{w}_{\text{col},j}^T \mathbf{q} + b_j\}$, $i = 1, \dots, (n_h d_h + d_a)$, $G(i)$ is the set of indices of the neurons in the same circuit as neuron i , and $\mathbf{w}_{\text{row},i}$ and $\mathbf{w}_{\text{col},i}$ are respectively the i -th row and column of matrix \mathbf{W} (in the shape of a column vector), which corresponds to the synapses connected to neuron i . b_i is the i -th element in vector \mathbf{b} .

So far we derived an approximate policy representation – by finding a \mathbf{q} that satisfies Eq.(7), we can get the \mathbf{q}_a within it, which is the policy distribution we need. Notice that Eq.(7) can be seen as an iteration process by regarding the \mathbf{q} on the right side as constant. So an iterative method can be used to find \mathbf{q} – initialize \mathbf{q} with random numbers, then repeat updating it with Eq.(7) until numeric convergence. Although there has not been a theoretical guarantee of convergence, the iteration converges in our experiments.

Policy Inference Based on RWTA Network

Now we show that this iterative method for inference can be realized with the RWTA network. Since the inhibitory neuron in a WTA circuit produces an overall firing rate $\hat{\rho}$ [11], a natural link is to let firing probabilities encode $\rho_i(l) = \hat{\rho} q_i$. This transforms Eq.(7) into (note that $1/Z(q_{G(i)}) = \exp\{-\log[Z(q_{G(i)})]\}$):

$$\rho_i = \hat{\rho} \exp\{\mathbf{w}_{\text{row},i}^T \mathbf{q} + \mathbf{w}_{\text{col},i}^T \mathbf{q} + b_i - \log \sum_{j \in G(i)} \exp\{\mathbf{w}_{\text{row},j}^T \mathbf{q} + \mathbf{w}_{\text{col},j}^T \mathbf{q} + b_j\}\}. \quad (8)$$

By replacing w_j with the synaptic weights w_{ij} in Eq.(3), and designing $\kappa(y)$ such that $\int_0^\infty \kappa(y) dy = 1/\hat{\rho}$, we can use the potential $u(l)$ to replace the firing probability q . This results in the following spike-based inference function:

$$\begin{aligned} \rho_i(l) &= \hat{\rho} \exp\{u_i(l) - \log \sum_{j \in G(i)} \exp(u_j(l))\}, \\ u_i(l) &= \sum_{j \in N(i)} w_{ij} \int_0^\infty \kappa(y) S_{ij}(l-y) dy + b_i. \end{aligned} \quad (9)$$

This function is biologically plausible since the firing probability of each neuron is determined by its potential, which is determined by the spike trains from neighboring neurons. The only difference from the LIF model Eq.(2, 3) is that the potential threshold u_{th} is replaced by a logarithm item, which can be made by introducing lateral inhibition neurons [14].

3.3 Policy Optimization

As has been introduced in section 2.1, the key to policy gradient is to derive the differential of $\log \pi_\theta(a_t|s_t)$. This means to derive the differential of each q_i regarding each parameter w_{jk} and b_j , using Eq.(7). In this section, we first derive the precise differential, then show how R-STDP can represent its first-order approximation.

Precise Differential

Theorem 1. *The precise differential of q_{ha} with reference to the synaptic weight w_{jk} and the self-activation parameter b_j is as follows:*

$$\begin{aligned} \frac{\partial q_{ha}}{\partial w_{jk}} &= \mathbf{M}(\mathbf{U}_{jk} + \mathbf{U}_{kj})\mathbf{q} + \mathbf{M}(\mathbf{W} + \mathbf{W}^T) \frac{\partial \mathbf{q}}{\partial w_{jk}}, & \frac{\partial q_{ha}}{\partial b_j} &= \mathbf{M}\mathbf{b} + \mathbf{M}(\mathbf{W} + \mathbf{W}^T) \frac{\partial \mathbf{q}}{\partial b_j}, \\ \mathbf{M} &= \text{diag}(\mathbf{q}_{ha})[-\mathbf{G}_{ha}\text{diag}(\mathbf{q}) + \mathbf{D}_{\text{sel}}], \end{aligned} \quad (10)$$

where \mathbf{G}_{ha} is a $(n_h d_h + d_a) \times N$ logical matrix of which 1 elements indicate the two neurons (column index and row index) are in the same circuit, \mathbf{D}_{sel} is a logical matrix that selects the first $(n_h d_h + d_a)$ elements in a vector with length N , i.e., $\mathbf{D}_{\text{sel}} = \begin{bmatrix} \mathbf{I}_{(n_h d_h + d_a)} & \mathbf{O}_{(n_h d_h + d_a) \times d_s} \end{bmatrix}$, and \mathbf{U}_{jk} is a $N \times N$ logical matrix with only the jk -th element being 1.

The proof is given in Appendix A.1. Theorem 1 reveals that $\partial q_{ha}/\partial w_{jk}$ can be calculated by solving the matrix equations in Eq.(10). Specifically, the calculation includes the matrix pseudo-inverse of $\mathbf{M}(\mathbf{W} + \mathbf{W}^T)$, whose shape is $(n_h d_h + d_a) \times N$. So the calculation complexity is more than $O((n_h d_h + d_a)^3)$, which can be intractable for large neuron numbers. More importantly, this solution does not reveal the link between local parameter updates and the global RL target, and is not biologically plausible.

First-Order Approximation of the Differential

Below we derive a first-order approximation to the differential, and show how it can be implemented by an R-STDP learning rule. As will be shown in the experiment section, this approximation does not hinder the network from solving complex RL tasks.

The main idea for the approximation is to treat the \mathbf{q} on the right side of the inference function Eq.(7) as a constant with reference to w_{jk} and b_j . By doing so, the differential only considers the last iteration in the inference process, where the status of each neuron is only affected by its neighboring neurons. Thus the differential only requires the properties of neighboring neurons, which enables the link to local learning rules like R-STDP.

To get the policy gradient $\nabla \log \pi(a_t|s_t)$, $\partial \log(q_i)/\partial w_{jk}$ and $\partial \log(q_i)/\partial b_j$ are required. We present the results in the following Theorem 2. The proof is given in Appendix A.2.

Theorem 2. *The approximate differentials of firing rate q_i with respect to w_{jk} and b_j in the matrix form are:*

$$\begin{aligned} \frac{\partial \log(q_i)}{\partial \mathbf{W}} &= (\mathbf{U}_{i:}\text{diag}(\mathbf{q}) + \text{diag}(\mathbf{q})\mathbf{U}_{:i}) - \text{diag}(\mathbf{q})(\mathbf{U}_{G(i):} + \mathbf{U}_{:G(i)})\text{diag}(\mathbf{q}), \\ \frac{\partial \log(q_i)}{\partial \mathbf{b}} &= \mathbf{u}_i - \text{diag}(\mathbf{q})\mathbf{u}_{G(i)}, \end{aligned} \quad (11)$$

where $i \in \{1, \dots, (n_h d_h + d_a)\}$, \mathbf{U} is a $N \times N$ logical matrix and \mathbf{u} is a length- N logical vector, whose subscripts indicates the positions of elements with value 1. $G(i)$ is the set of indices of neurons in the same circuit as neuron i ; “:” means the entire row/column.

With this result, when the firing state of the network is \mathbf{v} , the corresponding REINFORCE policy gradient is

$$\nabla J(\pi) = \sum_t \gamma^t r_t [\sum_{i=1}^{n_h} \mathbf{v}_{h_i}^T \nabla(\log \mathbf{q}_{h_i}) + \mathbf{v}_a^T \nabla(\log \mathbf{q}_a)], \quad (12)$$

where $\nabla \log \mathbf{q}_{h_i}$ and $\nabla \log \mathbf{q}_a$ are respectively the vectors of $\nabla \log q_{h_{ij}}$ and $\nabla \log q_{a_i}$. So far we can learn the RWTA network by the REINFORCE algorithm to optimize the RL target.

Policy Optimization Based on R-STDP

Notice that in the policy gradient Eq.(11), the terms of w only involve pre- and post-synaptic neurons, and the terms of b only involve connected neurons, which is similar to the R-STDP rules. Now we show how it can be linked to R-STDP.

Recall an R-STDP rule is determined by $\langle W_{\text{pre}}, W_{\text{post}}, A_+(w_{ij}), A_-(w_{ij}) \rangle$. By assuming a large enough spike step number in each RL step, we have that the averaged number of spikes in the spiking train S_i approximates the firing probability ρ_i . That is $\mathbb{E}[S_i(l)] = \rho_i$ and $\mathbb{E}[\int_0^\infty A_+(w_{ij}) S_i(l-y) dy] = A_+(w_{ij}) \rho_i$. By performing the following settings of STDP (element-wise notation for the synapse between neuron i and j):

$$\langle W_{\text{pre}}, W_{\text{post}}, A_+(w_{ij}), A_-(w_{ij}) \rangle = \langle v_i, v_j, -1/\hat{\rho}, -1/\hat{\rho} \rangle, \quad (13)$$

we get the following R-STDP rule:

$$\mathbb{E}[R(l)\text{STDP}(l)] = R[\rho_j(v_i - \rho_i/\hat{\rho}) + \rho_i(v_j - \rho_j/\hat{\rho})]. \quad (14)$$

Then, for the policy gradient we have just derived (Eq.11, 12), there is

$$\frac{\partial J(\pi)}{\partial w_{ij}} = \sum_t \gamma^t r_t [q_i(v_j - q_j) + q_j(v_i - q_i)], \quad (15)$$

which is the same as Eq.(14) (despite a constant $\hat{\rho}$). This shows that an R-STDP characterized by Eq.(13) can represent the policy gradient on synapse weights. As for the self-activation parameter b , a similar transformation exists as $\Delta b(l) = v_i + \int_0^\infty -(1/\hat{\rho}) S_i(l-y) dy$.

So far we get a variational policy gradient method where inference and optimization are implemented with the spiking RWTA network and an R-STDP rule. We name it *SVPG*.


Note that although our method is derived from the REINFORCE algorithm, it can be adapted to many other policy gradient algorithms. Some implementation details, including adaptation to Actor-Critic algorithms, are provided in Appendix B.

4 Experiments

4.1 Task Design

We selected two commonly used RL tasks: reward-based MNIST image classification (*MNIST*) and Gym InvertedPendulum (*GymIP*). *MNIST* is a one-step task featured by a high dimensional input (28×28). *GymIP* is a widely-used task in RL researches [6, 18, 21] featured by long-term decision sequences. Details of the task goals, reward settings, state processing, etc. are in Appendix C.1.

It has been shown that, compared to ANNs, SNNs trained with ANN2SNN methods can have better robustness to input and synapse weight noises [16, 22], and that SNNs trained



	SVPG	BP	BPTT	EP	ANN2SNN
MNIST	0.926± 0.001	0.933± 0.062	0.898± 0.067	0.971± 0.001	0.933± 0.062
GymIP	199.87 ± 0.27	199.95 ± 0.13	199.96 ± 0.12	N/A	190.79 ± 27.64

Figure 2: Two types of changes to GymIP. Table 1: Zero-noise performances (mean±std).

with BPTT can be more resilient to adversarial attacks [27]. Therefore it is worth checking whether SVPG brings better robustness than ANNs and other SNN learning methods. We test three types of disturbance – input noise, network parameter noise, and GymIP environmental variation. **1) Input noise** is independently added to each dimension of state observations. For MNIST, Gaussian, salt, salt&pepper, and Gaussian&salt noises are considered. For GymIP, Gaussian and uniform noises are considered. Some illustrations are in Appendix D.1. **2) Network parameter noise** is independently added to each of the learnable parameters in the policy networks. Gaussian and uniform noises are considered. **3) Environment variations in GymIP.** For GymIP, the pendulum’s length and thickness are modified during testing. In training the setting is $\langle \text{length}=1.5, \text{thickness}=0.05 \rangle$. In testing we change the length in the range of $[0.16, 4.9]$ and thickness of $[0.02, 0.30]$. Figure 2 illustrates these variations.


4.2 Compared Methods

We select four representative learning methods. 1) *ANN2SNN* with the methods from [24]. 2) *BP* [15] with the ReLU function for hidden layers, which is a conventional baseline ANN model using multilayer perceptrons architecture learning with error back-propagation. This serves as a usual approach in RL. 3) Fast sigmoid *BPTT* (backpropagation through time) from [36]. 4) *EP* (Equilibrium Propagation) [25], which is a biologically plausible algorithm with local learning rules, though not designed for SNNs. Note the EP is designed for supervised learning, so we only test it in MNIST using ground truths. For all these methods, the number of hidden layers is set to 1; the optimizer is set to stochastic gradient descent with zero momentum. For RL algorithms, we select REINFORCE [29] for the MNIST task and Advantage Actor-Critic [20] for the GymIP task. Hyper-parameters including the number of hidden neurons and RL discount factor γ are kept the same across methods. Note that the RWTA network in our method is fully-connected, which means it has more learnable parameters when having the same number of hidden neurons. To make the comparison fair we introduce a *SVPG-shrink* with fewer hidden WTA circuits, so that the number of learnable parameters are similar in all the models. Details are provided in Appendix C.2.

SVPG: spike coding versus rate coding


The spike train simulation in SVPG is computationally expensive for common GPUs. With Eq.(7) and Eq.(12), SVPG can also be implemented with rate-coded neuron states, which is more computationally efficient. To reproduce noises caused by spike-train simulation, we add Gaussian noise to rate coding after each iteration of Eq.(7). This results in the “Rate with noise” implementation.

To check whether rate coding alters the performance, we compare it with spike coding in the MNIST task. We consider two different STDP windows – double exponential and rectangle – and name them as “*Spike dexp*” and “*Spike rect*”. As in Figure 3 and Appendix D.2, the three implementations generate similar curves, which indicates the rate coding




(a) Input salt noise (b) Net Gaussian noise (a) Input Gaussian noise (b) Net Gaussian noise

Figure 3: SVPG: spike/rate coding. Figure 4: MNIST: Robustness.




(a) Input uniform noise (b) Net uniform noise (c) Pendulum length (d) Pendulum thickness

Figure 5: GymIP: Robustness.



(a) Net weights



(b) Noise responses

Figure 6: Visualizations.

implementation can be used as a replacement for spike coding. In the following experiments, we adopt the rate coding implementation.

4.3 Results

4.3.1 Training Performances

We perform 10 independent training trials (each with a different seed for the random number generator; each with the same number of episodes) for each method. We regularly save checkpoints during training and select the ones with the best zero-noise testing performance as the training results, which are summarized in Table 1. 1) In MNIST, SVPG achieves an accuracy close to BP and outperforms other RL-based methods. Note that although EP achieves high accuracy, it is trained using supervised signals, i.e., does not suffer from the distribution shift in RL. 2) In GymIP, SVPG achieves a near-optimal performance which is close to BP and BPTT. These results indicate that SVPG is able to solve both image-input and long-term decision tasks.

4.3.2 Robustness

We test the learned policies with different kinds of noises. To evaluate the performance, in MNIST, the testing dataset is used; in GymIP, 100 independently sampled episodes are used. We average the results across 10 training trials with different seeds, and present the standard deviations as shaded areas in the figures.

1) *Input noise*. The results of Gaussian noise in MNIST and uniform noise in GymIP are presented in Figure 4(a), 5(a). Others are in Appendix Figure 9, 11. For all the tested input noises, the performance of SVPG degrades the slowest as the amount/strength of noises increases. For 0.2 uniform noise in the GymIP task, SVPG even achieves more than twice the accuracy of all the compared methods.

2) *Network parameter noise*. The results of Gaussian noise in MNIST and uniform noise in GymIP are in Figure 4(b), 5(b). Other results are in Appendix Figure 10, 12. SVPG achieves slower degradation of performance as the amplitude of noises increases and has a better performance than all compared methods when the strength of noises is large.

3) *Environment variations in GymIP*. Results on variations in the pendulum’s length and thickness are in Figure 5(c), 5(d). Other variation results are in Appendix Figure 13. When the shape of the pendulum deviates from the one in training, the performance of all the compared methods degrades. However, the performance of SVPG degrades much slower than the baselines. Especially, when changing the pendulum’s length from 1.5 to 4, the performance of SVPG does not appear to degrade. This means the policy trained using SVPG naturally adapts to a larger range of pendulums with different shapes.

The above results show that 1) SVPG can solve classical RL tasks, even with high-dimensional state space. 2) SVPG provides better robustness to the mentioned different kinds of noises and environment variations than the compared methods.

Note that the results about *SVPG-shrink* and *SVPG* are similar. These indicate that it is not the larger number of learnable parameters in the RWTA network that brings the above better training performance and robustness.

4.3.3 Additional Results

We check some more properties of SVPG. **1) Network sparsity.** We plot the total distribution of synapse weights of 10 training trials (clipped at ± 0.15) in the MNIST task in Figure 6(a). Note that the number of synapse weights in SVPG-shrink and BP are close. As shown, there are more zero weights in SVPG than in BP. This indicates that SVPG tends to learn a more sparse network. **2) Hidden layer response to input noise.** Figure 6(b) shows the distribution of changes in the RWTA hidden neurons’ firing probabilities Δq_h when adding Gaussian noise to input MNIST images (std 0.02), and its relationship with SVPG’s inference iteration number. Results are gained using all 10k testing images and clipped at ± 0.3 . In early iterations, the firing probabilities are more affected by the noise as there are more of them that deviate from the no-noise firing probability; in the later iterations, they are less affected. This indicates that the recurrent design of the network should contribute to the robustness of SVPG. **3) Computational costs.** We summarize the times needed for inference and optimization stages in Appendix D.5. The rate coding version of SVPG is the fastest in optimization and is only slower than BP in inference.

5 Discussion and Conclusion

This paper proposes SVPG, a spiking-based variational policy gradient method with RWTA network and R-STDP. Current results reveal its potential for solving high-dimensional RL tasks and to have inherent robustness. More theoretical analysis on the cause of robustness, and more experiments in real-world control problems may improve this method and advance the research of biologically plausible methods for RL for applications in real-life scenarios.

Acknowledgements

This work was supported in part by National Natural Science Foundation of China Grant 62206151, China National Postdoctoral Program for Innovative Talents Grant BX20220167, Shuimu Tsinghua Scholar Program, Natural Science Foundation of Fujian Province of China Grant 2022J01656, and UK Royal Society Newton Advanced Fellowship Grant NAF-R1-191082. This work was undertaken in part on ARC3, part of the High Performance Computing facilities at the University of Leeds.

References

- [1] Hajar Asgari, Babak Mazloom-Nezhad Maybodi, Melika Payvand, et al. Low-Energy and Fast Spiking Neural Network for Context-Dependent Learning on FPGA. *IEEE Transactions on Circuits and Systems II: Express Briefs*, 67(11):2697–2701, 2020.
- [2] Jeffrey M. Beck, Katherine A. Heller, and Alexandre Pouget. Complex Inference in Neural Circuits with Probabilistic Population Codes and Topic Models. In *NeurIPS*, pages 3068–3076, 2012.
- [3] Guillaume Bellec, Franz Scherr, Anand Subramoney, et al. A Solution to the Learning Dilemma for Recurrent Networks of Spiking Neurons. *Nature Communications*, 11(1): 3625, 2020.
- [4] Marc G. Bellemare, Salvatore Candido, Pablo Samuel Castro, et al. Autonomous Navigation of Stratospheric Balloons Using Reinforcement Learning. *Nature*, 588 (7836):77–82, 2020.
- [5] Zhenshan Bing, Zhuangyi Jiang, Long Cheng, et al. End to End Learning of a Multi-Layered SNN based on R-STDP for a Target Tracking Snake-like Robot. In *ICRA*, pages 9645–9651, 2019.
- [6] Stephen Chung and Robert Kozma. Reinforcement Learning with Feedback-Modulated TD-STDP. *arXiv:2008.13044 [cs, stat]*, 2020.
- [7] A. Devo, G. Mezzetti, G. Costante, et al. Towards Generalization in Target-Driven Visual Navigation by Using Deep Reinforcement Learning. *IEEE Transactions on Robotics*, 36(5):1546–1561, 2020.
- [8] Nicolas Frémaux and Wulfram Gerstner. Neuromodulated Spike-Timing-Dependent Plasticity, and Theory of Three-Factor Learning Rules. *Frontiers in Neural Circuits*, 9, 2016.
- [9] Guillermo Gallego, Tobi Delbruck, Garrick Orchard, Chiara Bartolozzi, Brian Taba, Andrea Censi, Stefan Leutenegger, Andrew Davison, Jorg Conradt, Kostas Daniilidis, and Davide Scaramuzza. Event-based vision: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(1):154–180, 2022.
- [10] Wulfram Gerstner, Werner M. Kistler, Richard Naud, et al. *Neuronal Dynamics: From Single Neurons to Networks and Models of Cognition*. Cambridge University Press, 2014. ISBN 978-1-107-44761-5.
- [11] Shangqi Guo, Zhaofei Yu, Fei Deng, et al. Hierarchical Bayesian Inference and Learning in Spiking Neural Networks. *IEEE Transactions on Cybernetics*, 49(1):133–145, 2019.
- [12] Nicolas Heess, David Silver, and Yee Whye Teh. Actor-Critic Reinforcement Learning with Energy-Based Policies. In *Proceedings of the Tenth European Workshop on Reinforcement Learning*, pages 45–58, 2013.
- [13] Hyeryung Jang, Nicolas Skatchkovsky, and Osvaldo Simeone. VOWEL: A Local Online Learning Rule for Recurrent Networks of Probabilistic Spiking Winner-Take-All Circuits. In *ICPR 2020*, pages 4597–4604, 2020.

- [14] David Kappel, Robert Legenstein, Stefan Habenschuss, et al. A Dynamic Connectome Supports the Emergence of Stable Computational Function of Neural Circuits through Reward-Based Learning. *eNeuro*, 5(2):ENEURO-0301, 2018.
- [15] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553): 436–444, May 2015. ISSN 1476-4687.
- [16] Chen Li, Runze Chen, Christoforos Moutafis, et al. Robustness to Noisy Synaptic Weights in Spiking Neural Networks. In *IJCNN*, 2020.
- [17] Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, et al. Continuous Control with Deep Reinforcement Learning. In *ICLR*, 2016.
- [18] Ângelo Gregório Lovatto, Thiago Pereira Bueno, and Leliane Nunes de Barros. Analyzing the Effect of Stochastic Transitions in Policy Gradients in Deep Reinforcement Learning. In *BRACIS*, pages 413–418, 2019.
- [19] Camilo Andrés Manrique Escobar, Carmine Maria Pappalardo, and Domenico Guida. A Parametric Study of a Deep Reinforcement Learning Control System Applied to the Swing-Up Problem of the Cart-Pole. *Applied Sciences*, 10(24), 2020.
- [20] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, et al. Asynchronous Methods for Deep Reinforcement Learning. In *ICML*, pages 1928–1937, 2016.
- [21] Recep Özalp, Nuri Köksal Varol, Burak Taşci, and Ayşegül Uçar. *A Review of Deep Reinforcement Learning Algorithms and Comparative Results on Inverted Pendulum System*, pages 237–256. Springer International Publishing, 2020.
- [22] Devdhar Patel, Hananel Hazan, Daniel J Saunders, et al. Improved Robustness of Reinforcement Learning Policies upon Conversion to Spiking Neuronal Network Platforms Applied to Atari Breakout Game. *Neural Networks*, 120:108–115, 2019.
- [23] Danilo Rezende, Daan Wierstra, and Wulfram Gerstner. Variational Learning for Recurrent Spiking Networks. In *NeurIPS*, 2011.
- [24] Bodo Rueckauer, Iulia-Alexandra Lungu, Yuhuang Hu, et al. Conversion of Continuous-Valued Deep Networks to Efficient Event-Driven Networks for Image Classification. *Frontiers in Neuroscience*, 11:682, 2017.
- [25] Benjamin Scellier and Yoshua Bengio. Equilibrium Propagation: Bridging the Gap between Energy-Based Models and Backpropagation. *Frontiers in Computational Neuroscience*, 11:24, 2017.
- [26] John Schulman, Filip Wolski, Prafulla Dhariwal, et al. Proximal Policy Optimization Algorithms. *CoRR*, abs/1707.06347, 2017.
- [27] Saima Sharmin, Priyadarshini Panda, Syed Shakib Sarwar, Chankyu Lee, Wachirawit Ponghiran, and Kaushik Roy. A Comprehensive Analysis on Adversarial Robustness of Spiking Neural Networks. In *IJCNN*, 2019.
- [28] Ahmad Suhaimi, Amos Lim, Xin Chia, Chunyue Li, and Hiroshi Makino. Representation learning in the artificial and biological neural networks underlying sensorimotor integration. *Science Advances*, 8(22), 2022.

- [29] Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. MIT press, 2018.
- [30] Martin J. Wainwright and Michael I. Jordan. Graphical Models, Exponential Families, and Variational Inference. *Foundations & Trends in Machine Learning*, 1(1-2):1–305, 2008.
- [31] Timo Wunderlich, Akos F Kungl, Eric Müller, et al. Demonstrating Advantages of Neuromorphic Computation: A Pilot Study. *Frontiers in Neuroscience*, 13:260, 2019.
- [32] Angel Yanguas-Gil. Coarse Scale Representation of Spiking Neural Networks: Back-propagation through Spikes and Application to Neuromorphic Hardware. In *ICNS*, 2020.
- [33] Zhaofei Yu, Fei Deng, Shangqi Guo, Qi Yan, Jian K. Liu, and Feng Chen. Emergent Inference of Hidden Markov Models in Spiking Winner-Take-All Neural Networks. *IEEE Transactions on Cybernetics*, 2018.
- [34] Zhaofei Yu, Jian K. Liu, Shanshan Jia, Yichen Zhang, Yajing Zheng, Yonghong Tian, and Tiejun Huang. Toward the Next Generation of Retinal Neuroprosthesis: Visual Computation with Spikes. *Engineering*, 6(4):449–461, 2020.
- [35] Mengwen Yuan, Xi Wu, Rui Yan, et al. Reinforcement Learning in Spiking Neural Networks with Stochastic and Deterministic Synapses. *Neural Computation*, 31(12): 2368–2389, 2019.
- [36] Friedemann Zenke and Surya Ganguli. SuperSpike: Supervised Learning in Multilayer Spiking Neural Networks. *Neural Computation*, 30(6), 2018.