

Tuning Synaptic Connections instead of Weights by Genetic Algorithm in Spiking Policy Network

Duzhen Zhang^{1,2}, Tielin Zhang^{1,2†}, Shuncheng Jia^{1,2}, Qingyu Wang^{1,2} and Bo Xu^{1,2,3†}

¹School of Artificial Intelligence, University of Chinese Academy of Sciences, Beijing 100049, China.

²Institute of Automation, Chinese Academy of Sciences (CAS), Beijing 100190, China.

³Center for Excellence in Brain Science and Intelligence Technology, CAS, Shanghai 200031, China.

†Corresponding author

Abstract

Learning from interaction is the primary way that biological agents acquire knowledge about their environment and themselves. Modern deep reinforcement learning (DRL) explores a computational approach to learning from interaction and has made significant progress in solving various tasks. However, despite its power, DRL still falls short of biological agents in terms of energy efficiency. Although the underlying mechanisms are not fully understood, we believe that the integration of spiking communication between neurons and biologically-plausible synaptic plasticity plays a prominent role in achieving greater energy efficiency. Following this biological intuition, we optimized a spiking policy network (SPN) using a genetic algorithm as an energy-efficient alternative to DRL. Our SPN mimics the sensorimotor neuron pathway of insects and communicates through event-based spikes. Inspired by biological research showing that the brain forms memories by creating new synaptic connections and rewiring these connections based on new experiences, we tuned the synaptic connections instead of weights in the SPN to solve given tasks. Experimental results on several robotic control tasks demonstrate

that our method can achieve the same level of performance as mainstream DRL methods while exhibiting significantly higher energy efficiency.*

Keywords: Spiking Neural Networks, Genetic Evolution, Bio-inspired Learning, Agent & Cognitive Architectures, Robotic Control

1 Introduction

When we consider the nature of learning, the first thing that comes to mind is the idea of biological agents learning by interacting with their environment. Reinforcement learning (RL) is an algorithm that attempts to study this biological mechanism of how agents learn from their interactions, from a computational perspective [1]. In RL, an artificial agent interacts with the environment in a trial-and-error manner and learns an optimal policy by maximizing the accumulated rewards. In modern deep reinforcement learning (DRL), agents controlled by deep neural networks (DNNs) and optimized through gradient-based algorithms have proven to be capable of solving various tasks, ranging from video games [2, 3] to robotic control [4, 5, 6].

Although modern DRL has made remarkable progress, it still faces a challenge of high energy consumption. This includes both high inference and optimization energy consumption, which has been a major concern in real-world applications, especially for robot-control tasks that have limited onboard energy resources [7]. Specifically, the high inference energy consumption is due to the dense communication of continuous floating-point values, which is required by a deep policy network (DPN), a DNNs-based policy network, resulting in high computational costs. Moreover, the high optimization energy consumption mainly arises from performing backpropagation (BP) [8]. For instance, the computational cost of layer-by-layer BP for calculating gradients in gradient-based algorithms is significantly high, particularly when optimizing the parameters in a DPN using stochastic gradient descent/ascent.

In summary, despite its power, DRL is still far from being energy-efficient compared to the learning mechanism of biological agents. During learning from environmental interaction, biological agents typically exhibit ultra-low energy consumption [9, 10, 11]. Recent studies have shown that the biological brain can perform complex calculations with only 20 watts of energy [12, 13], while existing computing systems often require at least an order of magnitude higher energy consumption to achieve the same task [14]. Although it is not yet fully understood how biological agents learn so efficiently, we believe that the integration of spiking communication between neurons and biologically-plausible synaptic plasticity plays a fundamental role.

Drawing on biological intuition, we propose an energy-efficient alternative to DRL by optimizing a spiking policy network (SPN) using a gradient-free

*The source code, saved checkpoints and results are available at <https://github.com/BladeDancer957/SPN-GA>.

genetic algorithm (GA). Our SPN, which is used to infer actions from state observations, is inspired by the biological sensorimotor neuron pathway found in insects [15, 16]. The neurons in the SPN communicate via event-based binary spikes, which can reduce inference energy consumption. To reduce the optimization energy consumption, we design a gradient-free optimization paradigm based on GA, which directly conducts a SPN search. Furthermore, we highlight the importance of synaptic connection tuning as opposed to solely tuning synaptic weights, which is the traditional approach to optimizing a policy network. Inspired by biological research, where the brain forms memories by forming new synaptic connections and rewiring existing ones based on new experiences [17, 18, 19], we only tune the synaptic connections of the SPN to solve specific tasks, without adjusting the synaptic weights.

Our main contributions can be summarized as follows:

- Conventional DRL algorithms are computationally expensive and become more so as task complexity increases. To address this issue, we construct a biology-inspired SPN optimized by a gradient-free GA, providing an energy-efficient alternative to DRL. Our approach is innovative and significant for energy-efficient computation in the RL field, where the GA algorithm can replace big memory buffers to achieve a similar independent identically distributed distribution of training data.
- We adopt the viewpoint of direct synaptic connection tuning instead of weight tuning by GA in SPN to solve RL tasks. Our surprising finding is that modifying synaptic connections is not only more efficient but also more fitting for the SPN, where discrete instead of continuous signals are used for information representation.
- We present extensive experimental results on several robotic control tasks, demonstrating that our alternative method can achieve the same performance level as mainstream DRL methods while exhibiting significantly higher energy efficiency (saving about 1.1~1.3 times inference energy consumption and 6.5~87.9 times optimization energy consumption under various tasks). Further analysis verifies the superiority of connection tuning over weight tuning and the necessity of using GA to optimize the SPN.

2 Related Work

2.1 RL

Conventional RL algorithms excel at handling simple, tabular state spaces [1]. However, they face challenges when dealing with complex state spaces. To address this issue, modern DRL relies on powerful function approximators, such as DNNs, to learn a mapping function directly from the raw sensory state space to the action space. Two main families of gradient-based algorithms are commonly used in DRL: the Q-learning algorithm [20], represented by DQN [2], and the policy gradient algorithm [21], represented by PPO [22].

DQN [2] leverages DNNs to approximate the optimal Q function, resulting in a policy that selects the action that maximizes the Q-value for a given state. Two crucial techniques employed in DQN are the experience replay and the target network. The former stores the agent's experiences into a replay buffer, allowing them to be reused in many weight updates for greater data efficiency. The latter generates the targets in the Q-learning update and addresses the instability problem that arises when combining DNNs and Q-learning. DQN and its improved variants, such as Double DQN [23] and Dueling DQN [24], are currently advancing towards achieving human-level control in Atari video games [25] with discrete action spaces.

However, the value-based DRL algorithms discussed above have limited processing capacity for continuous action spaces, making them unsuitable for certain scenarios. To address this issue, the PG algorithm [21] was proposed to directly learn the parameters in a DPN. However, it still suffers from the high variance problem, partly due to the probabilistic output for taking actions. To address this, the actor-critic (AC) framework was introduced, where a deep actor/policy network infers an action from a given state, and a deep critic/-value network estimates the associated state-value or action-value to guide the actor during optimization. Several improved AC-based PG algorithms have been proposed to improve exploration and data efficiency, including A3C [26], TRPO [27], PPO [22], DDPG [5], TD3 [28], and SAC [29]. Among these algorithms, PPO is a more general and easier-to-implement algorithm that can be scaled to almost all task scenarios, including video games with discrete action spaces and robotic control with both discrete and continuous action spaces.

Despite the remarkable progress made by DRL algorithms described above, their solutions are often energy-intensive due to the reliance of deep policy networks on dense communication of continuous values and the gradient-based optimization paradigm. To address this issue, we propose a biologically-inspired SPN, optimized by a gradient-free genetic algorithm, as an energy-efficient alternative to DRL.

2.2 Integrating Spiking Neural Networks (SNNs) with RL

Recently, there has been a growing body of literature exploring the integration of SNNs [30, 31] with RL. Early approaches [32, 33, 34, 35, 36] were typically based on local plasticity rules, such as reward-modulated spike-timing-dependent plasticity [37, 38]. While they perform well in simple tasks such as logic gate functions and random walk problems [33], they fail to tackle more complex tasks such as video games or robotic control due to limited optimization capabilities.

To address this limitation, some methods have incorporated SNNs into gradient-based DRL optimization algorithms [39]. Using the Q-learning algorithm as a base, some approaches directly train deep spiking Q-networks [16, 40, 41, 42] or convert trained DQNs into SNNs [43, 44] to achieve competitive scores on Atari video games with discrete action spaces. Other methods have proposed a hybrid learning framework tested on robotic control tasks with

continuous action spaces. This framework involves training a population-coded spiking actor network in conjunction with deep critic networks using AC-based PG algorithms such as PPO [45, 7, 46, 47]. While these methods obtain a SNN-based policy network for inference to solve the high energy consumption of DPN inference, they still rely on gradient-based optimization algorithms, with the high optimization energy consumption overlooked.

2.3 Neuroscience

In neuroscience, a connectome is a comprehensive “wiring diagram” that maps all neural connections in the brain [48]. Currently, connectomes for simple organisms like fruit flies [49, 50] and roundworms [51, 52] have been constructed. The study of connectomes is motivated by the potential to guide future research on how the brain learns and forms memories through its connections [53]. It is clear that humans, as biological agents, form memories and acquire new skills by creating new synaptic connections, particularly during early childhood [54, 55]. The brain then rewires these connections based on new experiences [17, 18, 19, 56]. Drawing inspiration from this biological research, we integrate biologically plausible synaptic connection tuning with biologically inspired SPNs to achieve energy-efficient robotic control.

3 Method

3.1 Constructing a SPN

In RL tasks, an artificial agent interacts with its environment by taking a series of actions based on its state observations and receiving rewards. At each time step t , the agent chooses an action \mathbf{a}^t from a policy π based on the current state observation \mathbf{o}^t . The agent then receives a scalar reward r^{t+1} and a new state observation \mathbf{o}^{t+1} , which creates an episode $\epsilon = \mathbf{o}^0, \mathbf{a}^0, r^1, \dots, \mathbf{o}^{T-1}, \mathbf{a}^{T-1}, r^T$, where T is the length of the episode. The return of the episode ϵ is the sum of rewards: $R^\epsilon = \sum_{i=1}^T r^i$. To learn a mapping from state observations to actions, we construct a SPN that represents the policy π , which infers actions \mathbf{a}^t from observations \mathbf{o}^t .

Our SPN is inspired by the sensorimotor neuron pathway observed in insects, as described in previous research [15, 16]. This pathway involves sensory neurons receiving information from the external environment and transmitting it to spiking interneurons that have discrete dynamics. Once the afferent current reaches a threshold, the spiking interneurons generate spikes, which then carry the signal to non-spiking interneurons. The membrane potential of these non-spiking interneurons determines the input current of motor neurons, which are responsible for effective locomotion.

To align with the biological pathway, our SPN follows a three-layer architecture consisting of the sensory layer, middle processing layer, and motor layer. The correspondence between the sensorimotor neuron pathway and our SPN is depicted in Fig. 1. For instance, in the context of a robotic control

6 Tuning Synaptic Connections instead of Weights by GA in SPN

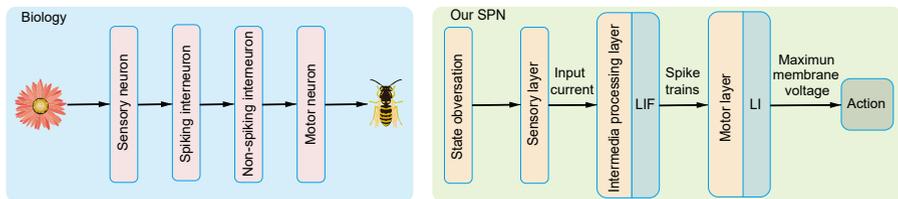


Fig. 1 The correspondence diagram between the sensorimotor neuron pathway and our SPN.

task, the sensory layer acquires a continuous state observation $\mathbf{o} \in \mathbb{R}^n$ that comprises information such as the angles, velocities, and external forces of each joint of the simulated robot. Here, n represents the dimensionality of the input observation. The sensory information is subsequently transmitted to the middle processing layer, where neurons are modeled using the leaky integrate-and-fire (LIF) [57], which is the most widely used artificial spiking neuron. In contrast, the motor layer neurons are modeled using the leaky integrate (LI) [16], which is a non-spiking neuron similar to LIF but without the capability to fire and reset. We elaborate on the fundamental neuron models employed in our SPN in the subsequent sections.

LIF.

The dynamic firing process at neuron j can be described as:

$$\begin{cases} \frac{dv_{j,\tau}}{d\tau} = -g(v_{j,\tau} - v_{reset}) + \sum_{i=1}^n (w_{i,j} \times x_{i,j})o_{i,\tau} \\ v_{j,\tau} = v_{reset}, s_{j,\tau} = 1 \quad \text{if } (v_{j,\tau} > v_{th}) \end{cases} \quad (1)$$

where $j = 1, \dots, h$, h is the number of neurons in the middle layer, $\tau = 1, \dots, T'$, τ is a simulation time step, T' is the number of simulation time step (i.e., time window), meaning the same observation will be fed into SPN for T' times, $v_{j,\tau}$ is the membrane potential of neuron j at time τ in the middle layer, g is the decay factor, v_{reset} denotes the rest membrane potential, $s_{j,\tau}$ means the output spike, $x_{i,j}$ is a binary variable indicating whether the synaptic connection from neuron i to j exists, $o_{i,\tau}$ is the input current from upstream neuron i ($i = 1, \dots, n$), v_{th} is the predefined firing threshold, v_{reset} denotes the predefined reset membrane potential, and $w_{i,j}$ is the synaptic weight between neuron i and j . In our method, we only tune synaptic connections instead of weights (See Section 3.2 for details).

The LIF neuron in equation (1) captures the temporal dynamics of membrane potential, emulating the behavior of biological neurons. The threshold firing property of LIF neurons results in sparse spike signals. Furthermore, the computation is event-driven, i.e., it is performed only when a spike arrives. This property can reduce the energy consumption during inference, especially when deployed on neuromorphic chips [58, 59].

Algorithm 1 The pseudocode of forward propagation in SPN

Synaptic weight matrices \mathbf{W}^1 and \mathbf{W}^2 for middle processing and motor layer (randomly initialized and fixed);

Score matrices corresponding to synaptic connections \mathbf{C}^1 and \mathbf{C}^2 for middle processing and motor layer (tuned by GA);

Initialize the score threshold s_{th} ;

Obtain the mask of synaptic connection $\mathbf{X}^1 = (\text{sigmoid}(\mathbf{C}^1) \geq s_{th})$ and $\mathbf{X}^2 = (\text{sigmoid}(\mathbf{C}^2) \geq s_{th})$ for middle processing and motor layer;

Initialize the potential decay factor g and firing threshold v_{th} ;

n -dimensional state observation, \mathbf{o} ;

for $\tau = 1$ to T'

$$\mathbf{v}_\tau^1 = g\mathbf{v}_{\tau-1}^1(1 - \mathbf{s}_{\tau-1}^1) + (\mathbf{W}^1 \times \mathbf{X}^1)\mathbf{o};$$

$$\mathbf{s}_\tau^1 = \mathbf{v}_\tau^1 > v_{th};$$

$$\mathbf{v}_\tau^2 = g\mathbf{v}_{\tau-1}^2 + (\mathbf{W}^2 \times \mathbf{X}^2)\mathbf{s}_\tau^1;$$

Generate output action, \mathbf{a} ;

if Continuous action space

$$\mathbf{a} = \max_{1 \leq \tau \leq T'} \mathbf{v}_\tau^2;$$

else

$$\mathbf{a} = \arg \max(\max_{1 \leq \tau \leq T'} \mathbf{v}_\tau^2);$$

LI.

The LI neuron is a special case of the LIF neuron without the firing and resetting mechanism. When the firing threshold v_{th} of a LIF neuron is set to infinity, it becomes a non-spiking neuron. The dynamic equation in equation (1) can be simplified as:

$$\frac{dv_{k,\tau}}{d\tau} = -g(v_{k,\tau} - v_{rest}) + \sum_{j=1}^h (w_{j,k} \times x_{j,k})s_{j,\tau} \quad (2)$$

where $k = 1, \dots, m$ and m is the number of neurons in the motor layer.

The LI neurons act as a spike decoder, converting spike trains from the previous layer into continuous values. The membrane potential of the LI neurons is recorded at each simulation time step τ , resulting in $\mathbf{v}_\tau \in \mathbb{R}^m$. The maximum membrane potential is then used to generate the output action:

$$\begin{cases} \mathbf{a} = \max_{1 \leq \tau \leq T'} \mathbf{v}_\tau, \text{ continuous action space} \\ \mathbf{a} = \arg \max(\max_{1 \leq \tau \leq T'} \mathbf{v}_\tau), \text{ discrete action space} \end{cases} \quad (3)$$

where $\mathbf{a} \in \mathbb{R}^m$ represents the torque applied to each joint of the robot for continuous action space, and a represents one of the predefined actions for discrete action space. The forward inference procedure for our SPN is summarized in Algorithm 1.

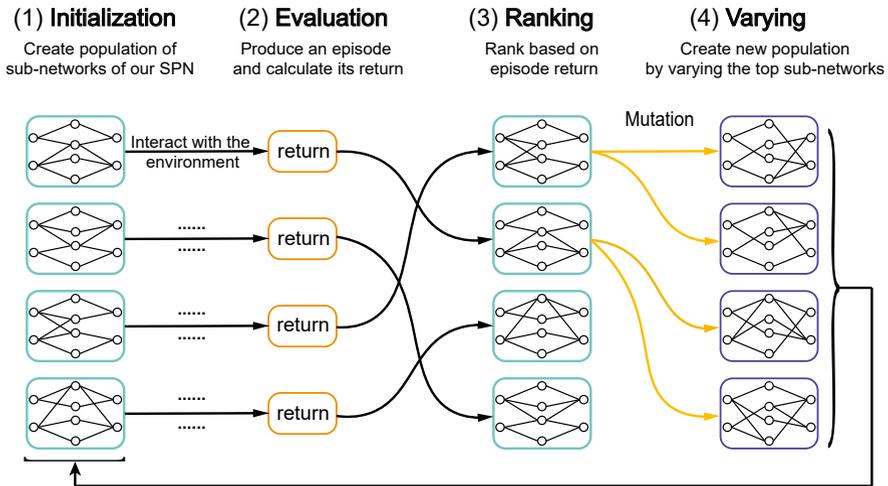


Fig. 2 Overview of the evolution of our SPN sub-networks.

3.2 Tuning Synaptic Connections by GA

Our method aims to evolve the sub-networks of the SPN in order to solve a given task for an artificial agent in a simulation environment. During this process, we focus solely on tuning the synaptic connections through GA, rather than placing emphasis on tuning synaptic weights. The evolution of the sub-networks within our SPN can be summarized as shown in Fig. 2.

Initialization

To begin, we generate the first generation population \mathcal{P}^0 consisting of N individuals, each of which represents a vector of SPN parameters θ_j^i . This population only contains scores that correspond to synaptic connections without synaptic weights, which are randomly initialized and fixed. In order to reduce variance, we employ mirrored sampling from previous evolution strategy literature [60, 61, 62]. Specifically, we use pairs of perturbations to produce individuals: $\theta_j^0 = \sigma\epsilon$ or $-\sigma\epsilon$, where $j = 0, \dots, N-1$, σ represents the standard deviation, and $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. Once this is complete, we normalize the connection scores between 0 and 1 using the sigmoid function, and then set a score threshold of s_{th} . If the normalized score is less than s_{th} , we remove the connection, otherwise we keep the connection, resulting in a sub-network. Each SPN sub-network can be considered as an individual.

Evaluation

At each generation i ($i = 0, \dots, G-1$), where G represents the total number of generations, we evaluate every sub-network to obtain a fitness score. By analyzing the interaction between the sub-network and the environment, we can generate an episode and utilize its return as the fitness score.

Algorithm 2 The optimization process of our GA version

Input: number of total generations G , population size N , standard deviation of perturbations σ , ratio of truncation selection η , and fitness score function F .

Output: Elites = [] (store the elite of each generation).

for $i = 0, \dots, G-1$ generations

 Create i generation population of N individuals \mathcal{P}^i .

 for $j = 0, \dots, N-1$ individuals

 if $i = 0$

Initialization. $\theta_j^{i=0} = \sigma\epsilon$ or $-\sigma\epsilon$, $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$

 {initialize scores corresponding to all synaptic connections of SPN.}

 else

Varying. $\theta_j^i = \theta_{j_\eta}^{i-1} + \sigma\epsilon$ or $-\sigma\epsilon$, $j = 0, \dots, N-1$, $\theta_{j_\eta}^{i-1} \in \mathcal{P}_\eta^{i-1}$.

 {Mutation.}

Evaluation. $F_j = F(\theta_j^i)$

 {Obtain the sub-network of SPN by sigmoid normalization of connection scores and filtering by score threshold. Based on the interaction between the sub-network and environment, obtain an episode and use its return as the fitness score.}

Ranking. Sort θ_j^i with descending order by F_j

 Perform truncation selection with a ratio of η to produce the next generation's parents \mathcal{P}_η^i .

 Set elite candidates for generation i (top 10):

$C^i \leftarrow \theta_{1\dots 10}^i$.

 Select elite for generation i :

 elite ^{i} $\leftarrow \arg \max_{\theta \in C^i} \frac{1}{10} \sum_{k=1}^{10} F(\theta)$.

 Elites += [elite ^{i}]

Return: Elites

Ranking

We rank each sub-network in descending order based on fitness. Then, we perform truncation selection with a ratio of η , where the top $(\eta * 100)\%$ of sub-networks will serve as the parents of the next generation. This group of parents is denoted as \mathcal{P}_η^i . To ensure that we obtain the optimal sub-network with greater reliability, we evaluate the top 10 sub-networks from each generation in 10 additional episodes (considering these interaction steps as consumed ones during optimization). The sub-network with the highest average return is referred to as the “elite” sub-network of the current generation.

Varying

Next generation population \mathcal{P}^{i+1} of N sub-networks is created by varying the parents obtained via truncation selection in the previous step. This process of variability is typically referred to as “Mutation”. Specifically, we generate offspring by adding perturbations produced by mirror sampling to the parents: $\theta_j^{i+1} = \theta_{j_\eta}^i + \sigma\epsilon$ or $-\sigma\epsilon$, where $j = 0, \dots, N-1$, and $\theta_{j_\eta}^i \in \mathcal{P}_\eta^i$ is randomly

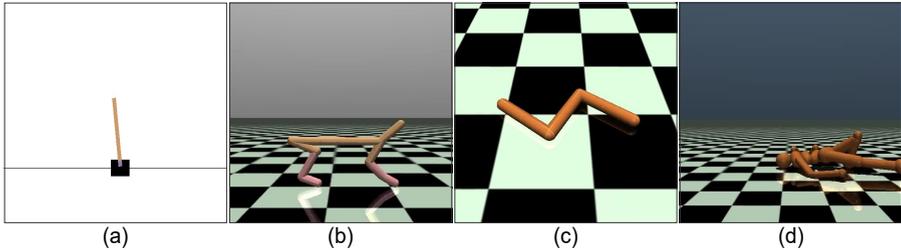


Fig. 3 Four robotic control tasks. **(a) CartPole-v1**: Observation dimension: $n = 4$, Action dimension: $m = 2$ (discrete), Goal: balance a pole on a cart; **(b) HalfCheetah-v2**: Observation dimension: $n = 17$, Action dimension: $m = 6$ (continuous), Goal: make a 2D cheetah robot run as fast as possible; **(c) Swimmer-v2**: Observation dimension: $n = 8$, Action dimension: $m = 2$ (continuous), Goal: make a 2D robot swim; **(d) HumanoidStandup-v2**: Observation dimension: $n = 376$, Action dimension: $m = 17$ (continuous), Goal: make a 3D two-legged robot standup.

selected with replacement. We can obtain new sub-networks (individuals) by normalizing the connection scores using the sigmoid function and filtering them by the score threshold s_{th} . We can repeat steps **Evaluation-Varying** for $G-1$ iterations.

The pseudocode for our GA version is summarized in Algorithm 2.

4 Experimental Settings

4.1 Evaluation Tasks

We evaluated our method on a benchmark of robotic control tasks interfaced by OpenAI Gym [63], as shown in Fig. 3. These tasks included classic discrete control tasks like CartPole-v1, as well as more challenging continuous control tasks such as HalfCheetah-v2, Swimmer-v2, and HumanoidStandup-v2, which were run in the fast physics simulator MuJoCo [64]. The original CartPole-v1 task had an upper limit return of 500, without any manual modification to the environment or reward.

4.2 Implement Details

On all tasks, our SPN was a multi-layer perceptron architecture with a 64-unit middle layer, i.e., (sensory layer, middle processing layer, LIF, motor layer, LI), where the number of simulation time step T' was 4, the decay factor g was 0.75, the rest and reset potential v_{rest} and v_{reset} were both 0, and the firing threshold v_{th} was 0.5. In our GA version, we set the number of generations G to 100, the number of individuals N in each generation to 200, the standard deviation σ of perturbations to 0.01, the ratio η of truncation selection to 0.25, and the connection score threshold s_{th} to 0.5.

For a fair comparison, on all tasks, the DPN was set up with the same number of parameters and architecture as our SPN but with different neuron types, i.e., (sensory layer, middle processing layer, ReLU, motor layer, Tanh).

It was optimized in conjunction with a deep value network by PPO [22] (one of the most general RL algorithms and the default RL algorithm at OpenAI). The architecture of the deep value network was (sensory layer, middle processing layer, ReLU, output layer), where the output layer with 1 unit estimated the value of the input observation. The PPO implementation and its hyper-parameter configurations were from OpenAI Spinning Up [65]. Our DPN baseline achieved similar or better performance than reported in the literature [22, 66]. The detailed hyper-parameter configurations and optimization process of PPO were shown in Appendix.

To address concerns regarding reproducibility [67], we used the original set of tasks from Gym [63] without modifying the reward or environment. We conducted 10 independent runs with different random seeds for all experiments. In GA, we ran each task for 100 generations and evaluated the “elite” sub-network of each generation with 10 episodes, reporting the average return over 10 runs. In PPO, following previous works [22, 66], we ran each task for $1e6$ steps and evaluated the final learned DPN with 10 episodes, reporting the average return over 10 runs as its performance level. All experiments were conducted on an AMD EPYC 7742 server.

5 Results and Discussions

Our experiments aimed to accomplish the following objectives: 1) To demonstrate the feasibility of our alternative method, SPN-Connections-GA, for tuning the synaptic connections of SPN by gradient-free GA, compared to the mainstream DRL method, DPN-Weights-PPO, for tuning the synaptic weights of DPN by gradient-based PPO [22] (Section 5.1); 2) To demonstrate the superiority of connection tuning over weight tuning (Section 5.2); 3) To demonstrate the necessity of using GA to optimize SPN (Section 5.3); 4) To showcase the dynamic evolution process of SPN-Connections-GA (Section 5.4).

5.1 SPN-Connections-GA vs. DPN-Weights-PPO

In this section, we compare our alternative method, SPN-Connections-GA, with the mainstream DRL method, DPN-Weights-PPO [22], in terms of the return gained, data efficiency, and energy efficiency.

5.1.1 Return Gained

Fig. 4 (a) shows that on the classical discrete control task CartPole-v1, SPN-Connections-GA reached the upper limit return of 500, preliminarily verifying the feasibility of our method. Furthermore, as demonstrated in Fig. 4 (b)–(d), on more challenging MuJoCo continuous control tasks HalfCheetah-v2, Swimmer-v2, and HumanoidStandup-v2, our SPN-Connections-GA was able to achieve the performance level of DPN-Weights-PPO [22] (the red horizontal line in the Fig.).

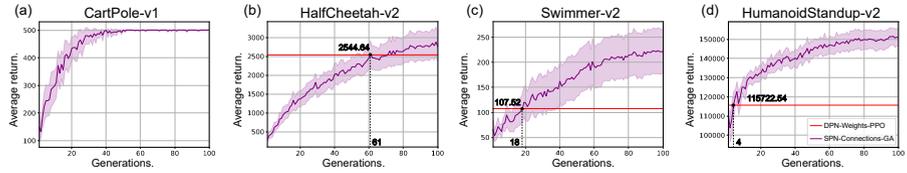


Fig. 4 Learning curves for the OpenAI Gym robotic control tasks. The purple areas represented the learning curves of SPN-Connections-GA during 100 generations of evolution, where the solid curves correspond to the mean and the shaded region to half a standard deviation over 10 runs. The red horizontal line represented the performance level of DPN-Weights-PPO [22].

Table 1 The exact data efficiency tradeoffs between SPN-Connections-GA and DPN-Weights-PPO [22]. \mathcal{R} denoted the performance level of DPN-Weights-PPO [22] after using $1e6$ steps of environment interaction. γ denoted a measure of data efficiency, defined in Equation (4).

Task	\mathcal{R}	\mathcal{G}	\mathcal{E}	γ
HalfCheetah-v2	2544.64	61	$1.83e7$	18.3
Swimmer-v2	107.52	18	$5.4e6$	5.4
HumanoidStandup-v2	115722.54	4	$1.2e6$	1.2

5.1.2 Data Efficiency

Although the data efficiency of GA is typically low, we found that the data efficiency γ of our SPN-Connections-GA was surprisingly good. In our approach, each data (step) that interacted with the environment was used only once. As Table 1 shows, using no more than 18.3x as much data (i.e., 18.3x on HalfCheetah-v2, 5.4x on Swimmer-v2, and 1.2x on HumanoidStandup-v2), our SPN-Connections-GA was able to achieve the same performance level as DPN-Weights-PPO [22], which used $1e6$ steps of environment interaction. However, DPN-Weights-PPO [22] was more data-efficient because it maintained a vast replay buffer for storing data interacting with the environment, which could be utilized multiple times during optimization. In contrast, our SPN-Connections-GA did not use data repeatedly, did not have a deep value network, and did not perform BP during optimization, resulting in much lower energy consumption, which offset the slight decrease in data efficiency. We defined the specific data efficiency calculation method as follows.

For each MuJoCo control task, we first calculate the steps \mathcal{E} of environment interaction required for SPN-Connections-GA to achieve the performance level \mathcal{R} of DPN-Weights-PPO [22]. Then, we calculate the ratio γ of \mathcal{E} to standard PPO steps ($1e6$) as a measure of data efficiency.

$$\begin{aligned}
 \mathcal{A} &= N \times T + 10 \times 10 \times T \\
 \mathcal{E} &= \mathcal{G} \times \mathcal{A} \\
 \gamma &= \mathcal{E}/1e6
 \end{aligned} \tag{4}$$

Table 2 The network inference energy consumption comparison of our SPN and DPN. The input layer in our SPN was non-spiking, so its energy was same as DPN. Unit: pJ.

Task	E_{DPN}^{infer}	E_{SPN}^{infer}	$\frac{E_{DPN}^{infer}}{E_{SPN}^{infer}}$
HalfCheetah-v2	6.77e3	5.2e3	1.3
Swimmer-v2	2.94e3	2.41e3	1.2
HumanoidStandup-v2	1.16e5	1.1e5	1.1

Table 3 The optimization energy consumption comparison of our GA and PPO. Unit: pJ

Task	E_{PPO}^{optim}	E_{GA}^{optim}	$\frac{E_{PPO}^{optim}}{E_{GA}^{optim}}$
HalfCheetah-v2	6.16e11	9.52e10	6.5
Swimmer-v2	2.85e11	1.3e10	21.9
HumanoidStandup-v2	1.16e13	1.32e11	87.9

where \mathcal{G} was the number of generations required to achieve the performance level of DPN-Weights-PPO [22] (inferred from the vertical dashed line in Fig. 4 (b) – (d)), \mathcal{A} denoted the number of interaction steps consumed per generation of optimization, T was the length of the episode ($T=1000$ in all the tasks we used), and $N=200$ was the number of individuals in each generation (each individual needed to interact with the environment to get an episode and use the return of the episode as its fitness score). In each generation, we needed to evaluate each of the top 10 individuals in 10 additional episodes to obtain the “elite” individual. Thus, \mathcal{A} was $3e5$. The exact data efficiency tradeoffs were listed in Table 1.

5.1.3 Energy Efficiency

In this section, we analyze the energy efficiency of our SPN-Connections-GA compared to DPN-Weights-PPO [22], considering both network inference and optimization. As shown in Table 2, our SPN’s network inference energy consumption is 1.3x, 1.2x, and 1.1x lower on HalfCheetah-v2, Swimmer-v2, and HumanoidStandup-v2, respectively, compared to DPN. For optimization, as shown in Table 3, our GA’s energy consumption is 6.5x, 21.9x, and 87.9x lower on HalfCheetah-v2, Swimmer-v2, and HumanoidStandup-v2, respectively, compared to PPO. The specific energy consumption calculation method was defined in the next section.

5.1.4 Theoretical Analysis of Energy Consumption

For inference, we first analyzed the computation patterns of DPN and SPN. For DPN, each operation involved a dot-product that required one floating-point multiplication and one accumulation (MAC). In contrast, for SPN, each operation involved only an accumulation (AC) due to the use of binary spikes (0 or 1). Following [68], we computed the inference energy consumption for DPN and SPN using 45nm CMOS technology. Previous literature has shown

Table 4 The specific forward and back propagation times of data in the network.

		Policy Network		Value Network	
		forward	backward	forward	backward
PPO		2.6e7	2.5e7	2.6e7	2.5e7
GA	HalfCheetah	1.83e7	–	–	–
	Swimmer	5.4e6	–	–	–
	HumanoidStandup	1.2e6	–	–	–

that the energy consumption for a 32-bit DNN MAC operation ($e_{MAC}=4.6pJ$) is 5.1 times more than that of an SNN AC operation ($e_{AC}=0.9pJ$) [69].

Following [68], we define the theoretical energy consumption for network inference as the product of the number of operations in each layer of the network and the energy required for a single operation, formulated as follows:

$$\begin{aligned} E_{DPN}^{infer} &= \sum_l \#OP_{DPN} \times e_{MAC} \\ E_{SPN}^{infer} &= \sum_l \#OP_{SPN} \times e_{AC} \end{aligned} \quad (5)$$

Following [68], in DPN, we defined the number of operations at layer l as:

$$\#OP_{DPN} = f_{in} \times f_{out} \quad (6)$$

where $f_{in}(f_{out})$ was the dimension of input (output) features at layer l . The number of operations at layer l in SPN was defined as:

$$\begin{aligned} \#OP_{SPN} &= rate_l \times \#OP_{DPN} \\ rate_l &= \frac{\#totalSpikes_l}{\#n_l} \end{aligned} \quad (7)$$

where $rate_l$ was the total spikes at layer l over time window averaged over the number of neurons n_l at layer l (averaged over all interaction steps).

We define the theoretical energy consumption of network optimization as the theoretical energy consumption of network inference multiplied by the number of network inferences during the optimization process:

$$\begin{aligned} E_{PPO}^{optim} &= E_{DPN}^{infer} \times (N_{DPN}^{forward} + N_{DPN}^{backward}) \\ &\quad + E_{DVN}^{infer} \times (N_{DVN}^{forward} + N_{DVN}^{backward}) \\ E_{GA}^{optim} &= E_{SPN}^{infer} \times N_{SPN}^{forward} \end{aligned} \quad (8)$$

Here, we calculate the number of network inferences during the optimization process. First, the PPO optimization process includes both the forward inference and backpropagation of the DPN, as well as the forward inference and backpropagation of the deep value network (DVN). To simplify the energy consumption calculations, we consider both the network backpropagation and forward inference to have the same energy consumption. Specifically, for each

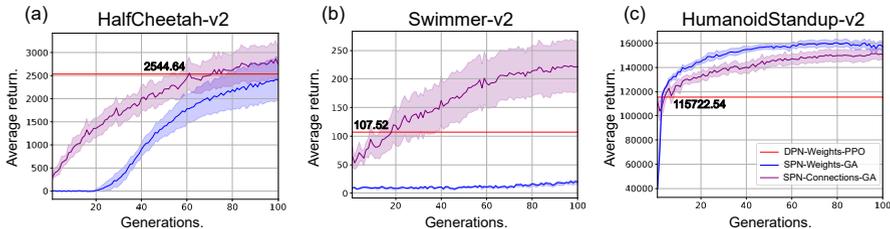


Fig. 5 Learning curves for the MuJoCo continuous control tasks. The purple areas represented the learning curves of SPN-Connections-GA during 100 generations of evolution, where the solid curves correspond to the mean and the shaded region to half a standard deviation over 10 runs. The blue areas represented the learning curves of SPN-Weights-GA during 100 generations of evolution, where the solid curves correspond to the mean and the shaded region to half a standard deviation over 10 runs. The red horizontal line represented the performance level of DPN-Weights-PPO [22].

task, the PPO optimization requires interaction with the environment to collect a total of $1e6$ data (steps). Each time a piece of data is collected, both the DPN (inferred an action based on observation) and the DVN (estimated the value of the observation) need to conduct a forward inference. For every $1e3$ data collected, 25 epochs of optimization (including forward inference and backpropagation of DPN and DVN) are conducted based on these data. During GA optimization, there is only forward inference of SPN, and the number of forward inferences corresponds to the number of data ($N_{SPN}^{forward}$ in Equation (8) equals \mathcal{E} in Equation (4)). The forward and backpropagation times of data in the network are listed in Table 4. It’s worth noting that we calculated the optimization energy consumption of our SPN-Connections-GA when achieving the same performance level as DPN-Weights-PPO [22].

5.2 SPN-Connections-GA vs. SPN-Weights-GA

In this section, we verified the superiority of connection tuning over weight tuning by comparing the performance of SPN-Connections-GA with SPN-Weights-GA.

Fig. 5 (a)–(c) shows that SPN-Connections-GA achieved a higher average return on HalfCheetah-v2 and Swimmer-v2 compared to SPN-Weights-GA. Additionally, SPN-Connections-GA achieved a comparable average return on HumanoidStandup-v2. It’s worth noting that even after 100 generations of evolution, SPN-Weights-GA could not achieve the same performance level as DPN-Weights-PPO [22] on HalfCheetah-v2 and Swimmer-v2. In fact, SPN-Weights-GA struggled to learn and was only able to achieve a shallow level of performance on Swimmer-v2. These results suggest that an “elite” sub-network of SPN with random weights can achieve a comparable or better return than the full network with learned weights tuned by the same GA.

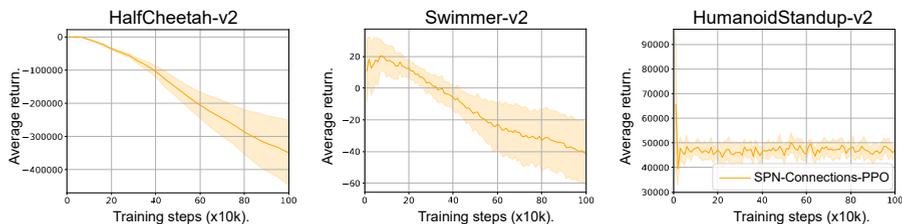


Fig. 6 Learning curves for the MuJoCo continuous control tasks. The yellow areas represented the learning curves of SPN-Connections-PPO, where the solid curves correspond to the mean and the shaded region to half a standard deviation over 10 runs.

5.3 SPN-Connections-GA vs. SPN-Connections-PPO

In this section, we aimed to verify the necessity of using GA to optimize SPN by comparing the performance of SPN-Connections-GA with SPN-Connections-PPO. We observed that directly replacing DPN with our SPN and integrating it into the gradient-based PPO optimization framework did not produce satisfactory results. As shown in Fig. 6, the learning curve of SPN-Connections-PPO did not converge, and its average return even decreased over time during training. On the other hand, as demonstrated in Fig. 4, SPN-Connections-GA performed well on the same tasks. The suboptimal performance of SPN-Connections-PPO might be attributed to the fact that our SPN, which relied on discrete spike communication, was a low-precision network. As a result, using a gradient-based PPO algorithm to optimize such a network could produce biased low-precision gradient estimates, leading to sub-optimal performance. In contrast, black-box optimization algorithms like GA might be better suited for optimizing low-precision networks such as SPN.

5.4 The evolution process of SPN-Connections-GA

For MoJoCo continuous control tasks HalfCheetah-v2 and Swimmer-v2, we randomly sampled one of 10 independent evolutionary runs and visualized the dynamic evolution process of the sub-network of SPN. The video URLs of dynamic evolution were listed in Table 5.

Table 5 The video URLs of dynamic evolution.

Task	URL
HalfCheetah-v2	https://www.youtube.com/watch?v=_o0v3LtmU-U
Swimmer-v2	https://www.youtube.com/watch?v=JmxR4deP4Cc

6 Conclusion

In this paper, we propose a novel approach for optimizing a biologically-inspired SPN using a gradient-free GA as an energy-efficient alternative to DRL. Unlike previous methods that primarily focused on tuning synaptic weights, we adopt a new perspective of optimizing synaptic connections to evolve “elite” sub-networks of the SPN capable of solving complex tasks. Our extensive experiments on several challenging robotic control tasks demonstrate that our approach can achieve the performance level of mainstream DRL methods while significantly reducing energy consumption. Moreover, our analyses confirm the superiority of connection tuning over weight tuning and the necessity of using GA to optimize SPN.

7 Acknowledgments

This work was supported by the Beijing Nova Program (Grant No. 20230484369), the Strategic Priority Research Program of Chinese Academy of Sciences (Grant No. XDA27010404), the Shanghai Municipal Science and Technology Major Project (Grant No. 2021SHZDZX), and the Youth Innovation Promotion Association of the Chinese Academy of Sciences.

References

- [1] Sutton, R.S., Barto, A.G.: Reinforcement Learning: An Introduction, (2018)
- [2] Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A.A., Veness, J., Bellemare, M.G., Graves, A., Riedmiller, M., Fidjeland, A.K., Ostrovski, G., *et al.*: Human-level control through deep reinforcement learning. *nature* **518**(7540), 529–533 (2015)
- [3] Vinyals, O., Babuschkin, I., Czarnecki, W.M., Mathieu, M., Dudzik, A., Chung, J., Choi, D.H., Powell, R., Ewalds, T., Georgiev, P., *et al.*: Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature* **575**(7782), 350–354 (2019)
- [4] Duan, Y., Chen, X., Houthoofd, R., Schulman, J., Abbeel, P.: Benchmarking deep reinforcement learning for continuous control. In: International Conference on Machine Learning, pp. 1329–1338 (2016). PMLR
- [5] Lillicrap, T.P., Hunt, J.J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., Wierstra, D.: Continuous control with deep reinforcement learning. In: ICLR (Poster) (2016)
- [6] Zhao, X., Zhang, D., Han, L., Zhang, T., Xu, B.: ODE-based Recurrent Model-free Reinforcement Learning for POMDPs. *Advances in Neural Information Processing Systems* **25**, 27159–27170 (2023)
- [7] Tang, G., Kumar, N., Yoo, R., Michmizos, K.: Deep Reinforcement Learning with Population-Coded Spiking Neural Network for Continuous Control. In: Conference on Robot Learning, pp. 2016–2029 (2021). PMLR

- [8] Rumelhart, D.E., Hinton, G.E., *et al.*: Learning representations by back-propagating errors. *Nature* **323**(6088), 533–536 (1986)
- [9] Salem, M.S.Z.: Biological networks: an introductory review. *Journal Of Proteomics And Genomics Research* **2**(1), 41–111 (2018)
- [10] Blanchini, F., Franco, E.: Structurally robust biological networks. *BMC systems biology* **5**(1), 1–14 (2011)
- [11] Whitacre, J.M.: Biological robustness: paradigms, mechanisms, and systems principles. *Frontiers in genetics* **3**, 67 (2012)
- [12] Attwell, D., Laughlin, S.B.: An energy budget for signaling in the grey matter of the brain. *Journal of Cerebral Blood Flow & Metabolism* **21**(10), 1133–1145 (2001)
- [13] Howarth, C., Gleeson, P., Attwell, D.: Updated energy budgets for neural computation in the neocortex and cerebellum. *Journal of Cerebral Blood Flow & Metabolism* **32**(7), 1222–1232 (2012)
- [14] Cox, D.D., Dean, T.: Neural networks and neuroscience-inspired computer vision. *Current Biology* **24**(18), 921–929 (2014)
- [15] Bidaye, S.S., Bockemühl, T., Büschges, A.: Six-legged walking in insects: how CPGs, peripheral feedback, and descending signals generate coordinated and adaptive motor rhythms. *Journal of neurophysiology* **119**(2), 459–475 (2018)
- [16] Chen, D., Peng, P., Huang, T., Tian, Y.: Deep Reinforcement Learning with Spiking Q-learning. *arXiv preprint arXiv:2201.09754* (2022)
- [17] Black, J.E., Isaacs, K.R., Anderson, B.J., Alcantara, A.A., Greenough, W.T.: Learning causes synaptogenesis, whereas motor activity causes angiogenesis, in cerebellar cortex of adult rats. *Proceedings of the National Academy of Sciences* **87**(14), 5568–5572 (1990)
- [18] Dayan, E., Cohen, L.G.: Neuroplasticity subserving motor skill learning. *Neuron* **72**(3), 443–454 (2011)
- [19] Kleim, J.A., Barbay, S., Cooper, N.R., Hogg, T.M., Reidel, C.N., Remple, M.S., Nudo, R.J.: Motor learning-dependent synaptogenesis is localized to functionally reorganized motor cortex. *Neurobiology of learning and memory* **77**(1), 63–77 (2002)
- [20] Watkins, C.J., Dayan, P.: Q-learning. *Machine learning* **8**(3), 279–292 (1992)
- [21] Sehnke, F., Osendorfer, C., Rückstieß, T., Graves, A., Peters, J., Schmidhuber, J.: Parameter-exploring policy gradients. *Neural Networks* **23**(4), 551–559 (2010)
- [22] Schulman, J., Wolski, F., Dhariwal, P., Radford, A., Klimov, O.: Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347* (2017)
- [23] Van Hasselt, H., Guez, A., Silver, D.: Deep reinforcement learning with double q-learning. In: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 30 (2016)
- [24] Wang, Z., Schaul, T., Hessel, M., Hasselt, H., Lanctot, M., Freitas, N.: Dueling network architectures for deep reinforcement learning. In: *International Conference on Machine Learning*, pp. 1995–2003 (2016). PMLR

- [25] Bellemare, M.G., Naddaf, Y., Veness, J., Bowling, M.: The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research* **47**, 253–279 (2013)
- [26] Mnih, V., Badia, A.P., Mirza, M., Graves, A., Lillicrap, T., Harley, T., Silver, D., Kavukcuoglu, K.: Asynchronous methods for deep reinforcement learning. In: *International Conference on Machine Learning*, pp. 1928–1937 (2016). PMLR
- [27] Schulman, J., Levine, S., Abbeel, P., Jordan, M., Moritz, P.: Trust region policy optimization. In: *International Conference on Machine Learning*, pp. 1889–1897 (2015). PMLR
- [28] Fujimoto, S., Hoof, H., Meger, D.: Addressing function approximation error in actor-critic methods. In: *International Conference on Machine Learning*, pp. 1587–1596 (2018). PMLR
- [29] Haarnoja, T., Zhou, A., Abbeel, P., Levine, S.: Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In: *International Conference on Machine Learning*, pp. 1861–1870 (2018). PMLR
- [30] Maass, W.: Networks of spiking neurons: the third generation of neural network models. *Neural Networks* **10**(9), 1659–1671 (1997)
- [31] Zhang, D., Zhang, T., Jia, S., Wang, Q., Xu, B.: Recent advances and new frontiers in spiking neural networks. In: *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022*, pp. 5670–5677 (2022). <https://doi.org/10.24963/ijcai.2022/790>. <https://doi.org/10.24963/ijcai.2022/790>
- [32] O’Brien, M.J., Srinivasa, N.: A spiking neural model for stable reinforcement of synapses based on multiple distal rewards. *Neural Computation* **25**(1), 123–156 (2013)
- [33] Yuan, M., Wu, X., Yan, R., Tang, H.: Reinforcement learning in spiking neural networks with stochastic and deterministic synapses. *Neural computation* **31**(12), 2368–2389 (2019)
- [34] Mahadevuni, A., Li, P.: Navigating mobile robots to target in near shortest time using reinforcement learning with spiking neural networks. In: *2017 International Joint Conference on Neural Networks (IJCNN)*, pp. 2243–2250 (2017). IEEE
- [35] Doya, K.: Reinforcement learning in continuous time and space. *Neural computation* **12**(1), 219–245 (2000)
- [36] Frémaux, N., Sprekeler, H., Gerstner, W.: Reinforcement learning using a continuous time actor-critic framework with spiking neurons. *PLoS computational biology* **9**(4), 1003024 (2013)
- [37] Florian, R.V.: Reinforcement learning through modulation of spike-timing-dependent synaptic plasticity. *Neural computation* **19**(6), 1468–1502 (2007)
- [38] Frémaux, N., Gerstner, W.: Neuromodulated spike-timing-dependent plasticity, and theory of three-factor learning rules. *Frontiers in neural circuits* **9**, 85 (2016)

- [39] Bellec, G., Scherr, F., Subramoney, A., Hajek, E., Salaj, D., Legenstein, R., Maass, W.: A solution to the learning dilemma for recurrent networks of spiking neurons. *Nature communications* **11**(1), 1–15 (2020)
- [40] Kim, J., Kwon, D., Woo, S.Y., Kang, W.-M., Lee, S., Oh, S., Kim, C.-H., Bae, J.-H., Park, B.-G., Lee, J.-H.: On-chip trainable hardware-based deep Q-networks approximating a backpropagation algorithm. *Neural Computing and Applications* **33**(15), 9391–9402 (2021)
- [41] Akl, M., Sandamirskaya, Y., Walter, F., Knoll, A.: Porting Deep Spiking Q-Networks to neuromorphic chip Loihi. In: *International Conference on Neuromorphic Systems 2021*, pp. 1–7 (2021)
- [42] Liu, G., Deng, W., Xie, X., Huang, L., Tang, H.: Human-Level Control through Directly-Trained Deep Spiking Q-Networks. *arXiv preprint arXiv:2201.07211* (2021)
- [43] Patel, D., Hazan, H., Saunders, D.J., Siegelmann, H.T., Kozma, R.: Improved robustness of reinforcement learning policies upon conversion to spiking neuronal network platforms applied to Atari Breakout game. *Neural Networks* **120**, 108–115 (2019)
- [44] Tan, W., Patel, D., Kozma, R.: Strategy and Benchmark for Converting Deep Q-Networks to Event-Driven Spiking Neural Networks. In: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, pp. 9816–9824 (2021)
- [45] Tang, G., Kumar, N., Michmizos, K.P.: Reinforcement co-learning of deep and spiking neural networks for energy-efficient mapless navigation with neuromorphic hardware. In: *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 6090–6097 (2020). IEEE
- [46] Zhang, D., Zhang, T., Jia, S., Cheng, X., Xu, B.: Population-coding and Dynamic-neurons improved Spiking Actor Network for Reinforcement Learning. *arXiv preprint arXiv:2106.07854* (2021)
- [47] Zhang, D., Zhang, T., Jia, S., Xu, B.: Multi-Sacle Dynamic Coding Improved Spiking Actor Network for Reinforcement Learning. In: *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelveth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22 - March 1, 2022*, pp. 59–67. AAAI Press, ??? (2022). <https://ojs.aaai.org/index.php/AAAI/article/view/19879>
- [48] Seung, S.: *Connectome: How the Brain’s Wiring Makes Us Who We Are*. HMH, ??? (2012)
- [49] Eichler, K., Li, F., Litwin-Kumar, A., Park, Y., Andrade, I., Schneider-Mizell, C.M., Saumweber, T., Huser, A., Eschbach, C., Gerber, B., *et al.*: The complete connectome of a learning and memory centre in an insect brain. *Nature* **548**(7666), 175–182 (2017)
- [50] Takemura, S.-y., Aso, Y., Hige, T., Wong, A., Lu, Z., Xu, C.S., Rivlin, P.K., Hess, H., Zhao, T., Parag, T., *et al.*: A connectome of a learning and memory center in the adult *Drosophila* brain. *Elife* **6**, 26975 (2017)

- [51] Varshney, L.R., Chen, B.L., Paniagua, E., Hall, D.H., Chklovskii, D.B.: Structural properties of the *Caenorhabditis elegans* neuronal network. *PLoS computational biology* **7**(2), 1001066 (2011)
- [52] White, J.G., Southgate, E., Thomson, J.N., Brenner, S., *et al.*: The structure of the nervous system of the nematode *Caenorhabditis elegans*. *Philos Trans R Soc Lond B Biol Sci* **314**(1165), 1–340 (1986)
- [53] Anwar, A., Cardwell, S.G., Musuvathy, S., Severa, W.M., Vineyard, C.M.: Evolving Spiking Circuit Motifs Using Weight Agnostic Neural Networks. *CSRI*, 3 (2020)
- [54] Huttenlocher, P.R.: Morphometric study of human cerebral cortex development. *Neuropsychologia* **28**(6), 517–527 (1990)
- [55] Tierney, A.L., Nelson III, C.A.: Brain development and the role of experience in the early years. *Zero to three* **30**(2), 9 (2009)
- [56] Bruer, J.T.: Neural connections: Some you use, some you lose. *The Phi Delta Kappan* **81**(4), 264–277 (1999)
- [57] Dayan, P., Abbott, L.F., *et al.*: Theoretical neuroscience: computational and mathematical modeling of neural systems. *Journal of Cognitive Neuroscience* **15**(1), 154–155 (2003)
- [58] Akopyan, F., Sawada, J., *et al.*: Truenorth: Design and tool flow of a 65 mw 1 million neuron programmable neurosynaptic chip. *TCAD* **34**(10), 1537–1557 (2015)
- [59] Davies, M., Srinivasa, N., *et al.*: Loihi: A neuromorphic manycore processor with on-chip learning. *IEEE Micro* **38**(1), 82–99 (2018)
- [60] Geweke, J.: Antithetic acceleration of Monte Carlo integration in Bayesian inference. *Journal of Econometrics* **38**(1-2), 73–89 (1988)
- [61] Brockhoff, D., Auger, A., Hansen, N., Arnold, D.V., Hohm, T.: Mirrored sampling and sequential selection for evolution strategies. In: *International Conference on Parallel Problem Solving from Nature*, pp. 11–21 (2010). Springer
- [62] Salimans, T., Ho, J., Chen, X., Sidor, S., Sutskever, I.: Evolution strategies as a scalable alternative to reinforcement learning. *arXiv preprint arXiv:1703.03864* (2017)
- [63] Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., Zaremba, W.: Openai gym. *arXiv preprint arXiv:1606.01540* (2016)
- [64] Todorov, E., Erez, T., Tassa, Y.: Mujoco: A physics engine for model-based control. In: *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 5026–5033 (2012). IEEE
- [65] Achiam, J.: *Spinning Up in Deep Reinforcement Learning* (2018)
- [66] Kangin, D., Pugeault, N.: On-policy trust region policy optimisation with replay buffers. *arXiv preprint arXiv:1901.06212* (2019)
- [67] Henderson, P., Islam, R., Bachman, P., Pineau, J., Precup, D., Meger, D.: Deep reinforcement learning that matters. In: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32 (2018)
- [68] Rathi, N., Roy, K.: DIET-SNN: A low-latency spiking neural network

22 REFERENCES

- with direct input encoding and leakage and threshold optimization. *IEEE Transactions on Neural Networks and Learning Systems* (2021)
- [69] Horowitz, M.: 1.1 computing's energy problem (and what we can do about it). In: *2014 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*, pp. 10–14 (2014). IEEE

Appendix

The Hyper-parameter Configurations and Optimization Process of PPO

For each MuJoCo continuous control task, we ran $1e3$ iterations, $1e3$ steps per iteration, and a total of $1e6$ steps of environment interaction. The deep policy network was trained in 25 epochs per iteration with a mini-batch size of 100, using Adam optimizer with a learning rate of $1e-5$. The training configurations of the deep value network were the same as the deep policy network, but the learning rate was $1e-4$. The discount factor γ for reward was 0.99, and the discount factor λ used in generalized advantage estimation was 0.95. The clip ratio ϵ_{clip} for PPO was 0.2. We presented the optimization process of PPO in Algorithm 3.

Algorithm 3 The optimization process of PPO.

Initialize deep policy network $\pi(\mathbf{a}|\mathbf{o})$ and deep value network $V(\mathbf{o})$ with parameter ϕ_π and ϕ_V .

for $iter = 1$ to $1e3$

Run π_{ϕ_π} to collect a set of episodes $\mathcal{D} = \{\epsilon^k\}$ containing $|\mathcal{D}|$ episodes, each ϵ^k is an episode contain $|\epsilon^k|$ data (samples or steps), $\epsilon^k := \{(\mathbf{o}^{k,t}, \mathbf{a}^{k,t}, r^{k,t+1}, \mathbf{o}^{k,t+1})\}$, $t \in [|\epsilon^k|]$.

Compute cumulative reward $R^{k,t}$ for each step t in every episode k using discount factor γ .

Update deep value network by minimizing the mean-square error:

$$\phi_V \leftarrow \arg \min_{\phi_V} \frac{1}{\sum_k |\epsilon^k|} \sum_{\epsilon^k \in \mathcal{D}} \sum_{t=0}^{|\epsilon^k|} (V(\mathbf{o}^{k,t}) - R^{k,t})^2$$

Estimate advantage $A^{k,t}$ for each step t in every episode k using generalized advantage estimation and deep value network $V_{\phi_V}(\mathbf{o})$.

Update the deep policy network by minimizing the following objective (the minimization is solved using Adam optimizer):

$$\phi_\pi \leftarrow \arg \min_{\phi'_\pi} \frac{1}{\sum_k |\epsilon^k|} \left[\sum_{\epsilon^k \in \mathcal{D}} \sum_{t=0}^{|\epsilon^k|} \min(r_{\phi'_\pi}(\mathbf{a}^{k,t}|\mathbf{o}^{k,t})A^{k,t}, g(r_{\phi'_\pi}(\mathbf{a}^{k,t}|\mathbf{o}^{k,t}))A^{k,t}) \right]$$

where $r_{\phi'_\pi}(\mathbf{a}^{k,t}|\mathbf{o}^{k,t}) := \frac{\pi_{\phi'_\pi}(\mathbf{a}^{k,t}|\mathbf{o}^{k,t})}{\pi_{\phi_\pi}(\mathbf{a}^{k,t}|\mathbf{o}^{k,t})}$, $g(r) := \text{clip}(r_{\phi'_\pi}(\mathbf{a}^{k,t}|\mathbf{o}^{k,t}), 1 - \epsilon_{\text{clip}}, 1 + \epsilon_{\text{clip}})$.

Return: Final learned deep policy network.



Duzhen Zhang received the B.Sc. degree in software engineering from Shandong University, China in 2019. He is a Ph.D. candidate in both the Institute of Automation Chinese Academy of Sciences and the University of Chinese Academy of Sciences. His current interests include theoretical research on reinforcement learning, natural language processing, and Spiking Neural Networks.

E-mail: zhangduzhen2019@ia.ac.cn ORCID iD: 0000-0002-4280-431X



Tielin Zhang received the Ph.D. degree from the Institute of Automation Chinese Academy of Sciences, Beijing, China, in 2016. He is an Associate Professor in the Research Center for Brain-Inspired Intelligence, Institute of Automation, Chinese Academy of Sciences, Beijing, China. His current interests include theoretical research on neural dynamics and Spiking Neural Networks (more information is in <https://bii.ia.ac.cn/~tielin.zhang/>).

E-mail: tielin.zhang@ia.ac.cn (Corresponding author) ORCID iD: 0000-0002-5111-9891



Shuncheng Jia is a Ph.D. candidate in both the Institute of Automation Chinese Academy of Sciences and the University of Chinese Academy of Sciences. His current interests include theoretical research on neural dynamics, auditory signal processing, and Spiking Neural Networks.

E-mail: jiashuncheng2020@ia.ac.cn



Qingyu Wang is an M.D. candidate in both the Institute of Automation Chinese Academy of Sciences and the University of Chinese Academy of Sciences. His current interests include theoretical research on neural dynamics and Spiking Neural Networks.

E-mail: wangqingyu2022@ia.ac.cn



Bo Xu is a professor, the director of the Institute of Automation Chinese Academy of Sciences, and also deputy director of the Center for Excellence in Brain Science and Intelligence Technology, Chinese Academy of Sciences. His main research interests include brain-inspired intelligence, brain-inspired cognitive models, natural language processing and understanding, brain-inspired robotics.

E-mail: xubo@ia.ac.cn (Corresponding author) ORCID iD: 0000-0002-1111-1529