

Hardware implementation of photonic neuromorphic autonomous navigation

Yonghang Chen

Xidian University

Shuiying Xiang

`syxiang@xidian.edu.cn`

Xidian University

Xintao Zeng

Xidian University

Mengting Yu

Xidian University

Tao Zou

Xidian University

Shangxuan Shi

Xidian University

Xingxing Guo

Xidian University

Yanan Han

Xidian University

Yahui Zhang

Xidian University

Yue Hao

Xidian University

Article

Keywords: Photonic neuromorphic computing, spiking reinforcement learning, photonic spiking neuron, autonomous navigation

Posted Date: March 11th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-8874344/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.
[Read Full License](#)

Additional Declarations: No competing interests reported.

Hardware implementation of photonic neuromorphic autonomous navigation

Yonghang Chen¹, Shuiying Xiang^{1*}, Xintao Zeng¹, Mengting Yu¹, Tao Zou¹,
Shangxuan Shi¹, Xingxing Guo¹, Yanan Han¹, Yahui Zhang^{1*}, and Yue Hao¹

¹State Key Laboratory of Integrated Service Networks, Xidian University, Xi'an 710071, China.

*Correspondence and requests for materials should be addressed to S.Y. X. (email: syxiang@xidian.edu.cn) and/or Y.H. Z. (email: zhangyahui@xidian.edu.cn)

KEYWORDS: Photonic neuromorphic computing, spiking reinforcement learning, photonic spiking neuron, autonomous navigation

ABSTRACT: Reinforcement learning (RL) is a core technology enabling the transition of artificial intelligence (AI) from perception to decision-making, but its deployment on conventional electronic hardware suffers from high latency and energy consumption imposed by the von Neumann architecture. Here, we propose a photonic spiking twin delayed deep deterministic policy gradient (TD3) reinforcement learning architecture for neuromorphic autonomous navigation and experimentally validate it on a distributed feedback laser with a saturable absorber (DFB-SA) array. The hybrid architecture integrates a photonic spiking Actor network with dual continuous-valued Critic networks, where the final nonlinear spiking activation layer of the Actor is deployed on the DFB-SA laser

array. In autonomous navigation tasks, the system achieves an average reward of 58.22 ± 17.29 and a success rate of $80\% \pm 8.3\%$. Hardware-software co-inference demonstrates an estimated energy consumption of 0.78 nJ/inf and an ultra-low latency of 191.20 ps/inf, with co-inference error rates of 0.051% and 0.059% in task scenarios with and without obstacle interference, respectively. Simulations for error-activated channels show full agreement with the expected responses, validating the dynamic characteristics of the DFB-SA laser. The architecture shows strong potential for integration with large-scale photonic linear computing chips, enabling fully-functional photonic computation and low-power, low-latency neuromorphic autonomous navigation.

1. Introduction

Reinforcement learning (RL) is a class of machine learning methods in which agents learn optimal policies through continuous interaction with their environment. Its theoretical foundation combines the trial-and-error learning mechanism from behavioral psychology with the reward optimization principle from optimal control theory ^[1-4]. As a core paradigm for sequential decision-making, RL has driven major advances across domains such as autonomous driving ^[5-7], strategic gaming ^[8-11], robotic control ^[12-14], and many other applications ^[15-20]. Recent algorithmic developments, including Twin Delayed Deep Deterministic Policy Gradient (TD3) ^[21], Proximal Policy Optimization (PPO) ^[22], Deep Deterministic Policy Gradient (DDPG) ^[23], and Deep Q-Network (DQN) ^[15], have significantly improved the capability of RL

to handle high-dimensional state spaces, enhance sample efficiency, and stabilize policy learning. However, the deployment of these algorithms on conventional von Neumann architectures remains limited by the intrinsic separation of memory and computation, which introduces substantial latency and energy overheads, critically constraining the real-time performance of RL in applications such as autonomous navigation ^[24].

As an emerging paradigm for high-performance computing, photonic computing has demonstrated remarkable potential due to its ultrafast signal propagation and inherent parallelism ^[25-28]. By utilizing photons instead of electrons as information carriers, photonic computing enables highly energy-efficient and low-latency linear operations through the intrinsic speed and parallel nature of light ^[29-43]. Existing photonic linear computing implementations primarily include microring resonator (MRR)-based weight banks ^[31-33], cascaded Mach-Zehnder modulators (MZMs) ^[34, 35], phase-change material (PCM) crossbar arrays ^[36, 37], and Mach-Zehnder interferometer (MZI) meshes ^[38-43]. However, despite recent advances ^[44-46], realizing a fully-functional photonic computing system remains challenging—particularly in implementing nonlinear operations. Against this backdrop, photonic spiking neural networks (PSNNs) have emerged as a promising computational framework that integrates the high-speed, energy-efficient processing of photonic computing with the event-driven and sparse dynamics of SNNs ^[47-55]. In 2019, Feldmann et al. experimentally demonstrated the first fully integrated optical spiking neuromorphic network based on PCMs, achieving pattern recognition on optically encoded letters with an energy consumption of approximately 60 pJ per synaptic weight update ^[47]. In

2020, Robertson et al. reported a photonic spiking neuron based on a vertical-cavity surface-emitting laser (VCSEL), achieving ultrafast detection of both oriented and omnidirectional edges in digital images ^[48]. In 2021, Xiang et al. proposed an all-optical SNN entirely based on VCSELs for supervised learning, theoretically demonstrating optical digit classification while eliminating the laterally inhibitory mechanism that is difficult to realize in optical systems ^[49]. In 2022, Lee et al. proposed and implemented an event-driven optoelectronic spiking neuron based on the Izhikevich model, achieving 97% accuracy and an energy efficiency of 50 trillion operations per second per joule (TOPS/J) on the MNIST handwritten digit recognition task ^[50]. In 2023, Xiang et al. fabricated and experimentally validated a photonic spiking neuron chip based on a Fabry-Perot laser with a saturable absorber (FP-SA), which realized single-neuron pattern classification through hardware-software co-computation ^[51]. In the same year, the team further developed and demonstrated a photonic synaptic core chip based on a DFB-SA laser, which, for the first time, simultaneously implemented linear weighting and nonlinear spiking activation within a single device, achieving 87% accuracy on the MNIST dataset ^[52]. More recently, Talukder et al. constructed a large-scale PSNN comprising 40,000 neurons, achieving bio-inspired excitability using a modified Ikeda map with slow inhibitory feedback. The network reached 77.5% accuracy on MNIST using only 8.5% of the neurons and represented the first hardware-compatible demonstration of simultaneous perturbation stochastic approximation (SPSA) training ^[53]. These advances collectively demonstrate the superior energy efficiency, processing speed, and sparsity of PSNNs. Nevertheless, their application to autonomous navigation

remains unexplored. Gao et al. introduced a photonic spiking neuron-based obstacle-avoidance system employing a FP laser, which used rate coding to emulate biological monocular vision and experimentally verified the feasibility of ultrafast, low-power robotic obstacle avoidance ^[54]. However, the decision-making process of this system relies on predefined thresholds, lacking autonomous interaction and online learning between the RL agent and the environment, thus preventing full-process neuromorphic autonomous navigation. Integrating the ultrafast perception-decision capability of PSNNs with the sequential decision-optimization ability of RL is therefore expected to provide a promising pathway toward intelligent, self-adaptive, low-power and low-latency neuromorphic autonomous navigation systems.

This paper presents a photonic spiking TD3 reinforcement learning architecture for neuromorphic autonomous navigation, enabling highly energy-efficient intelligent decision-making through hardware-software co-inference. The main contributions of this work are as follows:

1. We propose a photonic spiking TD3 reinforcement learning framework that integrates a spiking-based Actor network with dual continuous-valued Critic networks. The algorithm is optimized under the hardware constraints of the photonic neuron chips, where the Actor network adopts positive-weight and bias-free structures to ensure hardware compatibility. In autonomous navigation tasks, the proposed architecture achieves an average reward of 58.22 ± 17.29 and a navigation success rate of $80\% \pm 8.3\%$ in software simulation, demonstrating the feasibility of employing PSNNs for autonomous navigation.

-
2. The final nonlinear spiking activation layer of the Actor network is deployed on a DFB-SA laser array. The system achieves an estimated inference energy of 0.78 nJ/inf and an ultra-low latency of 191.20 ps/inf, highlighting the intrinsic low-power and low-latency advantages of photonic computing for reinforcement learning.
 3. Hardware-software co-inference experiments are conducted under two scenarios—with and without obstacle interference. The co-inference error rates are 0.051% and 0.059%, respectively, validating the high accuracy and reliability of PSNN-based reinforcement learning for neuromorphic autonomous navigation. Additionally, simulations based on the Yamada model using algorithm-generated inputs for error-activated channels show full agreement with the expected spiking responses, thereby theoretically validating the dynamic characteristics of the DFB-SA laser.

2. Methods

2.1. Photonic Spiking TD3 Reinforcement Learning Architecture

The proposed photonic spiking reinforcement learning architecture is based on the TD3 algorithm, an enhanced variant of the DDPG method. TD3 is designed to mitigate the overestimation bias of Q-values introduced by function approximation errors in the Actor-Critic framework, thereby enhancing both algorithmic stability and overall performance^[21]. This is achieved through three key innovations: employing twin Critic networks (clipped double Q-learning), introducing target policy smoothing regularization, and adopting delayed policy updates. These mechanisms collectively stabilize training, improve policy optimization accuracy, and enhance performance in

complex continuous control tasks. Owing to its superior stability and scalability, TD3 has been widely applied in domains requiring precise decision-making within high-dimensional state-action spaces, including robotic control^[56, 57], autonomous driving^[58, 59], and Unmanned Aerial Vehicles (UAV) communications^[60, 61].




Figure 1. Overall architecture and implementation details of the photonic spiking TD3 reinforcement learning system. (a) Overall architecture of the photonic spiking TD3 reinforcement learning system. (b) Hardware-software co-inference framework. (c) Real-time navigation path map. (d) Onboard camera view of the robot.

Figure 1(a) illustrates the proposed photonic spiking TD3 reinforcement learning architecture. The reinforcement learning agent comprises a PSNN-based Actor network and dual Artificial Neural Network (ANN)-based Critic networks. The photonic spiking

Actor network generates two continuous actions—the linear and angular velocities of the robot—based on the current state, while the Critic network evaluates the corresponding state-action Q-values. Using an ANN-based Critic offers more stable gradient estimation and faster convergence, which complements the event-driven and energy-efficient characteristics of the photonic spiking Actor network, resulting in a balanced learning framework. We considered a continuous control task in the Gazebo simulation environment ^[62], where the robot perceives obstacles through Light Detection and Ranging (LiDAR) sensors and performs mapless dynamic obstacle avoidance and goal-directed navigation through reinforcement learning. The overall system comprises three core modules: the agent, the Gazebo simulation environment, and the differential-drive control system, interconnected through the Robot Operating System (ROS) framework to enable LiDAR-based, mapless dynamic navigation ^[63]. Detailed configurations of the simulation environment, state space, and action space are provided in Supplementary Section 1.

In the autonomous navigation tasks, the photonic spiking Actor network processes real-time state data from the Gazebo simulation environment. The static state inputs are first transformed via spike encoding by expanding them along the temporal dimension and replicating them over T timesteps to match the time-series processing paradigm of SNNs. The encoded sequences are then propagated through the photonic spiking Actor network, where each hidden layer employs leaky integrate-and-fire (LIF) neurons to produce temporal spiking responses. These spiking outputs are subsequently decoded by temporal aggregation (i.e., averaging over time) to recover static representations,

which are normalized using a Tanh function to generate continuous velocity commands. During training, Gaussian noise is injected into actions to encourage exploration. In addition, a forced exploration strategy is adopted: when LiDAR detects nearby obstacles, the agent executes random actions with a certain probability to escape local optima. After each interaction, the environment returns rewards, next states, and termination signals, which are stored in a replay buffer. For policy evaluation, mini-batches randomly sampled from the replay buffer are processed by dual-ANN Critic networks ($800 \times 600 \times 1$) with a twin-Q architecture. Target networks are employed to stabilize Q-value estimation by minimizing the temporal-difference error. Target actions are generated by an independent target photonic spiking Actor network with clipped noise to mitigate Q-value overestimation; this target Actor follows the same spike encoding–LIF activation–spike decoding pipeline as the primary Actor. The Actor network is updated every two training steps by maximizing the expected state–action value using gradients provided by the Critic. This delayed policy update strategy, together with soft target network synchronization, ensures stable and efficient convergence. The pseudocode of the complete algorithm workflow is provided in Supplementary Section 2.

2.2. Hardware-Aware Algorithm Optimization and Photonic Implementation

To ensure hardware compatibility and general applicability, the photonic spiking Actor network is co-designed with hardware-aware constraints for deployment on generic photonic linear computing platforms. While current photonic linear processors

have experimentally demonstrated matrix operation scales up to 128×128 ^[43], we configure the Actor network to $24 \times 128 \times 128 \times 2$ as a representative and hardware-compatible example. This configuration illustrates the scalability and implementability of our approach within the capabilities of existing photonic linear computing architectures, while preserving high control accuracy. During software pretraining, the hidden-layer weights are constrained to be positive, and all bias terms are removed, ensuring that the trained parameters can be directly mapped onto photonic linear computing hardware. The final nonlinear spiking activation layer is deployed in hardware using a DFB-SA laser array, which performs optical nonlinear computation through laser spiking dynamics. This design enables seamless integration with large-scale photonic matrix-vector multiplication chips, realizing a fully-functional photonic computing architecture in which both linear and nonlinear operations are executed entirely in the optical domain. Simulation results verify the performance and stability of this architecture, providing theoretical guidance for subsequent hardware deployment and co-inference experiments.

2.3. Hardware-Software Co-Inference Architecture

To validate the effectiveness of the proposed photonic spiking TD3 reinforcement learning architecture, a hardware-software co-inference scheme is employed for the photonic spiking Actor network. Figure 1(b) presents the overall co-inference architecture. The process begins with software pretraining, where the photonic spiking Actor network is trained from scratch using a surrogate gradient method with a single

time step ($T = 1$) to obtain a pretrained model that already incorporates the hardware constraints outlined in Section 2.2. Following pretraining, the final nonlinear spiking activation layer is deployed in hardware using a DFB-SA laser array, which realizes optical nonlinear computation via laser spiking dynamics. The corresponding input-output data from the pretrained software model are fed into the photonic hardware to evaluate its response. The experimentally measured optical outputs are compared with the software-generated counterparts to quantify the co-inference error rate, validating the consistency between optical and digital inference. Figures 1(c) and 1(d) show the real-time navigation trajectory and onboard camera view of the robot, respectively, providing an intuitive visualization of the simulation environment and the autonomous navigation behavior of the robot.

2.4 Experimental Setup

The final nonlinear spiking activation layer of the photonic spiking Actor network is experimentally deployed using a DFB-SA laser array. The fabrication detail of the DFB-SA laser is like that reported in Ref.[64]. Figure 2(a) illustrates the experimental setup of this layer, including a micrograph of the DFB-SA laser chip used in the experiment. In the schematic, red lines denote optical paths, while blue dashed lines correspond to electronic paths.




Figure 2. Experimental setup and properties of a DFB-SA laser-based photonic spiking neuron. (a) Experimental setup of the final activation layer in the photonic spiking Actor network. DC: digital computer; FPGA: field-programmable gate array; EA: electronic amplifier; TL: tunable laser; MZM: Mach-Zehnder modulator; VS: voltage source; PC: polarization controller; OC: optical coupler; PD: photodetectors; OSC: oscilloscope; CIRC: circulator; LDC: laser diode controller. (b) The PI curve of DFB-SA laser. (c) The optical spectra of the TL and DFB-SA laser. (d) The optical spectrum of the DFB-SA laser. (e1) - (g2) The neuron-like response of the DFB-SA laser.

A tunable laser (TL) provides the optical carrier, which is injected into the Mach-Zehnder modulator (MZM, Fujitsu FTM7928FB), while a voltage source (VS) supplies the driving voltage. Electronic signals generated by a field-programmable gate array

(FPGA), controlled by a digital computer (DC), are amplified through an electronic amplifier (EA) and injected into the MZM for modulation. The modulated optical signal is then injected into the DFB-SA laser via a three-port optical circulator (CIRC), with a polarization controller (PC) adjusting the polarization state. A laser diode controller (LDC, ILX Lightwave LDC3724B) provides high-precision temperature control and low-noise bias currents to the gain region of the DFB-SA laser, while a voltage source applies a reverse bias voltage to the saturable absorber. The output of the DFB-SA laser is detected by a photodetector (PD, Agilent HP11982A) and recorded with an oscilloscope (OSC, Keysight DSOZ592A) and is subsequently analyzed along with the modulated optical signal tapped from the optical coupler (OC). Figure 2(b) shows the PI curve of the DFB-SA laser, with a threshold current of 18 mA. Figure 2(c) shows the spectra of the free-running tunable laser and the DFB-SA laser, from which the side-mode suppression ratios (SMSRs) are measured to be approximately 55.58 dB and 46.48 dB, respectively, demonstrating excellent single-mode performance and beam quality, which are critical for the experiment. Figure 2(d) shows the output spectrum of the DFB-SA laser with injections generated from TL. Figures 2(e1)-(g2) demonstrate the key neuro-mimetic dynamics of the DFB-SA laser—threshold, temporal integration, and refractory period—thereby validating its capability to emulate complex photonic spiking neuron behaviors. During the experiment, the DFB-SA laser temperature is maintained at 22 °C, the reverse bias voltage of the saturable absorber is set to 0.76 V, the bias current of the gain region is set to 27.9 mA, and the injection power P_{inj} is 97.96 μ W.

2.5 Theoretical Model of DFB-SA Laser Based on Yamada Model

In 2002, Minoru Yamada proposed a theoretical analysis of self-sustained pulsation phenomena and investigated the application of rate equations in narrowband-structured semiconductor lasers. The model can be used to analyze the dynamic characteristics of the DFB-SA laser^[65]. Here, we modify the Yamada model to simulate a photonic spiking neuron based on the DFB-SA laser. The mathematical formulation of the Yamada model can be expressed as follows:

$$\frac{dN_1}{dt} = \frac{I}{eV_1} - \frac{N_1}{\tau_s} - \frac{\Gamma_1 g_1}{V_1} (N_1 - N_{g_1}) S \quad (1)$$

$$\frac{dN_2}{dt} = -\frac{N_2}{\tau_s} - \frac{\Gamma_2 g_2}{V_2} (N_2 - N_{g_2}) S \quad (2)$$

$$\frac{dS}{dt} = [\Gamma_1 g_1 (N_1 - N_{g_1}) + \Gamma_2 g_2 (N_2 - N_{g_2})] S - \frac{S - S_{\text{ext}}}{\tau_{\text{ph}}} + \beta_{\text{sp}} \frac{N_1 V_1}{\tau_s} \quad (3)$$

Here, subscripts 1 and 2 refer to the gain region and saturable absorber, respectively. N denotes the carrier density, g the gain coefficient, Γ the confinement factor, V the cavity volume, and N_g the transparency carrier density. S_{ext} denotes the number of externally injected photons, β_{sp} is the spontaneous emission coefficient, τ_{ph} is the photon lifetime, τ_s is the carrier lifetime, and I represents the bias current applied to the gain region. For simplicity, carrier diffusion between the gain region and saturable absorber is not considered in the model. The corresponding parameter values are listed in Table 1.

Table 1. The corresponding parameter values of DFB-SA laser

Parameters	Gain region	Saturable absorber
Gain coefficient g	$1.8794 \times 10^{11} \text{ m}^3/\text{s}$	$0.594 \times 10^{-11} \text{ m}^3/\text{s}$
Confinement factor Γ	0.08	0.08
Cavity volume V	$172.8 \times 10^{-18} \text{ m}^3$	$7.2 \times 10^{-18} \text{ m}^3$
Transparency carrier density N_g	$1.4 \times 10^{24} \text{ m}^{-3}$	$1.6 \times 10^{24} \text{ m}^{-3}$
Photon lifetime τ_{ph}	$2.56 \times 10^{-3} \text{ ns}$	$2.56 \times 10^{-3} \text{ ns}$
Carrier lifetime τ_s	1 ns	1 ns

3. Results

In this section, we present the experimental results of neuromorphic autonomous navigation using the photonic spiking TD3 reinforcement learning architecture with hardware-software co-inference. Figures 3(a1) and 3(a2) show the average reward curve and the navigation success rate curve of the baseline software algorithm, respectively. Since the training curves exhibit oscillations after convergence, the mean and standard deviation are computed over the last 20% of the training data, yielding an average reward of 58.22 ± 17.29 and a success rate of $80\% \pm 8.3\%$. The environment configuration and training parameters used in the algorithm are provided in Supplementary Section 3. After convergence, two task scenarios—with and without obstacle interference—are selected for experimental implementation of the final nonlinear spiking activation layer in the photonic spiking Actor network. Because the initial and target positions of the robot are randomized, the number of state-action sample pairs varies between tasks. Demonstration videos for both task scenarios (with and without obstacle interference) are provided in Supplementary Section 7.1 and Supplementary Section 7.2, corresponding to Supplementary Video 1 and Supplementary Video 2, respectively.

Figures 3(b1) and 3(b2) illustrate the start and goal positions, complete navigation trajectories, and the distributions of discrepancies between the expected and experimentally measured spiking activation outputs, where a value of 1 represents a false-positive spike (overactivation) and -1 represents a false-negative spike (missed activation). The experimental results confirm that the activation outputs of the DFB-SA array closely match the expected

activations. For clarity, the scenario without obstacle interference is referred to as Task 1 and that with obstacle interference as Task 2.




Figure 3. Software training results and experimental nonlinear outputs for autonomous navigation tasks. (a1) - (a2): Average reward and success rate for software training. (b1) - (b2): Distributions of discrepancies between the expected spiking activation outputs and the experimentally measured spiking activation outputs for two task scenarios. (c1) - (f1): Comparative examples of experimental

inputs and outputs for two task scenarios. The red curves represent the experimental inputs, the blue curve represent the experimentally measured spiking activation outputs, and the orange curve represents the expected spiking activation outputs.

In Task 1, only 3 discrepancies were observed among 5,888 samples (128×46), yielding a hardware-software co-inference error rate of 0.051%. In Task 2, 7 discrepancies occurred among 11,904 samples (128×93), corresponding to an error rate of 0.059%. Figures 3(c1) - (f2) present representative input-output comparisons for the two task scenarios. For Task 1, Figs. 3(c1) - (c2) display a correctly matched spiking activation in channel 1, while Figs. 3(d1) - (d2) show a mismatch in channel 69, where the DFB-SA laser failed to generate a spike at $0.43 \mu\text{s}$. For Task 2, Figs. 3(e1) - (e2) illustrate a correctly matched spiking activation in channel 1, whereas Figs. 3(f1) - (f2) reveal an overactivated spike at $1.12 \mu\text{s}$ and a missing spike at $1.42 \mu\text{s}$ in channel 114. The proposed architecture demonstrates good scalability and can be extended to a wide range of navigation tasks beyond the two tested scenarios. Complete input-output comparisons for all mismatched activations are provided in Supplementary Section 4.

Furthermore, we evaluate the inference latency and energy consumption of the proposed photonic spiking TD3 reinforcement learning architecture. The power consumption of a single DFB-SA laser chip is estimated as $1.097 \times 27.9 + 0.76 \times 1.93 = 32.07 \text{ mW}$, corresponding to a total power consumption of $P_{\text{DFB-SA}} = 4.10 \text{ W}$ for 128 devices, which demonstrates its low-power advantage^[52]. The self-pulsation frequency of a single DFB-SA laser chip is 5.23 GHz, corresponding to an electro-optic delay of $\tau_{\text{DFB-SA}} = 1 / 5.23 \text{ GHz} = 191.20 \text{ ps}$ ^[66]. Accordingly, the photonic spiking TD3 reinforcement learning architecture achieves an inference latency of 191.20 ps per inference, with an energy consumption of $E_{\text{inf}} = P_{\text{DFB-SA}} \times \tau_{\text{DFB-SA}} = 0.78 \text{ nJ/inf.}$

Compared with the 2083 ps/inf inference latency of the Eyeriss architecture ^[67] and the 11.98 $\mu\text{J}/\text{inf}$ energy consumption of the PopSAN architecture ^[68], the proposed approach exhibits substantially reduced latency and markedly improved energy efficiency.




Figure 4. Simulation results of the DFB-SA laser corresponding to the input data of error-activated channels for autonomous navigation tasks. (a1) - (a4): simulation results for Channel 69 in Task 1; (b1) - (b4): simulation results for Channel 114 in Task 2. (a1), (b1): externally injected optical signals; (a2), (b2): the outputs of the DFB-SA laser; (a3), (b3): the carrier density in the gain region; (a4), (b4): the carrier density in the saturable absorber.

Finally, based on the experimental results, we further simulate the algorithm input data corresponding to the error-activated channels using the Yamada model to theoretically validate the dynamic characteristics of the DFB-SA laser. Representative simulation results are shown in Fig. 4, where Figs. 4(a1) - (a4) correspond to Channel 69 in Task 1, and Figs. 4(b1) - (b4) correspond to Channel 114 in Task 2. As can be seen, the simulated spiking activations exhibit full agreement with the expected outputs, confirming the theoretical feasibility of achieving accurate DFB-SA laser activation. Complete simulation results for all error-activated channels in both tasks are provided

in Supplementary Section 5.

4. Discussion

In this section, we further investigate the performance of PSNNs with larger network scales and longer timesteps. To ensure a fair comparison, only the hidden-layer dimension and timestep length of the photonic spiking Actor network are varied, while all other parameters are kept identical, and the positive-weight constraint applied during photonic hardware-compatible training is removed. For the network-scale analysis, smaller networks showed stronger oscillations and poorer convergence (100×100); therefore, training are conducted with hidden-layer dimensions of 128×128 , 256×256 , and 512×512 under a fixed timestep of $T = 1$ for comparison. For the timestep analysis, the hidden-layer dimension was fixed at 128×128 , and the timestep was varied as $T = 1, 2$, and 4 . The average reward and navigation success rate curves of the baseline software algorithm are presented in Fig. 5.




Figure 5. Average reward and navigation success rate of software training under different network scales and timesteps. Hidden-layer dimensions: (a1) - (a2), 128×128 ; (b1) - (b2), 256×256 ; (c1) -

(c2), 512×512. Timesteps: (d1) - (d2), T = 1; (e1) - (e2), T = 2; (f1) - (f2), T = 4.

Subsequently, we analyze the training results. Because oscillations appear in the curves after convergence, the mean and standard deviation are computed over the last 20% of the training data, and the results are summarized in the form of “mean ± standard deviation,” as shown in Table 2. The results show that neither increasing the hidden-layer dimension nor extending the timestep effectively improves the performance of the PSNN. On the contrary, both factors tend to destabilize the training process, resulting in more pronounced oscillations in the reward curves. Possible explanations include gradient vanishing or explosion during backpropagation, overfitting caused by the larger number of network parameters, and degraded performance due to hardware resource limitations. Although larger networks and longer timesteps occasionally achieve higher peak performance—reaching up to 100% navigation accuracy in certain epochs—they also yield lower reward values during the early stages of training. In contrast, a hidden-layer dimension of 128×128 with a time step of T = 1 achieves stable and efficient training for autonomous navigation tasks while minimizing hardware resource demands.

Table 2. The training results for different hidden-layer dimensions and timesteps

Ablation study	Parameter	Average reward	Success rate
Scale	100 × 100	37.73 ± 20.58	69.7% ± 10.1%
	128 × 128	58.22 ± 17.29	80% ± 8.3%
	256 × 256	44.60 ± 20.97	74% ± 10.3%
	512 × 512	40.25 ± 19.26	71.6% ± 9.5%
T	1	58.22 ± 17.29	80% ± 8.3%
	2	48.82 ± 18.71	75.8% ± 9.1%
	4	35.36 ± 20.92	69.4% ± 10%

5. Conclusion

This work proposed and experimentally validated a photonic spiking TD3 reinforcement learning architecture for neuromorphic autonomous navigation, enabling highly energy-efficient intelligent decision-making through hardware-software co-inference. The architecture integrated a spiking-based Actor network with dual continuous-valued Critic networks and successfully completed navigation tasks in the Gazebo simulation environment, achieving an average reward of 58.22 ± 17.29 and a navigation success rate of $80\% \pm 8.3\%$. The final nonlinear spiking activation layer of the photonic Actor network was implemented on a DFB-SA laser array. The proposed system achieved an estimated inference energy of 0.78 nJ/inf and an ultra-low latency of 191.20 ps/inf , highlighting the intrinsic low-power and low-latency advantages of photonic computing for reinforcement learning. Hardware–software collaborative inference further demonstrated high accuracy, with error rates as low as 0.051% and 0.059% under obstacle-free and obstacle-interference scenarios, respectively. In addition, simulations based on the Yamada model were conducted for error-activated channels using algorithm-generated, showing full agreement with the expected spiking responses and thereby theoretically validating the dynamic characteristics of the DFB-SA laser. These results establish a solid foundation for integrating the proposed architecture with large-scale photonic linear computing chips to achieve fully-functional photonic computation. Ablation studies on hidden-layer dimensions and timesteps further confirmed the rationality of the adopted network configuration. While

the present work focuses on relatively simple navigation tasks, future efforts will target large-scale system integration, multi-task generalization, and optoelectronic co-packaging. Leveraging the inherent parallelism and ultrafast response of photonic computing, the proposed approach offers a promising pathway toward next-generation real-time neuromorphic autonomous navigation systems.

Author contributions

This work is conceptualized by Shuiying Xiang (SYX) and Yonghang Chen (YHC), and the methodology is dominated by YHC, Yahui Zhang (YHZ), Xingxing Guo (XXG) and Yanan Han (YNH). The simulation environment is set up by Tao Zou (TZ) and Shangxuan Shi (SXS). Experimental investigations are executed by YHC, Xintao Zeng (XTZ) and Mengting Yu (MTY). The whole work is supervised by SYX and Yue Hao (YH). The original draft is written by YHC and SYX. All authors read and approved the final manuscript.

Acknowledgements

This study was funded by the National Natural Science Foundation of China (No. 62535015, 62575231), the Fundamental Research Funds for the Central Universities (QTZX23041), and Xidian University Specially Funded Project for Interdisciplinary Exploration (TZJH2024009).

Competing interests

All authors declare no financial or non-financial competing interests.

Data availability

The datasets generated and/or analysed during the current study are not publicly available due to intellectual property considerations of the research project but are available from the corresponding author on reasonable request.

Code availability

The underlying code for this study is not publicly available but may be made available to qualified researchers on reasonable request from the corresponding author.

Corresponding Author

*Email: syxiang@xidian.edu.cn

*Email: zhangyahui@xidian.edu.cn

References

- [1]. Sutton, R. S. & Barto, A. G. Reinforcement learning. *J. Cogn. Neurosci.* **11**, 126-134 (1999).
- [2]. Sutton, R. S. & Barto, A. G. Reinforcement Learning: An Introduction. (MIT Press, 1998).
- [3]. Wang, X. et al. Deep reinforcement learning: a survey. *IEEE Trans. Neural Netw. Learn. Syst.* **35**, 5064-5078 (2022).
- [4]. François-Lavet, V. et al. An introduction to deep reinforcement learning. *Found. Trends Mach. Learn.* **11**, 219-354 (2018).
- [5]. Aradi, S. Survey of deep reinforcement learning for motion planning of autonomous vehicles.

-
- IEEE Trans. Intell. Transp. Syst.* **23**, 740-759 (2020).
- [6]. Isele, D. et al. Safe reinforcement learning on autonomous vehicles. In *2018 IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, 1-6 (IEEE, 2018).
- [7]. Dinneweth, J. et al. Multi-agent reinforcement learning for autonomous vehicles: a survey. *Auton. Intell. Syst.* **2**, 27 (2022).
- [8]. Silver, D. et al. Mastering the game of Go with deep neural networks and tree search. *Nature* **529**, 484-489 (2016).
- [9]. Silver, D. et al. Mastering the game of Go without human knowledge. *Nature* **550**, 354-359 (2017).
- [10]. Silver, D. et al. A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science* **362**, 1140-1144 (2018).
- [11]. Mnih, V. et al. Human-level control through deep reinforcement learning. *Nature* **518**, 529-533 (2015).
- [12]. Kormushev, P. et al. Reinforcement learning in robotics: applications and real-world challenges. *Robotics* **2**, 122-148 (2013).
- [13]. Brunke, L. et al. Safe learning in robotics: from learning-based control to safe reinforcement learning. *Annu. Rev. Control Robot. Auton. Syst.* **5**, 411-444 (2022).
- [14]. Garaffa, L. C. et al. Reinforcement learning for mobile robotics exploration: a survey. *IEEE Trans. Neural Netw. Learn. Syst.* **34**, 3796-3810 (2021).
- [15]. Mnih, V. et al. Playing Atari with deep reinforcement learning. *arXiv:1312.5602* (2013).
- [16]. Lillicrap, T. P. et al. Continuous control with deep reinforcement learning. *arXiv:1509.02971* (2015).
- [17]. Frémaux, N. et al. Reinforcement learning using a continuous time actor-critic framework with spiking neurons. *PLoS Comput. Biol.* **9**, e1003024 (2013).
- [18]. Charpentier, A. et al. Reinforcement learning in economics and finance. *Comput. Econ.* **62**, 425-462 (2023).
- [19]. Shinn, N. et al. Reflexion: language agents with verbal reinforcement learning. *Adv. Neural Inf. Process. Syst.* **36**, 8634-8652 (2023).
- [20]. Lin, E. et al. Deep reinforcement learning for imbalanced classification. *Appl. Intell.* **50**, 2488-2502 (2020).
- [21]. Fujimoto, S. et al. Addressing function approximation error in actor-critic methods. In *Proc. 35th Int. Conf. Mach. Learn.*, 1587-1596 (PMLR, 2018).
- [22]. Schulman, J. et al. Proximal policy optimization algorithms. *arXiv:1707.06347* (2017).
- [23]. Silver, D. et al. Deterministic policy gradient algorithms. In *Proc. 31st Int. Conf. Mach. Learn.*, 387-395 (PMLR, 2014).
- [24]. Xiang, S. et al. A review: photonics devices, architectures, and algorithms for optical neural computing. *J. Semicond.* **42**, 023105 (2021).
- [25]. Shen, Y. et al. Deep learning with coherent nanophotonic circuits. *Nat. Photonics* **11**, 441-446 (2017).
- [26]. Wetzstein, G. et al. Inference in artificial intelligence with deep optics and photonics. *Nature* **588**, 39-47 (2020).
- [27]. Miller, D. A. B. Attojoule optoelectronics for low-energy information processing and communications. *J. Lightwave Technol.* **35**, 346-396 (2017).
- [28]. Shastri, B. J. et al. Photonics for artificial intelligence and neuromorphic computing. *Nat.*

-
- Photonics* **15**, 102-114 (2021).
- [29].Feldmann, J. et al. Parallel convolutional processing using an integrated photonic tensor core. *Nature* **589**, 52-58 (2021).
- [30].Xu, X. et al. 11 TOPS photonic convolutional accelerator for optical neural networks. *Nature* **589**, 44-51 (2021).
- [31].Tait, A. N. et al. Neuromorphic photonic networks using silicon photonic weight banks. *Sci. Rep.* **7**, 7430 (2017).
- [32].Xu, S. et al. High-order tensor flow processing using integrated photonic circuits. *Nat. Commun.* **13**, 7970 (2022).
- [33].Bai, B. et al. Microcomb-based integrated photonic processing unit. *Nat. Commun.* **14**, 66 (2023).
- [34].Xu, S. et al. Optical coherent dot-product chip for sophisticated deep learning regression. *Light Sci. Appl.* **10**, 221 (2021).
- [35].Lin, Z. et al. 120 GOPS photonic tensor core in thin-film lithium niobate for inference and in situ training. *Nat. Commun.* **15**, 9081 (2024).
- [36].Zhou, W. et al. In-memory photonic dot-product engine with electrically programmable weight banks. *Nat. Commun.* **14**, 2887 (2023).
- [37].Wei, M. et al. Monolithic back-end-of-line integration of phase change materials into foundry-manufactured silicon photonics. *Nat. Commun.* **15**, 2786 (2024).
- [38].Tian, Y. et al. Scalable and compact photonic neural chip with low learning-capability-loss. *Nanophotonics* **11**, 329 – 344 (2022).
- [39].Wan, Y. et al. Efficient stochastic parallel gradient descent training for on-chip optical processor. *Opto-Electron. Adv.* **7**, 230182 (2025).
- [40].Xu, Z. et al. Large-scale photonic chiplet Taichi empowers 160-TOPS/W artificial general intelligence. *Science* **384**, 202-209 (2024).
- [41].Ouyang, J. et al. 16-channel photonic solver for optimization problems on a silicon chip. *Chip* **4**, 100117 (2025).
- [42].Hua, S. et al. An integrated large-scale photonic accelerator with ultralow latency. *Nature* **640**, 361-367 (2025).
- [43].Ahmed, S. R. et al. Universal photonic artificial intelligence acceleration. *Nature* **640**, 368-374 (2025).
- [44].Wu, J. et al. Analog optical computing for artificial intelligence. *Engineering* **10**, 133-145 (2022).
- [45].Huang, C. et al. Prospects and applications of photonic neural networks. *Adv. Phys. X* **7**, 1981155 (2022).
- [46].Shi, Y. et al. Photonic integrated spiking neuron chip based on a self-pulsating DFB laser with a saturable absorber. *Photonics Res.* **11**, 1382-1389 (2023).
- [47].Feldmann, J. et al. All-optical spiking neurosynaptic networks with self-learning capabilities. *Nature* **569**, 208-214 (2019).
- [48].Robertson, J. et al. Image edge detection with a photonic spiking VCSEL-neuron. *Opt. Express* **28**, 37526-37537 (2020).
- [49].Xiang, S. et al. Computing primitive of fully VCSEL-based all-optical spiking neural network for supervised learning and pattern classification. *IEEE Trans. Neural Netw. Learn. Syst.* **32**, 2494-2505 (2020).

-
- [50]. Lee, Y. J. et al. Photonic spiking neural networks with event-driven femtojoule optoelectronic neurons based on Izhikevich-inspired model. *Opt. Express* **30**, 19360-19389 (2022).
- [51]. Xiang, S. et al. Hardware-algorithm collaborative computing with photonic spiking neuron chip based on an integrated Fabry-Perot laser with a saturable absorber. *Optica* **10**, 162-171 (2023).
- [52]. Xiang, S. et al. Photonic integrated neuro-synaptic core for convolutional spiking neural network. *Opto-Electron. Adv.* **6**, 230140 (2023).
- [53]. Talukder, R. et al. A spiking photonic neural network of 40,000 neurons, trained with latency and rank-order coding for leveraging sparsity. *Neuromorph. Comput. Eng.* (2025).
- [54]. Gao, S. et al. Hardware implementation of ultra-fast obstacle avoidance based on a single photonic spiking neuron. *Laser Photonics Rev.* **17**, 2300424 (2023).
- [55]. Zhang, Q. et al. All-optical digital logic and neuromorphic computing based on multi-wavelength auxiliary and competition in a single microring resonator. *Opto-Electron. Sci.* (2025).
- [56]. Huang, B. et al. Inspection robot navigation based on improved TD3 algorithm. *Sensors* **24**, 2525 (2024).
- [57]. Wu, J. et al. A-TD3: an adaptive asynchronous twin delayed deep deterministic for continuous action spaces. *IEEE Access* **10**, 128077-128089 (2022).
- [58]. Xu, T. et al. End-to-end autonomous driving decision method based on improved TD3 algorithm in complex scenarios. *Sensors* **24**, 4962 (2024).
- [59]. Liu, K. et al. Enhancing mixed traffic stability with TD3-driven bilateral control in autonomous vehicle chains. *Sustainability* **17**, 4790 (2025).
- [60]. Gao, X. et al. The optimal strategies of maneuver decision in air combat of UCAV based on the improved TD3 algorithm. *Drones* **8**, 501 (2024).
- [61]. Nguyen, T. H. et al. TD3-based optimization framework for RSMA-enhanced UAV-aided downlink communications in remote areas. *Remote Sens.* **15**, 5284 (2023).
- [62]. Koenig, N. & Howard, A. Design and use paradigms for Gazebo, an open-source multi-robot simulator. In *2004 IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, 2149-2154 (IEEE, 2004).
- [63]. Takaya, K. et al. Simulation environment for mobile robots testing using ROS and Gazebo. In *2016 20th Int. Conf. Syst. Theory Control Comput. (ICSTCC)*, 96-101 (IEEE, 2016).
- [64]. Zeng, X. et al. Photonic spiking neural network based on DML and DFB-SA laser chip for pattern classification. *Opt. Express* **33**, 12045-12058 (2025).
- [65]. Yamada, M. A theoretical analysis of self-sustained pulsation phenomena in narrow-stripe semiconductor lasers. *IEEE J. Quantum Electron.* **29**, 1330-1336 (2002).
- [66]. Xiang, S. et al. Nonlinear photonic neuromorphic chips for spiking reinforcement learning. *arXiv:2508.06962* (2025).
- [67]. Xiang, S. et al. Review of ASIC accelerators for deep neural network. *Microprocess. Microsyst.* **89**, 104441 (2022).
- [68]. Tang, G., Kumar, N. & Yoo, R. et al. Deep reinforcement learning with population-coded spiking neural network for continuous control. *Conf. Robot. Learn.*, 2016 (2021).

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [SupplementaryVideo1.mp4](#)
- [SupplementaryVideo2.mp4](#)
- [SupportingInformation.docx](#)