

Decision-making and control with diffractive optical networks

Jumin Qiu,^a Shuyuan Xiao^{b,c}, Lujun Huang,^{d,*} Andrey Miroshnichenko,^e Dejian Zhang,^a Tingting Liu^{b,c,*} and Tianbao Yu^{a,*}

^aNanchang University, School of Physics and Materials Science, Nanchang, China

^bNanchang University, School of Information Engineering, Nanchang, China

^cNanchang University, Institute for Advanced Study, Nanchang, China

^dEast China Normal University, School of Physics and Electronic Science, Shanghai, China

^eUniversity of New South Wales Canberra, School of Physics and Electronic Science, Canberra, Australia

Abstract. The ultimate goal of artificial intelligence (AI) is to mimic the human brain to perform decision-making and control directly from high-dimensional sensory input. Diffractive optical networks (DONs) provide a promising solution for implementing AI with high speed and low power-consumption. Most reported DONs focus on tasks that do not involve environmental interaction, such as object recognition and image classification. By contrast, the networks capable of decision-making and control have not been developed. Here, we propose using deep reinforcement learning to implement DONs that imitate human-level decision-making and control capability. Such networks, which take advantage of a residual architecture, allow finding optimal control policies through interaction with the environment and can be readily implemented with existing optical devices. The superior performance is verified using three types of classic games: tic-tac-toe, Super Mario Bros., and Car Racing. Finally, we present an experimental demonstration of playing tic-tac-toe using the network based on a spatial light modulator. Our work represents a solid step forward in advancing DONs, which promises a fundamental shift from simple recognition or classification tasks to the high-level sensory capability of AI. It may find exciting applications in autonomous driving, intelligent robots, and intelligent manufacturing.

Keywords: diffractive optical networks; optical computing; deep learning; reinforcement learning.

Received Feb. 7, 2024; revised manuscript received Mar. 29, 2024; accepted for publication May 13, 2024; published online May 30, 2024.

© The Authors. Published by SPIE and CLP under a Creative Commons Attribution 4.0 International License. Distribution or reproduction of this work in whole or in part requires full attribution of the original publication, including its DOI.

[DOI: [10.1117/1.APN.3.4.046003](https://doi.org/10.1117/1.APN.3.4.046003)]

1 Introduction

Artificial intelligence (AI) is to imitate the functions of neurons in performing decision-making by creating hierarchical artificial neural networks. It has found many exciting applications in computer vision,^{1,2} natural language processing,^{3,4} and data mining.⁵ Except for electronics and computer science applications, artificial neural networks have been applied to optimize the design of photonic devices, including metamaterials and metasurfaces, significantly facilitating the performance of photonic devices beyond the conventional inverse design strategy.⁶⁻¹³

Recently, optical neural networks have drawn tremendous attention because they provide a compelling route for processing information at the speed of light,¹⁴⁻¹⁹ with low energy consumption and massive parallelism compared with the electronic-circuit-based neural networks. In the pioneering work of Lin et al.,²⁰ diffractive optical networks (DONs, also known as diffractive deep neural networks, D²NN) consisting of multilayer of three-dimensional printed diffractive optical elements operating at terahertz were first proposed for inference and prediction through parallel computation and dense interconnection at the speed of light. Later, DONs were extended to various nanostructures for implementation. Such an architecture has been effectively validated in performing specific inference functions, such as image classification,²¹⁻²⁴ saliency detection,²⁵ and logic

*Address all correspondence to Lujun Huang, ljhuang@phy.ecnu.edu.cn; Tingting Liu, tliu@ncu.edu.cn; Tianbao Yu, yutianbao@ncu.edu.cn

operation.²⁶ More recently, a reconfigurable DON based on optoelectronic fused computing architecture has been proposed,²⁷ which can perform different neural networks and achieve a high model complexity with millions of neurons. Although DONs have witnessed significant progress in the past few years, their functions mainly focus on image classification and object recognition without involving any interaction with the environment. To our knowledge, human-level AI based on DONs that can perform decision-making and control has not yet been developed.

In this work, we bring the capability of decision-making and control directly from high-dimensional sensory inputs to the DON. The networks build upon deep reinforcement learning to interact with a simulated environment for optimal control policies. The training process of policy is based solely on deep reinforcement learning from self-play without a data set or guidance. A phase profile mapping features each layer of the DON and thus can be immediately implemented by optical modulation devices. The effectiveness of the proposed DON is validated with three typical games: tic-tac-toe, Super Mario Bros., and Car Racing. We also provide a direct experimental demonstration of such a DON capable of playing tic-tac-toe. Excellent agreement can be found between theoretical prediction and experimental measurements. This work enables a fundamental shift from the target-driven control of a predesigned state for simple recognition or classification tasks to human-imitative AI, revealing the potential of optoelectronic AI systems to solve complex real-world problems. We envision that such DONs will find promising applications in autonomous driving, industrial robots, and intelligent manufacturing, enhancing human life in every aspect.

2 Methods

The working principle of the DON for decision-making and control is illustrated in Figs. 1(a)–1(c), using an example of playing Nintendo's classic video game Super Mario Bros. In general, a human player goes through seeing, understanding, and making a decision in each step, and these perception and control behaviors loop until the game is over. To play games in a human-like manner, the network necessitates the sensory capability to capture continuous, high-dimensional state spaces and the controllable execution ability of sequences of different behaviors. The DON shown in Fig. 1(b) comprises the specific free-space configuration: an input layer with images encoded using an optical modulation device, multiple hidden layers encoding phases of transmitted waves, and an output layer into which the computational results are imaged.

More importantly, the proposed framework for decision-making and control integrates the deep reinforcement learning and the DON into a training procedure, allowing interaction between the game and the agent to learn control policies that can be implemented through the optical computing platform. The method observes each state within the game environment and chooses a particular action through a learned control policy for each situation. Then, the changed environment generates observation of the new state, makes the following action, and continuously updates the control policy in the loop. Unlike the previous optical networks, the input images from each video game frame are continuous high-dimensional sensory data. Furthermore, the execution procedure, such as playing games, is essentially a type of interactive control rather than the one-way recognition for a single objective, such as written digits or fashion items.

To address the complexity of imitating human players on the optical platform, we develop the training framework of policy and network shown in Fig. 1(d), using a combination of novel and existing general-purpose techniques for neural network architectures. As shown in the middle block of Fig. 1(d), central to the architecture is a control policy $\pi_{\theta}(a|s)$, which is represented by a convolutional neural network (CNN) with parameters θ that makes states s as inputs and takes actions a as outputs by optimizing the reward of games of self-play. Note that the training epoch of deep reinforcement learning is markedly more than that of the DONs due to the training of policies starting from entirely random behavior. Thus, we developed the training process approach with two main phases to eliminate unnecessary computations. First, deep reinforcement learning through an agent interacts with a simulated game environment to find a near-optimal control policy to meet the specified goals. Second, the control policy updates the DON by the error backpropagation algorithm.

In the first phase, a deep reinforcement learning algorithm collects data to find a control policy concerning the specific reward function through interaction with the game environment, thereby achieving the desired outcome. The states of these games need to satisfy the Markov property that the information of a particular state contains all relevant histories. Thus, it is possible to perform actions in the current state and move to the next state without considering the previous states. The agent interacts with the environment through a sequence of observations, actions, and rewards. At each step of interaction, the agent observes the state of the environment to decide on an action to take and then gives rewards based on the game result. The neural network decides the best action for each step based on the reward. It continuously updates the policy using proximal policy optimization²⁸ to find the optimal action. After testing, the trained policies can all complete the respective game. Compared with previous studies, the algorithm only requires game rules without the need for human data, guidance, or domain knowledge, avoiding the performance's dependence on the data set's quality.

In the second phase, the control policy is transferred onto the DON. The optimal control policy modeled using a CNN is utilized as the ground truth during the learning procedure. Meanwhile, following the forward propagation model based on Huygens' principle and Rayleigh–Sommerfeld diffraction, the encoded input light can be directed into any desired location at the output layer via the learnable transmission coefficients, that is, phase profiles of hidden layers in the network. The energy distributions clustered in the target detection region imply the prediction results. The transmission coefficients at each diffractive layer should be adequately trained via the error backpropagation algorithm and a loss function with mean square error, which is defined to evaluate the performance between the output intensities and the ground-truth target. The adaptive moment estimation,²⁹ an algorithm for first-order gradient-based optimization of stochastic objective functions, is adopted to reduce the loss function. Then, the gradient of the loss function concerning all the trainable network variables is backpropagated to iteratively update the network during each cycle of the training phase until the network converges.

Once the training is completed, the target phase profiles of the diffractive layers are determined, which are ready to connect the physical and digital worlds for optical neuromorphic computing. Here, we choose an approach similar to the diffractive

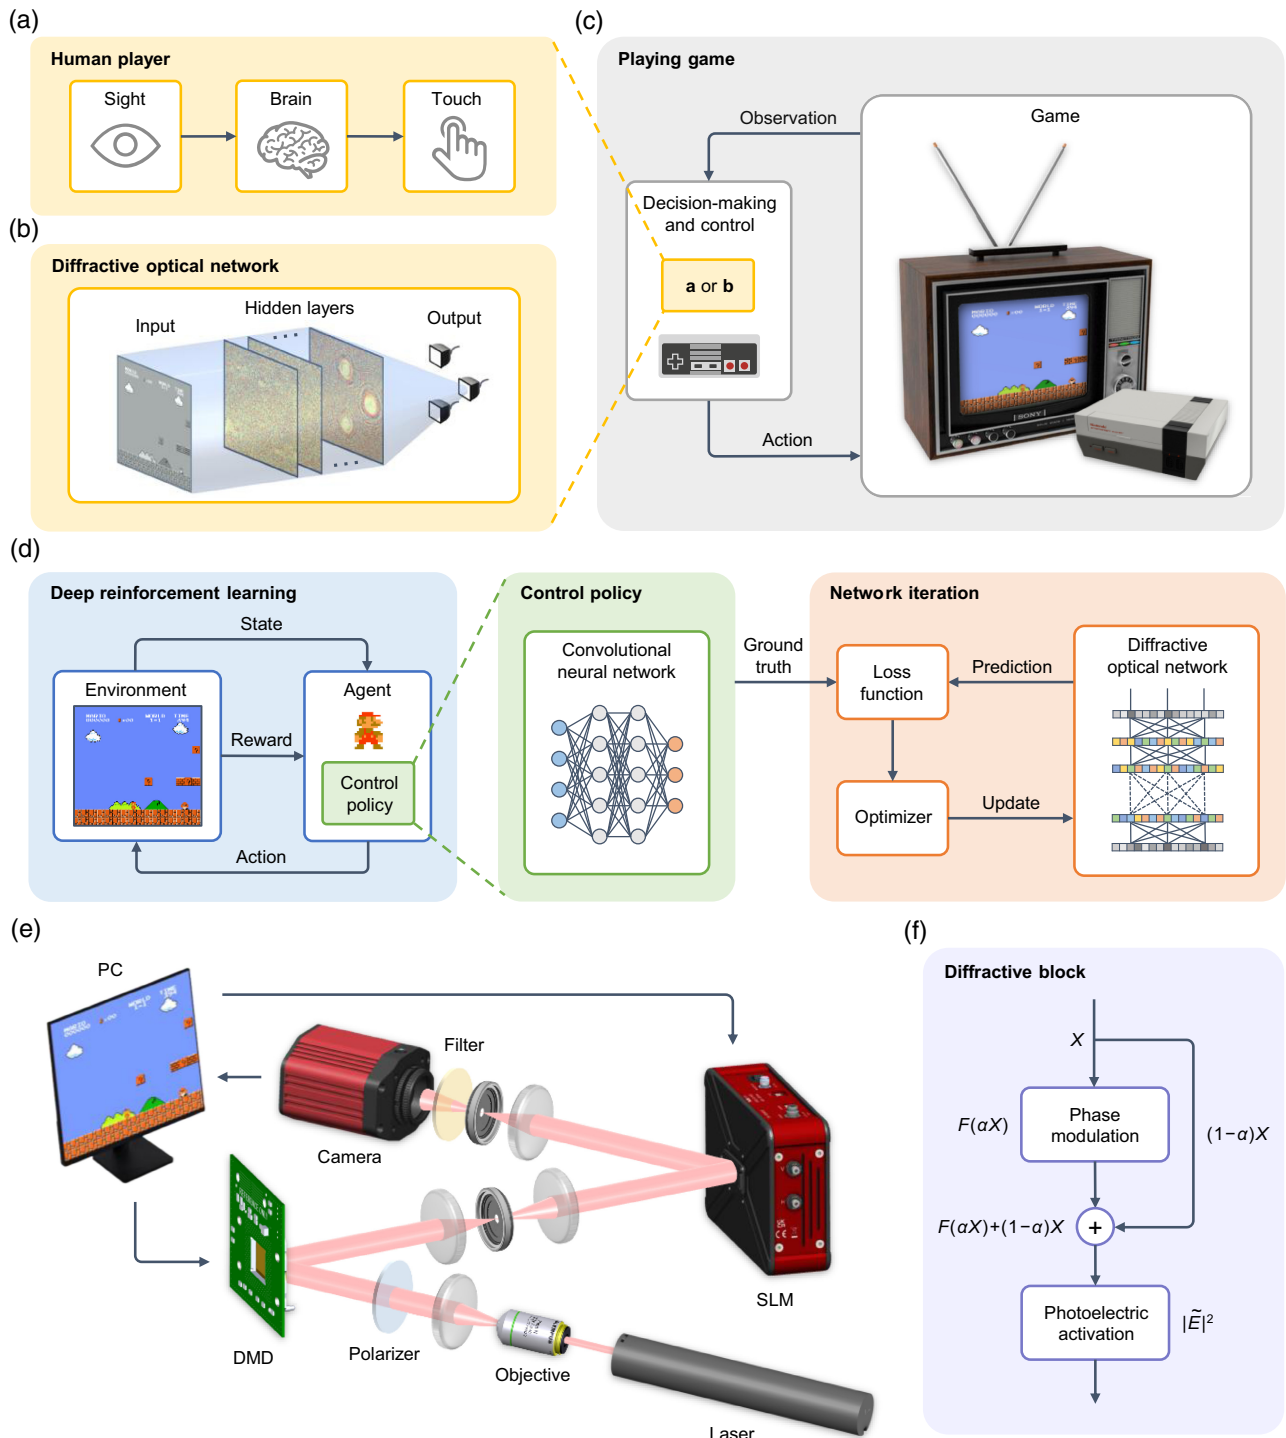


Fig. 1 DON for decision-making and control. (a)–(c) The proposed network plays the video game of Super Mario Bros. in a human-like manner. In the network architecture, an input layer captures continuous and high-dimensional game snapshots (seeing), a series of diffractive layers choose a particular action through a learned control policy for each situation faced (making a decision), and an output layer maps the intensity distribution into preset action regions to generate the control signals in the games (controlling). (d) Training framework of policy and network. Deep reinforcement learning through an agent interacts with a simulated environment to find a near-optimal control policy represented by a CNN, which is employed as the ground truth to update the DON by an error backpropagate algorithm. (e) The experimental setup of the DON for decision-making and control. (f) The building block of the DON.

processing unit²⁷ to build the network because of its reconfigurability and ability to support millions of neurons for computation. The experimental setup of the DON is shown in Fig. 1(e). A laser beam with a working wavelength of 632.8 nm is expanded using a microscope objective and lens, while a linear polarizer can be embedded to adjust the incident light intensity, which is then projected onto a digital micromirror device (DMD). The input image data are optically encoded and modulated by the DMD, followed by two relay lenses to adjust the image to the appropriate size and projected onto the spatial light modulator (SLM) for phase modulation. The optical iris is used to filter out high-order diffractions and stray light. The diffraction pattern is imaged onto the camera; then, the output image is input to the DMD for the next diffractive layer until the end of the network computation. After that, the optical intensities in the predefined detection zones are extracted from the output image, and the predicted results are decoded to generate the control signals in the games. Then, the new frame image of the video game stimulates the new process procedure, and the updated results control the game until the end. In addition, because the DMD is a binary device, the training process needs to simulate the fast rotation of the micromirror when displaying gray-scale images to make the training results more practical. We adapt the previously trained phase profiles for the DMD, as detailed in the [Supplementary Material](#). The entire computing process is primarily optical, except for the dataflow control. These light modulation devices are very fast and therefore allow for real-time computation.

Such an experimental system allows for a deep residual framework that can overcome the vanishing gradients problem by introducing shortcut connections between layers, and the architecture has become one of the cornerstones of neural networks.³⁰ Figure 1(f) demonstrates a block that composes the DON. First, when there is an angle between the polarization direction of incident light and the extraordinary axis of the liquid crystal of SLM, some light will not be modulated and reflected directly to the camera, thus creating a shortcut connection. Formally, the incident light is denoted as X , the diffraction computation is denoted as $F(X)$, and the original mapping can be recast into $F(\alpha X) + (1 - \alpha)X$, where α is the modulation ratio of the SLM, which can be fine-tuned by rotating laser and polarizer to change the polarization direction (or adding a half-wave plate). Compared with previous research,³¹ this approach does not require introducing additional optical devices, providing a free improvement. In addition, the approach lowers the bar for the polarization state of light, and partially polarized light can be used in the network. Then, we use the photoelectric effect occurring at each image sensor pixel to implement the activation function of diffractive neurons, denoted as $|\tilde{E}|^2$. In addition, to some extent, the exposure of the camera and the differences in resolution among various devices can be analogized to the layer normalization and downsampling operations of neural networks, respectively. Unlike previous studies that used complex network structures, we stack the block to build the DON.

3 Results

3.1 Playing Tic-Tac-Toe

In our first implementation, we perform the decision-making and control for tic-tac-toe. This classic game is played on a 3×3 grid of cells where each player places their mark, an X or an O, in an empty cell. The first player to place three of their

marks in a row vertically, horizontally, or diagonally wins the game. If all cells are filled, and neither player has three marks in a row, the game is declared a draw. There are 255,168 possible ways to play this game, and we use the proposed network architecture to capture the effective policies to make the most optimal move in every possible situation.

To play this game, the network composed of three diffractive blocks is designed by the above training algorithm. The input images carrying the information of the current states are encoded into the amplitude of the input field to the network. The network is trained to map the incident energy into nine cells corresponding to the grid (labeled by the numbers 1 to 9), where the received energy distribution at each region reveals the current state and predicts the probability of the player's next move, as shown in Fig. 2(a). Since the observed state and the action are both discrete in this game, tic-tac-toe can be considered to demonstrate our method for a collection of tasks with discrete state and action spaces.

Note that the first player (X) and the second player (O) have different control policies; specifically, X tries to win, and O tries to draw in the ideal case. After training, the X moves of each turn are illustrated in Fig. 2(b). In the first turn, two possible positions, 1 and 5, are predicted as shown on the output. However, the starting position of 5 is finally chosen because of the maximum energy intensity among these positions. After O responds to X, the input image changes, and in the second turn, the intensity distributions change as well so that the predicted move of X at the position of 1 is determined by extracting the maximum signal among the unoccupied positions. It is also noted that the output intensity is focused not only on the predicted position but also on the current states with occupied positions. Following this prediction and control procedure, the first player wins in the fourth turn. Following the same principle, the O moves are predicted and controlled, as shown in Fig. 2(c). It can be observed that O responds to a corner opening with a central mark and chooses moves next to X to avoid the opponent having three marks in a row. In such a way, O prevents X from winning. This policy is successfully used in the proposed network, and a draw game result is shown in Fig. 2(c), while O can win if X plays weakly in some exceptional cases.

While, in general, a human player aims to win the game, tic-tac-toe will end in a draw if both players play their best because it is a zero-sum game. To evaluate the accuracy and effectiveness of our proposed network in playing tic-tac-toe, we use the sum of the win and draw rate as the accuracy rate. After the self-play training per game rules, we numerically test the design of the DON with all possible states, as shown in Fig. 2(d). The policy we trained is optimal and will only choose the best moves, so only 670 states will appear. Among them, the accuracy rate of X is 100%, the accuracy rate of O is 90.56%, and the average rate is 91.79%. The accuracy of the network in predicting O shows a slight degradation relative to that of X due to factors such as more complex policy and more states of O.

In addition, we evaluate the dependence of the prediction accuracy on the number of hidden layers in Fig. 2(e). It can be seen that the accuracy of the network is greatly improved when changing from 2 to 3 layers because if there are not enough layers in the network, the shortcut connections between layers may not be fully computed, thus affecting the results. However, the accuracy does not show a noticeable change when the layer

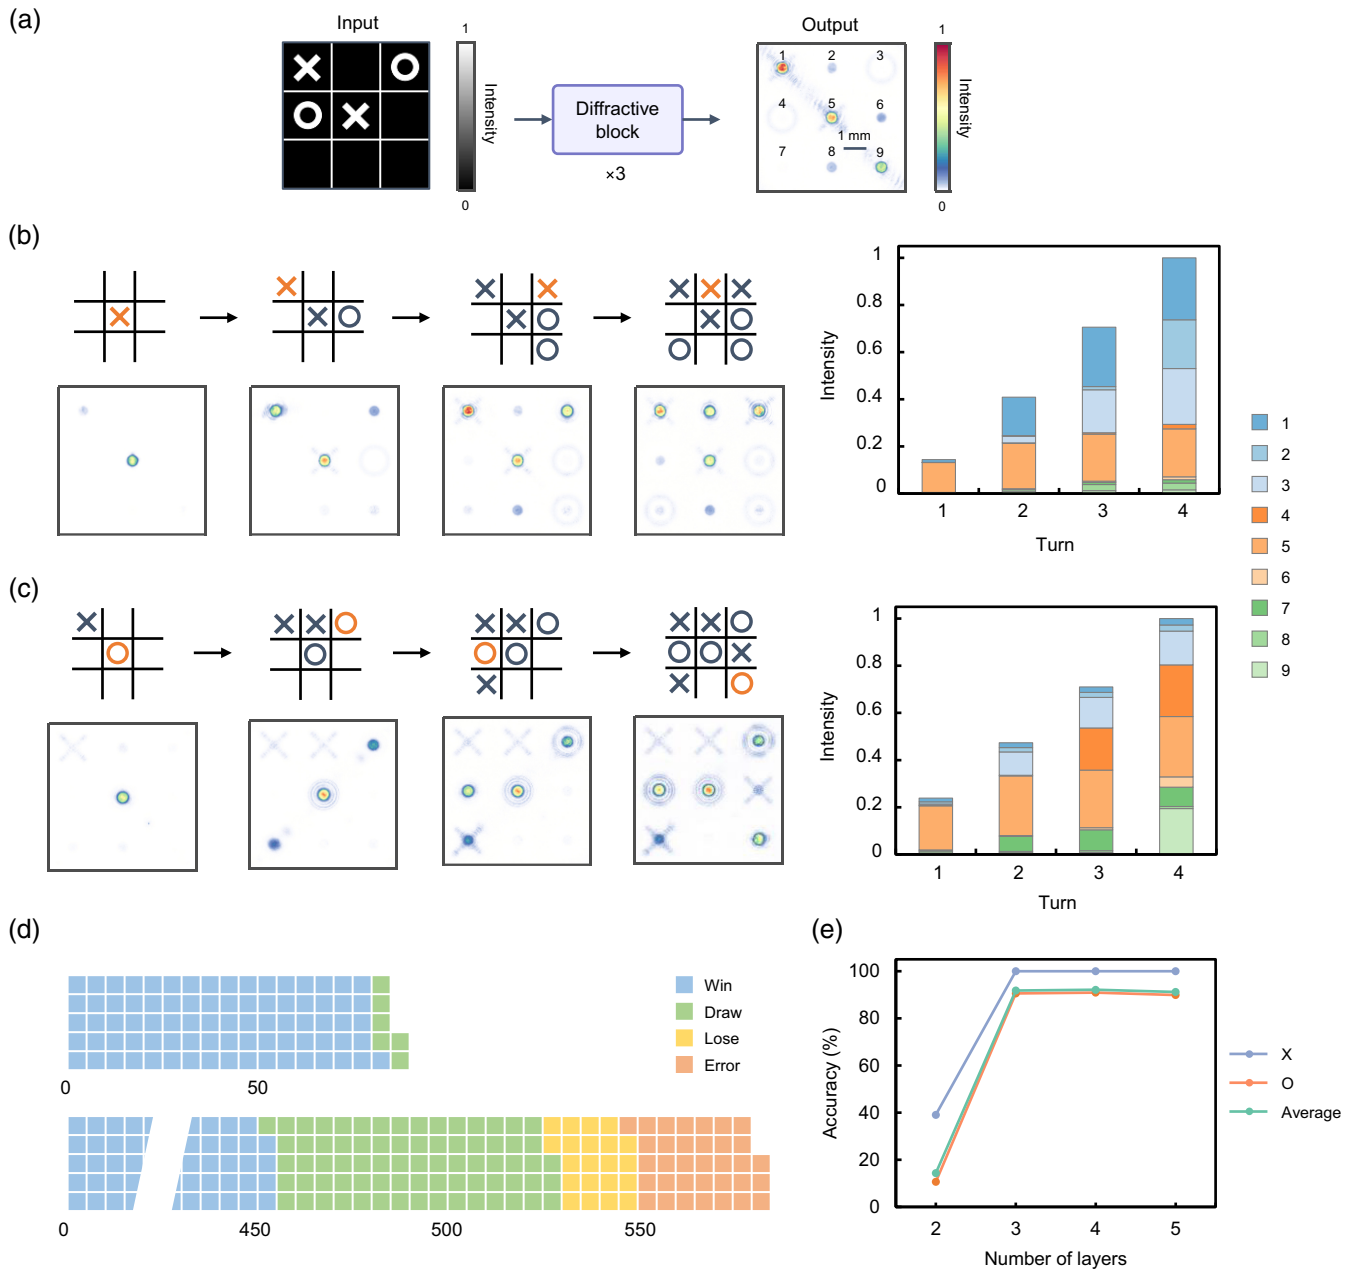


Fig. 2 Playing tic-tac-toe. (a) The schematic illustration of the DON composed of an input layer, hidden layers of three cascaded diffractive blocks, and an output layer for playing tic-tac-toe. (b) and (c) The sequential control of the DON in performing gameplay tasks for X and O. (d) The accuracy rate of playing tic-tac-toe. There is a collection of 87 games utilized for predicting X, obtaining 81 wins and 6 draws in these games. In the rest of the 583 games, O obtains 454 wins, 74 draws, and 21 losses. When previous moves have occupied the predicted position at a turn, such a case is counted as a playing error and occurs 34 times. (e) Dependence of the prediction accuracy on the number of hidden layers.

number continues to increase from 3, which may be due to the following reasons. First, the DON is unsuitable for predicting states with high similarity;³² see the [Supplementary Material](#) for detailed derivation. In addition, DONs have a global perceptual property similar to a multilayer perceptron (MLP), which can capture features at given spatial locations. However, it is difficult to capture features among different spatial locations.³³ We will discuss this point later in the paper.

3.2 Playing Super Mario Bros

In our second implementation, the world 1-1 of the original Super Mario Bros. game is used to demonstrate the validity of the DON. Unlike the tic-tac-toe on a square-divided board, Super Mario Bros. is a video game with continuous high-dimensional state inputs. The gameplay consists of moving the player-controlled character, Mario, through two-dimensional levels to get to the level's

end, traversing it from left to right, avoiding obstacles and enemies, and interacting with game objects. In the game, the player controls Mario to take discrete actions: run, jump, and crouch. Under these considerations, this game can be an example of continuous state space and discrete action space for testing the proposed network.

Figure 3(a) illustrates the DON for playing Super Mario Bros. The network consists of an input layer carrying the optical field encoded from each video game frame, hidden layers composed of three cascaded diffractive blocks trained by the same algorithm, and the output layer mapping the intensity distribution into preset regions. It is clear that the input images from the

game scene consisting of moving backgrounds and different objects are more complex compared with the tic-tac-toe with a regular pattern. In addition, the game images are similar to adjacent ones and constantly changing due to the gameplay on a side-scrolling platform, which challenges the DON in processing highly similar input states for choosing optimal actions.

After training with the control policy, the network makes decisions for Mario's optimal action. It achieves accurate control to reach the end of the level until taking down the flag raised above the castle, as shown in Video 1. Specifically, at any given state, the most optimal action that Mario chooses to take is predicted by the maximum action signal. In the examples of

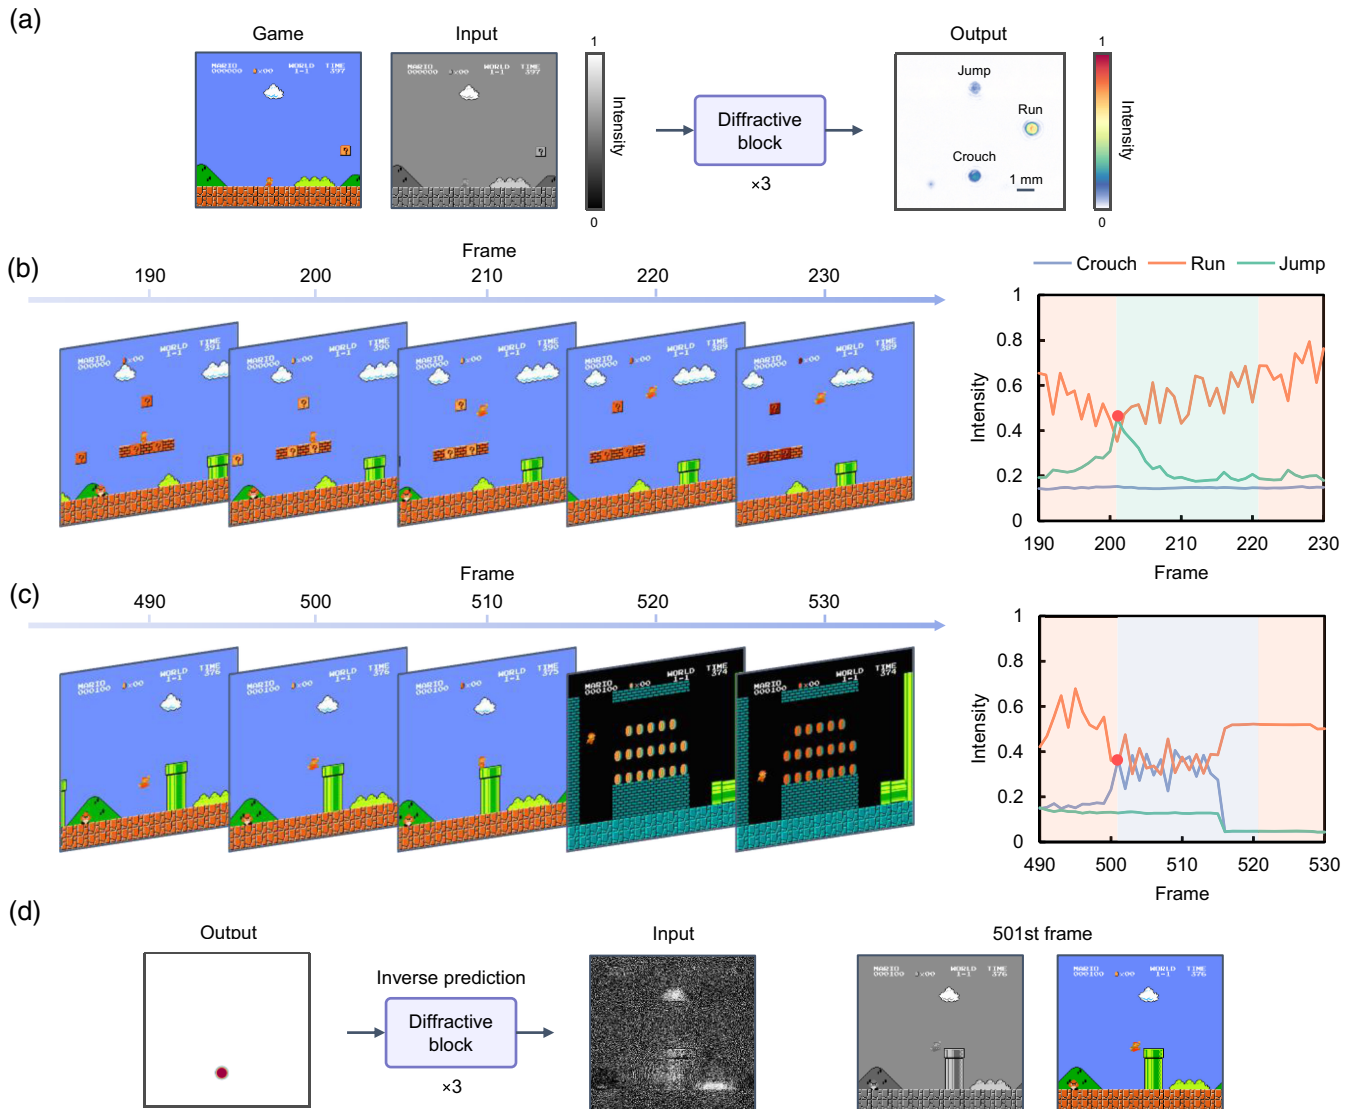


Fig. 3 Playing Super Mario Bros. (a) The layout of the designed network for playing Super Mario Bros. (b) and (c) Snapshots of Mario's jumping and crouching actions by comparing the output intensities of actions. The output intensity of the jump is maximum at the 201st frame, so the predicted action is jump, and Mario is controlled to act, as shown in panel (b). A similar series of prediction and control for another crouch action can also be observed in panel (c). (d) The inverse prediction result. Considering the predicted crouch at the current state is crucial for updating Mario's action, we use the maximized output intensity of the crouch as input, ignoring the simultaneous output of other actions (Video 1, MP4, 19.8 MB [URL: <https://doi.org/10.1117/1.APN.3.4.046003.s1>]).

Figs. 3(b) and 3(c), we take some snapshots from Video 1 to analyze the decision-making and control of Mario's actions in complex and time-varying configurations. Since the goal of our network is to finish the level as quickly as possible successfully, Mario should maintain the run action until the end while

choosing to jump or crouch to overcome the challenges at certain states. Thus, the output intensity of run remains high throughout the game, while the intensity of jump and crouch shows smaller fluctuations, verified by Figs. 3(b) and 3(c). Although this significant intensity triggers the prediction only

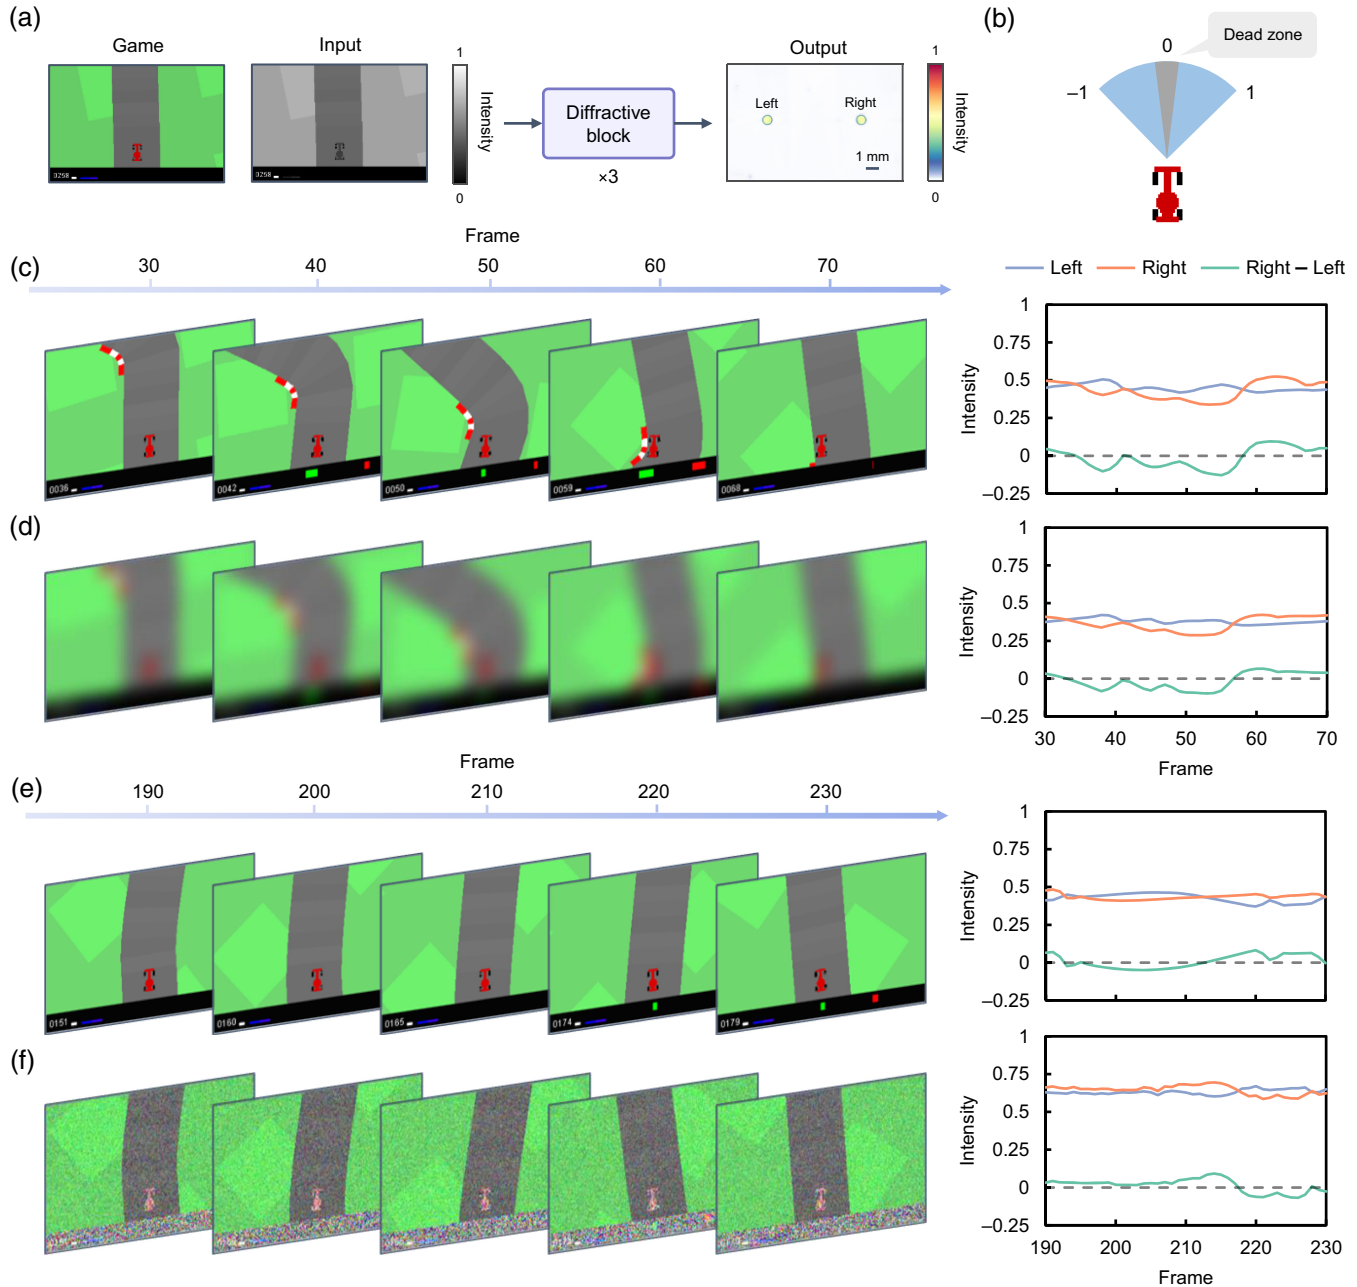


Fig. 4 Playing Car Racing. (a) The layout of the designed network for playing Car Racing. (b) The control of the steering direction and angle of the car with respect to the difference value between the intensities at the current state, normalized between -1 and 1 . (c)–(f) Snapshots of controlling the car steering. When the car is facing a left-turn track in panel (c), the output intensity on the left keeps the value greater than the right intensity, allowing continuous control in updating the rotation angle of the left-turn action. A similar control process can also be performed for the right-turn track in panel (e). In addition, the anti-disturbance of the network is validated by introducing (d) the Gaussian blur and (f) Gaussian noise to the game images (Video 2, MP4, 8.36 MB [URL: <https://doi.org/10.1117/1.APN.3.4.046003.s2>]; Video 3, MP4, 6.78 MB [URL: <https://doi.org/10.1117/1.APN.3.4.046003.s3>]; Video 4, MP4, 16.8 MB [URL: <https://doi.org/10.1117/1.APN.3.4.046003.s4>]).

at a particular frame, this control signal is intentionally set to last for 20 frames to ensure that Mario finishes the entire action. It is worth noting that the intensity-frame curve remains relatively stable during the 516th to 530th frame, which can be understood with the static and high-contrast background images after Mario enters the pipe, as shown in Fig. 3(c).

To gain an insight into how the DON makes decisions, we investigate the network's perception capability, employing inverse prediction in Fig. 3(d). We demonstrate what the network has learned from the high-dimensional sensory input to perform the crouch action corresponding to the 501st frame image. We use the error backpropagation algorithm in a retrained network to inversely predict the input image at this moment, where $\alpha = 1$ in the network to avoid the effect of the residual structure;

see the [Supplementary Material](#) for detailed derivation. The inversely predicted image matches the original input image of the 501st frame, especially the background, such as clouds and grasses. When humans play the game, they may ignore these backgrounds and focus only on the critical parts, such as Mario, enemies, and pipes. The inverse prediction of the whole scene highlights the capability of the network to extract global features instead of local ones; the property is the same as MLP, further verifying the perception capability of the network in capturing the global features to make decisions.

3.3 Playing Car Racing

In our third implementation, we demonstrate the proposed network capability in Car Racing, which requires perceiving the

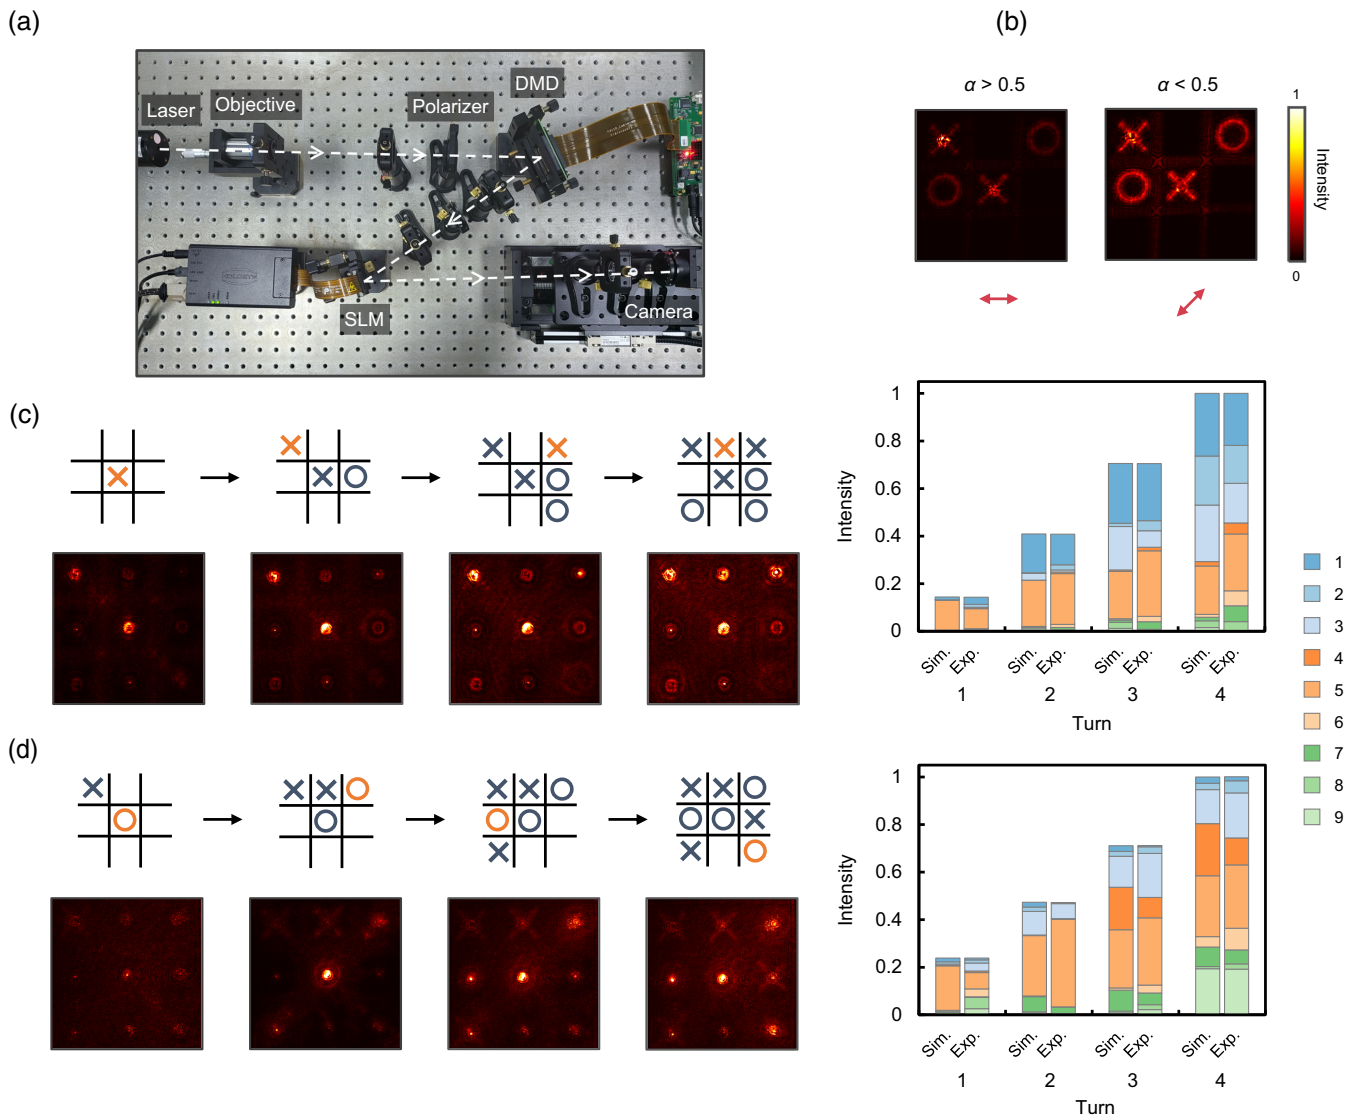


Fig. 5 Experimental demonstration of the DON for tic-tac-toe. (a) The photo of the experimental system, where the unlabeled devices are lenses, a spatial filter is used to remove the unwanted multiple-order energy peaks, and a filter is mounted on the camera. (b) The output of the first layer of the sample in Fig. 2(a), and the red arrows represent the polarization direction of the incident light. (c) and (d) The sequential control of the DON in playing the same two games as in Figs. 2(b) and 2(c), respectively. The experimental results are normalized based on simulation results. Sim., simulation result; Exp., experimental result.

game environment using continuous high-dimensional inputs and making decisions to control the car by performing continuous steering actions. The game's control policy is trained based on the rules of keeping the car within the track by controlling its rotation, and the car is set to increase the speed once the game starts continuously. The DON architecture is shown in Fig. 4(a) and is similar to previous examples. The input energy of the optical field is redistributed through three diffractive blocks into the two designated regions on the left and right of the output layer. The difference value between the intensities at the current state controls the steering direction and angle of the car shown in Fig. 4(b). In addition, just as in the steering dead zone in real vehicles, a slight difference value would not lead to steering action to avoid disturbance.

The successful network implementation in Car Racing is illustrated in Video 2, where the car is controlled in the center of the track almost within the whole lap. For the two basic actions of the left and right turn, some exemplary snapshots are provided in Figs. 4(c) and 4(e). Specifically, the negative difference values in Fig. 4(c) predict the left turn of the car wheel, while the larger absolute values indicate sharper turns. It is also observed that sometimes the difference values approach zero, and a rotation angle of 0 is predicted to keep the car moving in the direction of the current state. Due to the larger turning angle of the track, the intensity difference of the left turn shows a more drastic change. It is also intriguing that although the steering of the car in the left turn is somewhat unsmooth so that it does not appear in the middle of the track in certain states, the controlled action that is updated in the following state leads to successful gameplay. This real-time feedback and updating feature shows the great potential of the architecture for challenging auto-driving almost at the speed of light,³⁴ such as dealing with sudden obstacles.

To validate the anti-disturbance ability of the proposed approach, we introduce two crucial randomization disturbance mechanisms to the frame image of the game and then test the network performance in controlling Car Racing. With the same network previously trained, Gaussian blur and Gaussian noise are added to the frames; the control results are shown in Figs. 4(d) and 4(f), respectively. Although the introduction of disturbances, including blur and noise, causes the quality decline of the input image, the car can still maintain accurate and effective control to successfully complete the game, as verified by Videos 3 and 4. Compared with the normal cases in Figs. 4(c) and 4(e), the output intensity curves in Figs. 4(d) and 4(f) show similar trends to control the left or right turning actions. However, the curves are less smooth, with more amplitude fluctuations, indicating unsmooth steering angle control. The successful control in the cases with the randomization disturbance reveals the great perception of the game environment, especially the full access to the global features.

3.4 Experimental Demonstration of Playing Tic-Tac-Toe

Finally, to evaluate the actual experimental performance of the DON, we built an experimental system using off-the-shelf optical modulation devices. It realizes residual architecture with only one path of light, reducing additional devices and easier alignment. We tested it by playing tic-tac-toe, the experimental system, as shown in Fig. 5(a).

We first tested our proposed residual architecture. Fig. 5(b) shows the effect of our proposed residual architecture, which is

the output of the first layer of the sample in Fig. 2(a). It can be seen that the value of α varies with the polarization direction of the incident light. This shows that our proposed residual architecture is valid and can easily adjust the ratio of modulation and residual channels to flexibly adapt to various tasks.

After that, we tested the same two games as in Figs. 2(b) and 2(c); the experimental results are presented in Figs. 5(c) and 5(d). It can be seen that the intensity distribution of the output changes as the input game state changes. Due to the unavoidable physical error in the experimental system, the experimental results are different from the simulated ones, but the overall intensity changes are very similar. The maximum intensity distributions occur at the same positions, and the same games are successfully completed.

4 Conclusion

We have demonstrated DONs for decision-making and control. The optimal control policy enables this technique through a harmonious combination of deep reinforcement learning and the DON architecture. Based solely on reinforcement learning from self-play, the control policy of the training algorithm is flexible, as demonstrated by successfully learning to play the three types of classic games. In addition, we further exploit the potential of the photoelectric fusion DON by introducing a free residual architecture that achieves excellent performance in the simplest network structure.

It is worth noting that tic-tac-toe does not achieve perfect results despite the definite rules and optimal control policy, just like Super Mario Bros. and Car Racing. There are several possible reasons for this result: Playing tic-tac-toe needs to strategically handle different states and a more significant number of output signals. The gameplay of tic-tac-toe requires correct predictions at each state, while the other two games show that better error tolerance and accidental mistakes do not necessarily affect the results. In addition, using the difference as a mechanism to trigger actions improves the network's performance in Car Racing to some extent. Since the DON is not good at extracting local features, the differences in intensity distributions between the adjacent input board images are challenging to detect for tic-tac-toe.

By testing our proposed DON on the challenging domain of classic games, we demonstrate its ability to master difficult game control policies for playing games for the first time on an optical platform. This work bridges the gap between optical and digital neural networks aiming to achieve human-level AI. The most important aspect is that the decision-making and control process is implemented in optical devices at the speed of light by imitating human competence. Another ideal platform for implementing DONs is metasurface. Metasurfaces provide an unprecedented ability to manipulate the wavefront of light and are widely used to implement sophisticated functions such as holography and computational imaging.^{35–38} Therefore, driven by the demand for all-optical on-chip integration of AI systems, some recent studies have introduced optical metasurfaces consisting of an array of subwavelength meta-atoms to replace bulky diffractive optical devices for high-density integration.^{22,23,39–41} The working mechanism and design principle of our proposed DONs are universal, and thus can be generalized to nanostructures. We have also implemented the above network on metasurfaces; see the [Supplementary Material](#) for details. Therefore, a metasurface-based DON can be envisaged

and will serve as a very promising candidate for photonic integrated circuits.

Despite the exciting results of playing games, the DON currently has limitations for handling more complex tasks. First, for the sake of the computational requirements of optical forward propagation, we deploy a two-phase training architecture to obtain the policy model before iterating the DON instead of end-to-end learning in this work. Combining the two steps may reduce errors and make it easier to use. Second, ideally, the last layer of the network should not have a shortcut connection, which can be improved by modifying the experimental system. In addition, given the similar properties of the DON and MLP, the introduction of MLP-based attention mechanisms^{33,42} into the field of optics could be considered. Moreover, the inference and control capability of DONs could be improved by introducing methods such as nonlinear optical effects,^{43–46} multichannel structures,⁴⁷ and Fourier space²⁵ in the future, leading to a variety of new applications. While preliminary, this research suggests that the DON has great potential for processing complex visual inputs and tasks. It could provide a promising avenue for an optical computing system for decision-making and control, which would be a fruitful area for next-generation AI.

Disclosures

The authors declare no competing interests.

Code and Data Availability

Data underlying the results presented in this paper may be obtained from the authors upon reasonable request. The codes used in this work are available in the GitHub repository at <https://github.com/qiujumin/DON-RL>.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (Grant Nos. 12064025, 12264028, 12364045, and 12304420), the Natural Science Foundation of Jiangxi Province (Grant Nos. 20212ACB202006, 20232BAB201040, and 20232BAB211025), the Shanghai Pujiang Program (Grant No. 22PJ1402900), the Australian Research Council Discovery Project (Grant No. DP200101353), the Interdisciplinary Innovation Fund of Nanchang University (Grant No. 2019-9166-27060003), and the China Scholarship Council (Grant No. 202008420045).

References

1. A. Krizhevsky, I. Sutskever, and G. E. Hinton, “Imagenet classification with deep convolutional neural networks,” *Commun. ACM* **60**, 84–90 (2017).
2. O. Russakovsky et al., “Imagenet large scale visual recognition challenge,” *Int. J. Comput. Vision* **115**, 211–252 (2015).
3. Q. Chen et al., “Enhanced LSTM for natural language inference,” in *Proc. 55th Annu. Meeting of the Assoc. for Comput. Linguist. (Vol. 1: Long Pap.)*, Association for Computational Linguistics, Vancouver, Canada, pp. 1657–1668 (2017).
4. J. Devlin et al., “BERT: pre-training of deep bidirectional transformers for language understanding,” in *Proc. 2019 Conf. of the North Am. Chap. of the Assoc. for Comput. Linguist.: Hum. Lang. Technol., Vol. 1 (Long and Short Pap.)*, Association for Computational Linguistics, Minneapolis, Minnesota, pp. 4171–4186 (2019).
5. A. Grover and J. Leskovec, “Node2vec: scalable feature learning for networks,” in *Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discov. and Data Mining, KDD '16*, Association for Computing Machinery, New York, pp. 855–864 (2016).
6. W. Ma et al., “Deep learning for the design of photonic structures,” *Nat. Photonics* **15**, 77–90 (2021).
7. O. Khatib et al., “Learning the physics of all-dielectric metamaterials with deep lorentz neural networks,” *Adv. Opt. Mater.* **10**(13), 2200097 (2022).
8. O. Khatib et al., “Deep learning the electromagnetic properties of metamaterials—a comprehensive review,” *Adv. Funct. Mater.* **31**(31), 2101748 (2021).
9. L. Huang, L. Xu, and A. E. Miroshnichenko, “Deep learning enabled nanophotonics,” Chap. 4 in *Advances and Applications in Deep Learning*, M. A. Aceves-Fernandez, Ed., IntechOpen, Rijeka (2020).
10. C. C. Nadell et al., “Deep learning for accelerated all-dielectric metasurface design,” *Opt. Express* **27**, 27523–27535 (2019).
11. L. Xu et al., “Enhanced light–matter interactions in dielectric nanostructures via machine-learning approach,” *Adv. Photonics* **2**(2), 026003 (2020).
12. P. R. Wiecha and O. L. Muskens, “Deep learning meets nanophotonics: a generalized accurate predictor for near fields and far fields of arbitrary 3D nanostructures,” *Nano Lett.* **20**(1), 329–338 (2020).
13. P. Dai et al., “Accurate inverse design of Fabry–Perot-cavity-based color filters far beyond sRGB via a bidirectional artificial neural network,” *Photonics Res.* **9**, B236–B246 (2021).
14. Y. Shen et al., “Deep learning with coherent nanophotonic circuits,” *Nat. Photonics* **11**, 441–446 (2017).
15. J. Feldmann et al., “All-optical spiking neurosynaptic networks with self-learning capabilities,” *Nature* **569**, 208–214 (2019).
16. R. Hamerly et al., “Large-scale optical neural networks based on photoelectric multiplication,” *Phys. Rev. X* **9**, 021032 (2019).
17. H. Zhang et al., “An optical neural chip for implementing complex-valued neural network,” *Nat. Commun.* **12**, 457 (2021).
18. J. Liu et al., “Research progress in optical neural networks: theory, applications and developments,” *Photonix* **2**, 5 (2021).
19. Z. Wu et al., “Neuromorphic metasurface,” *Photonics Res.* **8**, 46–50 (2020).
20. X. Lin et al., “All-optical machine learning using diffractive deep neural networks,” *Science* **361**(6406), 1004–1008 (2018).
21. H. Chen et al., “Diffractive deep neural networks at visible wavelengths,” *Engineering* **7**(10), 1483–1491 (2021).
22. X. Luo et al., “Metasurface-enabled on-chip multiplexed diffractive neural networks in the visible,” *Light: Sci. Appl.* **11**, 158 (2022).
23. C. Liu et al., “A programmable diffractive deep neural network based on a digital-coding metasurface array,” *Nat. Electron.* **5**, 113–122 (2022).
24. H. Zheng et al., “Meta-optic accelerators for object classifiers,” *Sci. Adv.* **8**(30), eabo6410 (2022).
25. T. Yan et al., “Fourier-space diffractive deep neural network,” *Phys. Rev. Lett.* **123**, 023901 (2019).
26. C. Qian et al., “Performing optical logic operations by a diffractive neural network,” *Light: Sci. Appl.* **9**, 59 (2020).
27. T. Zhou et al., “Large-scale neuromorphic optoelectronic computing with a reconfigurable diffractive processing unit,” *Nat. Photonics* **15**, 367–373 (2021).
28. J. Schulman et al., “Proximal policy optimization algorithms,” <https://arxiv.org/abs/1707.06347> (2017).
29. D. P. Kingma and J. Ba, “Adam: a method for stochastic optimization,” in *3rd Int. Conf. Learn. Represent., ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conf. Track Proc.*, Y. Bengio and Y. LeCun, Eds. (2015).
30. K. He et al., “Deep residual learning for image recognition,” in *IEEE Conf. Comput. Vision and Pattern Recognit. (CVPR)*, pp. 770–778 (2016).
31. H. Dou et al., “Residual D²NN: training diffractive deep neural networks via learnable light shortcuts,” *Opt. Lett.* **45**, 2688–2691 (2020).

32. S. Zheng, S. Xu, and D. Fan, "Orthogonality of diffractive deep neural network," *Opt. Lett.* **47**, 1798–1801 (2022).
33. I. O. Tolstikhin et al., "MLP-mixer: an all-MLP architecture for vision," in *Advances in Neural Information Processing Systems*, M. Ranzato et al., Eds., Vol. **34**, pp. 24261–24272, Curran Associates (2021).
34. S. P. Rodrigues et al., "Weighing in on photonic-based machine learning for automotive mobility," *Nat. Photonics* **15**, 66–67 (2021).
35. A. H. Dorrah and F. Capasso, "Tunable structured light with flat optics," *Science* **376**(6591), eabi6860 (2022).
36. L. Li et al., "Machine-learning reprogrammable metasurface imager," *Nat. Commun.* **10**, 1082 (2019).
37. T. Liu et al., "Phase-change metasurfaces for dynamic image display and information encryption," *Phys. Rev. Appl.* **18**, 044078 (2022).
38. W. J. Padilla and R. D. Averitt, "Imaging with metamaterials," *Nat. Rev. Phys.* **4**, 85–100 (2022).
39. Z. Wang et al., "Arbitrary polarization readout with dual-channel neuro-metasurfaces," *Adv. Sci.* **10**(5), 2204699 (2022).
40. C. Qian et al., "Dynamic recognition and mirage using neuro-metamaterials," *Nat. Commun.* **13**, 2694 (2022).
41. C. He et al., "Pluggable multitask diffractive neural networks based on cascaded metasurfaces," *Opto-Electron. Adv.* **7**(2), 230005 (2024).
42. M.-H. Guo et al., "Beyond self-attention: external attention using two linear layers for visual tasks," *IEEE Trans. Pattern Anal. Mach. Intell.* **45**(5), 5436–5447 (2023).
43. C. Schlickriede et al., "Nonlinear imaging with all-dielectric metasurfaces," *Nano Lett.* **20**(6), 4370–4376 (2020).
44. Y. Xiao, H. Qian, and Z. Liu, "Nonlinear metasurface based on giant optical Kerr response of gold quantum wells," *ACS Photonics* **5**, 1654–1659 (2018).
45. M. Akie et al., "GeSn/SiGeSn multiple-quantum-well electroabsorption modulator with taper coupler for mid-infrared Ge-on-Si platform," *IEEE J. Sel. Top. Quantum Electron.* **24**(6), 1–8 (2018).
46. Y. Zuo et al., "All-optical neural network with nonlinear activation functions," *Optica* **6**, 1132–1137 (2019).
47. Z. Xu et al., "A multichannel optical computing architecture for advanced machine vision," *Light: Sci. Appl.* **11**, 255 (2022).

Jumin Qiu received his MS degree from Nanchang University, China, in 2023, where he is currently pursuing a PhD in physics at the School of Physics and Materials Science. His research interests include optical computing, nanophotonics, and computational imaging.

Shuyuan Xiao obtained his PhD from Huazhong University of Science and Technology, China, in 2018. He joined Nanchang University and was a research fellow at the Institute for Advanced Study from 2019

to 2024. He is currently a research fellow at the School of Information Engineering. He is a member of the *Journal of Optics* editorial board and a frequent reviewer for leading publications in physics, optics, and materials science. His current research interests focus on metasurfaces and nanophotonics.

Lujun Huang is a professor in the School of Physics and Electronic Science at East China Normal University (ECNU). He received his PhD in material science and engineering from North Carolina State University in 2017. Then, he was a research associate at the University of New South Wales from 2018 to 2021. He has been a full professor at ECNU since July 2022. His current research interests focus on resonant nanophotonics and resonant acoustics as well as light-matter interaction of 2D materials.

Andrey Miroshnichenko received his PhD from the Max-Planck Institute for Physics of Complex Systems in Dresden, Germany, in 2003. In 2004, he moved to Australian National University, Australia. During that time, he made fundamentally important contributions to the field of photonic crystals and brought the concept of the Fano resonances to nanophotonics. In 2017, he moved to the University of New South Wales Canberra, Australia. The topics of his research are nanophotonics, nonlinear and quantum optics, and resonant interaction of light with nanoclusters, including optical nanoantennas and metamaterials.

Dejian Zhang is a lecturer at the School of Physics and Materials Science, Nanchang University, China. He received his PhD in optics from the Department of Physics, Beijing Normal University in 2016, China. His current research interests include computational optical imaging and optical neural networks.

Tingting Liu received her PhD in information and communication engineering from the Huazhong University of Science and Technology, China. Then, she joined the Hubei University of Education, China, where she was granted tenure and promoted to an associate professor. Next, she moved to the University of Shanghai for Science and Technology, China, and was a research fellow at the Institute of Photonic Chips. Currently, she is an associate professor at the Nanchang University, China. Her research interests focus on light field manipulation and optical neural networks.

Tianbao Yu is a professor and a PhD supervisor at the School of Physics and Materials Science, Nanchang University, China. He has been a professor since December 2013. He received his PhD in optics from the Department of Information Science and Electronic Engineering, Zhejiang University, China, in 2007. He is the author of more than 120 journal papers. His current research interests include nanophotonics and computational optical imaging.