Visual Explanation using Attention Mechanism in Actor-Critic-based Deep Reinforcement Learning

Hidenori Itaya, Tsubasa Hirakawa, Takayoshi Yamashita, Hironobu Fujiyoshi,

Chubu University

1200 Matsumotocho, Kasugai, Aichi, Japan

{itaya, hirakawa}@mprg.cs.chubu.ac.jp, {takayoshi, fujiyoshi}@isc.chubu.ac.jp

Komei Sugiura Keio University 3-14-1 Hiyoshi, Kohoku, Yokohama, Kanagawa, Japan komei.sugiura@keio.jp

Abstract-Deep reinforcement learning (DRL) has great potential for acquiring the optimal action in complex environments such as games and robot control. However, it is difficult to analyze the decision-making of the agent, i.e., the reasons it selects the action acquired by learning. In this work, we propose Mask-Attention A3C (Mask A3C), which introduces an attention mechanism into Asynchronous Advantage Actor-Critic (A3C), which is an actor-critic-based DRL method, and can analyze the decision-making of an agent in DRL. A3C consists of a feature extractor that extracts features from an image, a policy branch that outputs the policy, and a value branch that outputs the state value. In this method, we focus on the policy and value branches and introduce an attention mechanism into them. The attention mechanism applies a mask processing to the feature maps of each branch using mask-attention that expresses the judgment reason for the policy and state value with a heat map. We visualized mask-attention maps for games on the Atari 2600 and found we could easily analyze the reasons behind an agent's decisionmaking in various game tasks. Furthermore, experimental results showed that the agent could achieve a higher performance by introducing the attention mechanism.

I. INTRODUCTION

Reinforcement learning (RL) problems seek optimal actions to maximize cumulative rewards. Unlike supervised learning problems, RL problems collect training data by exploring the environment. Therefore, RL has achieved high performance in specific tasks (e.g., controlling autonomous systems [1]–[3] and video games [4]–[6]) in which it is difficult to create training data. In Go, AlphaGo has defeated a professional Go player [7]. In 2015, the deep Q-network (DQN), a method that combines Q-learning [8] and deep neural network (DNN), achieved a score higher than human players on the Atari 2600 [9]. Since the advent of DQN, deep RL (DRL), a method that combines deep learning and RL, has become mainstream, and it is now possible to solve problems featuring a huge number of states, such as images.

In general, deep learning can solve complex tasks by training using a large number of network parameters. However, it is difficult to understand the reasoning behind the decisionmaking of the trained network because the number of network parameters used to make the decision is enormous. This problem occurs in DRL as well. The reason for judging the acquired action is unclear, since agents collect training data by searching the environment and the calculation inside the network is complicated. Therefore, in order to prove that the trained network is sufficiently reliable, it is important to analyze the reason for the judgment of the action that it outputs.

One approach to interpreting the decision-making of a network, visual explanation, has been studied in the field of computer vision [10]–[12]. Visual explanations analyze the factors of the network output by using an attention map that highlights the important regions in an input image. Visual explanation methods have also been applied to DRL models to help with understanding the decision-making of an agent [13]. [14]. These methods can be categorized into two approaches: bottom-up and top-down. Bottom-up visual explanations compute attention maps by using the gradient information of a network. Because the bottom-up approach does not need to retrain a network, it can be applied to any trained network and is commonly used in computer vision and DRL. The attention maps obtained by the bottom-up approach are based on the input data and response values calculated from each layer. The bottom-up approach highlights local textural context. Topdown visual explanations generate attention maps by using the response values in a network. In contrast to the bottom-up approach, the attention maps of the top-down approach are output for the current network output.

In this paper, we propose Mask-Attention A3C (Mask A3C), which introduces an attention mechanism into Asynchronous Advantage Actor-Critic (A3C), an actor-critic-based DRL method. Mask A3C calculates a mask-attention that is an attention map of the policy and state value, and then a visual explanation for these values is achieved by visualizing the created mask-attention. Our method also learns the policy and state value while considering mask-attention by implementing the attention mechanism, thereby improving the performance of the agent.

1) Contributions: The main contributions of this paper are as follows.

- We propose a top-down visual explanation method that implements an attention mechanism in the DRL model. In the proposed method, mask-attention, which is an attention map for the outputs, can be obtained simply by forward pass.
- In the proposed method, the decision-making of the agent after learning can be analyzed by visualizing the acquired mask-attention. We conducted an experiment with games on the Atari 2600 and analyzed which information influences the agent's decision-making.
- By implementing the attention mechanism in the policy branch and value branch of the actor-critic method, a different mask-attention can be obtained depending on the policy and state value. In this way, it is possible to analyze an agent's decision-making from the two viewpoints of policy value and state value.
- The proposed method outputs the control value of the agent while considering mask-attention by implementing the attention mechanism. Therefore, the performance of the agent can be improved by emphasizing the area related to the control value.

The remainder of this paper is organized as follows. Section II reviews the related work. Section III introduces the maskattention A3C. Section IV shows experimental results on Atari 2600 environments. We conclude the paper in Section V.

II. RELATED WORKS

A. Deep reinforcement learning

The deep Q-network (DQN) [9], which is a typical method of DRL, expresses the action value function $Q(a|s;\theta)$ by using a neural network and acquires the optimum action by training the network parameters θ . DRL methods that learn the optimal action by a value function such as DQN are called value-based DRL, and have been studied extensively [15]-[18]. There is also a policy-based DRL that directly learns the policy by expressing the policy $\pi(a|s;\theta)$ with a neural network [19]– [22]. The actor-critic method [23], which is a policy-based method, consists of an actor that outputs the policy $\pi(a|s;\theta)$ and a critic that outputs the state value $V(s; \theta)$. Here, the state value $V(s; \theta)$ numerically expresses how the current state s contributes to the reward. The actor selects and performs an action according to a policy $\pi(a|s;\theta)$ that is a probability distribution from a state s to an action a. The critic estimates the state value $V(s; \theta)$ as the evaluation value of the policy $\pi(a|s;\theta)$ that is output by the actor. To update the network parameters in the actor-critic method, the actor parameter update by the policy gradient method and the critic parameter update by the TD error are performed in parallel.

Other approaches include distributed DRL, which improves learning efficiency by constructing multiple environments and agents [24]–[26]. A3C [27] is a distributed DRL method on the basis of the actor-critic method. A3C introduces Asynchronous, which is an asynchronous parameter update in distributed learning and Advantage, which that learns while considering rewards several steps ahead. Experiments with the Atari 2600 showed that A3C could achieved a high score in a short training time by executing the generation of experiences used for learning in parallel.

In this study, we acquire mask-attention, attention maps for policy and state value, by implementing an attention mechanism in A3C. By visualizing mask-attention at the time of inference, the decision-making of the agent acquired by learning is analyzed from the visual explanation of the policy and state value.

B. Visual explanations

1) Visual explanations in image recognition: In the field of image recognition, several methods utilizing an attention map have been proposed for analyzing the reason for judgments on the inference result of the network. An attention map visualizes the network attended area at the time of inference. Zhou et al. proposed a class activation mapping (CAM) [10], which acquires the attention map of a specific class from the response value of the convolutional layer and the weight of the fully connected layer. However, the recognition performance of CAM deteriorates because it is necessary to change a part of the network structure, such as by introducing global average pooling (GAP) between the convolution and fully connected layers. For that problem, Selvaraju et al. proposed gradientweighted CAM (Grad-CAM) [11], which acquires an attention map by using the response value of the convolutional layer during forward pass and the gradient during back-propagation. Grad-CAM avoids the deterioration of the recognition performance deteriorates by generating an attention map from the gradient information. In image recognition, the recognition accuracy is known to improve by using an attention map during learning. Fukui et al. proposed an attention branch network (ABN) [12] that applies the attention map to the attention mechanism. This method provides a visual explanation of the reason for judgment by the attention map and simultaneously improves the recognition accuracy.

2) Visual explanations in deep reinforcement learning: In DRL, several works for visual explanation of DRL models have examined. Sorokin *et al.* proposed the deep attention recurrent Q-network (DARQN) [13], where an attention mechanism is implemented in DQN, a representative value-based method. Manchin *et al.* introduced a self-attention to a policy-based DRL method [28], in order to improve the score along with policy analysis. This method analyzes the agent's decision-making by using an attention map for the policy. Our method differs in that it can improve the interpretability of the agent's decision-making in the actor-critic-based DRL method by simultaneously acquiring different attention maps for policies and state values.

Greydanus *et al.* acquired a saliency map in A3C by calculating a perturbation image utilizing an applied Gaussian filter from the gradient during back-propagation [14]. This method takes a bottom-up approach, similar to Grad-CAM, and therefore it is necessary to perform back-propagation to acquire the saliency map. In contrast, our method takes a top-



Fig. 1: Detailed network structure of Mask-Attention A3C.

down approach that implements an attention mechanism in the network structure of A3C.

Zhang *et al.* proposed attention guided imitation learning (AGIL) [29], which guides the focus area of a network on the basis of human gaze information. They trained a model to replicate human attention with supervised gaze heat maps. The input state was then augmented with this additional information. This style of attention fundamentally differs from that used in our work in that it incorporates hand crafted features as input.

The most similar work to ours was conducted by Mott *et al.* [30]. Their method acquires two attentions ("what" and "where") by using query-based attention in an actor-criticbased DRL method. This method requires an attention query to be generated, which means significant changes must be made to the network architecture (e.g., keys, values). In contrast, our method has a simple structure in which an attention mechanism is implemented in the policy and value branches and no significant changes to the network architecture are required. Also, they obtained different attentions related to "what" and "where" by generating an attention query. In contrast, our method improves interpretability by acquiring different attentions toward policy and state value, which are the outputs of actor-critic-based DRL methods.

The above visual explanations for DRL generates attention maps from the low-level feature maps extracted from early or middle convolutional layers, or output for action. In DRL, a *strategy* is an important clue to solve a given task and environment. From the viewpoint of the strategy, the state value of actor-critic based model plays a crucial role because the state value is the expected value from current to future states and affects the future action selections. However, the existing methods output attention maps with respect to the instantaneous action selection. Our method based on an actorcritic outputs two attention maps from both the policy and value branches. By considering both attention maps, we can clarify the basis of an agent's decision-making in more detail.

III. MASK-ATTENTION A3C

We propose Mask-Attention A3C (Mask A3C), which introduces an attention mechanism into A3C, an actor-critic-based distributed DRL method. In Mask A3C, by implementing an attention mechanism for the policy branch and the value branch, we can acquire an attention map that expresses the focus area of the network for the output of each branch. In our study, this attention map is called mask-attention. By visualizing the mask-attention of each branch, we obtain a visual explanation of the reason for the judgment on the policy and state value. In addition, our method learns while considering the mask-attention by implementing the attention mechanism, which improves the performance of the agent.

A. Overview of Mask A3C structure

Figure 1 shows the network structure of the proposed Mask A3C. It consists of a feature extractor, policy branch, value branch, and attention mechanism. Here, let s_t be a state at time t. First, the feature extractor extracts a feature map $F_{fe}(s_t)$ from the given state s_t . The feature extractor consists of convolutional layers. In A3C, temporal information can be considered by utilizing LSTM, which greatly improves the performance of the agent. However, LSTM cannot consider the spatial information of the input image, so if it is used in Mask A3C, mask-attention cannot be calculated. We therefore utilize convolutional LSTM (ConvLSTM) [31], which can consider spatiotemporal information.

The extracted feature map is fed into the policy and value branches. The policy branch outputs policies, and the value branch outputs the state value function. At these branches, the proposed method adds mask-attention module. Each branch takes the feature map extracted from the feature extractor and computes a new feature map $F_v(\mathbf{s}_t)$ and $F_p(\mathbf{s}_t)$ by applying a convolution and a ReLU activation. Meanwhile, the proposed method also generates mask-attention. We denote $M_v(\mathbf{s}_t)$ and $M_p(\mathbf{s}_t)$ as the mask-attentions for value and policy branches, respectively. The mask-attention can be generated by applying one convolutional layer of $1 \times 1 \times \#$ of channels and a sigmoid function to the feature map \mathbf{s}_t .

These feature maps and mask-attention are used for attention mechanism, the each branch output policy and state value, respectively.

B. Attention mechanism

Mask A3C implements an attention mechanism in the policy branch and value branch so that the policy and state value functions are learned in consideration of the acquired maskattention. The attention mechanism performs mask processing on the feature map of the middle layer in each branch by using mask-attention. With this mask processing, the area that contributes to the optimum action and state value can be emphasized. By using mask-attention for the feature map, the mask processing for each branch $F'_v(\mathbf{s}_t)$ and $F'_p(\mathbf{s}_t)$ are calculated as follows:

$$F'_{v}(\mathbf{s}_{t}) = F_{v}(\mathbf{s}_{t}) \cdot M_{v}(\mathbf{s}_{t}), \qquad (1)$$

$$F'_p(\mathbf{s}_t) = F_p(\mathbf{s}_t) \cdot M_p(\mathbf{s}_t).$$
⁽²⁾

The masked feature maps $F'_v(\mathbf{s}_t)$ and $F'_p(\mathbf{s}_t)$ are then fed into the output layer and we obtain state value and policy. By using the masked feature map, the agent focuses on the highlighted region and selects the optimal action.

IV. EXPERIMENTS

To evaluate the effectiveness of Mask A3C, we conducted experiments using the game task of OpenAI gym [32]. Three games were used: "Ms. Pac-Man", "Space Invaders", and "Seaquest". The comparison methods were A3C, Policy Mask A3C, Value Mask A3C, and Mask A3C, for a total of four patterns. Policy Mask A3C and Value Mask A3C refer to a Mask A3C in which the attention mechanism is implemented in only one branch (i.e., policy branch or value branch). The learning conditions were 35 for the number of workers, 0.0001 for the learning coefficient, and 0.99 for the discount rate. The termination condition of learning was when the global steps reached 1.0×10^8 . The termination condition of an episode was the end of one play in the game and the case where the number of steps reached 1.0×10^4 . We used the following four evaluation methods.

- Visual explanations using mask-attention
- Score comparison on the Atari 2600
- Score comparison by inverting the gaze area of maskattention
- · Reaction of mask-attention to new states

A. Implementation details

The input was a grayscale image of the game screen and the output was the action in each game. The image used as input was resized to 80×80 . The output dimension of the feature extractor was 32-dimensional for the first two convolutional

layers and 64-dimensional for the one remaining convolutional layer. For the hidden state in ConvLSTM, the output dimension was 64-dimensional. The policy branch consisted of one convolutional layer, one fully connected layer, and a softmax function. The output dimension of the convolution layer was 32-dimensional and the number of output units of the fully connected layer was the number of actions in each game. The value branch consisted of one convolutional layer with a 32-dimensional output dimension and one fully connected layer with one output unit. The network structure of A3C in this experiment was the structure excluded the attention mechanism from Mask A3C.

B. Visual explanations using mask-attention

Figures 2 and 3 show visualization examples of maskattention in Atari 2600. Hereafter, we discuss the obtained mask-attentions shown in these figures.

1) Ms. Pac-Man: Ms. Pac-Man is a game in which the player collects scattered cookies while avoiding enemies. The actions of the agent that is the Pac-Man are "Noop", "Up", "Down", "Left", "Right", "Up + Left", "Up + Right", "Down + Left", and "Down + Right". In Fig. 2(a), the agent of Frame 1 was attending around Pac-Man. The agent of Frame 2 was attending to the cookies remaining on the screen. In Frame 3, Pac-Man moved to the gazed area of Frame 2 and acquired the cookies. These results demonstrate that the agent controlled Pac-Man toward the cookies while simultaneously attending to the surroundings of Pac-Man. In Fig. 3(a), Point 1, which was the beginning of the game, the agent was attending to the entire screen. In contrast, Point 2 is reduced the gaze area in accordance with the decreases of cookies. Also, from Point 1 to Point 2, the state value decreased as the cookies on the screen decreased. These results demonstrate that the agent recognized the cookies as the score source.

2) Space Invaders: Space Invaders is a shooting game in which the player repels the enemy invaders. The actions of the agent that is the cannon are "Noop", "Left", "Right", "Attack", "Left + Attack", and "Right + Attack". In Fig. 2(b), the agent in Frame 1 is attended to the invaders and the action was "Attack". In Frame 2, we can see that the beam in Frame 1 was heading toward the invader, and in Frame 3, the beam was repelling the invaders that were attended to in Frame 1. In all frames, the agent attended around itself while avoiding the defensive walls. From these results, we can see that the agent repelled the invaders while simultaneously avoiding the defensive wall. In Fig. 3(b), the agent in Point 1 was attending to all of the invaders and the agent in Point 2 was shrinking the gaze area in accordance with the decreasing number of invaders. In addition, from Point 1 to Point 2, the state value is decreased according to the number of invaders. These results demonstrate that the agent recognized the invaders as the score source.

3) Seaquest: Seaquest is a game in which players use submarines to rescue divers and destroy enemies and fish. The actions of the agent that is the submarine are "Noop", "Up", "Down", "Left", "Right", and "Attack". In Fig. 2(c), the agent







(b) **Space Invaders**: The red circles in Frame 2 show the beam that is the attack of the agent in Frame 1, and the red circles in Frame 3 show the destroyed invaders. The white arrow shows the direction of travel of the beam that is the attack of the agent.



(c) Seaquest: The red circle in Frame 2 shows the submarine replenishing oxygen.

Fig. 2: Visualization example of mask-attention in policy: The controller in "Image with mask-attention" is an action that is output by the DRL model. The green broken line in the State value shows the transition to the next stage.



Fig. 3: Visualization example of mask-attention in state value: The green broken line in the graph shows the transition to the next stage.

in Frame 1 is attending to the oxygen gauge and the action was "Up". In Frame 2, the submarine operated by the agent is floating on the surface of the sea. When the amount of oxygen is reduced in Seaquest, it can be replenished when the submarine rises to the surface of the sea. In Frame 3, we can see that the oxygen gauge is full due to the rise to the sea level in Frame 2, and there is no gaze on the oxygen gauge. From these results, we can see that the agent recognizes the oxygen gauge is low and controls the submarine to rise to sea level. In Fig. 3(c), the agent in Point 1 was attending to the entire oxygen gauge. On the other hand, in Point 2, where the amount of oxygen is low, the gaze area of the oxygen gauge was reduced in accordance with the amount of oxygen. The state value also decreased from Point 1 to Point 2. These results demonstrate that the agent recognized the oxygen as the score source.

4) Discussion: We obtained different mask-attentions in the policy and state value by implementing an attention mechanism for each branch. The policy represents the probability distribution of the action selection in the current state. Therefore, the mask-attention of the policy indicates the area that contributes to the action of the agent. The state value represents the expected value of the return in the current state. Here, return is the sum of rewards during one episode. Therefore, the mask-attention of the state value indicates the area that represents the property of the game.

C. Score comparison on the Atari 2600

Table I shows the max and mean scores over 100 episodes on the Atari 2600. As we can see, the mean score in Breakout was lower for Policy Mask A3C, Value Mask A3C, and Mask A3C than for A3C. In contrast, the max score in Breakout was 864 for all methods. This score is the best score that can

TABLE I: Max and mean scores over 100 episodes on Atari 2600: Scores of models that had the highest average score among five trials in each method are shown.

Att. mechanism		Breakout		Ms. Pac-Man		Space	Invaders	Beamrider		Fishing Derby		Seaquest	
Policy	Value	max	mean	max mean		max	mean	max	mean	max	mean	max	mean
-		864	662.0	5380	4573.3	19505	18531.8	34748	28341.1	41	32.1	2760	2728.2
\checkmark		864	595.8	6330	4833.8	19860	19102.8	32604	28495.3	41	37.5	2820	2784.0
	\checkmark	864	606.9	4830	4044.5	19675	18537.8	35108	28205.7	43	36.1	2820	2786.4
\checkmark	\checkmark	864	640.0	6610	5314.1	19810	19212.5	34448	27671.1	41	34.3	17150	6701.9

TABLE II: Score comparison by inverting the gaze area of mask-attention: Normal and inverse are the scores when the gaze area was not inverted and is inverted, respectively. Random is the score when the action was randomly selected. Max / mean = maximum and average scores over 100 episodes.

Att. mechanism			Breakout		Ms. Pac-Man		Space Invaders		Beamrider		Fishing Derby		Seaquest	
Policy	Value		max	mean	max	mean	max	mean	max	mean	max	mean	max	mean
		normal	864	595.8	6630	4833.8	19860	19102.8	32604	28495.3	41	37.5	2820	2784.0
v		inverse	4	2.2	290	268.9	805	306.9	4996	1554.2	-49	-75.7	280	158.2
\checkmark	1	normal	864	640.0	6610	5314.1	19810	19212.5	34448	27671.1	41	34.3	17150	6701.9
	v	inverse	5	1.8	410	194.4	915	420.2	6380	2063.9	-49	-74.7	420	280.6
random		5	1.2	1080	247.8	460	142.1	852	356.5	-85	-93.1	300	82.8	

TABLE III: Decrease rate of score due to inverse of gaze area in mask-attention (%): Max / mean = maximum and average scores over 100 episodes.

Att. mechanism		Breakout		Ms. Pac-Man		Space Invaders		Beamrider		Fishing Derby		Seaquest	
Policy	Value	max	mean	max	mean	max	mean	max	mean	max	mean	max	mean
\checkmark		99.53	99.63	95.41	94.43	95.94	98.39	84.67	94.54	98.90	99.12	90.07	94.31
\checkmark	\checkmark	99.42	99.71	93.79	96.34	95.38	97.80	81.47	92.54	98.90	99.09	97.55	95.81

be obtained in Breakout. Breakout is a simple game with no external factors: it simply consists of the player hitting the ball back with a paddle. Therefore, we presume that A3C and Mask A3C got the same score. In Ms. Pac-Man, the max and mean scores of Policy Mask A3C and Mask A3C improved compared to those of A3C. The control of the agent in Ms. Pac-Man is complex because it is necessary to select an action while considering external factors (e.g., the enemy). Policy Mask A3C and Mask A3C, which implement an attention mechanism on the policy branch, can emphasize the areas that contribute to the action (e.g., cookies and enemies), which is why Policy Mask A3C and Mask A3C obtained a higher score than A3C. In Space Invaders, Policy Mask A3C and Mask A3C improved the max and mean scores compared to A3C. Also, the max and mean scores of Value Mask A3C were almost the same as those of A3C. The agent in Space Invaders needs to select actions in consideration of external factors (e.g., enemies), the same as in Ms. Pac-Man. Policy Mask A3C and Mask A3C, which implement an attention mechanism in the policy branch, can emphasize the areas that contribute to the action (e.g., defensive walls and invaders), which is why they obtained a higher score than A3C. There was no significant difference among the scores for Beamrider. In Beamrider, two kinds of enemies exist from the agent's point of view: one, an enemy that the agent should defeat, and two, the necessity that the agent should avoid collisions. These enemies are like similar in appearance. Since our attention mechanism and the mask-attention that highlights the enemies are insufficient to distinguish these enemies, they do not contribute to the score improvement. In Fishing Derby, Policy Mask A3C, Value Mask A3C, and Mask A3C improved the mean scores compared to A3C. In Fishing Derby, there are many fish that are score sources, and getting the closest fish is the fastest way to get a point. Policy Mask A3C, Value Mask A3C, and Mask A3C, which implements an attention mechanism, can emphasize the fish closest to the player. This is why the mask-attention methods obtained a higher score than that A3C. In Seaguest, Mask A3C improved both the maximum and average scores compared to the other methods. This is because only the Mask A3C agent could learn the action of replenishing oxygen. Seaquest features an oxygen gauge at the bottom of the screen, and the game ends when the oxygen is gone. Mask A3C, which implements an attention mechanism in the policy and value branches, can emphasize the oxygen gauge, which is why Mask A3C obtained a higher score than the other methods.

D. Comparison of scores by inverting gaze area in maskattention

In visually explaining the decision-making of an agent using Mask A3C, we want to verify whether mask-attention represents an effective gaze area for the optimum action. In a case where the game score obtained by inverting maskattention does not change, it means the mask-attention does not contribute to the agent's action. In contrast, if the game score significantly decreases, it means the mask-attention largely contributes to the agent's actions acquiring the game score. For this verification method, we created a map in which the gaze area of the mask-attention in the policy branch is inverted and then calculated the score on Atari 2600 when the created map was used for the attention mechanism. We investigated whether mask-attention is effective for the visual explanation of an action by comparing the scores when the gaze area was inverted and when it was not inverted. The map in which the gaze area of the mask-attention is inverted was created by

$$M_{\rm inverse}(\mathbf{s}_t) = 1 - M(\mathbf{s}_t),\tag{3}$$

where \mathbf{s}_t is the state (grayscale image in the experiment), $M(\mathbf{s}_t)$ is mask-attention, and $M_{\text{inverse}}(\mathbf{s}_t)$ is the map after inverting the gaze area of mask-attention.

Table II shows the score comparison by inverting the gaze area of mask-attention. (Also, Table III shows the decrease ratio from the normal score to the inverse score.) As shown in Table II, the inverse score was significantly lower than the normal score in all games. In addition, the inverse score in Breakout was equivalent to a random score, and the inverse score of Mask A3C in Ms. Pac-Man was 53.4 lower than random. In contrast, we can see that the inverse score in Space Invaders, Beamrider, Fishing Derby, and Seaquest was higher than a random score. However, from Table III, we can see that the decrease ratio in the mean of Space Invaders, Beamrider, Fishing Derby, and Seaquest was more than 90%, as with the other games. Therefore, we conclude that the gaze area of mask-attention in the policy branch can represents a useful area for action to obtain a high score.

E. Reaction of mask-attention to new states

We want to verify whether mask-attention and the behavior of an agent are affected when changes are made to the object that mask-attention was attending to. In a case where the behavior of the agent significantly changes by making changes to the object being attended, the object is an important component that contributes to the agent's behavior. In this experiment, we focused on the fish and oxygen gauges in Seaquest. For this verification method, we added the fish to the image when evaluating Mask A3C in Seaquest. The frame to add fish is the frame in which the agent destroys the fish and the fish disappears from the image. After adding the fish to the image, we investigated the behavior of the agent and the change in mask-attention. Similarly, we added the oxygen gauge full of oxygen to the image. The frame to add the oxygen gauge is before the agent replenishes oxygen.

Figure 4a shows the changes in agent behavior and maskattention due to the addition of fish. Here, the fish was added after Frame 2. From the Figure 4a, we can see that the policy / value mask-attention after Frame 3 is attending to the fish. At this time, the agent is moving to the right in Frame 3, but is attacking the fish added in Frame 4. These results demonstrate that the fish in Seaquest is an object that greatly contributes to the behavior of the agent. Figure 4b shows changes in agent behavior and mask-attention due to the addition of oxygen gauge. Here, the oxygen gauge is added after Frame 2. From the figure 4b, we can see that the policy mask-attention after Frame 3 is not attending to the oxygen gauge. Also, the value mask-attention of Frames 3 and 4 in Fig. 4b is not attending to the entire oxygen gauge, unlike Fig. 3c Point 1. However, the value mask-attention of Frames 5 and 6 is attending to the entire oxygen gauge, similar to Fig. 3c Point 1. At this time, the agent's behavior rose to sea level in Frames 3 and 4, but descended in Frames 5 and 6. These results demonstrate that the amount of oxygen in the oxygen gauge in Seaquest is an object that greatly contributes to the behavior of the agent.

The policy mask-attention in the experiment of adding the oxygen gauge is not attending to the oxygen gauge in Frame 3. Here, Frame 3 is the frame immediately after adding the oxygen gauge. In other words, the policy mask-attention immediately reflects the effect of adding the oxygen gauge. On the other hand, in the value mask-attention in the same experiment, the entire oxygen gauge is not attended to in Frames 3 and 4, but it is attended to in Frames 5 and 6. These results demonstrate that the change in mask-attention due to the addition of an oxygen gauge is different for the policy and value mask-attention. Therefore, we conclude that policy mask-attention shows the area that contributes to the behavior in the current state, and value mask-attention shows the area related to the characteristic of the game considering the time series.

V. CONCLUSION

In this paper, we proposed Mask-Attention A3C (Mask A3C), which introduces an attention mechanism into Asynchronous Advantage Actor-Critic (A3C). In Mask A3C, we acquire a mask-attention that expresses the important area for the policy and state value by implementing an attention mechanism in the policy and value branches. This enables a visual explanation of the judgment reason in the decisionmaking of the agent, from the two viewpoints of policy and state value, by visualizing mask-attention. We also emphasize which areas contribute to the optimal action and state value by implementing an attention mechanism, which simultaneously improves the performance of the agent.

Experiments with the Atari 2600 confirmed the acquisition of different mask-attentions in the policy and value branches. The results demonstrate that the mask-attention of the policy branch indicates the area that contributes to the action while that of the value branch indicates the area that expresses the property of the game. We provided a useful analysis for the decision-making of agents in game tasks from the two viewpoints of policy and state value by visualizing these maskattentions. A comparison of game scores showed that the score improved when the attention mechanism was implemented in the policy branch. However, our experiments were conducted with game tasks that are easy to analyze visually. Our future work will entail the visual analysis of agents using maskattention is a future work for complex tasks (e.g., robot control and autonomous driving).

ACKNOWLEDGMENT

This paper is based on results obtained from a project, JPNP20006, commissioned by the New Energy and Industrial Technology Development Organization (NEDO).



(b) Reaction to oxygen-related new states

Fig. 4: Visualization example of mask-attention to new states

REFERENCES

- J. Kober, J. A. Bagnell, and J. Peters, "Reinforcement learning in robotics: A survey," *The International Journal of Robotics Research*, vol. 32, no. 11, pp. 1238–1274, 2013.
- [2] S. Gu, E. Holly, T. Lillicrap, and S. Levine, "Deep reinforcement learning for robotic manipulation with asynchronous off-policy updates," in *Proceedings of international conference on robotics and automation* (*ICRA*), 2017, pp. 3389–3396.
- [3] A. Rajeswaran, V. Kumar, A. Gupta, G. Vezzani, J. Schulman, E. Todorov, and S. Levine, "Learning Complex Dexterous Manipulation with Deep Reinforcement Learning and Demonstrations," *Proceedings* of Robotics: Science and Systems (RSS), 2018.
- [4] C. Tessler, S. Givony, T. Zahavy, D. J. Mankowitz, and S. Mannor, "A deep hierarchical approach to lifelong learning in Minecraft," AAAI Conference on Artificial Intelligence, 2017.
- [5] N. Justesen, P. Bontrager, J. Togelius, and S. Risi, "Deep Learning for Video Game Playing," arXiv preprint, arXiv:1708.07902, 2017.
- [6] K. Shao, Z. Tang, Y. Zhu, N. Li, and D. Zhao, "A survey of deep reinforcement learning in video games," arXiv preprint arXiv:1912.10944, 2019.
- [7] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot *et al.*, "Mastering the game of Go with deep neural networks and tree search," *Nature*, vol. 529, no. 7587, pp. 484–489, 2016.
- [8] C. J. Watkins and P. Dayan, "Q-Learning," *Machine learning*, vol. 8, no. 3-4, pp. 279–292, 1992.
- [9] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski *et al.*, "Human-level control through deep reinforcement learning," *Nature*, vol. 518, no. 7540, pp. 529–533, 2015.
- [10] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba, "Learning deep features for discriminative localization," in *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2016, pp. 2921–2929.
- [11] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual Explanations from Deep Networks via Gradient-based Localization," *Proceedings of the IEEE international conference on computer vision (ICCV)*, pp. 618–626, 2017.
- [12] H. Fukui, T. Hirakawa, T. Yamashita, and H. Fujiyoshi, "Attention Branch Network: Learning of Attention Mechanism for Visual Explanation," in *Proceeding of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2019.
- [13] I. Sorokin, A. Seleznev, M. Pavlov, A. Fedorov, and A. Ignateva, "Deep attention recurrent q-network," *Proceedings of neural information processing systems (NeurIPS) workshops*, 2015.
- [14] S. Greydanus, A. Koul, J. Dodge, and A. Fern, "Visualizing and understanding Atari agents," *Proceedings of international conference on machine learning (ICML)*, vol. 80, pp. 1792–1801, 2018.
- [15] H. Van Hasselt, A. Guez, and D. Silver, "Deep Reinforcement Learning with Double Q-Learning," in the AAAI conference on artificial intelligence (AAAI), 2016.
- [16] Z. Wang, T. Schaul, M. Hessel, H. Hasselt, M. Lanctot, and N. Freitas, "Dueling Network Architectures for Deep Reinforcement Learning," in *International conference on machine learning (ICML)*, 2016, pp. 1995– 2003.
- [17] M. G. Bellemare, W. Dabney, and R. Munos, "A distributional perspective on reinforcement learning," in *International conference on machine learning (ICML)*, 2017, pp. 449–548.
- [18] M. Hessel, J. Modayil, H. Van Hasselt, T. Schaul, G. Ostrovski, W. Dabney, D. Horgan, B. Piot, M. Azar, and D. Silver, "Rainbow: Combining improvements in deep reinforcement learning," in *the AAAI conference on artificial intelligence (AAAI)*, 2018.
- [19] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, "Continuous control with deep reinforcement learning," in *International conference on learning representations* (*ICLR*), 2016.
- [20] J. Schulman, S. Levine, P. Abbeel, M. Jordan, and P. Moritz, "Trust Region Policy Optimization," in *International conference on machine learning (ICML)*, 2015, pp. 1889–1897.
- [21] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," in *arXiv preprint* arXiv:1707.06347, 2017.

- [22] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, "Soft actor-critic: Offpolicy maximum entropy deep reinforcement learning with a stochastic actor," in *International conference on machine learning (ICML)*, 2018.
- [23] V. R. Konda and J. N. Tsitsiklis, "Actor-Critic Algorithms," in Proceeding of neural information processing systems (NeurIPS), 2000, pp. 1008–1014.
- [24] A. Nair, P. Srinivasan, S. Blackwell, C. Alcicek, R. Fearon, A. De Maria, V. Panneershelvam, M. Suleyman, C. Beattie, S. Petersen *et al.*, "Massively parallel methods for deep reinforcement learning," in *International conference on machine learning (ICML) deep learning workshop*, 2015.
- [25] M. Jaderberg, V. Mnih, W. M. Czarnecki, T. Schaul, J. Z. Leibo, D. Silver, and K. Kavukcuoglu, "Reinforcement Learning with Unsupervised Auxiliary Tasks," in *International conference on learning representations (ICLR)*, 2017.
- [26] S. Kapturowski, G. Ostrovski, J. Quan, R. Munos, and W. Dabney, "Recurrent Experience Replay in Distributed Reinforcement Learning," in *International conference on learning representations (ICLR)*, 2019.
- [27] V. Mnih, A. P. Badia, M. Mirza, A. Graves, T. Lillicrap, T. Harley, D. Silver, and K. Kavukcuoglu, "Asynchronous Methods for Deep Reinforcement Learning," in *International conference on machine learning* (*ICML*), 2016, pp. 1928–1937.
- [28] A. Manchin, E. Abbasnejad, and A. van den Hengel, "Reinforcement learning with attention that works: A self-supervised approach," in *In International Conference on Neural Information Processing. Springer, Cham.*, 2019, pp. 223–230.
- [29] R. Zhang, Z. Liu, L. Zhang, J. A. Whritner, K. S. Muller, M. M. Hayhoe, and D. H. Ballard, "AGIL: Learning Attention from Human for Visuomotor Tasks," in *Proceedings of the european conference on computer vision (ECCV)*, 2018, pp. 663–679.
- [30] A. Mott, D. Zoran, M. Chrzanowski, D. Wierstra, and D. Jimenez Rezende, "Towards interpretable reinforcement learning using attention augmented agents," *Advances in Neural Information Processing Systems*, vol. 32, pp. 12350–12359, 2019.
- [31] S. Xingjian, Z. Chen, H. Wang, D.-Y. Yeung, W.-K. Wong, and W.-c. Woo, "Convolutional LSTM Network: A Machine Learning Approach for Precipitation Nowcasting," *Proceedings of neural information processing systems (NeurIPS)*, pp. 802–810, 2015.
- [32] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, "Openai gym," arXiv preprint, arXiv:1606.01540, 2016.