

Using Part-based Representations for Explainable Deep Reinforcement Learning

Manos Kirtas^[0000-0002-8670-0248], Konstantinos Tsampazis, Loukia Avramelou, Nikolaos Passalis^[0000-0003-1177-9139], and Anastasios Tefas^[0000-0003-1288-3667]

Computational Intelligence and Deep Learning Research Group
School of Informatics, Aristotle University of Thessaloniki, Greece.
{eakirtas, tsampaka, avramell, passalis, tefas}@csd.auth.gr

Abstract. Utilizing deep learning models to learn part-based representations holds significant potential for interpretable-by-design approaches, as these models incorporate latent causes obtained from feature representations through simple addition. However, training a part-based learning model presents challenges, particularly in enforcing non-negative constraints on the model’s parameters, which can result in training difficulties such as instability and convergence issues. Moreover, applying such approaches in Deep Reinforcement Learning (RL) is even more demanding due to the inherent instabilities that impact many optimization methods. In this paper, we propose a non-negative training approach for actor models in RL, enabling the extraction of part-based representations that enhance interpretability while adhering to non-negative constraints. To this end, we employ a non-negative initialization technique, as well as a modified sign-preserving training method, which can ensure better gradient flow compared to existing approaches. We demonstrate the effectiveness of the proposed approach using the well-known Cartpole benchmark.

Keywords: Part-based Learning · Explainable Reinforcement Learning · Non-negative Constraints · Proximal Policy Optimization

1 Introduction

Deep Reinforcement Learning (RL) has achieved state-of-the-art performance in various applications, including robotics [1, 2]. However, the use of RL agents in critical environments, where safety is highly prioritized, is hindered due to the limited transparency of the models. Extracting the rationale of a deep learning (DL) model in a human-interpretable way remains a challenging task, but doing so would be highly useful for improving both the performance and trustworthiness of the model, as well as preventing failures [3]. To this end, *post-hoc* explanation methods have been extensively studied over the years, providing rationales for the predictions of the model [4, 5]. However, such approaches cannot always provide a reliable explanation [6, 7], with *pre-hoc* methods for explainable AI gaining increasing attention recently [8]. The *pre-hoc* approaches aim

to design inherently explainable models, providing a transparent mechanism to the decision-making process in such a way that one can calibrate user trust and predict the system’s capabilities.

To this end, extracting a part-based representation of deep learning models provides great potential for interpretable-by-design approaches, since they are based on the simple addition of latent causes acquired from feature representations, making models easily interpretable by human actors due to the elimination of canceling neurons [9, 10]. However, training a part-based learning model is challenging since it requires non-negative constraints to the model’s parameters, leading to training difficulties, such as instabilities and convergence issues [11]. Furthermore, existing approaches for part-based learning are limited, e.g., applied solely on autoencoders [10, 12], and models that are not usually used in DL models, such as Pyramid Neural Networks [13, 14], resulting in a significant performance degradation [11], making them unsuitable for RL.

In this work, we propose a non-negative training approach for actor models in RL approaches, allowing for extracting part-based representations that can provide increased interpretability, while also building upon non-negative constraints that are known to be conceptually tied to human cognition [15, 16]. To this end, the proposed method employs a non-negative initialization method, along with an appropriately modified sign-preserving training method. More specifically, we propose using an exponential distribution-based non-negative initialization method for the actor model. Then, we introduce a sign-preserving alternative of Stochastic Gradient Ascent (SGA) that is used to train the actor model in a non-negative manner. The proposed optimization method allows better gradient flow, compared to existing clipping-based approaches, reducing the phenomenon of vanishing gradients and increasing the stability of the training process. As a result, the proposed pipeline enables more efficient training of inherently explainable models that are based on the non-negative part-based representation of the actor. Note that even though the proposed method is presented within the Proximal Policy Optimization (PPO) [17] algorithm, this is without loss of generality and could be readily adapted to any other Deep RL approach. We demonstrate the effectiveness of the proposed method in a traditionally used benchmark, named Cartpole, in a high-fidelity 3D robotics simulation.

The remainder of this paper is structured as follows. The proposed method is introduced and described in detail in Section 2, while the experimental evaluation is provided in Section 3. Finally, conclusions are drawn in Section 4.

2 Proposed Method

In this work, we focus on training non-negative agents using policy gradient-based approaches, such as the PPO algorithm [17], but without loss of generality, since the proposed method can also be directly applied to other RL methods as well, such as Q-learning based approaches. More specifically, PPO utilizes actor-critic networks, where the actor model decides which action should be taken, with its parameters denoted as θ . On the other hand, the critic network, equipped

with parameters $\tilde{\theta}$, informs the actor about the quality of its actions and guides the actor on how to adjust them during training. The PPO method trains the actor based on the policy gradient approach, while the critic evaluates the actions by computing the corresponding state/action values. For simplicity, we assume that both models have the same number of layers.

First, we propose an initialization scheme for the parameters of the actor model using an exponential distribution to ensure a positive-only initialization. Then, to train the actor model in a part-based representation manner, we propose a non-negative optimization approach based on Stochastic Gradient Ascent (SGA), ensuring that the canceling neurons of the network will be diminished by constraining parameters to the non-negative space, making them more easily interpretable by humans. Thus, only the actor network is trained in a non-negative manner, since it is responsible for the actions made by the agent during the deployment and, as a result, it is the one that needs to be explainable during deployment. Constraining to positive values only the parameters of the actor model, without similarly restricting the critic model, allows for reducing the risk of convergence issues that usually arise in non-negative neural networks [11]. At the same time, this does not reduce the interpretability of the actor model during deployment, since the critic is only used during the training process.

In the case of RL approaches with a policy gradient, the actor model is trained to learn a policy, $\pi(a|s)$, by observing the state s and returning the probability of selecting the action a . The groundbreaking results of the PPO are attributed to the constraints utilized in the policy parameter steps. More precisely, PPO employs the action probability ratio between the policy parameterization formulated as:

$$r_t(\theta) = \frac{\pi_{\theta}(a|s_t)}{\pi_{\theta_{old}}(a|s_t)} \in \mathbb{R}, \quad (1)$$

where $\pi_{\theta}(a|s_t)$ is the probability that policy π , with actor’s parameters θ , selecting an action a when the agent observes environment state s_t at time step t and the previous step parameters are denoted as θ_{old} . Using the clipped version of the ratio $r_t(\theta)$ around the value of 1 within ϵ , the policy exploration can be constrained to the close vicinity of the parameter space. The clipped policy ratio is defined as:

$$r_t^{clip}(\theta) = \text{clip}(r_t(\theta), 1 - \epsilon, 1 + \epsilon) \in [1 - \epsilon, 1 + \epsilon], \quad (2)$$

where ϵ is the constraint range of policy update and by default is set to $\epsilon = 0.2$, while the clip function is defined as:

$$\text{clip}(x, m, M) = \max(\min(x, M), m) \in [m, M]. \quad (3)$$

The final objective function of the PPO is defined as:

$$L^{actor}(s_t; \theta, \tilde{\theta}) = \mathbb{E}_t \left[\min \left(r_t^{clip}(\theta) A_t(\tilde{\theta}), r_t^{clip}(\theta) A_t(\tilde{\theta}) \right) \right] \in \mathbb{R}, \quad (4)$$

where $A_t(\tilde{\theta})$ is the advantage. In this work, we use the General Advantage Estimation (GAE) [18] approach. The Temporal Difference (TD) residual for each

time step t is calculated as:

$$\delta_t(\tilde{\theta}) = R_t + \gamma V_{\tilde{\theta}_t}^\pi(\mathbf{s}_{t+1}) - V_{\tilde{\theta}_t}^\pi(\mathbf{s}_t) \in \mathbb{R}, \quad (5)$$

where R_t is the reward the agent receives at time step t , $V_{\tilde{\theta}_t}^\pi(\mathbf{s}_t)$ is the value estimation predicted by the critic policy π for current state \mathbf{s}_t based on critic parameter $\tilde{\theta}_t$, γ is the discount factor and λ is the smoothing parameter. In this work, we use $\gamma = 0.99$ and $\lambda = 0.95$. Then, the advantage A_t is defined as:

$$A_t(\tilde{\theta}) = \sum_{i=0}^{n-t} \gamma^i \lambda^i \delta_{t+i}(\tilde{\theta}) \in \mathbb{R}, \quad (6)$$

where n is the total number of steps within an episode and t is the time step.

On the other hand, the critic network is typically trained to minimize the temporal difference between the returns and it is formulated as:

$$L^{critic} = \mathbb{E}_t[\delta_t(\tilde{\theta})^2] \in \mathbb{R} \quad (7)$$

Traditionally used initialization schemes, such as Kaiming [19] and Xavier [20], oriented to ANNs that apply the ReLU activation function, initialize the parameters of the k -th layer around zero, drawing values from a Gaussian distribution, $\theta \sim \mathcal{N}(0, \sigma_k)$, where θ denotes a parameter from θ and the standard deviation depends on the number of inputs and neurons. To this end, even by applying optimization methods that constrain parameters to positive values, it results in a very low variance of parameters during the first epochs of training that can lead to convergence difficulties or even halt the training process. To this end, inspired by [12], we propose to initialize actor parameters θ in the positive domain using an exponential distribution given by:

$$\theta \sim \text{Exp}(\lambda) = \frac{\ln(U(0, 1))}{\lambda} \in \mathbb{R}_+, \quad (8)$$

where $U(0, 1)$ is a uniform distribution between $(0, 1)$, \mathbb{R}_+ denotes the set of positive real values, and λ is the rate parameter of the distribution and it is a hyperparameter that by default is set to $\lambda = 100$. Even though initializing the actor parameters allows us to obtain a part-based representation before training, the traditionally used optimization algorithms, such as SGA, do not allow one to preserve the initial sign of the parameters. Therefore, we propose a sign-preserving optimization method that is based on the SGA. More specifically, we propose a sign-preserving alternative of SGA that modifies the update term, which is attributed to the sign change, ensuring that the trainable parameters will remain non-negative during the training phase. To this end, the actor's parameters are updated as:

$$\theta = \left| \theta_{old} + \eta_a \frac{\partial L^{actor}}{\partial \theta_{old}} \right|, \quad (9)$$

where the η_a denotes the learning rate of the actor and $|\cdot|$ the absolute value operator.

On the other hand, critic model is trained as usual employing the standard Stochastic Gradient Descent (SGD) denoted as:

$$\tilde{\theta} = \tilde{\theta}_{old} - \eta_c \frac{\partial L^{critic}}{\partial \tilde{\theta}_{old}}, \quad (10)$$

where the η_c defines the learning rate of the critic model and L^{critic} denotes the loss function of the critic. We call the proposed sign preservation method *Absolute Stochastic Gradient Ascent* (ASGA), since it employs the absolute value operator to preserve the positive sign of parameters during the training process.

The proposed training method is presented algorithmically in Algorithm 1. More specifically, it receives as input the initialization hyperparameters for the actor model, which is initialized with the exponential distribution using $\lambda \in \mathbb{R}$, and for the critic model, which is initialized by a normal distribution using $\sigma \in \mathbb{R}^n$, where n is the number of layers. The former network is optimized with the proposed sign-preserving alternative of SGA using learning rate η_a and the latter using the SGD optimizer with learning rate η_c . The algorithm outputs both actor’s θ and critic’s $\tilde{\theta}$ parameters, ensuring that the actor’s parameters will be non-negative. Firstly, the proposed method initializes the parameters for both models (lines 2-3), using an exponential distribution for the actor (line 3), constraining to positive-only parameters, and a normal distribution for critic parameters (line 4). In turn, the method iterates over the environment for $T_{episodes}$ episodes (lines 6-12), applying the obtained policy ($\pi_{\theta_{old}}$) for T_{steps} steps (lines 6-8), collecting the states, trajectories and rewards (line 7). Then, the PPO algorithm calculates the GAE (line 8) and the proposed method trains both actor and critic for T_{epochs} (line 9-12) using mini-batches (line 10-12). More precisely, the actor’s parameters are updated using the proposed sign-preserving gradient ascent (line 11), while the actor is trained as usual by applying the SGD optimizer (line 12). This results in non-negative trained parameters for the actor model, making the part-based representation of the actor feasible.

3 Experimental Evaluation

We experimentally evaluated the proposed method on the typical RL benchmark Cartpole implemented using the Deepbots framework [21, 22]. More specifically, the simulated environment is composed of a four-wheeled cart that has a long pole attached to it by a free hinge. On the top of the pole, there is a sensor to measure their vertical angle. The pole acts as an inverted pole pendulum and the goal is to keep it vertical by moving the cart forward and backward.

We applied the PPO algorithm to all evaluated cases, with the actor network getting the observations of the agent as input, consisting of two hidden layers of 10 neurons each, and outputs the action of the agent. Similarly, the critic network gets the observations as input, and it outputs the advantage of each state. The critic network also consists of two hidden layers of 10 neurons. On both networks, we employed the ReLU activation function in the hidden layers, with the actor

Algorithm 1: Non-Negative Actor PPO Training

Input : $\lambda \in \mathbb{R}$: the rate parameter of the exponential distribution,
 $\sigma \in \mathbb{R}^n$: a vector containing the standard deviation for the Gaussian distribution for each layer,
 $\eta_a \in \mathbb{R}$ and $\eta_c \in \mathbb{R}$: learning rates for actor and critic.

Output: θ : actor’s parameters, and
 $\tilde{\theta}$: critic’s parameters.

```

1 begin
2   for  $k = 1 \dots n$  do
3      $\theta^{(k)} \sim \text{Exp}(\lambda)$  ; // Initialize Actor’s Parameters
4      $\tilde{\theta}^{(k)} \sim \mathcal{N}(0, \sigma_k)$  ; // Initialize Critic’s Parameters
5   for  $i = 1 \dots T_{episodes}$  do
6     for  $j = 1 \dots T_{iter}$  do
7       Apply policy  $\pi_{\theta_{old}}$  and collect state, actions and rewards;
8        $A_j(\tilde{\theta}) = \sum_{i=0}^{n-t} \gamma_i \lambda_i \delta_{t+i}(\tilde{\theta})$  ; // Compute Advantage estimates
9     for  $j = 1 \dots T_{epochs}$  do
10    for every batch do
11       $\theta = \left| \theta + \eta_a \frac{\partial L^{actor}}{\partial \theta} \right|$  ; // Update Actor’s Parameters
12       $\tilde{\theta} = \tilde{\theta} - \eta_c \frac{\partial L^{critic}}{\partial \tilde{\theta}}$  ; // Update Critic’s Parameters
13    return  $\theta, \tilde{\theta}$ ;
```

model setting an upper bound to ReLU at 2 (ReLU2) for better visualization of the model. The observations contain the cart position, velocity on the x-axis, vertical angle of the pole, and pole’s velocity at its tip. The available discrete actions at each step are either to move forward or backward. For each step, the agent is rewarded with +1 and each episode ends after $T_{steps} = 195$ steps or earlier if the pole has fallen $\pm 15^\circ$ off vertical or if the cart has moved more than ± 39 centimeters on the x-axis. The networks are optimized for $T_{episodes} = 10^4$ episodes with learning rates equal to $\eta_a = 0.1$ and $\eta_c = 0.003$ for actor and critic, respectively. The PPO iterates for $T_{iter} = 5$ over batches of 8 collected samples.

We evaluate the proposed method on the aforementioned setup against two baseline methods using: a) Clipping Stochastic Gradient Ascent (CSGA) with Kaiming initialization and b) CSGA with Xavier initialization. These are based on the clipping approach proposed in [12, 11], after appropriate adaptation for use in PPO. More specifically, the CSGA applies a clipping function to the typical stochastic gradient ascent and it is formulated as:

$$\theta' = \max \left(0, \theta - \eta \frac{\partial J}{\partial \theta} \right), \quad (11)$$

where the $\max(\cdot)$ operator is applied element-wise. In all cases, the evaluated optimization methods are used on the actor model, ensuring the non-negativity




Fig. 1. On the left, the figure depicts the obtained reward during training that is smoothed using a moving average filter with a window of 100. On the right, the action probabilities for each method are depicted using the same moving average setting.

of its parameters, with the critic model being trained as usual with the SGD optimizer.

Table 1. Average and variance of rewards both for training and evaluation phase over 5 runs.

Method	Training	Evaluation
CSGA (Kaiming Init.)	62.83 ± 39.64	89 ± 98.59
CSGA (Xavier Init.)	53.67 ± 35.47	58.2 ± 78.4
Proposed	89.45 ± 1.04	140.4 ± 43.9

We conducted five evaluation runs for each method using different seeds during the training phase. In Table 1 the average and variance of the rewards over 5 runs are reported for all the evaluated methods. Both baselines lead to an unstable training process, resulting in high variance reward values at the end of the training. In contrast, the proposed method offers significantly more consistent training, resulting in low-variance rewards after training, as well as higher performance. This behavior is also highlighted in the evaluation performance, where the proposed method holds the pole for more than 50 steps in the average case contrary to the other evaluated baselines.

To highlight such different behavior, we report in Figure 1, the average reward between different runs and the average action probability during training. For each case, we smooth the reported result using a moving average filter, setting the window value equal to 100. Similarly to the results reported in Table 1, we observe that during training, the two baselines are highly unstable, resulting in a poor local minimum and, as a result, significantly lower rewards. On the other hand, the proposed method allows for more consistent training, achieving significantly higher performance than the baselines.

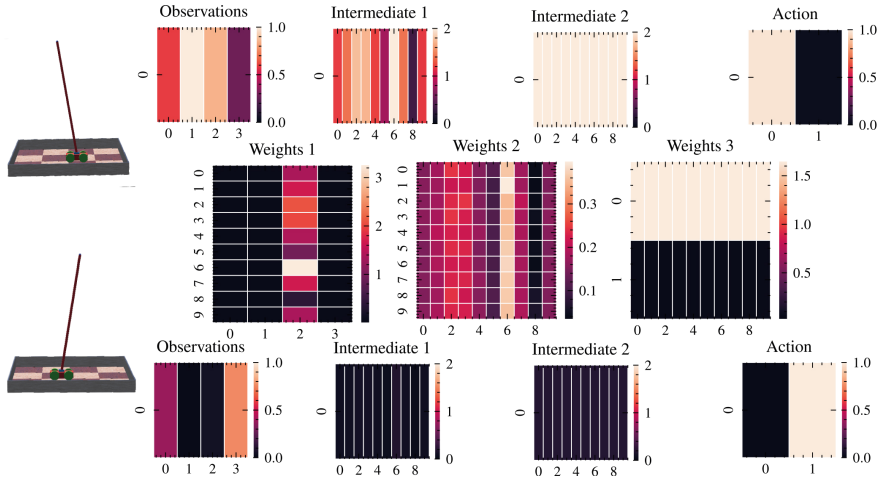


Fig. 2. Part-based representation of the actor model. At the top and bottom rows, the input and responses of each layer are depicted. In the central row, the weights of the actor model are depicted. The biases are omitted for simplicity.

This instability in training can also be observed in the action probability plot, where the baselines lead to increases at a high rate during the initial stage of training, before subsequently gradually decreasing the corresponding probabilities. This indicates that baselines converge quickly to a bad local policy, with the clipping update term introducing difficulties in the training process. Such difficulties can be attributed to the fact that the clipping method zeros out synapses when they try to change sign, reducing the learning capacity of the model. This can also lead to vanishing gradient phenomena, which in turn can lead to bad local minima or even halt the training process [23]. On the other hand, the proposed optimization method ensures that the parameters will remain non-negative without suppressing weights to zero, allowing gradients to flow through the network since the absolute value operator has a non-zero derivative both for positive and negative values. In this way, it provides a smooth training process and consistent results, as well as a more explainable representation for the actor model, allowing one to extract rationales of the agent due to the non-negative constraints that diminish canceling neurons and leading to part-based representations.

To highlight this, in Figure 2 we provide an example of the part-based representation acquired using the proposed method. More specifically, two cases are presented: a) when the pole is dropped from the front of the cart (top row) and b) when the pole is falling from the rear side of the cart (bottom row). On the middle row, the weights of the model are depicted, omitting biases for simplicity. The observations are presented on the first plot from the left and are normalized between $[0, 1]$, with columns on the plot denoting from 0 to 3: 0) cart position, 1) cart velocity, 2) pole angle and 3) endpoint velocity. The action made by the

actor is represented on the right plot with the zeroth neuron response denoting the forward move of the agent (left as depicted in the figure) and the first neuron the backward move (right as depicted in the figure).

As depicted in Figure 2, applying part-based learning by using the proposed method allows one to extract visually meaningful representations. More precisely, when the pole is dropped from the positive side of the x-axis, meaning on the front side of the cart, and has large values for the pole angle, the agent moves the cart toward the same direction as the falling pole trying to keep it vertical, as depicted in the first row of the figure with the lighter heatmaps. On the other hand, when the pole has lower values on the pole angle (darker colors in the observations’ heatmap are used to denote this), the agent moves the cart backward. Lighter heatmaps when multiplied with the final layer’s weights (weights 3 in the figure) of the actor lead to firing the zeroth neuron of the layer that is translated on moving the cart forward. We can safely conclude that the agent has high confidence in its decision by comparing the intermediate representation of two cases after the second layer, since the two representations are significantly different. As depicted, in the first case (first row of the figure), the last layer zeroes out the response of the first neuron, maintaining the large values on the input of the zeroth neuron. On the contrary, in the second case, the input of the last layer already has values close to zero, with the first layer increasing the values of the input to fire the zeroth neuron. We have to mention that the action heatmap of the figure is after the softmax activation function, translating the response of the last layer to action probabilities.

We can also extract the rationale for the agent based on other observations. For example, the second observation (meaning the column denoted by 1 in the figure) refers to the cart velocity (normalized between range $[0, 1]$). This means that if the value of cart velocity is lower than 0.5, then the cart has a direction to the negative side of the x-axis (backward movement). Accordingly, if the cart velocity value is greater than 0.5, then the cart is moving forward, towards the positive side of the x-axis. As depicted in the observations’ heatmap, the agent moves toward the opposite direction of the direction that the pole is falling. This is expected since the observations are extracted from an irrational agent, where the actor parameters have been randomly initialized. However, the trained agent using the proposed method outputs actions that moves the cart in the same direction as the falling pole to keep the pole vertical. This can also be observed in the heatmap of the weights of the first layer, where the third column has an inverse color with respect to the second column. Indeed, there is an inverse correlation between the cart position and the pole’s angle.

We extend our analysis to obtain further insight regarding the response of the actor model, leveraging the advantages provided by the part-based representation of the model. To this end, we optimized the observation vector keeping the trained weights frozen in order to minimize the distance between the output of the model and a given action. The inputs are optimized for 5 epochs using SGD optimizer with the proposed sign-preserving update function applying mean squared error loss. We applied the proposed optimization method to ensure




Fig. 3. Optimized input of actor model to maximize the action probability of a given action using three different initialization. In the first column, the observations are optimized to maximize the forward action probability. In the second column, the observations are optimized to maximize the backward action probability. A different initialization of the observation vector is used for each row.

that the values of the observations will remain non-negative according to the actual normalized observations acquired from the environment. The observations obtained are presented in Figure 3, where in the left column they are optimized to maximize the action probability of the forward action, and, respectively, on the right side they are optimized to maximize the action probability for the backward action. For each case, we report both the initial vector of observation and the one after the optimization process.

To evaluate which of the observations is more significant for each action, we initialize each observation to 0.5 and then optimize them to maximize the respective action probability, depicted in the (a) row. As expected, for both actions, the observation that changed significantly is the pole angle (in the third column).

Indeed, when the observations are optimized in order to obtain a forward action from the actor model, the pole angle value converges to its maximum, which is value one. On the other hand, when the observations are optimized in order to obtain a backward action, the pole angle value converges to a value close to zero. In both cases, the other observations remain close to the initialized values, and it seems that they are not significantly affected by the optimization process. This is an expected behavior since the pole angle is the most significant indicator, revealing the direction in which the pole is falling, and as a result the one that significantly contributes to the prediction of the actor model.

We repeat the aforementioned experimental setup by applying different initializations on the observations vector and reporting the results at the (b) and (c) rows. More specifically, we randomly draw values from a Gaussian distribution with 0.5 mean and 0.2 standard deviation. As shown, except for the pole angle observation, which is significantly changed in most cases, the rest of the observations are not affected by the optimization process. Regarding the observation of the pole angle, when it has opposite direction with the given action, the obtained direction of the pole angle after optimization is inverted, such in case (b.2). In cases where the initialized observation vector has the same direction of the pole angle with the given action, as in cases of (b.1), the optimization process slightly changes the initial vector and the value of the pole angle. Finally, it is observed that when the initialized pole is vertical, such as in (a) and (c) cases, the optimization process leads to maximizing or minimizing the pole angle value according to the given action, resulting in the same direction with the action.

4 Conclusions

In this study, we have introduced a novel training approach that focuses on non-negativity in deep RL using PPO. The proposed approach enables the extraction of part-based representations, which offers enhanced interpretability while following non-negative constraints associated with human cognition. To achieve this objective, the proposed method employs a non-negative initialization technique, followed by a modified sign-preserving training method. More specifically, we proposed employing an exponential distribution-based non-negative initialization method for the actor model and then using an appropriately modified sign-preserving alternative to Stochastic Gradient Ascent (SGA) for training the actor model in a non-negative manner. By adopting the proposed method, we can mitigate issues related to the reduction of the learning capacity of models and the vanishing gradients due to the use of clipping mechanisms involved in existing approaches. This helps mitigate issues such as vanishing gradients and enhances training stability. Consequently, the proposed pipeline enables more efficient training of inherently explainable models based on the non-negative part-based representation of the actor. To validate the effectiveness of the proposed method, we conducted experiments on the well-established Cartpole benchmark. The results demonstrate the effectiveness of the proposed method in achieving su-

perior performance and showcasing the advantages of the proposed non-negative training methodology.

The promising results reported in this paper highlight several interesting future research directions. First, the proposed method can also be extended to handle value-based RL approaches, such as DQN [24]. Furthermore, extending the part-based representation learning to the actor model could also provide further insight into the training dynamics of the RL process, as well as allow for better adapting it to the task at hand, potentially leading to more robust algorithms. Finally, combining the proposed method with distillation approaches that can transfer knowledge from the intermediate layers of traditional DL models, for example [25], could potentially allow for better guidance of the optimization process and learning more accurate policies.

Acknowledgments

This work was supported by the European Union’s Horizon 2020 Research and Innovation Program (OpenDR) under Grant 871449. This publication reflects the authors’ views only. The European Commission is not responsible for any use that may be made of the information it contains.

References

1. 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. B. R. Kiran, I. Sobh, V. Talpaert, P. Mannion, A. A. A. Sallab, S. Yogamani, and P. Pérez, “Deep reinforcement learning for autonomous driving: A survey,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 6, pp. 4909–4926, 2022.
3. B. Hayes and J. A. Shah, “Improving robot controller transparency through autonomous policy explanation,” in *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction, HRI ’17*, (New York, NY, USA), p. 303–312, Association for Computing Machinery, 2017.
4. M. T. Keane and E. M. Kenny, “How case-based reasoning explains neural networks: A theoretical analysis of xai using post-hoc explanation-by-example from a survey of ann-cbr twin-systems,” in *Case-Based Reasoning Research and Development* (K. Bach and C. Marling, eds.), (Cham), pp. 155–171, Springer International Publishing, 2019.
5. S. Bach, A. Binder, G. Montavon, F. Klauschen, K.-R. Müller, and W. Samek, “On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation,” *PloS one*, vol. 10, no. 7, p. e0130140, 2015.
6. D. Slack, S. Hilgard, E. Jia, S. Singh, and H. Lakkaraju, “Fooling lime and shap: Adversarial attacks on post hoc explanation methods,” in *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society, AIES ’20*, (New York, NY, USA), p. 180–186, Association for Computing Machinery, 2020.
7. Y. Zhou, S. Booth, M. T. Ribeiro, and J. Shah, “Do feature attribution methods correctly attribute features?,” *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, pp. 9623–9633, Jun. 2022.

8. C. Rudin, C. Chen, Z. Chen, H. Huang, L. Semenova, and C. Zhong, “Interpretable machine learning: Fundamental principles and 10 grand challenges,” *Statistics Surveys*, vol. 16, no. none, pp. 1 – 85, 2022.
9. A. Lemme, R. F. Reinhart, and J. J. Steil, “Online learning and generalization of parts-based image representations by non-negative sparse autoencoders,” *Neural Networks*, vol. 33, pp. 194–203, 2012.
10. E. Hosseini-Asl, J. M. Zurada, and O. Nasraoui, “Deep learning of part-based representation of data using sparse autoencoders with nonnegativity constraints,” *IEEE Transactions on Neural Networks and Learning Systems*, vol. 27, no. 12, pp. 2486–2498, 2016.
11. J. Chorowski and J. M. Zurada, “Learning understandable neural networks with nonnegative weight constraints,” *IEEE Transactions on Neural Networks and Learning Systems*, vol. 26, no. 1, pp. 62–69, 2015.
12. B. O. Ayinde and J. M. Zurada, “Deep learning of constrained autoencoders for enhanced understanding of data,” *IEEE Transactions on Neural Networks and Learning Systems*, vol. 29, no. 9, pp. 3969–3979, 2018.
13. M. Ferro, B. Fernandes, and C. Bastos-Filho, “Non-negative structured pyramidal neural network for pattern recognition,” in *Proceedings of the International Joint Conference on Neural Networks*, pp. 1–7, 2018.
14. M. S. A. Ferro, B. J. T. Fernandes, and C. J. A. Bastos-Filho, “Non-negative pyramidal neural network for parts-based learning,” in *Proceedings of the International Joint Conference on Neural Networks*, pp. 1709–1716, 2017.
15. D. D. Lee and H. S. Seung, “Learning the parts of objects by non-negative matrix factorization,” *Nature*, vol. 401, no. 6755, pp. 788–791, 1999.
16. K. Tanaka, “Columns for Complex Visual Object Features in the Inferotemporal Cortex: Clustering of Cells with Similar but Slightly Different Stimulus Selectivities,” *Cerebral Cortex*, vol. 13, pp. 90–99, 01 2003.
17. J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, “Proximal policy optimization algorithms,” *arXiv preprint - 1707.06347*, 2017.
18. J. Schulman, S. Levine, P. Moritz, M. I. Jordan, and P. Abbeel, “Trust region policy optimization,” *arXiv preprint - 1502.05477*, 2015.
19. K. He, X. Zhang, S. Ren, and J. Sun, “Delving deep into rectifiers: Surpassing human-level performance on imagenet classification,” in *Proceedings of the IEEE International Conference on Computer Vision*, December 2015.
20. X. Glorot and Y. Bengio, “Understanding the difficulty of training deep feedforward neural networks,” in *Proceedings of the International Conference on Artificial Intelligence and Statistics* (Y. W. Teh and M. Titterton, eds.), vol. 9 of *Proceedings of Machine Learning Research*, (Chia Laguna Resort, Sardinia, Italy), pp. 249–256, PMLR, 13–15 May 2010.
21. M. Kirtas, K. Tsampazis, N. Passalis, and A. Tefas, “Deepbots: A webots-based deep reinforcement learning framework for robotics,” in *Proceedings of the International Conference on Artificial Intelligence Applications and Innovations*, pp. 64–75, 2020.
22. M. Kirtas, K. Tsampazis, P. Tosidis, N. Passalis, and A. Tefas, “Chapter 21 - deep learning for robotics examples using opendr,” in *Deep Learning for Robot Perception and Cognition* (A. Iosifidis and A. Tefas, eds.), pp. 579–596, Academic Press, 2022.
23. M. Kirtas, N. Passalis, G. Mourgiyas-Alexandris, G. Dabos, N. Pleros, and A. Tefas, “Robust architecture-agnostic and noise resilient training of photonic deep learning models,” *IEEE Transactions on Emerging Topics in Computational Intelligence*, 2022.

24. V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller, "Playing atari with deep reinforcement learning," *arXiv preprint arXiv:1312.5602*, 2013.
25. N. Passalis and A. Tefas, "Learning deep representations with probabilistic knowledge transfer," in *Proceedings of the European Conference on Computer Vision*, pp. 268–284, 2018.