




## Article

# Deep Reinforcement Learning with Interactive Feedback in a Human–Robot Environment

Ithan Moreira <sup>1,†</sup>, Javier Rivas <sup>1,†</sup>, Francisco Cruz <sup>1,2,\*</sup> , Richard Dazeley <sup>2</sup> , Angel Ayala <sup>3</sup>  
and Bruno Fernandes <sup>3</sup> 

<sup>1</sup> Escuela de Ingeniería, Universidad Central de Chile, Santiago 8330601, Chile; ithan.moreira@alumnos.ucentral.cl (I.M.); javier.rivas@alumnos.ucentral.cl (J.R.)

<sup>2</sup> School of Information Technology, Deakin University, Geelong 3220, Australia; richard.dazeley@deakin.edu.au

<sup>3</sup> Escola Politécnica de Pernambuco, Universidade de Pernambuco, Recife 50720-001, Brasil; aaam@comp.poli.br (A.A.); bjtf@comp.poli.br (B.F.)

\* Correspondence: francisco.cruz@deakin.edu.au

† These authors contributed equally to this work.

Received: 7 July 2020; Accepted: 10 August 2020; Published: 12 August 2020



**Abstract:** Robots are extending their presence in domestic environments every day, it being more common to see them carrying out tasks in home scenarios. In the future, robots are expected to increasingly perform more complex tasks and, therefore, be able to acquire experience from different sources as quickly as possible. A plausible approach to address this issue is interactive feedback, where a trainer advises a learner on which actions should be taken from specific states to speed up the learning process. Moreover, deep reinforcement learning has been recently widely used in robotics to learn the environment and acquire new skills autonomously. However, an open issue when using deep reinforcement learning is the excessive time needed to learn a task from raw input images. In this work, we propose a deep reinforcement learning approach with interactive feedback to learn a domestic task in a Human–Robot scenario. We compare three different learning methods using a simulated robotic arm for the task of organizing different objects; the proposed methods are (i) deep reinforcement learning (DeepRL); (ii) interactive deep reinforcement learning using a previously trained artificial agent as an advisor (agent-IDeepRL); and (iii) interactive deep reinforcement learning using a human advisor (human-IDeepRL). We demonstrate that interactive approaches provide advantages for the learning process. The obtained results show that a learner agent, using either agent-IDeepRL or human-IDeepRL, completes the given task earlier and has fewer mistakes compared to the autonomous DeepRL approach.

**Keywords:** robotics; interactive deep reinforcement learning; deep reinforcement learning; domestic scenario

## 1. Introduction

Robotics has been getting more attention since new researched advances have introduced significant improvements to our society. For instance, for many years, robots have been installed in the automotive industrial area [1]. However, the current technological progress has allowed expanding the robot's applications domain in areas such as medicine, military, search and rescue, and entertainment. In this regard, under current research, another challenging application of robotics is its integration to domestic environments, mainly due to the presence of many dynamic variables in comparison to industrial contexts [2]. Moreover, in domestic environments, it is expected that humans regularly interact with robots and that the robots can understand and respond accordingly to the interactions [3,4].

Algorithms such as reinforcement learning (RL) [5] allow a robotic agent to autonomously learn new skills, in order to solve complex tasks inspired by the way humans do, through trial and error [6]. RL agents interact with the environment in order to find an appropriate policy that meets the problem aims. To find the appropriate policy, the agent interacts with the environment by performing an action  $a_t$  and, in turn, the environment returns a new state  $s_{t+1}$  with a reward  $r_{t+1}$  for the performed action to adjust the policy. However, an open issue in RL algorithms is the time and the resources required to achieve good learning outcomes [7,8], which is especially critical in online environments [9,10]. One of the reasons for this problem is that the agent, at the beginning of the learning process, does not know the environment and the interaction responses. Thus, to address this problem, the agent must explore multiple paths to refine its knowledge about the environment.

In continuous spaces, an alternative is to recognize the agent's state directly from raw inputs. Deep reinforcement learning (DeepRL) [11] is based on the same RL structure but also adds deep learning to process the function approximation for the state in multiple abstraction levels. An example of DeepRL implementations is by convolutional neural networks (CNN) [12] which can be modified to be used for DeepRL, e.g., DQN [13]. CNNs have brought significant progress in the last few years in different areas such as image, video, audio, and speech processing, among others [14]. Nevertheless, for a robotic agent working in highly dynamic environments, DeepRL still needs excessive time to learn a new task properly.

## 2. Related Works

In this section, we review previously developed works considering two main areas. First, we address the deep reinforcement learning approach and the use of interactive feedback. Following, we discuss the problem of vision-based object sorting using robots in order to contextualize our approach properly. Finally, we describe the scientific contributions of this work.

### 2.1. Deep Reinforcement Learning and Interactive Feedback

Deep reinforcement learning (DeepRL) combines reinforcement Learning (RL) and deep neural networks (DNN). This combination has allowed the enhancement of RL agents when autonomously exploring a given scenario [15]. If an RL agent is learning a task, the environment gives the agent the necessary information on how good or bad the taken actions are. With this information, the agent must differentiate which actions lead to a better accomplishment of the task aims [5].

Aims may be expressed by a reward function that assigns a numerical value to each performed action from a given state. Additionally, after performing an action, the agent reaches a new state. Therefore, the agent associates states with actions to maximize  $r_0 + \gamma \cdot r_1 + \gamma^2 \cdot r_2 + \dots$ , where  $r_i$  is the obtained reward in the  $i$ -th episode and  $\gamma$  is the discount factor parameter that indicates how influential future actions are.

The fundamental base of RL tasks is the Markov decision process (MDP), in which future transitions and rewards are affected only by the current state and the selected action [16]. Therefore, if the Markovian property is present for a state, this state contains all the information needed about the dynamics of a task. For instance, chess is a classic example of the Markov property. In this game, it does not matter the history of plays in order to make a decision about the next movement. All the information is already described in the current distribution of pieces over the board. In other words, if the current state is known, the previous transitions that led the agent to that situation become irrelevant in terms of the decision-making problem.

Formally, an MDP is defined by a 4-tuple  $\langle S, A, \delta, r \rangle$  where:

- $S$  is a finite set of system states,  $s \in S$ ;
- $A$  is a finite set of actions,  $a \in A$ , and  $A_{s_t} \subseteq A$  is a finite set of actions available in  $s_t \in S$  at time  $t$ ;
- $\delta$  is the transition function  $\delta : S \times A \rightarrow S$ ;
- $r$  is the immediate reward (or reinforcement) function  $r : S \times A \rightarrow \mathbb{R}$ .

Each time-step  $t$ , the agent perceives the current state  $s_t \in S$  and chooses an action  $a_t \in A$  to be carried out. The environment gives a reward  $r_t = r(s_t, a_t)$  and the agent moves to the state  $s_{t+1} = \delta(s_t, a_t)$ . The functions  $r$  and  $\delta$  depend only on the current state  $s_t$  and action  $a_t$ , therefore, it is a no memory process. Over time, the agent tries to learn a policy  $\pi : S \rightarrow A$  which, from a state  $s_t$ , yields the greatest value or discounted reward [5], as shown in Equation (1):

$$Q^\pi(s_t, a_t) = r_t + \gamma \cdot r_{t+1} + \gamma^2 \cdot r_{t+2} + \dots = \sum_{i=0}^{\infty} \gamma^i \cdot r_{t+i+1} \quad (1)$$

where  $Q^\pi(s_t, a_t)$  is the action-value function following the policy  $\pi$  (e.g., choosing action  $a_t$ ) from a given state  $s_t$ . The discount factor  $\gamma$  is a constant ( $0 \leq \gamma < 1$ ), which determines the relative importance of immediate and future rewards. For instance, in case  $\gamma = 0$ , then the agent is short-sighted and maximizes only the immediate rewards, or in case  $\gamma \rightarrow 1$  the agent is more foresighted in terms of future reward.

The final goal of RL is to find an optimal policy ( $\pi^*$ ) mapping states to actions in order to maximize the future reward ( $r$ ) over a time ( $t$ ) with a discount rate ( $\gamma \in [0, 1]$ ), as shown in Equation (2). In the equation,  $\mathbb{E}_\pi[\cdot]$  denotes the expected value given that the agent follows policy  $\pi$  and  $Q^*(s_t, a_t)$  denotes the optimal action-value function [5]. Table 1 summarizes all the elements shown within Equation (2). In DeepRL, an approximation function, implemented by DNN, allows an agent to work with high-dimensional observation spaces, such as pixels of an image [13]:


$$Q^*(s_t, a_t) = \max \left( \mathbb{E}_\pi[r_t + \gamma r_{t+1} + \gamma^2 r_{t+2} + \dots | s_t = s, a_t = a, \pi] \right) \quad (2)$$

**Table 1.** Elements to compute the optimal action-value function.

Symbol	Meaning
$Q^*(s, a)$	Optimal action-value function
$\mathbb{E}$	Expected value following policy $\pi$
$\pi$	Policy to map states to actions
$\gamma$	Discount factor
$r_t$	Reward received at time step $t$
$s_t$	Agent's state at time step $t$
$a_t$	Action taken at time step $t$

Interactive feedback is a method that improves the learning time of an RL agent [17]. In this method, an external trainer can guide the agent's apprenticeship to explore more promising areas at early learning stages. The external trainer is an agent that can be a human, a robot, or another artificial agent.

There are two principal strategies for providing interactive feedback in RL scenarios, i.e., evaluative and corrective feedback [18]. In the first one, called reward-shaping, the trainer modifies or accepts the reward given by the environment in order to bias the agent's learning [19,20]. In the second one, called policy-shaping, the trainer may suggest a different action to perform, by replacing the one proposed by the policy [21,22]. A simple policy-shaping method involves forcing the agent to take certain actions that are recommended by the trainer [23,24]. For instance, a similar approach is used when a teacher is guiding a child's hand to learn how to draw a geometric figure. In this work, we use the policy-shaping approach since it has been shown that humans using this technique to instruct an agent provide advice that is more accurate, are able to assist the learner agent for a longer time, and provide more advice per episode. Moreover, people using policy-shaping have reported that the agent's ability to follow the advice is higher, and therefore, felt their own advice to be of higher accuracy when compared to people providing advice via reward-shaping [25]. The policy-shaping approach is depicted in Figure 1.



**Figure 1.** Policy-shaping interactive feedback approach. In this approach, the trainer may advise the agent on what actions to take in a particular given state.

There are different ways to support the agent's learning, which in turn may lead to other problems. For instance, if the trainer delivers too much advice, the learner never gets to know other alternatives because most of the decisions taken are given from the external trainer [26]. The quality of the given advice by the external trainer must also be considered to improve the learning. It has been shown that inconsistent advice may be very detrimental during the learning process, so that in case of low consistency of advice, autonomous learning may lead to better performance [27].

One additional strategy to better distribute the given advice is to use a budget [26]. In this strategy, the trainer has a limited amount of interaction with the learner agent, similar to the limited patience of a person for teaching. There are different ways of using the budget, in terms of when to interact or give advice, namely, early advising, alternating advice, importance advising, mistake correcting, and predictive advising. In this work, we use early advising allowing us to fairly compare interactive approaches using the different kinds of trainers used in the proposed methods, i.e., humans or artificial agents as trainers.

Although there have been some approaches addressing the interactive deep reinforcement learning problem, they have been mostly used in other scenarios. For instance, Ref. [28] presents an application to serious games, and Ref. [29] presents a dexterous robotic manipulation approach using demonstrations. In the game scenario [28], the addressed task presents different environmental dynamics compared to Human–Robot environments. Moreover, the authors propose undoing an action by the advisor, which is not always feasible. For example, in a Human–Robot environment, a robot might break an object as a result of a performed action, which is impossible to undo.

## 2.2. Vision-Based Object Sorting with Robots

The automation of sorting object tasks has been previously addressed using machine learning techniques. For instance, Lukka et al. [30] implemented a recycling robot for construction and demolition waste. In this work, the robot sorts the waste using a vision-based system for object recognition and object manipulation to control the movement of the robot in order to classify the objects presented on a moving belt properly. The authors did not present performance results since the approach is presented as a functional industrial prototype for sorting objects through images.

The object recognition problem is an extended research area that has been addressed by different techniques, including deep learning, as presented in [31]. This approach is a similar system to [30] in terms of proposing to sort garbage from a moving belt. The authors used a convolutional neural network, called Fast R-CNN, to obtain the moving object's class and location, and send the information to the robotic grasping control unit to grasp the object and move it. As the authors point out, the key problem the deep learning method tries to solve is the object identification. Moreover, another approach to improve the object recognition task is presented in [32], where the authors implemented a stereo

vision system to recognize the material and the clothes categories. The stereo vision system creates a 3D reconstruction of the image to process and obtains local and global features to predict the clothing category class and manipulate a robot that must grasp and sort the clothing in a preestablished box. These two systems, presented in [31,32], use a supervised learning method that requires prior training of items to be sorted, leading to low generalization for new objects.

Using RL to sort objects has also been previously addressed. For instance, in [33], a cleaning-table task in a simulated robotic scenario is presented. In order to complete the task, the robot needs to deal with objects such as a cup and a sponge. In this work, the RL agent used a discrete tabular RL approach complemented by interactive feedback and affordances [34]. Therefore, the agent did not deal with the problem of continuous visual inputs for state representation. Furthermore, in [35], an approach for robotic control using DeepRL is presented. In this work, a simulated Baxter robot learned autonomous control using a DQN-based system. When transferring the system to a real-world scenario, the approach failed. To fix this, the system ran replacing the camera images with synthetic images in order for the agent to acquire the state and decide which action to take in the real world.

### 2.3. Scientific Contribution

Although a robot may be capable of learning autonomously to sort objects in different contexts, current approaches address the problem using supervised deep learning methods with previously labeled data to recognize the objects, e.g., [31,32]. In this regard, using the DeepRL approach allows for classifying objects as well as deciding how to act with them. Additionally, if prior knowledge of the problem is transferred to the DeepRL agent, e.g., using demonstrations [36], the learning speed may also be improved. Therefore, using interactive feedback as an assistance method, we will be able to advise the learner agent during the learning process, using both artificial and human trainers, to evaluate how the DeepRL method responds to the given advice.

In this work, we present an interactive-shaping vision-based algorithm derived from DeepRL, referred to here as interactive DeepRL or IDeepRL. Our algorithm allows us to speed up the required learning time through strategic interaction with either a human or an artificial advisor. The information exchange between the learner and the advisor gives the learner a better understanding of the environment by reducing the search space.

We have implemented a simulated domestic scenario, in which a robot has to organize objects considering color and shape. The RL agent perceives the world through RGB images and interacts through a robotic arm while an external trainer may advise the agent a different action to perform during the first training steps. The implemented scenario is used for test and comparison between the DeepRL and IDeepRL algorithms, as well as the evaluation of IDeepRL using two different types of trainers, namely, another artificial agent previously trained and a human advisor.

Therefore, the contribution of the proposed method is to demonstrate that interactive-shaping advice can be efficiently integrated into vision-based deep learning algorithms. The interactive learning methodologies proposed in this work outperform current autonomous DeepRL approaches allowing for collecting more and faster reward using both artificial and human trainers.

## 3. Materials and Methods

### 3.1. Methodology and Implementation of the Agents

In this work, our focus is on assessing how interactive feedback, used as an assistance method, may affect the performance of a DeepRL agent. To this aim, we implement three different approaches for the RL agents:

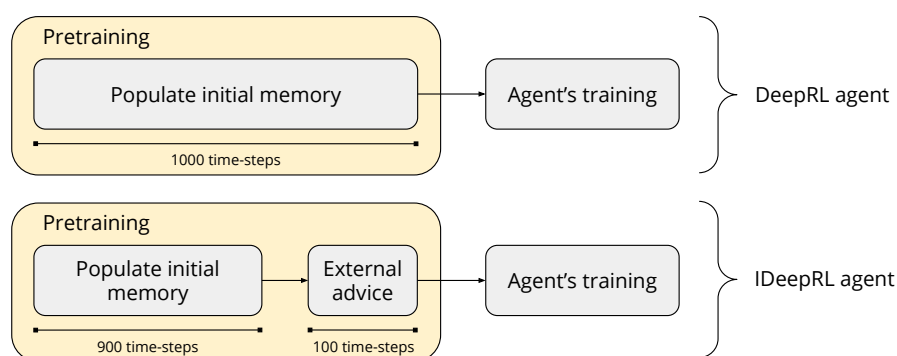
- i DeepRL: where the agent interacts autonomously with the environment;
- ii agent-IdEEPRL: where the DeepRL approach is complemented with a previously trained artificial agent to give advice; and
- iii human-IdEEPRL: where the DeepRL approach is complemented with a human trainer.

The first approach includes a standard deep reinforcement learning agent, referred to here as DeepRL, and is the basis of both of the interactive agents discussed subsequently. The DeepRL agent perceives the environment information through a visual representation [37], which is processed by a convolutional neural network (CNN) that estimates the Q-values. The deep Q-learning algorithm allows the agents to learn through actions previously experienced, using the CNN as a function approximator, allowing them to generalize states and apply Q-learning in continuous state spaces.

To save past experiences, the experience replay [38] technique is implemented. This technique saves the most useful information (experience) in memory, which is used afterward to train the RL agent. The neural network is responsible for processing the visual representation and gives the Q-value of each action to the agent, which decides what action to take. In order for the agent to balance exploration and exploitation of actions, the  $\epsilon$ -greedy method is used. This method includes an  $\epsilon$  parameter, which allows the agent to perform either a random exploratory action or the best known action proposed by the policy.

The learning process for the autonomous agent, i.e., DeepRL agent, is separated into two stages. The first pretraining stage consists of 1000 random actions that the agent must perform to populate the initial memory. In the second stage, the agent's training is carried out using the  $\epsilon$ -greedy policy and, after each performed action, the agent is trained using 128 tuples considering state, action, reward, and next state, as  $\langle s_t, a_t, r_t, s_{t+1} \rangle$  extracted from the memory.

Both IDeepRL approaches are based on autonomous DeepRL and include the interactive feedback strategy from an external trainer to improve the DeepRL performance. Therefore, the agents have the same base algorithm, adding an extra interactive stage. For the interactive agents, the learning process is separated into three stages. The first pretraining stage corresponds to 900 random actions that the agent must perform in order to populate the initial memory. In the second stage, the external trainer participates giving early advice about the environment dynamics and the intended task during 100 consecutive time-steps. In the third stage, the agent starts training using the  $\epsilon$ -greedy policy, and, following each action selected, the agent is trained with 128 tuples as  $\langle s_t, a_t, r_t, s_{t+1} \rangle$  saved previously in the batch memory. The learning process for both autonomous and interactive agents is depicted in Figure 2.



**Figure 2.** The learning process for autonomous and interactive agents. Both approaches include a pretraining stage comprising 1000 actions. For interactive agents, the final part of the pretraining is performed using external advice instead of random actions.

In the second stage of the IDeepRL approach, the learner agent receives advice either from a previously trained artificial agent or from a human trainer. The artificial trainer agent used in agent-IDEepRL is an RL agent that collected previous experience by performing the autonomous DeepRL approach using the same hyperparameters. Therefore, the knowledge is acquired by interacting with the environment in the same manner as the learner agent does. Once the trainer agent has learned the environment, it is used to then provide advice in agent-IDEepRL over 100 consecutive time-steps. Both the trainer agent and the learner agent perceived the environmental information through a visual representation after an action is performed.



Algorithm 1 shows the first stage for DeepRL and IDeepRL approaches, which corresponds to the pretraining stages to populate the batch memory. This algorithm also contains the second stage for interactive agents using IDeepRL represented with a conditional in line 4. Moreover, in Algorithm 2, the second stage for DeepRL is observed, which also corresponds to the third stage for IDeepRL, the training stage.

---

**Algorithm 1** Pretraining algorithm to populate the batch memory including interactive feedback.

---

```

1: Initialize memory M
2: Observe agent's initial state  $s_0$ 
3: while len(M)  $\leq$  1000 do
4:   if interaction is used AND length of M  $>$  900 then
5:     Get action  $a_t$  from advisor
6:   else
7:     Choose a random action  $a_t$ 
8:   end if
9:   Perform action  $a_t$ 
10:  Observe  $r_t$  and next state  $s_{t+1}$ 
11:  Add ( $< s_t, a_t, r_t, s_{t+1} >$ ) to M
12:  if  $s_t$  is terminal OR time-steps  $>$  250 then
13:    Reset episode
14:  end if
15: end while

```

---



---

**Algorithm 2** Training algorithm to populate and extract information from the batch memory.

---

```

1: Perform the pretraining Algorithm 1
2: for each episode do
3:   Observe state  $s_t$ 
4:   repeat
5:     Choose an action  $a_t$  using  $\epsilon$ -greedy
6:     Perform action  $a_t$ 
7:     Observe  $r_t$  and next state  $s_{t+1}$ 
8:     Add ( $< s_t, a_t, r_t, s_{t+1} >$ ) to M
9:     Populate randomly batch B from M
10:    Train CNN using data in B
11:     $s_t \leftarrow s_{t+1}$ 
12:     $\epsilon \leftarrow \epsilon * \epsilon\_decay$ 
13:  until  $s_t$  is terminal OR time-steps  $>$  250
14: end for

```

---

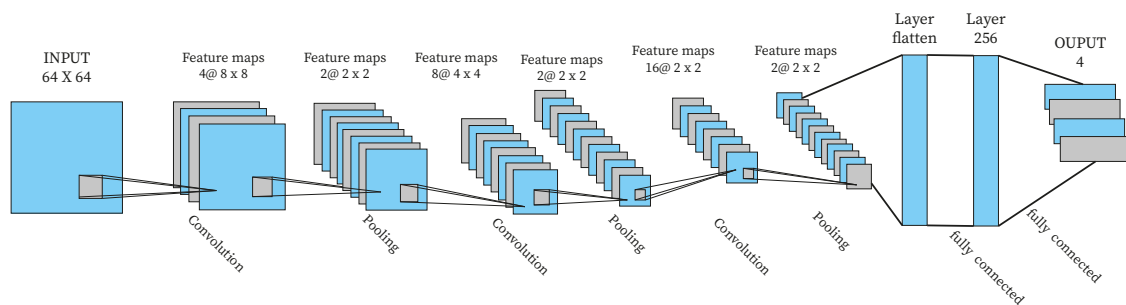
### 3.1.1. Interactive Approach

As previously discussed, the IDeepRL methods include an external advisor, which can be another already trained agent or human. In our scenario, the advisor uses a policy-shaping approach during the decision-making process, as previously shown in Figure 1. Moreover, between the different alternatives for interactive feedback, we use teaching on a budget with early advising [39]. This technique attempts to reduce the time required for an RL agent to understand the environment better, achieved by 100 early consecutive pieces of advice from the trainer, trying to transfer the trainer's knowledge of the environment as quickly as possible. For the different ways to implement the interactive approach, we use early advising for the training of the learner agent, using a limited consecutive amount of advice to be used by the trainer to help the agent.

### 3.1.2. Visual Representation

A visual representation for the deep Q-learning algorithm is used, which consists of Q-learning using a function approximator for the Q-values with a CNN. Additionally, it uses a memory with past experiences from where batches for the network training are taken.

Our architecture is capable of processing input images of  $64 \times 64$  pixels used in RGB channels for learning the image features. The architecture is inspired by similar networks used in other DeepRL works [13,14]. Particularly, in the first layer, an  $8 \times 8$  convolution with four filters is used, then a  $2 \times 2$  max-pooling layer followed by a  $4 \times 4$  convolution layer with eight filters, followed by another max-pooling with the same specification as the previous one. The network has a last  $2 \times 2$  convolution layer with 16 filters. After the last pooling, a flattened layer is applied, which is fully connected to a layer with 256 neurons. Finally, the 256 neurons are also fully connected with the output layer, which uses a softmax function, including four neurons to represent the possible actions. The full network architecture can be seen in Figure 3. Since this work is oriented to compare different learning methodologies, all agents were trained with the same architecture to compare them fairly.



**Figure 3.** Neural network architecture, with an input of a  $64 \times 64$  RGB image, and composed of three convolution layers, three max-pooling layers, and two fully connected layers, including a softmax function for the output.

### 3.1.3. Continuous Representation

Given the task characteristics, considering images as inputs to recognize different objects in a dynamic environment, it is impractical to generate a table with all the possible state-action combinations. Therefore, we have used a continuous representation combining two methods. The first method is a function approximator through a neuronal network for  $Q(s_t, a_t)$ , which allows us to generalize the states, in order to use Q-learning in continuous spaces and select which action is carried out. The second method is the experience replay technique, which saves into a memory a tuple of an experience given by  $\langle s_t, a_t, r_t, s_{t+1} \rangle$ . These data saved in memory are used afterward to train the RL agent.

## 3.2. Experimental Setup


We have designed a simulated domestic scenario focused on organizing objects. The agent aims to classify geometric figures with different shapes and colors and organize them in designated locations to optimize the collected reward. Classification tasks are widespread in domestic scenarios, e.g., organizing cloth. The object shape might represent different cloth types, while the color might represent whether it is clean or dirty.

In order to compare DeepRL and IDepRL algorithms, three different agents are trained in this scenario in terms of collected reward and learning time. The experimental scenario is developed in the simulator CoppeliaSim developed by Coppelia Robotics [40].

Three tables are used in the scenario; the first contains all the unsorted objects initially placed randomly on the table within nine preestablished positions. This represents the initial dynamic state. The second two tables are used to place the objects once the agent determines to which table the objects



belong. To perform this sort, a robotic manipulator arm with seven degrees of freedom, six axes, and a suction cup grip is used. The robotic arm is placed on another table along with a camera from where we obtain RGB images. The objects to be organized are cubes, cylinders, and disks in two different colors, red and blue, as are presented in Figure 4.



**Figure 4.** The simulated domestic scenario presenting six objects in different colors and the robotic arm. In the upper left corner, the camera signal is shown, which is a  $64 \times 64$  pixels RGB image for the state representation of the agent.

### 3.2.1. Actions and State Representation

The available actions for the agent are four and can be taken in an autonomous way or through given advice from the external trainer. The actions are the following:

- i Grab object: the agent grabs randomly one of the objects with the suction cup grip.
- ii Move right: the robotic arm is moved to the table on the right side of the scenario; if the arm is already there, do nothing.
- iii Move left: the robotic arm is moved to the table on the left side of the scenario; if the arm is already there, do nothing.
- iv Drop: if the robotic arm has an object in the grip and is located in one of the side tables, the arm goes down and releases the object; in case the arm is positioned in the center, the arm keeps the position.

For example, the actions required to correctly organize a blue cube from the central table consist of (i) grab object, (ii) move right, and (iii) drop. The robot low-level control to reach the different positions within the scenario is performed using inverse kinematics. Although to reach an object we use inverse kinematics, the CNN is responsible for deciding to perform the action *grab an object* through the Q-values, and if so, to decide where to place the object, based on the classification.

The state comprises a high-dimensional space, represented by a raw image captured by an RGB camera. The image presents a dimension of  $64 \times 64$  pixels from where the agent perceives the environment and chooses what action to take according to the network output. The input image is normalized to values  $\in [0, 1]$  to be presented to the convolutional neural network.

### 3.2.2. Reward

In terms of the reward function, there are different subtasks to complete an episode. To complete the task successfully, all the objects must be correctly organized. Organizing one object is considered a partial goal of the task. All of the objects are initially located in the central table to be sorted, and once placed in the side tables, they cannot be grasped again. If all the objects are correctly sorted, the reward

is equal to 1, and, for correctly organizing a single object, the reward is equal to 0.4. For example, if the classification of the six objects is correct, each of the first five organized objects leads to a reward of 0.4, and the last one obtains a reward of 1, summarizing a total reward of 3. Furthermore, to encourage the agent to accomplish the classification task in fewer steps, a small negative reward of  $-0.01$  per step is considered when the steps are more than 18, which is the minimal time-steps needed to complete the task satisfactorily. If an object is misplaced, the current training episode ends, and the agent receives a negative reward of  $-1$ . The complete reward function is shown in Equation (3):

$$r(s) = \begin{cases} 1 & \text{if all the objects are organized} \\ 0.4 & \text{if a single object is organized} \\ -1 & \text{if an object is incorrectly organized} \\ -0.01 & \text{if steps} > 18 \end{cases} \quad (3)$$

### 3.2.3. Human Interaction

In the case of human trainers giving advice during the second stage of the IDeepRL approach, a brief three-step induction is carried out for each participant:

- i The user reads an introduction to the scenario and the experiment, as well as the expected results.
- ii The user is given an explanation about the problem and how it can be solved.
- iii The user is taught how to use the computer interface to advise actions to the learner agent.

In general terms, the participants chosen for the experiment have not had significant exposure to artificial intelligence, and are not familiar with simulated robotic environments. The solution explanation is given to the participants in order to give to all of them an equal understanding of the problem and thus to reduce the time that they spend exploring the environment and focus on advising the agent. Each participant communicates with the learner agent using a computer interface while observing the current state and performance in the robot simulator. The user interface contains in a menu all possible actions that can be advised. These action possibilities are shown at the screen, and the trainer may choose any of them to be performed by the learner agent. There is no time limit to advise each action, but, as mentioned, during the second stage of IDeepRL, the trainer has a maximum of 100 consecutive time-steps available for advice.

## 4. Results

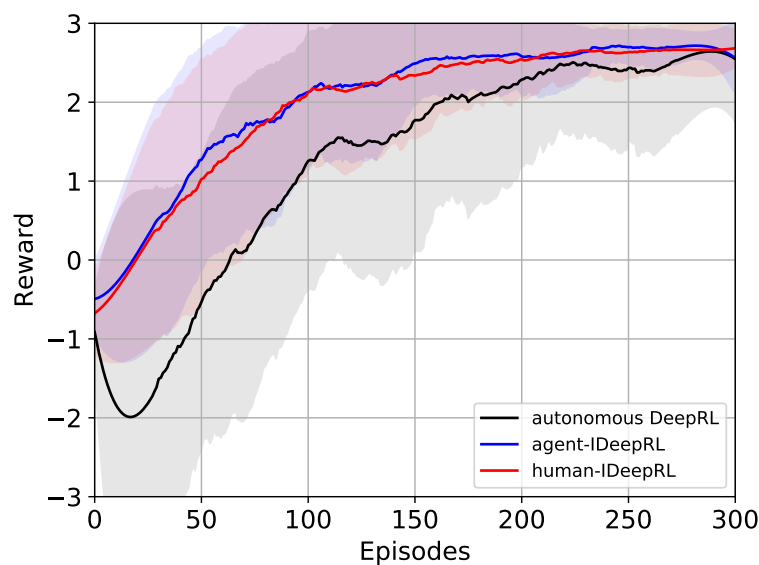
In this section, we show the experimental results obtained during the training of three different agents implemented with the three proposed methodologies, i.e., an autonomous agent scenario, a human–agent scenario, and an agent–agent scenario, namely, DeepRL, human–IDeepRL, and agent–IDeepRL. The methodologies are tested with the same hyperparameters, which have been experimentally determined concerning our scenario, as follows: initial value of  $\epsilon = 1$ ,  $\epsilon$  decay rate of 0.9995, learning rate  $\alpha = 10^{-3}$ , and discount factor  $\gamma = 0.9$  during 300 episodes.

As discussed in Section 3.1, the first methodology is an autonomous agent using DeepRL, who must learn how the environment works and how to complete the task. Given the complexity of learning the task autonomously, the time required for the learning process is rather high. The average collected reward for ten autonomous agents is shown in Figure 5 represented by the black line. Moreover, this complexity also increases the error rate or misclassification of the objects located in the central table.

Next, we perform the interactive learning approaches by using agent–IDeepRL and human–IDeepRL. The average obtained results for ten interactive agents are shown in Figure 5 represented by the blue line and the red line, respectively. The agent–IDeepRL approach performs slightly better than the human–IDeepRL approach, mainly because people needed more time to understand the setup and to react during the experiments. However, both interactive approaches obtain very similar results, achieving much faster convergence when comparing to the autonomous DeepRL approach.

Furthermore, the learner agents getting advice from external trainers make fewer mistakes, especially at the beginning of the learning process, and are able to learn the task in fewer episodes. On the other hand, the autonomous agent makes more mistakes at the beginning since it is trying to learn how the environment works and the aim that it has to be accomplished. This is not the case for the interactive agents since the advisors help them during this critical part of the learning. For a qualitative analysis of the obtained results, we have computed the total collected reward  $R_T$  defined as the sum of all individual rewards  $r_i$  received during the learning process, i.e.,  $N = 300$  for this case (see Equation (4)). The total collected rewards for the three methods are: autonomous DeepRL  $R_T = 369.8788$ , agent-IDeepRL  $R_T = 606.6995$ , and human-IDeepRL  $R_T = 589.5974$ . Therefore, the methods agent-IDeepRL and human-IDeepRL present an improvement, in terms of collected reward, of 64.03% and 59.40%, respectively, in comparison to the autonomous DeepRL used as a baseline:

$$R_T = \sum_i^N r_i \quad (4)$$




**Figure 5.** Average collected reward for the three proposed methods. The black line represents the autonomous (DeepRL) agent, which has to discover the environment without any help. The blue and red lines are the agents with an external trainer, namely an artificial advisor (agent-IDeepRL) and a human advisor (human-IDeepRL), respectively. The shadowed area around the curves shows the standard deviation for ten agents. The methods agent-IDeepRL and human-IDeepRL collect 64.03% and 59.40% more reward  $R_T$  when compared to the autonomous DeepRL baseline.

Due to the trainer having a budget of 100 actions to advise, the interactive feedback is consumed within the first six episodes, taking into account that the minimal amount of actions to complete an episode are 18 actions. Even with such a small amount of feedback, the learner agent receives an important knowledge from the advisor that is complemented with the experience replay method. During the human-IDeepRL approach, to give demonstrative advice, 11 people participated in the experiments, four males and seven females, with ages between 16 and 24 ( $M = 21.63, SD = 2.16$ ). The participants were explained how to help the robot to complete the task giving advice using the same script for all of them (see Section 3.1).


Figure 6 shows the collected rewards by an autonomous DeepRL agent and three learner agents trained by different people as examples. All human-IDeepRL approaches work much better than the autonomous DeepRL, making fewer mistakes and, therefore, collecting faster rewards. Between the interactive agents, there are some differences, especially at the beginning of the learning process, within the first 50 episodes. It is possible to observe the different strategies followed by the human trainers, for instance, the sixth trainer (yellow line, labeled as human-IDeepRL6) started giving

wrong advice, leading to less reward at the beginning, while the eighth trainer (green line, labeled as human-IDeepRL8) started giving almost perfect advice, experiencing a drop in the collected reward some episodes later. Nevertheless, all the agents, even the autonomous, managed the task reaching a similar reward well. As in the previous case, we have computed the total collected reward  $R_T$  (see Equation (4)) for each agent shown in Figure 6. The autonomous DeepRL agent collected 266.0905 of total reward, whereas the agents with the human-IDeepRL method using as trainers subjects 2, 6, and 8 collected 603.6375, 561.4384, and 630.4684, respectively. Although there are differences in the way that each trainer instructs the learner agent, yet they have accomplished an improvement in the total accumulated reward of 126.85%, 110.99%, and 136.94%.



**Figure 6.** Collected rewards for a selection of example interactive agents. The figure compares the learning process of agents trained by different people using human-IDeepRL (the black line is an autonomous agent that is included as a reference). The three human examples differ from each other by following different strategies to teach the learner agent, especially at the beginning. For example, the sixth trainer (yellow) started giving wrong advice, leading to less reward collected initially, while the eighth trainer (green) started giving much better advice, which is reflected in the accumulated reward at the beginning; however, it experienced a drop in the collected reward some episodes later. Although each person has initially a different understanding of the environment considering objectives and possible movements, all the interactive agents converge to similar behavior at the end of the learning process. Quantitatively, the interactive agents collected 126.85%, 110.99%, and 136.94% more reward  $R_T$  when compared to the autonomous DeepRL agent.

Figure 7 shows Pearson's correlation of the collected rewards for all the interactive agents trained by the participants in the experiment. Moreover, we include an autonomous agent ( $A_{Au}$ ) and an interactive agent trained by an artificial trainer agent ( $A_{AT}$ ) as a reference. It is possible to observe that all interactive agents, including the one using an artificial trainer agent, have a high correlation in terms of the collected reward. However, the autonomous agent shows a lower correlation in comparison to the interactive approaches.



**Figure 7.** Pearson's correlation between the collected rewards for different agents. The shown agents include an autonomous agent ( $A_{Au}$ ), an interactive agent trained by an artificial trainer ( $A_{AT}$ ), and the interactive agents trained by humans (from  $A_{H1}$  to  $A_{H11}$ ). The collected reward for all the interactive approaches, including the one using the artificial trainer, presents a similar behavior showing a high correlation. On the contrary, the collected reward by the autonomous agent shows a lower correlation in comparison to the interactive agents.

Additionally, we have computed the Student's  $t$ -test to test the statistical difference between the obtained results. When the autonomous DeepRL approach is compared to agent-IDeepRL and human-IDeepRL, it obtains a  $t$ -score  $t = 7.6829$  ( $p$ -value  $p = 6.3947 \times 10^{-14}$ ) and a  $t$ -score  $t = 7.0192$  ( $p$ -value  $p = 6.0755 \times 10^{-12}$ ), respectively, indicating that there is a statistically significant difference between the two approaches. On the other hand, comparing both interactive approaches between each other, i.e., agent-IDeepRL and human-IDeepRL, a  $t$ -score  $t = 0.8461$  ( $p$ -value  $p = 0.3978$ ) is obtained, showing that there is no statistical difference between the interactive methods. Table 2 shows all the obtained  $t$ -scores along with the  $p$ -values for each of them.

**Table 2.** Student's  $t$ -test for comparison of autonomous DeepRL, agent-IDeepRL, and human-IDeepRL.

DeepRL vs. Agent-IDeepRL	DeepRL vs. Human-IDeepRL	Agent-IDeepRL vs. Human-IDeepRL
$t = 7.6829$ $p = 6.3947 \times 10^{-14}$	$t = 7.0192$ $p = 6.0755 \times 10^{-12}$	$t = 0.8461$ $p = 0.3978$

In all the tested approaches, approximately since episode 150, the agent performs actions mainly based on its learning or training. In that episode, the value of  $\epsilon$  in the  $\epsilon$ -greedy policy decay to 1% of exploratory actions. Moreover, in all the approaches, the maximal collected reward fluctuates between 2.5 and 3. This is because the robot, with its movements, sometimes throws away another object from the table, different from the one being manipulated.

## 5. Conclusions

We have presented an interactive deep reinforcement learning approach to train an agent in a Human–Robot environment. We have also performed a comparison between three different methods

for learning agents. First, we implemented an autonomous version of DeepRL, which had to interact and learn the environment by itself. Next, we proposed an interactive version called IDeepRL, which used an external trainer to give useful advice during the decision-making process through interactive feedback delivered through early advising.

We have implemented two variations of IDeepRL by using previously trained artificial agents and humans as trainers. We called these approaches agent–IDeepRL and human–IDeepRL, respectively. Our proposed IDeepRL methods considerably outperform the autonomous DeepRL version in the implemented robotic scenario. Moreover, in complex tasks, which often require more training time, to have an external trainer giving supportive feedback leads to great benefits in terms of time and collected reward.

Overall, the interactive deep reinforcement learning approach introduces an advantage in domestic-like environments. It allows for speeding up the learning process of a robotic agent interacting with the environment and allows people to transfer prior knowledge about a specific task. Furthermore, using a reinforcement learning approach allows the agent to learn the task without the necessity of previously labeled data, such as the case for supervised learning methods. In this regard, the task is learned in such a way that the agent learns to recognize the state of the environment as well as to behave in regards to it, deciding where to place the different objects. Our novel interactive-shaping vision-based approach outperforms the current autonomous DeepRL method, used as a baseline in this work. The introduced approaches demonstrate that the use of external trainers, either artificial or human, lead to more and faster reward in comparison to the traditional DeepRL approach.

As future work, we consider the use of different kinds of artificial trainers to select possible advisors better. A bad teacher can negatively influence the learning process and somehow limit the learner by teaching a strategy that is not necessarily optimal. To select a good teacher, it is necessary to take into account that an agent that obtains the best results for the task, in terms of accumulated reward, is not necessarily the best teacher [27]. Instead, a good teacher could be one with a small standard deviation over the visited states. This would allow for advising the learner agent in more specific situations. Additionally, we plan to transfer and test the proposed approach in a real-world Human–Robot interaction scenario. In such a case, it is necessary to deal with additional environmental dynamics and variables that in simulated scenarios may be easier to control.

**Author Contributions:** Conceptualization, F.C.; Funding acquisition, F.C. and B.F.; Investigation, I.M. and J.R.; Methodology, A.A.; Supervision, F.C.; Writing—original draft, I.M. and J.R.; Writing—review and editing, F.C., R.D., A.A. and B.F. All authors have read and agreed to the published version of the manuscript.

**Funding:** This research was partially funded by Universidad Central de Chile under the research project CIP2018009, the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior—Brasil (CAPES)—Finance Code 001, the Brazilian agencies FACEPE, and CNPq.

**Conflicts of Interest:** The authors declare no conflict of interest.

## References

1. Shepherd, S.; Buchstab, A. Kuka robots on-site. In *Robotic Fabrication in Architecture, Art and Design 2014*; Springer: Berlin/Heidelberg, Germany, 2014; pp. 373–380.
2. Cruz, F.; Wüppen, P.; Fazrie, A.; Weber, C.; Wermter, S. Action selection methods in a robotic reinforcement learning scenario. In *Proceedings of the 2018 IEEE Latin American Conference on Computational Intelligence (LA-CCI)*, Guadalajara, Mexico, 7–9 November 2018; pp. 13–18.
3. Goodrich, M.A.; Schultz, A.C. Human–robot interaction: A survey. *Found. Trends® Hum.-Interact.* **2008**, *1*, 203–275. [[CrossRef](#)]
4. Churamani, N.; Cruz, F.; Griffiths, S.; Barros, P. iCub: Learning emotion expressions using human reward. *arXiv* **2020**, arXiv:2003.13483.
5. Sutton, R.S.; Barto, A.G. *Reinforcement Learning: An Introduction*; MIT Press: Cambridge, MA, USA, 2018.
6. Niv, Y. Reinforcement learning in the brain. *J. Math. Psychol.* **2009**, *53*, 139–154. [[CrossRef](#)]



7. Cruz, F.; Parisi, G.I.; Wermter, S. Multi-modal feedback for Affordance-driven interactive reinforcement learning. In Proceedings of the 2018 International Joint Conference on Neural Networks (IJCNN), Rio, Brazil, 8–13 July 2018; pp. 1–8.
8. Bignold, A.; Cruz, F.; Taylor, M.E.; Brys, T.; Dazeley, R.; Vamplew, P.; Foale, C. A conceptual framework for externally-influenced agents: An assisted reinforcement learning review. *arXiv* **2020**, arXiv:2007.01544.
9. Ayala, A.; Henríquez, C.; Cruz, F. Reinforcement learning using continuous states and interactive feedback. In Proceedings of the International Conference on Applications of Intelligent Systems, Las Palmas de Gran Canaria, Spain, 7–12 January 2019; pp. 1–5.
10. Millán, C.; Fernandes, B.; Cruz, F. Human feedback in continuous actor-critic reinforcement learning. In Proceedings of the 27th European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning ESANN, Bruges, Belgium, 24–26 April 2019; pp. 661–666.
11. Barros, P.; Tanevska, A.; Cruz, F.; Sciutti, A. Moody Learners—Explaining Competitive Behaviour of Reinforcement Learning Agents. *arXiv* **2020**, arXiv:2007.16045.
12. LeCun, Y.; Bengio, Y.; Hinton, G. Deep learning. *Nature* **2015**, *521*, 436–444. [[CrossRef](#)] [[PubMed](#)]
13. Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A.A.; Veness, J.; Bellemare, M.G.; Graves, A.; Riedmiller, M.; Fidjeland, A.K.; Ostrovski, G.; et al. Human-level control through deep reinforcement learning. *Nature* **2015**, *518*, 529. [[CrossRef](#)] [[PubMed](#)]
14. Krizhevsky, A.; Sutskever, I.; Hinton, G.E. Imagenet classification with deep convolutional neural networks. In Proceedings of the Twenty-sixth Annual Conference on Neural Information Processing Systems (NIPS), Lake Tahoe, NV, USA, 3–8 December 2012; pp. 1097–1105.
15. Van Hasselt, H.; Guez, A.; Silver, D. Deep reinforcement learning with double q-learning. In Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, Phoenix, AZ, USA, 12–17 February 2016; pp. 2094–2100.
16. Puterman, M.L. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*; Wiley: Hoboken, NJ, USA, 1994.
17. Suay, H.B.; Chernova, S. Effect of human guidance and state space size on interactive reinforcement learning. In Proceedings of the International Symposium on Robot and Human Interaction Communication (Ro-Man), Atlanta, GA, USA, 31 July–3 August 2011; pp. 1–6.
18. Najar, A.; Chetouani, M. Reinforcement learning with human advice. A survey. *arXiv* **2020**, arXiv:2005.11016.
19. Ng, A.Y.; Harada, D.; Russell, S. Policy invariance under reward transformations: Theory and application to reward shaping. In Proceedings of the Sixteenth International Conference on Machine Learning (ICML), Bled, Slovenia, 27–30 June 1999; Volume 99, pp. 278–287.
20. Brys, T.; Nowé, A.; Kudenko, D.; Taylor, M.E. Combining multiple correlated reward and shaping signals by measuring confidence. In Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence, Québec City, QC, Canada, 27–31 July 2014; pp. 1687–1693.
21. Griffith, S.; Subramanian, K.; Scholz, J.; Isbell, C.; Thomaz, A.L. Policy shaping: Integrating human feedback with reinforcement learning. In Proceedings of the Twenty-seventh Annual Conference on Neural Information Processing Systems (NIPS), Lake Tahoe, NV, USA, 5–10 December 2013; pp. 2625–2633.
22. Li, G.; Gomez, R.; Nakamura, K.; He, B. Human-centered reinforcement learning: A survey. *IEEE Trans. Hum.-Mach. Syst.* **2019**, *49*, 337–349. [[CrossRef](#)]
23. Grizou, J.; Lopes, M.; Oudeyer, P.Y. Robot learning simultaneously a task and how to interpret human instructions. In Proceedings of the 2013 IEEE Third Joint International Conference on Development and Learning and Epigenetic Robotics (ICDL-EpiRob), Osaka, Japan, 18–22 August 2013; pp. 1–8.
24. Navidi, N. Human AI interaction loop training: New approach for interactive reinforcement learning. *arXiv* **2020**, arXiv:2003.04203.
25. Bignold, A. Rule-Based Interactive Assisted Reinforcement Learning. Ph.D. Thesis, Federation University, Ballarat, Australia, 2019.
26. Taylor, M.E.; Carboni, N.; Fachantidis, A.; Vlahavas, I.; Torrey, L. Reinforcement learning agents providing advice in complex video games. *Connect. Sci.* **2014**, *26*, 45–63. [[CrossRef](#)]
27. Cruz, F.; Magg, S.; Nagai, Y.; Wermter, S. Improving interactive reinforcement learning: What makes a good teacher? *Connect. Sci.* **2018**, *30*, 306–325. [[CrossRef](#)]

28. Dobrovsky, A.; Borghoff, U.M.; Hofmann, M. Improving adaptive gameplay in serious games through interactive deep reinforcement learning. In *Cognitive Infocommunications, Theory and Applications*; Springer: Berlin/Heidelberg, Germany, 2019; pp. 411–432.
29. Rajeswaran, A.; Kumar, V.; Gupta, A.; Vezzani, G.; Schulman, J.; Todorov, E.; Levine, S. Learning complex dexterous manipulation with deep reinforcement learning and demonstrations. *arXiv* **2017**, arXiv:1709.10087.
30. Lukka, T.J.; Tossavainen, T.; Kujala, J.V.; Raiko, T. ZenRobotics recycler—Robotic sorting using machine learning. In Proceedings of the International Conference on Sensor-Based Sorting (SBS), Aachen, Germany, 11–13 March 2014.
31. Zhihong, C.; Hebin, Z.; Yanbo, W.; Binyan, L.; Yu, L. A vision-based robotic grasping system using deep learning for garbage sorting. In Proceedings of the 2017 36th Chinese Control Conference (CCC), Dalian, China, 26–28 July 2017; pp. 11223–11226.
32. Sun, L.; Aragon-Camarasa, G.; Rogers, S.; Stolkin, R.; Siebert, J.P. Single-shot clothing category recognition in free-configurations with application to autonomous clothes sorting. In Proceedings of the 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Vancouver, BC, Canada, 24–28 September 2017; pp. 6699–6706.
33. Cruz, F.; Magg, S.; Weber, C.; Wermter, S. Improving reinforcement learning with interactive feedback and affordances. In Proceedings of the 4th International Conference on Development and Learning and on Epigenetic Robotics (ICDL-EpiRob), Genoa, Italy, 13–16 October 2014; pp. 165–170.
34. Cruz, F.; Parisi, G.I.; Wermter, S. Learning contextual affordances with an associative neural architecture. In Proceedings of the 24th European Symposium on Artificial Neural Network, Computational Intelligence and Machine Learning ESANN, Bruges, Belgium, 27–29 April 2016; pp. 665–670.
35. Zhang, F.; Leitner, J.; Milford, M.; Upcroft, B.; Corke, P. Towards vision-based deep reinforcement learning for robotic motion control. *arXiv* **2015**, arXiv:1511.03791.
36. Vecerik, M.; Hester, T.; Scholz, J.; Wang, F.; Pietquin, O.; Piot, B.; Heess, N.; Rothörl, T.; Lampe, T.; Riedmiller, M. Leveraging demonstrations for deep reinforcement learning on robotics problems with sparse rewards. *arXiv* **2017**, arXiv:1707.08817.
37. Desai, N.; Banerjee, A. *Deep Reinforcement Learning to Play Space Invaders*; Technical Report; Stanford University: Stanford, CA, USA, 2017.
38. Adam, S.; Busoniu, L.; Babuska, R. Experience replay for real-time reinforcement learning control. *IEEE Trans. Syst. Man Cybern. Part C Appl. Rev.* **2012**, *42*, 201–212. [[CrossRef](#)]
39. Cruz, F.; Wüppen, P.; Magg, S.; Fazrie, A.; Wermter, S. Agent-advising approaches in an interactive reinforcement learning scenario. In Proceedings of the 2017 Joint IEEE International Conference on Development and Learning and Epigenetic Robotics (ICDL-EpiRob), Lisbon, Portugal, 18–21 September 2017; pp. 209–214.
40. Rohmer, E.; Singh, S.P.; Freese, M. V-REP: A versatile and scalable robot simulation framework. In Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Tokyo, Japan, 3–7 November 2013; pp. 1321–1326.



© 2020 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).