

Article

MarsExplorer: Exploration of Unknown Terrains via Deep Reinforcement Learning and Procedurally Generated Environments

Dimitrios I. Koutras ^{1,2}, Athanasios C. Kapoutsis ^{2,*}, Angelos A. Amanatiadis ³
and Elias B. Kosmatopoulos ^{1,2}

¹ Department of Electrical and Computer Engineering, Democritus University of Thrace, 671 00 Xanthi, Greece; dkoutras@iti.gr (D.I.K.); kosmatop@iti.gr (E.B.K.)

² The Centre for Research & Technology, Information Technologies Institute, Hellas, 570 01 Thessaloniki, Greece

³ Department of Production and Management Engineering, Democritus University of Thrace, 671 00 Xanthi, Greece; aamanat@pme.duth.gr

* Correspondence: athakapo@iti.gr

Abstract: This paper is an initial endeavor to bridge the gap between powerful Deep Reinforcement Learning methodologies and the problem of exploration/coverage of unknown terrains. Within this scope, MarsExplorer, an openai-gym compatible environment tailored to exploration/coverage of unknown areas, is presented. MarsExplorer translates the original robotics problem into a Reinforcement Learning setup that various off-the-shelf algorithms can tackle. Any learned policy can be straightforwardly applied to a robotic platform without an elaborate simulation model of the robot's dynamics to apply a different learning/adaptation phase. One of its core features is the controllable multi-dimensional procedural generation of terrains, which is the key for producing policies with strong generalization capabilities. Four different state-of-the-art RL algorithms (A3C, PPO, Rainbow, and SAC) are trained on the MarsExplorer environment, and a proper evaluation of their results compared to the average human-level performance is reported. In the follow-up experimental analysis, the effect of the multi-dimensional difficulty setting on the learning capabilities of the best-performing algorithm (PPO) is analyzed. A milestone result is the generation of an exploration policy that follows the Hilbert curve without providing this information to the environment or rewarding directly or indirectly Hilbert-curve-like trajectories. The experimental analysis is concluded by evaluating PPO learned policy algorithm side-by-side with frontier-based exploration strategies. A study on the performance curves revealed that PPO-based policy was capable of performing adaptive-to-the-unknown-terrain sweeping without leaving expensive-to-revisit areas uncovered, underlying the capability of RL-based methodologies to tackle exploration tasks efficiently.

Keywords: Deep Reinforcement Learning; OpenAI gym; exploration; unknown terrains



check for updates

Citation: Koutras, D.I.; Kapoutsis, A.C.; Amanatiadis, A.A.; Kosmatopoulos, E.B. MarsExplorer: Exploration of Unknown Terrains via Deep Reinforcement Learning and Procedurally Generated Environments. *Electronics* **2021**, *10*, 2751. <https://doi.org/10.3390/electronics10222751>

Academic Editors: Ahmad Taher Azar, Anis Koubaa, Alaa Khamis, Ibrahim A. Hameed and Gabriella Casalino

Received: 5 October 2021

Accepted: 8 November 2021

Published: 11 November 2021

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2021 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 (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

1.1. Motivation

At this very moment, three different uncrewed spaceships, PERSEVERANCE (USA), HOPE (UAE), TIANWEN-1 (China), are on the surface or in the orbit of Mars. Never before has such a diverse array of scientific gear arrived at a foreign planet at the same time, and with such broad ambitions [1]. On top of that, several lunar missions have been arranged for this year to enable extensive experimentation, investigation, and testing on an extraterrestrial body [2]. In this exponentially growing field of extraterrestrial missions, a task of paramount importance is the autonomous exploration/coverage of previously unknown areas. The effectiveness and efficiency of such autonomous explorers may significantly impact the timely accomplishment of crucial tasks (e.g., before the fuel depletion) and, ultimately, the success (or not) of the overall mission.

Exploration/coverage of unknown territories is translated into the online design of the path for the robot, taking as input the sensory information with the objective of mapping the whole area in the minimum possible time [3,4]. This setup shares the same properties and objectives with the well-known NP-complete setup of Traveling Salesman Problem (TSP), with the even more restrictive property that the area to be covered is discovered incrementally during the operation.

1.2. Related Work

The well-established family of approaches incorporates the concept of *next best pose* process, i.e., a turn-based, greedy selection of the next best position (also known as *frontier-cell*) to acquire measurement, based on heuristic strategy (e.g., [5–7]). Although this family of approaches has been extensively studied, some inherent drawbacks significantly constrain its broader applicability. For example, every deadlock that may arise during the previously described optimization scheme should have been predicted, and a corresponding mitigation plan should have been already in place [8]; otherwise, the robot is going to be stuck in this locally optimal configuration [9]. On top of that, to engineer a multi-term strategy that reflects the task at hand is not always trivial [10].

The recent breakthroughs in Reinforcement Learning (RL), in terms of both algorithms and hardware acceleration, have spawned methodologies capable of achieving above human-level performance in high-dimensional, non-linear setups, such as the game of Go [11], atari games [12], multi-agent collaboration [13], robotic manipulation [14], etc. A milestone in the RL community was the standardization of several key problems under a common framework, namely openai-gym [15]. Such release eased the evaluation among different methodologies and ultimately led to the generation of a whole new series of RL frameworks with standardized algorithms (e.g., [16,17]), all tuned to tackle openai-gym compatible setups.

These breakthroughs motivated the appliance of RL methodologies in the path-planning/exploration robotic tasks. Initially, the problem of navigating a single robot in previously unknown areas to reach a destination, while simultaneously avoiding catastrophic collisions, was tackled with RL methods [18–20]. The first RL methodology solely developed for exploration of unknown areas was developed in [21], and has successfully presented the potential benefits of RL. Recently, RL methodologies have been proposed that seek to leverage the deployment of multi-robot systems to cover an operational area [22].

However, Ref. [22] assumes only a single geometry for the environment to be covered and thus being prone to overfit, rather than being able to generalize in different environments. This drawback is mitigated by Ref. [21] by introducing a learning scheme with 30 different environments during the training phase. Although such a methodology can adequately tackle the generalization problem, the RL agent's performance is still bounded to the diversity of the human-imported environments.

1.3. Contributions

The main contribution of this work is to provide a framework for learning exploration/coverage policies that possess strong generalization abilities due to the procedurally generated terrain diversity. The intuition behind such an approach to exploration tasks is the fact that most areas exhibit some kind of structure in their terrain topology, e.g., city blocks, trees in a forest, containers in ports, office complexes. Thereby, by training multiple times in such correlated and procedurally generated environments, the robot will grasp/understand the underlining structure and leverage it to efficiently complete its goal, even in areas that it has never been exposed to.

Within this scope, a novel openai-gym compatible environment for exploration/coverage of unknown terrains has been developed and is presented. All the core elements that govern a real exploration/coverage setup have been included. MarsExplorer is one of the few RL environments where any learned policy can be transferred to real-world robotic platforms, providing that a proper translation between the proprioceptive/exteroceptive sensors' read-

ings and the generation of 2D perception (occupancy map) and also an integration with the existing robotic systems (e.g., PID low level control, safety mechanisms, etc.) are implemented.

Four state-of-the-art RL algorithms, namely A3C [23], PPO [24], Rainbow [25], and SAC [26], have been evaluated on the MarsExplorer environment. To better comprehend these evaluation results, the average human-level performance in the MarsExplorer environment is also reported. A follow-up analysis utilizing the best-performing algorithm (PPO) is conducted with respect to the different levels of difficulty. The visualization of the produced trajectories revealed that the PPO algorithm had learned to apply the famous space-filling Hilbert curve, with the additional capability of avoiding on-the-fly obstacles that might appear on the terrain. The analysis is concluded with a scalability study and a comparison with non-learning methodologies.

It should be highlighted that the objective is not to provide another highly realistic simulator but a framework upon which RL methods (and also non-learning approaches) will be efficiently benchmarked in exploration/coverage tasks. Although several wrappers are available for high-fidelity simulators (e.g., Gazebo [27], ROS [28]) that could be tuned to formulate an exploration coverage setup, in practice the required execution time for each episode severely limits the type of algorithms that can be used (for example PPO usually needs several millions of steps for environment interactions to converge). To the best of our knowledge, this is the first openai-gym compatible framework oriented for robotic exploration/coverage of unknown areas.

Figure 1 presents 4 sample snapshots that illustrate the performance of a trained RL robot inside the MarsExplorer environment. Figure 1a demonstrates the robot's entry inside the unknown terrain, which is annotated with black color. Figure 1b illustrates all the so-far gained "knowledge", which is either depicted with Martian soil or brown boxes to denote free space or obstructed positions, respectively. An attractive trait is depicted in Figure 1c, where the robot chose to perform a dexterous maneuver between two obstacles to be as efficient as possible in terms of numbers of timesteps for the coverage task. Note that any collision with an obstacle would have resulted in a termination of the episode and, as a result, an acquisition of an extreme negative reward. Figure 1d illustrates the robot's final position, along with all the gained information for the terrain (non-black region) during the episode.

All in all, the main contributions of this paper are:

- Develop an open-source (<https://github.com/dimikout3/MarsExplorer> (accessed on 4 November 2021)), openai-gym compatible environment tailored explicitly to the problem of exploration of unknown areas with an emphasis on generalization abilities.
- Translate the original robotics exploration problem to an RL setup, paving the way to apply off-the-shelf algorithms.
- Perform preliminary study on various state-of-the-art RL algorithms, including A3C, PPO, Rainbow, and SAC, utilizing the human-level performance as a baseline.
- Challenge the generalization abilities of the best performing PPO-based agent by evaluating multi-dimensional difficulty settings.
- Present side-by-side comparison with frontier-based exploration strategies.

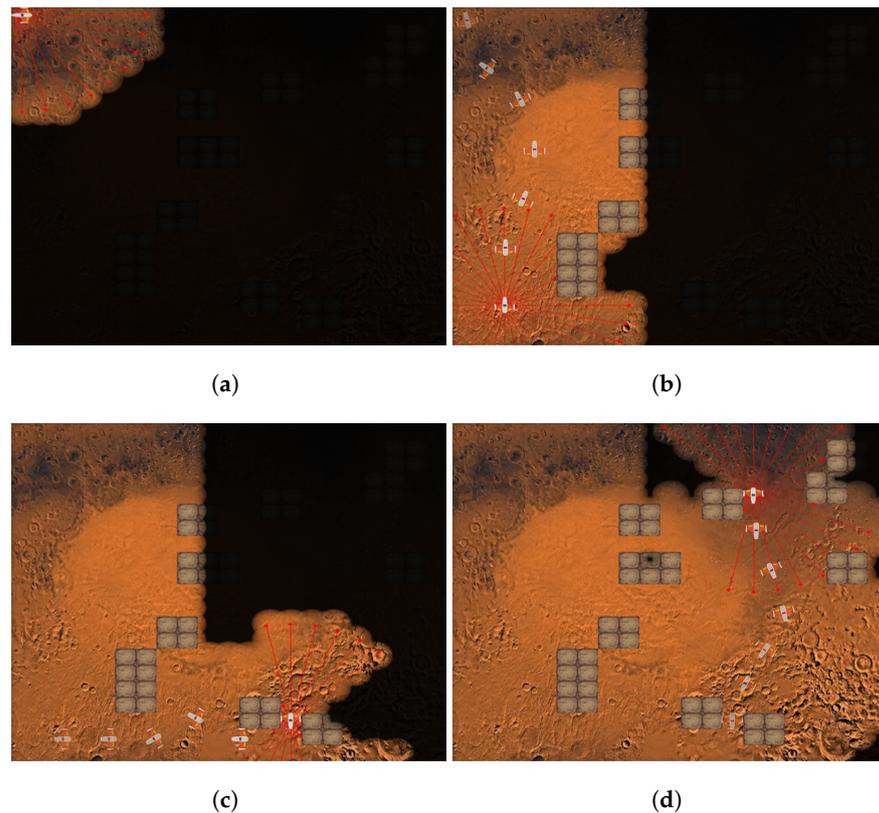


Figure 1. Indicative example: Trained RL agent executes exploration/coverage task in previously unknown and cluttered terrain utilizing MarsExplorer environment. (a) Initial timestep; (b) 30% progress; (c) 65% progress; (d) Final timestep.

1.4. Paper Outline

The remainder of this paper is organized as follows: Section 2 presents the details of the openai-gym exploration environment, called MarsExplorer, along with an analysis of the key RL attributes inserted. Section 3 presents the experimental analysis from the survey regarding the performance of the state-of-the-art RL algorithm to the evaluation against standard frontier-based exploration. Finally, Section 4 summarizes the findings the draws the conclusions of this study.

2. Environment

This section identifies the fundamental elements that govern the family of setups that fall into the coverage/exploration class and translates them to the openai gym framework [15]. In principle, the objective of the robot is to cover an area of interest in the minimum possible time while avoiding any non-traversable objects, the position of which gets revealed only when the robot's position is in close proximity [29,30].

2.1. Setup

Let us assume that the area to be covered is constrained within a square, which has been discretized into $n = rows \times cols$ identical grid cells:

$$\mathcal{G} = \{(x, y) : x \in [1, rows], y \in [1, cols]\} \quad (1)$$

The robot cannot move freely inside this grid, as some grid cells are occupied by non-traversable objects (obstacles). Therefore, the map of the terrain is defined as follows:

$$\mathcal{M}(q) = \begin{cases} 0.3 & \text{free space} \\ 1 & \text{obstacle} \end{cases} \quad q = (x, y) \in \mathcal{G} \quad (2)$$

The values of \mathcal{M} correspond to the morphology of the unknown terrain and are considered a priori unknown.

2.2. Action Space

Keeping in mind that the movement capabilities of the robot mainly impose the discretization of the area into grid cells, the action space is defined in the same grid context as well. The position of the robot is denoted by the corresponding x, y cell of the grid, i.e., $p_a(t) = [x_a(t), y_a(t)]$. Then, the possible next actions are simply given by the Von Neumann neighborhood [31], i.e.,

$$\mathcal{A}_{p_a} = \{(x, y) : |x - x_a| + |y - y_a| \leq 1\} \quad (3)$$

In the openai-gym framework, the formulation above is realized by a discrete space of size 4 (North, East, South, West).

2.3. State Space

With each movement, the robot may acquire some information related to the formation of the environment that lies inside its sensing capabilities, according to the following lidar-like model:

$$y_q(t) = \begin{cases} 1 & \text{if } \|p_a(t) - q\| \leq d \text{ AND} \\ & \exists \text{ line-of-sight between} \\ & p_a(t) \text{ and } q \\ 0 & \text{otherwise} \end{cases} \quad \forall q \in \mathcal{G} \quad (4)$$

where d denotes the maximum scanning distance.

An auxiliary boolean matrix $D(t)$ is introduced to denote all the cells that have been discovered from the beginning till t timestep. $D(t)$ annotates with one all cells that have been sensed and with zero all the others. Starting from a zero matrix $rows \times cols$, its values are updated as follows:

$$D_q(t) = D_q(t-1) \vee y_q(t), \quad \forall q \in \mathcal{G} \quad (5)$$

where \vee denotes the logical OR operator. The state is simply an aggregation of the acquired information over all past measurements of the robot (4). Having updated (5), the state $s(t_k)$ is a matrix of the same size as the grid to be explored (1), where its values are given by:

$$s_q(t) = \begin{cases} \mathcal{M}_q & \text{if } D_q(t) \\ 0 (= \text{undefined}) & \text{otherwise} \end{cases} \quad \forall q \in \mathcal{G} \quad (6)$$

Finally, the robot's position is declared by making the value of the corresponding cell equal to 0.6, i.e., $s_{q=p_a(t)}(t) = 0.6$. Overall, state $s(t)$ is defined as a 2D matrix, that takes values from the following discrete set: $\{0, 0.3, 0.6, 1\}$. Figure 2 presents an illustrative example of a registration between the graphical environment (Figure 2a) and the corresponding state representation (Figure 2b).

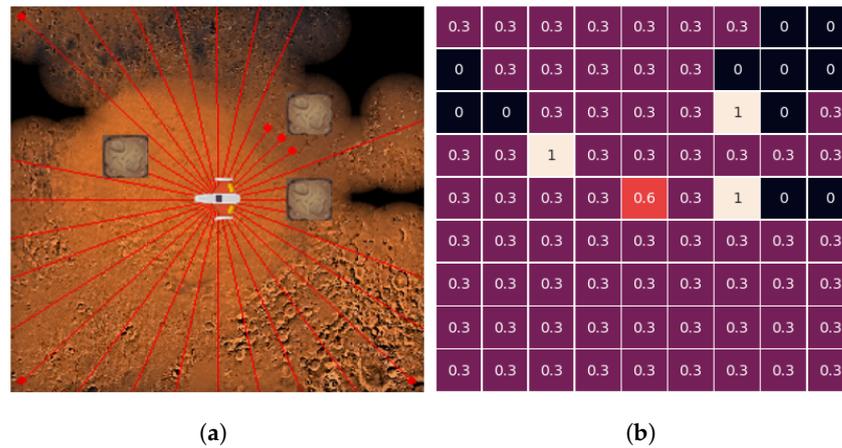


Figure 2. State encoding convention. (a) Graphical environment; (b) State $s(t)$ representation.

2.4. Reward Function

Having in mind that the ultimate objective is to discover all grid cells, the instantaneous core reward, at each timestep t , is defined as the number of newly explored cells, i.e.,

$$r_{explor}(t) = \sum_{q \in \mathcal{G}} D_q(t) - \sum_{q \in \mathcal{G}} D_q(t-1) \tag{7}$$

Intuitively, if $\sum_{k=0}^T r_{explor}(k) \rightarrow n$, then the robot has explored the whole grid (1) in T timesteps.

To force the robot to explore the whole area (7), while avoiding unnecessary movements, an additional penalty $r_{move} = 0.5$ per timestep is applied. In essence, this negative reward aims to distinguish among policies that lead to the same number of discovered cells but needed a different number of exploration steps. Please note that the value of r_{move} should be less than 1, to have less priority than the exploration of a single cell.

The action space, as defined previously, may include invalid next movements for the robot, i.e., out of the operational area (1) or crashing into a discovered obstacle. Thus, apart from the problem at hand, the robot should be able to recognize these undesirable states and avoid them at all costs. Towards that direction, an additional penalty $r_{invalid} = n$ is introduced for the cases where the robot’s next movement leads to an invalid state. Along with such a reward, the episode is marked as “done”, indicating that a new episode should be initiated.

At the other side of the spectrum, a completion bonus $r_{bonus} = n$ is given to the robot when more than $\beta\%$ (e.g., 95%) of the cells have been explored. Similar to the previous case, this is also considered a terminal state.

Putting everything together, the reward is defined as:

$$r(t) = \begin{cases} -r_{invalid} & \text{if next state is invalid} \\ r_{explor}(t) - r_{move} + \begin{cases} r_{bonus} & \text{if } \frac{\sum_{q \in \mathcal{G}} D_q(t)}{n} \geq \beta \\ 0 & \text{otherwise} \end{cases} & \text{otherwise} \end{cases} \tag{8}$$

2.5. Key RL Attributes

MarsExplorer was designed as an initial endeavor to bridge the gap between powerful existing RL algorithms and the problem of autonomous exploration/coverage of a previously unknown, cluttered terrain. This subsection presents the build-in key attributes of the designed framework.

Straightforward applicability. One of the fundamental attributes of MarsExplorer is that any learned policy can be straightforwardly applied to an appropriate robotic platform

with little effort required. This can be achieved by the fact that the policy calculates a high-level exploration path based on the perception of the environment (6). Thus, assuming that a smooth integration with the sensor's readings (for example, using a Kalman filter), can be used to represent the environment as in (6), no elaborate simulation model of the robot's dynamics is required to adjust the RL algorithm into the specifics of the robotic platform.

Terrain Diversity. For each episode, the general dynamics are determined by a specific automated process that has different levels of variation. These levels correspond to the randomness in the number, size, and positioning of obstacles, the terrain scalability (size), the percentage of the terrain that the robot must explore to consider the problem solved, and the bonus reward it will receive in that case. This procedural generation [32] of terrains allows training in multiple/diverse layouts, forcing, ultimately, the RL algorithm to enable generalization capabilities, which are of paramount importance in real-life applications where unforeseen cases may appear.

Partial Observability. Due to the nature of the exploration/coverage setup, at each timestep, the robot is only aware of the location of the obstacles that have been sensed from the beginning of the episode (5). Therefore, any long-term plan should be agile enough to be adjusted on the fly, based on future information about the unknown obstacles' positions. Such a property renders the acquisition of a global exploration strategy quite tricky [33].

Fast Evaluation. Disregarding the environment from any irrelevant physics dynamics and focusing only on the exploration/coverage aspect (1)–(8), MarsExplorer allows rapid execution of timesteps. This feature can be of paramount importance in the RL ecosystem, where the algorithms usually need millions of timesteps to converge, as it can enable fast experimental pipelines and prototyping.

3. Performance Evaluation

This section presents an experimental evaluation of the MarsExplorer environment. The analysis begins with all the implementation details that are important for realizing the MarsExplorer experimental setup. For the first evaluation iteration, 4 state-of-art RL algorithms are applied and evaluated in a challenging version of MarsExplorer that requires the development of *strong generalization* capabilities in a highly randomized scenario, where the underlying structure is almost absent. Having identified the best performing algorithm, a follow-up analysis is performed with respect to the difficulty vector values. The learned patterns and exploration policies for different evaluation instances are further investigated and graphically presented. The analysis is concluded with a scale-up study in two larger terrains and a comparison between the trained robot and two well-established frontier-based approaches.

3.1. Implementation Details

Aside from the standardization as an openai-gym environment, MarsExplorer provides an API that allows manually controlled experiments, translating commands from keyboard arrows to next movements. Such a feature can assess human-level performance in the exploration/coverage problem and reveal important traits by comparing human and machine-made strategies.

Ray/RLlib framework [34] was utilized to perform all the experiments. The fact that RLlib is a well-documented, highly-robust library also eases the build-on developments (e.g., apply a different RL pipeline), as it follows a common framework. Furthermore, such an experimental setup may also leverage the interoperability with other powerful frameworks from the Ray ecosystem, e.g., Ray/Tune for hyperparameters' tuning.

Table 1 summarizes all the fixed parameters used for all the performed experiments. MarsExplorer admits the distinguishing property of stochastically deploying the obstacles at the beginning of each episode. This stochasticity can be controlled and ultimately determines the difficulty level of the MarsExplorer setup. The state-space of MarsExplorer has a strong resemblance to thoroughly studied 2D environments, e.g., ALE [35], only with the key difference that the image is generated incrementally and based on the robot's actions.

Therefore, as it has been standardized from the DQN algorithm's application domain [12], a vision-inspired neural network architecture is incorporated as a first stage. Figure 3 illustrates the architecture of this pre-processor, which is comprised of 2 convolutional layers followed by a fully connected one. The vectorized output of the fully connected layer is forwarded to a "controller" architecture dependent on the RL algorithm enabled.

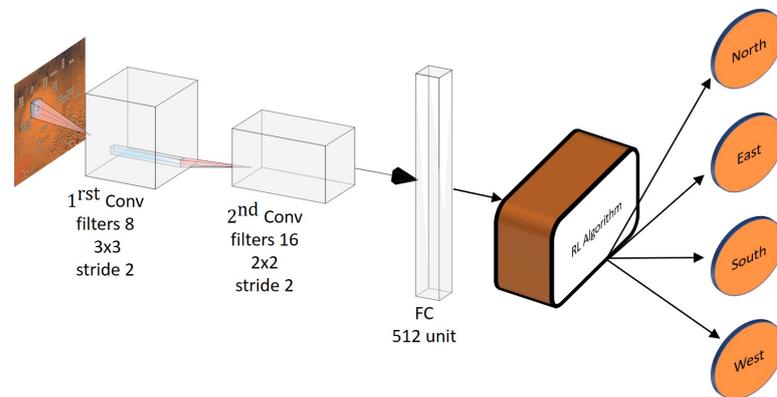


Figure 3. Overview of the experimental architecture.

Table 1. Implementation parameters.

Parameter	Value	Equation
Grid size	$[21 \times 21]$	(1)
Sensor radius	$d = 6$ grid cells	(4)
Considered done	$\beta = 99\%$	(8)

3.2. State-of-the-Art RL Algorithms Comparison

Apart from the details described in the previous subsection, for the comparison study, at the beginning of each episode, the formation (position and shape) of obstacles was set randomly. This choice was made to force RL algorithms to develop novel generalization strategies to tackle such a challenging setup. The list of studied RL algorithms is comprised by the following model-free approaches: PPO [24], DQN-Rainbow [25], A3C [23] and SAC [26]. All hyperparameters of these algorithms are reported in the Appendix A.

Figure 4 presents a comparison study among the approaches mentioned above. For each RL agent, the thick colored lines stand for the episode's total reward, while the transparent surfaces around them correspond to the standard deviation. Moreover, the episode's reward (*score*) is normalized in such a way that 0 stands for an initial invalid action by the robot, $r_{invalid}$ in (8), while 1 correspond to the theoretical maximum reward, which is the r_{bonus} in (8) plus the number of cells.

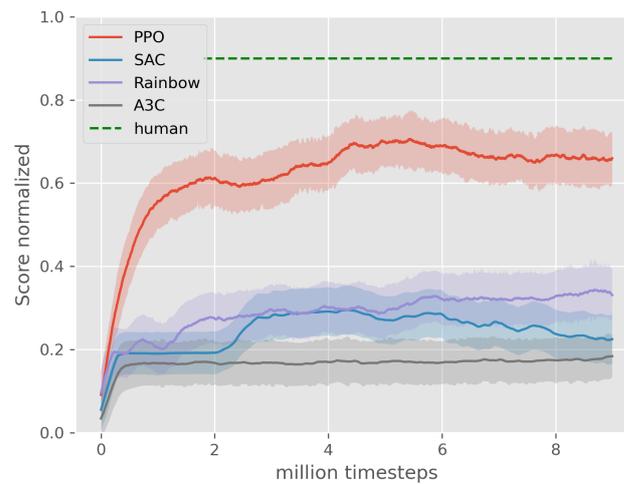


Figure 4. Learning curves for MarsExplorer with randomly chosen obstacles.

To increase the qualitative comprehension of the produced results, the average human-level performance is also introduced. To approximate this value, 10 players were drawn from the pool of *CERTH/ConvCAO* employees to participate in the evaluation process. Each player had an initial warm-up phase of 15 episodes (non-ranked), and after that, they were evaluated on 30 episodes. The average achieved score of the 300 human-controlled experiments is depicted with a green dashed line.

A clear-cut outcome is that the PPO algorithm achieves the highest average episodic reward, reaching an impressive 85.8% of the human-level performance. DQN-Rainbow achieves the second-best performance; however, the average is 50.04% and 42.73% of the PPO and human-level performance, respectively.

3.3. Multi-Dimensional Difficulty

Having defined the best performing RL algorithm (PPO), now the focus is shifted on producing some preliminary results, related to the difficulty settings of MarsExplorer. As mentioned in the definition section, MarsExplorer allows for setting the elements of difficulty vector independently. More specifically, the difficulty vector comprises 3 elements $[d_t, d_m, d_b]$, where:

- d_t denotes the **topology stochasticity**, which defines the obstacles' placement on the field. The fundamental positions of the obstacles are equally arranged in a 3 columns–3 rows format. The radius of deviation around these fundamental positions is controlled by d_t . As the value of d_t increases, the obstacles' topology has a more unstructured formation. d_t takes values from $\{1, 2, 3\}$ discrete set.
- d_m denotes the **morphology stochasticity**, which defines the obstacles' shape on the field. d_m controls the area that might be occupied from each obstacle. The bigger the value of d_m , the larger the compound areas of obstacles that might appear on the MarsExplorer terrain. d_m takes values from $\{1, 2\}$ discrete set.
- d_b denotes the **bonus rewards**, that are assigned for the completion (r_{bonus}) and failure ($r_{invalid}$) of the mission (8). For this factor only two values are allowed $\{1, 2\}$, that correspond to cases of providing and not-providing the bonus rewards, respectively.

Higher values in the elements of the difficulty vector correspond to less structured behavior in the formation of the obstacles. Thus, a trained agent that has been successfully trained in greater difficulty setups may exhibit increased generalization abilities. Overall, the aggregation of the aforementioned elements' domain generates 12 combinations of difficulty levels. Figure 5 shows the total average return of the evolution of the average episodic reward for each one of the 12 levels during the training of the PPO algorithm. To improve the readability of the graphs, the results are organized into 3 graphs, one for each level of d_t , with 4 plot lines each.

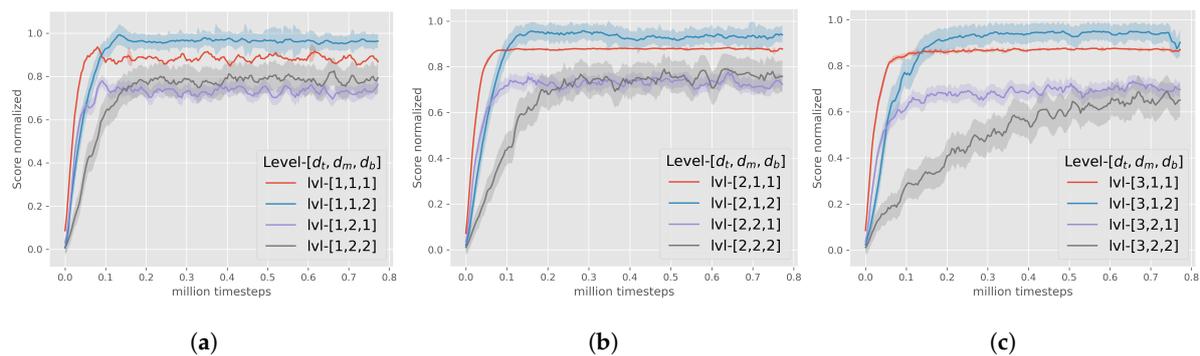


Figure 5. The sensitivity of PPO algorithm learning curves with respect to the different levels of multi-dimensional difficulty vector. (a) Topology stochasticity level $d_t = 1$; (b) Topology stochasticity level $d_t = 2$; (c) Topology stochasticity level $d_t = 3$.

A study on the learning curves reveals that d_m has the largest effect on the learned policy. Blue and red lines (cases where $d_m = 1$), in all three figures, demonstrate a similar convergence rate and also the highest-performance policies. However, a serious degradation in the results is observed in purple and gray lines ($d_m = 2$). As it was expected, when $d_m = 2$ and also $d_t = 3$ (purple and gray lines in Figure 5c), the final achieved performance reached only a little bit above 0.6 in the normalized scale. It seems that d_b does not affect the overall performance much, at least until this vector of difficulty, apart from the convergence rate depicted in the gray line of Figure 5c.

3.4. Learned Policy Evaluation

This section is devoted to the characteristics of the learned policy from the PPO algorithm. For each of the 12 levels of difficulty defined in the previous section, the best PPO policy was extracted and evaluated in a series of 100 experiments with randomly (controlled by the difficulty setting) generated obstacles. Figure 6 presents one heat map for each difficulty level. The blue colormap corresponds to the frequency of the robot visiting a specific cell of the terrain. The green colormap corresponds to the number of detected obstacles in each position during the robot's exploration.

A critical remark is that, for each scenario, the arrangement of discovered obstacles matches the drawn distribution as described in the previous subsection, implying that the learned policy does not have any “blind spots”.

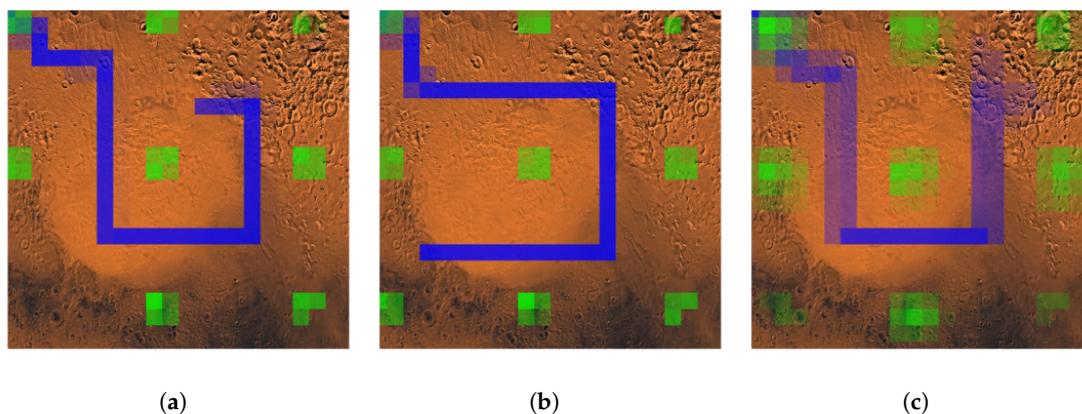


Figure 6. Cont.

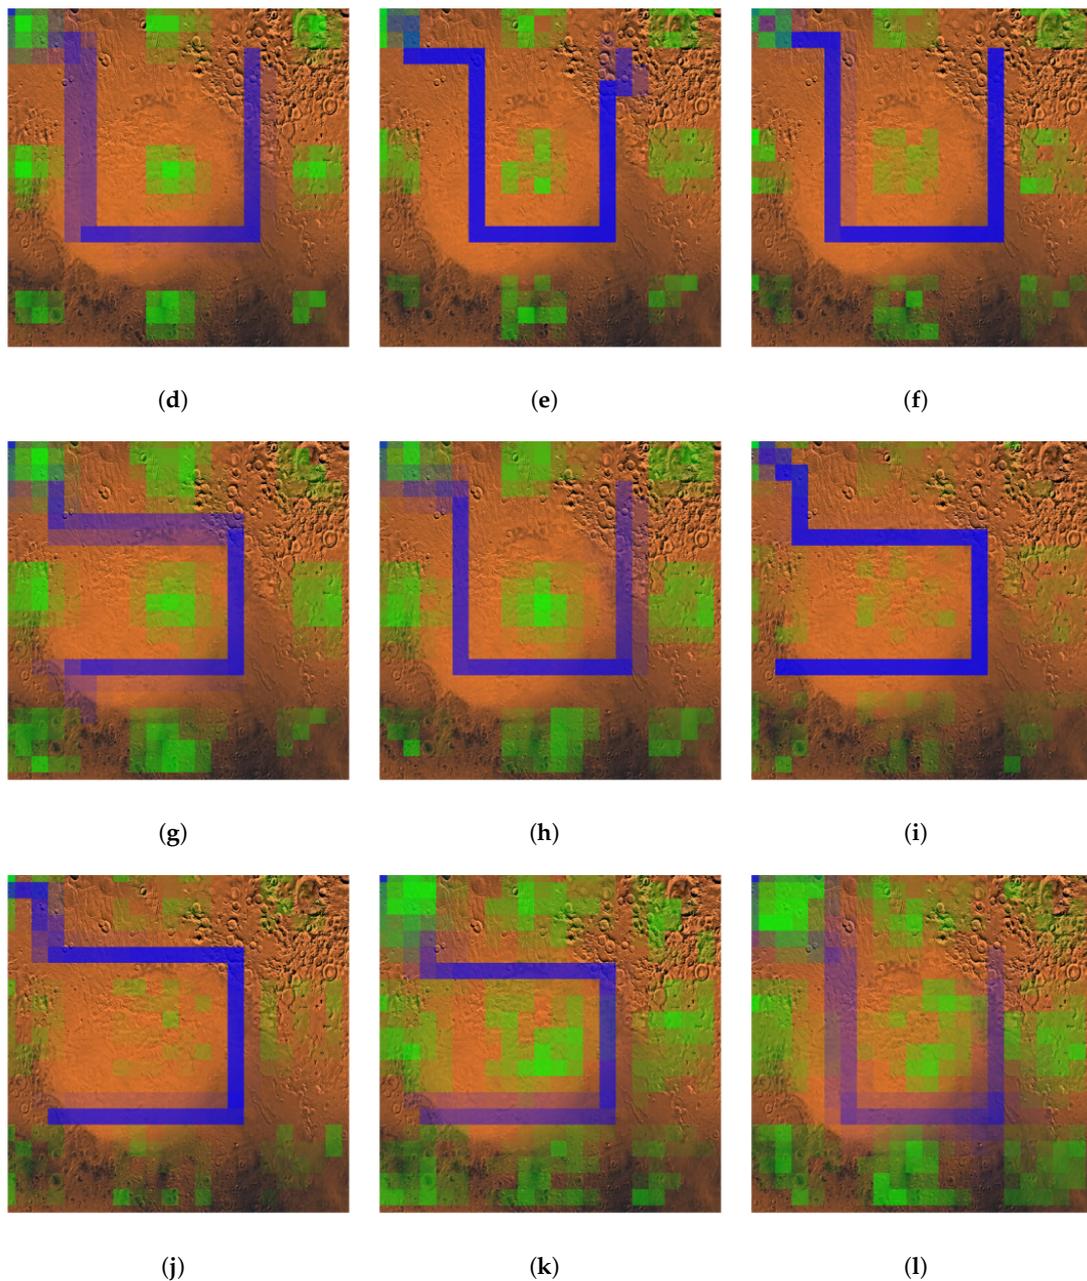


Figure 6. Heatmap of the evaluation results of the learned PPO policy. For each of the 12 difficulty levels, 100 experiments were performed, with the randomness in obstacles' formation as imposed by the corresponding level. The blue colormap corresponds to the frequency of cell visitations by the RL agent, while the green colormap corresponds to the location of the encountered obstacles for all the evaluations. (a) level-[1,1,1]; (b) level-[1,1,2]; (c) level-[1,2,1]; (d) level-[1,2,2]; (e) level-[2,1,1]; (f) level-[2,1,2]; (g) level-[2,2,1]; (h) level-[2,2,2]; (i) level-[3,1,1]; (j) level-[3,1,2]; (k) level-[3,2,1]; (l) level-[3,2,2].

Examining the heatmap of the trajectories in each scenario, it is crystal clear that the same family of trajectories has been generated in all cases and with great confidence. The important conclusion here is that this pattern is the first order of the Hilbert curve that has been utilized extensively in the space-filling domain (e.g., [36,37]). This highlights that such a pattern has not been imported to the simulator or rewarded when achieved from the RL algorithm; however, the algorithm learned that this is the most effective strategy by interacting with the environment.

It would be an omission not to mention the learned policy's ability to adapt to changes in the obstacles' distribution and, ultimately, find the most efficient obstacle-free route. This

trait can be observed more clearly in Figure 6k,l, where the policy needed to be extremely dexterous and delicate to avoid obstacles' encounters.

3.5. Comparison with Frontier-Based Methodologies for Varying Terrain Sizes

The analysis is concluded with a scalability study and comparison to non-learning methodologies. Two terrains with sizes $[42 \times 42]$ and $[84 \times 84]$ were used. The difficulty level was set to $[d_t, d_m, d_b] = [2, 2, 1]$, while 100 experiments were conducted for each scenario. Utility and cost-based frontier cell exploration methodologies [7] were enabled for positioning the achieved PPO policy in the context of non-learning approaches. In these frontier-based approaches, the exploration policy is divided into two categories based on the metric to be optimized:

- *Cost*: the next action is chosen based on the distance from the nearest frontier cell.
- *Utility*: the decision-making is governed by frequently updated information potential field.

Figure 7 summarizes the result of such evaluation study by presenting the average exploration time for each algorithm (PPO, cost frontier-based, utility frontier-based) over 100 procedurally generated runs. A direct outcome is that the learning-based approach requires the robot to travel less distance to explore the same percentage of terrain as the non-learning approaches. The final remark is devoted to the "knee" that can be observed in almost all the final stages of the non-learning approaches. Such behavior is attributed to having several distant sub-parts of the terrain unexplored, the exploration of which requires this extra effort. On the contrary, the learning-based approach (PPO) seems to handle this situation quite well, not leaving these expensive-to-revisit regions along its exploration path.

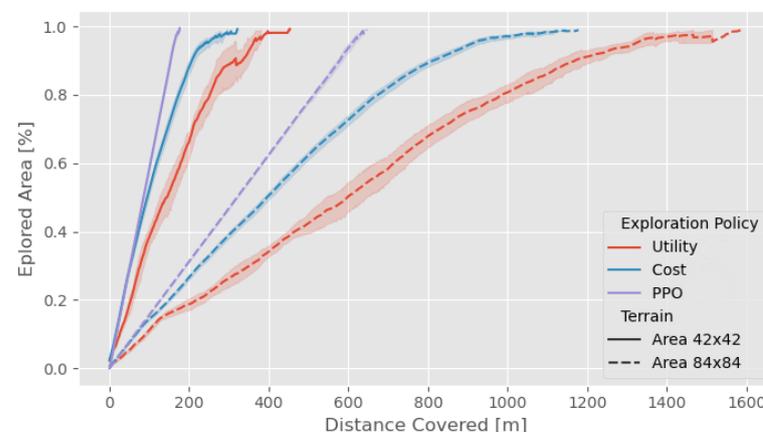


Figure 7. Comparison between 3 exploration methodologies, depicting the average and standard deviation over 100 procedurally generated environments. Red and blue colors correspond to the non-learning approaches, while purple corresponds to the performance of the PPO trained policy. Line type (solid or dashed) denotes the terrain size (42^2 or 84^2).

4. Conclusions

A new openai-gym environment called MarsExplorer that bridges the gap between reinforcement learning and the real-life exploration/coverage in the robotics domain is presented. The environment transforms the well-known robotics problem of exploration/coverage of a completely unknown region into a reinforcement learning setup that can be tackled by a wide range of off-the-shelf, model-free RL algorithms. An essential feature of the whole solution is that trained policies can be straightforwardly applied to real-life robotic platforms without being trained/tuned to the robot's dynamics. To achieve that, the same level of information abstraction between the robotic system and the MarsExplorer is required. A detailed experimental evaluation was also conducted and presented. 4 state-of-the-art RL algorithms, namely A3C, PPO, Rainbow, and SAC, were

evaluated in a challenging version of MarsExplorer, and their training results were also compared with the human-level performance for the task at hand. The PPO algorithm achieved the best score of 85.8%, which was the same as the human-level performance. Then, the PPO algorithm was utilized to study the effect of the multi-dimensional difficulty vector changes in the overall performance. The visualization of the paths for all these difficulty levels revealed a rather important trait. The PPO learned policy has learned to perform a Hilbert curve with the extra ability to avoid any encountered obstacle. Lastly, a scalability study clearly indicates the ability of RL approaches to be extended in larger terrains, where the achieved performance is validated with non-learning, frontier-based explorations strategies.

Author Contributions: Conceptualization, D.I.K., A.C.K., A.A.A. and E.B.K.; methodology D.I.K. and A.C.K.; software, D.I.K.; validation, D.I.K., A.C.K., A.A.A. and E.B.K.; formal analysis, A.C.K. and E.B.K.; investigation, D.I.K., A.C.K. and A.A.A.; resources, E.B.K.; data curation, D.I.K.; writing—original draft preparation, A.C.K. and A.A.A.; writing—review and editing, A.C.K., A.A.A. and E.B.K.; visualization, D.I.K.; supervision, E.B.K.; project administration, A.C.K. and E.B.K.; funding acquisition, E.B.K. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by European Commission Of European Union’s Horizon 2020 research and innovation programme under grant number 833464 (CREST). Also, we gratefully acknowledge the support of NVIDIA Corporation with the donation of GPUs used for this research.

Conflicts of Interest: The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript, or in the decision to publish the results

Appendix A

Table A1. PPO Hyperparameters.

Parameter	Value	Comments
γ	0.95	Discount factor of the MDP
λ	5×10^{-5}	Learning rate
Critic	True	Used a critic as a baseline
GAE l	0.95	GAE (lambda) parameter
KL coeff	0.2	Initial coefficient for KL divergence
Clip	0.3	PPO clip parameter

Table A2. DQN-Rainbow Hyperparameters.

Parameter	Value	Comments
γ	0.95	Discount factor of the MDP
λ	5×10^{-4}	Learning rate
Noisy Net	True	Used a noisy network
Noisy σ	0.5	initial value of noisy nets
Dueling Net	True	Used dueling DQN
Double dueling	True	Used double DQN
ϵ -greedy	[1.0, 0.02]	Epsilon greedy for exploration.
Buffer size	50,000	Size of the replay buffer
Prioritized Replay	True	Prioritized replay buffer used

Table A3. SAC Hyperparameters.

Parameter	Value	Comments
γ	0.95	Discount factor of the MDP
λ	3×10^{-4}	Learning rate
Twin Q	True	Use two Q-networks
Q hidden	[256, 256]	Hidden layer activation
Policy hidden	[256, 256]	Hidden layer activation
Buffer size	1e6	Size of the replay buffer
Prioritized Replay	True	Prioritized replay buffer used

Table A4. A3C Hyperparameters.

Parameter	Value	Comments
γ	0.95	Discount factor of the MDP
λ	1×10^{-4}	Learning rate
Critic	True	Used a critic as a baseline
GAE	True	General Advantage Estimation
GAE λ	0.99	GAE(lambda) parameter
Value loss	0.5	Value Function Loss coefficient
Entropy coef	0.01	Entropy coefficient

References

- Witze, A.; Mallapaty, S.; Gibney, E. All Aboard to Mars. 2020. Available online: <https://www.nature.com/articles/d41586-020-01861-0> (accessed on 4 November 2021).
- Smith, M.; Craig, D.; Herrmann, N.; Mahoney, E.; Krezel, J.; McIntyre, N.; Goodliff, K. The Artemis Program: An Overview of NASA's Activities to Return Humans to the Moon. In Proceedings of the 2020 IEEE Aerospace Conference, Big Sky, MT, USA, 7–14 March 2020; pp. 1–10.
- Shrestha, R.; Tian, F.P.; Feng, W.; Tan, P.; Vaughan, R. Learned map prediction for enhanced mobile robot exploration. In Proceedings of the 2019 International Conference on Robotics and Automation (ICRA), Montreal, QC, Canada, 20–24 May 2019; pp. 1197–1204.
- Kapoutsis, A.C.; Chatzichristofis, S.A.; Doitsidis, L.; de Sousa, J.B.; Pinto, J.; Braga, J.; Kosmatopoulos, E.B. Real-time adaptive multi-robot exploration with application to underwater map construction. *Auton. Robot.* **2016**, *40*, 987–1015. [[CrossRef](#)]
- Batinovic, A.; Petrovic, T.; Ivanovic, A.; Petric, F.; Bogdan, S. A Multi-Resolution Frontier-Based Planner for Autonomous 3D Exploration. *IEEE Robot. Autom. Lett.* **2021**, *6*, 4528–4535. [[CrossRef](#)]
- Renzaglia, A.; Dibangoye, J.; Le Doze, V.; Simonin, O. Combining Stochastic Optimization and Frontiers for Aerial Multi-Robot Exploration of 3D Terrains. In Proceedings of the 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Macau, China, 3–8 November 2019; pp. 4121–4126.
- Basilico, N.; Amigoni, F. Exploration strategies based on multi-criteria decision making for searching environments in rescue operations. *Auton. Robot.* **2011**, *31*, 401–417. [[CrossRef](#)]
- Palacios-Gasós, J.M.; Montijano, E.; Sagüés, C.; Llorente, S. Distributed coverage estimation and control for multirobot persistent tasks. *IEEE Trans. Robot.* **2016**, *32*, 1444–1460. [[CrossRef](#)]
- Koutras, D.I.; Kapoutsis, A.C.; Kosmatopoulos, E.B. Autonomous and cooperative design of the monitor positions for a team of UAVs to maximize the quantity and quality of detected objects. *IEEE Robot. Autom. Lett.* **2020**, *5*, 4986–4993. [[CrossRef](#)]
- Popov, I.; Heess, N.; Lillicrap, T.; Hafner, R.; Barth-Maron, G.; Vecerik, M.; Lampe, T.; Tassa, Y.; Erez, T.; Riedmiller, M. Data-efficient deep reinforcement learning for dexterous manipulation. *arXiv* **2017**, arXiv:1704.03073.
- Silver, D.; Huang, A.; Maddison, C.J.; Guez, A.; Sifre, L.; Van Den Driessche, G.; Schrittwieser, J.; Antonoglou, I.; Panneershelvam, V.; Lanctot, M.; et al. Mastering the game of Go with deep neural networks and tree search. *Nature* **2016**, *529*, 484–489. [[CrossRef](#)]
- 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*, 529–533. [[CrossRef](#)] [[PubMed](#)]
- Baker, B.; Kanitscheider, I.; Markov, T.; Wu, Y.; Powell, G.; McGrew, B.; Mordatch, I. Emergent tool use from multi-agent autocurricula. *arXiv* **2019**, arXiv:1909.07528.
- Zhu, H.; Yu, J.; Gupta, A.; Shah, D.; Hartikainen, K.; Singh, A.; Kumar, V.; Levine, S. The Ingredients of Real World Robotic Reinforcement Learning. *arXiv* **2020**, arXiv:2004.12570.
- Brockman, G.; Cheung, V.; Pettersson, L.; Schneider, J.; Schulman, J.; Tang, J.; Zaremba, W. Openai gym. *arXiv* **2016**, arXiv:1606.01540.

16. Dhariwal, P.; Hesse, C.; Klimov, O.; Nichol, A.; Plappert, M.; Radford, A.; Schulman, J.; Sidor, S.; Wu, Y.; Zhokhov, P. OpenAI Baselines. 2017. Available online: <https://github.com/openai/baselines> (accessed on 4 November 2021).
17. Liang, E.; Liaw, R.; Nishihara, R.; Moritz, P.; Fox, R.; Goldberg, K.; Gonzalez, J.; Jordan, M.; Stoica, I. RLLib: Abstractions for distributed reinforcement learning. In Proceedings of the International Conference on Machine Learning, Stockholm, Sweden, 10–15 July 2018; pp. 3053–3062.
18. Lei, X.; Zhang, Z.; Dong, P. Dynamic path planning of unknown environment based on deep reinforcement learning. *J. Robot.* **2018**, *2018*, 5781591. [[CrossRef](#)]
19. Wen, S.; Zhao, Y.; Yuan, X.; Wang, Z.; Zhang, D.; Manfredi, L. Path planning for active SLAM based on deep reinforcement learning under unknown environments. *Intell. Serv. Robot.* **2020**, *13*, 263–272. [[CrossRef](#)]
20. Zhang, K.; Niroui, F.; Ficocelli, M.; Nejat, G. Robot navigation of environments with unknown rough terrain using deep reinforcement learning. In Proceedings of the 2018 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR), Philadelphia, PA, USA, 6–8 August 2018; pp. 1–7.
21. Niroui, F.; Zhang, K.; Kashino, Z.; Nejat, G. Deep reinforcement learning robot for search and rescue applications: Exploration in unknown cluttered environments. *IEEE Robot. Autom. Lett.* **2019**, *4*, 610–617. [[CrossRef](#)]
22. Luis, S.Y.; Reina, D.G.; Marín, S.L.T. A Multiagent Deep Reinforcement Learning Approach for Path Planning in Autonomous Surface Vehicles: The YpacaraC-Lake Patrolling Case. *IEEE Access* **2021**, *9*, 17084–17099. [[CrossRef](#)]
23. Mnih, V.; Badia, A.P.; Mirza, M.; Graves, A.; Lillicrap, T.; Harley, T.; Silver, D.; Kavukcuoglu, K. Asynchronous methods for deep reinforcement learning. In Proceedings of the International Conference on Machine Learning, New York, NY, USA, 20–22 June 2016; pp. 1928–1937.
24. Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; Klimov, O. Proximal policy optimization algorithms. *arXiv* **2017**, arXiv:1707.06347.
25. Hessel, M.; Modayil, J.; Van Hasselt, H.; Schaul, T.; Ostrovski, G.; Dabney, W.; Horgan, D.; Piot, B.; Azar, M.; Silver, D. Rainbow: Combining improvements in deep reinforcement learning. In Proceedings of the AAAI Conference on Artificial Intelligence, Virtually, 2–9 February 2018; Volume 32.
26. Haarnoja, T.; Zhou, A.; Abbeel, P.; Levine, S. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *arXiv* **2018**, arXiv:1801.01290.
27. Zamora, I.; Lopez, N.G.; Vilches, V.M.; Cordero, A.H. Extending the OpenAI Gym for robotics: A toolkit for reinforcement learning using ROS and Gazebo. *arXiv* **2016**, arXiv:1608.05742.
28. Lopez, N.G.; Nuin, Y.L.E.; Moral, E.B.; Juan, L.U.S.; Rueda, A.S.; Vilches, V.M.; Kojcev, R. gym-gazebo2, a toolkit for reinforcement learning using ROS 2 and Gazebo. *arXiv* **2019**, arXiv:1903.06278.
29. Kapoutsis, A.C.; Chatzichristofis, S.A.; Kosmatopoulos, E.B. A distributed, plug-n-play algorithm for multi-robot applications with a priori non-computable objective functions. *Int. J. Robot. Res.* **2019**, *38*, 813–832. [[CrossRef](#)]
30. Burgard, W.; Moors, M.; Fox, D.; Simmons, R.; Thrun, S. Collaborative multi-robot exploration. In Proceedings of the Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings (Cat. No. 00CH37065), San Francisco, CA, USA, 24–28 April 2000; Volume 1, pp. 476–481.
31. Gray, L.; New, A. A mathematician looks at Wolfram’s new kind of science. *Not. Am. Math. Soc.* **2003**, *50*, 200–211.
32. Cobbe, K.; Hesse, C.; Hilton, J.; Schulman, J. Leveraging procedural generation to benchmark reinforcement learning. In Proceedings of the International Conference on Machine Learning, Virtual, 13–18 July 2020; pp. 2048–2056.
33. Yin, H.; Chen, J.; Pan, S.J.; Tschitschek, S. Sequential Generative Exploration Model for Partially Observable Reinforcement Learning. In Proceedings of the AAAI Conference on Artificial Intelligence, Vancouver, BC, Canada, 2–9 February 2021; Volume 35, pp. 10700–10708.
34. Liang, E.; Liaw, R.; Nishihara, R.; Moritz, P.; Fox, R.; Gonzalez, J.; Goldberg, K.; Stoica, I. Ray rllib: A composable and scalable reinforcement learning library. *arXiv* **2017**, *85*, arXiv:1712.09381.
35. Bellemare, M.G.; Naddaf, Y.; Veness, J.; Bowling, M. The Arcade Learning Environment: An Evaluation Platform for General Agents. *J. Artif. Intell. Res.* **2013**, *47*, 253–279. [[CrossRef](#)]
36. Kapoutsis, A.C.; Chatzichristofis, S.A.; Kosmatopoulos, E.B. DARP: Divide areas algorithm for optimal multi-robot coverage path planning. *J. Intell. Robot. Syst.* **2017**, *86*, 663–680. [[CrossRef](#)]
37. Sadat, S.A.; Wawerla, J.; Vaughan, R. Fractal trajectories for online non-uniform aerial coverage. In Proceedings of the 2015 IEEE international conference on robotics and automation (ICRA), Seattle, WA, USA, 26–30 May 2015; pp. 2971–2976.