

Improving Generalization in Aerial and Terrestrial Mobile Robots Control Through Delayed Policy Learning

Ricardo B. Grando^{*12}, Raul Steinmetz^{*3}, Victor A. Kich⁴, Alisson H. Kolling², Pablo M. Furik¹, Junior C. de Jesus², Bruna V. Guterres¹, Daniel T. Gamarra³, Rodrigo S. Guerra², Paulo L. J. Drews-Jr²
¹Universidad Tecnologica del Uruguay - UTEC, ²Universidade Federal de Rio Grande - FURG,
³Universidade Federal de Santa Maria - UFSM, ⁴University of Tsukuba

Abstract—Deep Reinforcement Learning (DRL) has emerged as a promising approach to enhance motion control and decision-making through a wide range of robotic applications. While prior research has demonstrated the efficacy of DRL algorithms in facilitating autonomous mapless navigation for aerial and terrestrial mobile robots, these methods often grapple with poor generalization when faced with unknown tasks and environments. This paper explores the impact of the Delayed Policy Updates (DPU) technique on fostering generalization to new situations and bolstering the overall performance of agents. Our analysis of DPU for aerial and terrestrial mobile robots in four simulated environments reveals that this technique significantly curtails the lack of generalization and accelerates the learning process for agents, enhancing their efficiency across diverse tasks and unknown scenarios.

Index Terms—Deep Reinforcement Learning, Delayed Policy Updates, Mobile Robots, Learning Generalization.

I. INTRODUCTION

Reinforcement Learning (RL) and Deep Reinforcement Learning (DRL) have enhanced autonomous decision-making in machines. RL entails an agent learning to achieve a goal through actions in an environment, guided by feedback from rewards or penalties [1], [2]. By integrating deep neural networks with RL, DRL advances these capabilities, enabling the approximation of optimal policies for handling complex inputs and learning from large amounts of unstructured data. These advancements have spurred significant progress in gaming [3] and robotics [4], [5]. However, a notable challenge in model-free DRL is agents’ susceptibility to poor generalization in unfamiliar scenarios, limiting their task generalization effectiveness [6]–[8].

Generalization, the capacity of neural networks to apply learned knowledge to unseen data, is essential for developing neural networks, representing a primary objective in deep learning [9]. Insufficient generalization limits a model’s usefulness, resulting in subpar performance on data not encountered during training. Various techniques have been proposed to improve generalization, such as dropout [10], [11], normalization [12], [13], and ensemble methods [14], enhancing the robustness and effectiveness of the model.

Improving generalization in DRL is critical for ensuring robust performance across various unknown scenarios. This generalization is significant, as DRL agents often excel in their training environments, optimizing their behavior to such

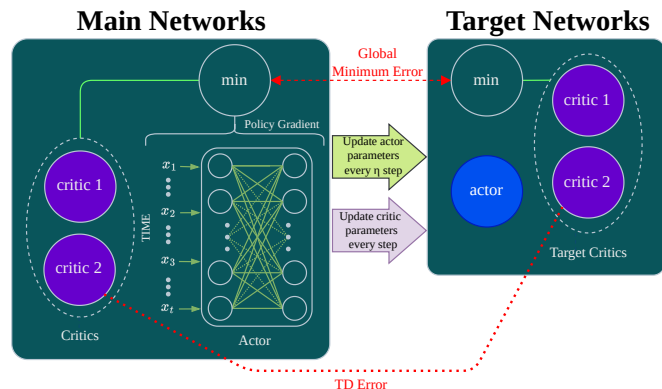


Fig. 1: Architecture of the Twin Delayed Deep Deterministic Policy Gradient (TD3) algorithm and its mechanisms.

an extent that they appear to overfit these specific scenarios. However, when faced with changes in the environment, as small as they may be (a new object to avoid in autonomous navigation, for example), their performance can drastically decrease, limiting task precision and their ability to complete it successfully [15]. Numerous methods have been developed to enhance generalization in deep reinforcement learning, including soft data augmentation [16], mixture regularization [17], and others.

Delayed Policy Updates (DPU) [18] is one of these techniques, and it aims to bolster learning stability by postponing the agent’s policy updates, allowing for learning from a steadier experience pool over extended periods. This study analyzes DPU’s effect on RL agents’ generalization in navigation tasks for mobile robots. Robots must adapt to new tasks or environments in this critical area while retaining learned knowledge. We demonstrate through experiments with varying DPU settings in different environments that DPU significantly enhances generalization. Our research evaluates DPU’s role in addressing generalization challenges within RL for both aerial and terrestrial robot navigation.

The present work specifically contributes to the following topics:

- We train three identical agents to navigate autonomously as aerial robots and three to navigate autonomously as terrestrial robots, each configured with different DPU values. Our analysis reveals that a greater delay contributes positively to the success rate for a navigation task assigned to an aerial or mobile robot.

* Ricardo B. Grando and Raul Steinmetz contributed equally.

- In a subsequent experiment within an unfamiliar scenario, we confirm the agents’ training effectiveness. The results demonstrate that an increased delay greatly enhances the agent’s generalization capabilities across different environments in both applications.

The work has the following structure: related works are addressed in Section II. Our approaches and tools used can be seen in Section III. The results achieved are presented in the Section IV. Section V shows our overall analysis of the results, and Section VI presents the perspective for future works.

II. RELATED WORK

Throughout the years, a vast number of techniques have been developed to enhance the generalization capabilities of DRL architectures. Karl Cobbe *et al.* [19] delved into the issue of overfitting within deep reinforcement learning, emphasizing the difficulty of achieving generalization across procedurally generated environments. By introducing CoinRun, a benchmark for RL generalization, they revealed that agents often exhibit overfitting when trained on extensive datasets. Their findings suggested that adopting deeper convolutional structures and incorporating methods commonly used in supervised learning, such as ℓ_2 regularization, dropout, data augmentation, and batch normalization, can significantly bolster generalization in RL contexts.

Building on this understanding, Farebrother *et al.* [20] also tackled the challenge of overfitting in DRL, focusing on DQN models applied to Atari 2600 games. They underscored the innovative application of game modes to assess the generalization capacity of DQN agents, showcasing how the integration of regularization techniques like dropout and ℓ_2 regularization can amplify the adaptability of DQN models across varied game scenarios. This work provides additional insight into addressing the generalization dilemma within RL environments.

In a complementary effort, Lee *et al.* [21] proposed a novel method to refine the generalization of DRL agents. Their technique, which utilizes randomized (convolutional) neural networks to alter input observations, encourages the formation of robust features consistent across a spectrum of randomized environments. Moreover, they introduced a Monte Carlo approximation-based inference strategy to reduce the variance associated with this input perturbation, further contributing to the toolkit for enhancing RL generalization.

Extending the exploration of generalization in DRL, Justesen *et al.* [22] investigated the efficacy of using procedurally generated levels during training. Their research demonstrated that such an approach facilitates the generalization of RL models to new, similarly distributed levels and improves efficiency by dynamically adjusting the difficulty of levels in response to the agent’s performance, highlighting a strategic method to optimize training processes.

Further advancing the field, Raileanu *et al.* [23] introduced automated techniques for identifying effective data augmentations across various tasks, markedly improving agent

generalization. This innovation circumvents the need for expert knowledge in choosing task-specific augmentations, addressing a significant challenge in scaling RL applications. Additionally, introducing two novel regularization terms for policy and value functions integrates data augmentation into actor-critic algorithms at a theoretical level, marking a significant step forward in the quest for enhanced RL generalization.

The Twin Delayed Deep Deterministic Policy Gradient (TD3) algorithm, introduced by Fujimoto *et al.* [18], advances the Deep Deterministic Policy Gradient (DDPG) framework [24] for improved performance in continuous control tasks within RL. By employing dual critics (critic 1 and critic 2) for Q-value estimation, TD3 strategically addresses the overestimation bias common in Q-learning algorithms. A key innovation of TD3 is the asynchronous update of policy (actor) and Q-value estimators (critics), introducing a deliberate delay in policy updates to mitigate the risk of overfitting premature Q-value estimates. This methodological refinement has seen application across various domains, notably in enhancing mobile robotics tasks [25]–[30], these works focus primarily on algorithmic or task-specific improvements rather than on the learning impact of the DPU technique itself.

This study aims to fill a gap in the literature by analyzing the effects of DPU within the context of continuous control in RL, particularly concerning learning generalization. We delve into how delayed policy updates influence learning processes and an agent’s capability to execute specific tasks, focusing on navigation tasks for aerial and terrestrial mobile robots. To the best of our knowledge, it consists of a novel application not yet explored in existing research.

III. METHODOLOGY

In this section, we present our DRL approach. We detail the network structure and the implementation of the TD3 algorithm. We also present the simulation details and the task used to perform the analysis.

A. Twin Delayed Deep Deterministic Policy Gradient

The Deep Deterministic Policy Gradient (DDPG) algorithm, introduced by Lillicrap *et al.* [24], marked a significant advancement in the field of reinforcement learning for continuous control tasks, particularly in robotics. Despite its achievements, DDPG is prone to overestimating Q-values, which can destabilize the learning policy and lead to suboptimal performance.

TD3 algorithm introduces several key improvements over its predecessor. By employing clipped double-Q learning, TD3 mitigates the overestimation bias by using two critic networks to calculate the Bellman error loss functions, taking the minimum value of their Q-value estimations. Furthermore, it incorporates target policy smoothing by adding Gaussian noise to the target action, enhancing exploration, and preventing policy exploitation. A distinctive feature of TD3 is its delayed policy updates strategy, wherein the policy network’s updates are less frequent than those of the critic networks. The delay




Fig. 2: Training and evaluation scenarios for aerial mobile robots, respectively.

parameter η represents the number of environment steps that should be executed before updating the target actor network for the learning calculations. This delay reduces the risk of policy overfitting, promoting more stable learning and better generalization across a broader range of tasks, as shown in Section IV. The implementation of the TD3 algorithm in this work incorporates these advancements, demonstrating its efficacy in complex control scenarios.

We use a fully connected Artificial Neural Network (ANN) as actor-network and denote them by ϕ and its copy ϕ' as actor target. The actor target chooses the action a' based on the state s' , and we add Ornstein-Uhlenbeck [31] noise to it. The double critic targets take the tuple (s', a') and return two Q-values as output. The minimum of the two target Q-values is considered the approximated value return. The loss is calculated with the Mean Squared Error of the approximate value from the target networks and the value from the critic networks. We use Adaptive Moment Estimation (Adam) to minimize the loss.

Three variants of the TD3 algorithm were used in the methodology. The first sets the parameter $\eta = 2$, the second uses $\eta = 4$, and the last uses $\eta = 8$. Each agent was trained and evaluated independently.

B. Aerial Robot Simulated Environments and Task Description

Our research utilized aerial mobile robots in simulations developed by Grando *et al.* [32]–[34], and Jesus *et al.* [35], utilizing the Gazebo simulator in conjunction with ROS. The RotorS framework [36] served as the foundation, facilitating the simulation of aerial vehicles with capabilities for various command levels, including angular rates, attitude, and location control, as well as wind simulation using the Ornstein-Uhlenbeck process.

We utilized two 10-meter square scenarios: the first features four symmetrically placed obstacles around the center, and the

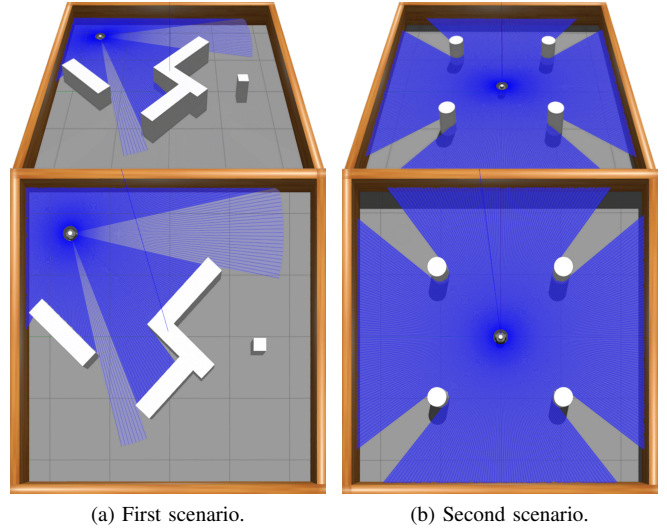


Fig. 3: Training and evaluation scenarios for terrestrial mobile robots, respectively.

second contains a single central obstacle. The first scenario was used for training purposes, while the second evaluated the model’s generalization capabilities and retention of learned behaviors. Fig. 2 depicts both scenarios.

In each scenario, the robot navigated from a fixed start point to a goal, requiring obstacle detection and avoidance facilitated by LIDAR data. During the training phase, the target goal was set aleatory inside the environment limits, while during the evaluation phase, four fixed positions were defined as targets. The environment state provided to the agent encompasses 26 dimensions: 20 from distance sensors, three from previous actions, and three detailing the target’s position (including relative position and angles in the x-y and z-range planes). Distance data was collected via LIDAR, spaced at 13.5-degree intervals across a 270-degree field.

Action scales range from 0 to 0.25 m/s for linear velocity, -0.25 to 0.25 m/s for altitude adjustments, and -0.25 to 0.25 radians for angular changes (Δ yaw), enabling effective navigation and obstacle avoidance.

A binary reward function was employed, offering a positive reward for successful navigation and a negative reward for collisions or reaching the 500-step episode limit:

$$r(s_t, a_t) = \begin{cases} r_{arrive} & \text{if } d_t < c_d \\ r_{collide} & \text{if } \min_x < c_o \parallel ep = 500, \end{cases} \quad (1)$$

where $r_{arrive} = 200$ is given to the agent when completed the task successfully, while $r_{collide} = -20$ penalizes collisions or timeouts, with both c_d and c_o set to 0.5 meters. The Reward and maximum episode values were defined empirically, and the model used in the evaluation was the one at step 5000000.

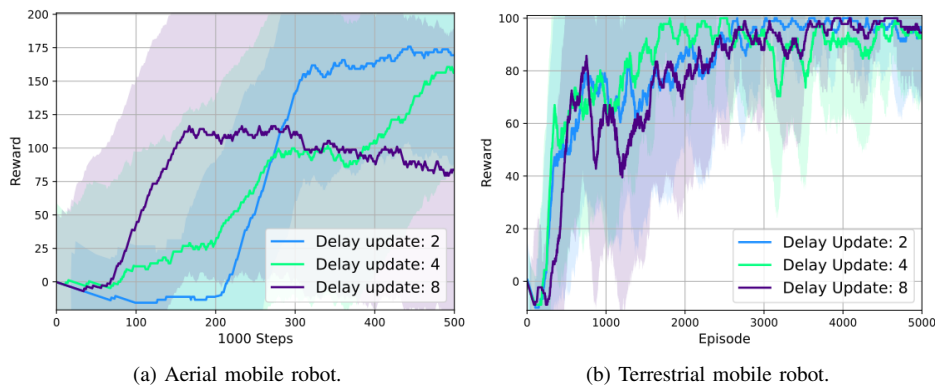


Fig. 4: Reward moving average over 5000000 steps of training (aerial) 5000 episodes of training (terrestrial).

C. Terrestrial Robot Simulated Environments and Task Description

Our experiments were conducted using a terrestrial mobile robot within simulated environments, utilizing the Gazebo simulator in conjunction with ROS2 and the Turtlebot3 robot.

We utilized two 5-meter square scenarios: the first presented three obstacles that constrained the robot’s movement through the center, while the second scenario arranged four obstacles symmetrically around the center. The first scenario was used for training purposes, while the second evaluated the model’s generalization capabilities and retention of learned behaviors. Fig. 3 depicts both scenarios.

Like the aerial simulations, the terrestrial robot’s task involved navigating from a start to a goal, relying on LIDAR for obstacle detection. The environment state fed to the agent included 14 dimensions: Ten sensor distance readings, distance and angle to target, and the robot’s current linear and angular velocity. Action scales range from 0 to 0.25 m/s for linear velocity and -0.25 to 0.25 m/s for angular velocity.

The same binary reward function was employed with a 250-step episode limit. $r_{arrive} = 100$ is given to the agent when completed the task successfully, while $r_{collide} = -10$ penalizes collisions or timeouts, with c_d set to 0.3 meters and c_o set to 0.19 meters. In this case, the Reward and maximum episode values were defined empirically, and the model used in the evaluation was the one at episode 5000.

IV. RESULTS

Regarding the aerial task, the reward trend over 500 episodes, depicted in Fig. 4a, illustrates that agents with higher DPU values learn more rapidly and attain a stabilized reward pattern sooner. This early stabilization suggests that higher DPU values enable agents to balance exploration with exploitation efficiently. However, in the known scenario, they may not reach the maximal reward potential obtained by agents with lower DPU values that exhibit more specialized behavior.

In the context of terrestrial robot navigation, the observed reward trend exhibited a high degree of similarity across all values of DPU, with each agent successfully attaining

TABLE I: Aerial mobile robot metrics for different updates delay (η) in known scenarios.

η	Success Rate (%)	ER Mean	ER Std.	ET Mean	ET Std.
2	82.98%	162.55	83.12	33.53	12.86
4	100.00%	200.00	0.00	27.25	5.90
8	100.00%	200.00	0.00	22.89	1.64

the maximum reward and sustaining knowledge retention, as depicted in Fig. 4b. Remarkably, towards the conclusion of the experiments, we observed that the systems characterized by $\eta = 2$ and $\eta = 4$ exhibited significant fluctuations in reward outcomes, demonstrating frequent failures. This phenomenon is further illustrated by the moving average, calculated over a reduced interval, which reveals the green (representing $\eta = 4$) and blue (representing $\eta = 2$) lines exhibiting pronounced oscillations in performance towards the latter stages of the study in stark contrast to the purple line (representing $\eta = 8$), which maintains a consistent level of stability.

To test agent performance under various DPU conditions, we employ metrics such as Success Rate, Episode Reward (ER) Mean, and Standard Deviation (Std.), along with Episode Time (ET) Mean and Std. The compiled outcomes for scenarios both recognized and unfamiliar are presented in Tables I and II for aerial tasks, as well as in Tables III and IV for terrestrial tasks. Data analysis reveals that elevated DPU values enhance performance in known environments while significantly boosting performance in unknown settings (suggesting that lower DPUs may tend to overfit). This fact is exemplified by the observation that with $\eta = 2$, performance in unfamiliar environments decreases dramatically, from 82.98% to 16.33% in aerial tasks and from 93.94% to 70.71% in terrestrial tasks. Conversely, a DPU of $\eta = 8$ demonstrates a minor decline from 100% to 98.67% in aerial scenarios and from 99% to 85% in terrestrial scenarios, thereby illustrating superior generalization capabilities.

Analysis of the aerial agent trajectories in the known and unknown environments (Fig. 5) provides valuable insights into the agents’ navigation strategies. In the known scenario, trajectories for higher DPU values ($\eta = 4$ and $\eta = 8$) reveal the more direct and efficient path to the goal, consistent with




Fig. 5: Trajectories achieved by the drone controlled by TD3 with different update delays (η) in known and unknown scenarios.

their 100%

When examining trajectories in the unknown environment, the higher DPU agent ($\eta = 8$) demonstrates remarkable adaptability, maintaining a clear and concise path to the goal, as illustrated in Fig. 5f, which is in line with its high success rate of 98.67%. Conversely, the agent with a lower DPU value ($\eta = 2$) shows significant difficulty in navigating the new environment, with erratic movements and a marked decrease in success rate to 16.33%, further emphasizing the importance of DPU in unfamiliar settings and the fact that low delays cause dramatic lack of generalization.

V. DISCUSSION

The experimental findings highlight DPU’s impacts on reinforcement learning agents’ learning dynamics and adaptability. Specifically, a higher DPU value correlates with accelerated learning in initial training phases and enhanced generalization to new scenarios. This fact suggests that delaying policy updates enables agents to develop more robust and flexible strategies that are effective across a variety of environmental conditions.

The superior generalization observed with higher DPU values, particularly the minimal performance degradation when transitioning from known to unknown scenarios, underscores the value of this approach in applications where agents must operate in diverse and dynamically changing environments, as noted in Tables II and IV. It contrasts with the tendency of agents trained with lower DPU values to become highly specialized in their training environments, which can limit their effectiveness in new or altered contexts.

This study’s insights have important implications for designing and training reinforcement learning agents, especially in domains requiring a balance between high performance in familiar settings and the ability to adapt to new situations. Optimizing the DPU parameter emerges as a crucial strategy for achieving this balance, providing a mechanism to tailor

TABLE II: Aerial mobile robot metrics for different updates delay (η) in unknown scenarios.

η	Success Rate (%)	ER Mean	ER Std.	ET Mean	ET Std.
2	16.33%	21.63	79.67	45.43	40.05
4	73.68%	142.10	97.39	22.33	10.12
8	98.67%	197.07	25.40	34.60	21.65

TABLE III: Terrestrial mobile robot metrics for different updates delay (η) in known scenarios.

η	Success Rate (%)	ER Mean	ER Std.
2	94	93.4	26.12
4	97	96.7	18.76
8	99	98.9	10.94

agent behavior to meet the specific challenges of the task at hand.

In conclusion, our work contributes to a deeper understanding of how Delayed Policy Updates influence the development of more adaptable and generalizable reinforcement learning agents, paving the way for future research to explore the optimal configuration of these parameters in various application domains.

VI. CONCLUSIONS

This study highlights the significant role that DPU plays in reinforcement learning, particularly within the ambit of robotic navigation. Our empirical analysis has underscored that DPUs influence RL agents’ learning speed and ability to generalize across diverse environments and optimize their navigation trajectories. A key takeaway from our experiments is the demonstrable advantage of higher DPU values in contexts that demand rapid adaptation from the agents.

Our investigation into the Delayed Policy Updates technique has shown its potential to enhance the continuous control capabilities of RL agents engaged in complex navigation tasks. Notably, we found that an agent configured with a more considerable DPU value ($\eta = 8$) exhibited a negligible drop in success rate when transitioning to an unfamiliar environment, a stark contrast to the performance of the agent with a lower DPU value ($\eta = 2$), which saw a dramatic decrease. These insights are pivotal, as they reveal how subtle adjustments in the learning algorithm’s parameters can yield substantial differences in an agent’s adaptability and competency.

TABLE IV: Terrestrial mobile robot metrics for different updates delay (η) in unknown scenarios.

η	Success Rate (%)	ER Mean	ER Std.
2	71	68.1	49.91
4	72	69.2	49.39
8	85	83.5	39.28

Our future directions will aim to understand other techniques that may interact with DPUs, particularly in addressing task generalization in reinforcement learning. By expanding our scope, we anticipate uncovering richer insights that could lead to more robust and versatile learning algorithms tailored to the multifaceted demands of real-world applications.

ACKNOWLEDGEMENT

This work was partly funded by the Brazilian funding agencies: CAPES, CNPq, PRH-ANP, and FINEP. The authors would like to thank the VersusAI team for their invaluable support. Additionally, the authors acknowledge using generative AI language tools for light editing, such as spelling and grammar corrections, in preparing this manuscript.

REFERENCES

- [1] R. S. Sutton and A. G. Barto, "Reinforcement learning: An introduction," *Robotica*, vol. 17, no. 2, pp. 229–235, 1999.
- [2] M. Morales, *Grokking deep reinforcement learning*. Manning Publications, 2020.
- [3] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski *et al.*, "Human-level control through deep reinforcement learning," *nature*, vol. 518, no. 7540, pp. 529–533, 2015.
- [4] S. Gu, E. Holly, T. P. Lillicrap, and S. Levine, "Deep reinforcement learning for robotic manipulation," *arXiv preprint arXiv:1610.00633*, vol. 1, p. 1, 2016.
- [5] S. Gu, E. Holly, T. Lillicrap, and S. Levine, "Deep reinforcement learning for robotic manipulation with asynchronous off-policy updates," in *2017 IEEE ICRA*. IEEE, 2017, pp. 3389–3396.
- [6] C. Packer, K. Gao, J. Kos, P. Krähenbühl, V. Koltun, and D. Song, "Assessing generalization in deep reinforcement learning," *arXiv preprint arXiv:1810.12282*, 2018.
- [7] R. B. Grando, J. C. de Jesus, V. A. Kich, A. H. Kolling, P. M. Pinheiro, R. S. Guerra, and P. L. Drews-Jr, "Deterministic and stochastic analysis of deep reinforcement learning for low dimensional sensing-based navigation of mobile robots," in *2022 IEEE LARS/SBR*. IEEE, 2022, pp. 193–198.
- [8] M. G. Mateus, R. B. Grando, and P. L. Drews-Jr, "Active perception applied to unmanned aerial vehicles through deep reinforcement learning," in *2022 LARS/SBR*. IEEE, 2022, pp. 1–6.
- [9] R. Novak, Y. Bahri, D. A. Abolafia, J. Pennington, and J. Sohl-Dickstein, "Sensitivity and generalization in neural networks: an empirical study," *arXiv preprint arXiv:1802.08760*, 2018.
- [10] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: a simple way to prevent neural networks from overfitting," *The journal of machine learning research*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [11] S. Wager, S. Wang, and P. S. Liang, "Dropout training as adaptive regularization," *Advances in neural information processing systems*, vol. 26, 2013.
- [12] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *IEEE ICML*. pmlr, 2015, pp. 448–456.
- [13] G. Klambauer, T. Unterthiner, A. Mayr, and S. Hochreiter, "Self-normalizing neural networks," *Advances in neural information processing systems*, vol. 30, 2017.
- [14] Y. Freund and R. E. Schapire, "A decision-theoretic generalization of on-line learning and an application to boosting," *Journal of computer and system sciences*, vol. 55, no. 1, pp. 119–139, 1997.
- [15] C. Zhang, O. Vinyals, R. Munos, and S. Bengio, "A study on overfitting in deep reinforcement learning," *arXiv preprint arXiv:1804.06893*, 2018.
- [16] N. Hansen and X. Wang, "Generalization in reinforcement learning by soft data augmentation," in *2021 IEEE ICRA*. IEEE, 2021, pp. 13 611–13 617.
- [17] K. Wang, B. Kang, J. Shao, and J. Feng, "Improving generalization in reinforcement learning with mixture regularization," *Advances in Neural Information Processing Systems*, vol. 33, pp. 7968–7978, 2020.
- [18] S. Fujimoto, H. Hoof, and D. Meger, "Addressing function approximation error in actor-critic methods," in *ICML*, 2018, pp. 1587–1596.
- [19] K. Cobbe, O. Klimov, C. Hesse, T. Kim, and J. Schulman, "Quantifying generalization in reinforcement learning," in *IEEE ICML*. PMLR, 2019, pp. 1282–1289.
- [20] J. Farebrother, M. C. Machado, and M. Bowling, "Generalization and regularization in dqn," *arXiv preprint arXiv:1810.00123*, 2018.
- [21] K. Lee, K. Lee, J. Shin, and H. Lee, "Network randomization: A simple technique for generalization in deep reinforcement learning," *arXiv preprint arXiv:1910.05396*, 2019.
- [22] N. Justesen, R. R. Torrado, P. Bontrager, A. Khalifa, J. Togelius, and S. Risi, "Illuminating generalization in deep reinforcement learning through procedural level generation," *arXiv preprint arXiv:1806.10729*, 2018.
- [23] R. Raileanu, M. Goldstein, D. Yarats, I. Kostrikov, and R. Fergus, "Automatic data augmentation for generalization in deep reinforcement learning," *arXiv preprint arXiv:2006.12862*, 2020.
- [24] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, "Continuous control with deep reinforcement learning," in *ICLR*, 2016.
- [25] J. Gao, W. Ye, J. Guo, and Z. Li, "Deep reinforcement learning for indoor mobile robot path planning," *Sensors*, vol. 20, no. 19, p. 5493, 2020.
- [26] R. B. Grando, J. C. de Jesus, V. A. Kich, A. H. Kolling, and P. L. J. Drews-Jr, "Double critic deep reinforcement learning for mapless 3d navigation of unmanned aerial vehicles," *Journal of Intelligent & Robotic Systems*, vol. 104, no. 2, pp. 1–14, 2022.
- [27] R. B. Grando, J. C. de Jesus, V. A. Kich, A. H. Kolling, P. M. Pinheiro, R. S. Guerra, and P. L. Drews-Jr, "Mapless navigation of a hybrid aerial underwater vehicle with deep reinforcement learning through environmental generalization," in *2022 LARS/SBR*. IEEE, 2022, pp. 1–6.
- [28] P. Li, Y. Wang, and Z. Gao, "Path planning of mobile robot based on improved td3 algorithm," in *2022 IEEE ICMA*. IEEE, 2022, pp. 715–720.
- [29] R. B. Grando, J. C. de Jesus, V. A. Kich, A. H. Kolling, R. S. Guerra, and P. L. Drews-Jr, "Doicl: Double critic deep reinforcement learning for mapless navigation of a hybrid aerial underwater vehicle with medium transition," *arXiv preprint arXiv:2308.09811*, 2023.
- [30] A. H. Kolling, V. A. Kich, J. C. de Jesus, A. C. Da Silva, R. B. Grando, P. L. J. Drews-Jr, and D. F. Gamarra, "Parallel distributional prioritized deep reinforcement learning for unmanned aerial vehicles," in *2023 IEEE LARS/SBR*. IEEE, 2023, pp. 95–100.
- [31] J. Nauta, Y. Khaluf, and P. Simoens, "Using the ornstein-uhlenbeck process for random exploration," in *4th International Conference on Complexity, Future Information Systems and Risk (COMPLEXIS2019)*. SciTePress, 2019, pp. 59–66.
- [32] R. B. Grando, J. C. de Jesus, V. A. Kich, A. H. Kolling, N. P. Bortoluzzi, P. M. Pinheiro, A. A. Neto, and P. L. Drews-Jr, "Deep reinforcement learning for mapless navigation of a hybrid aerial underwater vehicle with medium transition," in *2021 IEEE ICRA*. IEEE, 2021, pp. 1088–1094.
- [33] A. C. Horn, P. M. Pinheiro, R. B. Grando, C. B. da Silva, A. A. Neto, and P. L. Drews-Jr, "A novel concept for hybrid unmanned aerial underwater vehicles focused on aquatic performance," in *IEEE LARS/SBR*, 2020, pp. 1–6.
- [34] A. C. Horn, P. M. Pinheiro, C. B. Silva, A. A. Neto, and P. L. Drews-Jr, "A study on configuration of propellers for multicopter-like hybrid aerial-aquatic vehicles," in *19th ICAR*, 2019, pp. 173–178.
- [35] J. C. de Jesus, V. A. Kich, A. H. Kolling, R. B. Grando, R. S. Guerra, and P. L. Drews-Jr, "Depth-cuprl: Depth-imaged contrastive unsupervised prioritized representations in reinforcement learning for mapless navigation of unmanned aerial vehicles," in *2022 IEEE IROS*. IEEE, 2022, pp. 10 579–10 586.
- [36] F. Furrer, M. Burri, M. Achtelik, and R. Siegwart, "Rotors—a modular gazebo mav simulator framework," in *Robot Operating System (ROS)*, 2016, pp. 595–625.