

# ManeuverNet: A Soft Actor-Critic Framework for Precise Maneuvering of Double-Ackermann-Steering Robots with Optimized Reward Functions

Kohio Deflesselle<sup>1\*</sup>, Mélodie Daniel<sup>1\*</sup>, Aly Magassouba<sup>2</sup>, Miguel Aranda<sup>3</sup>, and Olivier Ly<sup>1</sup>

**Abstract**—Autonomous control of double-Ackermann-steering robots is essential in agricultural applications, where robots must execute precise and complex maneuvers within a limited space. Classical methods, such as the Timed Elastic Band (TEB) planner, can address this problem, but they rely on parameter tuning, making them highly sensitive to changes in robot configuration or environment and impractical to deploy without constant recalibration. At the same time, end-to-end deep reinforcement learning (DRL) methods often fail due to unsuitable reward functions for non-holonomic constraints, resulting in sub-optimal policies and poor generalization. To address these challenges, this paper presents ManeuverNet, a DRL framework tailored for double-Ackermann systems, combining Soft Actor-Critic with CrossQ. Furthermore, ManeuverNet introduces four specifically designed reward functions to support maneuver learning. Unlike prior work, ManeuverNet does not depend on expert data or handcrafted guidance. We extensively evaluate ManeuverNet against both state-of-the-art DRL baselines and the TEB planner. Experimental results demonstrate that our framework substantially improves maneuverability and success rates, achieving more than a 40% gain over DRL baselines. Moreover, ManeuverNet effectively mitigates the strong parameter sensitivity observed in the TEB planner. In real-world trials, ManeuverNet achieved up to a 90% increase in maneuvering trajectory efficiency, highlighting its robustness and practical applicability.

## I. INTRODUCTION

Recent advances in deep reinforcement learning (DRL) for robotics have mainly focused on navigation and locomotion tasks for holonomic robots such as omni-wheeled [1], quadruped [2], or biped robots [3]. By holonomic robots, we refer to the ones that can move freely in all directions [4]. In contrast, Ackermann-steering robots, like cars, represent a restrictive type of non-holonomic robots [5]. These robots cannot rotate without moving forward or backward, which reduces their number of degrees of freedom (DOF) [6]. Such constraints significantly limit their maneuverability and impose additional challenges for control [7]. In this work, we specifically focus on controlling a mobile platform with four-wheel steering (4WS), and a double-Ackermann-steering mechanism (see Fig. 1). Controlling such robots presents even more complex challenges compared to single-Ackermann-steering systems. The need to coordinate both




Fig. 1: Maneuver handling (left figure) with 4WS robots in DRL is challenging because it requires current reward loss (circled in red in the right figure), making classical approaches sub-optimal.

front and rear steering introduces additional non-holonomic constraints, making precise maneuvering, such as reversing or parallel alignment, even more difficult [8]. Despite these challenges, double-Ackermann-steering mobile robots (DASMRs) offer several advantages: improved maneuverability, stability, and energy efficiency [9]. These characteristics make them particularly suitable for applications in agriculture, autonomous navigation, and uneven terrain, where robustness and cost-effectiveness are essential [10], [11]. Enabling these robots to perform precise maneuvers significantly enhances their practical use and autonomy in real-world environments, such as parking in tight spaces, or positioning within recharging areas.

Classical trajectory optimization and control methods, such as the widely used Timed Elastic Band (TEB) planner [12], have been applied to Ackermann and double-Ackermann robots. While these methods can be effective in principle, they require tuning multiple parameters, making them highly sensitive to small variations in robot dynamics (e.g., tire pressure, load distribution). In practice, this limits their robustness and complicates deployment in real-world agricultural environments [13]. Therefore, DRL approaches represent a promising alternative. However, the maneuvering constraints of double-Ackermann-steering robots pose a major challenge for DRL agents [14]. In many scenarios, reaching a desired position requires the robot to perform a complex maneuver, such as initially moving away from the goal to later approach it with the correct orientation. Unfortunately, classic DRL approaches based on reward functions such as the Euclidean distance fail in such scenarios because they penalize the robot for increasing its distance from the desired position, resulting in a reward loss (cf. Fig. 1). As a result, the maneuvers needed to achieve successful positioning are often discouraged. This can lead to sub-optimal policies [15], where the agent becomes stuck near

\*These authors contributed equally.

<sup>1</sup>Univ. Bordeaux, CNRS, Bordeaux INP, LaBRI, UMR 5800, F-33400 Talence, France. <sup>2</sup>School of Computer Science, University of Nottingham, UK. <sup>3</sup>Instituto de Investigación en Ingeniería de Aragón (I3A), Universidad de Zaragoza, 50018 Zaragoza, Spain. Corresponding author: Mélodie Daniel, e-mail: melodie.daniel@u-bordeaux.fr.

the goal in an incorrect orientation, unable to complete the task without temporarily sacrificing reward. Achieving such maneuvers can also be challenging when controlled through teleoperation, as even human operators often struggle with precise navigation [16]. Therefore, designing specific DRL frameworks that accommodate such constraints is essential for successful learning in DASMRs.

In this paper, we address the control of DASMRs using DRL in a fully model-free setting without relying on expert demonstrations or handcrafted guidance. Our goal is to develop a DRL framework that enables robust and generalizable maneuver learning despite the inherent kinematic constraints of DASMRs. To that end, we propose ManeuverNet, a novel Soft Actor-Critic (SAC) framework for DASMR with optimized reward functions designed specifically to perform precise maneuvers. The contributions of this article can be summarized as follows.

- **End-to-End DRL Framework:** We propose a fully model-free, end-to-end DRL framework designed to teach DASMRs precise maneuvering without relying on expert data, predefined trajectories, or handcrafted guidance. ManeuverNet ensures robust learning across a variety of environments, and leverages the SAC [17] algorithm enhanced with CrossQ [18] for improved sample efficiency and stability during training.

- **Comprehensive Study of Reward Functions:** We conduct an in-depth study of reward functions for DASMRs, reviewing state-of-the-art reward functions and presenting four novel reward functions for enabling precise maneuvers in non-holonomic robots.

- **Experimental Validation:** We conduct extensive validation of ManeuverNet against a range of DRL and analytical baselines, demonstrating its superior efficiency and robustness. Furthermore, we demonstrate zero-shot transfer capabilities, with the robot consistently performing well in real-world environments across diverse terrains, without any fine-tuning or domain-specific adaptation.

## II. RELATED WORK

The control of non-holonomic mobile robots, including those with double-Ackermann-steering mechanisms, has traditionally been tackled using classical control techniques [8], [11], [19]. Among these, the TEB planner [12] is widely used to generate smooth trajectories and perform local obstacle avoidance. However, TEB is highly sensitive to parameter tuning, and small changes in robot dynamics, payload, or tire pressure often necessitate full recalibration [13]. Its local obstacle handling, based on RANSAC polygon approximations, can also be overly conservative, sometimes causing the robot to halt or oscillate between forward and backward movements rather than progressing toward the goal.

In recent years, DRL has shown promise for robotic control tasks, but its application has predominantly targeted holonomic systems [1]. These approaches typically use reward functions  $\mathcal{R}_{\text{Euclid}}$  based on Euclidean distance to the goal. While such rewards perform well in holonomic systems, they are not well suited for non-holonomic systems [6].

Among non-holonomic platforms, differential-drive robots have been the most studied in DRL research [20], [21]. Although these robots are non-holonomic, they can rotate in place [6], which helps reduce some of the challenges of maneuvering. To guide these robots more effectively, the exponential reward function  $\mathcal{R}_{\text{Exp}}$  was introduced, combining both the Euclidean distance and heading error [21]. However, this reward function is less effective for Ackermann-steering robots, as they cannot rotate in place and often need to move away from the goal temporarily in order to align correctly [9]. Classic reward functions penalize such maneuvers, resulting in sub-optimal policies. Furthermore, small heading misalignments when near the target can cause significant reward losses, which further contributes to the learned policy being sub-optimal.

Some works have proposed DRL solutions specifically for single-Ackermann-steering robots. A notable example is the FastRLap framework [22], which trains an agent to follow a predefined racing trajectory. Its reward function  $\mathcal{R}_{\text{FastRLap}}$  encourages fast forward motion along the track and handles complex situations using a finite-state machine and expert demonstrations. While effective for high-speed path-following, FastRLap does not directly address the challenge of maneuvering in constrained environments. Furthermore, its reliance on handcrafted guidance and expert input limits its adaptability and generalization across tasks and robot platforms [22], [23]. Other studies have addressed Ackermann-steering car parking scenarios using reward functions that combine distance and angle error terms [14], [24]. While these reward functions  $\mathcal{R}_{\text{Car}}$  effectively address alignment maneuvers when the car is initially perpendicular to the desired position, they fall short when dealing with more complex maneuvers. Specifically, they do not account for situations in which the robot must rotate and move along different axes to reach the goal.

To overcome local reward minima and sub-optimal policy convergence in DRL more broadly, two major strategies have been explored. The first is curriculum learning, which gradually increases task complexity to facilitate learning [25]. This often involves breaking down a complex goal into intermediate waypoints to help guide the agent. However, all methods that rely on guided training, whether via curriculum learning, imitation learning, or supervised learning, tend to be highly task and environment-specific [22], [23]. The second is Hindsight Experience Replay (HER) [26], which reframes failed trajectories as successful by redefining the goal retroactively. HER uses sparse reward functions  $\mathcal{R}_{\text{HER}}$ , avoiding penalizing exploratory behaviors that are essential for reaching the goal. HER has proven to be a powerful general-purpose strategy, but it is not tailored to the specific challenges of double-Ackermann-steering control.

Despite growing interest, to the best of our knowledge, no end-to-end DRL framework has been proposed specifically for DASMR, which are more difficult to maneuver. This highlights a critical gap that our work aims to address.




Fig. 2: DASMR rotating around an instantaneous center of rotation (ICR).

### III. PROBLEM STATEMENT

We consider a DASMR, controlled to make a maneuver to reach a desired 2D position  $\mathbf{X}_d = (x_d, y_d)$ . The DASMR is a non-holonomic platform. Its configuration includes  $(\mathbf{X}_c, \theta_c)$ , where  $\mathbf{X}_c = (x_c, y_c)$  denotes the center position and  $\theta_c$  the orientation of its longitudinal axis (cf. Fig. 2). However, only two DOF (forward motion and change of orientation) can be directly controlled, and the DASMR cannot rotate without moving forward or backward. In this paper, the double-Ackermann configuration is used in a symmetric, negative 4WS setup, where the front and rear wheel steering angles are mirrored, i.e., they are equal in magnitude but opposite in direction relative to the chassis-fixed frame as shown in Fig. 2. Although ManeuverNet can be applied to any kind of DASMR, in this study, we consider large and heavy ( $> 50$  kg) robots, which are generally used in agricultural settings. The robot is initially positioned at the center of a 8-meter side square workspace and can move freely in this environment. The objective is to control the spinning velocity  $\omega$  and steering angle  $\phi$  of the four robot wheels to reach  $\mathbf{X}_d$ .

A key challenge is enabling the robot to generate feasible maneuvers without relying on prior expert knowledge or pre-defined trajectories. In addition, we assume a model-free setting in which the robot's dynamics are unknown. The robot's wheel spinning acceleration and steering speeds are limited and empirically determined to ensure stability during motion. We train a DRL agent to generate control commands that respect these physical constraints while achieving the desired behavior. This agent interacts with a simulated environment to learn a policy  $\pi$  that maximizes cumulative rewards over time. The agent is designed according to the Markov Decision Process (MDP) formalism [27]. An MDP is defined as the tuple  $(\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R})$ , where  $\mathcal{S}$  is the state space,  $\mathcal{A}$  is the action space,  $\mathcal{P}$  is the state transition function, and  $\mathcal{R}$  is the reward function. At each discrete time step  $t$ , the agent observes a state  $s_t \in \mathcal{S}$ , selects an action  $a_t \in \mathcal{A}$  according to its current policy, and transitions to a new state  $s_{t+1}$  based on  $\mathcal{P}$ . It then receives a reward  $r_t = R(s_t, a_t)$  as feedback. The policy  $\pi : \mathcal{S} \rightarrow \mathcal{A}$  is deterministic and maps each state to a specific action. While the distance between the robot

and the target point is defined by  $d$ , the success of reaching  $\mathbf{X}_d$  is determined by a distance threshold  $d_{th}$ . Similar to many DRL frameworks [28], ManeuverNet is trained in a simulator.

### IV. METHOD

In our framework, a DRL agent controls  $\omega$  and  $\phi$ . At the beginning of each training episode, the robot is initialized in the center of a square obstacle-free workspace, and must reach  $\mathbf{X}_d$ .

#### A. RL Background

In robotics, actor-critic algorithms have been successfully used in several control tasks [3], [23], [29]. Actor-critic algorithms rely on the interaction between two dense neural networks (DNNs): an actor and a critic [23]. The actor  $\mu$ , also called the policy network, selects an action  $a_t$  based on the current state  $s_t$ , following  $a_t = \mu(s_t)$ . The critic  $Q$ , also called the Q-network, estimates the expected return of the state-action pair  $(s_t, a_t)$  by computing the Q-value  $Q^\pi(s_t, a_t)$ . The critic is updated using the temporal difference learning and the Bellman equation [27], with  $Q_t^\pi(s_t, a_t) = r_t + \gamma \mathbb{E}[Q_{t+1}^\pi(s_{t+1}, a_{t+1})]$ . The actor is updated by maximizing the expected Q-value.

#### B. DRL Algorithms

SAC, introduced in [17], optimizes both cumulative rewards and policy entropy to encourage exploration and improve stability. Additionally, it uses two critic networks, a technique from Twin Delayed Deep Deterministic Policy Gradient (TD3) [30], to mitigate Q-value overestimation bias.

Despite these improvements, SAC is still computationally expensive. To improve learning efficiency, the CrossQ algorithm [18] has been recently proposed as an SAC overlay that eliminates target networks [31]. CrossQ introduces batch normalization layers in the DNN of the actor and the critic, and employs wider critic layers. Experimental results in [18], [29] showed that CrossQ outperforms existing actor-critic algorithms while significantly reducing computational times and increasing sample efficiency, making it an interesting solution for robotics applications. Therefore, our DRL agent leverages the SAC algorithm, enhanced with the CrossQ overlay.

#### C. State Space

As described in Section III, the robot's center position is denoted by  $\mathbf{X}_c$  and its orientation by the yaw angle  $\theta_c$ . The spinning velocities of the left and right wheels are  $\omega_l$  and  $\omega_r$ . The steering angles of the left and right wheels are  $\phi_l$  and  $\phi_r$ , and the corresponding steering velocities  $\dot{\phi}_l$  and  $\dot{\phi}_r$ . The robot's center linear velocity is  $\mathbf{V}_c = (\dot{x}_c, \dot{y}_c)$ , and its angular velocity is  $\dot{\theta}_c$ . At the time step  $t$ , we define the DRL agent's current state  $s_t \in \mathcal{S}$  as  $(\mathbf{X}_c, \mathbf{X}_d, \theta_c, \omega_l, \omega_r, \phi_l, \phi_r, \dot{\phi}_l, \dot{\phi}_r, \mathbf{V}_c, \dot{\theta}_c) \in \mathbb{R}^{14}$ .

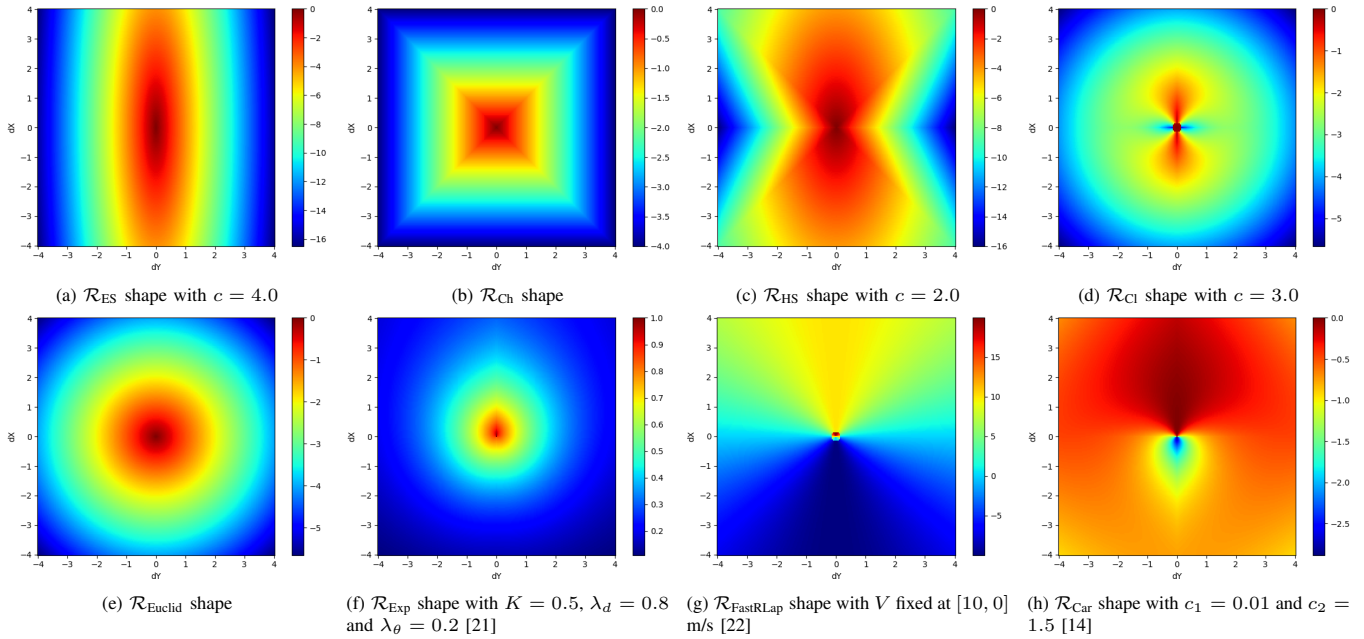


Fig. 3: Comparison between the shape of our reward functions (top) vs. classic state-of-the-art reward functions (bottom). X is the longitudinal axis, Y is the lateral axis, and the robot's center position is at the origin. Each point on the heatmaps represents the reward for a  $\mathbf{X}_d$  in the same position. Thus, the robot reaches  $\mathbf{X}_d$  when  $\|\mathbf{X}_d\| < d_{th}$ .

#### D. Action Space

When neglecting steering dynamics and assuming ideal Ackermann-steering, the kinematics of the platform can be approximated by a simplistic bicycle model, reducing front and rear pairs of wheels to a single pair of virtual central wheels [8]. Since we consider a negative symmetric 4WS configuration, the rear steering angles mirror the front steering angles. Let us define the spinning velocity and the steering angle of the virtual central wheels as  $\omega_c$  and  $\phi_c$ , respectively. At each time step  $t$ , the DRL agent outputs an action  $\mathbf{a}_t = (\omega_c, \phi_c) \in [-1, 1]$ , representing normalized commands for the wheel spinning velocity and the steering angle. These values are then scaled by the vehicle's respective actuation limits to obtain the actual control inputs. Following the double-Ackermann-steering geometry, the spinning velocities and steering angles of the left and right wheels can be calculated from  $\omega_c$  and  $\phi_c$ . For this, we compute the steering angles of the inner and outer wheels  $\phi_i$  and  $\phi_o$ , relative to the robot's current instantaneous center of rotation (ICR) shown in Fig. 2, as:

$$\phi_i = \tan^{-1} \frac{2L \cdot \sin \phi_c}{2L \cdot \cos \phi_c - W \cdot \sin \phi_c} \quad (1)$$

$$\text{and } \phi_o = \tan^{-1} \frac{2L \cdot \sin \phi_c}{2L \cdot \cos \phi_c + W \cdot \sin \phi_c}, \quad (2)$$

where  $L$  is the wheelbase of the vehicle and  $W$  is the track of the vehicle. Similarly, we can compute the spinning velocities of the inner and outer wheels  $\omega_i$  and  $\omega_o$  relative to the ICR as:

$$\omega_i = \omega_c \frac{\sqrt{[L \cdot \tan(\frac{\pi}{2} - |\phi_i|)]^2 + L^2}}{\sqrt{[L \cdot \tan(\frac{\pi}{2} - |\phi_i|) + \frac{W}{2}]^2 + L^2}} \quad (3)$$

$$\text{and } \omega_o = \omega_c \frac{\sqrt{[L \cdot \tan(\frac{\pi}{2} - |\phi_o|) + W]^2 + L^2}}{\sqrt{[L \cdot \tan(\frac{\pi}{2} - |\phi_o|) + \frac{W}{2}]^2 + L^2}}. \quad (4)$$

These inner and outer wheel spinning velocities and steering angles can then be applied to the left and right wheels based on the sign of  $\phi_c$ . If  $\phi_c \geq 0$ , the left wheel is the inner wheel, so that  $\omega_l = \omega_i$  and  $\phi_l = \phi_i$ , while the right wheel becomes the outer wheel with  $\omega_r = \omega_o$  and  $\phi_r = \phi_o$ . If  $\phi_c < 0$ , the roles are reversed, with the right wheel becoming the inner wheel.

#### E. Reward Functions

To address the issue of sub-optimal policies for DASMRs, we introduce the reward  $\mathcal{R}_{HS}$  (cf. Table I). This reward is designed to handle scenarios where the robot must temporarily deviate from the target to execute a successful maneuver. By prioritizing lateral ( $Y$ -axis) error over longitudinal ( $X$ -axis) error,  $\mathcal{R}_{HS}$  minimizes reward penalties when the robot temporarily moves away from the goal. As shown in Fig. 3c, this reward is shaped as an hourglass, which reduces penalization during maneuvers.

With the same objective, we also investigate alternative reward functions:  $\mathcal{R}_{ES}$ ,  $\mathcal{R}_{Ch}$ , and  $\mathcal{R}_{CI}$ , each defined in Table I. The reward  $\mathcal{R}_{ES}$  is a variation of the Euclidean distance that scales the error components asymmetrically, forming an ellipse. The reward  $\mathcal{R}_{Ch}$  is based on the Chebychev distance focusing on the maximum deviation along either axis, while  $\mathcal{R}_{CI}$  extends the exponential reward, proposed in [21], by incorporating a directional weighting term specifically tuned for DASMR dynamics. This reward penalizes large lateral deviations and misalignments only when the robot is close to the target  $\mathbf{X}_d$ . A comparative visualization of these reward functions alongside conventional formulations is provided in Fig. 3.

TABLE I: DASMR reward functions where  $\Delta x$  and  $\Delta y$  are the X and Y components of  $(\mathbf{X}_d - \mathbf{X}_c)$  and  $c$  a weighting parameter.

Rewards	Formulation
$\mathcal{R}_{HS}$	$-\sqrt{(\Delta x)^2 + (c \cdot (\Delta y \pm \max\{0,  \Delta y  -  \Delta x \}))^2}$
$\mathcal{R}_{ES}$	$-\sqrt{(\Delta x)^2 + (c\Delta y)^2}$
$\mathcal{R}_{Ch}$	$-\max\{ \Delta x ,  \Delta y \}$
$\mathcal{R}_{Cl}$	$\begin{cases} \tan^{-1}\left(\frac{\Delta y}{\Delta x}\right) \times ce^{d_{th}-d} & \text{if } \Delta y > d_{th} \\ -d & \text{otherwise} \end{cases}$

TABLE II: SAC and CrossQ parameters used on SBX during training and testing.

Parameter	Value
Nb. layers	2
Actor Hidden size	256
Critic Hidden size	1024
$\alpha_A = \alpha_C$	1e-3
Replay buffer size	1,000,000
Batch size $N$	256
$\gamma$	0.99
Training $d_{th}$	15 cm
Training random seed	9527
Total testing time steps	100,000
Nb. testing time steps limit per episode for settings (i), (ii) & (iii)	300
Nb. testing time steps limit per episodes for setting (iv)	500
Testing random seed	10, 52, 1234, and 9527
Testing $d_{th}$	10, and 15 cm

## V. EXPERIMENTAL RESULTS

### A. Environment Setup

We considered the Shadow Runner RR100 EDU rover as our mobile robot platform. The RR100 weighs approximately 100 kg and is 65 cm wide, 90 cm long, and 80 cm high. The DRL agent controlling the robot was trained using the PyBullet simulator, in an environment defined as follows: the robot is placed at coordinates  $(0, 0)$ , and its movements are limited to an  $8 \times 8 \text{ m}^2$  square workspace. When the environment is reset, the robot is reset to initial configuration, and a new  $\mathbf{X}_d = (x_d, y_d)$  is sampled from a  $4 \times 4 \text{ m}^2$  goal space. All values, including spaces and goal coordinates, are expressed in the robot's frame at reset. The rationale for using a small goal space is to maximize the number of  $\mathbf{X}_d$  requiring complex maneuvers. In real-world settings, the robot's workspace and goal space are constrained to the same  $4.2 \times 4.2 \text{ m}^2$  square, resulting in a more restrictive setup that better reflects practical operating conditions. Fig. 4 presents the environment setup used in both simulation and real-world experiments.

A training episode is considered successfully terminated when the DRL agent reaches  $\mathbf{X}_d$  within  $d_{th}$ , regardless of orientation. If the robot fails to reach  $\mathbf{X}_d$  within a time step limit or drives outside the bounds of its workspace, the episode is truncated, and the environment is consequently reset. We released the code at [https://github.com/MelodieDANIEL/4ws\\_actor\\_critic\\_maneuvering](https://github.com/MelodieDANIEL/4ws_actor_critic_maneuvering).

### B. Training Setup

The DRL agent was trained in a single, non-vectorized simulation environment for 600,000 time steps, where each episode was limited to 800 time steps before truncation. The agent interacted with its environment at a frequency of 40 Hz, resulting in 40 time steps per second. The training progress was monitored by logging the average




Fig. 4: Environment setup in simulation and real-world settings. The red square denotes the goal space. The blue square represents the robot's workspace. In real-world settings, both spaces coincide to impose stricter constraints on navigation and positioning.

episode reward and success rate (SR) every 10 episodes, both computed using a sliding window of 100 episodes. The full training process took approximately 1.5 hours. The training was carried out on a desktop computer equipped with an AMD Ryzen Threadripper PRO 7985WX 64-Cores (AMD Zen 4) CPU, 128 GB of memory, along with an Nvidia RTX 4090 GPU. The DRL algorithms were implemented using the Stable Baselines JAX (SBX) library, leveraging its CrossQ implementation built upon the SAC algorithm. Hyperparameters and network architecture details for both SAC and CrossQ are provided in Table II, and were kept consistent in all experiments unless otherwise specified. To ensure reproducibility, detailed network architecture and algorithm parameters are available in our GitHub repository.

### C. Benchmarking Against Baseline Approaches in Gazebo

To evaluate the effectiveness and generalizability of ManeuverNet, we conducted experiments in the Gazebo simulator, which offers a more realistic physics engine than PyBullet, thereby helping to assess robustness across different simulators (sim-to-sim gap). ManeuverNet combines the SAC framework with CrossQ, together with the reward  $\mathcal{R}_{HS}$  specifically designed to encourage maneuvers. We compared ManeuverNet with two representative DRL baselines: (i) a standard DRL agent trained with the SAC framework using a Euclidean distance-based reward commonly used for mobile robot navigation [32], and (ii) the single-Ackermann framework FastRLap [22] implemented with the SAC parameters reported in the original publication. To the best of our knowledge, FastRLap is the only available DRL framework for Ackermann mobile robots. In addition to these DRL-based methods, we benchmarked against the analytical TEB planner [12]. This closed-loop online local planner relies on several inputs, including the robot's current velocity, a local occupancy map, and its global pose estimation, to generate feasible trajectories in real time.

All evaluations were conducted using a different random seed than during training, ensuring exposure to 20 unseen

goals. Furthermore, unlike the training setup, the robot was not reset to its initial position after each episode. This forced the agent to handle successive tasks continually, without relying on episodic resets to simplify maneuvering. These evaluation conditions were designed to reflect real-world deployment scenarios and rigorously test the robustness of each method under dynamic and non-ideal circumstances.

Each approach was quantitatively evaluated with the success rate (SR), and the average distance error (AE) and the standard deviation to  $\mathbf{X}_d$ . We also evaluated the average Success weighted by (normalized inverse) Path Length (SPL) [33]. High SPL values indicate trajectories that are not only successful but also closely aligned with the shortest possible path. Achieving a high SPL is particularly challenging for DASMRs, as reaching the target position often requires complex maneuvers deviating from the optimal path.

As shown in Table III, although ManeuverNet does not reach the SR of the TEB planner, it consistently surpasses other DRL baselines. It demonstrates superior maneuvering proficiency while achieving shorter, more efficient trajectories, as evidenced by the average SPL metric. Interestingly, if we only focus on the successfully reached targets, ManeuverNet has a better SPL than TEB (0.89). The AE metric reveals that the standard DRL method gets relatively close to the target but it fails to attain precisely the desired position, primarily due to its limited maneuverability. Differently, FastRLap tends to drift away from the target and not maintain proximity, due to its emphasis on forward motion.

#### D. Comparative Study of the Reward Functions

To assess the effectiveness of our specially designed reward functions for DASMRs (see Section IV-E), we performed a comparative study under different testing conditions. This study emphasizes the reward functions that best support precise and robust maneuver learning. Additionally we evaluated three reward functions adapted from the literature: 1)  $\mathcal{R}_{\text{Exp}}$  [21] based on the direction and distance to the target goal; 2)  $\mathcal{R}_{\text{HER}}$  [26] based on a sparse binary reward function; 3)  $\mathcal{R}_{\text{Car}}$  [14] combining the direction and distance to the target goal for a parking task. In this case, we omitted the collision term as it was unnecessary in our setup.

In this study, the agents were tested in simulation using PyBullet under progressively challenging generalization settings, including: (i) using the same training seed and distance threshold  $d_{\text{th}}$ , (ii) using two unseen random seeds with  $d_{\text{th}} = 15$  cm, (iii) evaluating all seeds with  $d_{\text{th}} = 10$  cm. Setting (i) evaluated the agent under familiar conditions. In setting (ii), the sampled goal  $\mathbf{X}_d$  is altered, which typically requires the robot to perform different maneuvers, thereby testing its ability to generalize to unseen spatial configurations. Setting (iii) evaluated the precision of the learned policy, demanding finer control near the goal.

The corresponding results are summarized in Table IV. The results were generated over 100,000 time steps, with a new  $\mathbf{X}_d$  sampled either when the agent reached the goal or when the episode was truncated after 300 time steps. As shown in Table IV, all state-of-the-art reward functions either

TABLE III: Benchmarking results for an unseen seed and  $d_{\text{th}} = 15$  cm, with the success rate (SR) in %, the average error (AE) (the standard deviation  $\sigma$ ) in m, and the average SPL.

Approach Type	Approach	Test Unseen: seed = 10		
		SR $\uparrow$	AE( $\sigma$ ) $\downarrow$	SPL $\uparrow$
Analytical	TEB [12]	<b>100</b>	<b>0.05</b> (0.05)	<b>0.79</b>
	Standard DRL [32]	45	0.25 (0.17)	0.35
	FastRLap [22]	00	3.79 (2.11)	0.00
	ManeuverNet (ours)	<b>85</b>	<b>0.15</b> ( <b>0.04</b> )	<b>0.70</b>

TABLE IV: ManeuverNet simulation results for different reward functions on seen, unseen 1, and unseen 2 seeds, with the success rate (SR) in %, the average error (AE) (the standard deviation  $\sigma$ ) in m, and the average SPL.

Rew.	$d_{\text{th}}$ (cm)	Seen			Unseen 1			Unseen 2		
		SR $\uparrow$	AE( $\sigma$ ) $\downarrow$	SPL $\uparrow$	SR $\uparrow$	AE( $\sigma$ ) $\downarrow$	SPL $\uparrow$	SR $\uparrow$	AE( $\sigma$ ) $\downarrow$	SPL $\uparrow$
$\mathcal{R}_{\text{Car}}$	15	02	0.19(0.62)	0.00	02	0.17(0.75)	0.02	02	0.17(0.62)	0.02
	10	02	0.19(0.63)	0.00	02	0.18(0.76)	0.02	02	0.17(0.63)	0.02
$\mathcal{R}_{\text{Exp}}$	15	04	0.53(0.20)	0.03	04	0.55(0.25)	0.03	06	0.55(0.26)	0.04
	10	02	0.53(0.24)	0.01	02	0.56(0.24)	0.01	00	0.57(0.24)	0.00
$\mathcal{R}_{\text{HER}}$	15	55	0.25(0.21)	0.47	50	0.28(0.25)	0.40	55	0.30(0.26)	0.45
	10	48	0.21(0.19)	0.39	45	0.22(0.19)	0.35	47	0.28(0.28)	0.39
$\mathcal{R}_{\text{Cl}}$	15	72	0.27(0.25)	0.64	58	0.38(0.34)	0.50	60	0.34(0.29)	0.54
	10	62	0.26(0.28)	0.53	45	0.40(0.36)	0.39	47	0.36(0.32)	0.41
$\mathcal{R}_{\text{Ch}}$	15	92	<b>0.15</b> (0.03)	0.79	82	<b>0.16</b> (0.06)	0.70	88	0.17(0.12)	0.77
	10	81	0.13(0.06)	0.66	74	0.13(0.07)	0.62	79	0.13(0.13)	0.65
$\mathcal{R}_{\text{ES}}$	15	<b>97</b>	0.16(0.07)	0.81	96	0.17(0.16)	<b>0.82</b>	93	0.16(0.08)	0.80
	10	88	<b>0.12</b> (0.10)	0.71	<b>88</b>	0.12(0.11)	0.71	<b>85</b>	<b>0.12</b> (0.09)	<b>0.70</b>
$\mathcal{R}_{\text{HS}}$	15	<b>97</b>	0.16(0.22)	<b>0.82</b>	<b>97</b>	<b>0.16</b> (0.17)	<b>0.82</b>	<b>96</b>	<b>0.15</b> (0.03)	<b>0.84</b>
	10	<b>89</b>	0.13(0.24)	<b>0.73</b>	<b>88</b>	<b>0.11</b> (0.04)	<b>0.72</b>	<b>85</b>	<b>0.12</b> (0.07)	0.69

failed entirely or converged to sub-optimal policies when applied to DASMRs. For example,  $\mathcal{R}_{\text{Car}}$  and  $\mathcal{R}_{\text{Exp}}$  performed poorly, suggesting that they are ill-suited to the maneuvering constraints inherent to DASMRs. Similarly,  $\mathcal{R}_{\text{HER}}$  did not exceed a 55% SR, indicating limited effectiveness. While  $\mathcal{R}_{\text{Cl}}$  achieved moderately better results (72% and 62% of SR), it struggled to generalize to unseen environments, highlighting its lack of robustness.

In contrast, the rewards  $\mathcal{R}_{\text{HS}}$ ,  $\mathcal{R}_{\text{ES}}$  and  $\mathcal{R}_{\text{Ch}}$  demonstrated greater adaptability to complex maneuvering tasks. Among these, the reward  $\mathcal{R}_{\text{HS}}$  performed the best overall, achieving at least 96% of SR across all test scenarios with  $d_{\text{th}} = 15$  cm and also superior performance in terms of the SPL metric, highlighting its efficiency and reliability. This improvement stems from the stronger weighting of the lateral ( $y$ ) displacement in these rewards, which better captures the sideward maneuvering requirements inherent to DASMR.

Additional results under challenging conditions, specifically continuous goal targeting without pose reinitialization, are available in our GitHub repository. These findings align with the main results and further showcase the flexibility and robustness of our reward functions.

#### E. Real-World Results

We deployed ManeuverNet with  $\mathcal{R}_{\text{HS}}$  on a real robot using zero-shot transfer, without any fine-tuning. The current state was estimated using the robot’s sensors, which include an UM7 IMU, wheel encoders, and an RS-LiDAR-16 3D LiDAR. The robot was controlled via ROS within a real-world, square workspace measuring 4.2 meters per side, bounded by tables. Goal positions were randomly sampled within this area, and all evaluations were conducted using a fixed  $d_{\text{th}} = 15$  cm. As illustrated in Fig. 5, the DRL agent successfully controlled the robot with high precision in real-




Fig. 5: An example of a maneuver executed by ManeuverNet in both simulation and real-world settings is shown in the first two rows. Another example, showcasing multi-terrain performance, is presented in the final row of figures. The red dot represents  $X_d$ .

world conditions, despite the absence of any adaptation.

To further assess the generalization capabilities of ManeuverNet, we tested it on a variety of surface types, including vinyl, artificial grass, and carpet. An example is shown in Fig. 5. Despite significant differences in ground friction and contact dynamics, the policy maintained consistent and robust performance, underscoring its independence from specific physical parameters of the simulation. Moreover, the robot was not always initialized at the center of the workspace, yet the agent reliably handled arbitrary initial positions and orientations. This demonstrates ManeuverNet’s capacity to function effectively in unstructured and non-resettable environments.

To evaluate the effectiveness of ManeuverNet, we performed a comparative study with the TEB planner [12]. TEB was configured to generate smooth trajectories toward the target positions. To ensure a fair comparison, we sampled the same six successful target positions for both ManeuverNet and the TEB planner. To ensure consistency in evaluation, the desired robot orientation for the TEB planner was matched to the final pose attained by the DRL agent. The comparison in Table V demonstrates that, in all targets tested, ManeuverNet consistently produced shorter paths, achieving up to a 90% improvement in the SPL metric. On average, ManeuverNet outperforms TEB by 40% in SPL and reduces navigation time by 28%. Notably, similar trends have been observed in prior work [34], which reported superior performance of DRL over TEB in the context of differential-drive robots.

Although both approaches successfully reached the six target positions, several practical limitations help explain the performance differences reported in Table V. First, the TEB planner exhibits high sensitivity to the robot’s physical state: variations in dynamics, payload, or tire pressure often require fine-tuning and recalibration, which is impractical for real-world deployment. Second, the TEB planner includes a RANSAC-based algorithm for obstacle avoidance, which occasionally causes the robot to halt or oscillate between forward and backward motions, thereby impeding progress toward the goal. This behavior results in inefficient tra-

TABLE V: Real-world comparison of SPL and [navigation time] (in seconds) across six target positions. The final column reports the number of parameters requiring manual tuning for real-world deployment.  $\Delta^+$  denotes the performance improvement of ManeuverNet relative to the TEB planner.

Approach	Target ID						#Param.
	1	2	3	4	5	6	
TEB [12]	0.21 [9]	0.31 [4]	0.63 [7]	0.93 [4]	0.28 [7]	<b>0.37</b> [5]	> 10
ManeuverNet	<b>0.40</b> [4]	<b>0.32</b> [4]	<b>1.00</b> [3]	<b>0.99</b> [3]	<b>0.41</b> [5]	<b>0.37</b> [5]	<b>0</b>
$\Delta^+$ (%)	90 [56]	03 [00]	60 [57]	06 [25]	46 [29]	00 [00]	–

jectories. Finally, the reliance of the TEB planner on a sequential planning/control loop introduces latency, reducing robot responsiveness during navigation.

#### F. Limitations

Despite the strong performance of ManeuverNet, we have identified a few limitations. First, the current DRL approach does not account for obstacles. The agent is trained and evaluated in obstacle-free environments, which limits its applicability in cluttered or dynamic real-world scenarios where path planning and collision avoidance are critical. A promising solution to address this limitation is to integrate a higher-level planner such as A\*, which could generate a global collision-free path to the target. This was successfully tested, as shown in the video: <https://youtu.be/3-aarbuEOSY>. Alternatively, we can incorporate obstacle avoidance in the learning process by extending the state space and adapting the reward function, as explored in our preliminary work [35].

Second, ManeuverNet focuses solely on reaching a desired position and does not explicitly handle the robot’s final 2D pose. Nevertheless, our framework could be extended to handle precise final poses when needed by adapting the reward function to incorporate orientation constraints, similar to strategies used in goal-conditioned RL [36].

## VI. CONCLUSION

In this work, we proposed ManeuverNet, a DRL framework leveraging SAC and CrossQ, specifically tailored for the control of DASMRs. By designing novel reward functions that better take into account the maneuvering constraints of such robots, we addressed the limitations of

existing state-of-the-art methods, which often fail to generalize or converge to sub-optimal policies in this scenario. ManeuverNet was evaluated extensively in simulation under various generalization settings. These included unseen goal distributions, stricter success thresholds, and non-reset scenarios. Across all settings, ManeuverNet consistently outperformed baseline approaches, improving success rate by at least 40% while maintaining efficient trajectories. Furthermore, we demonstrated the framework's zero-shot transfer capabilities in real-world experiments, validating its robustness across diverse terrains and settings without requiring fine-tuning or expert demonstrations. Moreover, when compared to the widely used TEB planner, ManeuverNet improved the maneuvering trajectory efficiency by up to 90%. For future work, we plan to extend our framework by integrating obstacle-aware planning and orientation control, further broadening the versatility and applicability in real-world scenarios.

#### ACKNOWLEDGMENT

This work was supported by the French Government under the France 2030 program through the National Research Agency (ANR) grant reference ANR-24-PEAE-0002. It was also funded by the Nouvelle-Aquitaine Region through the MIRAE project. Author M. Aranda was supported through grant RYC2024-051408-I, funded by MICIU/AEI/10.13039/501100011033 and by ESF+.

#### REFERENCES

- [1] A. Mehmood, I. Shaikh, and A. Ali, "Application of Deep Reinforcement Learning for Tracking Control of 3WD Omnidirectional Mobile Robot," *Inf. Technol. Control*, vol. 50, no. 19, pp. 507–521, 2021.
- [2] Y. Zhang, J. Zeng, *et al.*, "Dual-Layer Reinforcement Learning for Quadruped Robot Locomotion and Speed Control in Complex Environments," *Appl. Sci.*, vol. 14, no. 19, 2024, Art. no. 8697.
- [3] C. Gaspard, G. Passault, *et al.*, "FootstepNet: An Efficient Actor-Critic Method for Fast On-line Bipedal Footstep Planning and Forecasting," in *IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, 2024, pp. 13 749–13 756.
- [4] J. Moreno, E. Clotet, *et al.*, "Design, Implementation and Validation of the Three-Wheel Holonomic Motion System of the Assistant Personal Robot (APR)," *Sensors*, vol. 16, no. 10, 2016, Art. no. 1658.
- [5] J.-S. Zhao, X. Liu, *et al.*, "Design of an Ackermann-type Steering Mechanism," *Proc. Inst. Mech. Eng.*, vol. 227, no. 11, pp. 2549–2562, 2013.
- [6] R. Siegwart and I. R. Nourbakhsh, *Introduction to Autonomous Mobile Robotics*. MIT Press, 2005.
- [7] C. Samson, "Control of Chained Systems Application to Path Following and Time-Varying Point-Stabilization of Mobile Robots," *IEEE Trans. Robot. (T-RO)*, vol. 40, no. 1, pp. 64–77, 1995.
- [8] L. Hultinen and J. Mattila, "Flow-Limited Path-Following Control of a Double Ackermann Steered Hydraulic Mobile Manipulator," in *IEEE/ASME Int. Conf. Adv. Intell. Mechatron. (AIM)*, 2020, pp. 625–630.
- [9] W. Yu, O. Y. Chuy, *et al.*, "Analysis and Experimental Verification for Dynamic Modeling of a Skid-Steered Wheeled Vehicle," *IEEE Trans. Robot. (T-RO)*, vol. 26, no. 2, pp. 340–353, 2010.
- [10] M. Deremetz, R. Lenain, *et al.*, "Path Tracking of a Four-Wheel Steering Mobile Robot: A Robust Off-Road Parallel Steering Strategy," in *Eur. Conf. Mob. Robot. (ECMR)*, 2017, pp. 1–7.
- [11] C. Cariou, R. Lenain, *et al.*, "Automatic Guidance of a Four-Wheel-Steering Mobile Robot for Accurate Field Operations," *J. Field Robot.*, vol. 26, no. 6-7, pp. 504–518, 2009.
- [12] C. Rösmann, F. Hoffmann, and T. Bertram, "Kinodynamic Trajectory Optimization and Control for Car-Like Robots," in *IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, 2017, pp. 5681–5686.
- [13] B. Magyar, N. Tsiogkas, *et al.*, "Timed-Elastic Bands for Manipulation Motion Planning," *IEEE Robot. Autom. Lett. (RA-L)*, vol. 4, no. 4, pp. 3513–3520, 2019.
- [14] L. Lazzaroni, F. Bellotti, *et al.*, "Deep Reinforcement Learning for Automated Car Parking," in *Springer Int. Conf. Appl. Electron. Perovading Ind. Environ. Soc. (ApplePies)*, vol. 1036, 2022, pp. 125–130.
- [15] M. Patil, B. Wehbe, and M. Valdenegro-Toro, "Deep Reinforcement Learning for Continuous Docking Control of Autonomous Underwater Vehicles: A Benchmarking Study," in *IEEE OCEANS*, 2021, pp. 1–7.
- [16] A. Datar, C. Pan, *et al.*, "Terrain-Attentive Learning for Efficient 6-DoF Kinodynamic Modeling on Vertically Challenging Terrain," in *IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, 2024, pp. 5438–5443.
- [17] T. Haarnoja, A. Zhou, *et al.*, "Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor," in *PMLR Int. Conf. Mach. Learn. (ICML)*, vol. 80, 2018, pp. 1856–1865.
- [18] A. Bhatt, D. Palenicek, *et al.*, "CrossQ: Batch Normalization in Deep Reinforcement Learning for Greater Sample Efficiency and Simplicity," in *PMLR Int. Conf. Learn. Represent. (ICLR)*, 2024.
- [19] R. R. Shamshiri, A. Azimi, *et al.*, "Online Path Tracking With an Integrated  $H_\infty$  Robust Adaptive Controller for a Double-Ackermann Steering Robot for Orchard Waypoint Navigation," *Springer Int. J. Intell. Robot. Appl. (IJIRA)*, vol. 9, no. 1, pp. 257–277, 2025.
- [20] W. Zhang, S. Wang, *et al.*, "DRL-DCLP: A Deep Reinforcement Learning-Based Dimension-Configurable Local Planner for Robot Navigation," *IEEE IEEE Robot. Autom. Lett. (RA-L)*, vol. 10, no. 4, pp. 3636–3643, 2025.
- [21] T. Soualhi, N. Crombez, *et al.*, "Leveraging Motion Perceptibility and Deep Reinforcement Learning for Visual Control of Nonholonomic Mobile Robots," *Robot. Auton. Syst.*, vol. 189, p. 104920, 2025.
- [22] K. Stachowicz, D. Shah, *et al.*, "FastRLAP: A System for Learning High-Speed Driving via Deep RL and Autonomous Practicing," in *PMLR Conf. Robot Learn. (CoRL)*, 2023, pp. 3100–3111.
- [23] M. Daniel, A. Magassouba, *et al.*, "Multi Actor-Critic DDPG for Robot Action Space Decomposition: A Framework to Control Large 3D Deformation of Soft Linear Objects," *IEEE Robot. Autom. Lett. (RA-L)*, vol. 9, no. 2, pp. 1318–1325, 2024.
- [24] L. Junzuo and L. Qiang, "An Automatic Parking Model Based on Deep Reinforcement Learning," *IOP Publishing J. Phys. Conf. Ser.*, vol. 1883, no. 1, p. 012111, 2021.
- [25] H. Xue, B. Hein, *et al.*, "Using Deep Reinforcement Learning with Automatic Curriculum Learning for Mapless Navigation in Intralogistics," *Appl. Sci.*, vol. 12, no. 6, 2022, Art. no. 3153.
- [26] M. Andrychowicz, F. Wolski, *et al.*, "Hindsight Experience Replay," *Adv. Neural Inf. Process. Syst.*, vol. 30, 2017.
- [27] R. S. Sutton and A. G. Barto, *Reinforcement Learning - An Introduction*. MIT Press, 2018.
- [28] W. Zhao, J. P. Queralta, and T. Westerlund, "Sim-to-Real Transfer in Deep Reinforcement Learning for Robotics: A Survey," in *IEEE Symp. Ser. Comput. Intell. (SSCI)*, 2020, pp. 737–744.
- [29] C. Gaspard, M. Duclusaud, *et al.*, "FRASA: An End-to-End Reinforcement Learning Agent for Fall Recovery and Stand Up of Humanoid Robots," in *IEEE Int. Conf. Robot. Autom. (ICRA)*, 2025, pp. 15 994–16 000.
- [30] S. Fujimoto, H. van Hoof, and D. Meger, "Addressing Function Approximation Error in Actor-Critic Methods," in *PMLR Int. Conf. Mach. Learn. (ICML)*, vol. 80, 2018, pp. 1582–1591.
- [31] T. P. Lillicrap, J. J. Hunt, *et al.*, "Continuous Control with Deep Reinforcement Learning," in *PMLR Int. Conf. Learn. Represent. (ICLR)*, 2016.
- [32] V. R. F. Miranda, A. A. Neto, *et al.*, "Generalization in Deep Reinforcement Learning for Robotic Navigation by Reward Shaping," *IEEE Trans. Ind. Electron.*, vol. 71, no. 6, pp. 6013–6020, 2024.
- [33] P. Anderson, A. X. Chang, *et al.*, "On Evaluation of Embodied Navigation Agents," *CoRR*, vol. abs/1807.06757, 2018.
- [34] D. Arce, J. Solano, and C. Beltrán, "A Comparison Study between Traditional and Deep-Reinforcement-Learning-Based Algorithms for Indoor Autonomous Navigation in Dynamic Scenarios," *Sensors*, vol. 23, no. 24, 2023, Art. no. 9672.
- [35] K. Deflesselle, M. Daniel, *et al.*, "Towards Safe Maneuvering of Double-Ackermann-Steering Robots with a Soft Actor-Critic Framework," in *SIIV-FM2L workshop organized within IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, 2025.
- [36] M. Liu, M. Zhu, and W. Zhang, "Goal-Conditioned Reinforcement Learning: Problems and Solutions," in *Int. Jt. Conf. Artif. Intell. (IJCAI)*, 2022, pp. 5502–5511.