

Deep Reinforcement Learning Based Tracking Control of an Autonomous Surface Vessel in Natural Waters

Wei Wang^{1,2,*}, Xiaojing Cao^{3,*}, Alejandro Gonzalez-Garcia^{2,4}, Lianhao Yin¹, Niklas Hagemann²,
Yuanyuan Qiao^{2,3}, Carlo Ratti², and Daniela Rus¹

Abstract—Accurate control of autonomous marine robots still poses challenges due to the complex dynamics of the environment. In this paper, we propose a Deep Reinforcement Learning (DRL) approach to train a controller for autonomous surface vessel (ASV) trajectory tracking and compare its performance with an advanced nonlinear model predictive controller (NMPC) in real environments. Taking into account environmental disturbances (e.g., wind, waves, and currents), noisy measurements, and non-ideal actuators presented in the physical ASV, several effective reward functions for DRL tracking control policies are carefully designed. The control policies were trained in a simulation environment with diverse tracking trajectories and disturbances. The performance of the DRL controller has been verified and compared with the NMPC in both simulations with model-based environmental disturbances and in natural waters. Simulations show that the DRL controller has 53.33% lower tracking error than that of NMPC. Experimental results further show that, compared to NMPC, the DRL controller has 35.51% lower tracking error, indicating that DRL controllers offer better disturbance rejection in river environments than NMPC.

I. INTRODUCTION

Autonomous Surface Vessels (ASVs) deliver solutions for a variety of applications, including environmental monitoring, security patrols, search and rescue, bridge and other infrastructure inspections, and military and naval operations [1]–[11]. Trajectory tracking is an essential capability for ASVs because many other higher-level robot tasks such as navigation, docking and multi-robot coordination require time-critical tracking.

Trajectory tracking is a difficult task because ASVs usually follow a predefined trajectory while being exposed to environmental disturbances caused by waves, winds, and ocean currents [12]. Several studies have addressed the ASV trajectory tracking control including the sliding mode method [13], input-output feedback linearization [14] and adaptive

control [15]. In recent years, Model Predictive Control (MPC) has been proposed to solve the trajectory tracking control problem [16] and proven to be the most efficient controller for ASV trajectory tracking [17]–[21]. However, MPC is susceptible to model uncertainty and external perturbations. While robust controllers can in some respect deal with time-varying disturbances and uncertainties, designing robust controllers for ASVs is still difficult. For example, robust controllers can cause instability with improper system models or lack of tuning.

Model-free Reinforcement Learning (RL) trains a control policy based on the interaction of agents with their environment [22] and does not require an accurate model. Therefore, RL has been proposed for ASV control problems [23]–[26]. However, only [23] provides an experimental evaluation of an RL approach, introducing sea trials for trajectory tracking of a fully-actuated ASV. Besides, the approach is model-based, learning model parameters, not only the control policy.

Moreover, Deep RL (DRL), learning from more complex and dynamic environments and scenarios [27], has led to a few interesting developments in the learning of control laws for ASVs [27]–[31]. The Deep Q-Network (DQN) is a powerful actor-critic DRL algorithm, and it has been applied for ASV path following in [27], and for obstacle avoidance in [29]. Nonetheless, DQN can only deal with discrete and low-dimensional action spaces. Hence, the Deep Deterministic Policy Gradient (DDPG) algorithm was developed to benefit from the actor-critic architecture, and to deal with continuous action spaces [32]. DDPG was then implemented in [30] for ASV path following control, presenting simulation and experimental results. Likewise, DDPG was used for obstacle avoidance in [31], illustrating an improvement when compared against DQN-based strategies.

Nevertheless, the vast majority of DRL controllers for ASVs are limited to simulation, as experimental evaluations are scarce. Furthermore, current DRL methods are often used to solve the path following problem. The trajectory tracking problem has rarely been tackled in the ASV literature, partly because it requires more control constraints that require a complicated multi-action joint policy in the DRL control system. Additionally, the performance of the current ASV DRL controller is usually not compared to traditional controllers, especially against techniques such as MPC.

Based on the above discussion, this paper focuses on solving the trajectory-tracking problem for a physical ASV using DRL. Physical robots typically exhibit complex behaviors in real environments in the presence of environmental

This work was supported by grant from the Amsterdam Institute for Advanced Metropolitan Solutions (AMS) in Netherlands. L. Yin was supported by Knut and Alice Wallenberg Foundation.

¹W. Wang, L. Yin, and D. Rus are with the Computer Science and Artificial Intelligence Lab (CSAIL), Massachusetts Institute of Technology, Cambridge, MA 02139 USA.

²W. Wang, A. Gonzalez-Garcia, N. Hagemann, Y. Qiao, and C. Ratti are with the SENSEable City Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139 USA.

³X. Cao, and Y. Qiao are with Intelligent Perception and Computing Research Center, the School of Artificial Intelligence, Beijing University of Posts and Telecommunications (BUPT), Beijing 100876, China.

⁴A. Gonzalez-Garcia is with MECO Research Team, Department of Mechanical Engineering, KU Leuven, Belgium and Flanders Make@KU Leuven, Belgium.

*These authors contribute equally to this work.

disturbances, noisy measurements, and non-ideal actuators. We propose a framework for creating an efficient DRL reward function that takes into account these complex behaviors.

The main contributions of our work are summarized as follows:

- A DRL controller for trajectory tracking of real ASVs;
- Numerical simulations considering model-based environmental disturbances which illustrate the potential of the DRL policy;
- Extensive natural-water experiments verifying the effectiveness of DRL controllers;
- Comparisons between DRL and MPC showing that: 1) DRL outperforms MPC for tracking error; 2) DRL does better than MPC at rejecting disturbances.

II. ASV OVERVIEW AND ITS DYNAMIC MODEL

In this section, we introduce the ASV prototype and its simplified mathematical model. The simplified model serves as the basis for the simulation environment of the proposed DRL controller in the next section.

A. Surface Vessel

The vessel has four thrusters around the hull to achieve holonomic motions as shown in Fig. 1. An Intel NUC is adopted




Fig. 1. The surface vessel. (a) The prototype (side view); (b) thruster configuration (bottom view). f_1 , f_2 , f_3 and f_4 denote the forces generated by left, right, anterior and rear thrusters, respectively.

as the main controller running the Robotic Operating System (ROS). Moreover, an auxiliary microprocessor (STM32F103) is used for real-time actuator control. The robot has multiple onboard sensors including a 3D LiDAR (Velodyne, Puck VLP-16), an IMU, a camera, and a GPS sensor. The camera and GPS are not used in this paper. The vessel weighs around 15 kg, and its dimensions are $0.90 \times 0.45 \times 0.15$ m. It is powered by an 11.1 V Li-Po battery which lasts around three hours. More details of the robot hardware can be found in [17].

B. Dynamic Model of ASV

The dynamics of our vessel is described by the following nonlinear differential equation [17]

$$\dot{\boldsymbol{\eta}} = \mathbf{T}(\boldsymbol{\eta})\mathbf{v} \quad (1)$$

$$\dot{\mathbf{v}} = \mathbf{M}^{-1}(\boldsymbol{\tau} + \boldsymbol{\tau}_{\text{env}}) - \mathbf{M}^{-1}(\mathbf{C}(\mathbf{v}) + \mathbf{D}(\mathbf{v}))\mathbf{v} \quad (2)$$

where $\boldsymbol{\eta} = [x \ y \ \psi]^T \in \mathbb{R}^3$ is the position and heading angle of the vessel in the inertial frame; $\mathbf{v} = [u \ v \ w]^T \in \mathbb{R}^3$ denotes the vessel velocity, which contains the surge velocity

(u), sway velocity (v), and yaw rate (w) in the body fixed frame; $\mathbf{T}(\boldsymbol{\eta}) \in \mathbb{R}^{3 \times 3}$ is the transformation matrix converting a state vector from body frame to inertial frame; $\mathbf{M} \in \mathbb{R}^{3 \times 3}$ is the positive-definite symmetric added mass and inertia matrix; $\mathbf{C}(\mathbf{v}) \in \mathbb{R}^{3 \times 3}$ is the skew-symmetric vehicle matrix of Coriolis and centripetal terms; $\boldsymbol{\tau}_{\text{env}} \in \mathbb{R}^3$ is the environmental disturbances from the wind, currents and waves; $\mathbf{D}(\mathbf{v}) \in \mathbb{R}^{3 \times 3}$ is the positive-semi-definite drag matrix-valued function; $\boldsymbol{\tau} = [\tau_u \ \tau_v \ \tau_w]^T \in \mathbb{R}^3$ is the force and torque applied to the vessel in all three DOFs, and is defined as follows

$$\boldsymbol{\tau} = \mathbf{B}\mathbf{u} = \begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ \frac{a}{2} & -\frac{a}{2} & \frac{b}{2} & -\frac{b}{2} \end{bmatrix} \begin{pmatrix} f_1 \\ f_2 \\ f_3 \\ f_4 \end{pmatrix} \quad (3)$$

where $\mathbf{B} \in \mathbb{R}^{4 \times 3}$ is the control matrix describing the thruster configuration and $\mathbf{u} = [f_1 \ f_2 \ f_3 \ f_4]^T \in \mathbb{R}^4$ is the control vector where f_1 , f_2 , f_3 and f_4 represent the left, right, anterior, and rear thrusters, respectively; a is the distance between the transverse thrusters and b is the distance between the longitudinal thrusters. We refer the readers to [17] for more details of \mathbf{M} , $\mathbf{C}(\mathbf{v})$ and $\mathbf{D}(\mathbf{v})$.

By combining (1) and (2), the complete dynamic model of the vessel is further reformulated as follows

$$\dot{\mathbf{q}}(t) = f(\mathbf{q}(t), \mathbf{u}(t)) \quad (4)$$

where $\mathbf{q} = [x \ y \ \psi \ u \ v \ w]^T \in \mathbb{R}^6$ is the state vector of the vessel, and $f(\cdot, \cdot) : \mathbb{R}^{n_q} \times \mathbb{R}^{n_u} \rightarrow \mathbb{R}^{n_q}$ denote the continuously differentiable state update function. The system model describes how the full state \mathbf{q} changes in response to applied control input $\mathbf{u} \in \mathbb{R}^4$. (4) will act as a simulator in the simulation environment to train or test the control policy.

III. DRL FRAMEWORK FOR ASV TRAJECTORY TRACKING

In this section, we propose a DRL framework to solve the trajectory tracking problem of a physical ASV.

A. Problem Formulation

Let $\boldsymbol{\eta}_d(t) = [x_d(t) \ y_d(t) \ \psi_d(t)]^T \rightarrow \mathbb{R}^3$ be a given sufficiently smooth time-varying desired trajectory with bounded time-derivatives. The trajectory tracking problem is: Design a controller such that all the closed-loop signals are bounded and the tracking error $\|\boldsymbol{\eta}(t) - \boldsymbol{\eta}_d(t)\|$ converges to a neighborhood of the origin that can be made arbitrarily small.

B. Method

We can formulate the trajectory tracking task as a discrete-time optimization problem and use DRL to solve this control task in which robots learn based on trial and error to modify their actions to maximize the gains from interaction with their environment [33]. In this paper, for verifying our proposed learning framework, we select DDPG network [32] as a DRL benchmark for ASV trajectory tracking due to its capability of continuous control and previous success in ASV control problems [30], [31]. Each actor and critic network has two fully connected hidden layers with 300 neurons. The activation function between the hidden layers is Relu, and that of the output layer is Tanh.

C. Control Architecture

Our goal is to learn a policy $\pi_\theta(\cdot)$ that takes robot state \mathbf{q} and reference commands $\boldsymbol{\eta}_d(t)$ as input and gives as output force commands \mathbf{u} , which are further converted into Pulse-Width Modulation (PWM) signals to drive the thrusters. Our DRL framework is illustrated in Fig. 2.




Fig. 2. Structure of the DRL controller.

Observation Space: The policy $\pi_\theta(\cdot)$ takes as input a history of previous observations (see in Fig. 3) and actions denoted by $\mathbf{s}_{t-H:t}$ where $\mathbf{s}_t = [\mathbf{q}_t, \mathbf{q}_t^d, \mathbf{a}_{t-1}]$ is the observation. $\mathbf{q}_t = [x(t) \ y(t) \ \psi(t) \ u(t) \ v(t) \ w(t)]^T \in \mathbb{R}^6$ is robot state at time t defined in Section II, $\mathbf{q}_t^d = [x_d(t) \ y_d(t) \ \psi_d(t) \ u_d(t) \ v_d(t) \ w_d(t)]^T \in \mathbb{R}^6$ is the reference state for tracking and \mathbf{a}_{t-1} is the action at the last time step $t-1$.




Fig. 3. Representation of a surface vessel in DRL trajectory tracking.

Action Space: The action, $\mathbf{a}_t \in \mathbb{R}^4$ is defined as $\mathbf{a}_t = \mathbf{u}_t$ where $\mathbf{u}_t = [f_1(t) \ f_2(t) \ f_3(t) \ f_4(t)]^T$ is expressed in (3).

Reward Function: Reward function follows with tracking task reward terms as well as a set of auxiliary terms for stability (velocity variation penalties), smoothness (action variation penalties), and minimal energy consumption.

We define Euclidean distance error as $e_p = [(x - x_d)^2 + (y - y_d)^2]^{1/2}$. Then, r_p is designed to minimize the position tracking error, denoted as follows:

$$r_p = 2^{-k_1(\mu\|e_v\|+1)}e_p - 1 \quad (5)$$

where μ and k_1 are constant parameters, and $e_v = [(u - u_d)^2 + (v - v_d)^2]^{1/2}$. e_v is used to smooth linear velocity variations.

Second, r_ψ is designed to minimize the heading error, expressed as follows:

$$r_\psi = \begin{cases} e^{-k_2\|\psi - \psi_s\|}, & \text{if } \|\psi - \psi_s\| \leq 0.5\pi \\ -e^{k_2(\|\psi - \psi_s\| - \pi)}, & \text{if } \|\psi - \psi_s\| > 0.5\pi \end{cases}, \quad (6)$$

where k_2 is a constant, $e_\psi = \psi - \psi_d$ is the heading error, and ψ_s is the sight heading which replaces the reference heading. The sight heading is calculated based on a line of sight (LOS) guidance method [34], shown in Fig. 4. We define $P_c(x, y)$,




Fig. 4. LOS guidance geometry for curved trajectory.

$P_d(x_d, y_d)$ and $P_s(x_s, y_s)$ are the current position, desired position and sight position of robot, respectively. The sight position $P_s(x_s, y_s)$ is L length forward along the tangent line at the reference position $P_d(x_d, y_d)$. Then the sight heading ψ_s is computed as the angle between the vector from $P_c(x, y)$ to $P_s(x_s, y_s)$ and X axis of the inertial frame.

Moreover, reward term r_w is designed to punish variations in angular velocity, denoted as:

$$r_w = e^{-k_3\|e_w\|} - 1 \quad (7)$$

where k_3 is a constant and $e_w = w - w_d$.

Furthermore, reward term r_a is proposed to suppress variations in actions, expressed as follows:

$$r_a = e^{-k_4\Delta_a} - 1 \quad (8)$$

where k_4 and is a constant, $\Delta_a(t) = \sum_{i=1}^4 (f_i(t) - f_i(t-1))$.

Finally, a reward term is mathematically defined to minimize the energy consumption of the propulsion system:

$$r_e = e^{-k_5E_p} - 1 \quad (9)$$

where k_5 is a constant, and $E_p = \sum_{i=1}^4 f_i^2$ represents the energy consumption of the thrusters. This reward term always punishes high energy costs-related activities, which leads to the convergence of control policies to minimal energy policies.

In summary, the whole reward function is formulated as follows

$$r = \begin{cases} \lambda_1 r_p + \lambda_2 r_\psi + \lambda_3 r_w + \lambda_4 r_a + \lambda_5 r_e, & \text{if } e_p \leq e_{\text{bound}} \\ -25, & \text{if } e_p > e_{\text{bound}} \end{cases}, \quad (10)$$

where e_{bound} is a boundary distance, $\lambda_1, \lambda_2, \lambda_3, \lambda_4, \lambda_5$ are weight coefficients for the five reward terms. If the vessel sails out of the boundary error, the total reward will turn into large punishment.

D. Simulation Environment

Computational fluid dynamics is a sufficiently accurate simulation environment but it is still too computationally expensive to be adopted during the training procedure. To balance accuracy and computing speed, we construct an ASV simulator using the differentiable nonlinear equation (4)

defined in Section II. The unknown hydrodynamic parameters \mathbf{M} , \mathbf{C} , and \mathbf{D} were identified by a nonlinear least squares method based on Trust Zone Reflection Algorithm [17]. The identified parameters for our vessel are found as $m_{11} = 12$ kg, $m_{22} = 24$ kg, $m_{33} = 1.5$ kg m², $d_{11} = 6$ kg s⁻¹, $d_{22} = 8$ kg s⁻¹, and $d_{33} = 1.35$ kg m² s⁻¹. At each time step, the simulator updates the vessel's state once it receives a new control action from the control policy.

E. Training

We train a control policy in simulation and transfer it to the real world without fine-tuning. The training contains random trajectory, initial position and measurement disturbance. A diversity of sinusoidal waves are selected as the training trajectories because they generate a coupled motion in the surge, sway, and yaw degrees of freedom, at time-varying velocities. The robot is randomly located near the initial target position, at the start of each episode. Meanwhile, the robot heading is also randomly generated. The episode will stop if 1) the tracking task fails which means that the vessel sails out of the boundary, i.e., $e_p > e_{\text{bound}}$ where $e_{\text{bound}} = 1$ m; 2) The runtime comes to its maximum value, which is set to $t_{\text{max}} = 30$ s.

The measured state \mathbf{q} of the robot are corrupted by zero-mean Gaussian noise $N(0, 0.1I_6)$, where I_6 is a 6D identity matrix. Moreover, the parameter values in the reward function (10) are listed in Table I. We add an Ornstein-Uhlenbeck

TABLE I
PARAMETER VALUES IN THE REWARD FUNCTION

Parameters	μ	k_1	k_2	k_3	k_4	k_5
Values	3	4	3	1	0.033	0.017
Parameters	λ_1	λ_2	λ_3	λ_4	λ_5	e_{bound}
Values	1.5	0.5	1	0.5	0.2	1 m

process [35] onto the action to improve the exploration with $\theta = 0.2$ and $\sigma = 0.15$. Policy weights were selected after around 3000 training episodes.

IV. SIMULATIONS AND EXPERIMENTS

In this section, extensive simulations and experiments are conducted with the surface vessel shown in Fig. 1 to validate the tracking capabilities of the developed DRL control strategy. Furthermore, we systematically compare the performance of the proposed DRL controller with NMPC in terms of tracking error, energy consumption, and force allocation strategy, and finally obtain several interesting capabilities of the DRL controller. Both simulation and experiments are executed on a small-form-factor barebone computer, Intel[®] NUC (NUC7i7DNH), which runs a robotics middleware, Robot Operating System (ROS).

A. Simulations

We use root mean square error (RMSE) to represent tracking error. In general, thruster power consumption is positively

correlated to its thrust. Hence, the mean power consumption E_{ave} of the controller is represented as follows:

$$E_{\text{ave}} \propto \sqrt{\frac{\sum_{i=1}^N (|f_1^i| + |f_2^i| + |f_3^i| + |f_4^i|)}{N}} \quad (11)$$

where N is the number of the data points.

1) *Ablation Study*: Typical reward functions for ASV tracking only consider the terms (5) and (6) to minimize the position and heading errors. By contrast, our DRL framework contains three more reward terms which largely guarantees that the trained control policy in simulation can be transferred to the real world without fine-tuning. An ablation study was conducted in simulation to demonstrate the benefits of the proposed reward function design in our DRL framework.

In particular, we compared DRL control tracking performance with two different reward functions. One policy uses a simple reward function that addresses the terms (5) and (6), while the other policy uses the developed reward function (10) and contains five terms. As shown in Fig. 5, it is clear that the control policy proposed with our developed reward function goes beyond the control policy with the simple reward function. The mean tracking errors of both control policies for




Fig. 5. Comparison of sinusoidal tracking performance in DRL using a simple reward function and our reward function. (a) Trajectory and (b) tracking error.

Euclidean distances are 0.0496 m and 0.0743 m, respectively, which means that the tracking errors of our control policy are 33.03% lower compared to a policy utilizing only simple reward functions. Fig. 6 compares the trajectory tracking control forces of the two policies. Our control policy reduces energy consumption by 37.07% and provides smooth actuation signals.

2) Controller Comparison With Realistic Disturbances:

We further compare the trained policy $\pi_{\theta}^*(\cdot)$ with the NMPC in [17]. We tested the controllers with different reference trajectories (C-Curve, Linear, and Sinusoidal) and both controllers can track the references well. Due to the limited space, the results of sinusoidal tracking are merely illustrated herein. Figure 7 compares sinusoidal tracking between DRL and NMPC, considering model-based disturbances from wind, waves, and

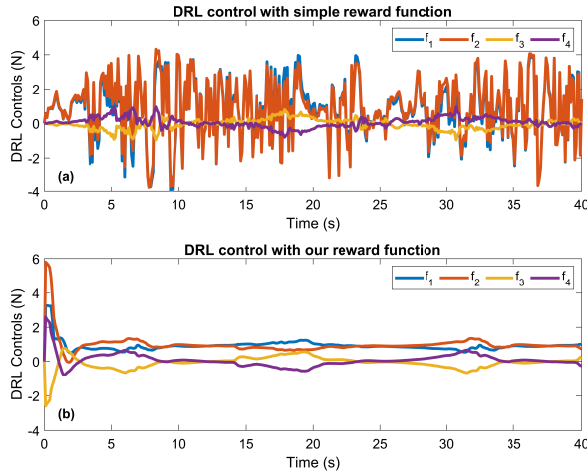


Fig. 6. Control forces during DRL tracking using the simple reward function and our reward function.

currents as $\tau_{\text{env}} = \tau_{\text{wind}} + \tau_{\text{waves}} + (\mathbf{C}(\mathbf{v}_c) + \mathbf{D}(\mathbf{v}_c))\mathbf{v}_c$. Here, τ_{wind} represents the wind originated forces and moments, τ_{wave} stands for the wave induced forces and moments, and \mathbf{v}_c the ocean current velocity. The disturbances dynamics are represented as follows:

Wind: Wind forces and moments depend on the wind speed V_w and its angle of attack γ_{rw} , described by

$$\boldsymbol{\tau}_{\text{wind}} = \begin{bmatrix} X_{\text{wind}} \\ Y_{\text{wind}} \\ N_{\text{wind}} \end{bmatrix} = \frac{1}{2} \rho_a V_{rw}^2 \begin{bmatrix} C_X(\gamma_{rw}) A_{FW} \\ C_Y(\gamma_{rw}) A_{LW} \\ C_N(\gamma_{rw}) A_{LW} L_{OA} \end{bmatrix} \quad (12)$$

where V_{rw} is the wind relative speed, which can be computed from the wind velocity vector $\mathbf{v}_w = [V_w \cos(\beta_{\text{wind}} - \psi), V_w \sin(\beta_{\text{wind}} - \psi), 0]^T$, $\gamma_{rw} = -\text{atan2}(v_{rw}, u_{rw})$, β_{wind} stands for the wind direction, ρ_a denotes the air density, $A_{FW} = 0.045$ is the projected frontal area, $A_{LW} = 0.09$ is the projected lateral area, and $L_{OA} = 0.9$ is the ASV length. See [36] for more details on the computation of wind coefficients C_X, C_Y, C_N .

Waves: Waves can be described combining first-order (zero-mean oscillatory motion) and second-order (wave drift) dynamics [37]. Then, the state-space equation follows an approximation of the Pierson-Moskowitz spectrum [37], where F_{wave} is the wave force, and N_{wave} is the wave moment. The equations depend on Gaussian noise signals (to introduce randomness into the system), the encountered wave spectrum peak frequency, and the wave direction. For more details regarding the state-space representation, see [36], [37]. Finally, the wave disturbance vector is

$$\boldsymbol{\tau}_{\text{wave}} = \begin{bmatrix} X_{\text{wave}} \\ Y_{\text{wave}} \\ N_{\text{wave}} \end{bmatrix} = \begin{bmatrix} F_{\text{wave}} \cos(\beta_{\text{wave}} - \psi) \\ F_{\text{wave}} \sin(\beta_{\text{wave}} - \psi) \\ N_{\text{wave}} \end{bmatrix} \quad (13)$$

Ocean Currents: Currents are usually simplified to be irrotational, and modeled by a Gauss-Markov process [37], such as

$$\dot{V}_c + \mu_c V_c = \omega_c \quad (14)$$

where V_c denotes the magnitude of the current speed at direction β_c , ω_c is a Gaussian white noise signal, and $\mu_c > 0$ is a constant. Then, the current velocity vector is $\mathbf{v}_c = [V_c \cos(\beta_c - \psi), V_c \sin(\beta_c - \psi), 0]^T$.

For this study, a wind speed of 4 knots was applied (a light breeze on the Beaufort scale [37]), while the waves provided forces of up to 1 N, and the currents a disturbance of up to 0.2 m/s. These magnitudes are enough to require around 25% more control effort. The simulations for DRL and NMPC were each performed three times, ensuring repeatability. The Gaussian noises for waves and currents varied between simulation runs, providing different random perturbations around the same magnitudes. It can be observed that both the DRL and NMPC

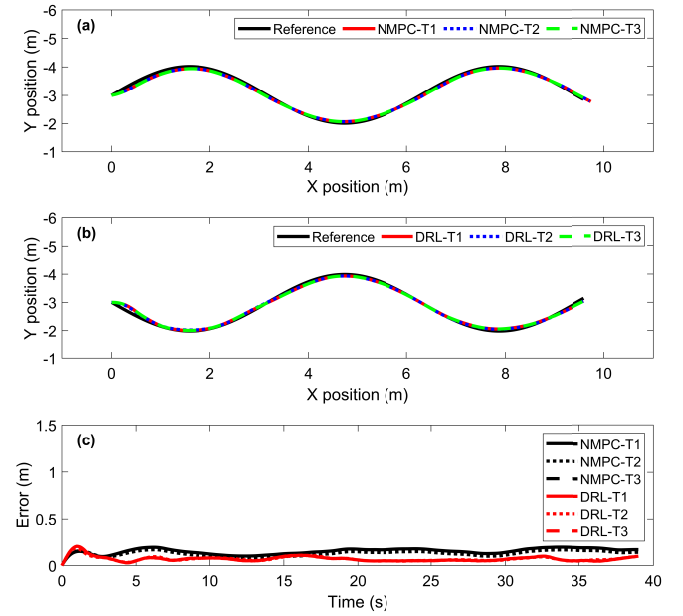


Fig. 7. Performance of sinusoidal tracking using DRL and NMPC. (a) Trajectory and (b) tracking error. T1, T2, and T3 represent first, second and third tests, respectively.

controllers can track well along the reference trajectory in simulation. The average tracking errors of the DRL controller and the NMPC controller in Euclidean distance are 0.068 m and 0.146 m, which means that the tracking error for the DRL is 53.33% less compared to that for the NMPC. Hence, NMPC has a larger error than the DRL in simulation, considering environmental disturbances. Moreover, Fig. 8 shows the control forces of DRL and NMPC. The average energy consumption for the NMPC and DRL are 2.018 N and 1.917 N respectively. Then, not only is the DRL more accurate than NMPC, but it also requires less energy to track the trajectory and reject the perturbations. Additionally, it is very interesting that the control forces behavior is similar between DRL and NMPC. This indicates that the DRL policy has learned a control performance similar to that of the NMPC controller, with enhanced robustness against external disturbances.




Fig. 8. Control force while tracking the sinusoidal trajectory. (a) Using NMPC controller; (b) using DRL controller.

B. Experiments

We deployed the trained control policy onto our surface vessel to track reference trajectories in the Charles River. We also used NMPC [17] as the benchmark controller and did the same tracking experiments as for the DRL. Both the DRL controller and the NMPC are run locally on the computer onboard the vessel. The execution time of the NMPC is approximately 1 ms while the execution time of the DRL is nearly 100 μ s. Moreover, the vessel runs a real-time SLAM [19], [38] using its onboard 3D LiDAR and IMU to estimate its pose η and velocity \mathbf{v} . All related parameters, including position, heading angle, velocity, and force of the robot, are recorded at 10 Hz during the experiments for analysis. During the field test, the wind on the Charles River was around 18 kph and the waves were less than one foot.

Figure 9 compared the trajectory tracking performance of the DRL and NMPC in the river with the aforementioned environmental disturbances. Tracking errors for the NMPC

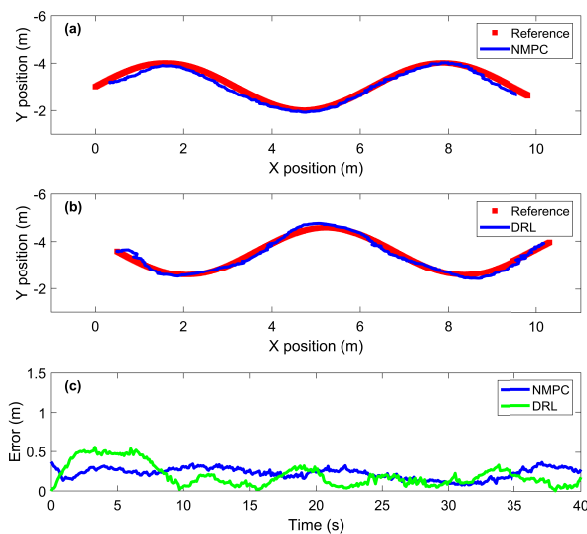


Fig. 9. Performance of sinusoidal trajectory tracking in the Charles River using DRL and MPC. (a) Trajectory and (b) tracking errors.

and the DRL are 0.2187 m and 0.1411 m, respectively, which

means that the tracking error for the DRL is 35.51% less compared to that for the NMPC. This suggests that the DRL controller has a better disturbance rejection ability than the NMPC. Furthermore, the control forces for the DRL and NMPC while tracking the reference are illustrated in Fig. 10. In contrast to the simulations, it is clear that the control forces




Fig. 10. Control forces for the DRL controller and the MPC while tracking a sinusoidal trajectory in the Charles River. (a) NMPC (b) DRL.

of the NMPC are smoother than those of the DRL controller. Likewise, the average energy consumption of the NMPC and the DRL is 3.978 N and 5.414 N, respectively. Then, the higher energy consumption and the larger variation within the actuator signals of the DRL is a consequence of its enhanced robustness ability.

V. CONCLUSION AND FUTURE WORK

In this paper, we propose a DRL approach for controlling the trajectory tracking of an ASV. A differentiable non-linear dynamic model is used to construct an efficient simulation environment. We use a deep DRL approach to train the tracking control policy in a simulation environment with a series of sinusoidal curves as a reference trajectory. To account for environmental disturbances, noisy measurements, and non-ideal actuators, multiple effective reward functions were delicately designed for the DRL tracking control policy. Numerical simulations subject to model-based environmental disturbances, and experiments in natural water show that the DRL policy can accurately track trajectories. In addition, in comparison to a state-of-the-art NMPC approach, the DRL performance presents a lower tracking error. In the experimental evaluation, DRL presents higher energy usage, a trade-off to counteract the external perturbations and achieve higher accuracy. This work will be extended in two directions. First, we will use the DRL approach to obtain a more accurate dynamic model for our ASV, which can improve the fidelity of the simulation environment and, thereby, further improve the tracking capabilities. Second, we will consider combining DRL and MPC into one control framework so that we might benefit from both the higher tracking accuracy and the low energy consumption.

REFERENCES

- [1] Z. Liu, Y. Zhang, X. Yu, and C. Yuan, "Unmanned surface vehicles: An overview of developments and challenges," *Annual Reviews in Control*, vol. 41, pp. 71–93, 2016.
- [2] M. Doniec, I. Vasilescu, C. Detweiler, and D. Rus, "Complete SE³ underwater robot control with arbitrary thruster configurations," in *2010 IEEE international conference on robotics and automation*. IEEE, 2010, pp. 5295–5301.
- [3] M. Jeong and A. Q. Li, "Risk vector-based near miss obstacle avoidance for autonomous surface vehicles," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 1805–1812.
- [4] P. Corke, C. Detweiler, M. Dunbabin, M. Hamilton, D. Rus, and I. Vasilescu, "Experiments with underwater robot localization and tracking," in *2007 IEEE International Conference on Robotics and Automation (ICRA)*, 2007, pp. 4556–4561.
- [5] M. Doniec, I. Vasilescu, C. Detweiler, and D. Rus, "Complete SE³ underwater robot control with arbitrary thruster configurations," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2010, pp. 5295–5301.
- [6] J. Curcio, J. Leonard, and A. Patrikalakis, "SCOUT - a low cost autonomous surface platform for research in cooperative autonomy," in *MTS/IEEE Proceedings of OCEANS*, 2005, pp. 725–729.
- [7] L. Paull, S. Saedi, M. Seto, and H. Li, "AUV navigation and localization: A review," *IEEE Journal of Oceanic Engineering*, vol. 39, no. 1, pp. 131–149, 2014.
- [8] A. Dhariwal and G. S. Sukhatme, "Experiments in robotic boat localization," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2007, pp. 1702–1708.
- [9] M. N. Azzeria, F. A. Adnanb, and M. Z. M. Zaina, "Review of course keeping control system for unmanned surface vehicle," *Jurnal Teknologi (Sciences & Engineering)*, vol. 5, no. 74, pp. 11–20, 2015.
- [10] L. Li, S. Wang, Y. Zhang, S. Song, C. Wang, S. Tan, W. Zhao, G. Wang, W. Sun, F. Yang *et al.*, "Aerial-aquatic robots capable of crossing the air-water boundary and hitchhiking on surfaces," *Science Robotics*, vol. 7, no. 66, p. eabm6695, 2022.
- [11] S. Yan, Z. Wu, J. Wang, Y. Huang, M. Tan, and J. Yu, "Real-world learning control for autonomous exploration of a biomimetic robotic shark," *IEEE Transactions on Industrial Electronics*, 2022.
- [12] H. R. Karimi and Y. Lu, "Guidance and control methodologies for marine vehicles: A survey," *Control Engineering Practice*, vol. 111, p. 104785, 2021.
- [13] N. Wang, H. R. Karimi, H. Li, and S. Su, "Accurate trajectory tracking of disturbed surface vehicles: A finite-time control approach," *IEEE/ASME Transactions on Mechatronics*, vol. 24, no. 3, pp. 1064–1074, 2019.
- [14] C. Paliotta, E. Lefeber, K. Y. Pettersen, J. Pinto, M. Costa, and J. T. de Figueiredo Borges de Sousa, "Trajectory tracking and path following for underactuated marine vehicles," *IEEE Transactions on Control Systems Technology*, vol. 27, no. 4, pp. 1423–1437, 2019.
- [15] M. Van, "An enhanced tracking control of marine surface vessels based on adaptive integral sliding mode control and disturbance observer," *ISA Transactions*, vol. 90, pp. 30–40, 2019.
- [16] H. Wang, Y. Tian, and C. Vasseur, "Optimal trajectory tracking controller design: Model based to model free methods," in *2014 International Conference on Mechatronics and Control (ICMC)*, 2014, pp. 617–622.
- [17] W. Wang, L. A. Mateos, S. Park, P. Leoni, B. Gheneti, F. Duarte, C. Ratti, and D. Rus, "Design, modeling, and nonlinear model predictive tracking control of a novel autonomous surface vehicle," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 6189–6196.
- [18] W. Wang, B. Gheneti, L. A. Mateos, F. Duarte, C. Ratti, and D. Rus, "Roboat: An autonomous surface vehicle for urban waterways," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Nov 2019, pp. 6340–6347.
- [19] W. Wang, T. Shan, P. Leoni, D. Fernández-Gutiérrez, D. Meyers, C. Ratti, and D. Rus, "Roboat ii: A novel autonomous surface vessel for urban environments," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 1740–1747.
- [20] W. Wang, Z. Wang, L. Mateos, K. W. Huang, M. Schwager, C. Ratti, and D. Rus, "Distributed motion control for multiple connected surface vessels," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 11 658–11 665.
- [21] W. Wang, N. Hagemann, C. Ratti, and D. Rus, "Adaptive nonlinear model predictive control for autonomous surface vessels with largely varying payload," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 7337–7343.
- [22] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. MIT Press, 1998.
- [23] A. B. Martinsen, A. M. Lekkas, S. Gros, J. A. Glomsrud, and T. A. Pedersen, "Reinforcement learning-based tracking control of usvs in varying operational conditions," *Frontiers in Robotics and AI*, vol. 7, 2020.
- [24] Q. Zhang, W. Pan, and V. Reppa, "Model-reference reinforcement learning for collision-free tracking control of autonomous surface vehicles," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 7, pp. 8770–8781, 2022.
- [25] N. Wang, Y. Gao, H. Zhao, and C. K. Ahn, "Reinforcement learning-based optimal tracking control of an unknown unmanned surface vehicle," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 32, no. 7, pp. 3034–3045, 2021.
- [26] N. Wang, Y. Gao, C. Yang, and X. Zhang, "Reinforcement learning-based finite-time tracking control of an unknown unmanned surface vehicle with input constraints," *Neurocomputing*, vol. 484, pp. 26–37, 2022.
- [27] Y. Zhao, X. Qi, Y. Ma, Z. Li, R. Malekian, and M. A. Sotelo, "Path following optimization for an underactuated usv using smoothly-convergent deep reinforcement learning," *IEEE Transactions on Intelligent Transportation Systems*, 2020.
- [28] T. Zhang, R. Tian, H. Yang, C. Wang, J. Sun, S. Zhang, and G. Xie, "From simulation to reality: A learning framework for fish-like robots to perform control tasks," *IEEE Transactions on Robotics*, pp. 1–18, 2022.
- [29] Y. Cheng and W. Zhang, "Concise deep reinforcement learning obstacle avoidance for underactuated unmanned marine vessels," *Neurocomputing*, vol. 272, pp. 63–73, 2018.
- [30] J. Woo, C. Yu, and N. Kim, "Deep reinforcement learning-based controller for path following of an unmanned surface vehicle," *Ocean Engineering*, vol. 183, pp. 155–166, 2019.
- [31] C. Zhou, Y. Wang, L. Wang, and H. He, "Obstacle avoidance strategy for an autonomous surface vessel based on modified deep deterministic policy gradient," *Ocean Engineering*, vol. 243, p. 110166, 2022.
- [32] 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," *arXiv preprint arXiv:1509.02971*, 2015.
- [33] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.
- [34] T. I. Fossen, M. Breivik, and R. Skjetne, "Line-of-sight path following of underactuated marine craft," *IFAC proceedings volumes*, vol. 36, no. 21, pp. 211–216, 2003.
- [35] G. E. Uhlenbeck and L. S. Ornstein, "On the theory of the brownian motion," *Physical review*, vol. 36, no. 5, p. 823, 1930.
- [36] A. Gonzalez-Garcia and H. Castañeda, "Adaptive integral terminal super-twisting with finite-time convergence for an unmanned surface vehicle under disturbances," *International Journal of Robust and Nonlinear Control*, vol. 32, no. 18, pp. 10 271–10 291, 2022.
- [37] T. I. Fossen, *Handbook of Marine Craft Hydrodynamics and Motion Control*. John Wiley & Sons, 2011.
- [38] T. Shan, B. Englot, D. Meyers, W. Wang, C. Ratti, and D. Rus, "Lio-sam: Tightly-coupled lidar inertial odometry via smoothing and mapping," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 5135–5142.