

A reinforcement learning control approach for underwater manipulation under position and torque constraints

Ignacio Carlucho^a, Mariano De Paula^b, Corina Barbalata^a, Gerardo G. Acosta^b

icarlucho@lsu.edu, mariano.depaula@fio.unicen.edu.ar, cbarbalata@lsu.edu, ggacosta@fio.unicen.edu.ar

^a Department of Mechanical Engineering, Louisiana State University, Baton Rouge, USA

^b INTELYMEC Group, Centro de Investigaciones en Física e Ingeniería del Centro CIFICEN – UNICEN – CICpBA – CONICET, 7400 Olavarría, Argentina

Abstract—In marine operations underwater manipulators play a primordial role. However, due to uncertainties in the dynamic model and disturbances caused by the environment, low-level control methods require great capabilities to adapt to change. Furthermore, under position and torque constraints the requirements for the control system are greatly increased. Reinforcement learning is a data driven control technique that can learn complex control policies without the need of a model. The learning capabilities of these type of agents allow for great adaptability to changes in the operative conditions. In this article we present a novel reinforcement learning low-level controller for the position control of an underwater manipulator under torque and position constraints. The reinforcement learning agent is based on an actor-critic architecture using sensor readings as state information. Simulation results using the Reach Alpha 5 underwater manipulator show the advantages of the proposed control strategy.

Index Terms—Underwater manipulation, Reinforcement learning, Neural networks, Intelligent control, Deep Deterministic Policy Gradient

I. INTRODUCTION

In the last decades the importance of underwater manipulators in marine operations has grown continuously. Most robotic underwater industrial applications are conducted with Remotely Operated Vehicles (ROV) where a human operator is tasked with the remote operation of the manipulator [1]. However, due to the limited number of expert operators and the high cost of operations, the industry is migrating towards Autonomous Underwater Vehicles (AUV) [2]. In this type of scenario, a manipulator, usually electric, is mounted on the AUV and operates autonomously, however, this requires a robust and adaptable control system. Furthermore, in autonomous missions different types of operational constraints may appear, such as specific joint constraints that must be followed in order to avoid collisions [3], or decreased joint torque due to faulty motors. These constraints increase the need for designing complex control systems for robust manipulation.

One of the most used low-level controllers for manipulators is the classical Proportional Integrative Derivative (PID) controller [4]. This is due mostly to its simplicity of use and low

computational requirements. However, when it is used for controlling manipulators arms, it must cope with highly non-linear systems. This issue is aggravated in underwater environments where unknown disturbances affect the behaviour of the arm. Furthermore, for underwater manipulators, controllers are used under the assumption that the arm will move slowly and as such it is possible to decouple each degree of freedom, something that is not true for every application [5]. Researchers have also turned to non-linear optimal control techniques as a viable option, since they allow to optimize a cost function under different metrics. One of these techniques is Model Predictive Control (MPC) [6], used successfully for controlling a different number of underwater robots [7], [8]. However, one of the drawbacks of this technique is that it requires an accurate model of the plant in order to work properly, not a trivial matter in underwater robotics [9].

Data driven control techniques have appeared as an alternative for systems with complex or unknown models. One of these techniques is Reinforcement Learning (RL) [10]. In the RL framework, the robot arm control problem can be formulated as a Markov Decision Process (MDP) [11]. Solving a RL problem consists in iteratively learning a task from interactions to achieve a goal. During learning, an artificial agent (controller) interacts with the target system (arm) by taking an action (torque command), that makes the robot evolve from its current state $x_t \in \mathbb{X} \subseteq \mathbb{R}^n$ to x_{t+1} . The agent then receives a numerical signal r_t , called reward, which provides a measure of how good (or bad) the action taken at time t is in terms of the observed state transition. Many works have used the RL paradigm for controlling AUVs in underwater environment [12]. However, this technique has not yet been applied to underwater manipulators.

The main contribution of this work is the development of a reinforcement learning based control system for the low-level control of an electric underwater manipulator under position and torque constraints. Our reinforcement learning formulation is based on the Deep Deterministic Policy Gradient (DDPG) algorithm [13]. The proposed method uses an actor critic structure, where the actor is a function that maps system

states to actions and the critic is a function that assess actions chosen by the actor. Deep Neural Networks (DNN) are used as function approximators for the actor and critic. Results in simulation show the advantages of our proposal when controlling a simulated version of the Reach 5 Alpha manipulator, shown in Fig. 1. The proposed controller is compared with a MPC, showing that the RL controller is able to outperform the MPC.

The article is structured as follows, Section II presents an overview of related works followed by Section III that introduces the basics of RL control utilized in our formulation. In Section IV the details of our implementations are described, in Section V we present the results obtained with our proposed control scheme, and finally Section VI presents the overall conclusions of the proposed work.

II. RELATED WORKS

Designing control systems under constraints considerations for robotic manipulators appeared due to the need of robots to interact with the environment. Some of the most fundamental approaches focused on designing motion/interaction control systems by using a hybrid control formulation [14], [15]. In this approach, the constraints are expressed based on the end-effector's working space, and are used to decide the type of control law (either a motion control law or a force regulator). Nevertheless, constraints are not imposed only by the interaction with the environment, but are also required for cases when the robotic manipulator has to adjust its working space due to obstacles in the environment, or faults in the robotic system. In [16] a passivity-based kinematic control law is proposed under joint velocity limits considerations. The method proposed can be adapted to different feedback control laws, but can be applied only for redundant systems. An adaptive neural-network control for robotic manipulators with parametric uncertainties and motion constraints is proposed in [17]. The simulation and experimental results with a 2 degree of freedom (DOF) planar manipulator show the velocity constraints always being respected, but steady-state errors are present. Deep neural-network approaches have become popular in the past years. An example is given in [18] where an obstacle avoidance control law is designed for redundant manipulators. The problem is reformulated as a Quadratic Programming (QP) problem in the speed level, and a deep recurrent neural network is designed to solve the QP problem in an online way. Although the simulation results show that the robot is capable of avoiding the obstacles while tracking the predefined trajectories, an experimental evaluation is not presented.

The requirements of adaptability and the difficulties with modeling have lead research towards intelligent control methods, such as reinforcement learning. Many RL methods have been previously applied to the control of manipulators [19]. In [20] asynchronous reinforcement learning was used to train a series of robotic manipulators to solve a door opening task. In [21] a mobile manipulator task is solved by Proximal Policy Optimization (PPO), a well known RL algorithm [22], in



Fig. 1. Reach 5 Alpha underwater manipulator [26]

combination with a Deep Neural Network. Specifically for the underwater environments, RL has been used as a control techniques in several previous works [23]–[25]. However, the majority of these works focus on the of AUVs. Works utilizing RL in underwater manipulators are lacking in the literature.

III. REINFORCEMENT LEARNING BASED CONTROL

From the point of view of classical and modern control, the design strategy of a control system is based on the complete knowledge of the dynamics of the system under study. Assuming that there is a model with the adequate capacity to describe the dynamics of the system, generally through a set of differential equations, the control problem is reduced to design a controller (or agent) capable of generating the adequate control actions for the system to achieve a given objective, goal, task or specific desired behavior. In this way, the performance capabilities of the conventional designed control systems are excessively dependent on the mathematical models used to describe the behavior of the dynamic systems to be controlled. However, underwater manipulation is a complex decision making problem in which, the presence of uncertainty in dynamics is ubiquitous and, consequently, it is of paramount importance designing and using controllers with suitable adaptation capabilities.

Markov decision processes are models for sequential decision making problems when outcomes are uncertain [11]. In our formulation, we consider a finite-horizon Markov decision process with a $1, 2, \dots, T$ decisions and $T - 1$ visited stages [27]. That is, the decision (action) at time t is made at the beginning of stage t which corresponds to the time interval from t to the next $t + 1$. So, at any stage, or discrete time, t , the system is at a state \mathbf{x}_t . In this sense, we have a finite set X of system states, such that $\mathbf{x}_t \in X, \forall t = 1, \dots, T$. The decision maker observes state $\mathbf{x}_t \in X$ at stage t and it may choose an action \mathbf{u}_t from the set of finite allowable actions

U generating cost $L(\mathbf{x}_t, \mathbf{u}_t)$. Moreover, we let $p(\cdot|\mathbf{x}_t, \mathbf{u}_t)$ denote the probability distribution or transition probabilities of obtaining states $\mathbf{x}' = \mathbf{x}_{t+1}$ at stage $t + 1$.

A deterministic Markovian decision rule at state \mathbf{x}_t is a function $\psi_t : \mathbf{x}_t \rightarrow \mathbf{x}_t$ which maps the action choice given at state \mathbf{x}_t . It is called deterministic because it chooses an action with certainty and Markovian (memoryless) since it depends only on the current system state. We let D_t denote the set of possible deterministic Markovian decision rules at stage t . D_t is a subset of more general rules where the action may depend on the past history of the system and actions may not be chosen with certainty but rather according to a probability distribution.

A policy or strategy specifies the decision rules to be used at all stages and provides the decision maker with a plan of which action to take given stage and state. That is, a policy π is a sequence of decision rules and we restrict ourselves to ranking policies π_t belonging to the set D_t of deterministic Markov policies (if randomized policies were included, the set of policies would not be countable). In some problems, the decision maker only focuses on this subset of policies, e.g. because randomized policies are hard to manage in practice or restrictions in management strategy. Moreover, if the states at a given time step t corresponds to different physical locations implementation of policies having a single action at each location may only be acceptable.

Under the above summarized framework of Markov decision processes, the reinforcement learning can be formalized where an RL agent located in its environment chooses an action, from the set of available ones, at every discrete time step t based on the current state of the system \mathbf{x}_t . In return, at each time step, the agent receives a reward signal that quantifies the quality of the action taken in term of the goal of the control task. In this paper we only consider one criterion of optimality, namely the expected total cost criterion, so the objective is to obtain an optimal policy π^* that satisfies:

$$L^* = \max L_\pi = \max_{\pi} \mathbb{E}\{R_t | \mathbf{x}_t = \mathbf{x}\} \quad (1)$$

where r_t is the instantaneous reward obtained at time step t and R_t is the cumulative reward, such that $R_t = \sum_{k=0}^{\infty} \gamma^k r_{t+k+1}$.

The basic RL algorithms discussed in the literature have been extensively developed to solve reinforcement learning problems without need of a dynamic model and when the state spaces and action are finite, which means that the value functions support a tabular representation [10], [28], [29]. In order to find an approximate solution to the control problem it is possible to obtain a discretized form of the space of states and/or actions [30]–[32] and then applying the RL algorithms that use discrete spaces. However, as the granularity of the representation increases, the computational implementations suffer the so-called curse of dimensionality, which consists of an exponential increase in the computational complexity of the problem due to the increase in the dimension (or number) of the state-action pairs to be selected. This makes it impossible to construct a value function for the problem in question, since

the agent has a low probability of "visiting" the same state, or a state-action pair, more than once depending on whether it is working with the state value function or value-action function, respectively.

In the underwater manipulation problem we have to deal with dynamic systems where the states and the applied actions are defined in real domains (continuous spaces) which imposes an important limitation for tabular representation of the value functions. To overcome this drawback, functional approximation techniques have emerged including inductive models to attempt generalizing the value function. Since a few years ago powerful brain inspired deep neural networks [33] have been introduced as functions approximations into the RL framework giving rise to deep reinforcement learning methodologies [34]. For instance, the Deep Deterministic Policy Gradient (DDPG) algorithm [13] is one of the most spread deep RL algorithms that utilizes an actor-critic formulation together with neural network as function approximators to obtain a deterministic optimal policy. In the actor-critic formulation, the role of the actor is to select an action based on the policy, such that $\mathbf{u} = \pi(\mathbf{x}_t)$. The critic on the other hand, gives feedback of how good or bad the selected action was. In the DDPG algorithm the state-action value function $Q(\mathbf{x}_t, \mathbf{u}_t)$ is used as a critic. This function is defined as:

$$Q^\pi(\mathbf{x}_t, \mathbf{u}_t) = \mathbb{E}\{R_t | \mathbf{x}_t, \mathbf{u}_t\} = \mathbb{E}\left\{\sum_{k=0}^{\infty} \gamma^k r_{k+t+1} | \mathbf{x}_t, \right\} \quad (2)$$

The update to the state-action value function can then be performed as:

$$Q^w(\mathbf{x}_t, \mathbf{u}_t) = \mathbb{E}\{r_{x_t, u_t} + \gamma Q^w(\mathbf{x}_{t+1}, \mathbf{u}_{t+1})\} \quad (3)$$

where Q^w is a differentiable parameterized function, so that $Q^w \approx Q^\pi$.

For the actor we consider a function π that parameterizes states directly into actions with parameters θ , thus $\pi(x_t|\theta)$. And we define a performance objective function $L(\pi_\theta) = \mathbb{E}\{r^\gamma | \mu\}$ and a probability distribution ρ , then the performance as an expectation can be written as:

$$L(\mu_\theta) = \int \rho^\mu r(\mathbf{x}_t, \mu) dx = \mathbb{E}[r(\mathbf{x}_t, \mu_\theta(\mathbf{x}_t))] \quad (4)$$

and by applying the chain rule to the expected return, we can then write:

$$\nabla_\theta L = \mathbb{E}[\nabla_\theta \mu_\theta(\mathbf{x}) \nabla_{\mathbf{u}} Q^\mu(\mathbf{x}, \mathbf{u})] \quad (5)$$

where Eq. (5) is the deterministic gradient of policies, as demonstrated in [35]. As indicated previously, both the actor and the critic can be represented by function approximators, where deep neural networks are commonly used since they allow to work with continuous state spaces. However, the nonlinearities of these networks and the training procedures used made it difficult for algorithms to converge, however, recent events have repaired these problems. The main causes of the lack of convergence were the correlation of the samples

used for training and the correlation between the updates of the network Q [36].

The first of these issues was addressed by implementing a replay buffer that stored state transitions, the actions applied, and the rewards earned. The agent is then trained using minibatches of transitions that are randomly selected from the replay buffer [37]. The second problem was solved by incorporating target networks, which are a direct copy of the actor's and critic's networks called π' and Q' , with the parameters θ' and ω' respectively, and which are periodically updated according to the parameter τ , so that $\theta' \leftarrow \theta\tau + (1 - \tau)\theta'$ and $\omega' \leftarrow \omega\tau + (1 - \tau)\omega'$ where $\tau \ll 1$.

IV. IMPLEMENTATION DETAILS

A simulated version of the Reach Alpha 5 manipulator is used. The Reach Alpha 5 is a 5DOF underwater manipulator, capable of lifting 2kg, and is able to operate in depths up to 300m. The manipulator is shown in Fig. 1.

For our proposed formulation, the state (\mathbf{x}_t) of the RL agent is determined by the joint position ($\mathbf{q} \in \mathbb{R}^n$) in [rads] and joint velocity $\dot{\mathbf{q}} \in \mathbb{R}^n$ in [rads/s], together with the desired joint position ($\mathbf{q}_{req} \in \mathbb{R}^n$) in [rad], such that the state is determined as: $\mathbf{x}_t = [\mathbf{q}_t, \dot{\mathbf{q}}, \mathbf{q}_{req}]$, with n being the DOFs of the manipulator. The goal of the agent is to achieve a determined joint position, where the request comes to the agent by higher a layer in the control hierarchy.

A fully connected feed forward network is used for both the actor and critic with two hidden layers of 400 and 300 units each. As activation function Leaky ReLus are used for the hidden networks, with Tanh used for the output neurons of the actor. The learning rate used is 0.0001 and 0.001 for the actor and critic respectively, with Adam used as an optimizer. A decay rate of 0.96 is applied to each learning rate after 100 thousands training steps.

In order to be able to achieve the required position and torque constraints, the reward function was developed as follows: If the position is within the allowed bounds the reward is a Gaussian function that penalizes the agent when the joint position is not close to the request and gives a positive reward when the position matches or is close to the request. On the other hand, when the agent goes over the allowed bounds it is penalized with a high negative number (-10). Formally the reward is defined as follows:

$$r_t = \begin{cases} -1 + e^{-\frac{1}{2}\left(\frac{\mathbf{x} - \mathbf{x}_{ref}}{\sigma}\right)^2}, & \text{if } \mathbf{x}_{min} < \mathbf{x}_t < \mathbf{x}_{max} \\ -10, & \text{otherwise} \end{cases} \quad (6)$$

with σ being a parameter that shapes the Gaussian function. For the experiments presented here we utilized $\sigma = 0.018$.

For the training of the agent, a different random goal was selected in the allowed work space for each epoch of training. The agent was trained for a total of 2000 epochs, with each epoch lasting 20 seconds of real time operation. Initially random noise is added to the actions of the agent to allow for exploration, such that $\mathbf{u}_t = \pi(\mathbf{x}_t) + \epsilon N$, with N being Ornstein-Uhlenbeck noise and ϵ linearly decaying over time

from 1 to 0.1. The minibatch size used for training is 64, with $\tau = 0.001$ and $\gamma = 0.99$.

V. RESULTS

For the presented results the agent was trained as previously stated and the policy is now executed without the addition of exploratory noise, effectively making $\epsilon = 0$. Furthermore, the nonlinear model has been degraded with some parameters changed randomly to test the adaptability of the agent and random noise is introduced to the velocity and position readings. While the Reach 5 Alpha has 5 DOF, the last degree of freedom corresponds to the gripper joint, which we are not interested in controlling, as such all results are shown for the first 4 DOF.

A simulation using the trained RL agent was ran on the Reach Alpha manipulator under normal operative conditions with a reference joint position of $\mathbf{x}_{ref} = (2.64, 0.26, -1.47, 0.82)$ [rad]. Fig. 2a shows the joint position while being controlled by the RL agent, Fig. 2b shows the control actions and Fig. 2c shows the position errors. In Fig. 2a it can be seen how the agent reaches the desired position in less than two seconds, without any overshoot, even when the requested position requires a long rotation of over two radians in Joint 1. The lack of overshoots demonstrates how the agent is capable of behaving without breaking any of the position constraints imposed during training. The agent utilizes the maximum torque initially available, as can be seen in Fig. 2b, and then utilizes small corrections to keep the joints in position. Fig. 2c shows how the errors are rapidly reduced, with no steady state error present.

Another example shows the behaviour of the arm when a completely different reference point is selected, $\mathbf{x}_{ref} = (-1.78, 0.11, -2.14, -2.26)$ [rad]. Fig. 3a shows the obtained position when using the RL controller. As can be seen the requested position is reached again in under two seconds, without any overshoot. Again the agent utilizes high levels of torque initially, with lower levels after the requested joints position have been reached, as depicted in 3b. Furthermore, no steady state error is present as illustrated in 3c.

The example presented here aims to test the behavior of the agent under torque constraints. In this example, the torque output of Joint 1 is reduced by 75%, with a desired requested position of $\mathbf{x}_{ref} = (2., -1., -1.75, 1.5)$ [rad]. In Fig. 4a the achieved position are shown, where it can be seen that the agent is capable of rapidly reaching the desired positions for Joints 2, 3 and 4, while Joint 1 takes around 5 seconds due to the new restrictions in torque, but no overshoot is present. This can also be seen in Fig. 4c where the errors are shown, where Joint 1 takes longer to reach the request, however it can also be seen that no steady state error or overshoot are present. Additionally, Fig. 4b shows the torque output, where the reduced torque of Joint 1 can be clearly seen.

A. Comparative results

In this section we introduce a comparison between the proposed RL agent and a MPC controller. The cost function

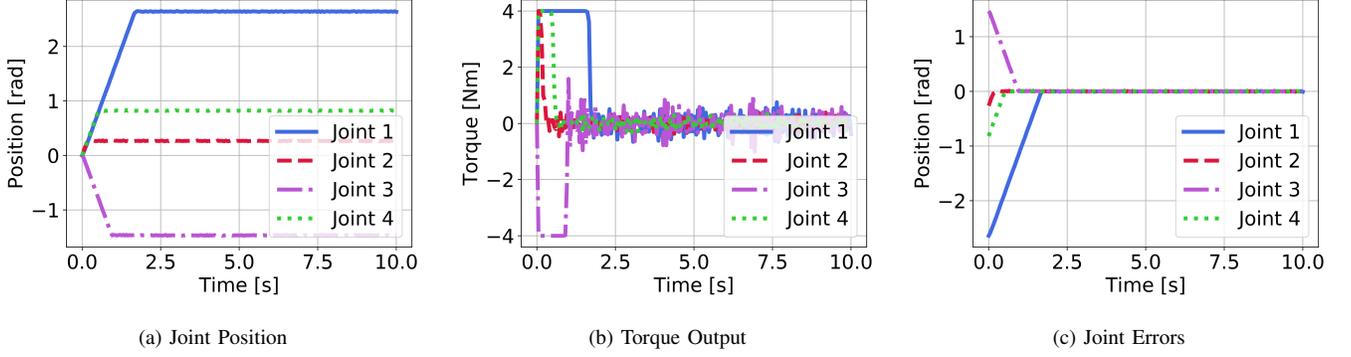


Fig. 2. Test 1 of the proposed RL algorithm with $\mathbf{x}_{ref} = (2.64, 0.26, -1.47, 0.82)$ [rad]

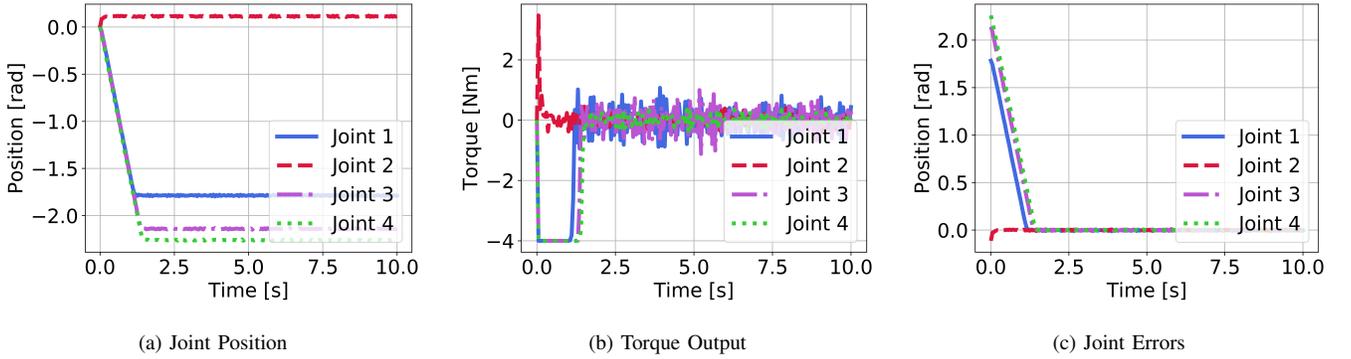


Fig. 3. Test 2 of the proposed RL algorithm with $\mathbf{x}_{ref} = (-1.78, 0.11, -2.14, -2.26)$ [rad]

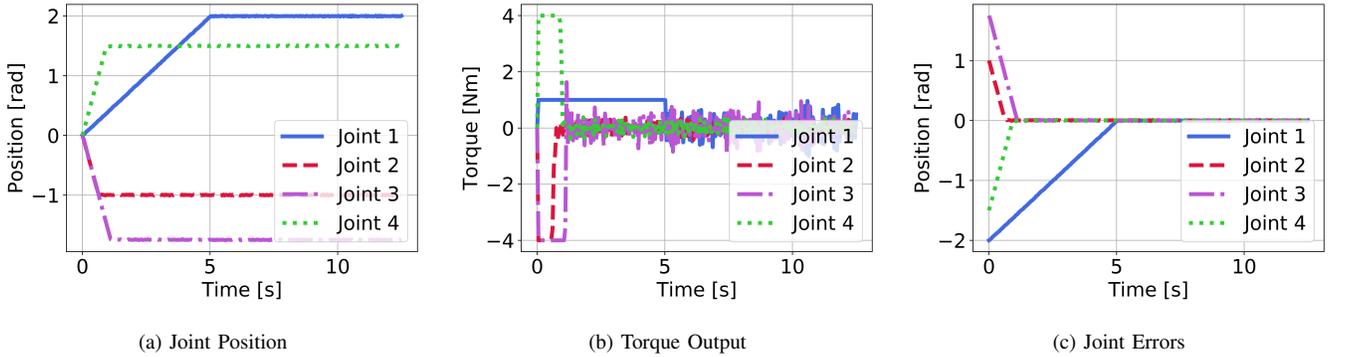


Fig. 4. Test 3 of the proposed RL algorithm under Torque constraints with $\mathbf{x}_{ref} = (2., -1., -1.75, 1.5)$ [rad]

of the MPC controller is $J = \sum_{k=0}^N \|\mathbf{x}_{ref} - \mathbf{x}_{t+k}\|_{\mathbf{Q}(t)}^2 + \|\Delta \mathbf{u}_{t+k}\|_{\mathbf{R}(t)}^2$. The gains \mathbf{Q} and \mathbf{R} of the cost function were tuned accordingly. As previously, the Reach 5 Alpha simulated arm is used where a desired position (\mathbf{x}_{ref}) should be attained.

An experiment is presented when a desired reference position of $\mathbf{x}_{ref} = (2.13, -0.74, -1.03, 2.51)$ [rad] is selected. Fig. 5 shows the obtained joint position when using the proposed RL agent, while Fig. 6 shows the results when utilizing the baseline MPC controller. While the MPC controller is able to reach the required position for Joint 1 in less than 2.5

seconds (Fig. 6a), it takes the rest of the joints longer. On the other hand, the RL agent is faster and presents no overshoot (Fig. 5a). In addition, the RL agent presents no steady state error, as can be seen in Fig. 5c, while the MPC shows some error in the steady state as Fig. 6c shows. While the control actions both utilize high levels of torque initially, the MPC shown in Fig. 6b seems to require less torque once the steady state is reached as compared to the RL agent, in Fig. 5b.

A series of experiments were performed in which random

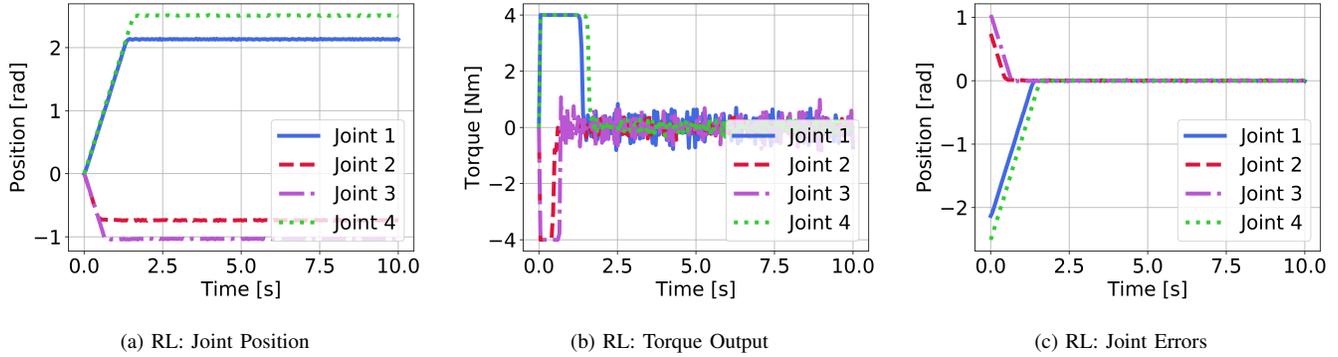


Fig. 5. Comparative results using RL for $\mathbf{x}_{ref} = (2.13, -0.74, -1.03, 2.51)$ [rad]

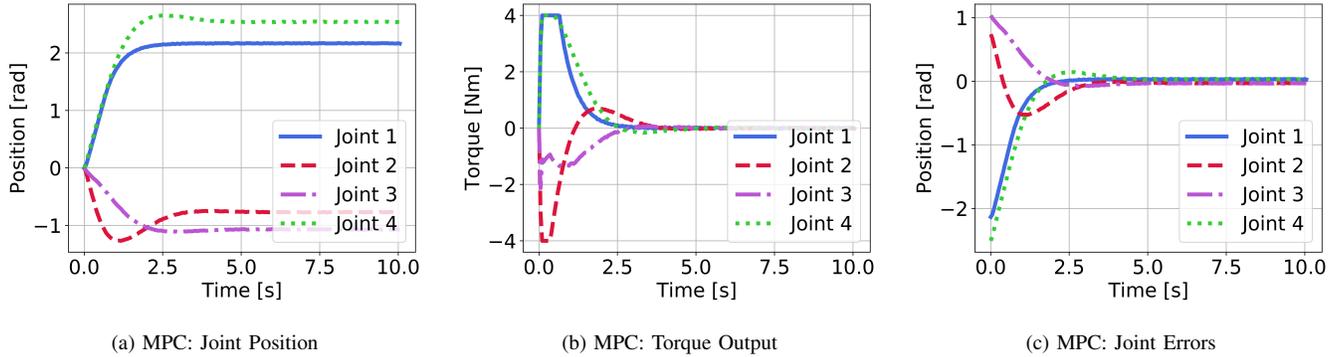


Fig. 6. Comparative results using MPC for $\mathbf{x}_{ref} = (2.13, -0.74, -1.03, 2.51)$ [rad]

positions were selected and a number of metrics were obtained in order to compare the performance of the two algorithms. These metrics include the average energy consumed (E) in Joules, the Root Mean Square Error (RMSE), Mean Integral Error (MIE), the Mean Steady Steate Error (MSSE), the Overshoot (OS) in percentage, and the Settling Time (ST) in seconds. A set of 20 experiments were conducted where the obtained results can be seen in Table I.

TABLE I
COMPARATIVE RL VS MPC

Algorithm	E [J]	RMSE	MIE	MSSE	OS[%]	ST [s]
RL	19.86	0.23	58.27	0.0033	1.43	6.26
MPC	13.8	0.27	97.64	0.033	18.21	7.96

The presented metrics show a much more favorable performance of the RL agent. Both MIE and MSSE are significantly lower, while the RMSE also presents lower values for the RL implementation. The OS is practically non existent in the RL agent as compared to the MPC, while the ST is over a second less for the RL, making it the faster solution. The only disadvantage is seen with regards to the energy consumption (E) which is lower for the MPC controller. However, the RL controller is not taking into account the energy consumption

as this was not the main focus of the proposal. Overall, the presented results show the benefits of the RL controller.

VI. CONCLUSIONS

In this article we presented a novel strategy for the low-level control of an underwater manipulator under position and torque constraints based on a reinforcement learning formulation. The actions selected by the actor are directly the torque commands sent to the manipulator, while the state is determined by the current position and velocity, together with the desired joint positions. By including the goal in the state we are able to generalize over different control requests, a fundamental requirement for a control system.

The data driven approach provided by RL avoids the use of complex models and is able to adapt to changes in the operative conditions. For instance, sudden changes that limit the normal operation of the system, such as obstacles in the working space, failure of any engine, and others, can cause reduced joint movement leading to limited range for joint positions and/or limited torque range. Such constraints can be difficult to surpass using classical controllers due to the lack of accurate model information and poor tuning of the controller. On the other hand, reinforcement learning controllers, are able to obtain highly non linear policies that can operate within the required boundaries and are able to adapt to new requirements.

As future works, the authors suggest the implementation of the algorithm in the Reach 5 Alpha arm as well as in other manipulators to test the adaptability of the proposal. Furthermore, investigating the possibility of reducing the higher energy consumption, when comparing with the MPC, could be of interest for autonomous operations.

REFERENCES

- [1] Satja Sivčev, Joseph Coleman, Edin Omerdić, Gerard Dooly, and Daniel Toal. Underwater manipulators: A review. *Ocean Engineering*, 163:431 – 450, 2018.
- [2] A.M. Yazdani, K. Sammut, O. Yakimenko, and A. Lammas. A survey of underwater docking guidance systems. *Robotics and Autonomous Systems*, 124:103382, 2020.
- [3] Xanthi S. Papageorgiou, Herbert G. Tanner, Savvas G. Loizou, and Kostas J. Kyriakopoulos. Switching manipulator control for motion on constrained surfaces. *Journal of Intelligent and Robotic Systems*, 62:217–239, 2011.
- [4] J. G. Ziegler and N. B. Nichols. Optimum Settings for Automatic Controllers. *Journal of Dynamic Systems, Measurement, and Control*, 115(2B):220, jun 1993.
- [5] Corina Barbalata, Matthew W. Dunnigan, and Yvan R. Petillot. Coupled and decoupled force/motion controllers for an underwater vehicle-manipulator system. 2018.
- [6] Carlos E. García, David M. Prett, and Manfred Morari. Model predictive control: Theory and practice—a survey. *Automatica*, 25(3):335 – 348, 1989.
- [7] C. Shen, Y. Shi, and B. Buckham. Trajectory tracking control of an autonomous underwater vehicle using lyapunov-based model predictive control. *IEEE Transactions on Industrial Electronics*, 65(7):5796–5805, 2018.
- [8] Guoxing Bai, Yu Meng, Li Liu, Weidong Luo, and Qing Gu. Review and comparison of path tracking based on model predictive control. volume 8, page 1077, 2019.
- [9] Thor I Fossen et al. *Guidance and control of ocean vehicles*, volume 199. 1994.
- [10] Richard S. Sutton and Andrew G. Barto. *Reinforcement learning: An Introduction*. MIT Press, 1998.
- [11] George E. Monahan. State of the Art—A Survey of Partially Observable Markov Decision Processes: Theory, Models, and Algorithms. *Management Science*, 28(1):1–16, January 1982.
- [12] Ignacio Carlucho, Mariano De Paula, Sen Wang, Yvan Petillot, and Gerardo G. Acosta. Adaptive low-level control of autonomous underwater vehicles using deep reinforcement learning. *Robotics and Autonomous Systems*, jun 2018.
- [13] D. Lillicrap, T.P. Hunt, J.J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., Wierstra. Continuous control with deep reinforcement learning. *Foundations and Trends® in Machine Learning*, 2(1):1–127, 2016.
- [14] James K Mills and Andrew A Goldenberg. Force and position control of manipulators during constrained motion tasks. *IEEE Transactions on Robotics and Automation*, 5(1):30–46, 1989.
- [15] Tsuneo Yoshikawa. Dynamic hybrid position/force control of robot manipulators—description of hand constraints and calculation of joint driving force. *IEEE Journal on Robotics and Automation*, 3(5):386–392, 1987.
- [16] Yinyan Zhang, Shuai Li, Jianxiao Zou, and Ameer Hamza Khan. A passivity-based approach for kinematic control of manipulators with constraints. *IEEE Transactions on Industrial Informatics*, 16(5):3029–3038, 2019.
- [17] Mingming Li, Yanan Li, Shuzhi Sam Ge, and Tong Heng Lee. Adaptive control of robotic manipulators with unified motion constraints. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 47(1):184–194, 2016.
- [18] Zhihao Xu, Xuefeng Zhou, and Shuai Li. Deep recurrent neural networks based obstacle avoidance control for redundant manipulators. *Frontiers in neurorobotics*, 13:47, 2019.
- [19] M. Deisenroth, C. Rasmussen, and D. Fox. Learning to control a low-cost manipulator using data-efficient reinforcement learning. In *Robotics: Science and Systems*, 2011.
- [20] S. Gu, E. Holly, T. Lillicrap, and S. Levine. Deep reinforcement learning for robotic manipulation with asynchronous off-policy updates. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3389–3396, 2017.
- [21] Cong Wang, Qifeng Zhang, Qiyan Tian, Shuo Li, Xiaohui Wang, David Lane, Yvan Petillot, and Sen Wang. Learning mobile manipulation through deep reinforcement learning. *Sensors (Basel, Switzerland)*, 20(3):939, Feb 2020. 32050678[pmid].
- [22] John Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal policy optimization algorithms. *ArXiv*, abs/1707.06347, 2017.
- [23] Ignacio Carlucho, Mariano De Paula, Sen Wang, Yvan Petillot, and Gerardo G. Acosta. Adaptive low-level control of autonomous underwater vehicles using deep reinforcement learning. *Robotics and Autonomous Systems*, 107:71 – 86, 2018.
- [24] Andres El-Fakdi and Marc Carreras. Two-step gradient-based reinforcement learning for underwater robotics behavior learning. *Robotics and Autonomous Systems*, 61(3):271 – 282, 2013.
- [25] G. Frost, F. Maurelli, and D. M. Lane. Reinforcement learning in a behaviour-based control architecture for marine archaeology. In *OCEANS 2015 - Genova*, pages 1–5, 2015.
- [26] Blueprint Lab. Reach alpha 5 (2020). URL <https://blueprintlab.com/products/reach-alpha/>, 2020.
- [27] K.M. Hee, van and L.B. Hartman. *Application of Markov decision processes to search problems*. Research Report. University of Waterloo, Department of Computer Science, 1992.
- [28] Ignacio Carlucho, Mariano De Paula, Sebastian A. Villar, and Gerardo G. Acosta. Incremental Q-learning strategy for adaptive PID control of mobile robots. *Expert Systems with Applications*, 2017.
- [29] Ignacio Carlucho, Mariano De Paula, and Gerardo G. Acosta. Double q-pid algorithm for mobile robot control. *Expert Systems with Applications*, 137:292 – 307, 2019.
- [30] William S. Lovejoy. Computationally feasible bounds for partially observed markov decision processes. *Operations Research*, 39(1):162–175, 1991.
- [31] L. Crespo and J. Q. Sun. Solution of fixed final state optimal control problems via simple cell mapping. *Nonlinear Dynamics*, 23:391–403, 01 2000.
- [32] LG Crespo and J. Q. Sun. Stochastic optimal control via bellman’s principle. *Automatica*, 39:2109–2114, 12 2003.
- [33] Yann LeCun, Yoshua Bengio, Geoffrey Hinton, Lecun Y., Bengio Y., and Hinton G. Deep learning. *Nature*, 521(7553):436–444, 2015.
- [34] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dhharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.
- [35] David Silver, Guy Lever, Nicolas Heess, Thomas Degris, Daan Wierstra, and Martin Riedmiller. Deterministic Policy Gradient Algorithms. *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, pages 387–395, 2014.
- [36] Scott Fujimoto, Herke van Hoof, and David Meger. Addressing Function Approximation Error in Actor-Critic Methods. In *ICML*, volume 80 of *Proceedings of Machine Learning Research*, pages 1582–1591. PMLR, 2018.
- [37] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proceedings of the 32Nd International Conference on International Conference on Machine Learning - Volume 37*, ICML 15, pages 448–456. JMLR.org, 2015.