

A preliminary experiment combining marine robotics and citizenship engagement using imitation learning

Angelo Odetti * Marco Bibuli * Gabriele Bruzzone *
Cristiano Cervellera * Roberta Ferretti * Mauro Gaggero *
Enrica Zereik * Massimo Caccia *

* *Institute of Marine Engineering, National Research Council of Italy,
Via de Marini 6, Genoa, Italy*
(e-mails: *angelo.odetti@inm.cnr.it, marco.bibuli@cnr.it,
gabriele.bruzzone@cnr.it, cristiano.cervellera@cnr.it,
roberta.ferretti@inm.cnr.it, mauro.gaggero@cnr.it, enrica.zereik@cnr.it,
massimo.caccia@cnr.it*)

Abstract: In this paper, we describe a preliminary experiment of citizenship engagement in the context of marine robotics using imitation learning to train a controller that mimics human behavior. The experiment has been carried out during the *Festival della Comunicazione* in Camogli, Italy, in September 2019. In more detail, citizens have been asked to pilot a small, light, and safe autonomous surface vehicle in front of a crowded public beach with the goal of performing an S-shaped path. The trajectories and controls performed by non-expert human operators have been recorded with the aim of training an imitation system that, after collecting a sufficient number of trajectory-control pairs, has been able to drive the vehicle without human intervention. To learn the human behavior, echo state networks have been employed as approximating architectures. The resulting controller turned out to be very effective in successfully performing the considered experiment with a reduced amount of training trajectories by imitating the human behavior also in unknown situations. The success of this experiment may pave the way to new research processes where citizens are actively engaged.

Keywords: Autonomous surface vehicles, citizen engagement, machine learning, ASV

1. INTRODUCTION

In recent years, the development of technology in the field of autonomous robotic vehicles has faced various legal and societal issues, related to public acceptance of the diffuse use of mobile robots operating in not professionally-structured environments and in the presence of daily activities carried out by human beings. This has pointed out bottlenecks in robot capabilities of understanding heavily unstructured environments and situations, where complex dynamics and interactions appear.

In the case of marine and maritime systems, where strategic research is driven by the lighthouse of the autonomous ship, the possibility of applying control techniques based on artificial intelligence (AI) has received greater and greater attention from the research community. In this perspective, AI is used to handle uncertain and heavily-constrained dynamic systems by providing the ability to adapt to changes in the environment and to implement efficient decisions. Reinforcement learning techniques have been applied, for instance, to ship berthing (Amendola et al., 2018), navigation in restricted waters (Pereira Figueiredo and Pereira Abou Rejaili, 2018), and steering of under-actuated ships (Tuyen et al., 2017). Owing to its capability in capturing helmsman behaviour, deep learning has been proposed to navigate autonomous vessels, and its application has been extended to situation awareness and collision avoidance (Perera, 2018).

However, operations of autonomous marine vehicles are not yet regulated by the International Maritime Organization, and several societal barriers continue to slow down the introduction of unmanned/autonomous robots in civil, commercial, and consumer domains. This is particularly evident when a closer interaction with human beings is required. Furthermore, citizens often perceive port, shipyards, and marine technology as a source of pollution or as a separator between city and sea rather than as an engine of eco-sustainable technological development. Thus, the dissemination of research activities in this field is fundamental to face the aforementioned issues. To this purpose, communication has to be regarded as “a strategically-planned process that starts at the outset of the action and continues throughout its entire lifetime, aimed at promoting the action and its results” (Scherer et al., 2018). In the meantime, the European Community is supporting citizenship engagement as an effective way to connect citizens, experts, and policy makers (Figueiredo Nascimento et al., 2016). In this context, citizenship engagement can pave the way to social acceptance of autonomous robots interacting with humans beings in their daily life.

In this paper, we present a preliminary experiment of citizenship engagement in the research process of a robotic marine vehicle, supported by AI-based approaches. In more detail, we report the results of an experimental setup placed in front of a public beach in the village of Camogli, located in North-West Italy, within the *Festival della Comunicazione* in September



Fig. 1. SWAMP at the *Festival della Comunicazione*, Camogli, Italy, September 14-15, 2019.

2019. This experiment involved a small, light, and safe autonomous surface vehicle (ASV) named SWAMP (see Fig. 1). Such vehicle has been designed and constructed in the laboratories of the Institute of Marine Engineering of the National Research Council of Italy (CNR-INM), to meet the requirements of accessing extremely shallow waters typical of wetlands (Odetti et al., 2018, 2019a). It is equipped with azimuth pumpjet actuators and soft foam-made hulls. The performed experiment consists in a situation where citizens actively contribute to data collection by driving and instructing the ASV. In more detail, volunteers piloted the ASV from a pier near the beach, with the goal of passing through two gates made up by four buoys in the same direction, thus describing an S-shaped path. The trajectories and controls performed by human beings were recorded with the aim of training an imitation system that, after collecting a sufficient number of trajectories and controls pairs, was able to drive the ASV without human intervention. A challenge was to perform training also exploiting inexperienced human pilots, such as kids or young students. A high interest of citizens has been recorded in remotely controlling the robot since this has been perceived as playing a video game.

The paradigm under which a robot learns how to perform tasks and complex manoeuvres exploiting examples from a reference controller, such as a human being, is a well-addressed research topic in the recent literature (see, e.g., Hussein et al., 2017 and the references therein). This paradigm is usually referred to as imitation learning. In its basic form, often called behavioral cloning, it consists in equipping the robot with an approximating architecture that is trained in order to replicate the actions performed by the reference controller. This is obtained through a classic supervised learning approach, in which the patterns are the states visited by the robot and the corresponding targets are the actions performed by the demonstrator, i.e., the human being. In general, such an approach is not the most efficient one in terms of performance optimization (Abbeel and Ng, 2004). However, its intuitive principle is easily understandable by a general audience, which makes it ideal in a scenario where citizenship engagement is a goal that is more important than the optimized execution of tasks.

To learn the human behavior, recurrent neural networks (RNNs) have been employed as approximating architectures. In more detail, we have employed the echo state network (ESN) model (Grigoryeva and Ortega, 2018), in which the weights of the input and recurrent layers are randomly extracted and not trained, and only the linear output layer is optimized, according to the so-called reservoir computing paradigm (Tanaka et al., 2019). This guarantees a very fast training procedure, which is an essential requirement due to the very limited available time

to process data and show the results. In fact, these operations were performed in real time in front of the audience. ESNs have been already applied to robotic control in the literature, including marine vehicles and imitation learning (see, e.g., Antonelo and Schrauwen, 2015; van der Zant et al., 2004).

Despite its overall simplicity, the combination of behavioral cloning and ESNs has allowed the trained robot to successfully perform the considered task, exhibiting a robust behavior and good generalization capabilities starting from various initial points.

The rest of this paper is organized as follows. A description of the SWAMP ASV is provided in Section 2, while the adopted imitation learning techniques for the purpose of control are presented in Section 3. Experimental results are reported in Section 4, and future research perspectives are discussed in Section 5.

2. THE SWAMP ASV

SWAMP¹ (Shallow Water Autonomous Multipurpose Platform) is a fully electric, portable, lightweight, safe, and highly-controllable catamaran. It is equipped with azimuth pumpjet actuators, and it is characterized by a soft-foam, unsinkable hull structure with high modularity and a flexible hardware/software architecture. It was designed and constructed from scratch in the laboratories of CNR-INM to operate in very shallow waters. These characteristics are crucial for its integration in the collective ambience of a beach village like Camogli. The soft hull is suitable for working in coastal areas in the bathing season since an impact with swimmers is much less dangerous with respect to the use of classic rigid structures. Moreover, the absence of protrusive rotating elements enables operating in crowded scenarios, so that people can approach the vehicle without the risk of hurts.

SWAMP is 1.25 m long and with a variable breadth between 0.8 m and 1.2 m. It is composed of two hulls that are double-ended to enhance manoeuvrability. They can be considered as a single vehicle since each one has its own control and propulsion units as well as power and communication systems. All the elements inside the single vehicle are connected via Wi-Fi. The weight of SWAMP is about 38 kg, with a draft of 0.1 m. The design maximum payload is 25 kg, with a consequent design maximum draft equal to 0.15 m. The bottom of the hull is flat and hosts four pumpjet 360 degrees azimuth thrusters specifically designed for this vehicle (Odetti et al., 2019b). One of the main peculiar aspects of SWAMP is the use of a light, soft, and impact-survival flexible structure made up with a sandwich of flexible, light-weight, closed-cell poly-ethylene foam, high density poly-ethylene plates, and pultruded bars. The foam of the hulls is drilled in order to make SWAMP a completely modular catamaran that is able to host various types of tools, such as intelligent systems, samplers, and sensors, together with thrusters. The main advantage of pumpjet motors is that they are flush with the hull, thus minimizing the risks of damages due to possible grounding. They are based on the principle of a vertical axis pump: an impeller sucks water from below the hull, and therefore both velocity and pressure is gained. The water is directed to an external volute and is pushed

¹ The SWAMP ASV was developed during the Ph.D. of A. Odetti at the University of Genoa, Italy. Special thanks go to Gi. Bruzzone, G. Camporeale, M. Giacomelli, and E. Spirandelli of CNR-INM, as well as to M. Altosole of the University of Naples, Italy, and M. Viviani of the University of Genoa.

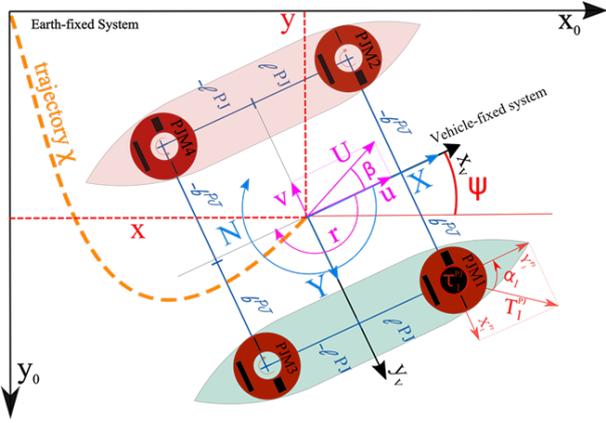


Fig. 2. The reference system for the thrust allocation of the SWAMP vehicle.

towards outlet nozzles in the 360 degrees steerable casing. The nozzles accelerate the flow and a jet of water produces thrust horizontally beneath the flat-bottomed hull.

2.1 SWAMP Dynamic Behavior

SWAMP dynamic behavior can be taken into account using classic ship manoeuvrability models starting from Newton's law. In more detail, the vehicle is considered as a rigid body with three degrees of freedom equipped with an Earth-fixed reference system and a vehicle-fixed mobile one, as shown in Fig. 2. We consider the forward motion along the longitudinal axis x_v , and the drift motion along the transverse axis y_v , as well as the yaw, i.e., the rotation ψ around the vertical axis.

The resultant actions (X_0, Y_0, N_0) on SWAMP in the Earth-fixed coordinate system are difficult to be computed. Thus, we compute rigid body dynamic equilibrium equations of the Earth-fixed system into mobile coordinates. This allows to uniquely identify the position of the ship system, i.e., the trajectory (given by x, y, ψ) in the Earth-fixed system. More specifically, we get

$$\begin{cases} X_0 = \Delta \ddot{x} \\ Y_0 = \Delta \ddot{y} \\ N_0 = I_{zz} \ddot{\psi} \end{cases} \Rightarrow \begin{cases} X = \Delta(\dot{u} - vr - x_G r^2) \\ Y = \Delta(\dot{v} + ur + x_G r^2) \\ N = I_{zz} \dot{r} + \Delta(\dot{v} + ru), \end{cases} \quad (1)$$

where u, v , and r are the forward, the drift, and the rotational speeds of the vehicle, respectively, x_G is the longitudinal coordinate of the center of gravity in the vehicle-fixed system, and Δ and I_{zz} are mass constants.

From the equilibrium equations, the global forces X and Y and the moment N acting on SWAMP in the vehicle-fixed system can be defined as the sum of external and internal contributes. The internal forces are the hydrodynamic response of the hull (denoted by X_H, Y_H , and N_H , respectively²), while the external ones are disturbances (denoted by X_D, Y_D and N_D and consisting in wind, current, and wave forces) and a sum of thrust-generated forces X_T, Y_T and moments N_T . In more detail, we can write

$$\begin{cases} X = X_H + X_T + X_D \\ Y = Y_H + Y_T + Y_D \\ N = N_H + N_T + N_D. \end{cases} \quad (2)$$

² At low speed, the structure of hydrodynamic forces can be described by a sum of products of velocity components (i.e., u, v , and r).

The thrust configuration of SWAMP is showcased in Fig. 2. We have a total amount of 4 thrusters, and denote by T_i^{PJ} and $\alpha_i = (0, 2\pi]$ the thrust intensity and the azimuth angle of the thruster i , respectively, with $i = 1, \dots, 4$. The i -th thruster creates a force component along the x -axis (denoted by X_i^{PJ}) and another one along the y -axis (denoted by Y_i^{PJ}), as well as a resulting moment (denoted by N_i^{PJ}) around the center of gravity. The thrust configuration on the vehicle-fixed coordinate system can be written as follows:

$$\begin{pmatrix} X_T \\ Y_T \\ N_T \end{pmatrix} = \begin{pmatrix} X_T^{PJ} \\ Y_T^{PJ} \\ N_T^{PJ} \end{pmatrix} = \begin{pmatrix} X_1^{PJ} + X_2^{PJ} + X_3^{PJ} + X_4^{PJ} \\ Y_1^{PJ} + Y_2^{PJ} + Y_3^{PJ} + Y_4^{PJ} \\ N_1^{PJ} + N_2^{PJ} + N_3^{PJ} + N_4^{PJ} \end{pmatrix}. \quad (3)$$

The contribution to the external force and torque given by each pumpjet motor is:

$$\tau_i^{PJ} = \begin{pmatrix} X_i^{PJ} \\ Y_i^{PJ} \\ N_i^{PJ} \end{pmatrix} = \begin{pmatrix} \cos \alpha_i \\ \sin \alpha_i \\ -b_i^{PJ} \cos \alpha_i + l_i^{PJ} \sin \alpha_i \end{pmatrix} T_i^{PJ}, \quad (4)$$

where l_i^{PJ} and b_i^{PJ} are the distances along the axes x_v and y_v of the thrusters' center from the origin, respectively.

As shown in Odetti et al. (2019b), each pumpjet module is composed by an azimuth motor and a main motor driving the pump impeller. The azimuth motor has an absolute encoder that allows to control the position with a high precision. Each $\alpha = 360$ degrees rotation corresponds to 495412 control steps. In the basic control architecture of SWAMP, the motor is driven by assigning an angle in degrees that is then converted into steps. The main motor drives the pumpjet thruster, and the amount of thrust is a function of the motor RPM. The percentage with respect to the maximum 1185 RPM ($RPM_{\%}$) is considered to assign a thrust intensity value. In more detail, T_i^{PJ} and $RPM_{\%}$ are related by a quadratic function, i.e., $T_i^{PJ} = 1.22 \times 10^{-3} \cdot (RPM_{\%})^2$.

A Logitech F310 joystick was adopted to drive the vehicle. The joystick levers were mapped as follows:

- the right lever was used to steer the pump-jet angle α , which was limited in the interval $[-\pi/2, \pi/2]$ linearly mapped on the $[-1, 1]$ output values of the lever;
- the left lever was used for powering pumpjet motors. The $RPM_{\%}$ was used to command the amount of thrust T_i^{PJ} by mapping $[0\%, 100\%]$ linearly into the $[0, 1]$ output values of the lever.

Both the angle and the thrust were applied only to the two bow thrusters, while the stern ones were turned off.

3. IMITATION LEARNING CONTROL SCHEME

In order to implement an imitation learning control scheme, each task demonstration from the human controller has been recorded at discrete time steps with sampling interval equal to Δt . In particular, the variables recorded at each time step k are $x(k), y(k), \psi(k), joy_{thrust}(k)$ and $joy_{angle}(k)$, i.e., the position and angle of the robot and the corresponding given joystick inputs (i.e., the variables $RPM_{\%}$ and α , respectively).

Let us denote the j -th trajectory performed by the demonstrator as $\chi_j = \{(s_j(1), a_j(1)), \dots, (s_j(K_j), a_j(K_j))\}$, where $s_j(k) = [x_j(k), y_j(k), \psi_j(k)]$ is the state at the discrete time stage k , $a_j(k) = [joy_{thrust}(k), joy_{angle}(k)]$ is the corresponding set of

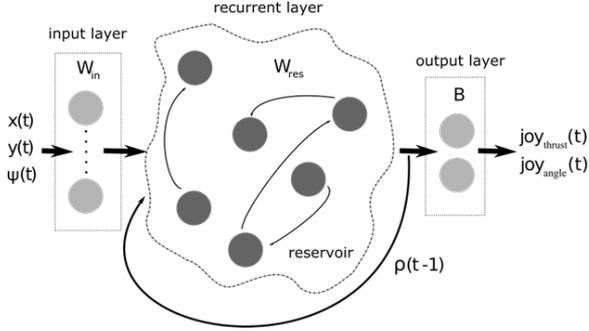


Fig. 3. The ESN for the generation of the joystick controls.

actions given to the robot and K_j is the total number of steps composing the trajectory.

As said, the machine learning model we considered is the ESN. This model takes as input the current observed state vector $s(k)$ and generates the corresponding action vector $a(k)$ as

$$a(k) = B \cdot \rho(k)^T \quad (5)$$

where \cdot denotes the dot product, B is a $(2 \times N)$ -dimensional matrix of output weights, and $\rho(k)$ is a N -dimensional vector representing the state of the reservoir. Its dynamics is ruled by the following relationship:

$$\rho(k) = h(c_{in} W_{in} \cdot s'(k)^T + c_{res} W_{res} \cdot \rho(k-1)^T), \quad (6)$$

where $s'(k) = [s(k), 1]$, W_{in} is the $(N \times 4)$ -dimensional matrix of the input weights, W_{res} is the $(N \times N)$ -dimensional matrix of the reservoir weights, c_{in} and c_{res} are real coefficients, and h is a sigmoidal activation function (e.g., the logistic function). Figure 3 illustrates the considered ESN and its reservoir network structure.

The main feature of the ESN is that both the inner weights W_{in} and the reservoir weights W_{res} are randomly extracted and kept fixed, while only the output weights B are trained. To this purpose, from a trajectory χ_j obtained by a human controller we can generate for each step k the corresponding $\rho_j(k)$ using (6) (initialized with $\rho_j(0)$ equal to the zero vector) eventually leading to a set of pattern/target pairs $\{(\rho_j(k), a_j(k))\}_{k=1, \dots, K_j}$. The collection of all the sets of such pairs for all the available trajectories corresponds to the training set over which we train the ESN in a classic pattern/target supervised fashion, i.e., by minimizing a mean square error between the output and the observed joystick controls. Given the linear structure of the output layer, this can be obtained very quickly by ridge regression, i.e., the optimal matrix B^* to use in (5) can obtained as follows:

$$B^* = (R^T \cdot R + \lambda I)^{-1} R^T A, \quad (7)$$

where R and A collect the pattern/target pairs from all the available human controller trajectories and λ is a regularization parameter.

4. EXPERIMENTAL RESULTS

In this section, we present the results of the experiment performed in Camogli during the *Festival della Comunicazione* 2019. To the best of the authors' knowledge, this is one of the first tests combining three main ingredients, i.e., marine robotics, citizenship engagement, and machine learning. In more detail, a piloting deck was devised on the *Rivo Giorgio* pier, and four buoys were positioned to mimic two gates in the stretch of sea in front of the pier. The SWAMP ASV was piloted

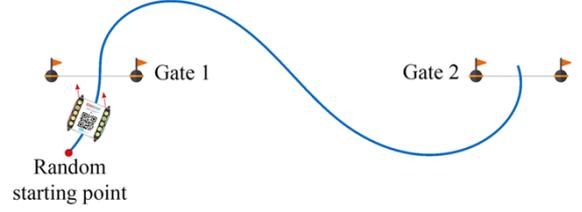


Fig. 4. The goal of the experiment performed in Camogli.

by citizens using a joystick, with the goal of performing a given trajectory: pass through the two gates in the same direction, thus describing an S-shaped path as shown in Fig. 4.

To test the ability of the imitation learning control scheme to successfully perform the desired task, a preliminary testing phase was carried out using a simulator of the dynamics of the vehicle, described by the equations reported in Section 2. In this case, the ESN controller was trained using trajectories collected when the authors piloted the vehicle in simulation. Such a simulation phase was also very useful to calibrate the controller parameters. In particular, we eventually set $N = 500$, $c_{in} = 1$, $c_{res} = 0.1$, and $\lambda = 10^{-3}$. After this off-line phase, the real experiment was carried out in Camogli. Not all the people were skilled enough to reach the goal, yet most of them succeeded in piloting SWAMP through the gates providing a wide range of “good” trajectories. In general, we noticed that most of people did not proportion the thrust amount, but rather they almost always used the full throttle. At the same time, only few people drove the vehicle exploiting all the possible angle positions of the joystick, inputting an almost constant 90 degrees thrust port or starboard command instead. The vehicle state s was measured by on-board global navigation satellite system and inertial measurement unit. In Fig. 5, a set of trajectories is reported on the left together with joystick inputs on the right. The above-mentioned driving styles can be easily identified in the figure. Notice that the 8-shaped path corresponds to the S-shaped goal task plus the path to get back to the gate starting line.

The trajectories used to train the AI-based controller were obtained by recording successful paths, starting a few seconds before driving through the starting gate, up to a few seconds after SWAMP passed through the goal gate. In more detail, a total amount of 16 trajectories, all recorded in one single morning, were employed to train the ESN controller. The 16 trajectories used for the training are reported in Fig. 6. In the same picture, three interesting examples of training trajectories are shown. It can be seen how such training trajectories are very different in quality. For instance, while in the first two the pilot went straight to the goal, the last one is quite chaotic. Yet, we decided to include all trajectories for the training to investigate robustness of the proposed approach.

After training the ESN parameters, the AI controller was tested by letting SWAMP perform the task autonomously while people were not driving. Fig. 7 reports three example trajectories performed by the AI controller. The first two trajectories show that the AI was able to achieve the goal of performing the S-shaped path and making SWAMP pass through the gates under standard initial conditions. The last trajectory is remarkable because the AI controller was activated while SWAMP was passing through the starting gate in the reverse direction with respect to the recorded training trajectories. Nevertheless, the AI drove the vehicle through the second gate successfully.

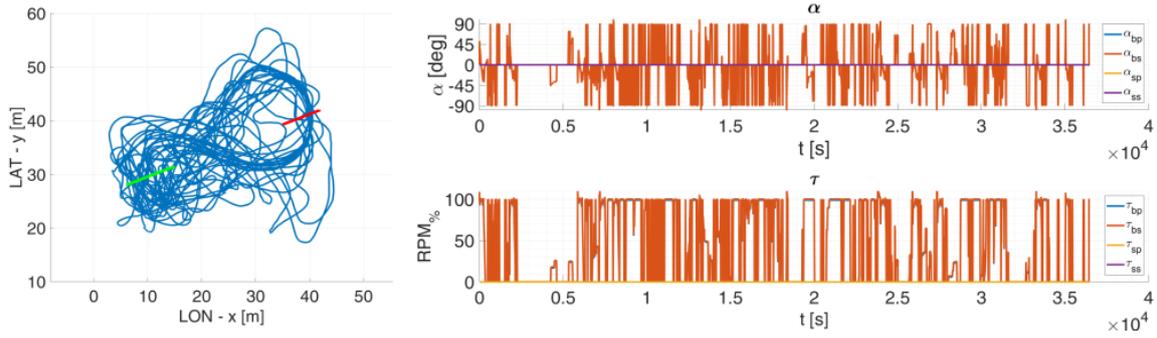


Fig. 5. Trajectories of SWAMP during the learning experiment and controls given by the human pilots.

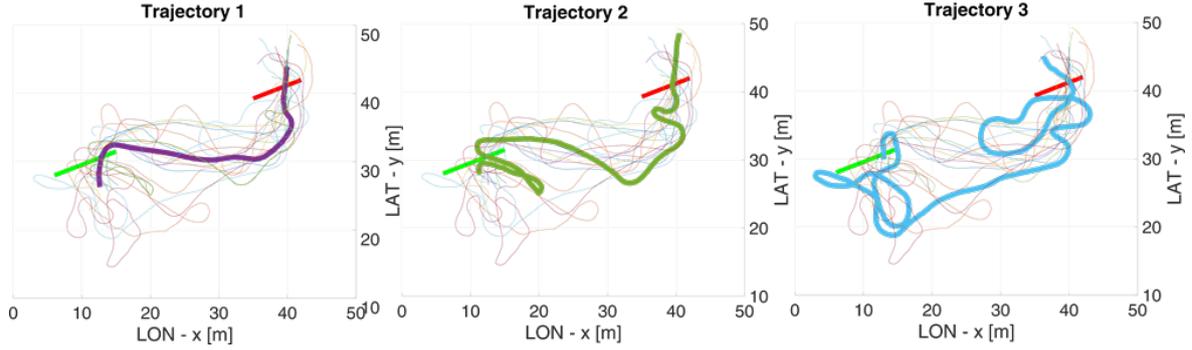


Fig. 6. Three examples of training trajectories. In the background, all the 16 trajectories used for the training are reported.

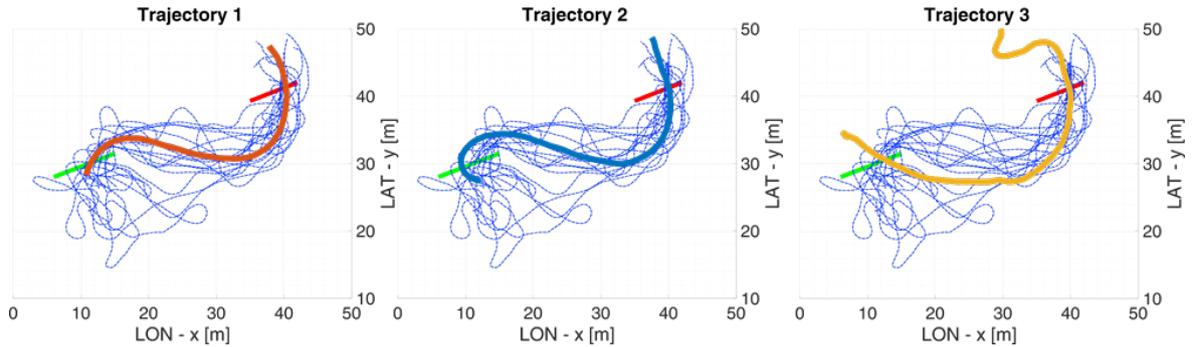


Fig. 7. Three trajectories performed by SWAMP with the AI piloting the vehicle. In the background, all the 16 trajectories used for the training are reported.

In general, the AI controller turned out to be very robust in the execution of the task. The AI correctly executed the required task all the times it was activated, also recovering from starting positions that were not seen in the training data. Sometimes, successful trajectories were accomplished under adverse sea conditions with waves that were not present during human demonstrations. On the one hand, this may be due to the relatively simple task that had to be performed. On the other hand, it may also be a hint that including “chaotic” training trajectories was a good choice, because it provided training data also in zones far from the “correct” path, which is very useful for recovery. Finally, this behavior also shows that the ESN structure, despite its simplicity, is able to generalize well on unseen data.

It is also interesting to notice that the control provided by the AI was in general smoother than the one made by humans. In particular, in Fig. 8 a comparison is shown. The human operator basically drove the robot through constant on-off type thrusts and set the angles at the maximum excursion. Instead, the AI

used less-nervous and jerky thrust than the human operator, and also smoother excursion angles. This suggests that the ESN model was able to avoid overfitting in learning the reference trajectories, which contributed to the good generalization capabilities that yielded a robust controller.

5. CONCLUSIONS

The preliminary experiment of citizenship engagement in robotics operations through imitation learning, carried out within the *Festival della Comunicazione 2019* in Camogli, has been a very interesting and successful experience. Many people have participated with enthusiasm and have shown interest in autonomous marine vehicles. The performed experiments allowed to obtain a first practical feedback of public reactions and participation in a scientific context involving autonomous robots and AI. In particular, the direct participation of citizens in robot training by imitation may serve as a practical mean to build public trust in robots performing autonomous operations

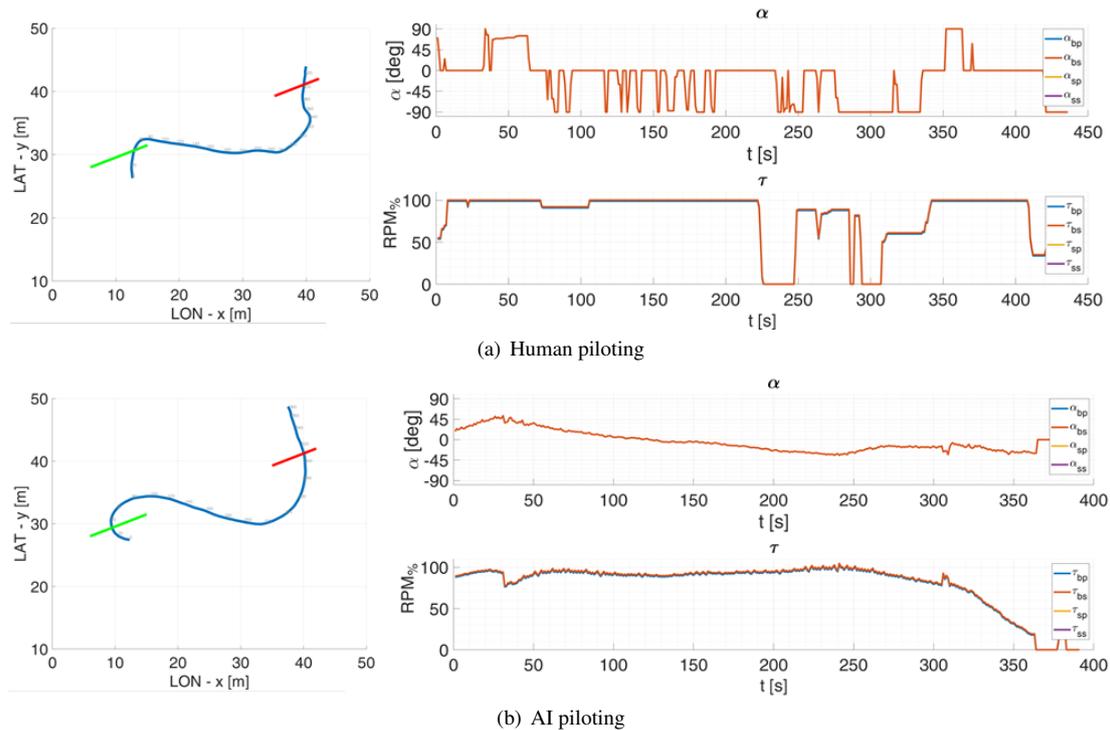


Fig. 8. A comparison of the controls given by human operator and by the AI-based controller.

in crowded areas, and a clear knowledge and understanding of the rules under which they operate.

From a technical viewpoint, the robust behavior of the AI controller operating the ASV in crowded areas opens various directions of research. To this purpose, future work will be aimed at evaluating effectiveness of the proposed approach using a larger amount of training and testing data, exploring different network architectures and more complex imitation learning paradigms, investigating the issue of selecting appropriate training trajectories, and exploiting the integration of simulated and real-world experiments.

REFERENCES

- Abbeel, P. and Ng, A. (2004). Apprenticeship learning via inverse reinforcement learning. In *Proc. of the 21st Int. Conf. on Machine Learning*.
- Amendola, J., Tannuri, E., Cozman, F., and Reali, A. (2018). Batch reinforcement learning of feasible trajectories in a ship maneuvering simulator. In *Anais do XV Encontro Nacional de Inteligência Artificial e Computacional*, 263–274.
- Antonelo, E. and Schrauwen, B. (2015). On learning navigation behaviors for small mobile robots with reservoir computing architectures. *IEEE Trans. on Neural Networks and Learning Systems*, 26(4), 763–780.
- Figueiredo Nascimento, S., Cuccillato, E., Schade, S., and Guimarães Pereira, A. (2016). Citizen engagement in science and policy-making. Technical report, Science for Policy report by the Joint Research Centre.
- Grigoryeva, L. and Ortega, J.P. (2018). Echo state networks are universal. *Neural Networks*, 108, 495–508.
- Hussein, A., Gaber, M., Elyan, E., and Jayne, C. (2017). Imitation learning: a survey of learning methods. *ACM Computing Surveys*, 50(2), 1–35.
- Odetti, A., Altosole, M., Caccia, M., Viviani, M., and Bruzzone, G. (2018). Wetlands monitoring: Hints for innovative autonomous surface vehicles design. *Proc. of 19th Int. Conf. on Ship and Maritime Research*, 1(1), 1014–1021.
- Odetti, A., Altosole, M., Caccia, M., Viviani, M., and Bruzzone, G. (2019a). A new concept of highly modular asv for extremely shallow water applications. In *Proc. 12th IFAC Conf. on CAMS, Robotics, Daejeon, Korea*.
- Odetti, A., Altosole, M., Bruzzone, G., Caccia, M., and Viviani, M. (2019b). Design and construction of a modular pump-jet thruster for autonomous surface vehicle operations in extremely shallow water. *J. Marine Sci. Eng.*, 7(7), 222.
- Pereira Figueiredo, J. and Pereira Abou Rejaili, R. (2018). Deep reinforcement learning algorithms for ship navigation in restricted waters. *Mecatrone*, 3(1).
- Perera, L.P. (2018). Autonomous Ship Navigation Under Deep Learning and the Challenges in COLREGs. In *Volume 11B: Honoring Symposium for Professor Carlos Guedes Soares on Marine Technology and Ocean Engineering*, Int. Conf. on Offshore Mechanics and Arctic Engineering.
- Scherer, J., Weber, S., Azofra, M., Ruete, A., Sweeney, E., Weiler, N., Sagias, I., Haardt, J., Cravetto, R., Spichtinger, D., and Ala-Mutka, K. (2018). Making the most of your H2020 project. Technical report, The European IPR Helpdesk.
- Tanaka, G., Yamane, T., Héroux, J., Nakane, R., Kanazawa, N., Takeda, S., Numata, H., Nakano, D., and Hirose, A. (2019). Recent advances in physical reservoir computing: A review. *Neural Networks*, 115, 100–123.
- Tuyen, L.P., Layek, M.A., Vien, N.A., and Chung, T. (2017). Deep reinforcement learning algorithms for steering an underactuated ship. In *IEEE Int. Conf. on Multisensor Fusion and Integration for Intelligent Systems*, 602–607.
- van der Zant, T., Bečanović, V., Ishii, K., Kobiálka, H.U., and Plöger, P. (2004). Finding good echo state networks to control an underwater robot using evolutionary computations. *IFAC Proc. Volumes*, 37(8), 215–220.