

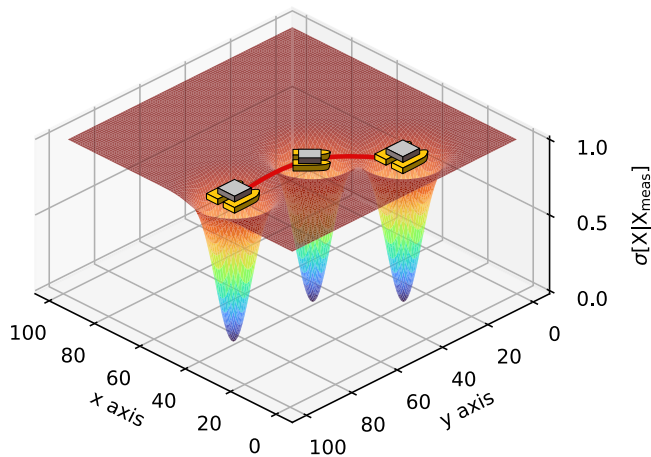
# Censored deep reinforcement patrolling with information criterion for monitoring large water resources using Autonomous Surface Vehicles



Samuel Yanes Luis\*, Daniel Gutiérrez-Reina, Sergio Toral Marín

Department of Electronics, University of Seville, Av. de Los Descubrimientos s/n, 41003, Sevilla, Spain

## GRAPHICAL ABSTRACT



## ARTICLE INFO

### Article history:

Received 4 April 2022

Received in revised form 20 November 2022

Accepted 23 November 2022

Available online 26 November 2022

MSC:

0000

1111

### Keywords:

Deep Reinforcement Learning

Autonomous Surface Vehicles

Information gathering

Environmental monitoring

Patrolling

## ABSTRACT

Monitoring and patrolling large water resources is a major challenge for nature conservation. The problem of acquiring data of an underlying environment that usually changes within time involves a proper formulation of the information. The use of Autonomous Surface Vehicles equipped with water quality sensor modules can serve as an early-warning system for contamination peak-detection, algae blooms monitoring, or oil-spill scenarios. In addition to information gathering, the vehicle must plan routes that are free of obstacles on non-convex static and dynamics maps. This work proposes a novel framework to obtain a collision-free policy using deterministic knowledge of the environment by means of a censoring operator and noisy networks that addresses the informative path planning with emphasis in temporal patrolling. Using information gain as a measure of the uncertainty reduction over data, it is proposed a Deep Q-Learning algorithm improved by a Q-Censoring mechanism for model-based obstacle avoidance. The obtained results demonstrate the effectiveness of the proposed algorithm for both cases in the Ypacaraí monitorization task. Simulations showed that the use of noisy-networks are a good choice for enhanced exploration, with 3 times less redundancy in the paths with respect to – greedy policy. Previous coverage strategies are also outperformed both in the accuracy of the obtained contamination model by a 13% on average and by a 37% in the detection of dangerous contamination peaks. Finally, the achieved results indicate the appropriateness of the proposed framework for monitoring scenarios with autonomous vehicles.

© 2022 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

## 1. Introduction

Water resources are fundamental to both human life and the economic development of communities. Large water bodies, such as rivers, reservoirs, and also lakes, are not only important for direct consumption but are the core of agriculture. Therefore, the conservation of these is vital for subsistence in many parts of the world. However, despite their importance, 80% of wastewater from cities and industries is discharged untreated into sea and rivers.<sup>1</sup> This situation has dramatic consequences for ecosystems.

A very frequent problem in situations of uncontrolled discharges of nutrients is the accelerated eutrophication of waters and the appearance of blue–green bacteria colonies. Whenever the water is nourished with ammonium and nitrites from agricultural fertilizers or human and animal fecal remains, large colonies of cyanobacteria emerge and populate the surface with a greenish mantle. The thick blooms prevent the sun from hitting the underwater flora, which is responsible for oxygenating the waters. As a result, the fish die of anoxia. Thus, the waters become muddy, with pestilential odors and a very poor quality for bathing, fishing, or human consumption. This is the situation in Lake Ypacaraí, Paraguay's largest water resource. The lake, once a tourist enclave, is now populated by colonies of blue–green algae with all its consequences. The situation is similar in places like Mar Menor (Spain) [1] or Lake Erie (USA) [2]. All these places have in common that they are very large extensions of water, and that it is difficult to obtain their updated contamination status, in terms of the relevant physico-chemical and water quality (WQ) variables (pH, dissolved oxygen, nitrites or ammonia concentration, etc.). Any solution to this contamination problem requires a comprehensive monitoring process. Monitoring is defined here as the sufficient and homogeneous collection of biological WQ data with which to form a physico-chemical model to serve as an early warning system for pollution peaks. However, this task is difficult due to the large extension size of the ecosystems under monitoring. Consequently, manual missions to measure the WQ are tedious and costly.

Moreover, installing a fixed sensor grid is a suboptimal solution as sampling locations cannot be varied and multiple battery replacements could be even more expensive than human-conducted missions. The use of Autonomous Surface Vehicles (ASVs) equipped with WQ sensor modules has been gaining momentum lately for some similar applications (see Fig. 1) [3]. These low-cost electric vehicles allow continuous sampling with routes that can be adapted according to various optimality criteria [4].

However, these robotic systems require intelligent perception and decision modules to perform monitoring tasks efficiently. When using these robots, aspects beyond information gathering have to be considered, such as physical battery limitations and nonnavigable terrain constraints. This transforms the monitoring problem into an Informative Path Planning (IPP) problem, which combines the challenge of obtaining the most informative path, and the compliance of ground restrictions and obstacle avoidance. This problem has been previously addressed for a wide range of applications: agricultural characterization [5], for the generation of water quality models [6], the search for gas leaks [7], or the location of contamination sources in radiological environments [8]. Thus, the problem under study has been treated from multiple perspectives in past works. Our approach intends to fill some research gaps present in these previous works or current literature, such as the dimensional complexity when monitoring big water resources, the formulation of the information gathering

in terms of the entropic uncertainty reduction or the need for an enhanced learning process using deterministic information of the non-convex scenario. Therefore, we present an algorithm that can be used without loss of generalization to all these aforementioned particular objectives, as the algorithm is model-free and the information formulation can be applied to the reality of these parameters under surveillance apart from the WQ measurement context.

We also present a novel approach to monitor dynamic scenarios, which is a major gap that previous works have not addressed after all. We present the problem called Informative Patrolling Path Planning (IPPP), which is an extension of the former IPP problem that addresses possible dynamics of the unknown environment. This situation gains importance when we want to measure eventual risks such as spills, new blooms or the time trend of certain pollution variables. This constitutes the base problem of detection and early warning of contamination hazards on long term missions.

To deal with the aforementioned monitoring scenarios, a stochastic framework is defined to model the information collected at each physical point by the ASV. Assuming that each possible point eligible to be sampled behaves as a Gaussian random variable  $p \in X \sim \mathcal{N}(\mu, \sigma)$ , a spatial correlation matrix  $\Sigma[X, X_{meas}]$  can be defined to model the statistic relationship between samples. This matrix indicates the level of uncertainty of each point, considering the locations where  $X_{meas}$  has been sampled. This approach makes use of a Radial Basis Function (RBF) to spatially correlate samples as a function of their adjacency, under the acceptable assumption that physically close samples will be more closely related. Given that, the IPP will consist of sequentially deciding the next physical point at which to take a sample, following the information gain maximization criterion. Following the Information Theory definition of information gain,  $\Sigma[X, X_{meas}]$  defines how informative a point is according to the entropy decrease of the model formed by the set of points of the lake  $X$  and the sample points  $X_{meas}$  [9]. Finally, the IPP problem to be solved here consists of minimizing the total entropy of the scenario  $H[X|X_{meas}]$ .

The dimension of this sequential problem explodes when the routes are large and movement restrictions are included. In the case of Lake Ypacaraí, large distances must be covered for the routes to decrease optimally the entropy, so the possibilities of movement are almost infinite. In addition, it is necessary for the vehicle to have a reactive policy, capable of adapting to different boundary conditions, i.e., arbitrary route starting points and temporal dependencies. This work proposes the use of Deep Reinforcement Learning (DRL) techniques to find an optimal policy, based on convolutional neural networks (CNNs). This convolutional network is able to interpret a graphical state of the problem and choose the next sample point according to a tailored reward function. This reward function should be designed to evaluate every action in terms of the information gathered. DRL techniques are useful because they do not require a prior model of the system and are able to adapt to a scenario with arbitrary boundary features through model-free interpretation of the state and reward function. These characteristics make DRL a suitable approach to solve the IPP and IPPP. Specifically, Deep Q-Learning (DQL) are applied here because of their ability to estimate the future reward given a state for each possible action in the possible set. To enhance the capability of the classical DQL algorithm, this paper further elaborates on the application of some useful techniques for this IPP problem, that have been proven effective in previous works of the DRL literature [10], but none in a path planning application: the use of a Prioritized Experienced Replay [11], a Dueling architecture in the Q neural network [12] and, finally, the use of noisy neural networks [13] to intrinsically motivate the exploration of the state–action domain. These publications revealed

\* Corresponding author.

E-mail address: syanes@us.es (S. Yanes Luis).

<sup>1</sup> <https://www.un.org/en/global-issues/water>



**Fig. 1.** ASV prototype for the monitorization of lake Ypacaraí (Paraguay). The ASV is equipped with water quality sensor modules, GPS and communications for remote planning.

the high learning capacity of algorithms such as DQL. However, in the application of them to complex problems, there are hard constraints that are known a priori or can be computed in a deterministic way. A difference between the proposed approach and the previous ones is in using the deterministic knowledge of the navigation map to increase the sampling efficiency. This is effectively a slight departure from the initial model-free paradigm of DQL. However, we propose the use of a censoring mechanism that increases the sample efficiency because it provides accurate information on illegal actions that would otherwise have to be inferred by trial and error. This algorithm, the Censoring-DQL, takes the deterministic information of the environment related to the obstacle avoidance task, and provides a way to neglect those actions that incurs in violations of the hard-constraints in navigation. As far as we understand, there is no apparent compromise on learning convergence when this mechanism is used and, moreover, it results in improved training efficiency.

The novelty of this work also lies in the formulation of the WQ monitoring problem as a sequential uncertainty minimization process, using the principles of DRL: The problem is defined as a discrete Markov decision process (MDP) in which the objective is to maximize information acquisition with a discounted finite horizon. Thus, it is proposed to use the information gain as the reward that reduces the total uncertainty under the assumption that the water quality parameters behave smoothly and as Gaussian random variables. As explained in [14], it is usually challenging to design the reward function to solve a real problem with as it must reflect the desires of the resulting behavior. In the particular IPP/IPPP case, this tailored reward function serves as a surrogate signal for a general objective: information acquisition, translated into model error reduction, maximal search, etc. This reward-engineering process contributes, to improve data collection systems and to obtain deep policies for decision making in the face of information uncertainty. With an appropriate definition of these learning parameters, the DRL is able to outperform previously proposed algorithms and heuristics, with more efficient and robust informative paths. In previous works, like in [15], the patrolling follows a criterion based on covered discrete positions and assuming no smooth useful property of the inner objective function under estimation. In addition, other similar algorithms, such as [4], use paths from side to side of the edges, without considering interior paths that do not intersect. Another disadvantage of these heuristics is the inability to act online, since

they base the optimization on a static criterion and, therefore, a single solution rather than a policy is obtained.

In summary, the contributions of this work are:

- The formulation of the Patrolling Problem in terms of the entropy information for water monitoring scenarios with a constrained search space.
- A model-free Deep Reinforcement Learning framework to solve the IPP and IPPP.
- The definition of a tailored reward function and a graphical state formulation policy optimization that is effective in both informative missions.
- The implementation of a Noisy + Censoring mechanism that allows for a more efficient learning of the deep policy using deterministic information from the environment map.

This article is organized as follows: in Section 2, there is a brief survey of previous works and other techniques that address similar problems. In Section 3, Materials and Method, the IPP and IPPP problem is formally stated, the algorithm is presented, and the Deep Reinforcement Learning framework described. In Section 4, the results of the training are analyzed and discussed. Finally, Section 5 closes with the conclusions of this work and future lines of work.

## 2. Previous works

The utilization of ASVs is an increasingly common option for autonomous operation in hydrological resources and harbors [6,16–21]. Applications range from bathymetry surveys [16], pollutant modeling and monitoring [17], or patrolling for early warning. In most applications, Artificial Intelligence or meta-heuristic optimization techniques are usually applied to solve complex NP-hard problems: Bayesian optimization [18], genetic algorithms (GA) [19], Swarm Intelligence [20], or DRL [21].

One of the most common approaches to characterize the environment is to use a surrogate model, such as Gaussian Processes (GPs) like in [6,18]. These approaches have been proposed for the monitorization of water quality variables in the Ypacaraí Lake itself, by means of a policy based on an Acquisition Functions (AF) using Bayesian Optimization classic theory. In [6], it is proposed a modification of the classic Expected Improvement (EI) acquisition function that can balance the exploration/exploitation of the uncertainty. This heuristic approach has been proven effective for

modeling WQ scalar fields. This work is cornerstone in informative path planning with maritime vehicles, nonetheless neglects any dynamics in the WQ profile and is more oriented to have a general model with low quantity of data. In the same line, in [18], authors proposed a multivariate extension for different objective functions to fit in a multivariate fashion. The same heuristics are implemented using the EI function as a policy for deciding the next point to sample. While efficient for low data regimes, authors do not consider that WQ parameters can change within the time of a mission. It does also bring up the question whether the use of the GP mean, used by the EI acquisition policy, really affects in the error metric or in the sampling redundancy. Our proposal tries to define an adaptive policy that is able to monitor efficiently the same scenarios using a reward function motivated only by the uncertainty perspective. This method is easily extended to the temporal case, which was not addressed by the aforementioned works, and serves the purpose of patrolling dynamic scenarios. In addition, our proposal's application considers a more intensive coverage with a higher number of samples and is not subjected to a particular regression method, like the EI is with the Gaussian Processes.

Another interesting example of IPP application can be found in [5], where it is implemented a method to gather information of a static scalar field using unmanned aerial vehicles with camera sensors. This method implements a Kalman Filter for data fusion while trying to minimize the information entropy of the environment. The path is optimized using a Covariance Matrix Adaptation Evolutionary Strategy (CMA-ES) with a receding optimization horizon. The algorithm decides every timestep a path from a limited number of global positions (up to 30 possible waypoints as decision variables), resulting in an online strategy that is difficult to scale for bigger resolutions and longer paths. This approach also proposes a definition of the uncertainty with zones that has no-importance, therefore do not contribute to the uncertainty reduction criterion. In terms of the informative path planning for WQ parameters, this cannot be applied, because we consider every zone is subjected to the same importance because of biological requirements for early-warning systems like this. Our work, by means of the Deep Reinforcement Learning algorithms and the capacity of such algorithms to work with high dimensions, is able to work with much higher state-action complexity and could be extended easily to a multiagent case. The difference with our proposal also resides in the fact that the Q-function proposed here codifies the reward function through the learning. Once an effective policy is learned, there is no need for long computation in the real mission beyond the pure inference of the next action. This is an advantage in terms of a future real-time application of the algorithm.

A similar design to [5] is proposed in [8]. This particular approach focus on modeling again static scalar fields of radiation levels in hazard scenarios. This work introduces a novelty respect to [5], which is the possibility of static obstacles that make the field non-convex respect to the vehicle's movements. We consider this aspect is important because in most IPP missions, there are obstacles that cannot be neglected and that affects to the complexity of the task. The performance of the proposed algorithm, nonetheless, is surpassed by other, simple but effective, heuristics like the lawn-mower algorithm or the random search. The results of balancing the exploration and exploitation of the scenario knowledge seems to be inconclusive when non-convex obstacles are present. This approach does not consider any kind of temporal dependency for the radiation scalar fields while using a Gaussian Process to conform the surrogate model for regression. No other regression methods were tested. Previous works like [22] have studied the use of the entropy for exploration in RL tasks. This work proposes an entropy measurement (state

entropy) to represent the uncertainty of the agent to select an action in a state. This way, the entropy is used for exploration of the action-state domain as an effective way to enhance the learning. While only validating the entropy effectiveness in a low complex problems, this work addresses the utility of the entropy and serves as a theoretical foundation for our work.

Reinforcement Learning, and more specifically, Deep Reinforcement Learning has been commonly used for exploration and information gathering lately [7,15,23]. One important example of use can be found in [15]. This work proposed the use of an on-policy algorithm for the error minimization in modeling of scalar fields using multiple vehicles. The work uses a visual representation of the state with recurrent networks, similar to [23] or like our work, with a Convolutional Neural Network to represent the policy. This work also proposed to use the Root Mean Squared Error (RMSE) directly as a reward function for the agents to learn how to maximize the accuracy of the estimation. This aforementioned work demonstrates that it is possible to train multiagent policies with DRL for this purpose. Nonetheless, directly using the regression error respect to the real scalar field under estimation undoubtedly implies knowing the exact ground truth at the time of learning, which is not possible in many cases nor in ours. To overcome this design problem, we propose modeling the uncertainty in a smooth way using a Radial Basis Function. This function not only allows a soft definition of the information (in [15] the information is defined to be binary – known or not known), but to have an estimation of the redundancy of the sampling. This is possible assuming that, in most of the estimation cases, it is not necessary to visit a place to estimate efficiently the real value if there is sufficient information of the surroundings. Our work also differs in the fact that we propose a method to avoid collisions using obstacle domain knowledge in a deterministic fashion. This censoring mechanism provides an enhanced learning, as it can be seen in the Results Section. Other similar approach using DRL for informative path planning is proposed in [7] for gas leaks localization. This approach used a differential model of the gases to provide a sufficiently accurate model for the deep agent to learn. This approach is quite effective when there is enough knowledge of the underlying dynamics to compose a model. However, this is not suitable for the problem we present as we do not have a mathematical model of the WQ parameters, only enough information to choose the kernel parameters based on the bandwidth of the benchmark function. This way, our approach works in a full model-free fashion respect to the information and do not rely on equations of the dynamic in the static nor in the non-stationary case. In [24,25], there are two similar examples of using the uncertainty as a robust measurement of the information. The uncertainty serves as a surrogate metric for the accuracy in localization tasks. The network that is used for the policy representation, is designed as a Siamese actor-critic network, in contrast to our work, where the Q-function is represented by a single CNN head.

Other works put the attention solely on coverage like in [4] which is a good example of the application of evolutionary algorithms for water resource coverage. This approach is based on solving a Travel Salesman Problem (TSP) using Eulerian cycles to maximize the unvisited area using a single vehicle with GA. This approach only uses edge-to-edge movements and is limited by the formulation of visitable nodes. Our algorithm proposes to overcome this method by allowing movements in different directions at each instant. Moreover, the information coverage is no longer based on an all-or-nothing basis, but follows the information criterion imposed by the RBF function, like it was aforementioned. Another example of full coverage algorithms is in [26], where a DRL algorithm for a transformable mobile robot with discrete actions is applied. In this work, the algorithm

called Actor–Critic with Experience Replay (ACER) and a visual formulation of the state with CNNs are used in the same way as in our proposal. However, the application again consists of a binary coverage (covered/uncovered regions) with no temporal criterion or useful redundancy factors in the coverage.

Regarding the particular use of DRL for path planning with non-convex scenarios, most of the works focus on simply finding obstacle-free routes [27] or low-level control [28]. In [27], the Twin Delayed Deep Deterministic policy gradient (T3D) algorithm is implemented for obtaining optimal obstacle avoidance policies for aerial vehicles with distance sensors. In [28], the DRL is applied in obtaining a low-level tracking policy for an Autonomous Underwater Vehicle (AUV) using visual inputs from a camera. By the means of the TinyYOLO architecture, an convolutional obstacle detection system, the AUV is able to track the desired paths without collisions. Another interesting application of DRL in vehicles can be found in [29], where a DRL policy is optimized to solve the vehicle routing and navigation task in an interactive and self-adapting manner. The difference in the sense of obstacle-avoidance between such approaches and ours is that the proposed censoring mechanism implies a simplification. The navigation map is known a priori, and the obstacles are observable. Then, the decisions related to the information acquisition are surrogated to the detection mechanism and therefore.

With regard to the patrolling and surveillance tasks, where there is a temporal criterion involve, there are previous examples [30,31] that use DRL with a visual formulation of the state. In [30], the use of DQL is proposed for the realization of partial coverage routes based on a battery budget with vehicle landing constraints. Our work differs in the sense of the resolution and the state–action domain, which is significantly bigger in our case. The work [30] is more similar to a low resolution task like [17]. In [31], an interesting example of nonhomogeneous temporal patrolling can be found, where emphasis is put on going over the most important areas with longer waiting time, in a similar way, but with a discrete redundancy criterion. The proposed work differs from [31] in the presence of nonnavigable zones. There is also a difference in the acquisition system. Whereas in [31] a camera is in charge of sensing, the proposed approach uses point samples and generates a model in between.

In summary, it can be observed that most of the previous works have worked in one or another of the main research lines in this work: informative path planning, non-convex planning, entropy reduction, etc. We have detected a research gap in those previous works in the sense that most applications of DRL in vehicles are limited to low-level control or path planning, with results similar to classical algorithms. This work proposes a way to optimize more complex problems, with a much higher resolution from previous approaches with entropy minimization as the optimality criterion for environmental monitoring. We propose a visual formulation of the process, which is not a novelty, but we go a step further from completely model-free DRL algorithms using deterministic knowledge of the environment for a more efficient learning. Another important improvement respect to some previous approaches is in the formulation of the information in terms of the predictive uncertainty. This formulation is able to take in account the redundancy of the information and a simple extension to the temporal patrolling. Finally, the novelty of this work relies in taking the ability of DRL to work with high dimensional decision problems (with almost infinite possible solutions) using a proper formulation of the information, which have only been treated using classical heuristics or computational expensive optimizers so far.

### 3. Materials and methods

This section introduces the mathematical framework of the monitoring problem and the assumptions taken to model the scenario. Both monitoring cases are considered: a static scenario, where the gathered information does not recover its importance once obtained, and a dynamic scenario, where the unvisited zones increase its uncertainty within time.

#### 3.1. Entropy framework

To present the IPP and IPPP problems, we begin with the definition a navigable space  $X \in \mathbf{R}^2$ , where the vehicle can take WQ samples. We also define the subset set of sampled locations  $X_{meas} \in \mathbf{R}^2$ , and a binary matrix  $M$  which is a navigability map such that  $X = \forall p := [p_x, p_y] \mid ; M(p) = 1$ . Now, we define that each visitable point  $p$  behaves as a Gaussian random variable of mean  $\mu$  and standard deviation  $\sigma$  such that  $p \approx \mathcal{N}(\mu, \sigma)$ . The space  $X$  is thus formed as a Multivariate Gaussian Distribution (MGD), where  $X \approx \mathcal{N}_n(\mu, \Sigma)$ , with  $\Sigma$  being the correlation matrix of  $X$ .

The measure of the correlation between points can be obtained by means of a function that indicates how closely related two variables are in the search space. For this application, a RBF kernel has been chosen, so that the correlation between two variables decreases smoothly and exponentially with the distance between them according to the Eq. (1). This choice makes sense from the point of view of water resources monitoring since physico-chemical variables usually have a smooth distribution.<sup>2</sup> Thus, two samples taken in close proximity will presumably be highly correlated and the information they both provide can be redundant. The parameter  $l$  will serve to scale how much correlated two measurements ( $p, p'$ ) are with each other and is usually chosen based on prior knowledge of the environment to be monitored [5,15] or intensity of the environment coverage.

$$RBF(p, p') = \exp\left(-\frac{(p_x - p'_x)^2 + (p_y - p'_y)^2}{2l^2}\right) \quad (1)$$

Thus, we can find the correlation matrix of the navigable space conditional on the sampled locations by the element-wise evaluation of Eq. (1). The conditional correlation matrix  $\Sigma[X|X_{meas}]$  will have the expression of the Eq. (2) [9].

$$\Sigma[X|X_{meas}] = \Sigma[X, X] - \Sigma[X, X_{meas}] \times \Sigma[X_{meas}, X_{meas}]^{-1} \times \Sigma[X_{meas}, X]^{-1} \quad (2)$$

Now, the monitoring objective involves decreasing the entropy associated with the conditional correlation. The information entropy  $H[X|X_{meas}]$  gives a measure of the uncertainty about the monitoring domain and the randomness of a sample at an arbitrary point in that space. The lower the entropy, the more confidence one has about the scenario. In this sense, under the assumption that the monitoring space  $X$  is an MGD, the entropy gives a lower bound on the error of an arbitrary estimator  $\hat{\mu}$  [32], as given in the Eq. (3).

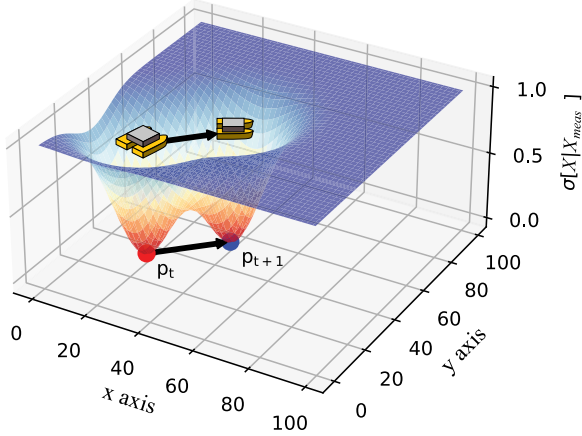
$$E[(\mu - \hat{\mu})^2] \geq \frac{1}{2\pi e} e^{2H(\mu|X_{meas})} \quad (3)$$

Finally, the entropy can be calculated as [9]:

$$H[X|X_{meas}] = \frac{1}{2} \log(|\Sigma[X|X_{meas}]|) + \frac{\dim(X)}{2} \log(2\pi e) \quad (4)$$

An example of how the min–max normalized standard deviation values  $\sigma$  of  $\Sigma[X|X_{meas}]$  (its diagonal) decrease when starting from a sample at  $p_t$  to  $p_{t+1}$  can be seen in Fig. 2. In the vicinity of those sample points,  $\sigma$  is not 1.0 at all since a smooth correlation is assumed with the RBF function.

<sup>2</sup> <https://marmenor.upct.es/maps/>



**Fig. 2.** Uncertainty conditioning process. When we go from  $p_t$  to  $p_{t+1}$  and incorporate a new sample to  $X_{meas}$  with the Eq. (2), we can obtain a new  $\Sigma[X|X_{meas}]$  matrix. It is observed that the standard deviation associated with  $\Sigma[X|X_{meas}]$  becomes zero at the sample points (assuming no sampling noise) and in their environments according to the RBF function. The entropy  $H[X|X_{meas}]$  of the process decreases accordingly.

### 3.2. Temporal-dependent entropy

In a more general case, where the variables to be monitored change during a mission, it is necessary to readjust the entropy formulation to include such factor. To do this, a monotonically increasing time-dependent function  $f(t)$  is included in the correlation matrix  $\Sigma[X|X_{meas}]$  calculus, such that the uncertainty of the measurements increases as a function of the time that has passed since each location  $X_{meas}$  was visited. Thus, the definition of  $X_{meas}$  is expanded to include the instant of visit of each sample  $t_{meas}$ . Similar to [33], the correlation matrix of samples  $\Sigma[X_{meas}, X_{meas}]$  is calculated by adding the temporal term to the diagonal:

$$\Sigma[X_{meas}, X_{meas}] = RBF(X_{meas}, X_{meas}) + \tau(t - t_{meas})^2 \mathbf{I} \quad (5)$$

This is analogous to increasing the noise of each particular sample quadratically with time. Old samples do not contribute to decreasing entropy, and over time, unvisited areas recover their original entropy. In the Eq. (5), the value  $\tau$  is a forgetting factor that modulates how fast the correlation in each zone regenerates. Thus, with this formulation, decreasing entropy involves not only visiting areas never visited before, but cyclically patrolling the areas to prevent entropy from increasing.

### 3.3. Deep Q-Learning

In this section it is presented the foundations of the selected DRL algorithm, Double Deep Q-Learning [34]. The different modules of the reinforcement learning paradigm are presented, such as the state representation, reward function, Deep Q-function, and the variations from the original approach in [34] that enhance the learning in this scenario.

Deep reinforcement learning (DRL) is a methodology for solving sequential problems such as the proposed one [17,21]. In DRL, a Markov Decision Problem (MDP) is described so that an agent (vehicle), which is in a state  $s_t$ , takes an action  $a_t$  according to a policy  $a_t = \pi(s_t)$ . The environment processes the action and returns the next state  $s_{t+1}$  and a reward value  $r$ . To solve a MDP one has to find the optimal policy  $\pi^*(s)$  that maximizes the discounted reward in a control horizon  $T$ :

$$\pi^*(s_t) = \max_{\pi(s_t)} \sum_{t=0}^T [R(s_t, a_t = \pi(s_t))] \quad (6)$$

Deep Q-Learning (DQL) algorithms within DRL are based on finding an action-value function  $Q(s, a; \theta)$ , represented by a Deep Neural Network (DNN) with  $\theta$  being the trainable parameter, which estimates the discounted future reward, given a state  $s$  and for each possible value of actions  $a_t \in A$  according to the Eq. (7). The value of  $\gamma \in [0, 1)$  discounts the value of the reward function  $r(s_t, a_t)$  over a finite decision horizon.

$$Q(s_t, a_t; \theta) = \mathbb{E} \left[ r(s_t, a_t) + \gamma \max_{a'} Q(s_{t+1}, a'; \theta) \right] \quad (7)$$

The Q function is estimated by collecting experiences  $E \sim \langle s_t, a_t, s_{t+1}, r_t \rangle$ , storing them in a experience-replay buffer, and training the DNN in batches of size  $|B|$ . The training procedure is based on updating the network parameters  $\theta$  in the direction of the descending gradient of the loss function. This loss function is computed as the Time Difference (TD) error between the chosen action of  $E$  and the best predicted values of a discounted target function  $Q^*(s', a, \theta')$  (see Eq. (8)). Thus, the parameters  $\theta$  of the network are updated by taking a stochastic gradient descent step with a learning rate of  $l_R$ .

$$\mathcal{L} = \left[ \underbrace{Q(s_t, a_t; \theta) - \left( r_t + \gamma \max_{a'} Q^*(s_{t+1}, a'; \theta') \right)}_{TD} \right]^2 \quad (8)$$

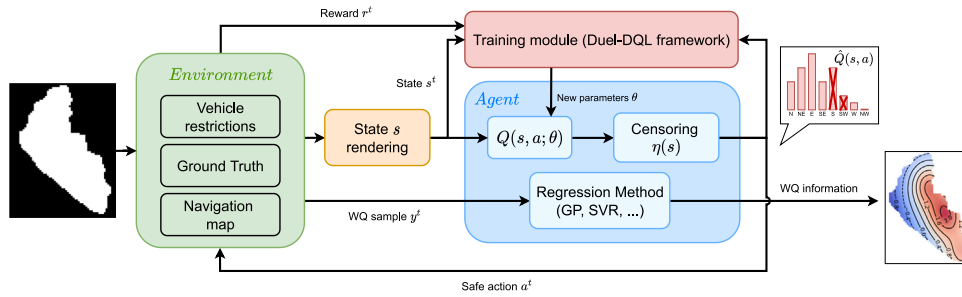
Through trial and error, the Q function is adjusted to the real value and the agent learns to select the best actions. The target function is periodically updated with the value of the Q function itself [34] to improve the stability of the training. To have a bias-free estimate of the Q function, it is necessary to collect both good and bad experiences. The balancing of the exploration of the action-state domain and the exploitation of what is learned is usually left to the DQL by a  $\epsilon$ -greedy policy. This policy balances exploration and exploitation by taking a random action to explore with  $\epsilon$ -probability. The value of  $\epsilon$  starts high at the beginning of training and decays as learning is completed to take advantage of the knowledge gained in the early stages.

### 3.4. Proposed framework

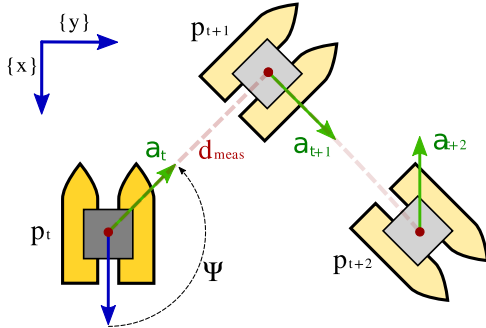
The proposed problem is based on sequentially choosing  $p_{t+1}$  that provides the most information from the perspective of entropy reduction. The formulation of the actions, the state, the deep policy architecture, and the reward function are described below. For a complete understanding of the system, see Fig. 1, which visually describes the framework based on the DRL Censoring method, its inputs and outputs (see Fig. 3).

#### 3.4.1. Actions

A discrete formulation of the action space is chosen to reduce the number of possible action-state combinations. Discretizing the action space in this way has proven to be sufficient in previous works [5] and allows for better convergence of policies [35]. Thus, the ASV agent can choose from up to 8 possible actions that will result in a movement in 8 different directions with respect to a fixed reference system parallel to the axes of the navigation map  $M$ . The possible angles of direction  $\Psi$  follow the bearings of a compass  $A := [S, SE, E, NE, N, NW, W, SW]$ . Note that the order of these angles does not intervene at all in the algorithm. An intuitive order has been chosen but, for the sake of the algorithm, an arbitrary order could have been chosen. Therefore, each action  $a$  leads the ASV to a movement in that direction until it reaches a point at a distance of  $d_{meas}$  with respect to the previous point. In Fig. 4 it is depicted the process of movement within its reference frame. Regarding the dangerous movements, any action that



**Fig. 3.** This diagram describes the general operation of the system. The algorithm is given a simulation environment based on plausible WQ data and a navigation map. The agent uses a deep policy  $Q(s, a)$  to decide the direction of movement in which to take the next sample. The censoring mechanism prevents invalid (hard-constrained) actions. Simultaneously, a WQ model is being shaped by a regression method (Gaussian Processes, Support Vector Regressor, etc.).



**Fig. 4.** The ASV is moves in 8 possible directions  $\Psi$ . Those directions are fixed respect to an inertial reference frame parallel to the navigation map  $M$ . Every action implies a travel in the selected direction for a distance  $d_{meas}$ . Therefore, every point  $p_t$  is separated from the previous one  $p_{t-1}$  by a distance of  $d_{meas}$ .

generates a waypoint outside the navigable areas is neglected and will not cause any movement. For any effective action, a sample of the WQ values is taken and the new position is updated in the set of visited zones such that  $X_{meas}^{t+1} \leftarrow X_{meas}^t \cup p_{t+1}$ . With this, the covariance matrix is recomputed according to Eq. (2) and the next action can be queried to the policy  $\pi(s)$ . The number of possible actions is determined by the length of the path  $D$ , which will depend on the problem we are solving and the ASV battery budget.

### 3.4.2. State

The state  $s_t$  represents the observable information of the scenario at a certain instant  $t$ . To fulfill the Markovian hypothesis, the future reward should depend exclusively on the current state and the action taken. Consequently, the state of the problem must be chosen carefully, including all necessary a priori information for entropy minimization. As the IPP is a problem with physical constraints dependent on the current position of the navigability map, in addition to the information collected, an image-like state with three channels is proposed: (i) the binary navigability map, (ii) a map of the path followed so far, and (iii) an uncertainty map in which each cell has a value equal to the value of  $\sigma$  according to the conditional correlation matrix  $\Sigma[X|X_{meas}]$  (see Fig. 5). This last state will depend on the kernel hyperparameters as it is mentioned in Section 3.1. and will change according to Eq. (2). This way, the state  $s = [s_1, s_2, s_3]$  is discretized into three images of  $75 \times 60$  pixels. These images represent all the information available for the scenario, which makes this problem a Fully Observable MDP (FOMDP) To avoid large input values to the Deep Q-network, the tree images are minmax-normalized. This way, every pixel value of the state is within  $[0, 1]$ .

It is worth mentioning that, to avoid violating the Markovian assumption, the path traveled by the ASV up to the current

instant is included as part of the image (see Fig. 5b). In order to distinguish the old positions from the current one, the oldest of the positions will have a value of 0 and the current one will have a value of 1.

### 3.4.3. Noisy neural architecture

To represent the function  $Q(s, a)$ , we propose the use of a CNN to process the state as an image. A Dueling architecture has been designed. With this architecture, we seek to compute two reward estimators: on the one hand, the value function  $V(s)$ , which represents the estimated future reward in the state  $s$ . On the other hand, the advantage function  $A(s, a)$ , which represents the estimated future reward for each action  $a$  with respect to the value of  $V(s)$ :

$$A(s, a) = Q(s, a) - V(s)$$

Thus, the calculation of  $Q(s, a)$  is performed such that:

$$Q(s, a) = V(s, a) + A(s, a) - \frac{1}{|A|} \sum A(s, a) \quad (9)$$

This is a common technique used in DQL approaches like [23] that allows to better generalize the learning of actions, in the presence of similar  $Q(s, a)$  states, as demonstrated in [12]. The proposed network is composed by two parts: the first is a visual feature extractor in charge of transforming the spatial relationships of the state (images) into a feature vector using CNNs. The feature extractor is composed of 3 consecutive convolutional blocks, with 64, 32 and 16 filters each, of size  $3 \times 3$ . Then, 3 layers of fully connected neurons are used. Finally, the 256 output values are processed in parallel by a value head  $V(s)$  and an advantage head  $A(s)$  with a size of 64 neurons each. Using Eq. (9), the final values of  $Q(s, a)$  are obtained. Every activation layer corresponds to a Rectified Linear Unit (ReLU), except for the output layer, which does not have any (see Fig. 6) to effectively represent the real values of  $Q$ .

Furthermore, it has been also implemented a technique useful in the efficient exploration of the state-action domain called Noisy Network [13]. The classic neurons are replaced by noisy neurons that vary their value at each evaluation in a stochastic manner. To the weights of each neuron ( $W, b$ ) are added the parameters ( $W_n, b_n$ ) which are trainable and are weighted with random values ( $\epsilon_w, \epsilon_b$ ) taken from Gaussian distributions ( $\xi_w, \xi_b$ ) (see Fig. 7). In this way, the neural network introduces an intrinsic method of exploration. This is different from classic dropout layers, implemented to avoid over-fitting. The noisy strategy embeds the exploration into the policy and lets the deep agent decide the level of noise in its actions as it learns. This method can be interpreted as a form of evolutionary strategy in which the exploration policy itself is embedded within the agent's action policy. Similar studies, such as [36], have shown that these techniques are able to return better solutions thanks to their enhanced exploration capability.

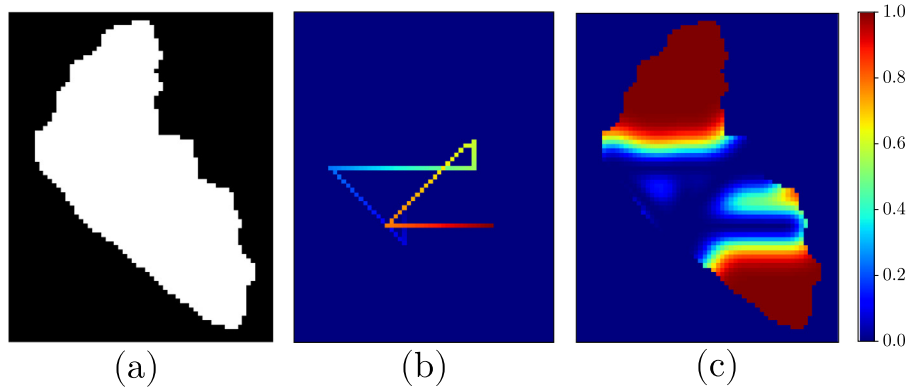


Fig. 5. The 3 channels of the state: (a) corresponds to the navigation map of the Ypacaraí Lake in Paraguay; (b) corresponds to the path traveled so far; (c) corresponds to the uncertainty map given by  $diag(\Sigma [X|X_{meas}])$ .

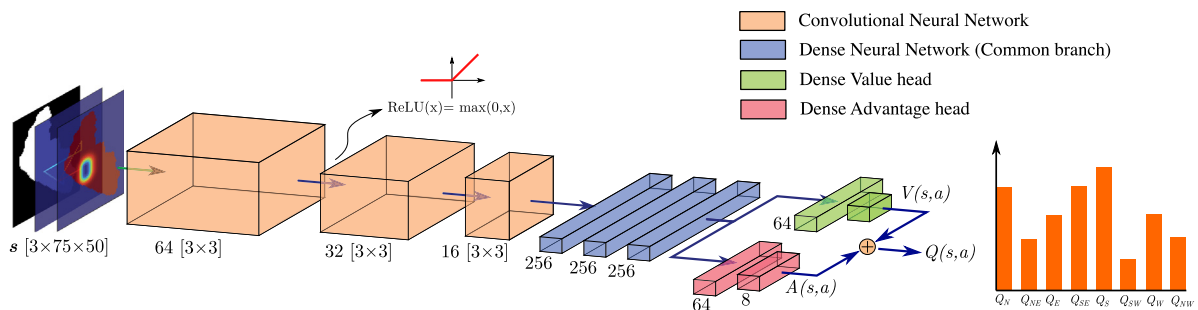


Fig. 6. CNN that represents the  $Q(s, a; \theta)$  function. The dueling-network processes the state to obtain the advantage value  $A(s, a)$  and the state-value  $V(s)$ . All the activation layers corresponds to ReLU function. The output of the network will be 8 Q-values corresponding to expected discounted accumulated reward for taking every action  $a \in A$  in the current state  $s$  (the input of the network). Note that these values does not take in account any invalid (hard-constrained) actions.

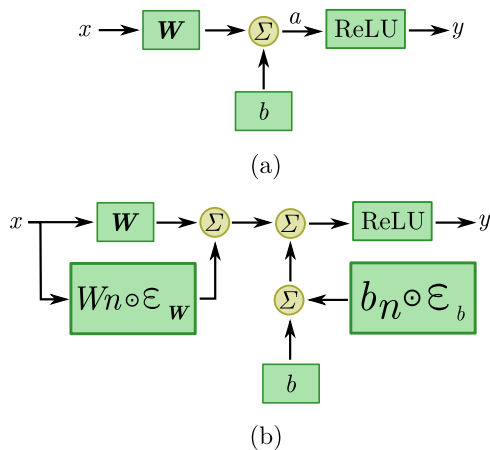


Fig. 7. In (a), it is depicted a classic neuron with ReLU activation function. In (b), a noisy neuron is presented. It is added to the weights and biases a stochastic value weighted by two trainable parameters ( $W_n, b_n$ ).

### 3.4.4. Prioritized experience replay

A fundamental part of DQL is the so-called Experience Replay (ER). As it is mentioned before, the ER consists of sampling batches of previous experiences, saved in a memory buffer  $\mathcal{M}$  as they occur, to fit the  $Q(s, a)$  function. It is commonplace to use uniform random sampling to batch the experiences to avoid correlated values that will lead the NN to overfitting easily. Nonetheless, with this method, every previous experience has the same probability to be sampled, without considering any further information like the knowledge  $Q(s, a)$  has already learned about similar states. To enhance this behavior, the work [11] proposes a

Prioritized Experience Replay (PER) method that emphasizes the learning of those experiences that have bigger TD error in terms of Eq. (8). The PER method imposes a probability  $P_i$  of sampling an experience  $i$  in the buffer proportional to its TD error

$$P_i = \frac{TD_i^\alpha}{(\sum_{\forall E \in Memory} TD)^\alpha} \quad (10)$$

with  $\alpha$  being a parameter of how uniform the sampling is. To avoid the bias generated by the prioritized selection of bigger TD values, the loss generated by these experiences is weighted using Eq. (11). In this importance-based sampling,  $\beta \in [0, 1]$  represents the level of compensation, and in this work it is annealed from a baseline value, to 1, as the learning progresses.

$$w_i = \left( \frac{1}{N} \cdot \frac{1}{P_i} \right)^\beta \quad (11)$$

### 3.4.5. Censoring Q-values

As learning simultaneously to deal with the entropy minimization task and avoiding invalid actions is an arduous task, a simple modification of the classical DQL algorithm can be implemented to overcome this situation. The navigation map can be used to obtain the actions that will cause a collision in a given state  $s$ . Once those invalid actions are calculated, a censoring function  $\eta(s, a) \in \mathbb{R}^{|A|}$  that represents whether an angle of  $\psi$  can be performed or not. Once the invalid actions are computed, the state  $s$  is processed by the CNN and the final censored Q values  $\hat{Q}$  are obtained. The agent can only choose those actions  $a$  that are not marked as invalid by  $\eta$ , as depicted in Eq. (12).

$$\eta(s, a) = \begin{cases} 1, & \text{if } a \text{ is valid} \\ -\infty, & \text{if } a \text{ not valid} \end{cases} \quad (12)$$

$$\hat{Q}(s, a) \leftarrow \eta(s, a) \circ Q(s, a)$$



This technique allows the agent to choose the best actions learned in  $Q(s, a)$  by greedily acting with basic knowledge of the scenario. The final algorithm can be seen in Algorithm 1. This algorithm is similar to the one implemented by [13] but with the censoring mechanism incorporated. When the agent selects a new action with a new noise distribution  $\epsilon$ , the censor mask  $\eta(s)$  is applied depending on its position. This will unequivocally lead to a valid movement without any penalization. When it is time to update the weights of the network using a batch of experiences  $\mathcal{B}$ , it is necessary to also apply the censoring mechanism to the target predictions  $\eta(s_t^{\mathcal{B}}, a') \circ Q^*(s_{t+1}^{\mathcal{B}}, a'; \theta', \epsilon)$ , so no invalid action is considered in greedy bootstrapping selection. This mechanism can be implemented in a straightforward fashion without modification of other parts of the algorithm, such as the PER, noisy networks, or the dueling architecture. This mechanism simplifies learning and reward shaping by eliminating the need for collision-related parameters, such as penalties and terminal conditions. The issue of reward shaping is discussed in [17] and involves not only the definition of penalization, but also the terminal condition of the episode. This condition determines how many collisions are permitted before the episode ends. If the memory buffer is filled with many collision experiences, the agent will overfit and will not learn efficient policies. On the other hand, if not enough collisions are experimented with, there is a high probability of collision, as there is a low probability of sampling collisions from the PER. This method allows all these aspects to be eliminated from the problem by means of a simple hard censoring approach. If navigable contours are available, the most efficient approach is not to learn those dynamics that are available all the time, which is the case for navigation boundaries.

---

#### Algorithm 1 Noisy Censoring-DQL

---

**Input:**  $Env, \alpha, \beta, \gamma, I_R, Memory, \varpi$   
**Output:**  $Q(s, a; \theta, \epsilon)$   
Initialize  $Q(s, a; \theta', \epsilon), Q^*(s, a; \theta'', \epsilon)$   
**while**  $episode \leq Learning\ budget$  **do**  
  Reset  $Env$   
   $s_t \leftarrow s_0$  from  $Env$   
  **while** not done **do**  
     $\epsilon \leftarrow \epsilon'$  ▷ Sample new noisy weights.  
     $\hat{Q}(s_t, a) \leftarrow \eta(s_t, a) \circ Q(s_t, a)$  ▷ Apply safe censoring.  
     $a_t \leftarrow \operatorname{argmax}_a \hat{Q}(s_t, a; \theta, \epsilon)$   
     $s_{t+1}, r_t = Env.step(a_t)$   
    Store  $\langle s_t, a_t, r_t, s_{t+1} \rangle$  in  $Memory$   
    Update  $Memory$  priorities using Eq. (10).  
     $\mathcal{B} \leftarrow ExpRep.sample$  ▷ Prioritized sample.  
     $\epsilon \leftarrow \epsilon''$  ▷ Sample another noisy weights.  
     $\mathbf{y} \leftarrow r_t^{\mathcal{B}} + \gamma \max_{a'} \eta(s_t^{\mathcal{B}}, a') \circ Q^*(s_{t+1}^{\mathcal{B}}, a'; \theta', \epsilon)$   
     $\mathcal{L} = w^{\mathcal{B}} \times [Q(s_t^{\mathcal{B}}, a_t^{\mathcal{B}}; \theta, \epsilon) - \mathbf{y}]^2$   
     $\theta \leftarrow \theta + I_R \times \frac{\partial \mathcal{L}}{\partial \theta}$  ▷ SGD step.  
     $\theta' \leftarrow (1 - \varpi) \times \theta' + \varpi \times \theta$   
    **if**  $Env.distance \geq distance\ budget$  **then**  
      done  $\leftarrow True$   
    **end if**  
  **end while**  
  episode  $\leftarrow episode + 1$   
**end while**

---

#### 3.4.6. Reward function

The design of an adequate reward function is fundamental for any DRL application [17]. The reward function  $r(s, a)$  will define how good an action  $a$  is in a state  $s$  from the perspective of the optimization objective. In this work, the final optimization goal is to reduce the entropy. As the entropy is proportional

to  $|\Sigma[X|X_{meas}]|$ , the reward function must evaluate the actions that decrease this determinant by defining an information gain measure. According to [37], the most effective way to reduce the determinant follows the A-optimal criterion, which is to minimize the average of the eigenvalues  $\lambda$  of  $\Sigma[X|X_{meas}]$ . This will lead to a definition of the information value  $I$  as seen in Eq. (13). This information measure has also been tested in other relevant works such as [5]. The A-optimal criterion consists of reducing the diagonal values of  $|\Sigma[X|X_{meas}]|$  by taking measurements following Eq. (2). This will be equivalent to weighting every action in terms of how much the action reduced the mean uncertainty in the monitoring domain  $X$ .

$$I = \sum_{i=1}^n \lambda_i = \operatorname{tr}(\Sigma[X|X_{meas}]) \quad (13)$$

In the end, the information gain  $\Delta I$  can be defined as the decrease in information when incorporating the next measurement.

$$\Delta I_{t+1|t} = I_t - I_{t+1} \quad (14)$$

If an action has less information than  $\Delta I_{min}$ , a useless movement is considered, and a penalization of  $\kappa < 0$  is applied. When collisions are considered, an action that aims to move the vehicle to an unvisitable zone, a penalty of  $c$ . The final reward function  $r(s, a)$  is represented in Eq. (15).

$$r(s, a) = \begin{cases} c, & \text{if } (s, a) \rightarrow \text{collision} \\ \kappa, & \text{if } \Delta I_{t+1|t} < \Delta I_{min} \\ \Delta I_{t+1|t}, & \text{otherwise} \end{cases} \quad (15)$$

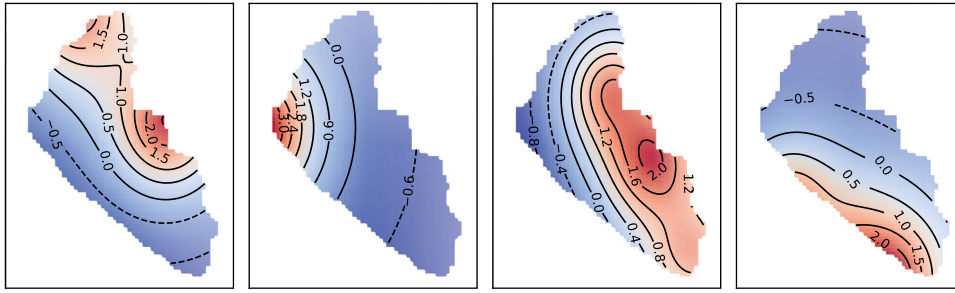
## 4. Results and discussions

This section describes the simulation settings, the assumptions and limitations, the performance metrics related to the monitoring task, the results of such simulations, and a comparative study between our proposed work and previous algorithms and heuristics.

### 4.1. Environment setting and assumptions

The proposed environment used for training will be the lake Ypacaraí (Paraguay, 60 km<sup>2</sup>), but the framework is not limited to this bounding condition. We assume that the vehicle will take one measurement at a time with a distance budget for static missions of 45 km. In the non-stationary case, since we are concerned with continuous patrolling of the waters, the total distance will be 112,5 km. This increase is comprehensible, since the vehicle needs to travel for longer periods to gather changes in the lake dynamics. In relation to WQ parameters, it is imposed that their behavior responds to a randomized benchmark function (Shekel function) previously used for monitoring tasks [18,21]. These functions represent a plausible characterization of real water resources such as the Mar Menor in Murcia.<sup>3</sup> We assume that in the static case the unknown distribution of data does not change, which is plausible for some WQ variables. On the other hand, in the dynamic case, we assume a random Brownian motion with a maximum speed of  $v_{max}$  (see Fig. 8), which is acceptable in terms of the expected changes in the dynamics of the algae blooms, wind conditions, etc. The information model parameters, like the length scale  $l$  of the RBF function or the forgetting factor  $\tau$  (see (1)) are chosen using the known data from the scenario [6], and can be adapted in terms of the data distribution for other use cases.

<sup>3</sup> <https://deap.readthedocs.io/en/master/api/benchmarks>



**Fig. 8.** Four different ground truths generated using a randomized Shekel function. The maxima of these ground truths can move with a constant speed of  $v$  in random directions in the dynamic case.

#### 4.2. Performance metrics

To measure the performance of the method, several metrics are used. These metrics are considered to take into account different aspects of the monitoring: the WQ parameter estimation, outlier detection, and covered area. These metrics will be used in the comparison between algorithms.

- $I_t$ : The non-gathered information, still available at time  $t$ . In the static case, the indicative metric of the performance is the information at the end of a mission  $I_T$ , and in the temporal case, as the  $I_t$  grows in the zones further from the actual position of the ASV, it is considered the average non-gathered information once the initial 45 km are traveled. The lower this metric, the less useful information remains non-gathered in the scenario.
- $A^I$ : Represents the useful information area in ( $\text{km}^2$ ) covered by the ASV. A zone  $x \in \mathbb{R}^2$  is considered covered if the uncertainty  $\sigma(x)$  is less than 0.05. The higher the value of  $A^I$ , the more thorough the coverage.
- **MSE**: Mean square error between a regression model and the ground truth of the WQ variables. Two regression methods have been tested: GP as in [6] and Space Vector Regressor (SVR) [38]. Both regression methods use the same kernel in Eq. (1). For the temporal case, the sample time  $t_{meas}$  is imposed as a variable like in [33]. The lower the MSE, the better the environmental model is obtained.
- $\xi$ : Peak detection rate. In the presence of  $k$  random peaks of algae blooms or contamination, the average rate of detected peaks is defined as  $\xi = \mathbb{E}[k_{detected}/k]$ . A local maximum is considered detected when the uncertainty  $\sigma(x)$  at its location  $x \in \mathbb{R}^2$  is less than 0.05. For the temporal case, to comply the dynamic scenario, the contamination peaks can move through the scenario with max. speed  $v_{max} = 0.3m \cdot s^{-1}$ , as explained in Section Section 4.1. A higher value of  $\xi$  indicates better detection of outliers and peaks.

#### 4.3. Learning settings

All simulations were run using PyTorch on an Ubuntu 20.04 server with an RTX A600 GPU (42 Gb VRAM), 192Gb RAM, and two Intel Xeon Gold 5220R 2.20 GHz. In both static and dynamic scenarios, the agent has been trained for  $1 \times 10^4$  episodes. Note that in the static case, at most and in the absence of collisions, the vehicle will take 67 water quality samples, considering the value of  $d_{meas} = 0.675$  km. In the case of the dynamic scenario, the episode will last 168 steps at most. Considering Algorithm 1, this implies a higher number of updates of the neural network weights. This is in line with the need for more training for the second problem, which has a higher dimensionality and therefore is more complex.

**Table 1**

List of environment parameters, simulation settings, and hyperparameters related to the proposed DQL algorithm.

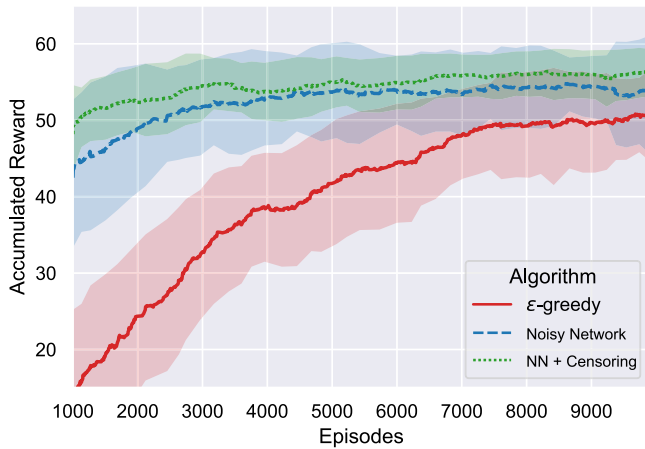
Learning hyperparameter	Value
Learning rate	$1 \times 10^{-4}$
Target update constant $\varpi$	$1 \times 10^{-4}$
Batch size $ \mathcal{B} $	64
Unbias $\beta$ interval	[0.5, 1]
Priority value $\alpha$	0.5
Discount factor $\gamma$	0.99
$\epsilon$ interval	[1, 0.05]
Collision penalty $c$	-1
Redundancy penalty $\kappa$	-0.5
Information threshold $\Delta I_{min}$	0.01
Environment parameter	Value
$d_{meas}$	0.675 km
RBF lengthscale $l$	1.125 km
Forgetting factor $\tau$	0.03
Max. peaks speed $v_{max}$	$0.3 \text{ m s}^{-1}$

The number of episodes has been chosen using a similar order to that of previous similar applications [17]. It has been observed that the number of episodes is sufficient so that, in both cases, the policy converges to a local optimum with a sufficiently adequate execution time. Regarding training times, it is important to note that the use of GPUs is highly recommended. In the static case, the simulation has a training time of approximately 4.5 h for the static case and 11.3 h for the dynamic case, which has a longer duration of episodes.

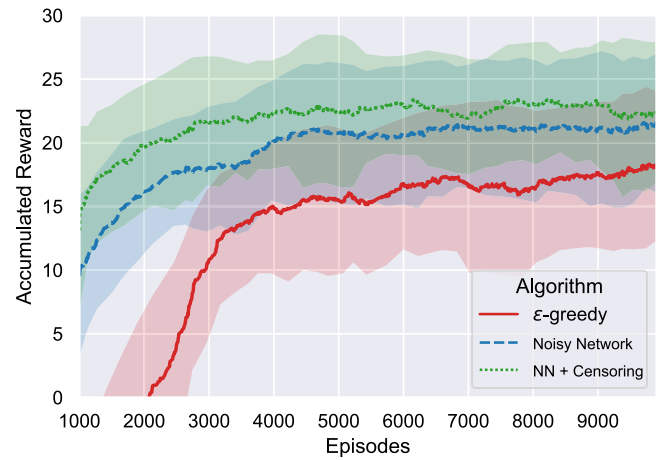
The SGD optimizer is Adam with a learning rate  $l_R$  of  $1 \times 10^{-4}$ . To compare the  $\epsilon$ -greedy exploration policy with the intrinsic exploration of the noisy neural network, separate experiments have been performed under the same conditions. In the case of the *epsilon*-greedy policy, the value of  $\epsilon$  has been imposed to decay within the interval  $[\epsilon_{max}, \epsilon_{min}] = [1, 0.05]$  over the training time. Then,  $\epsilon$  decays from the very beginning of the training to 3000 episodes. From this point, if the exploration phase is extended, no improvement has been observed. In any case, when evaluating the resulting policies, it is imposed a full-greedy action selection without any exploration  $\epsilon = 0$ . All other hyperparameters and other simulation settings are summarized in Table 1.

#### 4.4. Learning results

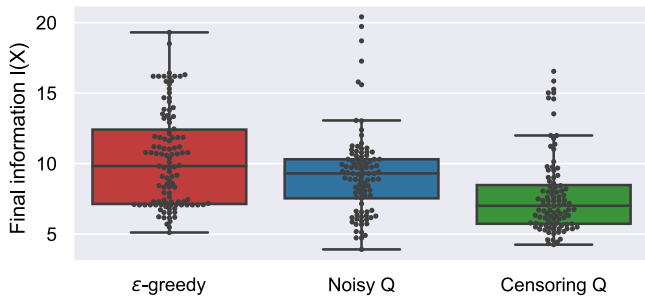
The learning results can be seen in Figs. 9 for the static case. The advantage of using a noisy network can be seen in terms of convergence and training speed. In this first scenario, the use of the noisy network will increase the learning efficiency and, at the same time, eliminate the need to design the  $\epsilon$ -greedy strategy and its hyperparameters. The noisy network converges to an explorative-exploitative balanced policy, which can be seen



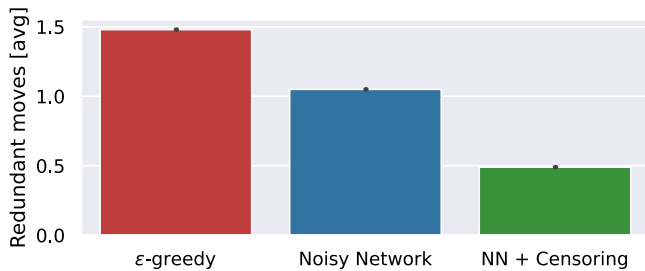
**Fig. 9.** Accumulated reward over training time in the static scenario. It represents the average reward with a moving window of 100 episodes with its standard deviation. Note that Q-Censoring has collision-free training and no penalization is applied.



**Fig. 12.** Accumulated reward over training time in the dynamic scenario. Note that the total reward changes from the static case, as the amount of information, once the map has been covered after a while, is much lower than in the static scenario.



**Fig. 10.** Box plot of information at the end of the mission, for 100 different starting points, in the static case.



**Fig. 11.** Number of useless moves on average, where it is applied a penalization of  $\kappa$  for the three variants of the DRL approach.

as incorporating the exploration behavior into the policy. The Noisy + Q-Censoring strategy, in the static case, slightly overcomes the average reward by approximately 30% with respect to the  $\epsilon$ -greedy counterpart and by 20% with respect to the only Noisy Network case. Although the improvement is not overly significant, the Q-Censoring strategy translates its advantage in the fact that the convergence is reached earlier and, once again, it eliminates the need of designing the collision penalty. In [17], it was studied how the penalty affects convergence and how it can cause poorer performance under certain conditions. In Fig. 10, the entropy is shown at the end of an episode after evaluating the proposed method 100 times starting from random points on the map at each run. The noisy policies return more informative paths (12% improvement), and the noncensored version still provides an effective policy to deal with collisions at the same time the

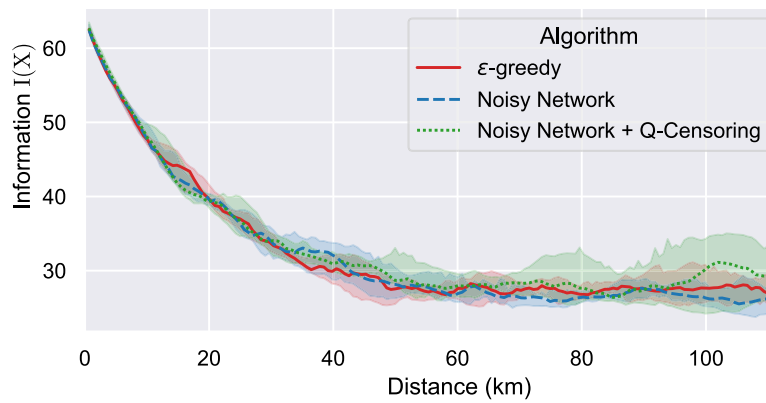
entropy is reduced. It is also noticeable that the number of average redundant/useless movements that produce  $I^t < I_{min}$  (see the reward function Eq. (15)) is 4 times higher in the  $\epsilon$ -greedy strategy and 2.75 times higher in the noncensored version than in the censored version (see Fig. 11). This indicates a better policy in terms of the information gathering capabilities of the proposed method.

Regarding the dynamic case, in Fig. 12 it can be seen how the Q-Censoring strategy overcomes the other methods. As the episode in this particular case is 2,5 times bigger, both the noisy and  $\epsilon$ -greedy strategies tend to generate collisions at a certain point. Using the Q-Censoring strategy, learning is simplified and allows for a higher reward in time with little effort. The Non-Censoring Noisy Network version can be seen as a curricular learning like in [39], where it is assimilated first through the boundaries of the map, and later, the entropy reduction policy is optimized. The censored strategy alleviates this condition, especially in longer episodes, where the possibility of collision is greater. In Fig. 13, it shows the result of the nongathered information  $I_t$  of 100 episodes using the three DRL approaches. Two important aspects must be remarked: (I) As the uncertainty grows within time, it is not possible to gather information with near-zero values, as in the static case. It is only possible to find an equilibrium where the patrolling path reduces the information to a certain point and avoids unnecessary redundant movement. (II), the three DRL methods result in valid temporal patrolling strategies, but the censoring version policy is obtained with 50% less episodes and therefore is the best candidate for more complex scenarios.

#### 4.5. Comparison with other methods

To evaluate the performance of the proposed method in IPP and IPPP, different algorithms have been tested in the Ypacaraí scenario. To fairly compare the ability of the other algorithms to reduce the available information, only safe actions can be performed, considering the navigation map:

- **Safe Random Agent:** Actions are randomly chosen from the action space  $\mathbf{A}$  if they are valid. This algorithm is useful for comparing the difficulty of the patrolling problem, in a way similar to [4].



**Fig. 13.** Comparison of non-gathered information  $I_t$  in the scenario for the three DRL algorithms. It represents the average of 100 missions with different starting points and its standard deviation.

**Table 2**

Metrics obtained by executing all algorithms after 100 episodes for the static case. Bold values correspond to the best performance, and red values to the second best performance.

Algorithm	$I_T$		$A'$		$\xi_{rate}$		$MSE_{GP}$		$MSE_{SVR}$	
	Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.
DRL	<b>7,722</b>	0,30	<b>52,73</b>	3,52	<b>0,602</b>	0,14	<b>0,019</b>	0,10	<b>0,032</b>	0,03
Lawn Mower	30,486	3,81	34,34	3,32	0,393	0,12	0,121	0,13	0,394	0,29
Random	40,302	6,77	22,36	5,52	0,256	0,15	0,741	0,40	0,755	0,57
NRRC	22,128	5,29	34,16	2,96	0,377	0,16	0,252	0,25	0,107	0,13
<i>I</i> -greedy	<b>15,524</b>	0,43	<b>47,18</b>	7,87	<b>0,541</b>	0,16	<b>0,132</b>	0,17	<b>0,116</b>	0,19
EI	29,34	16,19	33,03	15,87	0,34	0,22	0,310	0,22	0,481	0,53

- **Lawn Mower:** The Lawn Mower (LM) algorithm is a strategy that generates intensive coverage paths. The LM agent chooses a random direction to start and goes straight until it is not possible to go further. Then, the agent returns in a parallel straight line, one step forward from the previous path.
- **Non-Redundant Random Coverage:** The Non-Redundant Random Coverage (NRRC) algorithm randomly selects a direction of movement from  $\mathbf{A}$ , different from the previously selected to avoid retracing its steps. Once the vehicle cannot advance in this direction, a new direction is sampled. This is a variation of the Intelligent Random algorithm used for comparisons in [15].
- ***I*-greedy Strategy:** The *I*-greedy strategy selects the highest point of uncertainty on the map and, using a local path planner, guides the vehicle to this point without collisions.
- **Expected Improvement Acquisition Function:** This approach takes the Expected Improvement (EI) Acquisition Function as a policy that indicates the next valid best point to visit [6]. The EI is a function that tries to optimize the Gaussian Process mean by searching for a higher value than the maximum available. To adapt this algorithm to our framework, we maintain the movement restriction imposed by  $d_{meas}$ .

The results of the static case are presented in Table 2, for 100 different starting point routes. These results clearly show how the proposed method is able to obtain the most informative paths compared to the other algorithms. The information obtained at the end of the episode is 49% lower than the second best algorithm (*I*-greedy). The DRL agent is able to gather more information in much less time, as seen in Fig. 14. The DRL policy is also less affected by the starting point in terms of the standard deviation of the information, which can be seen as the robustness of the informative trajectories. The proposed algorithm also obtains a policy capable of detecting 60% randomly located contamination

peaks in the environment, which is a 51% improvement on average over the other algorithms. Improvement is also significant in terms of MSE using both regression methods. On the one hand, when using a GP, the MSE is reduced to 0.019, which represents an average error of 13.7% across the water surface (see Fig. 15) On the other hand, the SVR returns an average MSE of 0.032, again an average 17,8% error, lower than in any other acquisition strategy tested. The results of the EI approach indicates that the acquisition function is too greedy with respect to the GP mean. Even when is able to reach the maxima of the scalar field easily, as in Bayesian Optimization approaches [6,18], it is not explorative enough. This algorithm also suffers when imposing a truncated traveling distance  $d_{meas}$  without any prior hyperparameterization tailored to the ground truth value. Therefore, our proposal seems to be more robust in terms of the exploration as it does not depend on any regression values, only in the uncertainty.

It is noticeable that the acquisition path is relevant to the modeling problem and also to the regression method used. GPs obtained, on average, better solutions for the static case, independent of the path. Nevertheless, it is proven that the error in the model is conditioned on the information gathered, which is the reduction of the entropy. In this static case, the lower the entropy, the lower the error in the estimation for different regressors. This suggests that optimizing under this informative entropic criterion produces useful coverage paths for many purposes: model estimation, patrolling contamination peaks, complete coverage, etc.

Finally, if the final paths are studied as seen in Fig. 16, it is clear that they are less redundant in the proposed DRL policy. The paths generated by the LM approach tend to be too thorough and do not take into account the redundancy of the information in consecutive and near samples. In a different way, the greedy approach does not consider any further information than the maximum uncertainty position, which inevitably leads to redundant subpaths in the way of reaching the most informative point. The proposed algorithm takes into account not only the next point, but also the future discounted information available, and

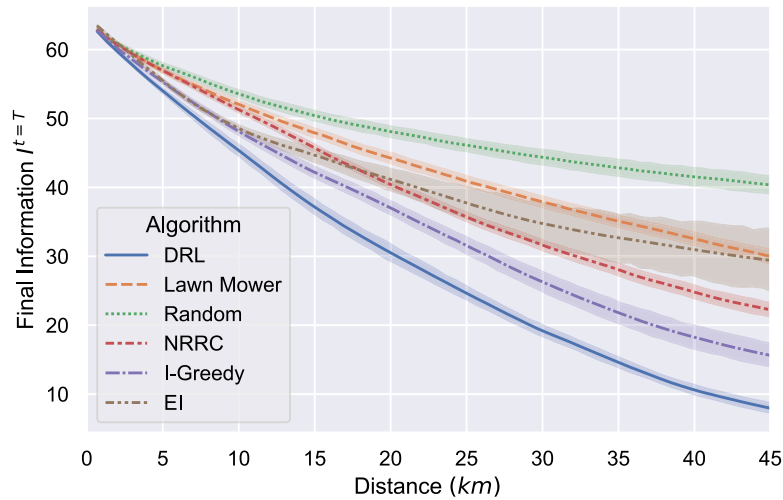


Fig. 14. Comparison of the average information in the static scenario for all algorithms after 100 missions from different starting points and its standard deviation.

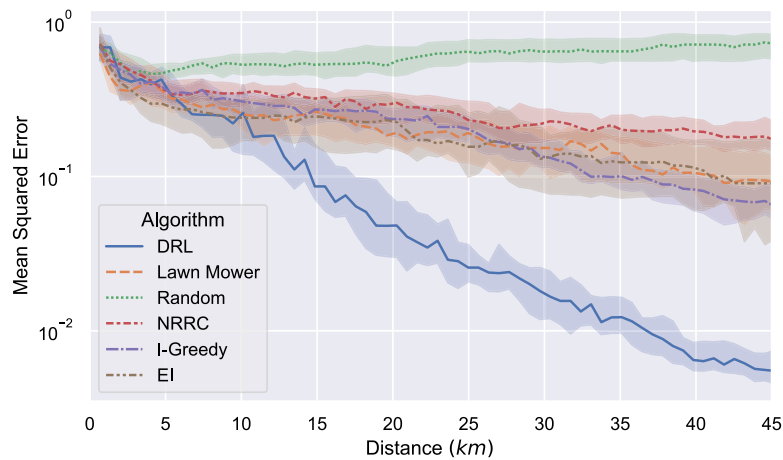


Fig. 15. Mean squared error using a GP with the proposed kernel after 100 missions (static case) with random ground truths for every algorithm. The average is represented within its standard deviation.

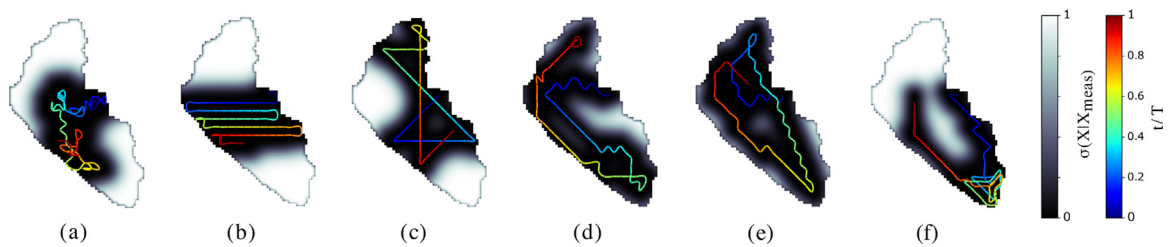


Fig. 16. Paths generated by common coverage path planning algorithms in the static scenario: (a) Random, (b) Lawn mower, (c) Nonredundant Random Coverage, (d) *I*-greedy, (e) the proposed DRL approach, and (f) Expected Improvement AF. All paths start in this figure from the same deploy zone for comparison. The path is colored according to the time  $t$  divided by the mission time  $T$ . The grayscale background color represents the predictive uncertainty  $\sigma(x)$  of a particular zone  $x$ .

obtains, even in the final steps when there is little information left, significant measurements.

In the temporal scenario, the results can be interpreted in a similar way (see Table 3 for complete results). The proposed algorithm is able to obtain 30% more informative paths on average than the other proposed heuristics. The DRL policy produces informative patrolling paths that can reduce entropy and maintain stability as information from the unvisited zone regenerates (see Fig. 17). When analyzing the ability of the algorithm to detect random peaks, the proposed DRL approach obtains similar but slightly better results than the second-best algorithm (LM). The LM algorithm, as it produces very extensive coverage paths, is

able to detect those moving peaks, but at the expense of the surrogate model accuracy, which is far from the one obtained using the DRL method (a 51% worse contamination model).

Both regression methods used in the temporal case are modified to consider the acquisition time of a sample as a measured parameter, as explained above (see (5)). The RBF kernels are consequently modified to include a time-dependent third dimension. It was also considered an estimation horizon of 67 samples in the past (the final number of samples in the static scenario), so the very first measurements were discarded for estimation, as they cannot be assumed valid anymore. With this simple temporal regression method, both GPR and SVR are able to converge in

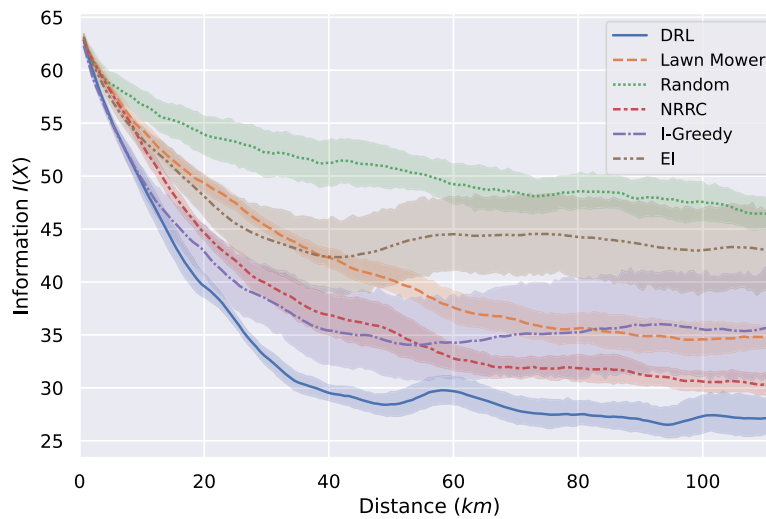


Fig. 17. Comparison of the average information in the temporal scenario for all algorithms after 100 missions from different starting points and its standard deviation.

Table 3

Metrics obtained by executing all algorithms after 100 episodes for the temporal case. Bold values correspond to the best performance, and red values to the second best performance. Those MSE values with  $> 1$  correspond to those cases in which the regression method fails to converge and returns a useless model with errors greater than 1.

Algorithm	$I_T$		$A^T$		$\xi_{rate}$		$MSE_{GP}$		$MSE_{SVR}$	
	Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.
DRL	<b>26.673</b>	4.75	<b>809.5</b>	32.4	<b>0.725</b>	0,10	<b>0.136</b>	0,20	<b>0.101</b>	0,07
Lawn Mower	36.765	4.15	<b>913.7</b>	96.0	<b>0.700</b>	0,20	0.342	0,22	> 1	> 1
NRRC	46.585	3.63	811.9	34.5	0.325	0,14	0.449	0,42	> 1	> 1
Coverage	<b>29.904</b>	2.52	673.2	36.6	0.565	0,16	<b>0.184</b>	0,20	<b>0.417</b>	0,20
I-greedy	37.755	11.77	587.1	163.9	0.525	0,20	0.471	0,65	> 1	> 1
EI	43.771	13.72	444.38	176.4	0,35	0,23	0,57	0,51	> 1	> 1

temporal-dependent models of the WQ variables with MSE values of 0,136 and 0,101, respectively, for each regressor. These values are obtained on average once the ASV has completed the first 45km (the distance limit in the first static scenario). As can be induced from Table 3, the MSE values are higher than in the static case, as the benchmark function changes over time. The model can only be adjusted when the old zones are revisited and the regressor updates its state information. In a further test, after 20 missions starting from a single starting point, the resulting MSE using GP is not only lower on average, but the lower bound is reached earlier in the DRL proposal (see Fig. 18). In the SVR, it is worth mentioning that not all policies result in a convergence of the surrogate model. In the case of the Random Agent, the NRRC, and the greedy approach, the measurements do not conform to an adequate set to converge over time. However, when converging, as happens in the DRL approach, the error can be even lower than in the GPR counterpart. With regards to the EI application, it suffers again from the things mentioned before. The routes seem to be more redundant, as the EI relies not only in the uncertainty but also in the mean of the GP. The EI-agent tries to revisit those places with higher mean instead of exploring high uncertainty values. This results in a lack of exploratory behavior that is not seen in, for example, the I-greedy agent or in the proposed DRL policy. It has to be further analyzed the sensitivity of the hyperparameters in the EI policy, nonetheless, it has been clear that it is not a suitable choice for extensive search but for sample-efficient missions without any movement restrictions imposed by  $d_{meas}$ . Nonetheless, sample-efficiency will always involves a trade-off between number of samples and robustness in the accuracy and anomalies detection (given the performance in the metric  $\xi_{rate}$ ).

The example paths in Fig. 19 show how the DRL policy tends to generate cyclic patrol routes to keep the growing uncertainty

low. LM algorithm, while intensive in the final steps, tends to generate high redundancy coverage. The greedy approach, as the maximum uncertainty is always growing, tends to oscillate and fails the patrolling task. The NRRC is the only algorithm that can compete in terms of the MSE model but easily generates long redundant paths.

#### 4.6. Discussion of the results

The proposed algorithm has been proven to be effective in both IPP and IPPP applications in terms of monitoring metrics and learning efficiency. It is important to highlight the following aspects of the results:

- The use of noisy neural networks allows for efficient exploration of the action-state domain, resulting in an improvement in the information collected on the routes (20% improvement). Additionally, the design of the *epsilon*-greedy strategy is no longer necessary.
- The use of Q-censoring allows for better results than baselines (8% improvement) in the static case and an increase in convergence speed (50% faster).
- The use of Q-censoring also allows the user to not impose any collision penalty, resulting in an improvement in policy performance in terms of information collected.
- The agent in the static case produces significantly better monitoring results than the other heuristics (13.7% on average among all algorithms in terms of MSE).
- In the temporal case, the improvement is not only significant in terms of model error (8% better than the best among the other algorithms), but also in other metrics, such as the detection rate of contamination agents (37% more detections on average among the other algorithms).

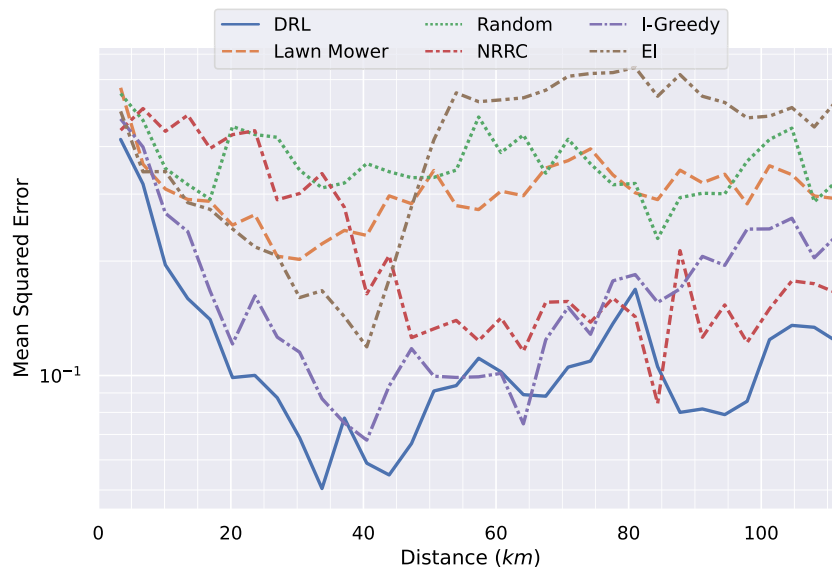


Fig. 18. Mean squared error using a GP with the proposed kernel after 20 missions starting from the same point (temporal case) with random ground truths. The standard deviation is omitted for a better visualization.

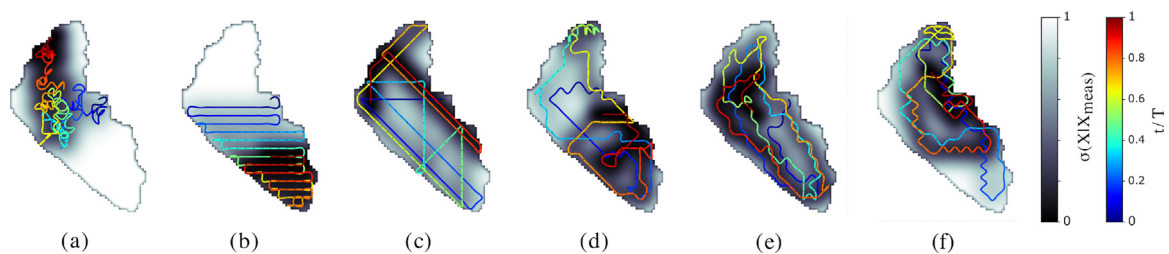


Fig. 19. Paths generated by common coverage path planning algorithms in the temporal scenario: (a) random, (b) lawn mower, (c) nonredundant random coverage, (d) *I*-greedy, (e) the proposed DRL approach, and (f) Expected Improvement AF. The path is colored according to the time  $t$  divided by the mission time  $T$ . The gray background color represents the uncertainty  $\sigma(x)$  of a particular zone  $x$ .

- On a qualitative level, the routes of both algorithms are less redundant and better understand the boundaries of the map and the information constraints on them.

### 5. Conclusions and future lines

The monitoring and patrolling task in big water resources requires an efficient path planning policy that gathers the most information available at the same time the navigation boundaries are considered. This work proposes a Deep Reinforcement Learning framework that optimizes a convolutional policy based on a visual state of the informative scenario using the entropy minimization as the surrogate objective for the patrolling task. Knowledge assimilation is enhanced using state-of-the-art methods such as Noisy Networks and Prioritized Experience Replay to dive into a more exploratory policy that can deal with two cases of monitoring: the static scenario and the temporal patrolling case. Ultimately, a Q-Censoring strategy has been proposed that permits scenario-gnostic learning to avoid dangerous movements towards the land. This approach has proven to be much more efficient than the classic  $\epsilon$ -greedy strategy and the noisy baseline, especially in the static case. The difference also lies in the convergence speed. As the Q-Censored approach does not deal with the collision problem, as in [7,15], consequently the speed of learning is significantly increased. The result of this framework will allow the establishment of efficient routes in the collection

of water quality data and in many other applications, such as the measurement of radiation levels, gas sources, etc., without loss of generality. This framework is particularly useful due to the vehicle safety criterion and due to the formulation of information based exclusively on uncertainty.

The algorithm, in the end, results in a more useful policy than other approaches like the Lawn Mower or *I*-greedy, since it learns how to plan in an optimization horizon according to the tailored reward function. There is also improvement from previous approaches using Expected Improvement as an acquisition function, since pursuing the reduction of the uncertainty seems to be a more robust criterion that balancing the exploitation/exploration like in Bayesian optimization approaches. This improvement in entropy minimization is translated into lower model errors and higher detection rates of contamination peaks. However, regression methods have been shown to be very sensitive to acquisition policies and require further research. The convergence of the regression is not guaranteed, and the estimation error could be used as part of the reward function. Another future line is the generalization of the algorithm to the multiagent paradigm. Like in [21], this approach can be applied to multiple cooperative agents using the same uncertainty map to efficiently collect information. In this regard, it is important to address the question of how different agents can effectively cooperate to cover the scenario. This raises the problem of not only static obstacle avoidance, but also between-agent collision avoidance. A future use of multiagent Q-censoring will be necessary.

## CRedit authorship contribution statement

**Samuel Yanes Luis:** Conceptualization, Methodology, Investigation, Software, Validation, Writing – original draft. **Daniel Gutiérrez-Reina:** Conceptualization, Methodology, Resources, Writing – review, Supervision, Project administration, Funding acquisition. **Sergio Toral Marín:** Conceptualization, Methodology, Resources, Writing – review, Supervision, Project administration, Funding acquisition.

## Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Samuel Yanes Luis reports administrative support was provided by University of Seville.

## Data availability

No data was used for the research described in the article.

## Acknowledgments

This work has been partially funded by the Spanish “Ministerio de Ciencia, Innovación y Universidades, Programa Estatal de I+D+i Orientada a los Retos de la Sociedad” under the Project “Despliegue Adaptativo de Vehículos no Tripulados para Gestión Ambiental en Escenarios Dinámicos RTI2018-098964-B-I00”, and under the PhD grant FPU-2020 (Formación del Profesorado Universitario) of Samuel Yanes Luis, by the regional government Junta de Andalucía, Spain under the Projects “Despliegue Inteligente de una red de Vehículos Acuáticos no Tripulados para la monitorización de Recursos Hídricos US-1257508” and “Despliegue y Control de una Red Inteligente de Vehículos Autónomos Acuáticos para la Monitorización de Recursos Hídricos Andaluces, Spain PY18-RE0009”.

## References

- [1] J. Álvarez Rogel, G. Barberá, et al., The case of mar menor eutrophication: State of the art and description of tested nature-based solutions, *Ecol. Eng.* 158 (2020) 106086, <http://dx.doi.org/10.1016/j.ecoleng.2020.106086>, URL <https://www.sciencedirect.com/science/article/pii/S09255857420303748>.
- [2] S.B. Watson, C. Miller, et al., The re-eutrophication of lake erie: Harmful algal blooms and hypoxia, *Harmful Algae* 56 (2016) 44–66, <http://dx.doi.org/10.1016/j.hal.2016.04.010>, URL <https://www.sciencedirect.com/science/article/pii/S1568988315301141>.
- [3] J. Sánchez-García, J. García-Campos, M. Arzamendia, D. Reina, S. Toral, D. Gregor, A survey on unmanned aerial and aquatic vehicle multi-hop networks: Wireless communications, evaluation tools and applications, *Comput. Commun.* 119 (2018) 43–65, <http://dx.doi.org/10.1016/j.comcom.2018.02.002>.
- [4] M. Arzamendia, D. Gregor, D. Gutierrez-Reina, S. Toral, An evolutionary approach to constrained path planning of an autonomous surface vehicle for maximizing the covered area of Ypacarai Lake, *Soft Comput.* 23 (5) (2019) 1723–1734.
- [5] M. Popovic, T. Vidal-Calleja, G. Hitz, et al., An informative path planning framework for UAV-based terrain monitoring, *Auton. Robots* 44 (2020) 889–911, <http://dx.doi.org/10.1007/s10514-020-09903-2>.
- [6] F. Peralta Samaniego, D.G. Reina, S.L. Toral Marín, D.O. Gregor, M. Arzamendia, A Bayesian optimization approach for water resources monitoring through an autonomous surface vehicle: The Ypacarai Lake case study, *IEEE Access* 9 (1) (2021) 9163–9179, <http://dx.doi.org/10.1109/ACCESS.2021.3050934>.
- [7] T. Wiedemann, C. Vlaicu, J. Josifovski, A. Viseras, Robotic information gathering with reinforcement learning assisted by domain knowledge - An application to gas source localization, *IEEE Access* 9 (1) (2021) 13159–13172, <http://dx.doi.org/10.1109/ACCESS.2021.3052024>.
- [8] Y. Brouwer, A. Vale, R. Ventura, Informative path planner with exploration-exploitation trade-off for radiological surveys in non-convex scenarios, *Robot. Auton. Syst.* 136 (2021) 103691, <http://dx.doi.org/10.1016/j.robot.2020.103691>, URL <https://www.sciencedirect.com/science/article/pii/S0921889020305315>.
- [9] C. Rasmussen, C. Williams, *Gaussian Processes for Machine Learning*, in: *Adaptive Computation and Machine Learning*, MIT Press, Cambridge, MA, USA, 2006, p. 248.
- [10] M. Hessel, J. Modayil, et al., Rainbow: Combining improvements in deep reinforcement learning, *CoRR abs/1710.02298*, 2017, [arXiv:1710.02298](https://arxiv.org/abs/1710.02298).
- [11] T. Schaul, J. Quan, I. Antonoglou, D. Silver, Prioritized experience replay, 2015, URL <http://arxiv.org/abs/1511.05952>. Published at ICLR 2016.
- [12] Z. Wang, N. de Freitas, M. Lanctot, Dueling network architectures for deep reinforcement learning, *CoRR abs/1511.06581*, 2015, [arXiv:1511.06581](https://arxiv.org/abs/1511.06581).
- [13] M. Fortunato, M.G. Azar, B. Piot, J. Menick, I. Osband, A. Graves, V. Mnih, R. Munos, D. Hassabis, O. Pietquin, C. Blundell, S. Legg, Noisy networks for exploration, *CoRR abs/1706.10295*, 2017, [arXiv:1706.10295](https://arxiv.org/abs/1706.10295).
- [14] R.S. Sutton, A.G. Barto, *Reinforcement Learning: An Introduction*, second ed., The MIT Press, 2018, URL <http://incompleteideas.net/book/the-book-2nd.html>.
- [15] A. Viseras, R. Garcia, DeepIG: Multi-robot information gathering with deep reinforcement learning, *IEEE Robot. Autom. Lett.* 4 (3) (2019) 3059–3066, <http://dx.doi.org/10.1109/LRA.2019.2924839>.
- [16] H. Ferreira, C. Almeida, A. Martins, J. Almeida, N. Dias, A. Dias, E. Silva, Autonomous bathymetry for risk assessment with ROAZ robotic surface vehicle, in: *Oceans 2009-Europe*, 2009, pp. 1–6, <http://dx.doi.org/10.1109/OCEANSE.2009.5278235>.
- [17] S. Yanes, D.G. Reina, S.L. Toral Marín, A deep reinforcement learning approach for the patrolling problem of water resources through autonomous surface vehicles: The Ypacarai Lake case, *IEEE Access* 6 (1) (2020) 1, <http://dx.doi.org/10.1109/ACCESS.2020.3036938>.
- [18] F. Peralta, D.G. Reina, S. Toral, M. Arzamendia, D. Gregor, A Bayesian optimization approach for multi-function estimation for environmental monitoring using an autonomous surface vehicle: Ypacarai Lake case study, *Electronics* 10 (8) (2021) 963.
- [19] M. Arzamendia, D. Gutierrez, S. Toral, D. Gregor, E. Asimakopoulou, N. Bessis, Intelligent online learning strategy for an autonomous surface vehicle in lake environments using evolutionary computation, *IEEE Intell. Transp. Syst. Mag.* 11 (4) (2019) 110–125.
- [20] M.J.T. Kathen, I.J. Flores, D.G. Reina, An informative path planner for a swarm of ASVs based on an enhanced PSO with Gaussian surrogate model components intended for water monitoring applications, *Electronics* 10 (13) (2021) 1605.
- [21] S. Yanes, D.G. Reina, S.L.T. Marín, A multiagent deep reinforcement learning approach for path planning in autonomous surface vehicles: The Ypacarai Lake patrolling case, *IEEE Access* 9 (2021) 17084–17099.
- [22] B. Xin, H. Yu, Y. Qin, Q. Tang, Z. Zhu, Exploration entropy for reinforcement learning, *Math. Probl. Eng.* 2020 (2020) 1–12, <http://dx.doi.org/10.1155/2020/2672537>.
- [23] M. Chen, A. Liu, W. Liu, K. Ota, M. Dong, N. N. Xiong, RDRL: A recurrent deep reinforcement learning scheme for dynamic spectrum access in reconfigurable wireless networks, *IEEE Trans. Netw. Sci. Eng.* 9 (2) (2022) 364–376, <http://dx.doi.org/10.1109/TNSE.2021.3117565>.
- [24] D. Zhang, Y. Fu, Z. Zheng, UAST: Uncertainty-aware siamese tracking, in: *Proceedings of the 39th International Conference on Machine Learning*, Vol. 162, PMLR, 2022, pp. 26161–26175.
- [25] D. Zhang, Z. Zheng, R. Jia, M. Li, Visual tracking via hierarchical deep reinforcement learning, *Proc. AAAI Conf. Artif. Intell.* 35 (4) (2021) 3315–3323, <http://dx.doi.org/10.1609/aaai.v35i4.16443>, URL <https://ojs.aaai.org/index.php/AAAI/article/view/16443>.
- [26] A. Krishna Lakshmanan, R. Elara Mohan, B. Ramalingam, A. Vu Le, P. Veerajagadeshwar, K. Tiwari, M. Ilyas, Complete coverage path planning using reinforcement learning for tetromino based cleaning and maintenance robot, *Autom. Constr.* 112 (May 2019) (2020) 103078, <http://dx.doi.org/10.1016/j.autcon.2020.103078>.
- [27] S. Zhang, Y. Li, Q. Dong, Autonomous navigation of UAV in multi-obstacle environments based on a deep reinforcement learning approach, *Appl. Soft Comput.* 115 (2022) 108194, <http://dx.doi.org/10.1016/j.asoc.2021.108194>, URL <https://www.sciencedirect.com/science/article/pii/S1568494621010383>.
- [28] P. Zieliński, U. Markowska-Kaczmar, 3D robotic navigation using a vision-based deep reinforcement learning model, *Appl. Soft Comput.* 110 (2021) 107602, <http://dx.doi.org/10.1016/j.asoc.2021.107602>, URL <https://www.sciencedirect.com/science/article/pii/S1568494621005238>.
- [29] S. Koh, B. Zhou, H. Fang, P. Yang, Z. Yang, Q. Yang, L. Guan, Z. Ji, Real-time deep reinforcement learning based vehicle navigation, *Appl. Soft Comput.* 96 (2020) 106694, <http://dx.doi.org/10.1016/j.asoc.2020.106694>, URL <https://www.sciencedirect.com/science/article/pii/S1568494620306323>.
- [30] M. Theile, H. Bayerlein, R. Nai, D. Gesbert, M. Caccamo, UAV coverage path planning under varying power constraints using deep reinforcement learning, in: *2020 IEEE/RISJ International Conference on Intelligent Robots and Systems, IROS, IEEE, 2020*, pp. 1444–1449.
- [31] C. Piciarelli, G.L. Foresti, Drone patrolling with reinforcement learning, *ACM Int. Conf. Proc. Ser.* 1 (1) (2019) 1–6, <http://dx.doi.org/10.1145/3349801.3349805>.



- [32] T.M. Cover, J.A. Thomas, *Elements of Information Theory*, in: *Wiley Series in Telecommunications and Signal Processing*, Wiley-Interscience, USA, 2006.
- [33] S. Morales-Enciso, J. Branke, Tracking global optima in dynamic environments with efficient global optimization, *European J. Oper. Res.* 242 (3) (2015) 744–755, <http://dx.doi.org/10.1016/j.ejor.2014.11.028>, URL <https://www.sciencedirect.com/science/article/pii/S0377221714009515>.
- [34] V. Mnih, K. Kavukcuoglu, D. Silver, et al., Human-level control through deep reinforcement learning, *Nature* 518 (7540) (2015) 529–533, <http://dx.doi.org/10.1038/nature14236>.
- [35] S. Yanes Luis, D. Gutiérrez-Reina, S. Toral Marín, A dimensional comparison between evolutionary algorithm and deep reinforcement learning methodologies for autonomous surface vehicles with water quality sensors, *Sensors* 21 (8) (2021) <http://dx.doi.org/10.3390/s21082862>, URL <https://www.mdpi.com/1424-8220/21/8/2862>.
- [36] X. Zhang, H. Duan, An improved constrained differential evolution algorithm for unmanned aerial vehicle global route planning, *Appl. Soft Comput.* 26 (2015) 270–284, <http://dx.doi.org/10.1016/j.asoc.2014.09.046>, URL <https://www.sciencedirect.com/science/article/pii/S1568494614004992>.
- [37] R. Sim, N. Roy, Global A-optimal robot exploration in SLAM, in: *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*, 2005, pp. 661–666, <http://dx.doi.org/10.1109/ROBOT.2005.1570193>.
- [38] X. Wang, L. Fu, C. He, Applying support vector regression to water quality modelling by remote sensing data, *Int. J. Remote Sens.* 32 (23) (2011) 8615–8627, <http://dx.doi.org/10.1080/01431161.2010.543183>.
- [39] J. Kang, M. Liu, A. Gupta, C. Pal, X. Liu, J. Fu, Learning multi-objective curricula for deep reinforcement learning, *CoRR* abs/2110.03032, 2021, [arXiv:2110.03032](https://arxiv.org/abs/2110.03032).