

# Responding to Illegal Activities Along the Canadian Coastlines Using Reinforcement Learning

Mohammed Abouheaf, Shuzheng Qu, Wail Gueaieb,  
Rami Abielmona, and Moufid Harb

Machine learning (ML) algorithms can prove to be instrumental in certain complex ill-conditioned systems when inserted as a middle layer to interface low-level hardware, such as sensors and actuators, and high-level decision-making kernels. Such an interface provides a secondary, or supervisory, conditioning layer that would enhance the system’s robustness in the face of various types of uncertainties and disturbances. This article elaborates how ML can leverage the solution of a contemporary problem related to the security of maritime domains.

The worldwide “Illegal, Unreported, and Unregulated” (IUU) fishing incidents have led to serious environmental and economic consequences which involve drastic changes in our ecosystems in addition to financial losses caused by the depletion of natural resources. The term “illegal” refers either to contravening the fishing regulations or conducting fishing activities unlawfully under the jurisdiction of a certain maritime territory. In this context, “unreported” indicates the failure to report to the relevant authorities, while “unregulated” describes fishing activities in areas with no applicable measures to control the catch. The Fisheries and Aquatic Department (FAD) of the United Nation’s Food and Agriculture Organization (FAO) issued a report which indicated that the annual losses due to IUU fishing reached \$25 Billion [1]. This imposes negative impacts on the future-biodiversity of the marine ecosystem and domestic Gross National Product (GNP). Maritime Domain Awareness (MDA) is concerned with the situational awareness relevant to the activities that have direct or indirect impacts on the operational, organizational, economical, and safety sides of the maritime domains [2, 3]. The assets in MDA support different missions, such as monitoring, risk assessment, and management. Hence, robust interception mechanisms are increasingly needed for detecting and pursuing the unrelenting illegal fishing incidents in maritime territories.

## Decision-Making System

The IUU fishing problem can be broken down into at least three sub-problems:

1. Recognizing an IUU activity, where radar and satellite monitoring systems are deployed to detect the unreported event and provide live information about its status, such as the vessels’ location measurements, for instance.

---

\*This work was partially supported by NSERC Grant EGP 537568-2018.

\*Mohammed Abouheaf, Shuzheng Qu and Wail Gueaieb are with the School of Electrical Engineering & Computer Science, University of Ottawa, Ottawa, Canada. E-mail: {mabouhea,fqu096,wgueaieb}@uottawa.ca. Rami Abielmona and Moufid Harb are with Larus Technologies, Ottawa, Ontario, Canada. E-mail: {rami.abielmona,moufid.harb}@larus.com.

2. Deciding on the resources (coast guard vessels, military planes, etc.) to dispatch in the pursuit of the IUU vessels. The decision may depend on many factors, including the locations of the pursuing vehicles, their speeds, cost of operation, and fuel consumption, to name a few. Herein, autonomous sea vessels are considered to pursue the illegal shipping activity.
3. Coordinating the motion of the fleet of pursuers to catch the IUU vessel while still in local waters.

In this article, only the latter task is addressed. It is formulated as a pursuer-evader problem that is tackled within an ML framework. One or more pursuers, such as law enforcement vessels, intercept an evader (i.e., the illegal fishing ship) using an online reinforcement learning mechanism that is based on a value iteration process. It employs real-time navigation measurements of the evader ship as well as those of the pursuing vessels and returns back model-free interception strategies. Means of adaptive critics are employed to approximate the learning process using actor-critic neural networks, where a gradient descent approach is followed to tune the weights of the neural networks. Test cases are simulated with the aid of Total::Perception™ (T::P) engine, a professional predictive analytics decision-making software that is developed to mimic the situational awareness in maritime as well as other domains [4, 5].

## Machine Learning Literature

The significant advances in the architectures of digital processing units, enabled the implementation of sophisticated artificial intelligence algorithms on integrated microprocessors [6]. Such a technological breakthrough opened a new window of engineering application opportunities that were not possible a few years ago. This includes, for example, the visual detection and autonomous refueling of UAVs [7], robot motion planing [8], brain MRI estimation methods [9], and latency measurement of distributed networks [10], to name a few. In a particular context, Approximate Dynamic Programming (ADP) is a machine learning tool that has been employed to tackle the computational challenges associated with the action and state spaces of the dynamic programming solutions. ADP can be classified into four classes, namely, Heuristic Dynamic Programming, Dual Heuristic Dynamic Programming, Action Dependent Heuristic Dynamic Programming, and Action Dependent Dual Heuristic Dynamic Programming. These classes are used to solve the Hamilton-Jacobi (HJ) or Hamilton-Jacobi-Bellman (HJB) equations for different optimization problems for single and multi-agent systems [11–13]. This is done in order to optimize a utility function associated to the dynamic system under consideration. Devising solutions for multi-agent optimization problems can be computationally challenging due to the coupling effects in the optimality equations [14]. Furthermore, the complexity increases even further when trying to solve that problem in a distributed fashion using only local information available to each individual agent.

Reinforcement Learning (RL) is a ML mechanism that is adaptive to the different ADP solutions [15]. It is concerned with selecting the best strategy-to-apply in a dynamic learning process that penalizes or rewards such strategy using a given cost function. This enables the system to transit from one state to another which represents a better choice in the path of finding an optimal solution to the problem in hand [16]. The RL process mimics the analytical solution environment of the optimization problems where the optimal strategy and its value are interrelated [11]. The RL platform employs temporal difference solutions arising from the various ADP classes. However, these solutions may get overwhelmingly complex especially in the case of multi-agent systems. Hence, solvers that are based on combined ideas of game theory, cooperative optimization, and RL

are used to provide distributed and real-time solutions for such problems. The RL solutions are implemented using one of two two-step techniques known as value iteration and policy iteration methods. They differ at how the strategy is evaluated and then updated accordingly [15]. Within the same context, the adaptive actor-critic structures are approximation tools used by the RL algorithms. The actor is a neural network that assesses the quality of the adopted policy, while the critic is a neural network that estimates the cost-to-go from each state to a final following a certain policy [17].

## Pursuer-Evader Problem Setup

Consider a system of one evader and  $N$  pursuers. Each vessel  $j$  is associated with an (ENU) Cartesian coordinate system  $\text{RF}_j$  pinned at the vessel's position, such that vector  $\mathbf{E}_j$  points to the vessel's heading and the 2D plane  $(\mathbf{E}_j, \mathbf{N}_j)$  is tangent to the earth surface, which is modeled as a sphere of radius  $R$  (see Figure 1). Let  ${}^i(\cdot)$  denote the quantity  $(\cdot)$  expressed in  $\text{RF}_i$ ; and  $({}^i x_k^j, {}^i y_k^j) \in \mathbb{R}^2$ ,  ${}^i v_k^j \in \mathbb{R}$ , and  ${}^i \theta_k^j \in SO(2)$ , be the x-y coordinates, linear surge speed (along  $\mathbf{E}_j$ ), and the orientation angle, respectively, of vessel  $j$  expressed in  $\text{RF}_i$ , where  $j = o, 1, 2, \dots, N$ . Vessel  $o$  denotes the evader, whereas vessels  $1, 2, \dots, N$ , denote the pursuers. The subscript  $k$  and  $SO(2)$  refer to the discrete-time index and the Special Orthogonal group in 2D, respectively. The decision-making mechanism is executed in a decentralized fashion where the RL algorithm is run locally by each pursuer  $j = 1, 2, \dots, N$ , with the goal of minimizing the magnitudes of the tracking errors  $e_k^{\ell,j} = \ell_k^j - \ell_k^o$ , for  $\ell \in \{x, y\}$ , and consequently  $e_k^{\theta,j} = \theta_k^j - \theta_k^o$ , as  $k \rightarrow \infty$ . This is done by computing two adjustment scalars  $u_k^{x,j}$  and  $u_k^{y,j}$  which are used to tune the surge velocity and heading of vessel  $j$ , such that

$${}^j v_{k+1}^j = {}^j v_k^j + \delta (u_k^{v,j} - u_{k-1}^{v,j}) \quad (1a)$$

$${}^j \theta_{k+1}^j = {}^j \theta_k^j + u_k^{\theta,j} \quad (1b)$$

where  $u_k^{v,j} = \sqrt{(u_k^{x,j})^2 + (u_k^{y,j})^2}$ ,  $u_k^{\theta,j} = \tan^{-1}(u_k^{x,j}/u_k^{y,j}) \in [-\pi/2, +\pi/2]$  rad, and  $\delta$  is a discount factor. Since the same RL algorithm is run by each vessel  $j$  and the output signals are computed with respect to  $\text{RF}_j$ , the reference frame and the vessel identifiers will no longer be explicitly indicated in the rest of the manuscript.




Figure 1: Earth and vessel coordinate systems

The corrective adjustments  $u_k^\ell$ ,  $\ell \in \{x, y\}$ , are set up as

$$u_k^\ell = \mathbf{K}_k^\ell \mathbf{E}_k^\ell \quad (2)$$

where  $\mathbf{E}_k^\ell = [e_k^\ell \ e_{k-1}^\ell \ e_{k-2}^\ell]^T \in \mathbb{R}^3$  is a window of the last three recorded tracking errors between the pursuer and the evader along the  $\ell$ -direction, while  $\mathbf{K}_k^\ell \in \mathbb{R}^{1 \times 3}$  represents the control gains which are computed in real-time by the RL algorithm.

One of the salient features of this setup is its flexibility to adapt to different problems of various complexities. For instance, for the application in hand, it was found sufficient to only include the last three tracking errors in  $\mathbf{E}_k^\ell$ . However, it is possible to add more error signals or remove some without having to redesign the control algorithm. The same remark is applicable for modeling the incremental difference in the surge speed and the orientation angle. In (1a) and (1b), such signals are shaped as  $v_{k+1} - v_k = \delta(u_k^v - u_{k-1}^v)$  and  $\theta_{k+1} - \theta_k = u_k^\theta$ , respectively. Changing them to some other forms, including  $(u_k^v, u_{k-1}^v, u_{k-2}^v, \dots)$ , and  $(u_k^\theta, u_{k-1}^\theta, \dots)$ , for example, does not require the restructure of the algorithm. This exhibits the flexibility inherited in the learning system where, using only few self-adapted control gains, a filtering behavior can be obtained during the online learning process.

## Temporal Difference Solution

For the RL framework to provide a solution for the pursuer-evader optimization problem, we need to derive the necessary temporal difference equations (i.e., Bellman equations) along with the required optimality conditions. The ultimate objective is to converge to the best policies  $u_k^{v*}$  and  $u_k^{\theta*}$  for each pursuer in order to intercept the evader ship. Hence, two convex objective functions are considered (each corresponds to the motion in one of the x-y directions)

$$U^\ell(\mathbf{E}_k^\ell, u_k^\ell) = \frac{1}{2} \left( \mathbf{E}_k^{\ell T} \mathbf{Q}^\ell \mathbf{E}_k^\ell + R^\ell (u_k^\ell)^2 \right)$$

where  $\ell \in \{x, y\}$ ,  $\mathbf{0} < \mathbf{Q}^\ell \in \mathbb{R}^{3 \times 3}$  (positive definite) and  $0 < R^\ell \in \mathbb{R}$ . The quality of the applied control strategies is assessed through a performance measure

$$J_k^\ell = \sum_{i=k}^{\infty} U^\ell(\mathbf{E}_i^\ell, u_i^\ell) \quad (3)$$

Taking advantage of the policy expression (2), the optimization process aims at minimizing the performance index  $J_k^\ell$  by optimizing the gain vector  $\mathbf{K}_k^\ell$ . This iterative process may lead the obtained policy to practically converge towards the optimal policy  $u_k^{\ell*} = \arg \min_{u_k^\ell} (J_k^\ell)$ , although it may not reach its theoretical value.

A common problem in RL is that  $J_k^\ell$  cannot be calculated. Motivated by the structure of the cost function  $U^\ell$ ,  $J_k^\ell$  is approximated by a value function  $V^\ell$  that is only dependent on two time-variant signals  $\mathbf{E}_k^\ell$  and  $u_k^\ell$ .

$$J_k^\ell \equiv V^\ell(\mathbf{E}_k^\ell, u_k^\ell) \quad (4)$$

From (3) and (4), the temporal difference equation is formed as

$$V^\ell(\mathbf{E}_k^\ell, u_k^\ell) = U^\ell(\mathbf{E}_k^\ell, u_k^\ell) + V^\ell(\mathbf{E}_{k+1}^\ell, u_{k+1}^\ell) \quad (5)$$

As with the cost function  $U^\ell$ , the value function is set as a quadratic function in the error vector  $\mathbf{E}_k^\ell$  and the control signal  $u_k^\ell$ , such that

$$V^\ell(\mathbf{E}_k^\ell, u_k^\ell) = \frac{1}{2} [\mathbf{E}_k^{\ell T} \quad u_k^{\ell T}] \mathbf{H}^\ell \begin{bmatrix} \mathbf{E}_k^\ell \\ u_k^\ell \end{bmatrix},$$

where  $\mathbf{H}^\ell$  is a symmetric positive definite matrix with the following block structure

$$\mathbf{H}^\ell = \begin{bmatrix} \mathbf{H}_{EE}^\ell & \mathbf{H}_{Eu}^\ell \\ \mathbf{H}_{uE}^\ell & \mathbf{H}_{uu}^\ell \end{bmatrix}$$

Note that the structure of the solving value function plays an important role in how a RL algorithm decides on the best strategies [14]. With the chosen quadratic structure, the pursuer-evader optimization problem is reduced to finding the optimal solving value functions, by optimizing  $\mathbf{H}^\ell$ , and consequently computing the optimal control gain  $\mathbf{K}_k^{\ell*}$  using Bellman optimality principles [11]. To this end, we can conclude that the optimal policy is

$$u_k^{\ell*} = \arg \min_{u_k^\ell} V^\ell(\mathbf{E}_k^\ell, u_k^\ell) = -(\mathbf{H}_{uu}^{\ell*})^{-1} \mathbf{H}_{uE}^{\ell*} \mathbf{E}_k^\ell \quad (6)$$

Relations (2) and (6) yield the optimal gain

$$\mathbf{K}^{\ell*} = -(\mathbf{H}_{uu}^{\ell*})^{-1} \mathbf{H}_{uE}^{\ell*}$$

Applying the optimal policy (6) in (5) yields the following Bellman optimality equation

$$V^{\ell*}(\mathbf{E}_k^\ell, u_k^{\ell*}) = U^{\ell*}(\mathbf{E}_k^\ell, u_k^{\ell*}) + V^{\ell*}(\mathbf{E}_{k+1}^\ell, u_{k+1}^{\ell*}) \quad (7)$$

## Online Reinforcement Learning Solution

In most real-world scenarios, it is impossible to find theoretical solutions for (6) and (7). To alleviate this problem, an online Value Iteration algorithm is employed herein to iteratively approximate such solutions.

### *Online Value Iteration*

Value Iteration is a two-step technique that is used to find the solutions of many temporal difference forms, such as the Bellman optimality equations (7). In the first step, the solving value function is updated as

$$V^{\ell(r+1)}(\mathbf{E}_k^\ell, u_k^\ell) = U^{\ell(r)}(\mathbf{E}_k^\ell, u_k^\ell) + V^{\ell(r)}(\mathbf{E}_{k+1}^\ell, u_{k+1}^\ell)$$

where  $r$  is the value-update index. The second step uses such update to suggest a new (improved) strategy, given by

$$u_k^{\ell(r+1)} = - \left[ (\mathbf{H}_{uu}^{\ell(r)})^{-1} \mathbf{H}_{uE}^{\ell(r)} \right]^{(r+1)} \mathbf{E}_k^\ell.$$

The algorithm cycles through this process until convergence is achieved. The Value Iteration technique is generally proven to converge. It yields a bounded sequence of non-decreasing solving value functions if the initial solving value function is selected to be positive definite [18]. The following development explains how the Value Iteration process approximates the solving value function  $V^\ell$  and the optimal strategy  $u^\ell$  using means of the adaptive actor-critics.

## Adaptive Actor-Critics Approximation

In this work, a critic neural network is applied to approximate the solving value function, whereas an actor neural network is designed to estimate the optimal strategy [19,20]. In the following, we will use the notation  $(\hat{\cdot})$  to denote an estimate of  $(\cdot)$ .

The approximation of the solving value function for each pursuing vessel is motivated by the following quadratic structure in the  $\ell$ -direction, where  $\ell \in \{x, y\}$ :

$$\hat{V}^\ell(\mathbf{E}_k^\ell, \hat{u}_k^\ell) = \frac{1}{2} \begin{bmatrix} \mathbf{E}_k^{\ell T} & \hat{u}_k^{\ell T} \end{bmatrix} \mathbf{W}^{[c]\ell} \begin{bmatrix} \mathbf{E}_k^\ell \\ \hat{u}_k^\ell \end{bmatrix},$$

with  $\mathbf{W}^{[c]\ell} \in \mathbb{R}^{4 \times 4}$  being the critic approximation weights matrix of the following block structure

$$\mathbf{W}^{[c]\ell} = \begin{bmatrix} \mathbf{W}_{EE}^{[c]\ell} & \mathbf{W}_{Eu}^{[c]\ell} \\ \mathbf{W}_{uE}^{[c]\ell} & \mathbf{W}_{uu}^{[c]\ell} \end{bmatrix}.$$

In a similar fashion, the optimized strategy is estimated by a form similar to that of (2) and (6),

$$\hat{u}_k^\ell = \mathbf{W}^{[a]\ell} \mathbf{E}_k^\ell,$$

where  $\mathbf{W}^{[a]\ell} \in \mathbb{R}^{1 \times 3}$  holds the actor approximation weights.

Let  $\varepsilon_k^{[c]\ell}$  and  $\varepsilon_k^{[a]\ell}$  be critic and actor weight approximation errors, respectively. Then,

$$\begin{aligned} \varepsilon_k^{[c]\ell} &= \frac{1}{2} \left[ \hat{V}^\ell(\mathbf{E}_k^\ell, \hat{u}_k^\ell) - \tilde{V}_k^\ell \right]^2 \\ \varepsilon_k^{[a]\ell} &= \frac{1}{2} \left[ \hat{u}_k^\ell - \tilde{u}_k^\ell \right]^2 \end{aligned}$$

where

$$\begin{aligned} \tilde{V}_k^\ell &= U^\ell(\mathbf{E}_k^\ell, \hat{u}_k^\ell) + \hat{V}^\ell(\mathbf{E}_{k+1}^\ell, \hat{u}_{k+1}^\ell) \\ \tilde{u}_k^\ell &= - \left[ (\mathbf{W}_{uu}^{[c]\ell})^{-1} \mathbf{W}_{uE}^{[c]\ell} \right] \mathbf{E}_k^\ell \end{aligned}$$

The update of the critic and actor weights is accomplished in real-time through a gradient descent approach, such that

$$\mathbf{W}^{[c]\ell(r+1)} = \mathbf{W}^{[c]\ell(r)} - \alpha_c \varepsilon_k^{[c]\ell} \begin{bmatrix} \mathbf{E}_k^\ell \\ \hat{u}_k^\ell \end{bmatrix} \begin{bmatrix} \mathbf{E}_k^{\ell T} & \hat{u}_k^{\ell T} \end{bmatrix} \quad (8)$$

$$\mathbf{W}^{[a]\ell(r+1)} = \mathbf{W}^{[a]\ell(r)} - \alpha_a \varepsilon_k^{[a]\ell} \mathbf{E}_k^\ell \quad (9)$$

where  $0 < \alpha_c, \alpha_a < 1$  are learning rate parameters and  $r$  is the index corresponding to the strategy update loop. The adaptive actor-critic implementation of the online Value Iteration process is detailed in Algorithm 1. The matrix magnitude used in the algorithm is defined as the maximum absolute value of the matrix elements. It is important to notice that the designed RL algorithm does not use any information about the dynamics of the pursuer or evader vessels. The strategies followed by each pursuer only depend on its own location measurements and those of the evader ship.

# Simulation Scenarios

## *Total::Perception Engine and Coordinate Transformations*

While the decision-making mechanism is realized through the RL technique listed in Algorithm 1, the vessels' dynamics and the evader ship behavior are simulated using Total::Perception™ (T::P), which is a state-of-the-art proprietary Intelligence, Surveillance and Reconnaissance (ISR) Systems Simulation Engine developed by Larus Technologies. It is capable of fusing large amounts of data from a multitude of sources, such as, active and passive sensing packages from GPS, Radar, and other sources. It is also able to learn complex relations among the tracked vessels to recognize any anomalous patterns in order to provide accurate information-processing strategies and continuously optimize the situational awareness mechanisms for decision support makers. Sea vessel objects in T::P are controlled through a variety of commands, such as the vessel's reference linear speed and orientation angle. The engine also enables chosen vessels to perform different missions like patrolling and interception. To assess the performance of the proposed solution, the heading angles and the linear speeds of the pursuer vessels are controlled by the RL approach described in Algorithm 1, while the evader's intelligence is provided by T::P.

Let  $\text{RF}_0$  be the (XYZ) Earth Centered Earth Fixed (ECEF) coordinates, as shown in Figure 1. The T::P engine operates with spherical coordinates in  $\text{RF}_0$ . If the spherical coordinates of a vector  $\mathbf{u}$  in  $\text{RF}_0$  are  $(\|\mathbf{u}\|, \theta, \gamma)$ , where  $\|\mathbf{u}\|$  stands for the Euclidean norm, and  $\theta$  and  $\gamma$  are the longitude and latitude angles, respectively, then the Cartesian coordinates of  $\mathbf{u}$  in  $\text{RF}_0$  are given by

$${}^0\mathbf{u} = \|\mathbf{u}\| \begin{bmatrix} \cos(\gamma) \cos(\theta) \\ \cos(\gamma) \sin(\theta) \\ \sin(\gamma) \end{bmatrix}$$

Converting this to its equivalent Cartesian coordinates  ${}^j\mathbf{u}$  in the (ENU) frame  $\text{RF}_j$  corresponding to vessel  $j$  is performed through the rotation matrix  ${}^j\mathbf{R}_0$ , where  ${}^j\mathbf{u} = {}^j\mathbf{R}_0 {}^0\mathbf{u}$  and

$${}^j\mathbf{R}_0 = \begin{bmatrix} -\sin(\theta) & \cos(\theta) & 0 \\ -\cos(\theta) \sin(\gamma) & -\sin(\theta) \sin(\gamma) & \cos(\gamma) \\ \cos(\theta) \cos(\gamma) & \sin(\theta) \cos(\gamma) & \sin(\gamma) \end{bmatrix}$$

To go back to  $\text{RF}_0$ , it is sufficient to multiply by the transpose of  ${}^j\mathbf{R}_0$  (a property of rotation matrices). In other words,  ${}^0\mathbf{u} = ({}^j\mathbf{R}_0)^T {}^j\mathbf{u}$ . For each pursuer  $j$ , the RL process generates the heading angle and surge speed variations or control decisions for that pursuer expressed in  $\text{RF}_j$ . They are transformed into their equivalent Cartesian coordinates in  $\text{RF}_0$  before passing them on to the T::P engine. The updated vessel poses generated by T::P are fed back to the RL algorithm after they are transformed back to their respective vessel frame  $\text{RF}_j$  coordinates. The T::P engine interacts with Google Earth KML tool to visualize the behavior of the vessels in real-time on Google Earth [4]. The simulation data flow is shown in Figure 2.

## *Test Results*

To demonstrate the adaptability of the decision-making mechanism to track IUU ships, three real-world scenarios are considered. They are realistic scenarios purposely designed by experts in the industry. They address typical situations where a suspicious IUU vessel is detected by radar and satellite monitoring systems. A resource allocation system identifies three coastguard vessels to participate in the pursuit. The command is then passed to the proposed RL-based algorithm to

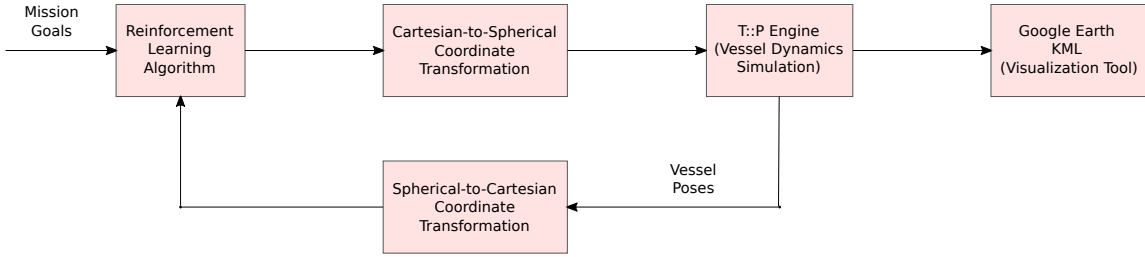


Figure 2: Simulation data flow

navigate the pursers in real-time to try to intercept the suspicious vessel as quickly as possible; preferably while it is still in action and within the Canadian marine borders. The simulation parameter values are listed in Table 1. The critic learning rate  $\alpha_c$  is set to a notably small value to compensate for the excessively large initial tracking errors which are in the range of hundreds of kilometers.

Table 1: Simulation parameters

$\delta$	$Q^\ell$	$R^\ell$	$\alpha_a$	$\alpha_c$	$\Delta$	$L$	$N_t$
0.99	$10^{-4} I_{3 \times 3}$	0.01	0.01	$10^{-6}$	$10^{-8}$	20	6,000

In Scenario 1, the three pursers (labeled 1, 2, and 3) are given a command to intercept an IUU fishing incident. They are initially located 556, 367, and 289 km, respectively, apart from the IUU ship. In this scenario, the T::P engine, which is in charge of controlling the evading strategy of the evader ship, commands the ship to evade with a constant speed towards the northeast direction. Figure 3 shows the trajectories of the pursuers until successfully intercepting the evader. The first vessel to reach the evader was pursuer 2, after 105 min, followed by pursuers 1 and 3, which took 140 min and 198 min, respectively. It is interesting to notice that the first interception was not accomplished by the vessel that was initially closest to the evader. This is due to the IUU ship’s evading maneuvers decided by the T::P engine.

In Scenario 2, the mission goals are set differently than in the previous scenario. Purser vessel 2 is now commanded to intercept the IUU ship, while vessels 1 and 3 are commanded to form a surveillance zone around the IUU ship by following it within a fixed pre-defined distance. This strategy can be useful to block the way in front of the evader to escape local waters, for instance. The vessels are initially located 744, 378, and 289 km away from the IUU ship. The evader is commanded by T::P to follow the same evading strategy as in the first scenario. The outcome of the pursuit is demonstrated in Figure 4. The pursuer manage to intercept the evader in 110 min, while the other two vessels draw the region which they prevent the evader to escape. This simulation shows the flexibility of the proposed RL technique in taking various forms of missions without having to modify the algorithm structure.

Finally, a third scenario is considered to test the robustness of the adaptive learning mechanism. To that end, Scenario 2 is repeated, except that this time, the vessel originally set to intercept the IUU ship (vessel 2) experiences a malfunction (at time 117 min) after intercepting the evader in 110 min, as shown in Figure 5. At this moment, vessel 1 is commanded to change its mission from surveilling to intercepting the evader (Figure 6). It took the new pursuer 213 min to lock down on the evader from the moment it switched its mission from surveillance to interception. This is another example on how the proposed decision-making mechanism can easily adapt to mission changes in real-time with no extra overhead to the operator.




Figure 3: Intercepting the IUU ship (Scenario 1).

## Conclusion

The article presented an online RL decision-making mechanism to tackle the negative impacts of IUU incidents along the Canadian coastlines. This is accomplished in real-time and in a completely decentralized fashion where the only feedback needed for each pursuing vessel is its own pose and that of the evader. No information is required about the dynamic models of the pursuing vessels or the evader ship. The adaptive learning solution is implemented using a Value Iteration process and means of neural network approximators. The solution is shown to be flexible and interactive to dynamic online variations to the pursuer missions. The proposed tool is being evaluated for the possible commercialization by Larus Technologies.

## References

- [1] FAO. (2016) The state of world fisheries and aquaculture: contributing to food security and nutrition for all. [Online]. Available: [www.fao.org/3/a-i5555e.pdf](http://www.fao.org/3/a-i5555e.pdf)
- [2] R. Abielmona, “Tackling big data in maritime domain awareness,” *Vanguard*, pp. 42–43, 2013.
- [3] P. J. B. Sánchez, M. Papaelias, and F. P. G. Márquez, “Autonomous underwater vehicles: Instrumentation and measurements,” *IEEE Instrumentation Measurement Magazine*, vol. 23, no. 2, pp. 105–114, 2020.
- [4] Larus Technologies. (2020) Larus technologies products (total::insight and total::perception). [Online]. Available: <https://www.larus.com/products/>




Figure 4: Intercepting the IUU ship (Scenario 2).

- [5] T. Akinbulire, H. Schwartz, R. Falcon, and R. Abielmona, “A reinforcement learning approach to tackle illegal, unreported and unregulated fishing,” in *2017 IEEE Symposium Series on Computational Intelligence (SSCI)*, 2017, pp. 1–8.
- [6] D. B. Fogel, “Machine intelligence,” *IEEE Instrumentation Measurement Magazine*, vol. 9, no. 3, pp. 12–16, 2006.
- [7] S. Sun, Y. Yin, X. Wang, and D. Xu, “Robust visual detection and tracking strategies for autonomous aerial refueling of uavs,” *IEEE Transactions on Instrumentation and Measurement*, vol. 68, no. 12, pp. 4640–4652, 2019.
- [8] M. Gomez Plaza, T. Martinez-Marin, S. Sanchez Prieto, and D. Meziat Luna, “Integration of cell-mapping and reinforcement-learning techniques for motion planning of car-like robots,” *IEEE Transactions on Instrumentation and Measurement*, vol. 58, no. 9, pp. 3094–3103, 2009.
- [9] P. Chaphekarande and D. Deshpande, “Machine learning based brain mri estimation method,” in *2019 2nd International Conference on Intelligent Computing, Instrumentation and Control Technologies (ICICICT)*, vol. 1, 2019, pp. 1423–1430.
- [10] S. A. Mohammed, S. Shirmohammadi, and S. Altamimi, “A multimodal deep learning-based distributed network latency measurement system,” *IEEE Transactions on Instrumentation and Measurement*, vol. 69, no. 5, pp. 2487–2494, 2020.
- [11] F. Lewis, D. Vrabie, and V. Syrmos, *Optimal Control*, 3rd ed. New York, USA: John Wiley, 2012.

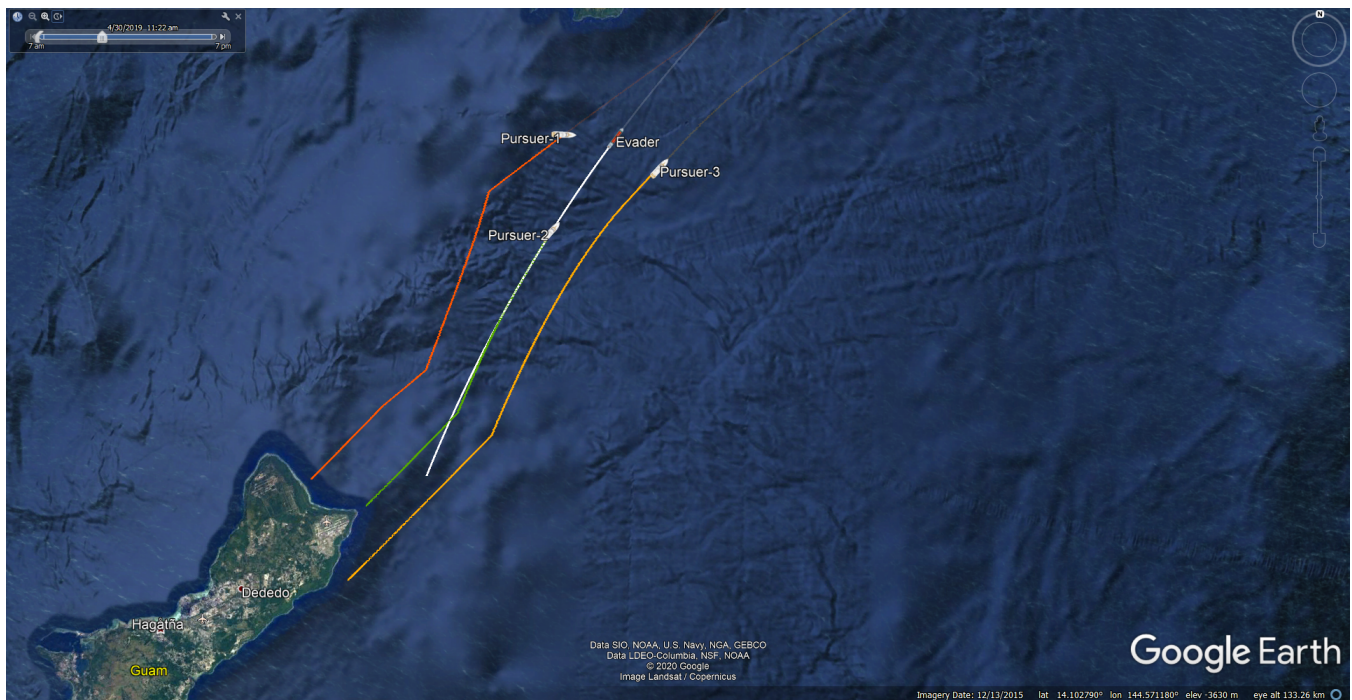


Figure 5: Vessel 2 malfunctions during the pursuit. Vessel 1 is commanded to change its mission from surveillance to interception of the evader (Scenario 3).

- [12] M. Abouheaf, W. Gueaieb, and D. Spinello, "Online multi-objective model-independent adaptive tracking mechanism for dynamical systems," *Robotics*, vol. 8, no. 4, p. 82, Sep 2019. [Online]. Available: <http://dx.doi.org/10.3390/robotics8040082>
- [13] M. Abouheaf and W. Gueaieb, "Multi-agent reinforcement learning approach based on reduced value function approximations," in *IEEE International Symposium on Robotics and Intelligent Sensors (IRIS)*, Oct. 2017, pp. 111–116.
- [14] M. Abouheaf, F. Lewis, M. Mahmoud, and D. Mikulski, "Discrete-time dynamic graphical games: Model-free reinforcement learning solution," *Control Theory and Technology*, vol. 13, no. 1, pp. 55–69, 2015.
- [15] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, 2nd ed., ser. Second. Massachusetts: MIT Press, 1998.
- [16] R. GhasemAghaei, M. A. Rahman, W. Gueaieb, and A. El Saddik, "Ant colony-based reinforcement learning algorithm for routing in wireless sensor networks," in *2007 IEEE Instrumentation Measurement Technology Conference IMTC 2007*, 2007, pp. 1–6.
- [17] M. Abouheaf, N. Q. Mailhot, W. Gueaieb, and D. Spinello, "Guidance mechanism for flexible-wing aircraft using measurement-interfaced machine-learning platform," *IEEE Transactions on Instrumentation and Measurement*, vol. 69, no. 7, pp. 4637–4648, 2020.
- [18] M. Abouheaf, W. Gueaieb, and F. Lewis, "Model-free gradient-based adaptive learning controller for an unmanned flexible wing aircraft," *Robotics*, vol. 7, no. 4, p. 66, 2018.




Figure 6: Pursuer vessel 1 intercepts the evader while vessel 3 surveils it from a distance (Scenario 3).

- [19] D. Bertsekas and J. Tsitsiklis, *Neuro-Dynamic Programming*, 1st ed. Massachusetts: Athena Scientific, 1996.
- [20] D. Prokhorov and D. Wunsch, "Adaptive critic designs," *IEEE Transactions on Neural Networks*, vol. 8, no. 5, pp. 997–1007, Sep. 1997.

---

**Algorithm 1** Adaptive Actor-Critic Implementation

---

**Input:**

- Positive definite  $\mathbf{Q}^\ell$  and  $R^\ell$ ,  $\ell \in \{x, y\}$
- Initial tracking error vectors  $\mathbf{E}_k^\ell$  and strategies  $\hat{u}_k^\ell$
- Convergence error threshold  $\Delta$
- Learning parameters  $\alpha_a$  and  $\alpha_c$
- Width  $L$  of a moving time window
- Maximum number of learning iterations  $N_t$

**Output:**

- Tuned neural network weights  $\mathbf{W}^{[c]\ell(*)}$  and  $\mathbf{W}^{[a]\ell(*)}$

```
1: stop  $\leftarrow$  false
2: while stop = false do
3:    $r \leftarrow 0$ 
4:    $\mathbf{W}^{[a]\ell(r+1)}, \mathbf{W}^{[c]\ell(r+1)} \leftarrow$  random values
5:   actor_converged  $\leftarrow$  false
6:   critic_converged  $\leftarrow$  false
7:   while  $r < N_t$  and stop = false do
8:      $r \leftarrow r + 1$ 
9:     Calculate  $U^\ell(\mathbf{E}_k^{\ell(r)}, u_k^{\ell(r)})$ 
10:    Compute  $\hat{V}^\ell(\mathbf{E}_{k+1}^{\ell(r)}, \hat{u}_{k+1}^{\ell(r)})$  and  $\hat{u}_{k+1}^{\ell(r)}$ 
11:    Calculate  $\varepsilon_k^{[c]\ell(r)}$  and  $\varepsilon_k^{[a]\ell(r)}$ 
12:    Update critic and actor weights ▷ using (8), (9)
13:    if  $r > L$  then
14:      if  $\|\mathbf{W}^{[c]\ell(r+1-j)} - \mathbf{W}^{[c]\ell(r-j)}\| \leq \Delta, \forall j \in \{0, 1, \dots, L\}$ , then
15:         $\mathbf{W}^{[c]\ell(*)} \leftarrow \mathbf{W}^{[c]\ell(r+1)}$ 
16:        critic_converged  $\leftarrow$  true
17:      end if
18:      if  $\|\mathbf{W}^{[a]\ell(r+1-j)} - \mathbf{W}^{[a]\ell(r-j)}\| \leq \Delta, \forall j \in \{0, 1, \dots, L\}$ , then
19:         $\mathbf{W}^{[a]\ell(*)} \leftarrow \mathbf{W}^{[a]\ell(r+1)}$ 
20:        actor_converged  $\leftarrow$  true
21:      end if
22:      if critic_converged and actor_converged then
23:        stop  $\leftarrow$  true
24:      end if
25:    end if
26:  end while
27: end while
28: return Tuned weights  $\mathbf{W}^{[c]\ell(*)}$  and  $\mathbf{W}^{[a]\ell(*)}$ 
```

---