Using Reinforcement Learning to Control Hydrofoils

Jurre van der Schaaf
University of Twente
P.O. Box 217, 7500AE Enschede
The Netherlands
j.j.vanderschaaf@student.utwente.nl

ABSTRACT

Hydrofoils on vessels are used to lift a vessel's hull above the water surface to reduce drag and increase efficiency and top speed. By changing the angles of the hydrofoils, more or less lift can be created. Controlling the different foil angles makes it possible to regulate the height and stability of the ship. However, these control systems can not always adapt to a new situation when the environment changes or when it is placed in a new environment, leading to a malfunctioning system. This research investigates the approach of a control system based on reinforcement learning. This is done by implementing Deep Q-Learning on a simulation of a hydrofoil boat. For this, different configurations have been tested and compared against a PID controller. These configuration differ in actions sets, reward functions and number of simultaneous agents. Examination of plots, as well as the standard deviation of the roll and pitch angles of the vessel, showed that using a dedicated agent per foil with a small action set with a reward based on height, the performance is comparable to a PID controller and considered stable. Other experiments with larger actions sets, one agent controlling two foils simultaneously or rewarding by roll angle showed that the system was in those cases not able to perform its task properly.

Keywords

Reinforcement learning, Hydrofoils, Hydrofoil control, Deep Q-Learning

1. INTRODUCTION

In maritime industries, hydrofoils are used more and more to reduce the amount of drag while sailing a vessel and subsequently to increase the maximum speed [5]. Hydrofoils are shaped like an airplane wing and mounted at the bottom side of a vessel. Due to their shape, the water flows with a different speed at the top and bottom side of the foil [6]. This causes a difference in pressure and thus an upward force is realized. The angle of attack, the angle under which the foil moves trough the water relative to the horizontal, has impact on the amount of lift. The upward force lifts the hull out of the water, which reduces the amount of drag.

36th Twente Student Conference on IT Febr. 4th, 2022, Enschede, The Netherlands.

There are different types of hydrofoils, of which one are fully submerged hydrofoils. This type of foils do not stabilize the vessel inherently like some other foil types do. An active control system has to take care of this stabilization task.

One application of hydrofoils is in solar boat racing, where teams build a solar powered vessel and race against each other[1], of which Solar Boat Twente[2] is one of the teams. These boats are relatively small and light. The boat of Solar Boat Twente (see Figure 1) has three fully-submerged hydrofoils, of which one is placed at the stern of the vessel. The other two are placed amidships, one at port side and one starboard side.

By making the angle of attack adjustable by means of an electromechanical system, the upward force can be regulated wile sailing. When a boat is equipped with three hydrofoils like the vessel of Solar Boat Twente, the stability of the ship can be controlled by changing the angle of attack of each hydrofoil. Both the rotation around the longitudinal axis (roll) and rotation around the lateral axis (pitch) can be controlled by creating more or less lift at a specific foil. A digital control system can be set up to adjust angles of attack based on sensor inputs like height measurements (height above the water surface) and gyroscopic forces. A proportional–integral–derivative (PID) controller is a possible control system to execute this task [3].

In dynamic environments or in new environments, such as changing wave heights, wavelengths, wind speed, or sailing speed, not all controllers are capable of keeping the system work correctly in the new situation. One solution can be to change parameters of the system to work in the new situation. This is however a difficult and time consuming task. As an alternative solution, this research proposes the use of reinforcement learning to control the angles of attack of hydrofoils. Because reinforcement learning learns from the feedback it gets from its environment [11], it may be capable of adapting to the changed characteristics of the environment.

1.1 Research question

This research investigates the basic application of reinforcement learning on controlling hydrofoils in a static environment, as a first step towards the goal of adapting in dynamic environments. The main research question is:

- RQ1: How can reinforcement learning be used to keep a boat equipped with hydrofoils lifted and stable above the waterline?

[1]https://solarsportone.org
[2]https://solarboattwente.nl

To answer this main research question, the following subquestions are defined:

- RQ1.1: What reward can be constructed to succeed in training a model?
- RQ1.2: How stable is the vessel when the reinforcement learning algorithm is used to control the hydrofoils?

These subquestions are answered by implementing several reward functions on a simulation. The movements of the simulated vessel are tracked and analyzed. This information is used to evaluate the performance in terms of height control and stability.

In next section, related work on hydrofoils and their control systems are reviewed. In section 3 the creation of the a simulation and the use of reinforcement learning is described. After this, the experimental setup to analyze the performance and the results of these experiments are given in section 4. Lastly, the paper is concluded in section 5.

2. RELATED WORK

In this section, some literature related to the topic is discussed. Research has been done in the past about hydrofoils [8] and their control systems [3, 5, 7].

Research about the hydrodynamics of hydrofoils, such as Ni et al., 2021 [8] describing the performance of a hydrofoil for different angles of attack, add to the general understanding of the way hydrofoils work and the hydrodynamic parameters.

The research of Bai and Kim [3] shows a comparison between three possible algorithms tested on a fully-submerged hydrofoil vessel: proportional– integral–derivative (PID), linear quadratic regulator (LQR) and sliding mode controls. They found the PID controller, which is considered least complex to implement, was most sensitive to large waves, but nevertheless had a "good overall performance". Sliding mode control was found to be the best, given accurate measurements of wave disturbances.

Recently, the research of Bencatel et al., 2021 [4] was published which describes a foiling control system for catamaran sailboats competing in the America's Cup. The paper shows the application of hydrofoil control systems, although the rules of the America's Cup do not allow automatic controlling systems during the race. The system is used to train sailors.

No research has been found specifically focused on controlling hydrofoils with machine learning or reinforcement learning particularly. Some recent research on machine learning combined with hydrofoils is restricted to analytics of the mechanical characteristics and hydrodynamics of the foils, but not on controlling the angle of attack.

Porter and Khaki-Sedigh [9] showed in 1988 that tunable and adaptive digital set-point tracking PID controllers can be used in changing environments, which can be seen as another solution for the problem described in section 1. This research included 10% noise which may not be representative for real-world changing environments such as going from a small lake without any waves to a big lake with short waves or to a sea with large wavelengths.

3. APPROACH

This section describes the simulation of a hydrofoil vessel and the training setup of a reinforcement learning system.

Figure 1. A picture of the vessel "Echo" of Solar Boat Twente, having 3 hydrofoils attached underneath. Photo by Solar Boat Twente.

Table 1. List of geometric values used during the creation of the simulation

Length	600 cm
Width	160 cm
Mass (including driver)	170 kg
Strut height	120 cm
Smallest foil angle	0 °
Largest foil angle	7 °
Surface area port and starboard foil	450 cm^2
Surface area stern foil	600 cm^2
Distance between center and port and starboard foil	70 cm
Distance between center and stern foil	300 cm

3.1 Simulation

A simulation of a hydrofoil-equipped vessel was created of which most characteristics are based on the boat of Solar Boat Twente (Figure 1). The simulation has three fully submerged hydrofoils, each foil can be adjusted separately. The current height of all three struts is kept in memory and changed based on calculations which include the angle of attack α. The geometric values can be found in Table 1. A constant speed of 7 m/s for all calculations is assumed. The update frequency is 5Hz, thus every step of the simulation simulates 200 ms.

3.1.1 Lift

The lift created by a foil [1] can be calculated using

$$F_L = \frac{1}{2} C_L \times \rho \times v^2 \times A$$

where

C_L = lift coefficient,
ρ = density of the medium [kg/m^3],
v = velocity through the medium [m/s],
A = surface area of the foil [m^3].

In this simulation, the medium taken is seawater and thus $\rho = 1025 kg/m^3$ [2].

C_L is a coefficient which depends on the exact shape of the hydrofoil and the angle of attack. For this simulation, a NACA0012 shape was chosen, which is one of the many possible shapes a foil can have. To compose the correct lift coefficients, a list has been created using the program XFOIL[3]. This program can generate lift coefficients, among other things, given a specific foil shape, angle of attack, Reynolds number and Mach number. The latter two have been taken as a constant (see Table 2). A list of C_L values has been composed for $\alpha \in \{0.0, 0.1, ..., 6.9, 7.0\}$.

The mass of the vessel with n number of foils is assumed to be evenly distributed over all hydrofoils. On each foil, we have a downward force due to gravity of $\frac{1}{n} m \times 9.81$.

[3]https://web.mit.edu/drela/Public/web/xfoil/

Table 2. Parameters used in XFOIL to construct a list of C_L values.

Reynolds number	3.4e−3
Mach number	395,000
Foil shape type	NACA0012

The netto force F_{netto} is then defined as the difference of upward force F_L and downward force by

$$F_{netto} = F_L - \frac{m \times 9.81}{n}$$

where

F_L = lift [N],
m = total mass of the vessel [kg],
n = number of foils attached to the vessel.

Since $s = \frac{1}{2}at^2$ (integral of the product of acceleration and time) and $a = F/m$ (Newton's Second Law), we can write the change in height (effectively a distance) at each foil position as

$$\Delta h = \frac{1}{2}\frac{F_{netto}}{(\frac{m}{n})}t^2$$

where

F_{netto} = netto force [N],
m = total mass of the vessel [kg],
n = number of foils attached to the vessel,
t = time step, for how long the current angle of attach is kept [s].

At all foil positions, the current height above the water of the vessel at that point is kept in the memory of the simulation. These can never be negative, as the buoyancy of the boat is such large that it cannot go under water due to any downward force of the foils. The foils are mounted using struts of which the length is 120 cm, thus this is the maximum height above the water.

3.1.2 Orientation

To evaluate the current orientation of the vessel, roll (rotation around the longitudinal axis) and pitch (rotation around the lateral axis) are calculated using trigonometry. Roll is determined by

$$\phi_{roll} = \tan(\frac{\frac{1}{2}(h_{starboard} + h_{port})}{d_1})$$

where

$h_{starboard}$ = current height of hydrofoil at starboard [cm],
h_{port} = current height of hydrofoil at port [cm],
d_1 = transverse distance from center of gravity to port and starboard hydrofoils [cm].

In the simulation, $d_1 = 70$ cm. Pitch is calculated with

$$\phi_{pitch} = \tan(\frac{\frac{1}{2}(\frac{1}{2}(h_{starboard} + h_{port}) - h_{stern})}{d_2})$$

where

$h_{starboard}$ = current height of hydrofoil at starboard [cm],
h_{port} = current height of hydrofoil at port [cm],
h_{stern} = current height of hydrofoil at stern [cm],
d_2 = longitudinal distance from center of gravity to stern hydrofoil [cm].

Here, $d_2 = 300$ cm.

3.1.3 Noise

In a real-world scenario, height measurements using sensors are not perfect. To simulate this, a random integer $i \in \{-3, ..., 3\}$ (cm) is added where a height measurement is taken place. However, the real heights (without noise) is used for all computations inside the simulation, for instance the calculation of ϕ_{roll} and ϕ_{pitch}.

3.1.4 Available values

The simulation outputs a state which consist of the following values:

- h_{port} (includes noise)
- $h_{starboard}$ (includes noise)
- h_{stern} (includes noise)
- ϕ_{roll}
- ϕ_{pitch}
- α_{port}
- $\alpha_{starboard}$
- α_{stern}
- $h_{average}$ (does not include noise)
- $h_{average,noisy}$ (includes noise)

3.1.5 Differences with real-world

The simulation is a simplified situation of a real-world environment, based on theoretical formulas. The simulation only includes some arbitrary measurement noise (see section 3.1.3) but does not include for instance wind, waves, mass of inertia and drag.

3.2 Reinforcement Learning

Reinforcement Learning works by collecting experiences and receiving rewards. Starting at a particular situation s, every action a the system takes will result in a new situation s', for which it receives a reward r. Every tuple (s, a, s', r) is saved into the *replay memory*.

The Deep Q-Learning algorithm is chosen to determine the best action to take. No other reinforcement learning algorithms has been evaluated in this research.

Deep Q-Learning [10] works with continuous states and a discrete set of actions. However, the choice of an angle of attack (α) between a minimum and maximum angle, is continuous. As the mechanical properties of hydrofoils do only allow for a certain significance, the continuous range can already been made discrete by allowing a maximum of one decimal. Thus $\{0.0, 0.1, ..., 6.9, 7.0\}$. Another mechanical constraint leads to an even smaller action set: within a time frame of 200 ms, α can be altered by a maximum of 2.0°. Because of this, the largest possible action set is $\{-2.0, -1.9, ..., +1.9, +2.0\}$. This last set defines the change in angle of attack, not the absolute angle of attack.

Deep Q-Learning makes use of a neural network. In this setup, three linear layers has been set, each with a size of 256 neurons.

The replay memory is used to store all experiences. At each learning step, a random batch of 128 samples is taken from the memory, which are used to update the network. The network is trained using the Adam optimizer (with a learning rate of 0.003) and mean square error as a loss function.

Every instance of the algorithm we call an agent. One or multiple agents can be used in parallel, although they don't work together. Each agent keeps track of it's own replay memory. The content of the states in the replay memory is dependend on the experiment.

The epsilon-greedy strategy uses a value ϵ that determines the probability to choose a random action from the action set, rather than choosing the best action according to the outcome of the neural network. This way, the algorithm will explore new situations. The value is decreased every learning step. For this research, an starting value of $\epsilon = 1$ is used, decreasing with 3.3e−5, until a minimum of $\epsilon = 0.2$ has been reached.

4. EXPERIMENTS

This section describes the experiments executed on the simulation and their results. The goal of the experiments is to validate the performance of the trained models created while training. During the experiments, it is the goal to lift the vessel out of the water and keep it stable. The goal height is set to 50 cm. The closer ϕ_{roll} and ϕ_{pitch} approach zero and the lower the standard deviation, the more stable the vessel is.

4.1 Experimental setup

In the experiments, each training runs 50,000 steps, simulating 167 minutes. After every step, the network is updated as described in section 3.2. After every 50 steps, the simulation is reset, thus all angles of attack are set to zero, as well as the height at each foil. This is done to let the system discover different situations, also when it came to a stable situation without any major disturbances. Different training setups are tried. There are four different properties that have varying values in the experiments, these are: reward function, number of agents and action set foil angle range. Table 3 shows each parameter for all experiments. These are explained in more detail in respectively sections 4.1.1, 4.1.2 and 4.1.3.

4.1.1 Reward functions

Different reward functions have been defined, with a slight difference. In the formula

$$r = \begin{cases} h, & \text{if } h < 25 \\ 25 - (h - 80), & \text{if } h > 80 \\ 100 + \frac{-4}{100}(h-50)^2 + \frac{-5}{1000}(h-50), & \text{otherwise} \end{cases} \tag{1}$$

the reward is based on h solely. h is the height at the foil the agent is controlling. For instance for the agent controlling the stern foil, $h = h_{stern}$. As $0 \leq h \leq 120$, the range of this function is $-15 \leq r \leq 100$. This function gives a maximum value when $h = 50$.

A second possible reward is based on both both height (h) and roll angle (ϕ_{roll}), defined as

$$r = \begin{cases} h, & \text{if } h < 30 \\ 25 - (h - 80), & \text{if } h > 70 \\ 125 - 10 \times min(5, |\phi_{roll}|) & \text{otherwise} \end{cases} \tag{2}$$

The function has a range of $-15 \leq r \leq 125$, reaching its maximum when $30 \leq h \leq 70$ cm and $\phi_{roll} = 0$. This means that to receive a maximum reward, the agent is free to let the boat sail between 30 cm and 70 cm, as long as the roll angle is close to 0.

The third reward function is

$$r = \begin{cases} h, & \text{if } h < 25 \\ 25 - (h - 80), & \text{if } h > 80 \\ 100 + \frac{-4}{100}(h-50)^2 + \frac{-5}{1000}(h-50) + p, & \text{otherwise} \end{cases} \tag{3}$$

with $p = 30 - 10 \times min(3, |\phi_{roll}|)$.

The idea is that the height is always affecting the score, just like Formula 1, however when $30 \leq h \leq 70$ the score is increased by a score based on the roll angle. The value is maximized when $h = 50$ cm and $\phi_{roll} = 0$.

4.1.2 Agents

Experiments with two or three agents are conducted. In the cases where three agents are used, the agent's action affects the angle of attack of one specific foil. When two agents are used, one agent is used to control the stern foil, and the other is used to control both the port foil as well as the starboard foil. As a consequence, the number of actions for this agent grows to n^2, where n is the number of actions per foil. The number of agents and action set per foil for each experiment are specified in Table 3.

The state s is dependent on the amount of agents, which foil(s) the agent controls and which reward is used. s is a subset of the values mentioned in section 3.1.4. It includes the heights of the foil(s) it controls, the angles of the foil(s) it controls and, if included in the reward function, the roll angle of the vessel.

4.1.3 Action set and foil angle range

As stated in section 3.2, there's already a limited amount of actions possible. However, in the experiments different subsets of the action set and foil angle range are tried. In some experiments, the size of the action set is smaller, such that fewer explorations are needed to build up a neural network.

4.1.4 PID control

As a PID control system is often used and has a "good overall performance" [3], an experiment using a PID controller was conducted with the simulation. The experiments A to E can not only be compared against each other, but also against the PID experiment. For the port foil and starboard foil, $K_p = 0.15$, $K_i = 0.1$ and $K_d = 0.06$ are used. For the stern foil, $K_p = 0.25$, $K_i = 0.1$ and $K_d = 0.06$. These values have been found to be best by trial and error with different values. To be able to compare the performance against the other experiments, this PID controller is also restricted to change α by a maximum of 2.0 to one decimal place.

4.2 Results

After training, the performance is evaluated using a run of 300 steps. This simulates 60 seconds of sailing. ϵ is set to 0, to prevent the system from taking random actions. At every step, $h_{average}$, ϕ_{roll} and ϕ_{pitch} are logged. Note that these do not include noise as described in section 3.1.3, as to evaluate the actual height, roll and pitch values of the vessel, and not distorted values.

The results of all experiments can be found in Table 4. A graph is available for each experiment. In all graphs, the green line represents the average height of the vessel $h_{average}$ without noise. The red and pink line show respectively ϕ_{roll} and ϕ_{pitch}.

4.2.1 PID control

The PID controller, shown in Figure 2, is able to keep the height on an average of 50 cm, which is equal to the goal height. Besides, the maximum roll angle is only 1.6° with standard deviation of 0.4°. It needs quite some steps (41) to reach a height of 40 cm, which corresponds to 8.2 s.

4.2.2 Experiment A

In experiment A, three agents are trained with a large angle range and an action set that allows the agent to control foils per 0.5°. The reward is only based on the height of each foil. The performance is shown in Figure 3.

Table 3. Properties of experiments A to E. The experiments differ in number of agents, action set, foil range and reward function.

Name	Number of agents	Action set	Foil angle range	Reward function
Experiment A	3	$\{-1.0, -0.5, ..., 0.5, 1.0\}$	$\{0.0, ..., 7.0\}$	Formula 1
Experiment B	3	$\{-1.0, -0.9, ..., 0.9, 1.0\}$	$\{0.0, ..., 7.0\}$	Formula 1
Experiment C	3	$\{-0.2, -0.1, 0, 0.1, 0.2\}$	$\{2.5, ..., 5.5\}$	Formula 1
Experiment D	2	$\{-0.2, -0.1, 0, 0.1, 0.2\}$	$\{2.5, ..., 5.5\}$	Formula 1 and Formula 2
Experiment E	3	$\{-0.2, -0.1, 0, 0.1, 0.2\}$	$\{2.5, ..., 5.5\}$	Formula 1 and Formula 3

Table 4. Results of the experiments A to E, and the PID control.

Experiment	**A**	**B**	**C**	**D**	**E**	**PID**
Maximum $\lvert\phi_{roll}\rvert$ (°)	6.2	16.8	1.6	66.2	25.7	1.6
Mean $\lvert\phi_{roll}\rvert$ (°)	2.5	6.2	0.4	61.3	10.0	0.5
Standard deviation $\lvert\phi_{roll}\rvert$ (°)	1.5	4.6	0.3	15.9	6.9	0.4
Steps to reach $h = 40$ cm	5	6	13	14	13	41
Minimum height[4] (cm)	43	44	42	34	45	41
Average height[4] (cm)	50	53	51	45	62	50
Maximum height[4] (cm)	58	63	56	51	72	57
Standard deviation height[4] (cm)	2.2	3.5	2.3	2.7	4.6	2.6

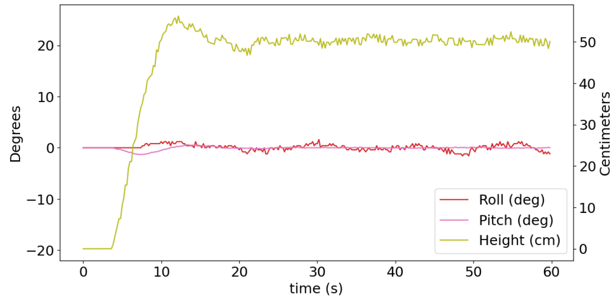

Figure 2. Performance with a PID controller used as a reference to compare experiments A to E against.

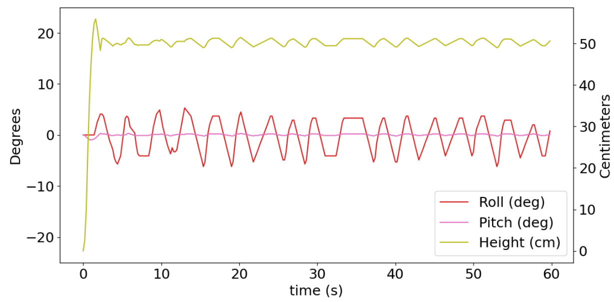

Figure 3. Experiment A, trained three agents with a reward function based on height, and action set $\{-1.0, -0.5, ..., 0.5, 1.0\}$

Although the height is comparable with the height of the PID controller, the roll angle has a maximum of 6.2°and a standard deviation of 1.5°. The peak corresponds to a difference in height between the port side and starboard side of 15 cm.

4.2.3 Experiment B

Experiment B has a larger action set compared to experiment A, allowing a more precise control in the same range. It uses three agents with a reward based on height only. The result can be seen in Figure 4. Although the average height is close to the target height of 50 cm, it can be seen in from the graph and the standard deviation (4.6°) that it fluctuates more compared to the height in experiment

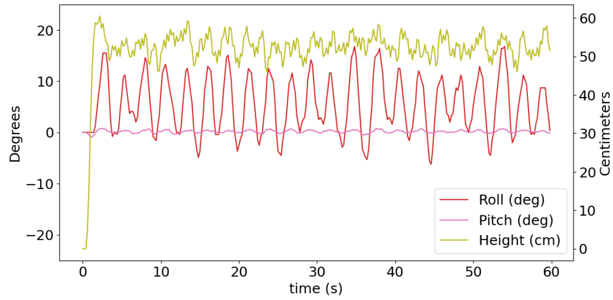

Figure 4. Experiment B, trained three agents with a reward function based on height, and action set $\{-1.0, -0.9, ..., 0.9, 1.0\}$

A.

4.2.4 Experiment C

In experiment C a smaller action set compared to experiment A and B has been used. Furthermore, the minimum and maximum foil angles have been restricted, meaning less options to set the angle to. In this experiment three agents are trained with a reward based on height only. The result is given in Figure 5. The height is kept close to 50 cm, with a standard deviation of 2.3 cm. It takes more steps to reach a height of 40 cm compared to experiment A and B, however, resulting in a roll angle with a maximum of 1.6°. This roll angle is equal to the PID controller, meaning this experiment matches the performance of the PID controller.

4.2.5 Experiment D

As an experiment to control the foil based on the roll angle, as this is an important metric in the performance, experiment D has been introduced. In this experiment only two agents are trained. One agent is controlling both port and starboard foil. The reward for this agent is based on both height and roll (Formula 2). The other agent is controlling the stern foil only, with a reward based on height only. The result is shown in Figure 6. Immediately after the run has started, the roll angle drops to -66.2°, which means the port side raise to its maximum value of 120 cm, whereas the starboard side kept at 0 cm. The agents

[4]Measured after $h = 40$ has been reached once.

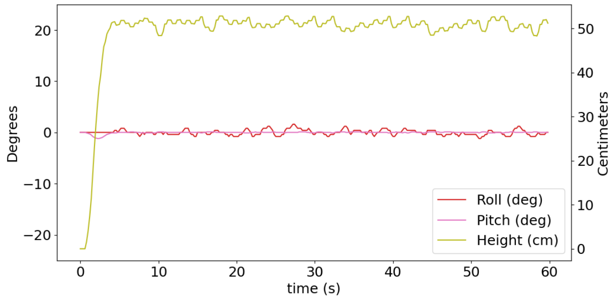

Figure 5. Experiment C, trained three agents with a reward function based on height, and action set $\{-0.2, -0.1, ..., 0.1, 0.2\}$

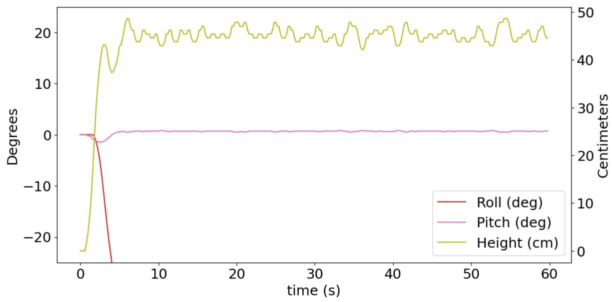

Figure 6. Experiment D, trained two agents with a reward function based on height and roll, and action set $\{-1.0, -0.9, ..., 0.9, 1.0\}$

seemed to over fit on the reward function, such that it would receive a constant reward by raising one side to it's maximum value and the other to the minimum value. After repeating the experiment 4 times, the result did not change significantly.

4.2.6 Experiment E

The last experiment conducted is an experiment where again the reward for the agents at port and starboard side are based on height and roll. This time three agents are trained. Figure 7 shows the performance of this experiment. Compared to experiment D, which also included roll angle in the reward function, performance is better due to the fact that both foil at port and starboard side rise to an average height of around 60 cm. Heights are fluctuating between 45 cm and 72 cm. The roll angle peaks to 25.7°, which means a difference of 59 cm between port and starboard. As it can be seen that the roll angle peaks to positive and negative values, both foils are constantly overcorrecting their current heights.

5. CONCLUSION & FURTHER RESEARCH

In this paper, different setups to use reinforcement learning to control hydrofoils are proposed. Using Deep Q-Learning and different reward functions, action sets and number of agents, we trained multiple setups and conducted experiments. Evaluating the performance of each setup with a run simulating 60 seconds, we conclude that the performance of a PID controller can be approached when using three agents, a small action set and a reward function based on height (experiment C). When using a larger action set or controlling multiple foils with one agent, the vessel can not be considered stable. The experiments using a reward function including the roll angle of the boat, were not considered stable as well.

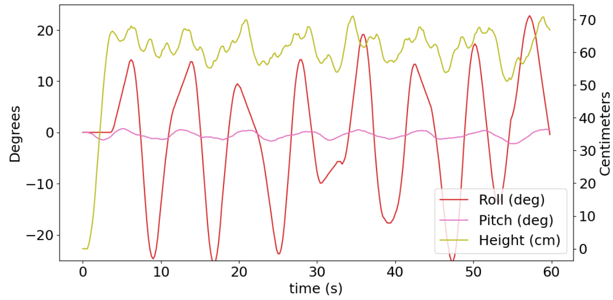

Figure 7. Experiment E, trained three agents with a reward function based on height and roll, and action set $\{-1.0, -0.9, ..., 0.9, 1.0\}$

In this research, only the Deep Q-Learning algorithm was used. Further research could be done on using other reinforcement learning algorithms or other reward functions. This might increase performance when controlling multiple foils with one agent. As controlling multiple foils with one agent might allow for controlling based on roll angle rather than height only, chances are that it is better capable of adapting to environments with waves, hence a larger chance to succeed in solving the problem described in section 1.

The simulation used for this research is a simplified representation, which may bias the performance. Further research could be performed using a simulation where more real-world properties are represented or waves are simulated, such to train agents on waves. As waves do have impact on the measures heights, a system trained on height measurements and not on roll angle, like in experiment C which had the best performance out of the experiments conducted in this research, we expect this will impact the stability as well. This can be validated by a research using a simulation including waves.

Lastly, the experiments done in this research could be performed on a real-world vessel. This way, the outcomes of the experiments of this research can be validated and the real-world usage of the approach suggested can be analyzed in more detail. This may make clear potential shortcomings of this research and the setups of the experiments.

6. REFERENCES

[1] The lift equation. https://www.grc.nasa.gov/www/k-12/airplane/lifteq.html. Accessed: 2021-12-06.

[2] Thermophysical properties of seawater. http://web.mit.edu/seawater/. Accessed: 2021-12-06.

[3] J. Bai and Y. Kim. Control of the vertical motion of a hydrofoil vessel. *Ships and Offshore Structures*, 5(3):189–198, 2010.

[4] R. Bencatel, S. Keerthivarman, I. Kolmanovsky, and A. R. Girard. Full state feedback foiling control for america's cup catamarans. *IEEE Transactions on Control Systems Technology*, 29(1):1–17, 2021.

[5] G. D. Bousquet, M. S. Triantafyllou, and J.-J. E. Slotine. Control of a flexible, surface-piercing hydrofoil for high-speed, small-scale applications. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4203–4208, 2017.

[6] O. M. Faltinsen. 6. hydrofoil vessels and foil theory, 2005.

[7] S. Liu, H. Niu, L. Zhang, and C. Xu. The longitudinal attitude control of the fully-submerged hydrofoil vessel based on the disturbance observer. In *2018 37th Chinese Control Conference (CCC)*, pages 397–401, 2018.

[8] Z. Ni, M. Dhanak, and T.-C. Su. Performance of a hydrofoil operating close to a free surface over a range of angles of attack. *International Journal of Naval Architecture and Ocean Engineering*, 13:1–11, 2021.

[9] B. Porter and A. Khaki-Sedigh. Design of digital set-point tracking pid controllers for hydrofoils. In *1988 International Conference on Control - CONTROL 88.*, pages 660–665, 1988.

[10] N. Sanghi. Deep reinforcement learning with python : with pytorch, tensorflow and openai gym, 2021.

[11] R. S. Sutton and A. G. Barto. *Reinforcement learning : an introduction*. The MIT Press, Cambridge, Massachusetts, 2018.