Investigating Aquatic Ecosystems with Computer Vision, Machine Learning and the Internet of Things

Tristan Y H Tay[1], Terence L Y Teo[1] and Kenneth Y T Lim[2[0000-0003-3756-6625]]

[1] independent author
[2] National Institute of Education, 1 Nanyang Walk, 637616 Singapore
kenneth.lim@nie.edu.sg

Abstract. This paper describes an independent research project undertaken by a pair of high school students in Singapore under the mentorship of a Research Scientist at the National Institute of Education. The importance of Artificial Intelligence (AI) and Machine Learning (ML) continues to increase every day in our constantly advancing society. This investigative study aims to uncover the practicality of both AI and ML by utilizing it in a real world context, more specifically, to monitor aquatic behaviour. The aim of this investigation focuses on how ML can be used alongside appropriate hardware to effectively monitor the behaviour of aquatic organisms in response to changes in environmental factors. This paper describes the design, construction, testing and design decisions behind the development of a module which allows us to use ML in tandem with IoT sensors to fulfill the above aim. This investigation ultimately concludes that we were able to successfully conceptualize and create two designs and accompanying prototypes where their strengths lie in the presence of an ecosystem in which sensor data measured can be easily compared to the number of fishes detected by our object detection model to draw relationships between aquatic behaviour and different environmental factors. However it is limited due to the main hindrance of the prototypes' incapability of maintaining clear visibility of fish in murky waters with high turbidity..

Keywords: Machine Learning, Arduino, Raspberry Pi.

1 Introduction

Machine Learning (ML) and Artificial Intelligence (AI) has increasingly found itself in the centre focus of the commercial and academic world, mainly because of the excitement and endless possibilities it brings for the world. When deliberating over what situation we should contextualize our usage of ML in, we stumbled upon a phenomenon at the Serangoon Gardens Secondary School pond. There, we noticed something curious: All the fishes were congregated together in one area of the pond and not at the other opposite end of the pond. We realised that periodically the fish would travel together between the two ends of the pond, and we wondered what factors were responsible for this phenomenon.

As such, this situation presented an opportunity for us to seize, in which we could utilize machine learning, accompanied with Internet of Things (IoT) sensors, to study aquatic behaviour in a localized environment, i.e. what are the microenvironmental factors which affect the behaviour of the fish in the pond.

Through this, we can contextualize our investigation in the field of marine biology, drawing relationships between abiotic factors, attained from sensors which monitor different environmental factors, and biotic behaviour, monitored using ML and its object recognition capabilities.

1.1 Objectives

In undertaking this study, we will be exploring how AI, or more specific to this application, ML, can be utilized and integrated with other data sensors to monitor aquatic biological behaviours, primarily, the behaviour of carp and other common fish used in Singapore school ponds. In doing so, we arrive at our main objectives for this investigation.

- designing an all-encompassing sensor module which comprises of image capturing and environmental monitoring capabilities, of which data recorded can be automatically uploaded onto the internet for analysis;
- training a object detection ML model so that the process of monitoring fish populations at opposite ends of the school pond can be automated by feeding the trained model with images captured by the sensor modules placed in the school pond; and
- utilize the sensor modules and trained ML model obtained from completing the two prior objectives in a real world application, placing them in the school pond to gain data on the microenvironmental factors as well as the respective fish populations at each end of the pond. Then, utilizing the data to draw relationships between the different microenvironmental factors and fish behaviour to arrive at conclusions as to how fish behaviour is affected by the present microenvironmental factors in the pond.

This paper will thus be organized accordingly, beginning with the literature which inspired us and directed us through this investigation, followed by the attempts at achieving the objectives listed above, with the main focus revolving around the hardware and software aspects of the prototypes we developed, in which we will be combining our methodology and results and discussion sections. We will then conclude with the general limitations with our investigation, as well as what our findings suggest for future research and prototyping.

2 Literature review

2.1 Necessity of investigation

Upon further research on the applicability of this investigation in real world contexts, it became apparent, from the Living Report 2020 by World Wildlife Fund (WWF), that the need to monitor aquatic behaviour in response to changing environmental factors is a must in the face of growing climate change and habitat destruction brought about by human activities, such as overfishing and coastal developments. Such human activities have led to the worsening of several environmental factors, such as the pH level, temperature and oxygen concentration of aquatic habitats, affecting the aquatic life which resides within them [1]. Hence this investigation is important and highly applicable in today's reality.

2.2 Application of Machine Learning

There have been existing studies which have already inculcated the usage of Artificial intelligence into aquatic biology monitoring systems. Salman *et al.* utilizes Gaussian Mixture Modelling (GMM), to model the background in a live video feed to detect objects such as fish which are not part of the background, optical flow, which detects movement in the video feed, and Region-based Convolutional Neural Network (R-CNN), which detects and proposes several regions where the object might be and limits object recognition to those areas for detection, with images from the Fish4Knowledge database for their usage in live fish sampling for estimation of fish populations [2]. Lu *et al.* on the other hand used transfer learning, utilizing pre-trained deep Convolutional Neural Network (CNN) models to train their own CNN model for fish species detection to detect different species of fish that have been caught to prevent overfishing [3].

However, an inherent flaw with the systems above is that despite them all having the ability to capture data (videos, pictures), the only data these systems are capable of collecting are pictures and videos. For commercial or recreational purposes, videos and pictures may be sufficient. But for research purposes, a lot more data is needed. In the context of biological or aquatic research, sensors like dissolved oxygen sensors, turbidity sensors, pH sensors, temperature sensors etc. are needed to understand and explain the biological behaviours of aquatic organisms in response to their surroundings.

Therefore, given the task, we would imagine an ideal version of a system, based on the above literature review, which would be: building a system that builds upon the above research and projects. Our system, which would be capable of remaining underwater for a sustained period of time (at least a week) autonomously, would include: a camera, multiple sensors that will work underwater which are relevant to aquatic research and the deployment of ML to help sift through all the data captured through the camera to aid identification of aquatic life.

From Fishcam [4] and Pipecam [5], we realised that a cylindrical pipe structure is commonly utilized for underwater monitoring of aquatic organisms and that Raspberry Pi is widely used in such underwater systems. Hence, we will be adapting both into our modules that we will be using for data collection.

2.3 Ideal aquatic environment

From literature review, we determined that the microenvironmental factors that make up good water quality for fish to reside in are > 4 mg/m3 of dissolved oxygen, pH level of around 7.5 - 8.5 as the average blood pH of fish is around 7.4, a turbidity value of 30 NTU (Nephelometric Turbidity Unit), and an optimum temperature which is dependent on the species of fish in question [6]. We realised that most of the fish in the school pond were koi fish, a subspecies of the common carp, hence the ideal temperature range would be around 20 - 25 °C [7].

These values will be kept in mind during our investigation, being useful in verifying the conclusions drawn from the obtained data, and should be considered if readers were to carry out similar investigations on carps as well.

3 Prototypes

3.1 Hardware

First prototype. We set out to create a pipe that is assembled from simple and cheap materials, as inspired by FishCam and PipeCam. For our module, we determined that we want to include a pH meter, dissolved oxygen sensor, temperature sensor, turbidity sensor. The dissolved oxygen sensor has to be calibrated for it to provide accurate readings. We have considered that should the sensor not be calibrated, the values obtained from the sensor would still be usable in our present context because we are merely comparing between dissolved oxygen readings. Their discrepancy from the actual value would be constant. However, to a researcher, accurate dissolved oxygen readings are critical. To calibrate the sensor, we needed Sodium Hydroxide (NaOH) solution. The manufacturer of the dissolved oxygen sensor is DFRobot, and it is specified that 0.5 Molar (M) of NaOH solution is required for proper calibration. The sensors would be connected to an Arduino board. The Arduino would be connected to a Raspberry Pi board, which controls the frequency of sensor readings being taken and images captured by the camera. The external shell of the module - which houses all the sensors and electronics - is a simple PVC pipe brought from a hardware store. To modify the pipe to fit our purpose, that is to allow the sensor to extend outside the PVC, we cut circular holes into the PVC pipe. In our attempt to make use, as much as possible, of equipment we have, we improvised and used a soldering tool to melt through the plastic and make holes in the PVC pipe. Following which, we fit the sensors and cameras into the pipe. For the camera we used a dome (improvised from Daiso's transparent christmas ball) to fit the camera within, allowing the camera to be adjusted in all x,y and z axis without the pipe itself needing to be tilted. The result was an integrated module which combines a multitude of sensors, cameras and internal computers. The pipe was sealed physically with duct tape and silicone sealant. We approached a Vocational Institute for help to professionalise the pipe construction. A rubber gasket was added, along with other waterproofing improvements. However, water could still enter the pipe after a period of time, rendering it incapable for long-term data collection. Although this pipe idea was, as one might say, a failure, we think that on the contrary, it is useful conceptually. This is because this present iteration is highly practical for research purposes as it is an integrated module for sensors and imaging capabilities. Should there be many data collection points, logistically it would be an issue for researchers. With a singular module that is capable of utilizing IoT, data collection can be done quicker and with greater frequency because data is uploaded to the cloud and transfer of equipment from one location to another is also efficient and easy. Hence, we believe that our hardware design works conceptually, the only problem is the lack of professional tools, knowledge, materials and machinery.

Second prototype. We were advised that, given our present situation, namely a lack of professional knowledge of waterproofing and the lack of equipment, an option would be to purchase IP68 boxes to house our components. This thought did occur to us at the very beginning, however we avoided it for our priority was to build an integrated module where all components fit into one common shell.

For the second prototype, we bought an IP68 waterproof box with a transparent cover, as well as a splash-proof box.

The choice of camera was also deliberated upon. We had two options: (1) Raspberry Pi Camera V2.1 is the standard camera that is typically used for Raspberry Pi units. This camera is equipped with the Sony IMX219 8-megapixel sensor, or (2) Raspberry Pi Infrared Camera Module is equipped with the 5-megapixel OV5647 sensor. This camera is unique as it emits infrared light, allowing it to capture images even in the dark.

We observed that the infrared camera had a wider field of view and could capture more in the picture it takes. We hence chose the infrared camera over the standard camera. However, the infrared light reflects off the transparent panel, resulting in glare in the picture. To resolve this problem, we unscrewed and removed the infrared lamps from the camera module itself. But, without the infrared light, the camera performs poorly in low light, thus needing a lot of light even in reasonably lit conditions.

After deciding on the most appropriate camera which should be utilized for our investigation, we faced another issue where we needed to ensure that the IP68 box which houses the camera module is able to sink to the bottom of the waterbody in which we are taking pictures of fish. It is here where buoyancy comes into play, given the equation from Archimedes Principle:

$$F_B = V\rho g$$

Where F_B represents buoyancy force, V the volume of liquid displaced by the object, ρ represents the density of the liquid, and g represents gravitational field strength. In order for the IP68 box to sink, its weight has to be greater than the buoyancy force, such that there is a net downwards force. Hence we get the following:

$$mg > V\rho g$$
$$m > V\rho$$

Where m represents the mass of the object, which in this case is the IP68 box. Given that the volume of the IP68 box is around 0.00181m3 and the density of water is approximately 1000kgm-3, the minimum mass needed for the IP68 box to sink would be 1.81kg. Hence we bought a weight which has a mass nearest to this value, which was 2.5kg, which we used to enable the IP68 box to sink.

Another problem we faced was imaging in harsh conditions. From our experience, murky water is an especially huge issue for the camera because the camera simply is not capable of capturing anything more than approximately 1 meter from it. Because our intention is to apply ML to recognise fishes, the images must be reasonably clear to at least distinguish vague lines of the fishes' silhouettes and outlines. To further investigate how water murkiness affects the image quality, we put our cameras to test in three locations, namely the pond in Serangoon Gardens Secondary, the Institute of Technical Education's (ITE) ecogarden pond - before and after pond cleaning, and an aquarium we set up at home just for this present testing purpose.

The pond in Serangoon Gardens secondary was extremely murky, with no fishes visible at all. The pond in ITE, prior to cleaning, was considerably murky but faint shapes of the fishes can be identified. After cleaning, more fish were visible. The home aquarium was the clearest, mostly because we can control the murkiness by replacing the water in the tank.

Thus for the investigation, the home aquarium served as the best-case scenario in which the water is not murky, and the pond at ITE was seen as a more realistic application of our prototype, where the water was murky but fish was still visible. The usage of the three different locations, in a sense, served as a litmus test for the effectiveness of our prototype, from which we realised that the largest limitation of our investigation would be that it is not applicable in water bodies that are too murky.

3.2 Software

Given the success of our more orthodox approach to this investigation, splitting the camera and sensors into two separate modules, we were able to obtain underwater pictures of fish for the training of our object detection model. Using about 200 pictures each, annotated with LabelImg, we trained two models, one to detect fish in the ITE pond, and the other to detect fish in the home aquarium. Training the models took about four hours each, but object detection itself was much faster, giving us images with drawn bounding boxes and labels around the fish in both locations (see Figs. 1 and 2 as examples).

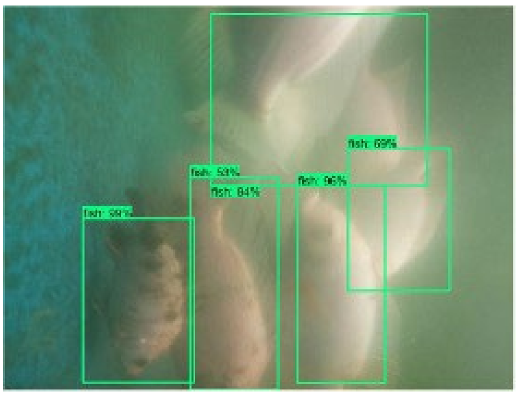

Fig. 1. Results with bounding boxes and labels for ITE pond

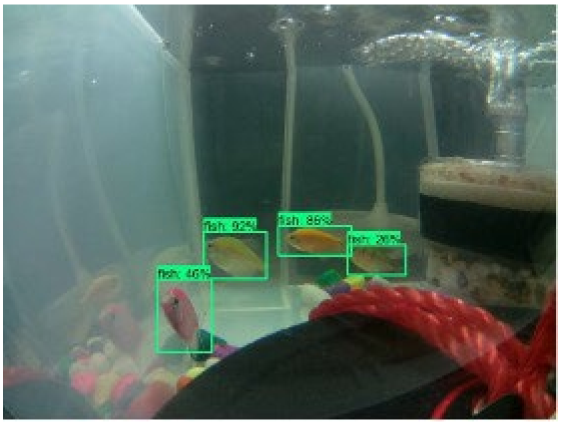

Fig. 2. Results with bounding boxes and labels for home aquarium

The bounding boxes also come with accompanying labels with the percentage match stated as well. The object detection model can be easily configured to only draw out bounding boxes for matches that are only above a fixed percentage match. For our investigation, and in the above images, we restricted the model to only draw bounding boxes for matches above a 20% percentage match. We also scripted the object detection model to output the number of objects detected in each image in an Excel spreadsheet so that easy comparisons, graphs and conclusions can be drawn from the fish population detected and the sensor readings.

We utilised the xlsxwriter python package to output the probability and number of fishes onto an excel file. Below is the function which draws the images and outputs the accompanying data into an excel sheet.

4 Concluding remarks

In conclusion, with respect to the three objectives mentioned at the beginning of this paper, we were successful in achieving two out of the three intended objectives for this investigation, namely designing an all-encompassing sensor module as well as training an object detection ML model to work in tandem with our sensor module. In the process, we were able to conceptualize two designs for the hardware, one being the cylindrical all-in-one pipe which would be successful given proper professional equipment and techniques for waterproofing, and the other being the two boxes one ecosystem model which relies more heavily on IoT for data from both environmental sensors and camera to be synced, taken simultaneously and uploaded onto the internet for analysis.

The remaining objective – applying the hardware and software developed from the preceding two objectives respectively into a real world application – could not be achieved due to the visibility of the water in the school pond. This problem arose from the murkiness of the water and the lack of professional equipment for enhanced visibility in such turbid environments. However, even with this setback, our investigation through the physical creation and testing of our conceptualized models still functions as a proof of concept, where both our hardware and software are able to function successfully together to allow for more fruitful investigations and research work in monitoring aquatic behaviour in response to environmental change.

The findings from our investigation also pose several future implications. Although the usage of AI and ML are achieving greater and greater prevalence in society, many still view its usage as complicated and too expensive, not to mention inculcating its usage in tandem with other supporting pieces of hardware to achieve a specific goal such as monitoring aquatic behaviour. Our investigation seeks to go against this perception, to show that creating hardware and utilizing AI alongside it can be as simple and affordable as creating a Do It Yourself (DIY) model with a PVC Pipe, or using an IP68 box and a socket box. The authorial team began this investigation without knowledge on AI and ML. But through guidance and independent learning, as well as a little of a maker spirit, we were able to accomplish the aforementioned objectives. However, with all projects comes its limitations, and some questions still have to be explored to affirm the validity of our achievements, and extend its usage beyond the localized context of school ponds and aquariums. How can the existing designs for monitoring aquatic behaviour be improved to allow for clearer imaging in turbid environments? How can such modules be utilized in response to the aforementioned issue of growing climate change?

References

1. Almond, R. E. A., Grooten, M., Petersen, T. (eds.) Living Planet Report 2020 – Bending the curve of biodiversity loss. World Wildlife Fund, Gland, Switzerland (2020).
2. Salman, A. *et al.*: Automatic fish detection in underwater videos by a deep neural network-based hybrid motion learning system. ICES Journal of Marine Science, 77, 1295-1307 (2019).
3. Lu, Y.-C., Tung, C., Kuo, Y.-F.: Identifying the species of harvested tuna and billfish using deep convolutional neural networks. ICES Journal of Marine Science, 77, 1318-1329 (2019).
4. Mouy, X., Black, M., Cox, K. Qualley, J., Mireault, C., Dosso, S., Juanes, F. Fishcam: a low-cost open source autonomous camera for aquatic research. HardwareX. 8. e00110. 10.1016/j.ohx.2020.e00110 (2020).
5. Fourie, F., Thomson, D. Pipecam: low-cost autonomous underwater camera. Hackaday.io, last accessed 2021/9/9.
6. Monitoring pond water quality to improve production. Thefishsite.com, last accessed 2022/1/12.
7. FAO: natural food and feeding habits. Fao.org, last accessed 2022/1/12.