Imperial College London

Industrial Placement
Final Report

22/04/2020 - 21/09/2020

Aida Manzano Kharman

CID: 01353594

Personal Tutor: Petar Kormushev

Word count: 1980

DIGITAL ADVERTISING RECYCLING PLATFORM

CONTENTS PAGE

EXECUTIVE SUMMARY

My role in the company was to develop the working prototypes of the two main products of U:CAAN: the Digital Advertising Recycling Pods (DARPs) and the LittaHunt App. I worked as a computer vision engineer intern where I learnt about computer vision, cloud computing, APIs, containers, microcontrollers, and implemented this knowledge in the working prototypes I produced for the company's minimum viable products. These deliverables that I provided, greatly surpassed my expectations for this placement and will give the company a stronger credibility as it continues to expand and pitch for further investment and market collaborations.

Design Engineering has equipped me with a set of skills that have enabled me to learn and adapt fast in a field of work that is highly specialised and in which I had no prior experience with, which is computer vision. I had this interest throughout my last two years of university, which I pursued in my Robotics projects, but it was through this experience in which I finally delved deeper into the theory and to put it into practise. The two most important skills I learnt from university that aided me in this placement were to A) know what I don't know and B) fail fast.

Through this placement I have also found a niche that I am passionate in pursing, and that is the intersection between sustainability, policy making, and STEM. It is the policies of governing bodies such as the EU that drive local authorities to shift towards more sustainable solutions, and it is tech's role to fill the gap between the need and the solution. I want my work to lie within that intersection, and I believe that as a Design Engineer, I have the skills to both understand and put in practise the technology to provide these solutions, but also to drive and design the projects that will move our society towards a more sustainable future.

CONTRIBUTIONS AND ACHIEVEMENTS

I worked on two projects during the placement: the Digital Advertising Recycling Pods (DARPs) and the LittaHunt App*. Given there was no physical prototype, I started by brainstorming the user journey and touch points of the customer with the guidance of my supervisor to create a technical specification for the system, shown in Figure 1.

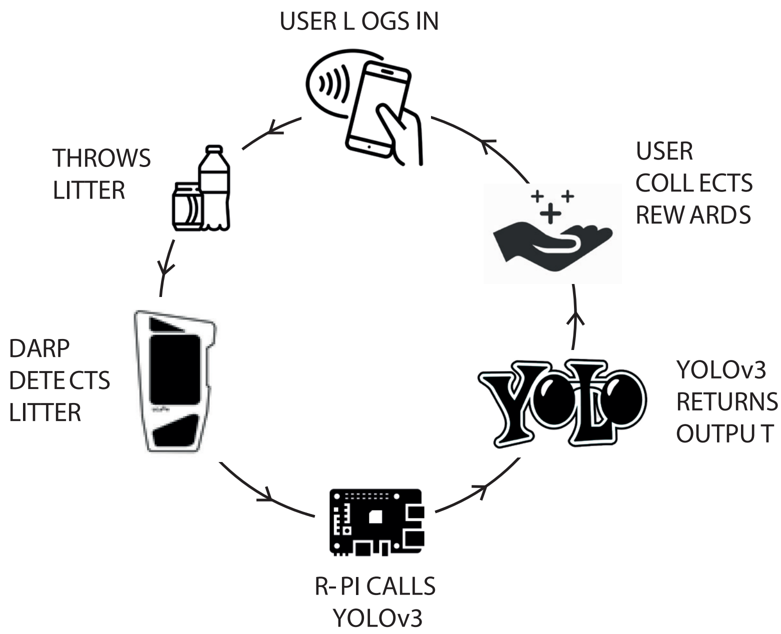

Figure 1: System design of the Digital Advertising Recycling Pods. (A.Manzano, 2020)

After the specifications were defined, I researched which image recognition algorithm was the most suitable one for this application and found it to be YOLOv3. I then trained it on a custom dataset to successfully detect aluminium cans, PET bottles, glass bottles and HDPEM bottles.

After this, I researched which microcontroller was the most affordable out of the ones that met the technical requirements to be able to run YOLOv3 and selected a Raspberry Pi (R-Pi) 4 to do so. I connected it a to a camera and optimised it to run neural network computations with a package called NNPACK and installed Darknet, the framework in which YOLOv3 was built and trained on.

By the end of the first half of the placement, I'd designed and implemented from scratch a system for which I created a working prototype, where the camera would capture an image of the litter, run the trained YOLOv3 and return the output of the prediction via the R-Pi, as shown in figures 2, 3 and 4.

For better understanding of the two projects I worked on please refer to the appendix.

Figure 2: HDPEM Milk bottles

Figure 3: PET Plastic bottles

Figure 4: Glass bottles

Figure 5: Aluminium Cans

My second main contribution was to the LittaHunt App. I also began this project by brainstorming the system design to establish the technical specifications.

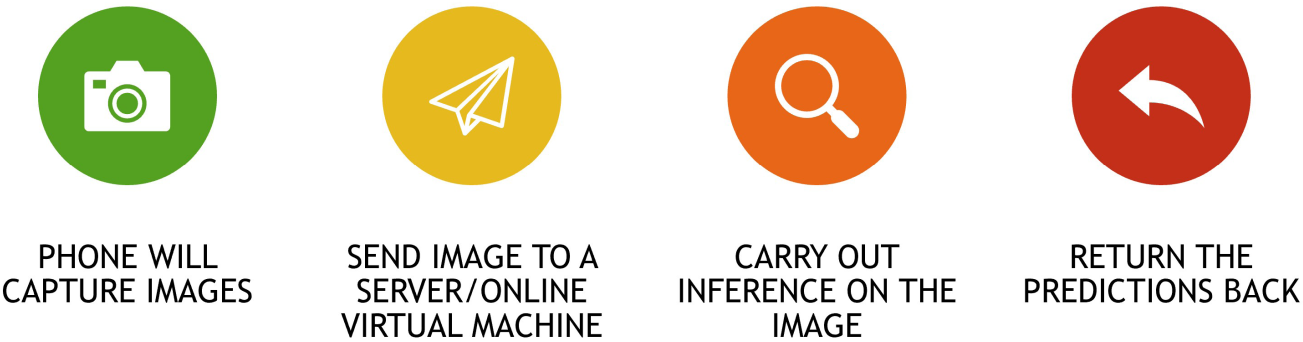

Figure 6: System design of the LittaHunt App. (A.Manzano, 2020)

I then researched how to best implement this system by looking further into APIs and client-server applications and found “Flask” to be the best method, given how simple and intuitive it is, and most importantly, its ability to scale up to complex applications *(The Pallets Projects, 2018)*. The latter is of utmost importance when thinking of the long term growth of the company. The final prototype I produced was a client that would post an image to the server, where the inference code was called to run the algorithm on the sent image, and then return the prediction back to the client.

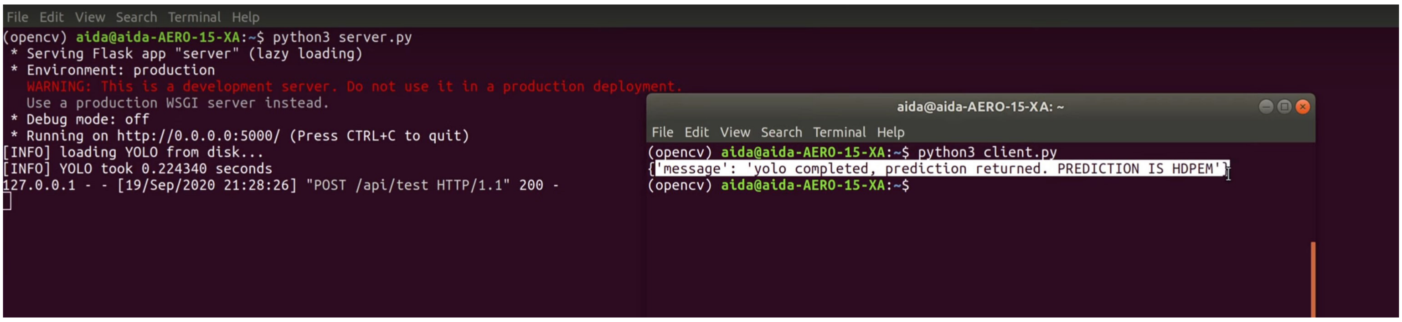

Figure 7: Working client and server system. (A.Manzano, 2020)
*Click **here** for video demo.*

Finally, I put this entire application on Docker, which is a program that allows developers to package up an application with all of the parts it needs, such as libraries and other dependencies, and deploy it as one package *(Opensource, 2017)*.

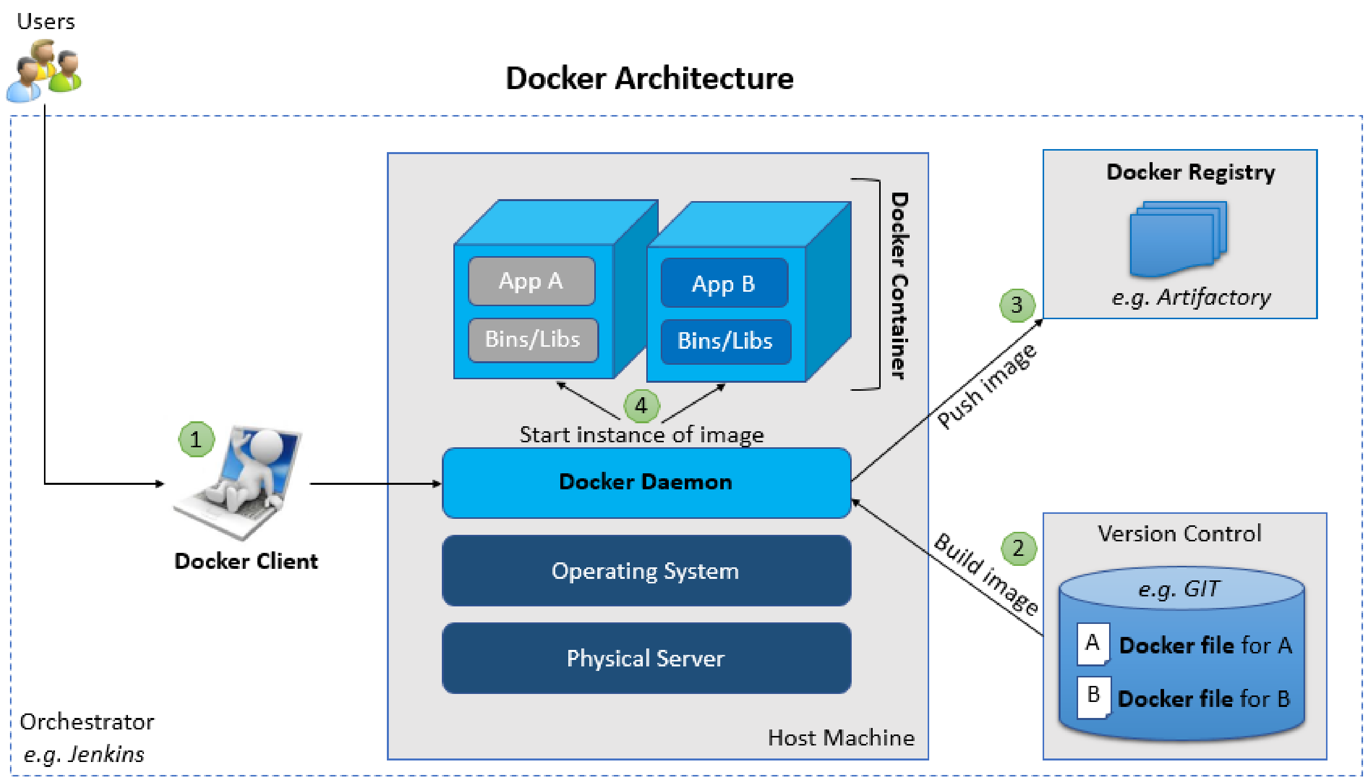

Figure 7: Architechture of Docker.
(Accenture 2018).

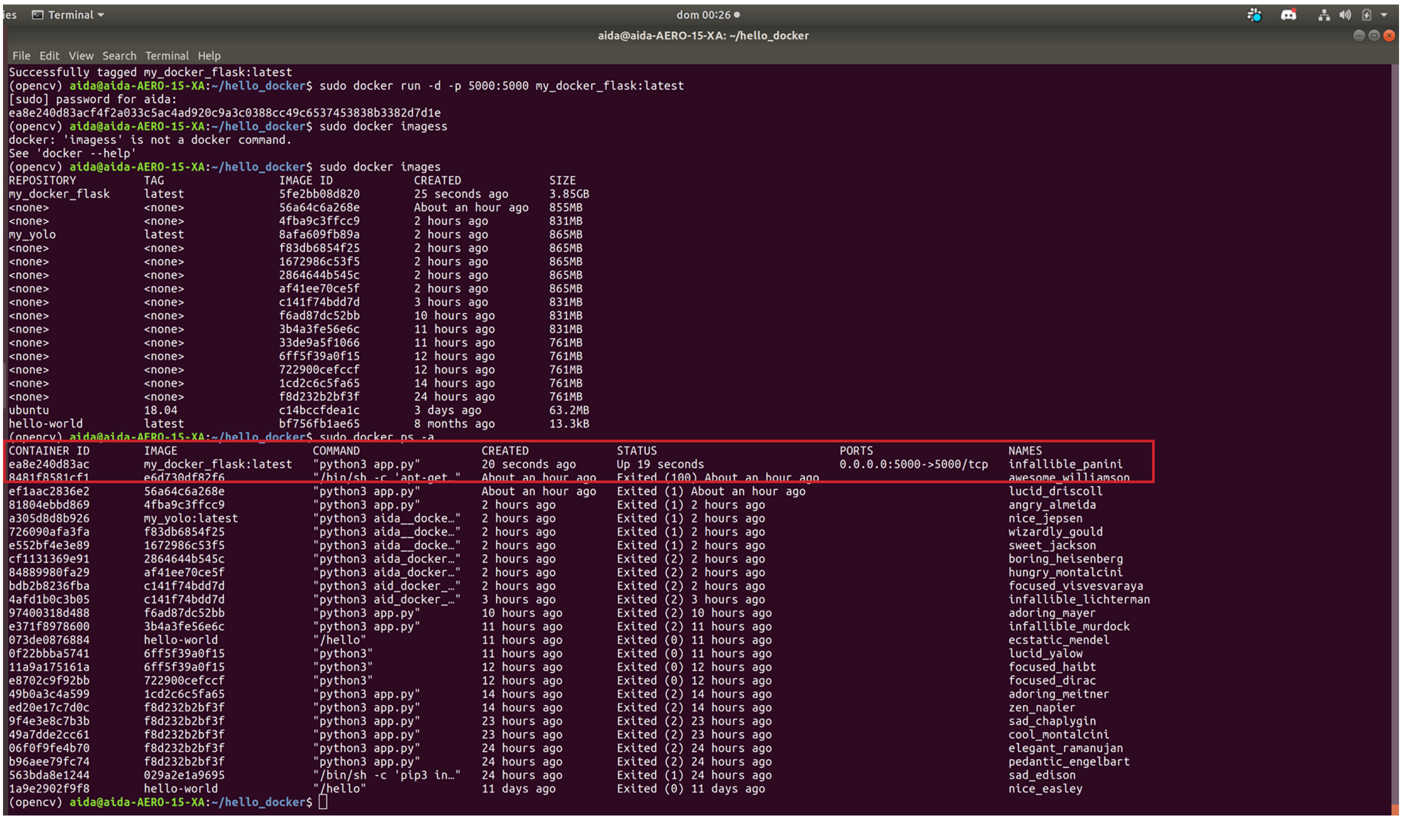

Figure 8: Flask client-server application running on Docker container, with port 5000 open. (A.Manzano, 2020)

This allows the application to be portable and easily deployable as well as scalable, and it will save the company's future software developers time and effort by allowing the development team and the testing and deployment team to work on the same environments and remove the common issues of missing dependencies, conflicting installations or deprecated packages.

Alongside the aforementioned, my research into Amazon Web Services' (AWS) cloud solutions has provided the company with a thorough insight as to which solution is the most cost effective for U:CAAN's applications whilst taking into account its future scalability. I set up a virtual machine instance on AWS Elastic Cloud Compute (EC2) with all the necessary requirements to run the trained YOLOv3, including CUDA, cuDNN, OpenCV, Darknet and of course the image recognition algorithm's trained weights. This will be crucial to U:CAAN as it begins deploying its products and its need for processing power and computing resources rapidly increases. My choice of using AWS has given U:CAAN a head start of £10,000 in credits, as AWS awards this to new start-ups.

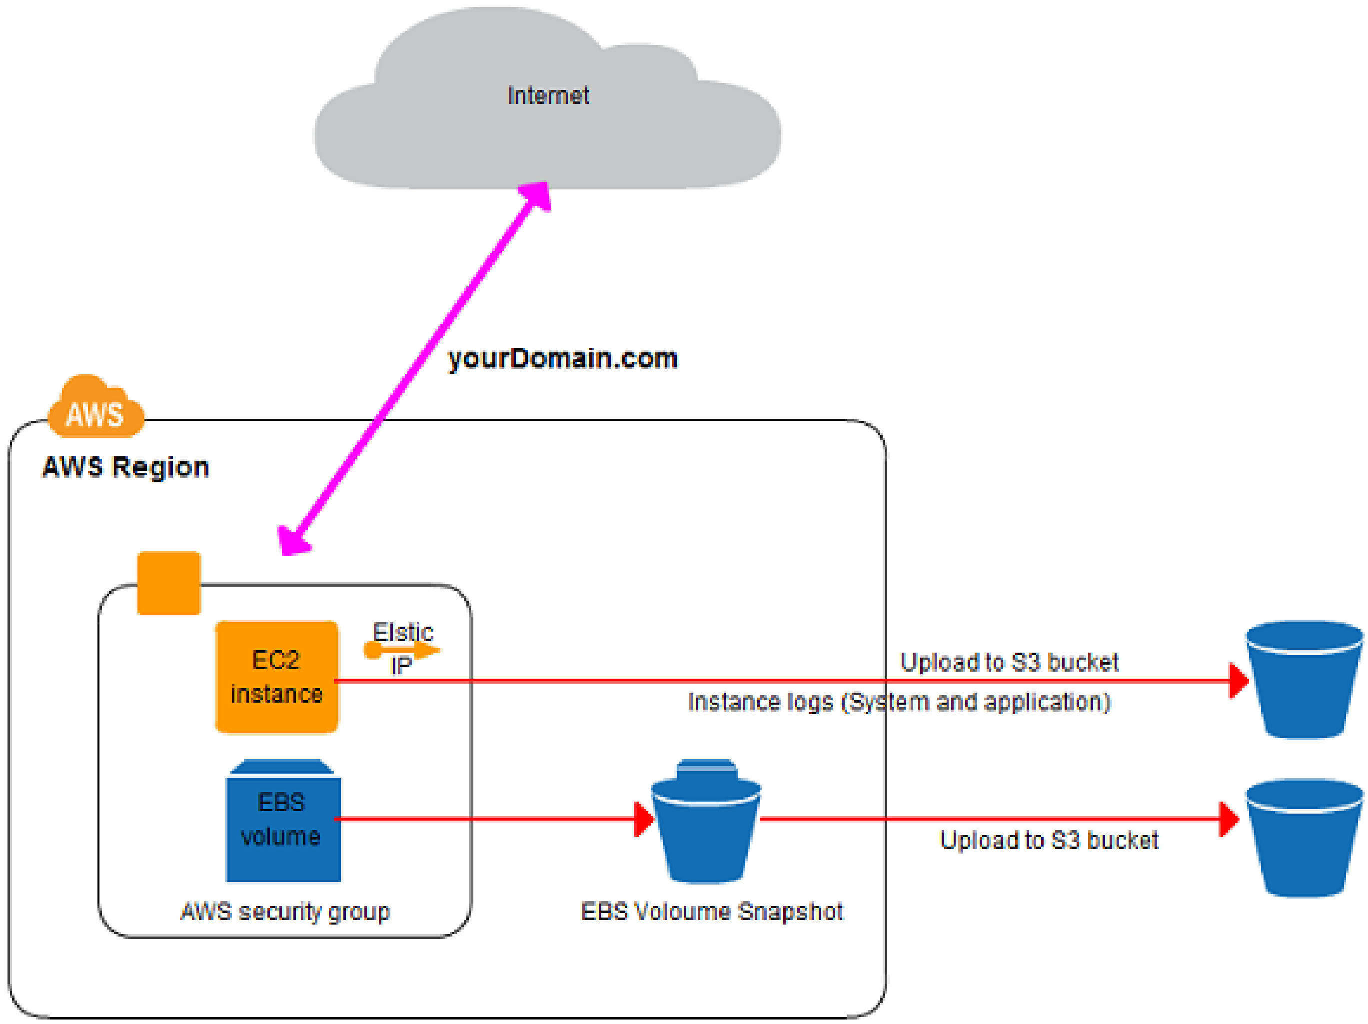

Figure 9: Basic structure of an Amazon Web Services Elastic Cloud Compute instance. (Creately, 2018).

My contributions are mainly technical deliverables that are of great value to the company to be able to present to grant applications, investors and potential market collaborations as proof of concept that their products are indeed viable and have been successfully prototyped and tested.

I have made sure to adopt the best practises in the field of computing that will save the company time and effort, and as a result, costs further down the line by planning in advance for the long term of the company's growth, and ensuring that my prototypes are scalable, portable and deployable.

A: Use a combination of general and specialist design engineering knowledge and understanding to optimise the application of existing and emerging technology:

Through the research of the user touch points and the development of the system design of the DARPs and the LittaHunt App, I applied both my design thinking skills as well as my electronics and computing skills to create the technical specifications necessary to produce the working prototype of both products.

I gained knowledge on the strengths and limitations of the available image recognition algorithms and their appropriate use cases as well as applying knowledge learnt in DE2 Big Data module about training models and overfitting. Furthermore, I acquired the skillset of implementing APIs by developing a functional client server application with Flask and containerising it on Docker.

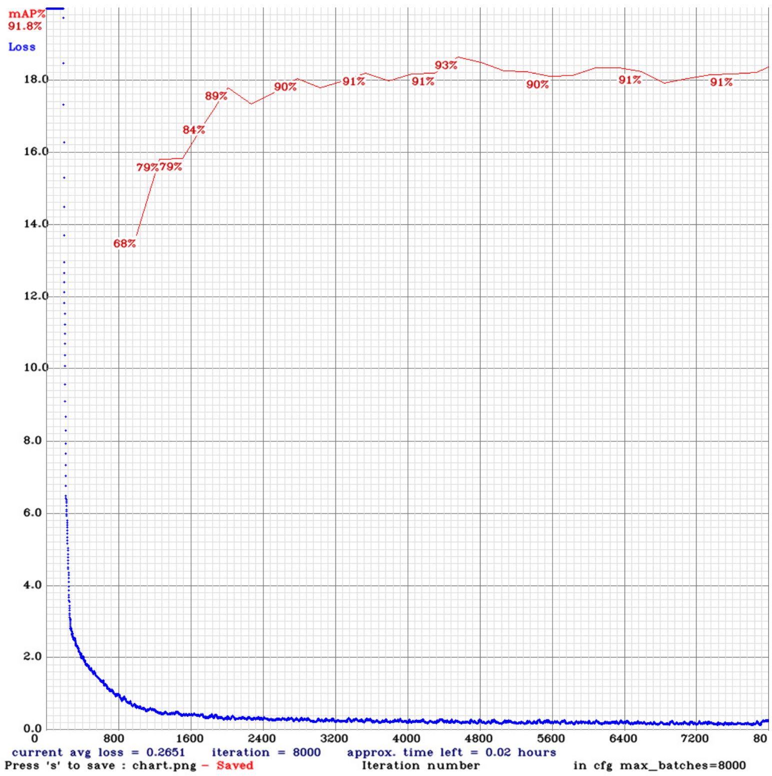

Figure 10: Training process of YOLOv3 on custom dataset. The graph shows how the model begins to overfit at around 5000 iterations. (A.Manzano, 2020)

Amongst my strengths I have found that I can work well with uncertainty and I am able to extract tangible milestones from a very open brief and then successfully achieve them. However, I have noticed that I struggle with the need to understand in detail the topic I am researching and found that this could lead me down a never ending cycle of delving deeper and deeper into the topic.

For my future personal development, I strive to continue to pursue my passion and curiosity of learning and truly understanding new topics I encounter, whilst also maintaining a balance with being time efficient and knowing when to start implementing the theory into practical applications.

B: Apply appropriate theoretical and practical methods to the analysis and solution of design engineering problems:

Applying the theoretical knowledge gained on existing image recognition algorithms, I selected the most appropriate algorithm given the specifications of this project and trained it on a custom dataset to solve the engineering problem of how to detect and classify litter.

Performance on the COCO Dataset

Model	Train	Test	mAP	FLOPS	FPS
SSD300	COCO trainval	test-dev	41.2	-	46
SSD500	COCO trainval	test-dev	46.5	-	19
YOLOv2 608x608	COCO trainval	test-dev	48.1	62.94 Bn	40
Tiny YOLO	COCO trainval	test-dev	23.7	5.41 Bn	244
SSD321	COCO trainval	test-dev	45.4	-	16
DSSD321	COCO trainval	test-dev	46.1	-	12
R-FCN	COCO trainval	test-dev	51.9	-	12
SSD513	COCO trainval	test-dev	50.4	-	8
DSSD513	COCO trainval	test-dev	53.3	-	6
FPN FRCN	COCO trainval	test-dev	59.1	-	6
Retinanet-50-500	COCO trainval	test-dev	50.9	-	14
Retinanet-101-500	COCO trainval	test-dev	53.1	-	11
Retinanet-101-800	COCO trainval	test-dev	57.5	-	5
YOLOv3-320	COCO trainval	test-dev	51.5	38.97 Bn	45
YOLOv3-416	COCO trainval	test-dev	55.3	65.86 Bn	35
YOLOv3-608	COCO trainval	test-dev	57.9	140.69 Bn	20
YOLOv3-tiny	COCO trainval	test-dev	33.1	5.56 Bn	220
YOLOv3-spp	COCO trainval	test-dev	60.6	141.45 Bn	20

Figure 11: Comparison of common image recognition algorithms' performance. and I chose YOLO mainly because its higher speed in predictions without too great of a sacrifice on accuracy. It also had the most updated and clear documentation on how to train on a custom dataset. (J. Redmon 2016)

To solve the problem I encountered of missing libraries or dependencies when trying to deploy my code, I researched how to maintain a consistent environment throughout development, testing and deployment. I found this issue to be common in the software development community, and that the best practise is to use Docker to make your application portable on a container. This allows the coding environment to be the same throughout the last three phases of the SLDC cycle and makes it very simple to package and deploy your code.

Figure 12: Docker provides a consistent computing environment throughout the whole SLDC. (Edureka, 2019).

I successfully containerised my prototype for the LittaHunt App on Docker, and thus established the precedent for good coding practise for the company's future software developers. These acquired skills and techniques are ones that I will be implementing in my future computing projects as they will save me and the teams I work with time and effort in the long term.

C: Provide technical and commercial leadership:

Through my advising of which cloud computing service was the most cost effective and best suited for U:CAAN's applications the company received £10,000 worth of credits from AWS for being a start-up. I then created a virtual machine on AWS that will be part of the necessary infrastructure that the company will require as it resource demand increases.

Furthermore, I also demonstrated commercial leadership by contacting the CEO of RECYCLEYE for potential collaborations as well as contacting researchers in the field of recycling and consumer products. Thanks to this, U:CAAN is now collaborating with Dr. Marco Auriscchio from Imperial College in two upcoming grant applications.

I have demonstrated an aptitude in networking and taking the initiative to contact researchers and entrepreneurs in the same field, which have had tangible, long term impact on the company's growth. For the future, I will strive to work on my ability to pitch ideas quickly, succinctly and without relying heavily on technical terms to convey the mission and vision of the company. This is a crucial skill to work on, as often, these ideas will be pitched to a non-scientific audience.

D: Demonstrate effective interpersonal skills:

Most of my communication during the placement was through online work channels and weekly calls where I would present my work for that week. I would discuss with my supervisor the research and the coding tasks I had been working on and the work I was to carry out the following week.

Overall, because all my interactions with my supervisor, potential investors, researchers and members of the open source community were entirely online, I learnt the importance of clear and concise communication. This was a skill I initially lacked, but it is one that is not only relevant to the situation brought by COVID 19, and it is one I aim to incorporate in my future professional career.

E: Demonstrate a personal commitment to professional standards, recognising obligations to society, the profession and the environment:

In order to maintain a high-quality professional standard, I realised that there is a clear bottleneck in the performance of the custom trained algorithm given the limitation of the available dataset.

Having tested it, I noticed it performs significantly better on videos than images, as can be seen in figures 13 to 17 shown below. I could have changed the system on the R-Pi and the client server Flask app to take an input of a short video instead of a single image, but I believe this would be a temporary fix that would incur extra costs to the company further down the line should they utilise EC2 or other cloud services, as the resource demand would increase and therefore so would the costs.

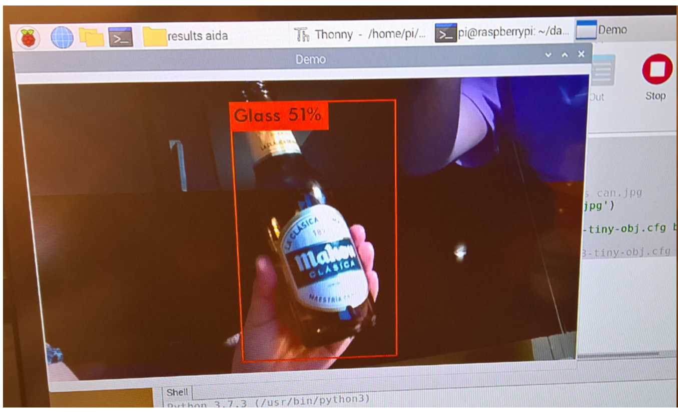

Figure 13: R-Pi real time results

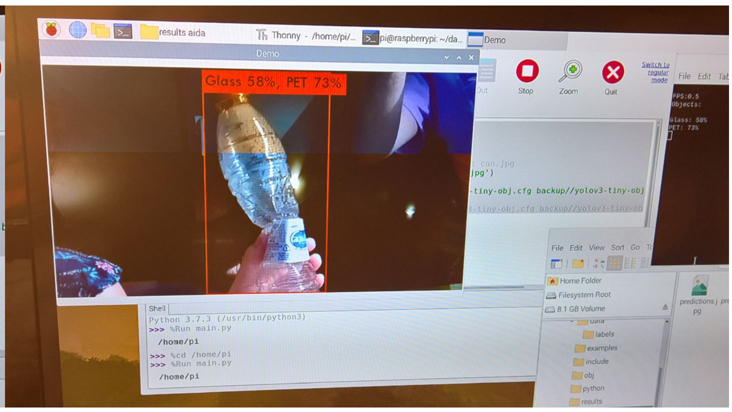

Figure 14: R-Pi real time results

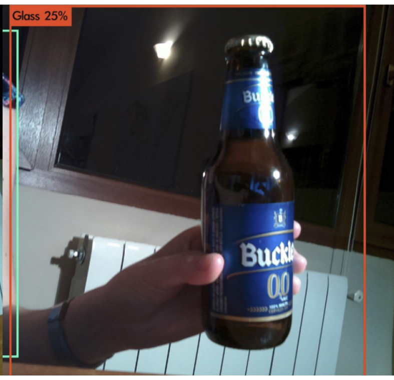

Figure 15 and 16: R-Pi results on a captured and stored image

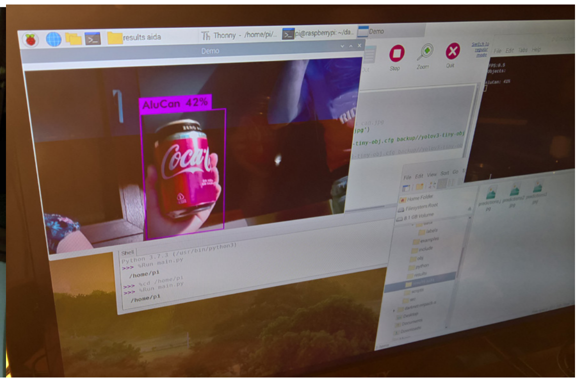

Figure 17: R-Pi real time results

Furthermore, estimates from the International Energy Agency suggest that the amount of power required to feed the world's need for data storage and processing in 2028 will be equivalent to the energy needs of the world's third most populous country, the United States *(IEA, 2019)*. As such, I want to be mindful of how the company utilizes this service so to not further contribute to the environmental strain that cloud computing creates, especially since U:CAAN is a sustainability driven company, focusing on reducing and minimising waste.

ACKNOWLEDGEMENTS

Thank you to Matt Pollen for taking me on board his team and for his continuous support and encouragement.

Thank you to Arkadiy Serezhkin for allowing me to use his dataset and for offering advice and guidance with the training process.

Thank you to Alexey AB for his extremely comprehensive guide on Darknet and how to train YOLOv3 on a custom dataset.

Thank you to Dr. Petar Kormushev for his support and mentorship throughout this placement.

APPENDIX

U:CAAN is a recycling start up that aims to incentivise recycling by gamifying the process and targeting a younger and wider audience of users. It intends to do so with two main products:

1. The Digital Advertising Recycling Pods (DARPs)
2. LittaHunt App

The DARPs are smart bins in which the user can log onto their U:CAAN app by scanning their phone using NFC, throw in the litter and the pod will detect what litter they threw in using a camera and a microcontroller that will call a trained image recognition algorithm. This algorithm will be used to detect the type of litter thrown and give the user reward points accordingly. This allows for the litter to be easily sorted within the pod, meaning that the process of recycling becomes much simpler, and it is also incentivised with the reward points that would be able to be exchanged for refreshments or discounts.

Aside from targeting individual customers, U:CAAN can also target larger Fast Moving Consumer Goods companies such as Coca Cola or Walkers Crisps for example. This is because a legislation that is already being widely adopted in Europe called the Extended Producer Responsibility (EPR) Act, will be adopted in the UK in 2023 *(UK GOV, 2019)*.

EPR uses financial incentives to encourage manufacturers to design environmentally friendly products by holding producers responsible for the costs of managing their products at end of life *(OECD, 2001)*. This means that the data collected through the smart bins of the amounts of packages, cans, bottles or other types of litter and their brands can be sold to their relative companies to demonstrate quantitative data to governing bodies on the amount of recycled packaging units, and thus be able to demonstrate their compliance with EPR.

The LittaHunt App is a different product, in which the user has a pet “litter monster” that they care for and feed with images of litter. The user takes images of litter on the streets and this litter feeds their monster as well creating a map of litter hotspots. The concept follows the very popular game of Pokemon Go and targets a much younger audience, and at the same time, provides U:CAAN with information on the accumulations of litter that can then be provided to institutional customers such as town councils. This allows them to save resources by optimising the garbage collection routes based on demand, as opposed to on arbitrarily allocated time intervals *(UK GOV, 2020)*.

REFERENCES

1. “Flask” The Pallets Projects, 2018. Available online at: https://palletsprojects.com/p/flask/

2. “What is Docker” Opensource, 2017. Available online at: https://opensource.com/resources/what-docker

3. “World Energy Outlook Report”, International Energy Agency, 2019. Available online at: https://www.iea.org/reports/world-energy-outlook-2019

4. “Extended Producer Responsibility: A Guidance Manual for Governments.” OECD. Paris, France. 2001.

5. “Consultation on reforming the UK packaging producer responsibility system” UK government. 2019.

6. “Guidance on prioritising waste collection services during coronavirus (COVID-19) pandemic” UK Government. 2020.