ASHESI UNIVERSITY

DESIGN AND IMPLEMENTATION OF A SPEED DETECTION SYSTEM USING RADAR AND IMAGE PROCESSING

CAPSTONE PROJECT

B.Sc. Computer Engineering

Tarek Nii Kpakpo Quao

2020

ASHESI UNIVERSITY

DESIGN AND IMPLEMENTATION OF A SPEED DETECTION SYSTEM USING RADAR AND IMAGE PROCESSING

CAPSTONE PROJECT

Capstone Project submitted to the Department of Engineering, Ashesi University in partial fulfilment of the requirements for the award of Bachelor of Science degree in Computer Engineering.

Tarek Nii Kpakpo Quao

2020

DECLARATION

I hereby declare that this capstone is the result of my own original work and that no part of it has been presented for another degree in this university or elsewhere.

Candidate's Signature:

……………**T.N.K.Q.**……………………………………………………………………

Candidate's Name:

……………Tarek Nii Kpakpo Quao………………………………………………

Date:

………………**May 29, 2020**……………………………………………………………

I hereby declare that the preparation and presentation of this capstone were supervised in accordance with the guidelines on the supervision of capstone laid down by Ashesi University.

Supervisor's Signature:

………………………………………………………………………………………

Supervisor's Name:

………………………………………………………………………………………

Date:

………………………………………………………………………………………

Acknowledgments

To my supervisor, Dr. Robert Sowah whose encouragement and academic advice helped me undertake this project.

To Mr. Francis Gatsi who suggested the project idea and served as a source of advice and guidance before I got in touch with Dr. Sowah

Abstract

The issue of vehicle accidents occurring on the roads on a daily basis, which only seems to be increased by the minute is what this paper researches and seeks to tackle. With the simple technology that has been made available globally, a cost-effective system can be developed to fight overspeeding, reduce road accidents and save lives. As a result, in this project, an Automatic Number Plate Recognition program and a web service for speed detection were implemented. This was done with the aim of monitoring frequently used roads, and those areaswhere the accidents normally take place in order to sanction drivers guilty of driving above set speeding limits. This made use of machine learning-based Optical Character Recognition program couples with a web application to issue speeding fines dynamically to offenders via email. In the course of this project a validation accuracy of 97% was achieved for the training done with Nigerian License plate characters. The mailing was fully functionally at the end of the project. Though, a few hitches were experienced on the license plate detection phase as some license plate shapes vary slightly. This proved the need to improve the license plate detection to detect modded license plates in addition to the regular ones.

Table of Contents

DECLARATION i

Acknowledgments ii

Abstract iii

List of Tables of Figures vi

List of Figures vii

Chapter 1: Introduction 1

1.1 Background 1

Chapter 2: Literature Review 5

Chapter 3: Design 9

3.1 Requirements Specifications 9

3.1.1 Product perspective 9

3.1.2 User interfaces 9

3.1.3 Hardware interfaces 10

3.1.4 Software interfaces 10

3.1.5 Communications interfaces 10

3.1.6 Operations 11

3.1.8 System Requirements 12

3.1.9 User Requirements 14

3.1.10 Constraints 15

Chapter 4: Implementation and Methodology 24

4.1 Hardware Setup 24

4.2 Software 24

4.2.1 Web interface 24

4.2.2 Backend 25

4.3 Doppler Effect Phenomenon 26

4.4 Data Preparation 28

4.5 Training The Model 29

4.6 Implementation 33

Chapter 5: Results and Discussion 38

5.1 Observations 38

5.2 Discussion 43

5.3 System Integration 43

Chapter 6: Conclusions, Limitations and Future works 45

6.1 Conclusions and Limitations..45
6.2 Future works..46

List of Tables of Figures

Table 3.1: Pugh matrix..12
Table 4.1: Adminstrative features and their priorities 25
Table 4.2: Comparison between the different classifiers explored....................................33

List of Figures

Figure 1.1: Actions carried out towards enforcing road safety [4] 3
Figure 3.1: A flow of the operation of the speed detection system setup 14
Figure 3.2 Block diagram of the intended speed detection system....................................15
Figure 3.3: UML database notation 18
Figure 3.4: High-level system architecture Software Architecture....................................19
Figure 3.5: Hardware block diagram 20
Figure 3.6: Activity Diagram 20
Figure 4.1: A visual representation of the doppler effect phenomenon [19] 27
Figure 4.2: Architecture of ANPR 30
Figure 4.3: The Flow of the License plate's detection and character extraction 30
Figure 4.4: The accuracy of the character recognition model's training and testing..........31
Figure 4.5: The loss of the Character Recognition Model's Training and Testing.............31
Figure 4.6:Admin sign-in page 35
Figure 4.7:Dashboard...36
Figure 4.8: The speeding Fine report template 37
Figure 5.1: Image of vehicle with the license plate number "GC 2259-09"run through the trained Automatic License Plate Recognition program 39
Figure 5.2: Result of the test with the license plate number "GC 2259-09 39
Figure 5.3: Image of vehicle with license plate number "Gx 1047-18" run through the trained License Plate Recognition program 40
Figure 5.4: Result of the test with license plate number "Gx 1047-18" 40
Figure 5.5: Image of vehicle with Nigerian license plate 41
Figure 5.6: The succesful output of the linear-SVM text recognition model 41
Figure 5.7: Failed license plate localization for "GC 2259-09" from earlier testing..........42

Chapter 1: Introduction

1.1 Background

Vehicle accidents in recent years have been occurring frequently, claiming the lives of hundreds of thousands in Africa. One problem that needs solving is the large number of road accidents in African countries. As a result of limited time and resources available, the scope of this problem will be narrowed down to Ghana, but particularly in the capital city of Accra and its surrounding towns. Accidents are noted to be one of the main causes of death among Ghanaians. WHO confirms this as it found out that "93% of the world's fatalities on the roads occur in low- and middle-income countries, even though these countries have approximately 60% of the world's vehicles." [1]

According to the World Bank, "there is general recognition of road traffic injuries and fatalities, little is known about the link between road traffic injuries and economic growth." [2] They also pointed out that "investments in road safety are also an investment in human capital," which are the main contributors to how well or poorly a country's economy performs. Also, it claims that the majority of people involved in accidents, who die or become disabled as a result, make up a large number of the working force. Hence, the more people we have with disabilities, especially those with severe permanent injuries, the fewer people a country has available for economic-related activities. Apart from injuries and deaths, accidents can have psychological effects on victims.

According to the World Health Organisation, "an increase in average speed is directly related both to the likelihood of a crash occurring. Hence this affects the severity of the consequences of the crash" [1]. World Health Organisation puts this in the sense that "every 1% increase in mean speed produces a 4% increase in the fatal crash risk and a 3% increase in the serious crash risk" [1]. The study also confirmed that "the death risk for pedestrians hit by car fronts rises rapidly (4.5 times from 50 km/h to 65 km/h)".[1] As cited by Warburton and Chen, in 2003, Runge pointed out that unsafe speed contributed to 31% of all fatal crashes in the United States, causing 13,380 deaths. According to the National Highway Traffic Safety Agency, the cost of speed-related collisions to American society is over $40billion per year[3].

Solving this problem would help cut down the high mortality rate in the country, which is caused by road accidents since Ghana falls within the group of low and middle-income country brackets. When more skilled human lives are sustained, these individuals can work within the jobs they are qualified for, which eventually improves the economic state of the country. Apart from having vital lives retained, the enormous costs incurred by government bodies whenever accidents take place in the country can be saved and channeled into more productive areas of national development[3].

So far, the attempts the Ghanaian state has made to tackle road accidents have been mainly in the form of educating road users, particularly the drivers. Though this has had some positive impact but not enough to significantly address the issue of road accidents.

A study conducted by Bonnet, Lechat, and Riddle in 2018 brought out information categorizing the efforts made by various African countries to combat Road Traffic Accidents. The respective percentages the categories represent out of the overall findings were also stated. The categories include Data Collection, Health Education, Road Safety Policy, and Safety Equipment[4]. Below is the Pie Chart, showing the distribution:

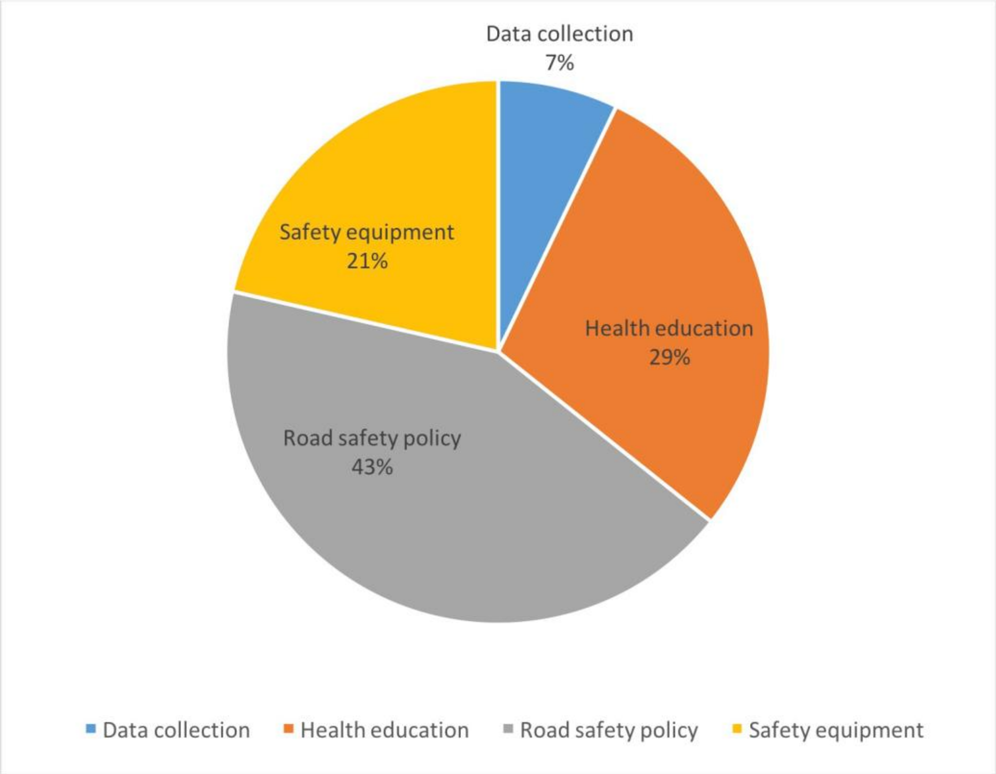

Figure 0.1: Actions carried out towards enforcing road safety [4]

In this light, efforts have been made to educate both drivers and pedestrians on driving safely and at acceptable speeds. These have taken the form of one-on-one interactions with road users. In contrast, the media has been used to disseminate bits of information to educate the public, through short texts and video infomercials.

The proposed solution to this problem is to develop a speed detection system with Automatic Number Plate Recognition, to be used on Ghanaian roads, particularly highways and motorways since much has not been done by the authorities to curb how often road accidents take place in the country. In doing this, it was realized that there is also a need to modify the solution slightly from the usual radar/laser implementation, which is more hardware-based/driven, to a more software-based approach, by way of image processing. This system would detect and track vehicle objects on the camera feed. An algorithm will be developed, which will be able to determine the speed of moving vehicles within the camera's field of view.

The reason for this solution is because of the positive effects it has had on other countries, by decreasing the rate at which accidents occur and, in the long run, drastically reduced the number of injuries and deaths that resulted from road accidents. A proof of this is seen in Warburton and Chen's research which noticed "average speeds dropped by about 7 percent (or 2.4 miles per hour), and (adjusted for trend and other changes) there was an estimated 33 percent reduction in injury/fatality collisions (4,030 fewer) and a 40 percent reduction in deaths.[3] They also proved that the cost the state incurs on road accidents is less than the amount of money spent by the government in generating speed tickets and sanctioning speeders in court.[3]

Though this is not an entirely new solution, it is relatively new within the Ghanaian context. With it, drivers in a bid to comply with driving regulations will be compelled to drive within safe speed limits. This system will, in turn, reduce over-speeding and resulting in road accidents.

It would also help the government save money that is used to repair state property, like street lights and traffic lights that get damaged when these road accidents happen. As of the year 2000, Bonnet *et al.* recorded that road accidents were noted to form 1% cost of the Gross Domestic Product (GDP) of most low-income countries, of which several are located in Africa [4]. Hence, the GDPs of these countries can grow to an extent if this problem is controlled. In addition to this, this system makes valuable data available that can be used to facilitate transport planning [5].

Chapter 2 : Literature Review

According to data collected by the Ghana Open Data Initiative from 1991 to 2011, the number of casualties and fatalities from road accidents has steadily grown during a large chunk of years in both urban and non-urban areas. Though there were years where these numbers reduced, in both areas, the number of fatalities in urban areas in 2011 (668 people) became over 400 more people than it was in 1991 (266 people). Rural areas, on the other hand, had become just under 1000 more people in 2011 (1531 people) than it was in 1991 (644) [6]. The high numbers measured in rural areas, though unexpected, make sense since many of these areas are used in traveling between some of the major cities in Ghana. Also, in grouping these numbers in the later years, 2010 and 2011, in the twenty-year period, according to regions, the major regions recorded some of the highest figures. Some of these regions, like the Greater Accra and Ashanti regions, are some of the most developed and populated areas. As a result, this makes such regions the most likely to have more vehicle users and pedestrians altogether; hence, it is, one way or the other, reasonable to expect such figures and trends. As of 1991 Ashanti region had recorded 183 traffic fatalities, which had increased to 474 in 2011, while Greater Accra had 126 and 408 fatalities in 1991 and 2011, respectively[7].

All but one of the studies showed the effectiveness of cameras up to three years or less after their introduction; one study showed sustained longer-term effects, about 4.6 years after introduction [8]. Reductions in outcomes across studies ranged from 5% to 69% for collisions,12% to 65% for injuries, and 17% to 71% for deaths near camera sites. The reductions over more extensive geographical areas were of a similar order of magnitude [8].

They could then conclude that the introduction of speed cameras is an effective intervention in reducing road traffic collisions and related casualties. Apart from them, Graham *et al.* [9] also developed a "causal inference framework to quantify the effects of speed cameras" in the same year. They pointed out that, on the other hand, some people speed cameras have no effect or that they even increase road collisions. They reported that road collisions went down in target areas by 14.4% to 15.5%. From the kind results they obtained following the methods used, Graham *et al.* believe such a framework can also help develop a more effective way of establishing speed cameras in various areas [9].

Many people in Canada, also feel that speed enforcement in the form of the introduction of speed cameras is a government scheme to make money off citizens who drive as a form of government revenue. Tay carries out a study to relate the number of crashes reported to authorities, hence recorded, per month after some years [10]. He also looks at the number of tickets issued to offenders throughout the time the speed cameras had been operated. At the end of the study, some conclusions that could be made are when speed cameras function for more hours, and fewer car crashes are recorded. The study also proved that when more tickets were issued to drivers, the deterrent effect it has on drivers' speed proved to be higher than the deterrent effect police officials stationed near roads had. Fortunately, the money obtained from the fines helps to improve the efficiency of road safety schemes, including the speed camera programme.

Much research has been conducted to find more accurate and cost-effective alternatives for how speed detection systems are implemented. One common approach that has been considered quite frequently is Image Processing. Image processing is among several technologies that are making it easier to develop an automated system for a wide range of purposes. Image processing itself has helped analyse traffic situations to modify how traffic systems function in controlling traffic flow. This is mainly based on what the camera is "*can see*" on the road, at varying times of the day. This project work is aimed at looking at how image processing can be adopted in traffic speed detection systems. When fully implemented, it will ultimately put away the sensors that have been utilized in time past in the detection systems, namely Radar and lidar sensors.

Ibrahim, ElGendy, and ElShafee strongly believe that a speed detection system developed primarily using image processing techniques on video streams of moving vehicles is more cost-effective than with the usage of radars and laser sensors [11]. These sensors have their weaknesses in the form of radar detectors and laser jammers that are used by some drivers to prevent these sensors from correctly detecting the actual speed with which they are driving. That, in effect, makes laser/radar speed detection systems ineffective in such situations.

Balasubramani [12] outlines various image processing techniques that can be used for reading the speed of moving vehicles. Among these are background image subtraction, which includes separating the foreground elements of an image from its background, and edge extraction, which focuses on the objects outlines to do some form of separation from the rest of the image. These approaches solely make use of the software without the need for any other device like the RADAR. The software used to execute computations involved with these techniques is MATLAB, which is used for data analysis of various forms through programming and built-in toolboxes.

It achieves its speed calculation by developing an algorithm to detect vehicles moving in the camera's field of view. It then rapidly captures frames while tracking the vehicle as it moves from one end of the camera to the other. It records the frame number that the vehicle entered the camera's scene as well as the frame number it left the scene. The speed is finally calculated by computing the number of frames it took the object to pass through the scene. Since the time of each frame is known, that knowledge is used to obtain the time taken to pass through the entire view of the camera.

Though the image processing-based system is considered as a better alternative to the RADAR implementation, in terms of the various costs involved, it is preferred in this case because of the limited scope of the project and funds available to carry out the project. Apart from that, the image processing implementation, though cheap, concerning components and maintenance, is more complex and requires some advanced knowledge in video signal processing.

In some Western/developed countries, laser speed cameras are used to measure the speed at which drivers drove their vehicles, while others make use of Radar in place of the laser. The problem with this particular mechanism of detecting the speed of moving vehicles is with the aid of laser detectors. Some drivers can disrupt the signal reflected from the laser receiver, making it unable to have an accurate speed reading of the target vehicle. Besides this, it is also challenging to be able to track multiple vehicles that are moving above the specific speed limit. Through the use of image processing techniques, each vehicle can be highlighted in a frame, and the speed at the instance the photo is captured can be noted down and considered in issuing a particular fine on the speeding ticket. The maintenance cost of this would be slightly less.

Chapter 3: Design

3.1 Requirements Specifications

3.1.1 Product perspective

The system falls into two main parts, namely the hardware and software aspects. The hardware part is represented by a camera being upheld on the side of a road, which detects the speeds vehicles are moving at. Unlike the already implemented version of this system in many countries, this system includes a specialised software which facilitates the creation and issuing of speed tickets for speeding drivers. The software also helps to keep track of all current speed offenders, and also has sections that differentiate drivers with outstanding tickets to be paid from those who have cleared on ticket fines.

3.1.2 User interfaces

The system is secured with a limited administrator account and login feature. This grants specific people with access to sensitive driver data.

There is also a comprehensive dashboard on the homepage after a successful login. This dashboard displays the different parts of the software system that are usable by the administrator, which will mainly be an official from the Driver and Vehicle Licensing Authority (DVLA). Though, to facilitate enforcement of these rules, The Police will need to be a vital part of the bigger system, though they will not be directly holding a speed detection system in hand. One key reason for making the DVLA a central part of this is, once offenders have not paid their fines, they will be denied access to their driver’s licenses other vital documents, upon expiry.

3.1.3 Hardware interfaces

A USB port is required to connect the camera to the controller of the system to allow for communication between the camera and the software.

3.1.4 Software interfaces

This system consists of some software components that range from databases to libraries utilized to achieve the desired functionalities. These include:

a. MongoDB database: this database is more flexible and scalable, in the sense that it grows to suit increasing data volume needs as compared to some other database frameworks

b. The Python OpenCV library is used to carry out image processing on video stream and the image of identified speeding vehicles captured. The character recognition for the system needs to be based on the machine learning algorithm that will be used with the OpenCV. A number of these will be explored, with the most efficient and accurate getting used in the final system. o is the K-Nearest Neighbour (KNN) algorithm. Different versions supervised learning classifiers were made with three different K values (3, 4 and 5) to observe their varying performances. A simple convolutional neural network(CNN)-based model was also implemented and tested for the text recognition aspect, this particular.

The operating system to be used is the Raspbian, which is the default operating system pre-installed on any raspberry pi board (including the versions of software component used).

3.1.5 Communications interfaces

- FTP: for the transmission of images of speeding vehicles (i.e., those with their license plates picked out) captured from streets to temporary storage points.

- Email protocol: for sending a ticket to vehicle owner's address
- Phone protocol (or API): for sending ticket info as SMS to drivers
- HTTP protocol for uploading tickets issued to an online server/database

3.1.6 Operations

The vehicle speed detection software for this system will run on a Raspberry Pi 3, which is managed by the Raspbian operating system. The software will be set up to interact with a software database as well as a web app (for managing cleared and uncleared speeding offenses that have taken place).

The software will also interact with the laser sensor, which is one of the means of determining the speed of vehicles.

The system will need to be active during both daytime and nighttime. Accidents occur both during the daytime and nighttime; hence none should be neglected.

The system will, like some already existing speed detection system, will be installed on the sides of roads. This will allow it to be useful during night time when Traffic officials are unavailable.

The main areas where the system will most likely dominate are highways, towns, and outskirts of towns, specifically roads connecting two major towns or regions in Ghana. This is because several accidents have occurred when people travel on such roads due to overspeeding a few other minor factors like the poor state of some of these roads. The various areas where the system will determine the kind of power to be supplied to it. Those that will be in towns and some highways will receive power from the main grid. The ones that will be found on the outskirts of towns will have to rely on solar power system since the national grid in such areas is not accessible.

Table 0.1: Pugh matrix

Criteria		**Baseline**	**Weight**	**A (Solely image processing)**	**B (Image processing & Radar)**
Safety		0	3	1	1
Ease set up		0	1	1	0
Start Cost		0	2	-1	1
Maintenance cost		0	2	1	0
Usability at night		0	3	0	0
Reliability (Distortion)		0	4	1	1
TOTAL			15	8	9

In the Pugh Matrix above, the baseline is considered as a speed detection system that operates solely using radar emitter and receiver.

Concept A depicts a system that does all object, tracking, detection, and speed estimation using image processing techniques alone, as well as for number plate detection and extraction.

Concept B depicts a system that uses a laser to detect vehicle speed and uses image processing to detect the number plate of a vehicle from a snapshot taken by the camera, extract the info for ticket generation

3.1.8 System Requirements

1. The system should be able to monitor the movement of vehicles across its field of view.

2. It should be able to measure the speed of the moving vehicle.

3. It should be able to capture images of vehicle objects in the camera's field of view, moving at speed [significantly] over the specified speed limit.
4. It should be able to extract number plate info from an image captured of a speeding vehicle.
5. The system should be able to locate the extracted number plate info with the driver/vehicle owner's information within a database.
6. The system should be able to autonomously generate a speeding ticket with details about the offense, such as fine amount, date, and time. (and send it digitally to the vehicle owner, which may not apply to all drivers).
7. The system should work when disconnected from the national grid.
8. The system should not be bulky.
9. The system should be able to have a visible view of the roads and any vehicle at night.
10. Each speed detection camera set should have unique ids (which they attach to speed ticket) to provide the location (or street) where it is sited, or even more simply, each camera should provide its location [hard coded].
11. All camera locations should be viewable from a central management point/center in an office.
12. The object detection software should be able to identify the number plate on any image of a vehicle taken by the camera.
13. The system should be able to keep track of the time, in this case, the GMT, to add to proof of offense, facilitating court proceedings.

14. The system should be able to access the Internet to upload images of speeding vehicles taken.
15. The system should be able to share an error or distress signal whenever a fault occurs with a laser sensor.

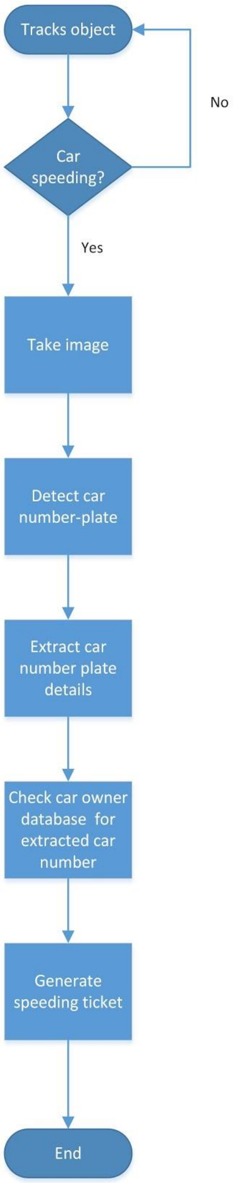

Figure 0.1: A flow of the operation of the speed detection system setup

3.1.9 User Requirements

1. The user should be able to view ticket payment statuses of all offenses.

2. The user should be able to send a reminder/warning to the offender if payment has not been made after a specified period.

3. The user should be able to print out a hard copy of the speeding ticket generated for mailing.

3.1.10 Constraints

Due to financial limits, the camera used might be unable to measure certain extremely high speeds due to a low camera frame rate.

3.2 Block Design

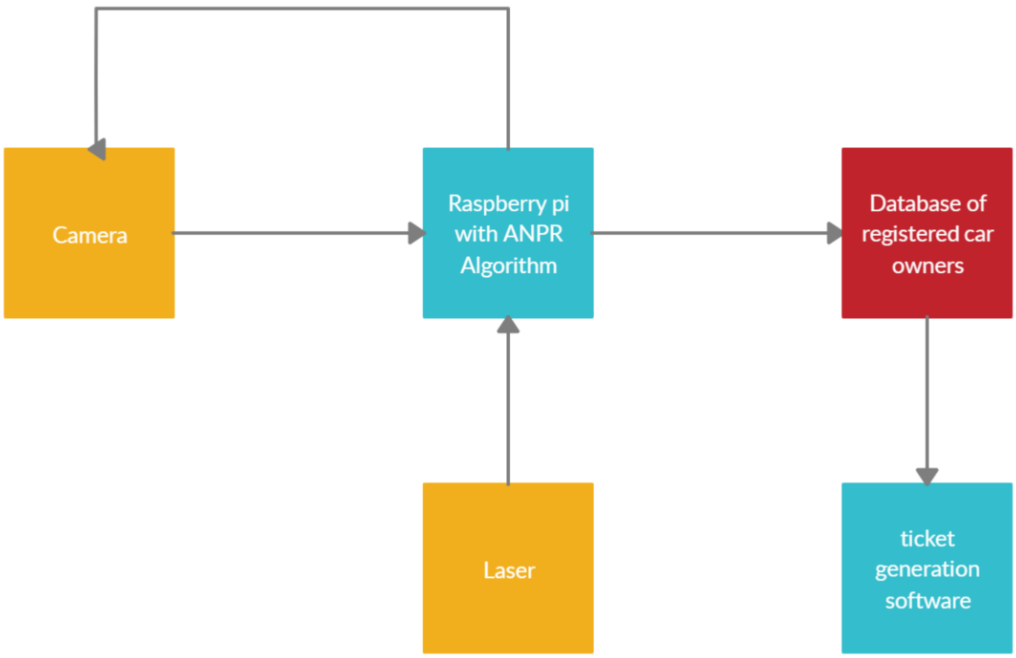

Figure 0.2 Block diagram of the intended speed detection system

The central/key block of the designed block diagram of the desired system is the ANPR program. The program will be running on the raspberry pi, which will have the radar sensor and the camera connected to it as well.

3.3 The Database

The database being used is the MySQL framework within SQLAlchemy, which is a SQL toolkit within the Python programming languages. It is also referred to as an Object Relational Mapper (ORM). SQLAlchemy provides the interface for creating an SQL-based database without the need to write any SQL statements. [13] This makes it possible to use Python code in any project or system to map data from the database schema to the applications' python objects. As a result of this, this made way for me to have easily focused my efforts on how the web service would work while a simple Python code dealt with any database interactions that needed to take place. In all, it reduces the amount of code implemented as well.

The choice of the database management system is MySQL, which is preferred because of its high scalability and satisfactory security capabilities. It being a relational database makes it easier to establish orderly and meaningful relationships between multiple database tables easily. This is the particular form of SQL database that is used underneath SQLAlchemy with the Flask web development framework. This makes it less intensive to retrieve data from the database for any use by the system while producing the implemented code base smaller than the case with actual SQL queries.

The database will exist in two main database tables. One table contains details of all registered vehicle owners, including the license plate number assigned to the owner by the DVLA, owner's name as well as one or two contact details. The vehicle owners' contact details needed include their phone number and email address. These will be used to notify them of their offense and the fine they are to pay. Subsequently, when offenders delay in paying, these details will serve as a way to remind them of the pending payment while stating the consequences/downsides they could suffer if they fail to make payment. One measure ensuring this is the inability to renew a driver's license and other documents upon their expiry.

The second table of the database will include the details of any speed ticket that is generated and issued. Each speed ticket record in the database will be referenced to its respective vehicle owner to easily modify a vehicle owner's fine payment status of all speeding offences.

The choice of what data will be present in the database table of all registered car owners was made based on the information that DVLA require applicants of vehicle ownership to provide on the forms, which are available on the agency's website, as well as the details used in creating an online profile with DVLA

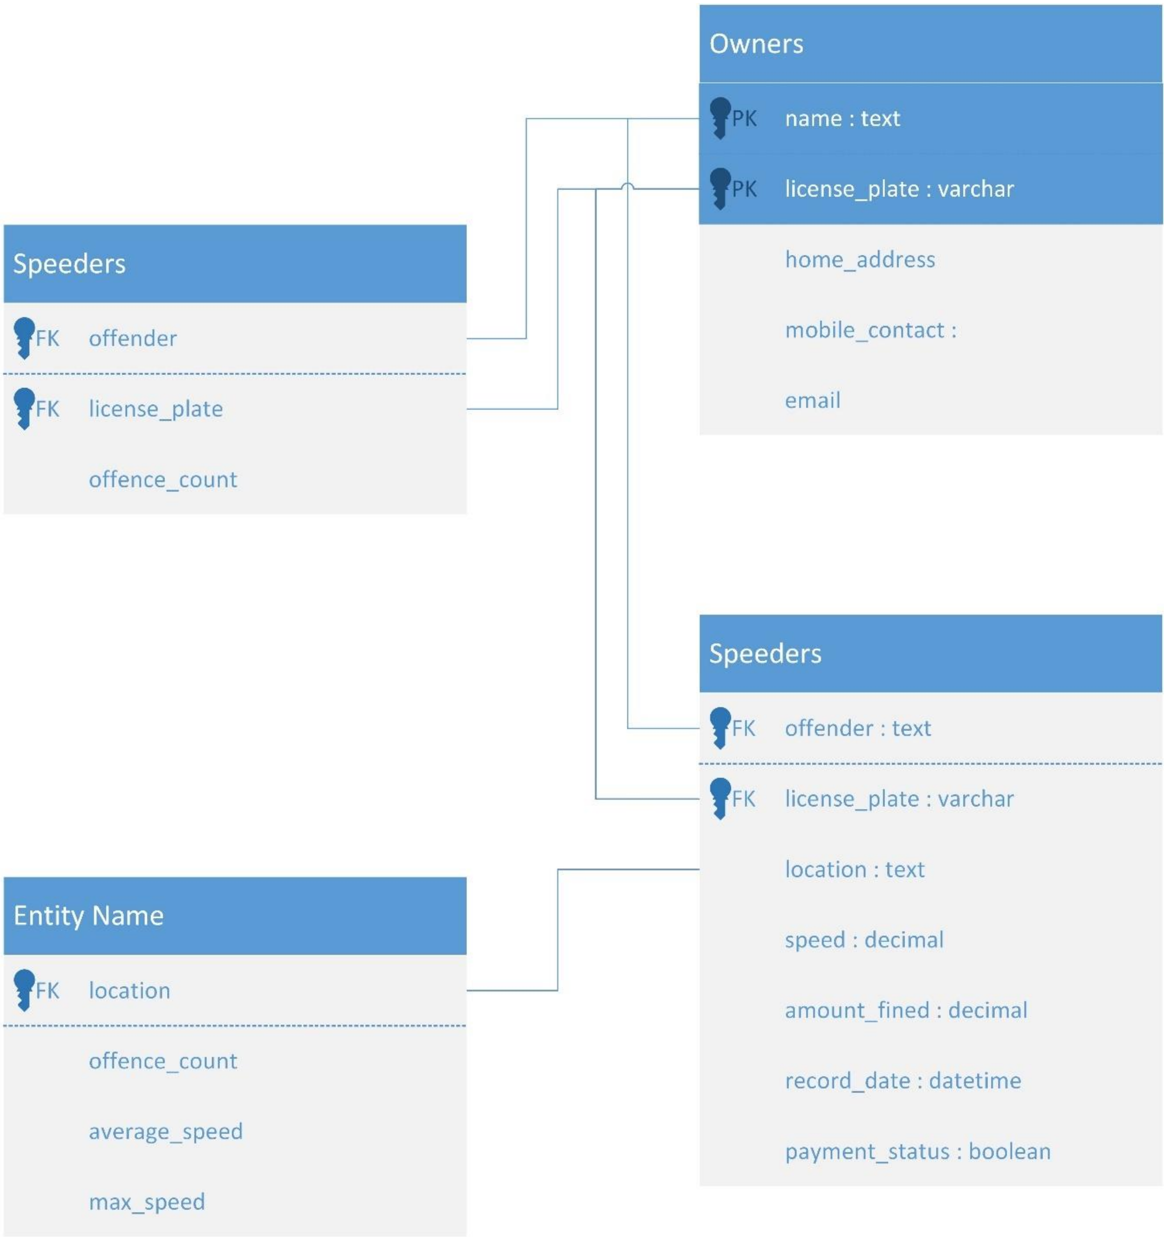

Figure 0.3: UML database notation

The above database notation simply describes the attributes of each of the two database tables as well as the relationships that exist between the database tables of the system.

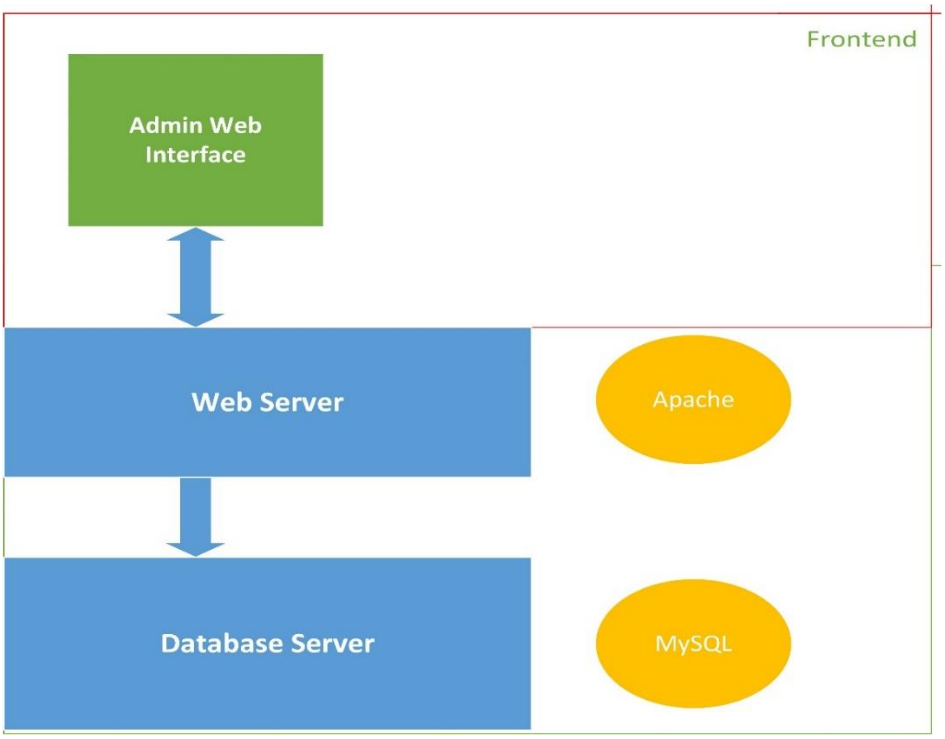

Figure 0.4: High-level system architecture Software Architecture

The diagram above depicts the software architecture and points out the various interactions that take place between the software components, as well as the directions of these interactions.

3.4 Hardware

Initially, a laser sensor, that is, emitter and receiver, was to be used to measure the speed of any moving vehicle closest to the system since it had better accuracy in speed detection. Its accuracy and precision are attributed to its project laser beam having a more direct target whose speed it reads. Unfortunately, due to the possible harmful effects, the laser emissions can have on people, the RADAR was used has been opted for, in its place for the meantime.

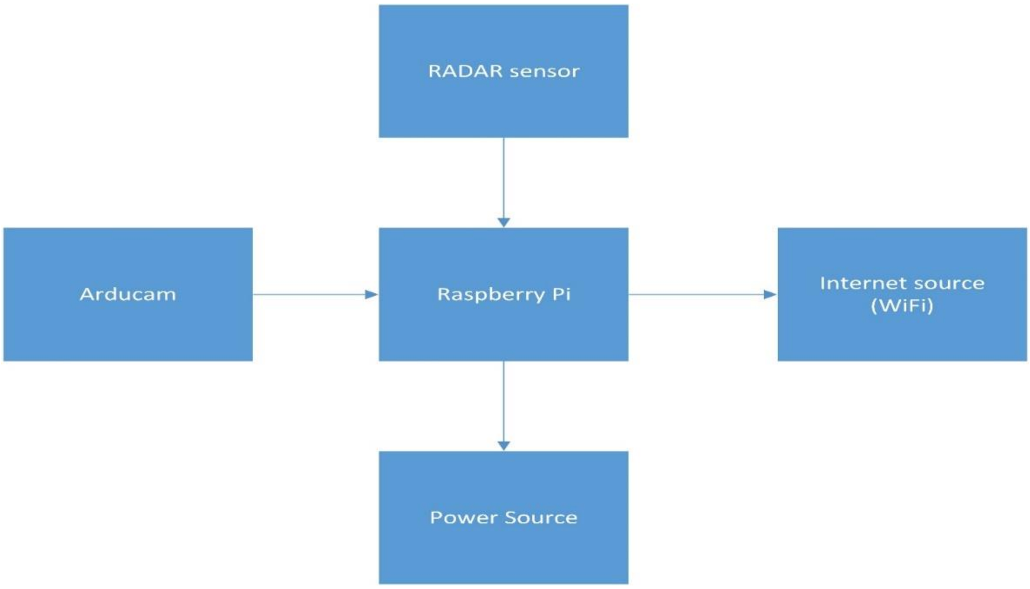

Figure 0.5: Hardware block diagram

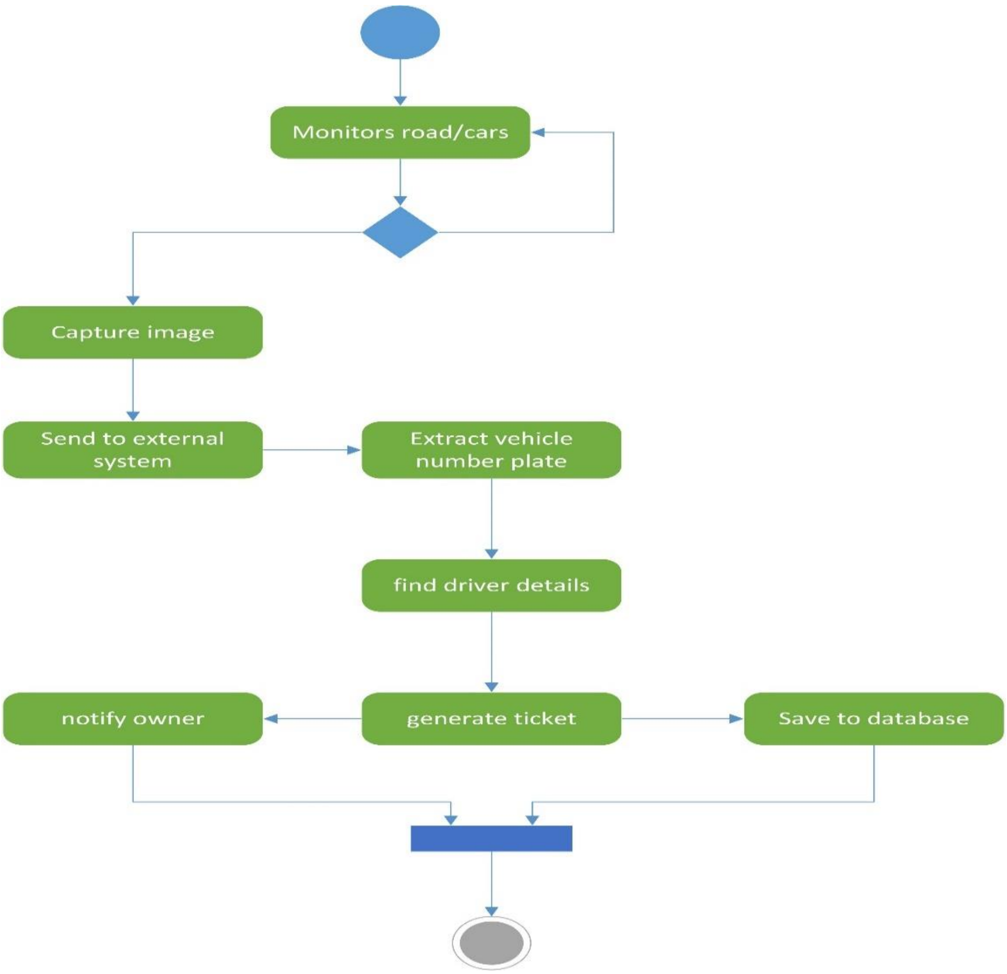

Figure 0.6: Activity Diagram

Figure 3 represents what happens within the system when it is operated. This points out the snapshot taken of the oncoming vehicle and stored on the microcontroller to be processed by the Automatic Number Plate Recognition program.

3.5 Alternative Considerations

An alternative to this design is to implement the system on a smartphone, but slightly different. This way, the motion of the phone within the car can be used for the speed measurement in the form of some up that uploads the information. This approach will make use of how quickly a phone's location changes on the surface of the road. This would not be viable since it doesn't guarantee trustworthy information. Such that a user could be dishonest by changing their needed vehicle details within the app and not get caught when speeding. Also, once a phone is switched off, speed measurement wouldn't be possible.

This method would also require that a driver downloads the app, which some people could dishonestly avoid, making the system ineffective. Since not all drivers have access to smartphones or personal computers, notifications will be sent to offenders via SMS in addition to emails.

Apart from the phone usage's negatives above, one advantage will be that the location of the speeding won't be dependent on location tag embedded on the system local (the hardcoded way). But would instead use the phone's inbuilt GPS tracker to provide location information while analyzing future changes that might have to be made on roads usually sped on by drivers. Some of these changes may include the introduction of speed ramps on a stretch of road to curb the incidence of overspeeding.

3.6 Automatic Number Plate Recognition Design

The Automatic Number Plate Recognition (ANPR) algorithm/program, also referred to as Automatic License Plate Recognition (ALPR), is a system that can be developed with different programming languages but commonly implemented in Python and C++. These two languages are widely used for computer vision using the OpenCV tool as a basis, as compared to MATLAB.

There are a few reasons that cut across both languages in addition to the advantages unique to each of these. One main reason is the OpenCV tool is a tool developed solely for OpenCV purposes. As a result of this, it has several packages that are easier to utilize when developing any computer vision application. Also, the Computer vision tools for both these languages are free, unlike MATLAB, which is itself pricey as a tool and requires priced license renewals for updates. OpenCV is slower when prototyping Computer vision programs, but it is more efficient in its execution than MATLAB [14].

What makes Python more favorable for such an application is its vast range of libraries available that can be utilized for computer vision joined in the scope of performing machine learning applications like image processing [15]. This helps computer vision developers save a lot more time, as a lot of what they wish to do has been handled in these libraries. C++, on the other hand, has fewer machine learning libraries, which is a disadvantage since computer vision applications require in-depth machine learning routines as it offers better methods for image processing and better object focus [15]. Also, Python has “more mature” web frameworks like Django and flask that allow for direct integration into any web application, as a functional component of the web service [16]. Since this system has a web application developed using the Flask framework, Python is the better option. In addition to this, the easy-to-read nature of the Python programming language code makes computer vision implementation equally easy. Apart from being able to implement Python codebase with relative ease, it also more comfortable to debug and additionally provides visualization libraries like matplotlib. Though one significant advantage C++ has over Python is that it is a compiled language and has shorter run times, thereforeto integrate a machine learning algorithms into this ANPR, the advantages Python has, including better machine learning support, makes Python a better option [14].

Chapter 4: Methodology

4.1 Hardware Setup

This part of the overall system is made up of a RADAR system, a camera for capturing images all which will be interfaced with a raspberry pi 3 model B+ microcontroller. This will serve as the central node for interaction within the system and the main hub for data processing and communication. The Raspberry Pi will be the component that transfers information, particularly the details extracted from the captured image of the vehicle to the webserver, and then forwarded to the database. It will also store all images captured by the system primarily but will back up these images to the cloud using WiFi. The Radar system does the real-time speed measurement of oncoming vehicles and interacts with the raspberry pi in the process.

4.2 Software

4.2.1 Web interface

The web interface provides administrative users with a dashboard to view all drivers that have been issued speeding tickets by the system. It allows a user to manually confirm a user's pending ticket payment each time the offender pays the fine. Through the web interface, the administrator(s) can also keep track of each driver's ticket record. All these were achieved through essential web technologies HTML and CSS for the appearance of the website and Flask micro web framework for the website's interaction with the established database.

React, a javascript library for web development, was also adopted for the user interface available to the system's administrator. This is useful in efficiently creating an interface for the real-time updates of some statistics related to speeding offences without the need to reload the webpages.

4.2.2 Backend

Automatic Number Plate Recognition (ANPR) program is written in Python and rendered on the raspberry pi microcontroller. Th will have a specific speed limit to which the real-time speed from the radar system is compared. When a "target" vehicle's speed exceeds this speed limit, the raspberry pi triggers the camera to take a snapshot of the car based on speed readings it is receiving from the radar sensor. The image is then processed by the ANPR program to extract the license plate details to prepare a custom speed ticket to be sent to the vehicle owner's details found to match the obtained license plate from the database.

Table 0.1: Administrative features and their priorities

Process	**Description**	**Priority**
Ticket Issue	This enables the administrator to manually send the notice for a speeding ticket to the registered owner whose details have been captured by the system.	High
Categorised Records Search	This enables the administrator to narrow down any necessary searches of a specific speeding offence or to verify payment of speeding ticket that has been issued. Some of the categories include location, dates, and names of offenders.	Low
Analytics	This will enable the administrator to visually keep track of all the "numbers" related to the speeding offences captured by the system	Medium
Ticket Cancellation	This will enable the administrator to manually change an offender's ticket status for each speeding ticket that has been issued for each vehicle owner	High

Admin addition	An already existing admin can create a new admin account. This is only possible when the admin has logged into the system	Low

4.3 Doppler Effect Phenomenon

According to Camuffo, the Doppler effect is defined as “the change in the observed frequency of an electromagnetic wave due to relative motion of the source and observer.” [17]. In this case, the source body is the monostatic radar sensor, whereas the observer is any moving vehicle. According to Chen *et al., m*ono-static for this application implies the device serves as both a transmitter and a receiver. [18]

The Radar, being the stationary body, will emit its specified frequency as the target vehicle introduces the relativity of the motion by its movement. This movement can either be towards the Radar or away from it. For this speed detection application, the vehicle’s movement towards the Radar is of interest. Specific unique characteristics make it possible to know the vehicle’s direction based on the reflected frequency. The Radar, together with the microcontroller, can be programmed to focus on the change resulting from motion towards it. One key characteristic that makes this detection is, as a body moves towards the source of the frequency, the difference in frequency, particularly the strength, is a positive one. Hence, all negative changes in frequency are simply ignored by the system.

The following diagram clearly depicts the nature of the reflected frequency concerning the direction of the target object’s motion

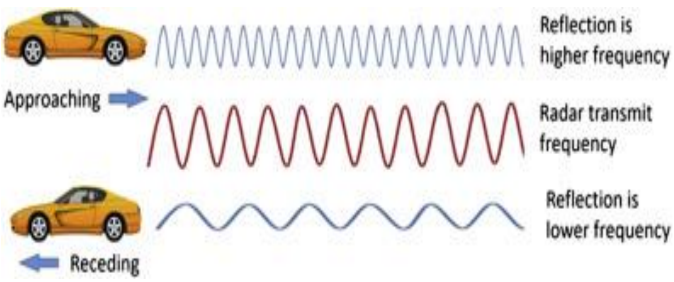

Figure 0.1: A visual representation of the doppler effect phenomenon [19]

A basic formula is applied, which expresses the relationship between the received frequency and the transmitted frequency. These two parameters are needed to estimate the speed the target object is moving at, towards the Radar.

Below is the formula for this relationship:

$$f = \left(\frac{c+v_T}{c+v_s}\right) f_0 = \left(1 + \frac{v_T - v_s}{c+v_s}\right) f_0 \qquad (1)$$

where c is the velocity of radio waves in the medium
v_T is the velocity of the receiver relative to the medium
v_s is the velocity of the source relative to the medium.

In finding the velocity of the moving target object, the following formula can be derived from the above one:

$$v_T = c\left(\frac{f}{f_0} - 1\right) \qquad (2)$$

Despite its usefulness in measure speed, the Doppler effect shift has a few limitations in successfully doing this. When measuring low speeds or low source frequencies, the shift is relatively small. More accurate results can be obtained when measuring the speed of fast-moving objects and higher frequencies that produce a proportionally, more significant change. However, since the purpose of the desired speed detection system is meant for vehicles moving at speeds fast enough to be considered dangerous, this limitation will have no significant effect on the system's efficiency. According to Forinash and Wisman, "a single frequency source is preferable as it allows the shift to more easily be determined" [20]. This is necessary not to have the receiver module of the Radar taking wrong or undesired readings. Another limitation is acceleration. The shift can be easily determined if the movement is at a constant velocity.

4.4 Data Preparation

A few images of parts of the vehicle that are not license plates were cropped out and used in the plate detection model of the system. Photographs of typical license plates were also used to train the model to teach it how it can identify a license plate in any image. This was done to enhance the number plate detection model.

The YOLOv3 object detection architecture, developed by Joseph Redmon, Ali Farhadi, Ross Girshick, and Santosh, was implemented in this training process. YOLO here stands for You Only Look Once. It's deeper neural networks makes it more accurate due to its ability to be developed and trained with large amounts of data. It also expresses object detection as regression problems rather than something that needs classifiers [21].

YOLO is based on convolutional neural networks, which, according to Suryani *et al.*, makes algorithms it is based on for text recognition perform better than "randomly initialized filters [22]. YOLO can predict various bounding boxes and "class probabilities for these boxes" and can train on full images, which directly optimizes its performance.

Some images of the alphabets and number digits, A-Z and 0-9, in different fonts were acquired and stored in folders sorted in alphabetical and numerical order. This was done to ensure the desired font type of the license plate could be easily detected, segmented, and its characters extracted to be applied by the system. The dataset consisted of characters popular on license plates, the United States, and on other objects.

4.5 Training The Model

When the image is obtained for text recognition and extraction, the image is, first of all, preprocessed. Processing is done to improve the appearance of the image by suppressing the noise present in the image [23]. As a result, one vital preprocessing technique vital for this application was the conversion of the image from colour to grayscale. A grayscale scale image helps simplify the algorithm to be used in the image processing and reduces computational requirements [24], helping to do an excellent job while limiting the computer resource usage

As a result, when the vehicle's image is passed through the ANPR, the license plate is first detected and segmented. This makes it easier to have a better view and focus of the object of interest and to easily identify the various characters inscribed on the surface. This segmentation is followed by the extraction of the characters identified. The system outputs a string of characters it has identified in terms of how closely the characters on the image match with those ones it was trained with.

The characters, most importantly, are classified by both numbers (digits) and alphabets to enable better recognition and determination of text. In the training process, a particular font type of character was used to aid in identifying the characters from the cropped-out license plate image correctly. This is because all the Ghanaian license plates issued by the DVLA share a unique font type with very minimal variations.

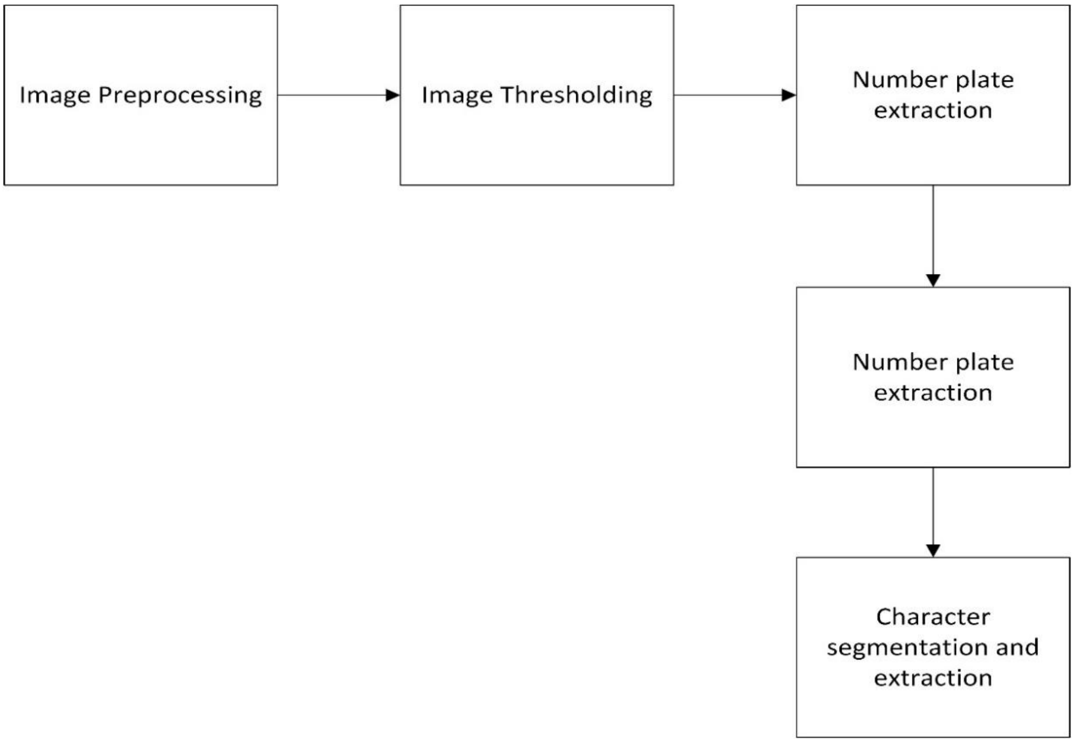

Figure 0.2: Architecture of ANPR

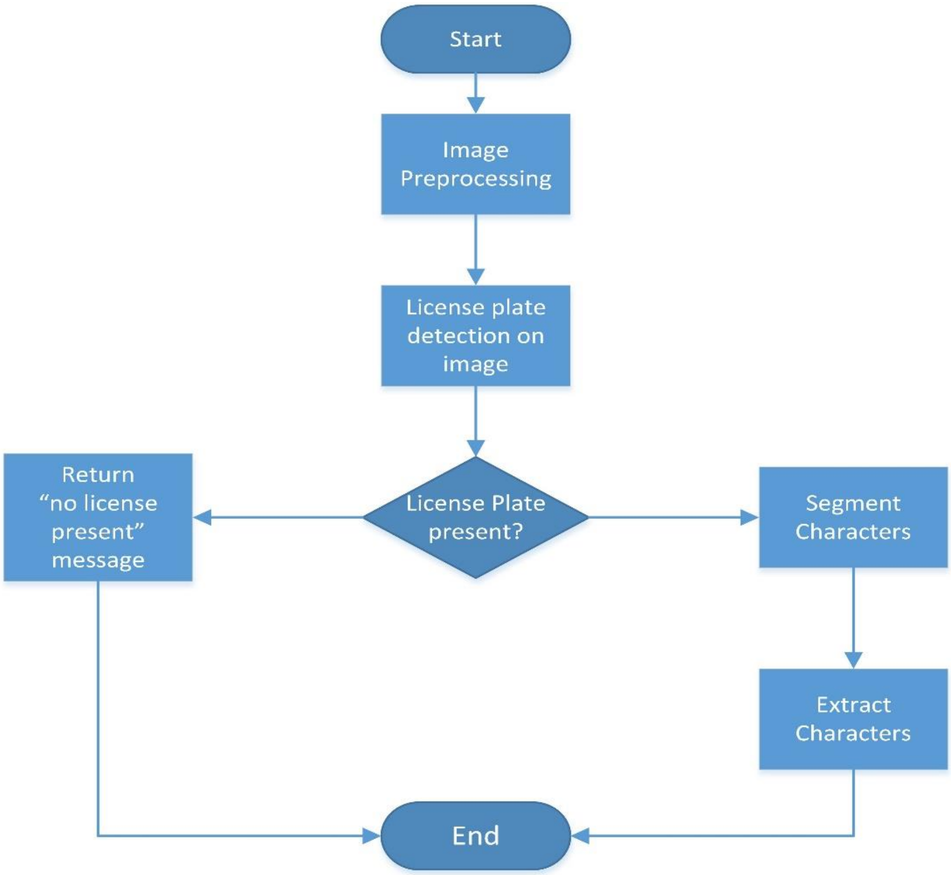

Figure 0.3: The Flow of the License plate’s detection and character extraction

Below are the training graphs from training the CNN-based character recognition model.

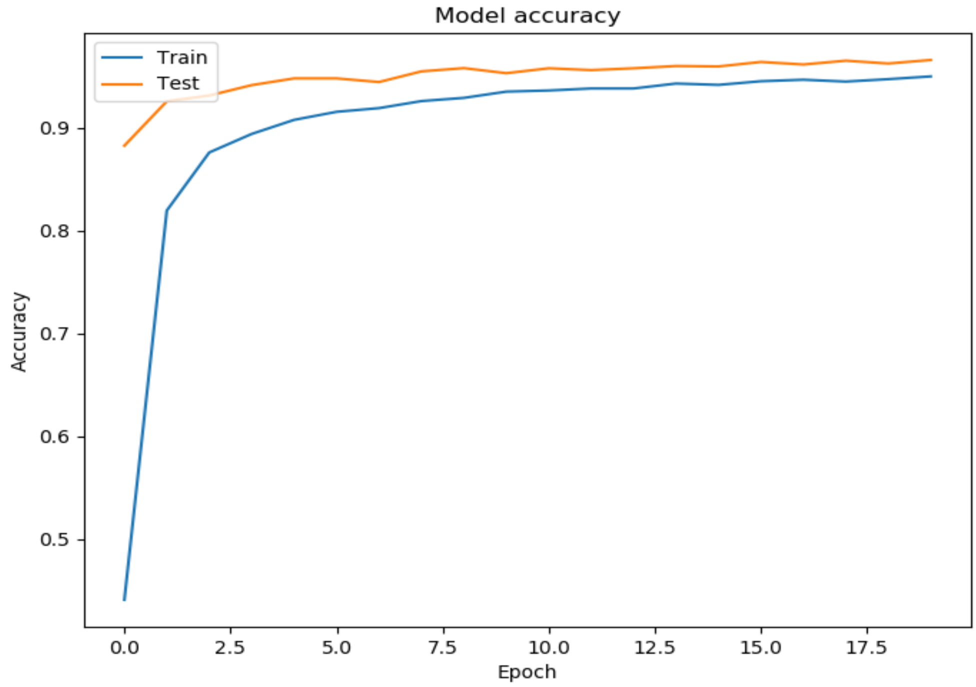

Figure 0.4: The accuracy of the character recognition model's training and testing

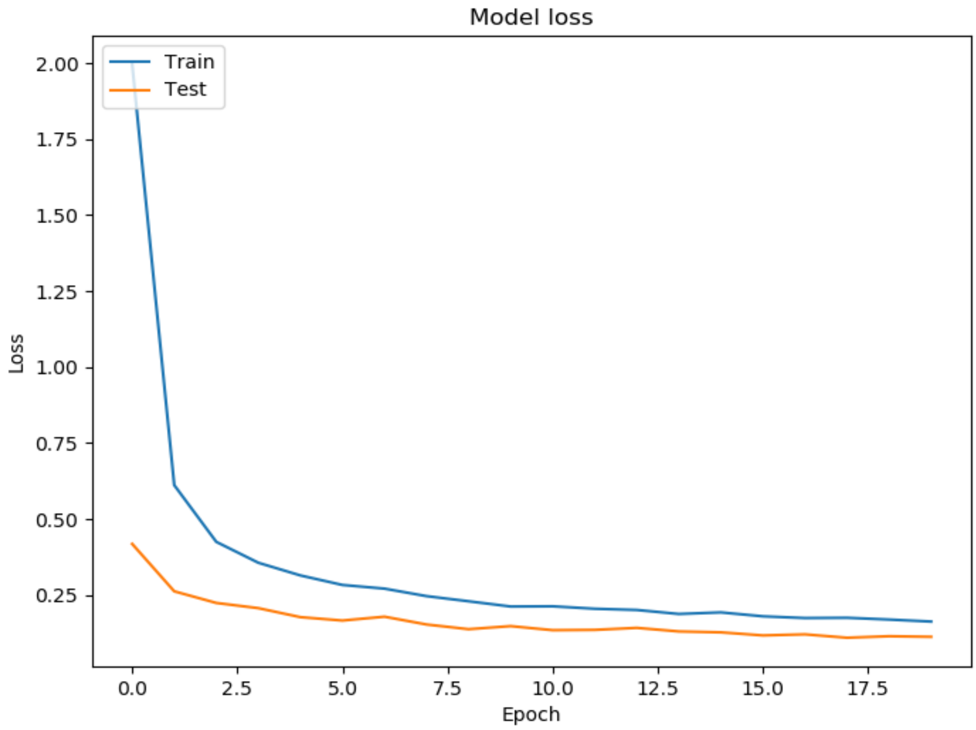

Figure 0.5: The loss of the Character Recognition Model's Training and Testing

From these diagrams, the accuracy of the model can be seen to have started oscillating after reaching its peak point. This can be attributed to the number of images used for the training. Also, the fact that since each character had about 1000 pictures, some fonts were similar but not the same. Hence making the model start to doubt its confidence in having figured out which character belongs to a particular font type, the model reads any time.

Despite the high accuracy level achieved above, the text recognition of the ANPR program developed the above way produced poor results. The confidence levels of some predicted characters were low as these texts had close to five other characters the algorithm saw to resemble the character of choice.

Another approach to train the character recognition model included the use of supervised learning classifiers. The best of these classifiers were settled upon to be integrated into the entire system. The first of these was the KneighboursClassifier in Python, which is based on the k-nearest neighbours vote. Another classifier which was also trained and tested with the dataset available was the Support Vector Machines (SVMs) which are able to, on its own, categorise new text or data, after being given some training data. Two different kernels of the SVM model were used, namely: Polynomial, Radial Basis Function (RBF) SVM and the Linear kernel SVM. Another classifier that was trained and tested was the random forest classifier [25]. A different dataset was used in training the KNeighbours classifier as well as the two kernel implementations of the SVM. This was because the data used to train the CNN model was subsequently noted to not have variations of noise present and varying orientations.

The supervised learning classifiers, rather than using YOLO for detecting the license plates, used template matching. Template matching compares a template of a license plate, with the different objects that appear in an image. This helps in efficiently localizing the license plate for character segmentation and recognition. It achieves this by feeding the model with the unique patterns license plates have, enabling license plates to be recognized in images captured in varying environments, including those with license plates attached to cars with minimal human intervention [26], [27].

Table 0.2: Comparison between the different classifiers explored

Classifier	**Validation Accuracy** %	**Testing Time (s)**	**Sets. of false predictions**
Linear SVM	97.647	0.351	Three
RBF SVM	94.706	0.280	Five
Poly SVM	93.823	0.279	Five
3-neighbour KNN	94.412	0.140	Five
4-neighbour KNN	95.294	0.148	Five
5-neighbour KNN	92.647	0.144	Seven
Random forest	95.882	0.929	Four

4.6 Implementation

The speed detection system will be set up alongside a stretch of road, preferably a highway between major cities in Ghana. This highway is assumed to have a speeding limit of 60km/h. It will be positioned such that its camera and radar components will face the direction care on its lane are coming from to have the desired positive speed readings based on the doppler effect, fed to the microcontroller by the Radar.

To have maximum efficiency, the Radar continuously emits its unique frequency towards oncoming vehicles. It simultaneously records the speed the target vehicle is moving at, which on a busy road is the closest vehicle. These speed values are continuously compared with the speeding limit. For the system to be fair, an allowance of ***+10km/h*** is set for any driver who exceeds the set speed limit. This would give room for drivers to adjust their speeds immediately they realise the speeding limit has been exceeded.

Once this allowance is exceeded, the camera is autonomously triggered to capture an image of the target vehicle with the license plate in view. Once this image is taken, it is stored in a folder based on the date on which it was captured. While the image is being stored on the raspberry pi, the ANPR program on the Raspberry Pi processes the just captured image to extract the license number of the vehicle. This is the setting for any unit of the system positioned along roads unbusy.

For bustling stretches of road, the images will be captured and collected in a single folder for the day. This means that all pictures captured on a particular date will share a single folder on the Pi. At a set time of the day, all images are processed by the ANPR to extract the license number of each image in the folder. This data is used as a key to acquire the needed details of the vehicle owner from the "owners" database and logged in a new database table. The reason for this implementation is attributed to the limited processing power of the raspberry, which might make it miss some speeders, assuming it was to process each image right after it is captured immediately.

Each extracted license number is checked against the database to acquire the needed details of the vehicle owner from the "owners" database and logged in a new database table

The administrator will then to manually issue the notice of all details concerning the speeding offence tickets to the respective owners. The system handles the notification process on its own as long as the "Generate" command has been clicked. Each offender receives alert, depending on the contact details used in their vehicle registration process.

The alert that is issued to the offender will either be an SMS, containing just the details of the offence as regular texts or an email with a pdf report attached. This pdf will be dynamically generated to contain the details of each driver's offence. The details include the location of the offence, the speed at which the image was captured, and the amount to be paid. The notice will be sent via SMS or email medium.

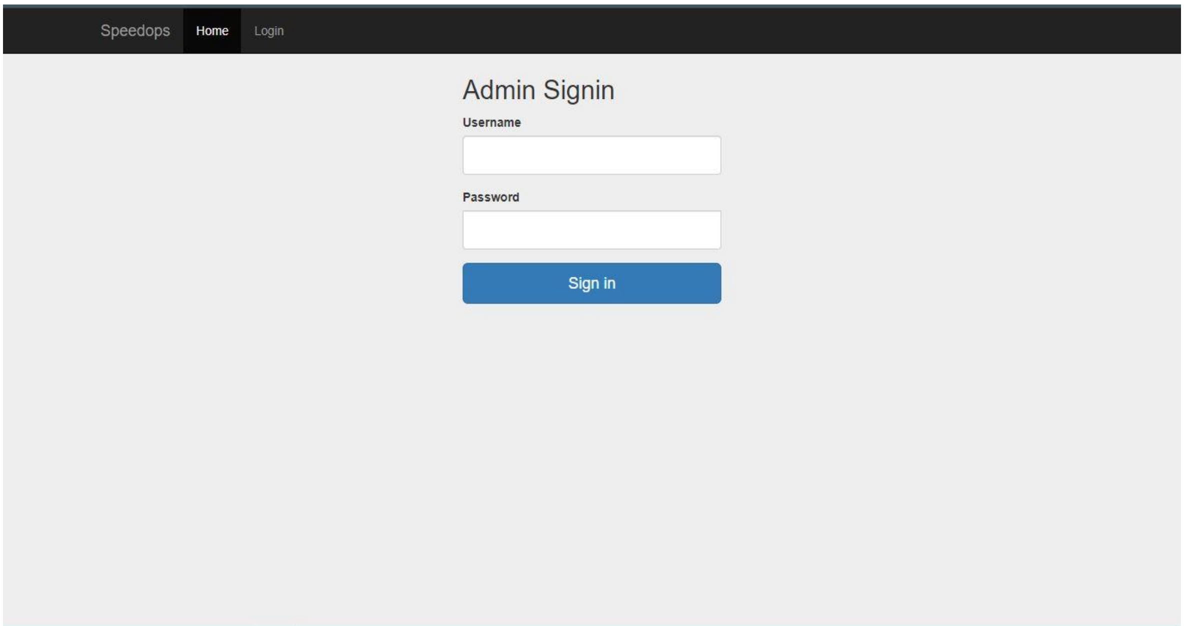

Figure 0.6: Admin sign-in page

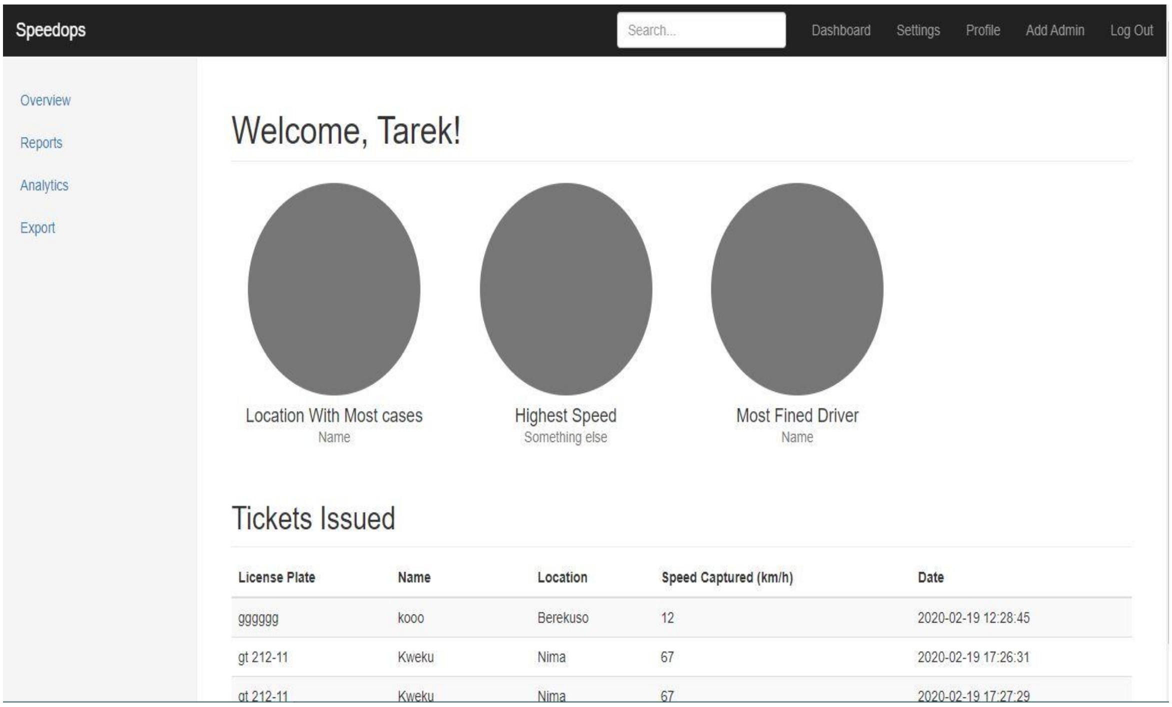

Figure 0.7: Dashboard

The above diagram displays speeding offences recorded by the system at various dummy locations in the country. In the situation where drivers make payments, the system's administrator can use a search bar to locate a specific driver or car owner in the list of offenders. This becomes helpful when the number of offences recorded becomes a lot.

From the table above, placeholders can be seen, which have all been given unique labels. In [28], maximum and average speed readings like one of the placeholders above can help to optimize traffic by using the information to reduce traffic rush on roads.

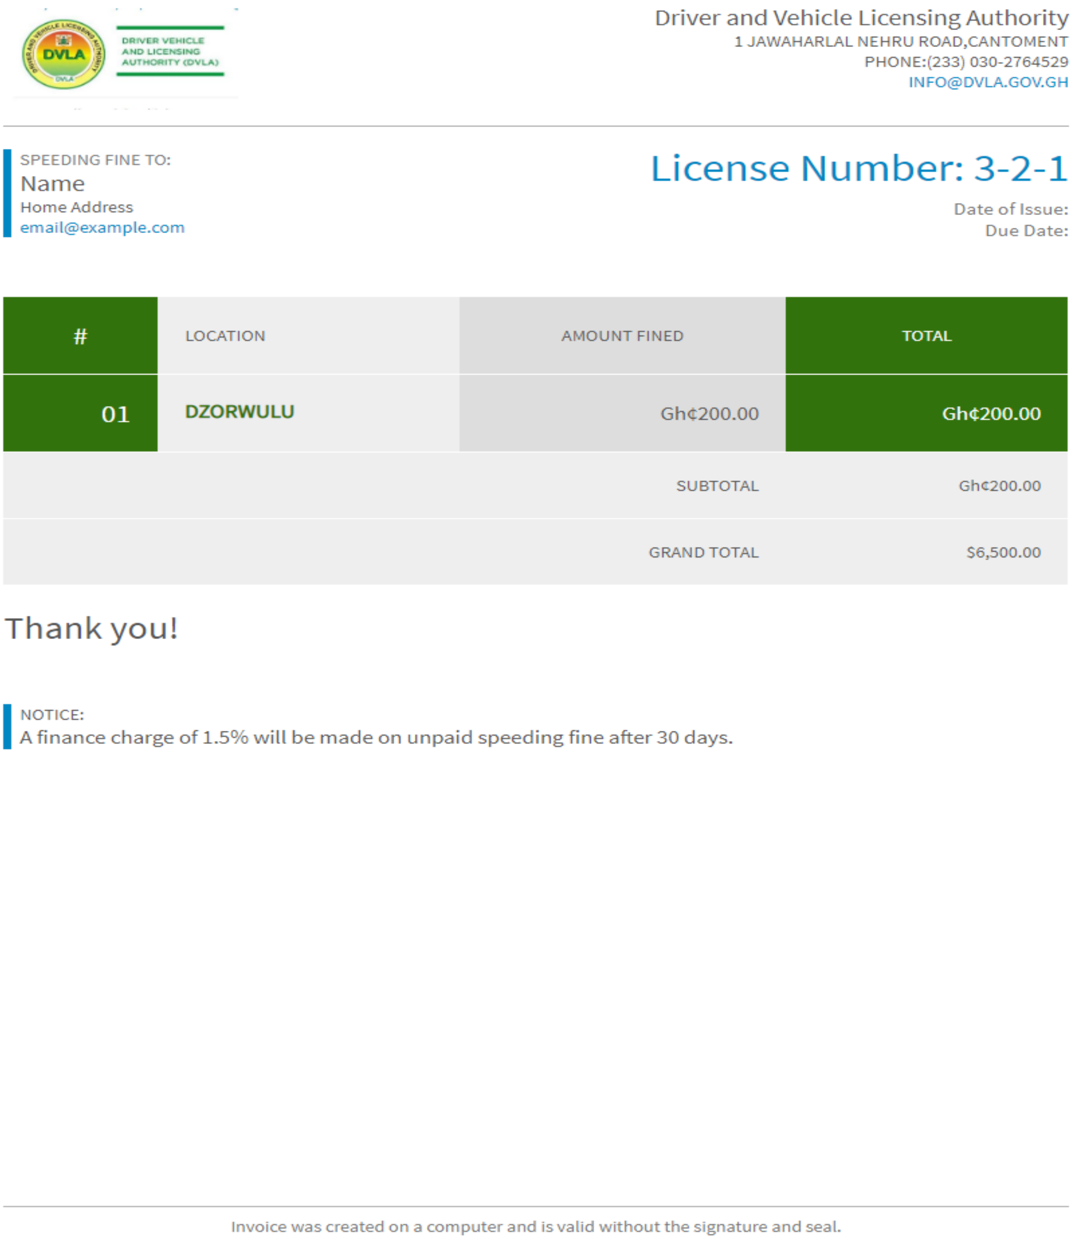

Driver and Vehicle Licensing Authority
1 JAWAHARLAL NEHRU ROAD,CANTOMENT
PHONE:(233) 030-2764529
INFO@DVLA.GOV.GH

SPEEDING FINE TO:
Name
Home Address
email@example.com

License Number: 3-2-1

Date of Issue:
Due Date:

#	LOCATION	AMOUNT FINED	TOTAL
01	DZORWULU	Gh¢200.00	Gh¢200.00
		SUBTOTAL	Gh¢200.00
		GRAND TOTAL	$6,500.00

Thank you!

NOTICE:
A finance charge of 1.5% will be made on unpaid speeding fine after 30 days.

Invoice was created on a computer and is valid without the signature and seal.

Figure 0.8: The speeding Fine report template

Chapter 5: Results and Discussion

The first phase of testing just involved the Automatic Number Plate Recognition section of the entire system. This did not include the verification side of the system, where the extracted license plate number is checked against a database containing the details of vehicle registration with their associated owners' details. Sample images of random cars with Ghanaian vehicle license plates attached were taken. The aim of this was to confirm that the system is at least able to carry out two basic but critical tasks successfully. That is, to detect the position of the license plate and to also identify the set of characters on the license plate, which is unique to every registered vehicle.

Using the Python programming language, the KNN algorithm, and the vehicle license plate dataset developed by Chris Dahms, the Automatic Number Plate Recognition program was designed and developed. Some already existing images of car license plates were downloaded from the Internet and run through the ANPR program to test the accuracy of the system developed. By using a phone, an image of a car parked on the Ashesi campus was taken using a smartphone's camera and run through the ANPR system developed.

5.1 Observations

Upon processing the image with the ANPR, some images had slightly different text(characters) extracted for each test run. That showed the imperfections of the training given to the system and the need to expand the training by including some other font type since not all license plates had the same font type. Also, hyphens will have to be included in the training set as they were all excluded by the algorithm.

Below is a result of one of the test cases:

Figure 0.1: Image of the vehicle with the license plate number "GC 2259-09" run through the trained Automatic License Plate Recognition program

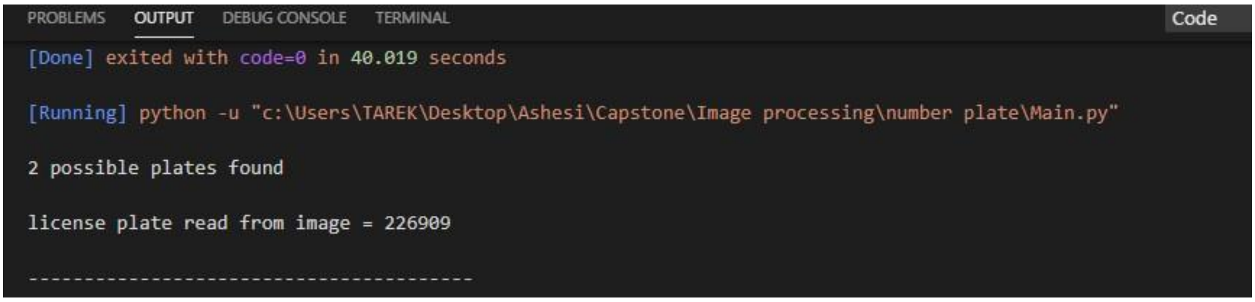

Figure 0.2: Result of the test with the license plate number "GC 2259-09

As seen above, the recognition is working to an extent but not perfectly since '5' was mistaken to be '6' by the text recognition algorithm.

.

Figure 0.3: Image of the vehicle with license plate number "Gx 1047-18" run through the trained License Plate Recognition program

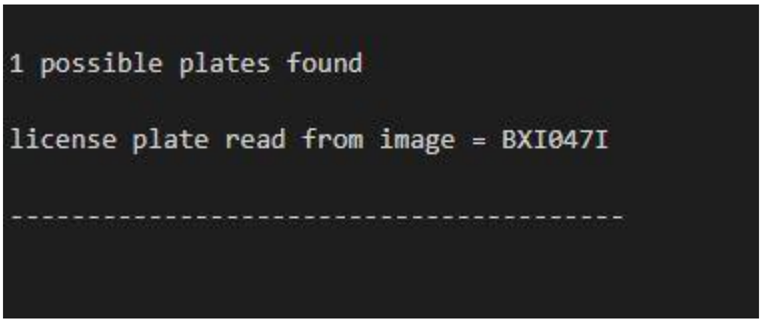

Figure 0.4: Result of the test with the license plate number "Gx 1047-18."

Just like the first test case, another character on the license plate was mixed up with a different style. Further research and training were carried out to improve upon the accuracy of the Automatic Number Plate Recognition program to get rid of the errors observed in the test cases above. In the course of this research, the classifiers mentioned in Chapter 5 were discovered, trained, and tested out.

Now using the final best performing classifier, that is the SVM model using the "linear: kernel for classification, new images were tested with the model to verify it's text recognition accuracy.

The first image is an image of a vehicle with a Nigerian license plate attached to it since the dataset mainly consists of Nigerian license plates font.

Figure 0.5: Image of a vehicle with a Nigerian license plate

Below is the output supplied by the model within the Pycharm console. The output can be seen in the red boundary in the upper-left corner of the image.

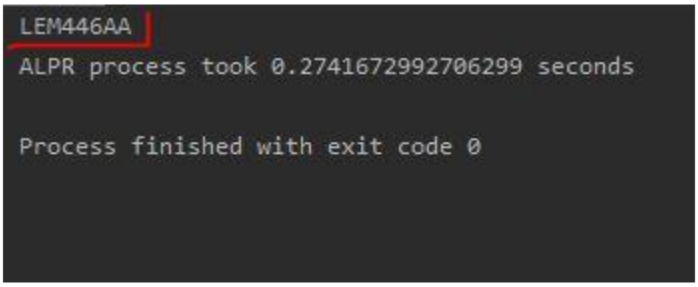

Figure 0.6: The successful output of the linear-SVM text recognition model

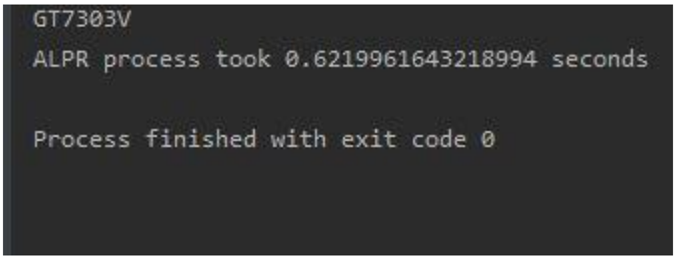

```
GT7303V
ALPR process took 0.6219961643218994 seconds

Process finished with exit code 0
```

Figure 0.7 The successful output of the linear-SVM text recognition model

Figure 0.8: Vehicle whose output can be seen above

```
  File "C:/Users/TAREK/PycharmProjects/ANPR_Naija/full.py", line 102, in <module>
    execute_ALPR()
  File "C:/Users/TAREK/PycharmProjects/ANPR_Naija/full.py", line 56, in execute_ALPR
    license_plate = license_plate_extract(plate_like_objects, pre_process)
  File "C:/Users/TAREK/PycharmProjects/ANPR_Naija/full.py", line 28, in license_plate_extract
    "Plate Localization", wx.OK | wx.ICON_ERROR)
wx._core.PyNoAppError: The wx.App object must be created first!

Process finished with exit code 1
```

Figure 0.9: Failed license plate localization for "GC 2259-09" from earlier testing

5.2 Discussion

The inability of the program to properly preprocess the image, detect and separate, the license plate tends to affect the resulting character segmentation and recognition processes, hence returning inaccurate results. This can be attributed to excess light surrounding the license plate, or even the light the license plate is reflecting. Another reason why some license plates may not be appropriately localized is that some are not perfect rectangles but also have a square-like shape. [29]

In testing the image of the vehicle with license plate number “GC 2259-09” again with this model, an issue was encountered with first detecting the number plate. This is evident in the fact that the outline of the license plate differs significantly from that of the Nigerian license plate images in the dataset used for the training.

This highlights issues that may be faced and will have to be dealt with if this is deployed into the real world. Such a problem is expected to crop up as many Ghanaian vehicles tend to customize license plate frames a lot

5.3 System Integration

The Radar continuously reads the speed any oncoming vehicle is moving at till that vehicle passes by. In the case where the car is overspeeding above the tolerance level coded in the system, the camera is triggered to capture the vehicle’s image with its license plate in view, and the specific speed at that moment is recorded.

On the Radar’s interaction with an oncoming vehicle, the speed recorded when the target vehicle’s image was captured is applied in a model designed to determine the amount to culprit should be fined. The resulting amount the owner is to be penalized with is stored temporarily in the program’s runtime.

Holistically, the ANPR program was set up as the first contact point of the system. Images in the specific image directory are processed by the ANPR program to extract the license number. This is imported and run in a “main.py” script, which as a function which returns license plate characters in the way they appear on the plate. The returned details are passed to another function, which verifies if the license plate exists in the database of registered vehicles. If that license plate is found, the owner’s name, home address, and email/phone number associated with that license plate.

These details, except for the contact details, but together with the speed of the moving vehicle and the speeding fine to be issued, are used to generate a speeding fine report dynamically. This report is emailed to owners whose email is present in the database. For owners with only phone numbers, these details are sent to the owner of the vehicle via SMS. After this, DVLA simply awaits payment, which simply cannot be ignored since a driver will not be issued any document renewals such as the roadworthy sticker and driver’s license.

Chapter 6: Conclusions, Limitations and Future works

6.1 Conclusions and Limitations

During the development of this system for detecting overspeeding on roads and issuing fines to guilty parties, it was discovered that certain parts of the final system were not as accurate as intended. One such component is the Automatic Number Plate Recognition module, which is a vital aspect of the program. In the course of the research and development, when the ANPR in the United States and some European countries were studied. It was realized that developers of these ANPR programs had access to large datasets of license plates in the country. Unfortunately, Ghana has no such available datasets, and this made it difficult for a well-functioning ANPR system to be developed.

The model for the Automatic Number Plate Recognition and its text extraction component needed to be trained with data, in the form of pictures of characters in license plates issued by the Driver and Vehicle Licensing Agency. As this data was unavailable to train and the test model, the need arose for several pictures of license plates of Ghanaian vehicles to be taken, enough to train the model well enough to achieve a reasonably good level of accuracy. The reason why license plate images unique to Ghana alone are needed is attributed to the fact that different font types have been used by various countries to print the characters on national license plates. Also though around 10 sample license plates on vehicles were found scattered on the Internet, these were very old license plates whose fonts differ from those that currently make up the majority of license plates being used.

Following this, another setback was faced due to the spread of the novel coronavirus right before the images could be taken of the license plates of cars that could be captured. The severity of the epidemic caused the government to call for a lockdown, restricting the movement of a vast majority of people, vehicles, and goods. The government urged citizens of the country to stay home to reduce the spread of the virus and to save lives. The only unfortunate thing about this was fewer images could be captured with a camera and used for training purposes. As stated earlier in this paper, a few other ones were randomly downloaded from the Internet and were used to test the model.

6.2 Future works

Moving forward, any license number that is identified to not exist in the nationwide database of all registered car owners can be logged into a separate database, and an alert created raising awareness for a vehicle with fake or forged license plates. This can aid the police in investigating such situations since they are real activities people engage in, especially with stolen or illegal vehicles, and better apprehend culprits.

Concerning stolen vehicles, the system can be extended to have a live continuous video stream of roads instead of the camera being used only to capture images at specific times. For this to be effective, any civilian can report the details of the stolen vehicle to the DVLA officials. The most useful information needed will be the license plate details. With that, a command can be issued to all the speed detection units to look out for the specified license plate, and to return an alert whenever spotted. The units will then continuously scan the live video feeds for the license plates. Once the target license plate is identified, an image can be taken and stored in a unique directory, as the DVLA and police are alerted of the finding.

This system can also be trained in the future to differentiate between the regular official license plates and the DV/DP Trade Licence Plates. According to the DVLA, these plates are issued to vehicles that are unregistered but which, for various reasons, have to ply the roads to transport their cars from the port to their dealership in search of a potential purchaser. This can be useful in identifying drivers who ply the streets with such plates beyond the time stipulated by the DVLA from when these were issued, which is usually at the ports.

The idea of identifying and tracking stolen vehicles and vehicles with DP/DV license plates can be implemented more effectively and more successfully when video streams are used. This is to say that it will be better to identify the license plate details in real-time rather than merely taking snapshots at some time intervals. This will require microcontrollers with more excellent processing capabilities, larger storage capacities as well as very stable internet connectivity for improved intercommunication in the network.

The location embedded in each speed detection unit will give investigators a fair idea of areas to start or focus their searches within. In making the location tracking feature more efficient, a GPS module can be installed and program in the future stages of development that will eliminate the need to hard code each unit's locations. This way, whenever the need arises to repair or replace a unit, the GPS module will automatically pick up the new site it has been positioned.

References

[1] WHO, "Road Traffic Injuries," *WHO*, 07-Dec-2018. [Online]. Available: https://www.who.int/news-room/fact-sheets/detail/road-traffic-injuries.

[2] WorldBank, "Road Deaths and Injuries Hold Back Economic Growth in Developing Countries," *World Bank*, 09-Jan-2018. [Online]. Available: https://www.worldbank.org/en/news/press-release/2018/01/09/road-deaths-and-injuries-hold-back-economic-growth-in-developing-countries.

[3] G ChenChen, G., and Warburton, R. (2006). Do speed cameras produce net benefits? Evidence from British Columbia, Canada. Journal of Policy Analysis and Management, 25(3), pp.661-678.

[4] E. Bonnet, L. Lechat, and V. Ridde, "What interventions are required to reduce road traffic injuries in Africa? A scoping review of the literature," *PLOS ONE*, vol. 13, no. 11, p. e0208195, Nov. 2018.

[5] D. Bhardwaj and S. Mahajan, "Review Paper on Automated Number Plate Recognition Techniques," 2015.

[6] "Annual Distribution of Fatalities by Road Environment - Road Accidents and Casualties in Ghana | Ghana Open Data Initiative," *Ghana Open Data Initiative*. [Online]. Available: https://data.gov.gh/dataset/road-accidents-and-casualties-ghana/resource/7da3d348-40a7-4e94-bef0-65e2c1e048a5#{}.

[7] "Annual Distribution of Traffic Fatalities by Region - Road Accidents and Casualties in Ghana | Ghana Open Data Initiative," *Ghana Open Data Initiative*. [Online]. Available: https://data.gov.gh/dataset/road-accidents-and-casualties-ghana/resource/8eef50bc-d0e9-435d-870b-360314db64f3

[8] P. Pilkington and S. Kinra, "Effectiveness of speed cameras in preventing road traffic collisions and related casualties: a systematic review," BMJ Online, 2019. Available: https://www.researchgate.net/publication/8077124_Effectiveness_of_speed_cameras_in_preventing_road_traffic_collisions_and_related_casualties_Systematic_review.

[9] D. Graham, C. Naik, E. McCoy, and H. Li, "Do speed cameras reduce road traffic collisions?", *PLOS ONE*, vol. 14, no. 9, p. e0221267, 2019. Available: 10.1371/journal.pone.0221267.

[10] R. Tay, "Speed Cameras: Improving Safety or Raising Revenue?", Journal of Transport Economics and Policy, Vol. 44, No. 2, pp. 247-257

[11] O. Ibrahim, H. ElGendy, and A. ElShafee, "Speed Detection Camera System using Image Processing Techniques on Video Streams"," International Journal of Computer and Electrical Engineering, pp. 771-778, 2011. Available: 10.7763/ijcee.2011.v3.418.

[12] R. Balasubramani, "A Review on Vehicle Speed Detection using Image Processing," Researchgate, 2011. Available: https://www.researchgate.net/publication/327337312_A_Review_on_Vehicle_Speed_Detection_using_Image_Processing.

[13] "SQLAlchemy - The Database Toolkit for Python," *SQLAlchemy*. [Online]. Available: https://www.sqlalchemy.org/.

[14] S. Mallick, "OpenCV (C++ vs Python) vs MATLAB for Computer Vision | Learn OpenCV." https://www.learnopencv.com/opencv-c-vs-python-vs-matlab-for-computer-vision/.

[15] "Machine Learning in Computer Vision," *Full Scale*, May 08, 2019. https://fullscale.io/machine-learning-computer-vision/.

[16] "Advantages of Using Python for Computer Vision," *Full Scale*, May 02, 2019. https://fullscale.io/advantages-using-python-computer-vision/.

[17] D. Camuffo, "Chapter 13 - Measuring Wind and Indoor Air Motions," in *Microclimate for Cultural Heritage (Second Edition)*, D. Camuffo, Ed. Boston: Elsevier, 2014, pp. 471–491.

[18] V. C. Chen, F. Li, S.-S. Ho, and H. Wechsler, "Micro-Doppler effect in radar: phenomenon, model, and simulation study," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 42, no. 1, pp. 2–21, Jan. 2006, DOI: 10.1109/TAES.2006.1603402

[19] Y. Jia, L. Guo, and X. Wang, "4 - Real-time control systems," in *Transportation Cyber-Physical Systems*, L. Deka and M. Chowdhury, Eds. Elsevier, 2018, pp. 81–113

[20] K. Forinash and R. Wisman, "Doppler Measurement of Speed," in Indiana University SE

[21] J. Redmon and A. Farhadi, "YOLO9000: Better, faster, stronger," *Proc. - 30th IEEE Conf. Comput. Vis. Pattern Recognition, CVPR, 2017*, vol. 2017-Janua, pp. 6517–6525, 2017, DOI: 10.1109/CVPR.2017.690.

[22] D. Suryani, P. Doetsch, and H. Ney, "On the Benefits of Convolutional Neural Network Combinations in Offline Handwriting Recognition," in *2016 15th International Conference on Frontiers in Handwriting Recognition (ICFHR)*, Oct. 2016, pp. 193–198, DOI: 10.1109/ICFHR.2016.0046.

[23] M. Sonka, V. Hlavac, and R. Boyle, "Image preprocessing," in *Image Processing, Analysis and Machine Vision*, M. Sonka, V. Hlavac, and R. Boyle, Eds. Boston, MA: Springer US, 1993, pp. 56–111.

[24] C. Kanan and G. W. Cottrell, "Color-to-Grayscale: Does the Method Matter in Image Recognition?," *PLOS ONE*, vol. 7, no. 1, p. e29740, Jan. 2012, DOI: 10.1371/journal.pone.0029740.

[25] "1. Supervised learning — sci-kit-learn 0.22.2 documentation," *sci-kit-learn*. https://scikit-learn.org/stable/supervised_learning.html#supervised-learning (accessed May 11, 2020).

[26] P. R. S. Swaroop and N. Sharma, "An Overview of Various Template Matching Methodologies in Image Processing," 2016, DOI: 10.5120/ijca2016912165.

[27] R. Brunelli, *Template Matching Techniques in Computer Vision: Theory and Practice*. John Wiley & Sons, 2009

[28] A. B. H.g, Arpitha, B. S. Barkuru, Hemashree.J, and D. S. M. Badhusha, "Review Paper on Automated Number Plate Recognition Techniques," *International Journal of Engineering Research & Technology*, Jul. 2018, Accessed: May 11, 2020. [Online]. Available: https://www.ijert.org/research/review-paper-on-automated-number-plate-recognition-techniques-IJERTCONV6IS15086.pdf, https://www.ijert.org/review-paper-on-automated-number-plate-recognition-techniques.

[29] C. S. Patel, D. Shah, and A. Patel, "Automatic Number Plate Recognition System (ANPR): A Survey," 2013, DOI: 10.5120/11871-7665.

[30] DVLA, "DVLA | News." http://dvla.gov.gh/singlenews.php?news=af3e133428b9e25c55bc59fe534248e6a0c0f17b

[31]Algorithm borrowed from https://github.com/femioladeji/License-Plate-Recognition-Nigerian-vehicles

Appendix

Source code: https://github.com/Kpakpo/SpeedOps_ANPR