UNIVERSITY OF TWENTE.

**Faculty of Electrical Engineering,
Mathematics & Computer Science**

A Low-Budget, End-To-End Warning System for Bicycles using Monocular Vision and Vibrating Handlebars

Marrit Schellekens

M.Sc. Thesis assignment
Augustus 2021

Supervisor:
dr.ing. Y. Huang
Graduation committee:
ir. E. Molenkamp
prof.dr. P.J.M. Havinga (Paul)

Pervasive Sensing Group
Faculty of Electrical Engineering,
Mathematics and Computer Science
University of Twente
P.O. Box 217
7500 AE Enschede
The Netherlands

Abstract

Awareness of surrounding traffic is crucial for safety. Nowadays warning and information technology in cars is often used to supplement people's awareness of their surroundings. This thesis presents such a warning system, but designed for bicycles instead. An end-to-end, low budget system was developed, that warns cyclists of upcoming traffic. The only sensor input to the system is a monocular camera, and the processing is done on a cheap 100$ computer. For processing the choice was made to use the Yolov4 neural network, in combination with custom algorithms for tracking and finding the direction of travel. Additionally, a survey was performed to explore people's acceptance of such a system, and find the necessary performance such a system must have. The survey indicated such a system must have warning times of at least 3-4 seconds, which in practice was only reached less in less than 50% of cases. Additionally parked cars often created false warnings. The main obstacle in having a faster warning time or less false positives was the performance of the tracking algorithm.

List of acronyms

ADAS Advanced driver assistance systems

YOLO 'You only look once' object detection

HUD head-up display

FOV field of view

PWM pulse-width modulation

csp scaling cross stage partial network

List of Figures

Contents

Chapter 1

Introduction

Cycling has many advantages. It keeps people healthy, reducing healthcare costs and sickness absence [2]–[5]. It prevents emissions of CO2 and reduces air and noise pollution [6]–[8]. It also leads to more efficient use of infrastructure (i.e. less parking space and traffic jams) [9]. In order to encourage more people to take up cycling, and to reduce injuries and deaths, improving cycling safety is important. Especially since a perceived lack of safety has been found to be a deterrent for cycling [10], [11]. The Netherlands is already performing very well in this area. Together with Denmark, it is the safest country to cycle with the lowest amounts of deaths per km biked [12]. This is in fact a surprisingly recent evolution. Since the 1970s, the Netherlands has achieved an 80% reduction in cyclists' fatality, in spite of elderly people becoming a much bigger percentage of all cyclists [13]. However, in the last two decades, these numbers have stabilized, while car fatalities are still decreasing [14]. This while the Netherlands still have big ambitions regarding cycling safety and promotion. The Dutch government has as goal to halve all traffic deaths in 2030, and even aim for 0 deaths in 2050 [15]. With 203 cycling deaths in 2019, clearly cycling safety has to be improved further to reach these goals [14]. The previous safety improvements have mostly been achieved by focusing on safe infrastructure and traffic regulations in favor of the cyclist [13]. This shows a clear difference with cars, where technological advances regarding the car itself (such as the inclusion of seat belts and airbags) have played a big part in the increased safety the last few decades. To improve biking safety, it might be good to draw inspiration from cars and look at the bicycle itself, instead of the environment around it.

While airbags and seat belts do not naturally translate to a cycling context, a new generation of technology provides more opportunities. Advanced driver assistance systems (ADAS) are already commonplace in cars, such as automatic lighting, parking assistance and cruise control. More advanced technologies are also migrating from high-end cars to the cheaper models, such as forward collision warning and autonomous emergency braking systems. What defines an ADAS is that it has access to outside information through sensors in the car, and uses this information to aid the user. On bikes, there is only one ADAS commercially available. This is the Garmin Varia. It is a warning system that warns bikers of approaching traffic from behind. Such rear-warning systems might become more important with quiet electric vehicles becoming more mainstream.

The e-bike has been steadily increasing in popularity. In 2019 a fifth of all bicycle trips in the Netherlands were done by e-bike [16] and half (45%) of all people aged 65-75 owned an e-bike [17]. The recent Corona pandemic seems to be giving an extra boost to this already growing industry [18]. The current high rise in e-bike numbers consists mostly out of elderly users who want to keep cycling, since the age group 65+ is responsible for half of all kilometers ridden on an e-bike [19]. A second potential future market is people using the e-bike for their commute instead of a car. The government is actively creating more opportunities to lease a bike from your company, and several municipalities offer subsidies to promote e-bikes as a commute alternative [20]. The current rise of e-bikes could be the perfect opportunity for bringing smart technology to the bike, as has already been done in cars. The e-bikes battery can supply power for smart electronics. The e-bike also has a higher selling price than a conventional bike, meaning that smart electronics build into the bike will have a relatively smaller impact on the price.

1.1 Problem Statement

Despite the potential benefit for improving cycling safety, very little academic work has been done regarding incorporating an ADAS on a bicycle. One cannot simply copy the ADAS as is from the car to the bicycle, since the bicycle has its own set of unique issues. First of all there is the power consumption issue. Although the rise of e-bikes reduces the severity of the issue, power consumption still is a major limit on introducing technology to the bicycle. Coupled to this is the issue of weight. Current e-bikes are still quite heavy, while bicycles ideally should be easy to move around and transport. Because of this limit one cannot simply add more battery capacity. There is also the matter of cost. Bicycles are much cheaper than cars. Expensive technology will be a proportionally smaller part of the cost for cars than for bicycles. One cannot simply put a high-end processor in a bicycle, because the cost increase would likely be too high. Another difference between cars and bicycles is the choice of output type. Cars have their own private ‘audio space‘, making sound a natural choice. But this choice might not transfer well to the cycling environment.

The three scientific papers that have worked on ADAS for bicycles mostly focus on the technological requirements [21]–[23]. The question of how and when to provide this information to the user hasn't been investigated at all. The typical output methods for a ADAS in a car are often not appropriate for the bicycle since the environment of the bicycle is distinctly different from the car. In a car the enclosed environment makes auditory feedback a much more straightforward solution. In contrast, bicycles are open to the world. Noises can be heard everywhere. Sound from the environment might cover up the sound of the ADAS. Additionally, the sound from the ADAS can be heard by everyone around the bicycle, possibly leading to frustration and confusion. Many cyclists also listen to music on their earphones, making it harder to hear the sounds from the ADAS. Running the sound through their headphones could be a possible solution, but cycling while listening to music is also dangerous ([24], [25]) and should probably not be encouraged. Finally, most current users of the e-bike will be elderly, and might suffer from hearing impairment. It is this target group that could especially benefit from a warning system since they cannot hear the cars approaching. Making the warning also auditory is not helpful for this target group. The solution used most commonly for bicycles is by displaying the output on a small screen attached to the handlebar. Although more functional, this also comes with its own problems. Looking at the screen could distract from traffic, and again could be harder to read for the elderly. Additionally, many screens can be hard to read during sunny weather.

1.2 Goal(s) of the assignment/Research question(s)

The proposed project is to implement an ADAS that warns cyclists of passing traffic from the rear. It will be a monocular camera system, inspired by the Hindsight paper (which will be discussed later on). However unlike the Hindsight prototype it will run on limited embedded hardware with limited power consumption. This will give an indication whether the mono-camera approach using neural networks is feasible in a resource-constrained environment. Output will be done through vibrating handlebars. It will also be tested in real traffic situations, and a detailed analysis will be done on the performance of such a system. Additionally, the desired performance of the system and potential acceptance will be investigated through a survey.

Chapter 2

State of the Art

An ADAS needs three different steps to function - acquiring the sensor data, processing the sensor data, and providing output to the user. Most papers focus either on the first two steps, or only on the last one. Therefore the choice was made to have a similar structure while discussing these papers. First the work done on acquiring and processing the data will be discussed, and then the work on outputting information to cyclists. Finally, the findings of both sections will shortly be summarized.

2.1 Sensors and processing

Using sensors to aid and enhance the driving experience is already a very common practice in the automotive industry. Parking aid technology is very common, and advanced emergency braking systems, which automatically brake when a collision is predicted, will become obligatory for new cars in the EU and Japan from 2022 [26]. It is safe to say that in the automotive world ADAS are a mature and common technology.

While electronics have become commonplace in car (to the disgruntlement of some car enthusiasts, who derisively refer to modern cars as 'driving computers'), the same cannot be said for bicycles. The current bicycle is in almost all aspects similar to the bicycle of more than a 100 years ago, as can be seen in an advertisement from 1897 (figure 2.1). The bicycle has stayed a mostly mechanical affair avoiding the impact of computerization which has changed so many other aspects of our life in a fundamental way. However, in recent years several companies have tried to bring smart technology to the biking context. One such product is VanMoof's smart bike. It has built in anti-theft alarms, automatic gears and will track the bike through GPS so it can be located when stolen. Another common product are the various 'smart helmets' that have popped up ([27], [28] and [29]).

Most of these commercial products have only limited access to information from the outside world. This means the products cannot respond flexibly to, or inform the user about the environment around them. The exception is the Garmin Varia, which warns the biker of upcoming traffic using radar technology. However, In the unique biking context of the Netherlands where many cycling lanes are separated from cars, a lot of the warnings of the Varia are not useful since the cyclist and car are not on the same road. A user of the Garmin Varia in the Netherlands noticed that from the 27 cars the system picked up, only 10 were actually on the same road as him [30].

In [21] they developed a driving assistance system for bikes. They used a radar-camera fusion system. In addition, they used an accelerometer to determine the bike's speed, and a camera facing the biker to determine where the biker was looking. 'You only look once' object detection (YOLO) was used to classify the camera objects, while MATLAB multi object tracking was used for the sensor fusion. They found that complex convolutional networks, as are used in object detection for cars, draw too much power for a bike. To work around this, they created two separate modes, an accurate one with high power needs, and a less accurate one with lower power needs. Depending on the risk-level of the situation, different sections of the view are assigned to a particular mode. The risk-level of a situation is determined based on current sensor data and navigational data. In [22] they also used two different modes, but in that case it was regarding the radar settings. In *rural mode* they used radar with a narrow beam to detect other road users earlier, while in *city mode* they used a wide radar to also recognize danger from the side. In [23] A 360-degree video camera was attached to a helmet. No radar was used. Classification was done using the YOLO v2 framework. The speed of an approaching vehicle was approximated by looking at the increase or decrease of the bounding box around a vehicle, and the distance is approximated by using assumptions on the typical size and relative velocity.

2.2 Feedback to bikers

The warning system developed by Garmin warns the user either visually on an led-matrix or screen (compatible with many bike-computers and can be developed for by third parties), or through audio. Both of these outputs have limitations. Sound might be hard to hear in the open context of a bike and it could pose problems for elderly with hearing disabilities. Screens on the other hand can distract from the traffic around the cyclist, are often hard to read in sunlight and again could pose problems for elderly with vision issues. To bring the visual information closer to the natural line-of-sight of a biker, a possible solution could be using projection on the road or on a transparent screen in front of the bike. The projection could inform the biker, but it could also aid in providing information to others around you as is done with Beryl's laserlight [31]. Their product projects a green bike in front of the bike, so that other traffic participants are aware of the bike even when it is in their blind spot. In [32] they projected a map on the road for navigation, as well as on a plastic screen in front of the bike, creating so a head-up display (HUD). They found that the HUD was preferred over the road projection because it was closer to the natural line of sight. A more elegant approach would be a HUD integrated into the helmet, such as the skully motorcycle helmet [33]. However, as the technology exists today, it is not suitable for the cycling context since it requires a full helmet, while most bikers in the Netherlands don't even wear a standard biking helmet. It could be revisited when AR glasses become more mature. Another method would be peripheral lights, as presented in [34]. However, this suffers from the same downside as the heads up display making it less suitable to biking.

A completely novel manner of alerting the user is through sonification. Relevant objects outside of the field of view, are transformed into audio cues. This was used in the Hindsight project to warn bikers of oncoming vehicles [23]. Bone conduction headphones were used, so that normal hearing was preserved.

Figure 2.1: Old advertisement dating from 1897, showing how little the bicycle has changed over the last hundred years [1]

Another approach could be to use wearables that provide haptic feedback to the user. This could be implemented in the smartwatch, where the intensity of the vibrations could slowly be increased. The advantage of this is that it is a piece of hardware many consumers already posses. However, the disadvantage is that it can be hard to represent complex or multiple sources of information, because only the intensity and frequency of vibration can be used. In most situations it would be a welcome addition to have an added dimension by having haptic feedback on both the left and right sides. This could be done with multiple wearables, such as developed in [35].

However, another solution would be to provide tactile feedback through the handlebars. This would prevent users from having to worry about bringing an extra accessory with them. In [36] a navigation system was developed that relied on giving directions through vibration motors inside the handlebars. This concept has also been implemented in a commercial application called smrtGRiPs [37].

2.3 Summary and choice of System input and output

Only limited work has been done on developing an ADAS suitable for bikes. All of the work found focuses on warning the user of upcoming traffic from the back. Using radar is a proven method of doing this, used in both [21], [22] and in the Garmin Varia. While radar is robust, camera has the advantage that it tends to give higher-resolution information, as well as colour and texture data. Therefore it is used either in sensor-fusion with radar in [21], or stand-alone in [23]. The disadvantage of including camera is that you need more processing power for either the sensor fusion or the object detection algorithms. This increases the systems complexity and energy usage. However, without camera it becomes very hard to get information on traffic signs, road markings, and type of environment (rural or city). When looking to new, imaginative applications it seems camera processing is essential. The work done in [23] shows that a monocular camera is all you need to create an effective warning system. However, their work was done on a high-performance laptop, which was to be carried in a back-pack. It is therefore unknown if such an approach is also possible for a limited embedded system. This is one of the research questions of this project, and therefore the choice was made to use a monocular camera.
More work has been done on designing new output methods of information for cyclists. The current standard solution in smart tech for bikes is through screens, but this distracts the user from the traffic around them. Therefore the goal has been to design a new output method that keep the users sight on the traffic. This is done either by moving the visual information closer to the natural field of view, or by using different sensory modes such as hearing or touch. Since both vision and hearing are already used by the cyclist to react to traffic, it might be a good idea not to use these modes to avoid distraction. Additionally, those who have impaired vision or hearing could especially benefit from a artificial warning system. Reusing those same senses for the warning output would make it less useful for the most important target group. Taking all this into consideration, the choice was made to use vibrating handlebars for the system's output.

Chapter 3

Background

This section gives some background on the neural network used. This information is however not necessary to understand the work done and can be skipped.

3.1 YOLO Neural network

Neural networks are a subset of machine learning. They are comprised of various node layers, which are connected together. The number of layers, and various architectures of connections can vary immensely, as well as the purposes for which the neural network is trained. All of these differences mean nowadays there is a large amount of neural networks available. Object detection is a very classical use for neural networks. It consists of identifying one or more objects in an image, by drawing a *bounding box* around the object, and giving it a label (making it a more complex problem than simply image classification). Various trained neural networks are available for object detection. During the inference stage, there is often a trade-off between speed, accuracy and computing power needed. Neural networks that are especially fast (between 10 and 45 images per second) without needing excessive computing power are labeled as 'real-time', because they are often used in (embedded) applications that have strict real time requirements. Three real-time neural networks that are often used are YOLO, MobileNet and SSD. For this project the choice was made to use YOLO (this choice is explained in section 4.1.3).

Object detection algorithms can generally be divided into *two-stage* or *one-stage* networks. The first stage in a two-stage network is to detect regions of interest that might contain an object. In the second stage this region is then labeled or *classified* as a certain object. This means there are multiple passes through the classification network. One-stage networks only require one pass through the entire neural network. These one-stage networks are much faster as a result. YOLO, as suggested by its name (You only look once) is such a one-stage network. YOLO has several network sizes. The choice of the network size is a trade-off between speed and accuracy.

In this project the choice was made for a version of YOLO that implements scaling cross stage partial network (csp). Csp splits the base feature map at the beginning of a stage into two paths. One of these paths is linked directly to the end of the stage, where the two paths are fused. This reduces the computations needed, without reducing accuracy [38].

3.2 Neural network optimization

TensorRT is developed by NVIDIA (which is the same company that has created the hardware used for this project: the Jetson Nano). It is an optimization library for neural networks. It does this optimization through various methods, including adjusting the data types and data representation, fusing nodes and layers, as well as hardware-specific adjustments.

Chapter 4

Methods

In the following chapters, the abbreviation FOV will be used to indicate the area that is visible to the camera. The camera is mounted at the back of the bike, looking towards the rear. ‘Approaching the cyclist‘ or ‘forward direction‘ are used to indicate a vehicle coming from the back of the cyclist, which will slowly increase in size in the camera’s FOV until it leaves the FOV to pass the cyclist. ‘Moving away from the cyclist’ and ’backwards direction‘ are used to indicate a vehicle that comes from the front of the cyclist and passes the cyclist before entering the FOV in large size, and then slowly decreasing in size as it moves away from the cyclist. This chapter is divided into three sections. The first section describes how the approaching vehicles are detected. The second section describes the vibrating handlebars, and the third section describes the survey that was developed. The code can be found on the utwente gitlab repository. The precise link is found in appendix A.

4.1 Approaching vehicle detection

4.1.1 Overview of the System

An overview of the detection system can be found in figure 4.1. The only sensor input is the video frames collected from a camera. A neural network is used to detect cars in the frame. The cyclist should only be warned about vehicles in the forward direction. Because of the monocular camera, we do not have the velocity information, in contrast to systems that use radar or stereo camera input. Instead the increase or decrease in size is used to determine whether the vehicle is approaching or moving away from the cyclist. To find this difference in size we need to be able to id and track a vehicle for multiple frames. For this a simple matching algorithm is used based on a cost matrix. An estimate of how close the vehicle is to the cyclist is based on the width of the bounding box. This works because nearly all cars, vans and trucks have a very similar width and because we can assume the camera

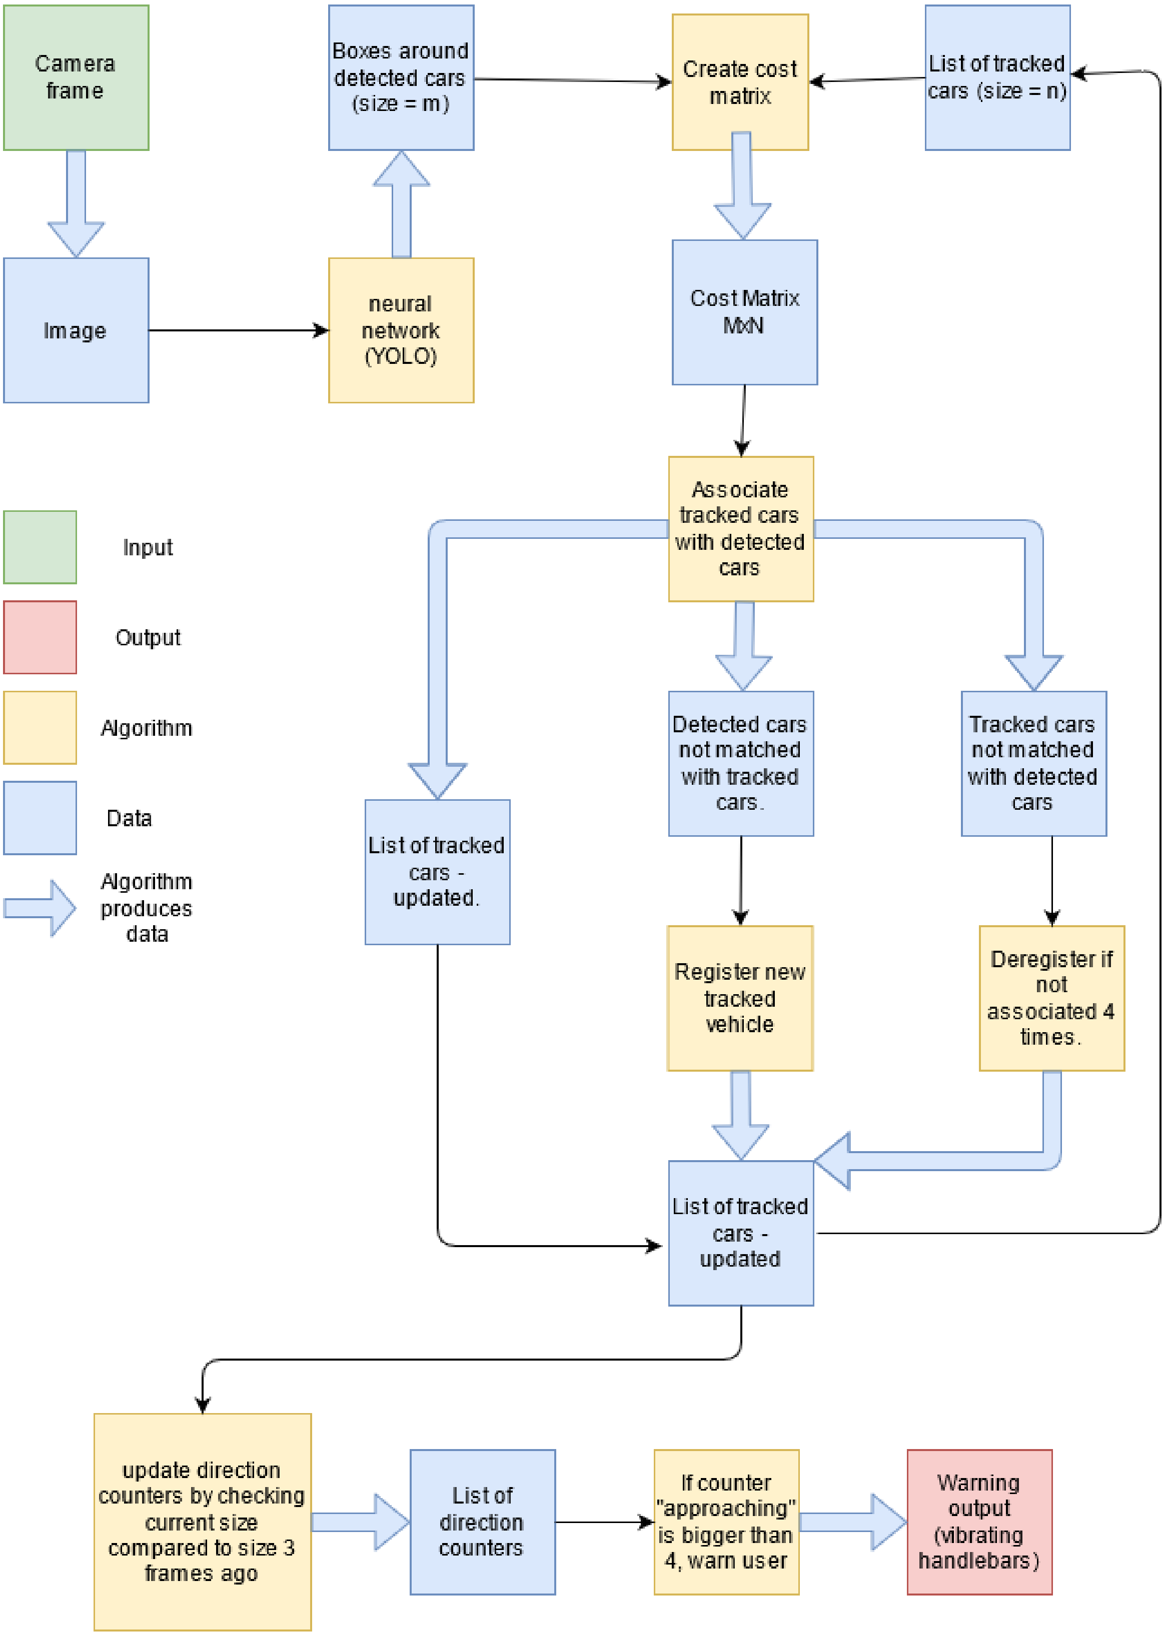

Figure 4.1: System flowchart

position is fixed. The program was developed with Python.

4.1.2 Computing Hardware

The hardware platform chosen was the Jetson Nano (4GB version) developer kit. This hardware platform is a good choice because it is cheap (around 100 euros), while still providing enough computational power to run neural networks. Maybe even more important is the fact that the product is also targeted at the hobbyist target group, and therefore has a lot of documentation, tutorials and hardware-specific libraries. This significantly reduces the development time of the prototype. Given that a bicycle is a power-limited environment, the Jetson was always run at the lower power consumption mode of 10W. For general setup, the blog by JK Jung was very helpful [39]. The camera used is the raspberry pi camera module (Verson 2).

4.1.3 Neural network

The choice of neural network to use was important. It must be able to perform in real-time on limited hardware. JK Jung has created code for several neural networks to be optimized for the Jetson Nano. He has also provided a table with the FPS and mAP values. I estimated that for good performance, a minimum fps rate of 6 was needed. Because the neural network would not be the only code running, and a lower power mode is used, an additional margin was used, leading to a minimum of 9 fps. Given that requirement, the chosen network was yolov4-csp-256 [40]. In implementing the neural network on the Jetson Nano I used the work of JK Jung, which optimizes the YOLO network using Nvidia's TensorRT inference optimizer [41]. For future work it should be noted that a reasonable accuracy in the placement of the bounding boxes is important for the direction detection algorithm. I also tried the system out with tinyYolo, but found that the box placement was too inaccurate to get a good estimate of direction of travel.

4.1.4 Tracking Algorithm

Most current tracking algorithms use neural networks, feature detection or a Kalman filter. The Kalman filter is often used for radar-based tracking. However, typically a Kalman filter makes use of (noisy) velocity information for its updates. Since a camera does not provide this information, it was decided not to try this approach. Running one neural network for object detection is already a big load on the computing power available, so it was preferable to not do this for tracking too. There are several tracking algorithms that work with finding 'points of interests', such as SIFT or

SURF. However, during early experiments it was noted that such methods performed very poorly when objects were far away (that is, low resolution). Finally, a custom algorithm was developed. Given M detected objects in the current frame, and N previously tracked objects. A NxM matrix was created with dissimilarity scores. These scores were based on three different metrics. The first, most important metric, was the euclidean distance between the detected and tracked vehicle. This on its own already performed quite well, but suffered from occlusion issues. To solve this the next two metrics were added. The second metric was the difference between the average color in the bounding boxes. (In a first approach the dominant color was extracted instead of the average color. However, this took significant extra computing time, and average color seems to perform just as well). The third metric is a combination of direction of travel and size difference. It checks whether the size difference between the detected and tracked vehicle, matches with the expected size difference given the direction of travel. The pseudo code this third metric can be seen in listing 4.1. This is helpful since most occlusions in traffic situation occur because two cars with opposite directions of travel pass each other.

```
size_diff = absolute(size_tracked - size_detected)
if direction_tracked is unknown:
    score = size_diff
else if direction_tracked is coming_from_back:
    if size_detected > size_tracked:
        score = - size_diff
    else if size_detected < size_tracked:
        score = size_diff
else if direction_tracked is coming_from_front:
    if size_detected > size_tracked
        score = size_diff
    else if size_detected < size_tracked:
       score = - size_diff
```

Listing 4.1: Pseudocode for the size/direction metric of the tracking algorithm

4.1.5 Direction detection algorithm

To find the direction two different counters are used. One for the forward direction (approaching the cyclist), and one for the backwards direction (moving away from the cyclist). These counters are incremented if the increase or decrease of the size between frames is bigger than a predetermined constant. When the counter is not incremented, it is always decremented, which acts as the extinction function. The counters are also limited on both sides to preserve responsiveness of the system. When the counter reaches a certain threshold (in this case the same as the maximum threshold, but this does not necessarily have to be so), the tracked vehicle will be assigned the appropriate direction. The pseudocode of the code can be seen in pseudocode listing 4.2. Despite using the word 'size', in practice only the width of the vehicle is used. This is because the width of most major vehicles (cars, vans, busses and trucks) are similar, as a direct consequence of the traffic infrastructure. In contrast, the height can differ a lot between these various types. The reader might notice that the pseudocode contains three different kinds of directions: forward, unknown and backwards. Unknown and backwards are treated the same by the rest of the system, but are differentiated in the code for ease of debugging and for flexibility of future extension. Two variables have the most impact on the balance of the system. The first is the amount of frames between the new size and old size. More frames means that smaller changes in size can be detected, and therefore that warnings can be given earlier. However, it also means the system is less responsive to changes of direction. Additionally, if a mismatch occurs between a tracked and detected vehicle, this often results in an incorrectly assigned orientation, because of the big jump in sizes. Having fewer frames between the old and new sizes means the system can recover more quickly from these kind of mismatches. The second important variable is the threshold the counters must reach before a direction is defined. Here a similar balance exists. A lower threshold means quicker assignment of direction, and therefore earlier warnings. However, a threshold that is too low can lead to reduced accuracy because of a higher susceptibility to noise. On the other hand, if the threshold is set too high, the system will respond slow to changes in directions of travel, and will recover slower from inaccurate direction assignments.

```
if newSize > SIZE_DELTA + oldSize:
    forwardDirectionCounter += 1
else:
    forwardDirectionCounter -= 1

    if newSize + SIZE_DELTA < oldSize:
        backwardDirectionCounter += 1
    else:
        backwardDirectionCounter -= 1


if forwardDirectionCounter >= MAX_COUNT_DIRECTION:
   forwardDirectionCounter = MAX_COUNT_DIRECTION
   direction = FORWARD
else if forwardDirectionCounter <= 0:
    forwardDirectionCounter = 0
    if direction is FORWARD:
        direction = UNKNOWN

if backwardDirectionCounter >= MAX_COUNT_DIRECTION:
```

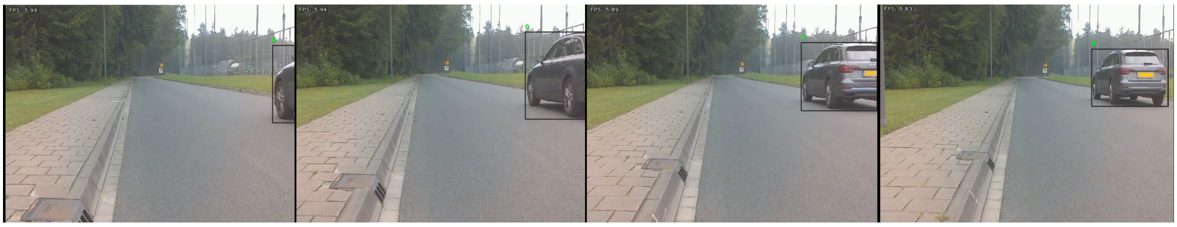

Figure 4.2: Sequence of frames that shows how a car entering the FOV will intially increase in width, before it will be decreasing again as it drives further away.

```
   backwardDirectionCounter = MAX_COUNT_DIRECTION
   direction = BACKWARD
else if backwardDirectionCounter <= 0:
    backwardDirectionCounter = 0
    if direction is BACKWARD:
        direction = UNKNOWN
```

Listing 4.2: Pseudocode for direction detection algorithm

4.1.6 Vehicles entering and leaving the FOV

The above direction detection algorithm assumes a car will increase in size as it approaches the cyclist. However, there is another common occurrence in which the size will increase. When a vehicle enters the FOV, at first it will only be partially visible. In the next couple of frames more of the vehicle will be revealed, which to the system reads as an increase in size. This can be seen in figure 4.2. A small addition to the system was made to account for this case. It was noted that when vehicles entered or left the FOV, the YOLO network actually returned bounding box coordinates that were slightly outside of the image. The exact reason for this is unknown, but it proved useful. When the coordinates were outside of the image, it was added to either a list of *leaving* or *entering* vehicles. The correct list was determined by seeing on which side of the image the vehicles were. (This has to be flipped if the system was used in a place where vehicle drive on the left of the road, such as the UK). If the vehicles are on the *entering* list, the direction detection algorithm will be skipped. If the vehicles were on the *leaving* list, the tracked vehicle will be deregistered from the tracking algorithm in the next iteration.

4.1.7 Warning system evaluation

The dataset needed for finetuning the system and evaluation has some difficult requirements. Because the tracking and direction algorithm requires multiple subsequent frames, the dataset has to consist of videos, not images. Additionally, it has to be from the viewpoint of the back of the bicycle. Additionally, although not necessary, it is strongly preferred that the video images are taken in The Netherlands. The Netherlands has a unique cycling infrastructure, which could impact the functioning of the system. No dataset was found which fulfilled all these requirements, so instead the video material was gathered as part of this project. Initially the idea was to process the images on the go, and save the processed images to analyze later. However, it was noted that saving the images severely impacted the systems performance, dropping the fps below 5. Instead the choice was made to record videos at 6 fps, and then process these videos later. The disadvantage of this method is that the effect of lag spikes is not taken into account in this evaluation. Additionally, it was decided that annotating the taken videos with the ground truth, and writing code that includes the wanted parameters, would be more time consuming than simply finding these parameters manually. In a spreadsheet the following parameters were noted: when the vehicle was first seen with the eye, when was the first the neural network detected it, when it was marked as coming towards the user, and when it passed the user. Additionally a record was kept of errors such as when cars were incorrectly evaluated as coming towards the user, or when a car coming towards the user lost that designation for multiple frames. The choice was made to not include bicyclers, scooters or motors. This is because they were far less common in the video material, making it hard to correctly fine-tune the code for them and evaluate them. In future iterations controlled experiments will probably have to be done to successfully include these in the software. Additionally, cars that were on a different road than the cyclist were not included in the spreadsheet until they turned on the same road as the user (unless their direction was incorrectly classified).

4.2 Vibrating Handlebars

For this project a pair of bicycle handlebars was used, which were made of quite flexible rubber. This allowed them to be partly cut open, and a small vibration motor (275 mW) was embedded inside the rubber. These motors are controlled by pulse-width modulation (PWM) signals. Although the Jetson Nano should be able to output a PWM, the process is fairly complicated [42]. To avoid delays an Arduino Uno was used to drive the motor, which was interfaced with the Jetson Nano through serial communication. The Arduino could most likely be replaced by a small and inexpensive PWM driver.

Figure 4.3: Prototype

4.3 Integration in prototype

The prototype can be seen in figure 4.3. The Jetson Nano was placed in a metal case for protection, correct camera angle and ease of fastening. This was attached to a bicycle carrier using plant twist tie. The handlebar with vibration motor was put on the bike. A layer of duct tape was used to waterproof the vibration motor. Two wires run from the vibration motor to the Arduino. This is duct taped on some places as necessary. The arduino is fastened using a tie wrap.

4.4 Evaluation metrics

4.4.1 Questionnaire

Normally a user study using the actual prototype would be the best way to obtain user feedback on a system such as this, but because of time limitations and the Corona pandemic, this was not possible. Instead, a questionnaire was used to elicit feedback from a distance. Because access to the prototype itself was not possible, demo videos were used to demonstrate the system. The questionnaire itself consisted of four parts:

- Questions about demographics and bicycle usage.
- Investigation in desired timing of warning.

- Question about preferred vibration warning pattern.
- Investigation into vibrating handlebars and system acceptance

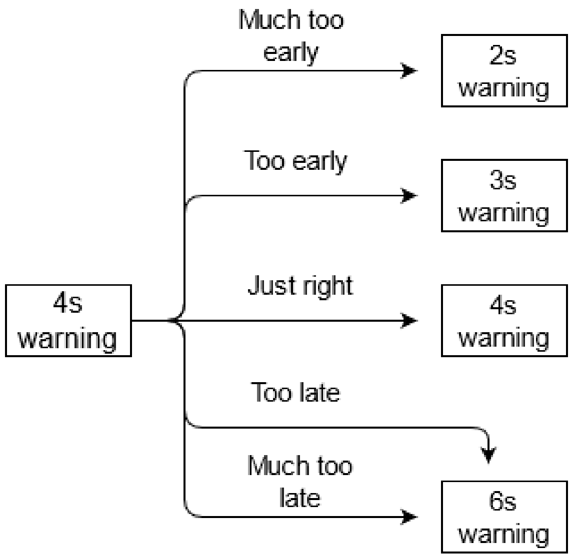

Figure 4.4: flow chart of questionnaire passage

In the second part the user was presented 8 demo video, taken from the point of view of a cyclist. A vibration noise was used to simulate the vibrating handlebars. Users were asked about the timing of the vibration using a 5-point Likert scale. Four separate clips were used. These were taken from the video "Fietsen van Heerhugowaard naar Alkmaar" posted by the Youtube user wierpmesch under a Creative Commons license [43]. Four different vibration timings were used. 2,3,4, and 6 seconds before the car starts passing the user. The first video for everyone uses a 4 second warning. The next video is then based on the user's answer. A flowchart of this can be seen in 4.4. This method is used for the following 2 clips too. In this way the first cycle of four clips is finished. Then a second cycle of another four clips is started, using the same method. This time the cycle starts with a video with a timing of 3 seconds. In this way the preferred timing of the respondent is found with as few videos as possible, while at the same time evaluating their consistency. In the third part a demo video is shown with five different vibration patterns, and respondents are asked about their preferences. The intensity of the vibration was mapped to volume of the buzzing sound in the questionnaire. The first the vibration patterns consisted of constant buzzing, which slowly increased in volume as the vehicle neared. The second pattern consisted of short pulses, which increased in frequency as the vehicle neared. The third pattern consisted of longer pulses, that increased in frequency and became shorter as the vehicle neared. The last two were simply the two pulse patterns, but with increasing volume too. Finally in the last part some more questions are asked about their opinion on the system and its usefulness, as well as their experiences with bicycle computers. The survey has a combination of multiple-choice questions for numerical analysis, as well as open-ended questions for a more qualitative evaluation. The open-ended questions act as a replacement for the semi-structured interviews I would normally hold before designing a survey or while doing a user study, but which was not really practical because of the corona pandemic.

The links to the questionnaires can be found in appendix A.

Chapter 5

Results

5.1 Video Analysis

First an analysis of the attainable fps was made. This was done using 5214 measurements of how long it takes to analyze a video frame. 3 different test-videos were used, 1 with low traffic density and two with high traffic density. The videos were projected on the wall, filmed with the camera and processed by the Jetson Nano. 80% of all measured values are higher than 6 fps. Therefore 6 fps was chosen as a good value to run the further tests on. The distribution of fps values reached can be seen in figure 5.1.

35 minutes of video were analyzed. This included 58 cars approaching the cyclist and 65 cars moving away from the cyclist, as well as an unknown number of parked cars. For the analysis of the time between the warning and the car passing the cyclist, the data was cleaned up. Of the 58 cars approaching the cyclist, 9 never passed the cyclist, and 1 slowed down for a long time behind the cyclist. These were removed from the dataset. Additionally there was one car that was never marked as approaching the cyclist. 2 cars were marked as approaching very early as a tracking error. This initial marking was ignored in the analysis. The final distribution of warning time can be seen in figure 5.2. The median value is 3.2 seconds. 12.5% of all warnings are less than 2 seconds, and 44% is less than 3 seconds. The cars were generally detected very early, but were only marked as approaching significantly later. The median time between a car being detected and marked as approaching was 9 seconds.

These measurements were derived from real-life traffic situations. Traffic can be complex and unexpected. Therefore an attempt was made to remove the traffic cases that are 'not standard', to see how this would affect the timings. Of the 48 cars, another 14 were removed. 2 cars that entered the road closely behind the cyclist, 6 cars that were occluded while close behind the cyclist, and 6 cars that slowed down before passing the cyclist. The new total of 36 cars was again analyzed, and can

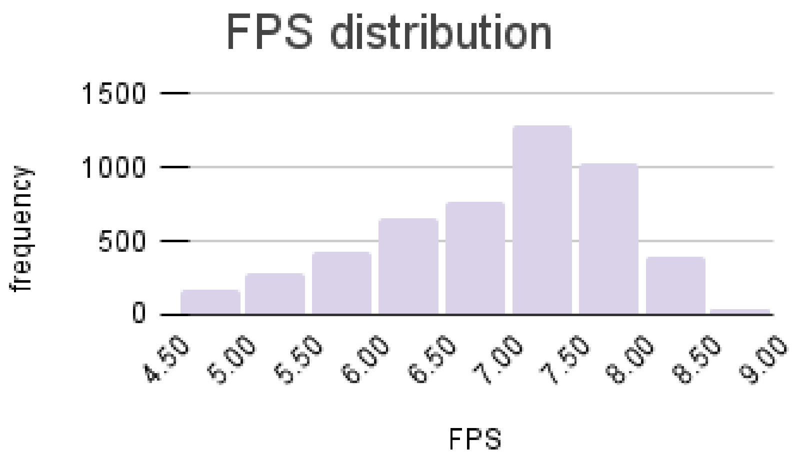

Figure 5.1: Distribution of fps values attained when running the system in real-time

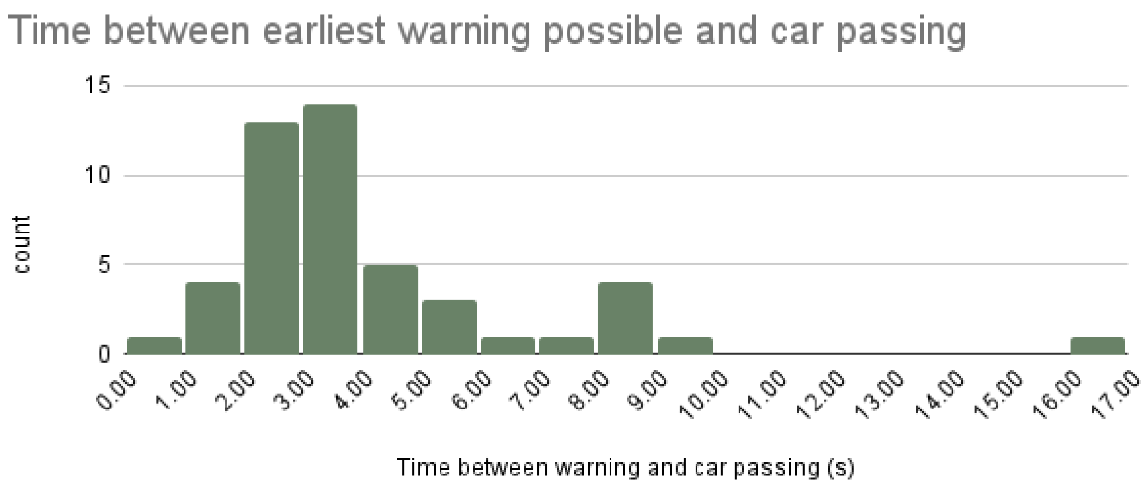

Figure 5.2: Time between warning and car passing. A total of 48 cars was analyzed

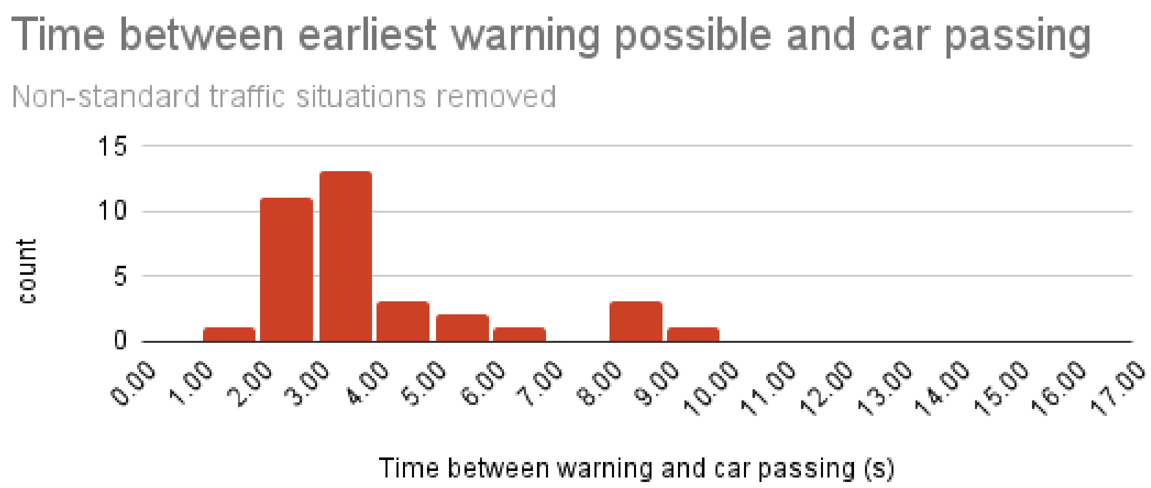

Figure 5.3: Time between warning and car passing. Non-standard traffic situations that could impair the detection and/or tracking were removed. A total of 36 cars was analyzed

be seen in figure 5.3. The removal of non-standard situations did not impact the median, which stayed at 3.2. It did lower the standard deviation from 2.8 to 2, but this could be explained by the reduced sample size. Therefore no conclusions can be drawn without more advanced statistical analysis.

The next step was looking at false positives, that is, cars that are not approaching the cyclist, yet are marked so by the system. Of the 65 cars that were driving in the opposite direction of the bicycle, 5 cars were marked incorrectly in such a way. 4 of these were marked so for a little less than a second, and 1 for a bit more. All of these were caused by an incorrect assignment from the tracker. Additionally, 2 cars that were moving laterally to the cyclist were marked incorrectly. One because of a tracker error, and one because it slowly entered the camera's FoV. The results are much worse for parked cars. In the 20 minutes of video that had a more urban environment with parked cars along the road, 41 parked vehicles were marked incorrectly as approaching the cyclist. In figure 5.4 can be seen how long the incorrect marking remained, and therefore how long the user was wrongly warned of an upcoming car. 85% was marked incorrectly for less than a second. Additionally, an analysis was made what regarding the cause of the incorrect marking. This was done by observing the bounding boxes and assigned tracking number. The results can be seen in 5.5. It is clear that the majority of false positives are caused by wrong assignments by the tracker.

5.2 Questionnaire

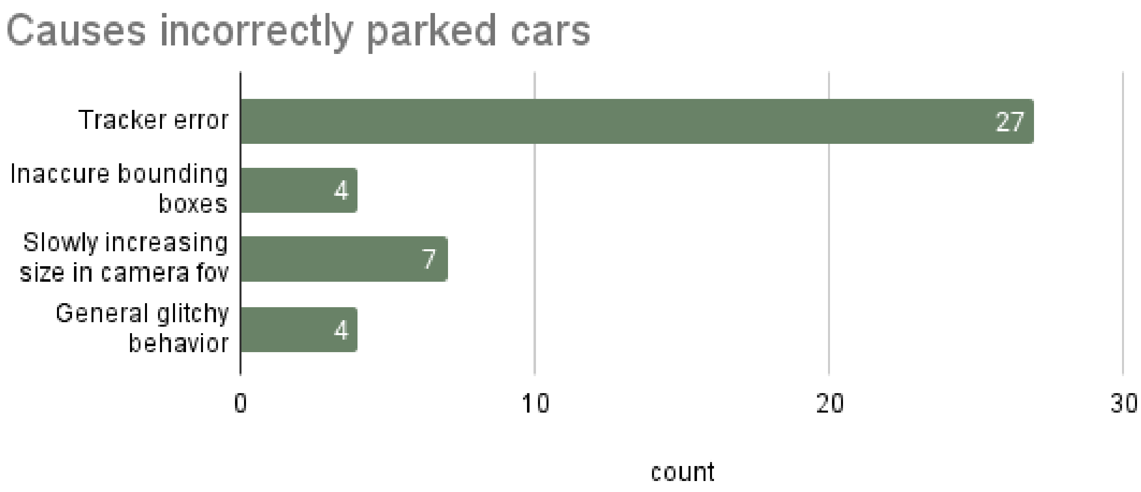

Figure 5.5: Causes of the incorrect marking. In a couple of cases there were multiple causes, so the bars do not add up to a 100%

A total of 41 responses had been recorded in the questionnaire. Most responses had beeng gathered by solliciting friends and family. Special attention had been paid to gather a larger percentage of elderly respondents, since they are probably the target group most likely to suffer from impaired hearing or vision. Because of this a peculiar age distribution occurred, where most respondents are either young adults or elderly, as can be seen in figure 5.6. For each respondent it was noted which of the four timing options (2,3,4 and 6) they had selected 'just right', either once or multiple times. In figure 5.7 the distribution of this can be seen. It is clear that 3 or 4 seconds is the preferred timing for the majority of respondents.

For the vibration pattern, a small problem occurred. Because testers had problems hearing the videos on phones, the volume of the videos was increased. However, this had as a result that the gradual volume increase of the pulse videos was barely noticeable. This has probably impacted the results. The preferred vibration pattern can be seen in figure 5.8.

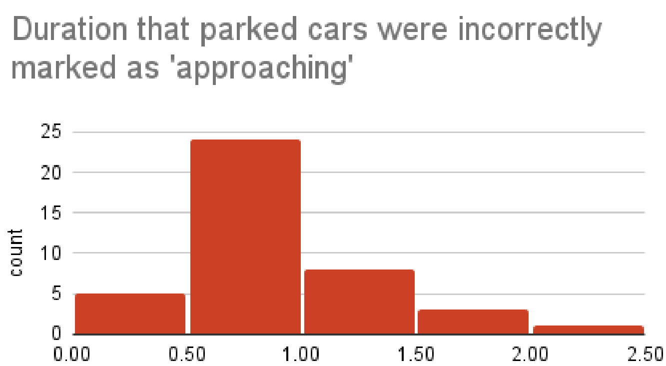

Figure 5.4: Duration that parked cars were incorrectly marked as approaching the cyclist. n = 41

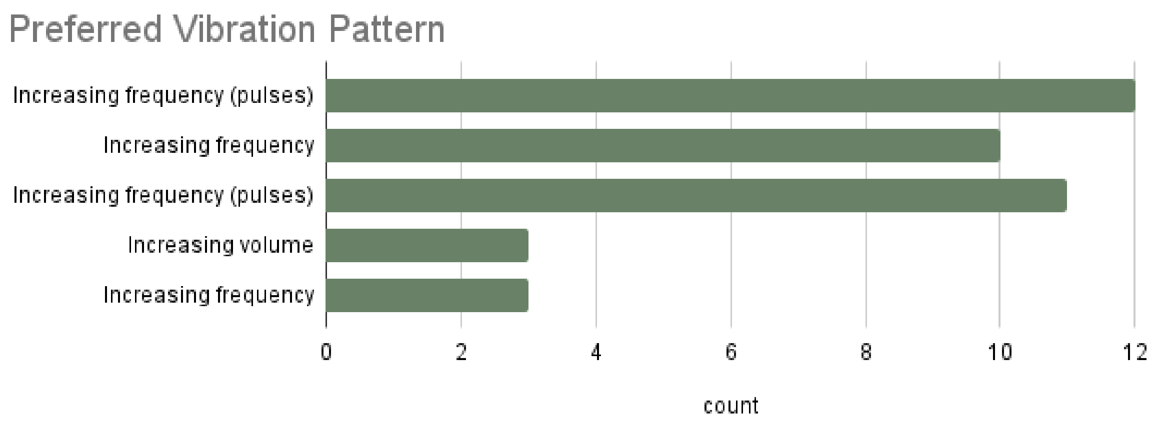

Figure 5.8: Vibration pattern preferred by respondents. n = 41

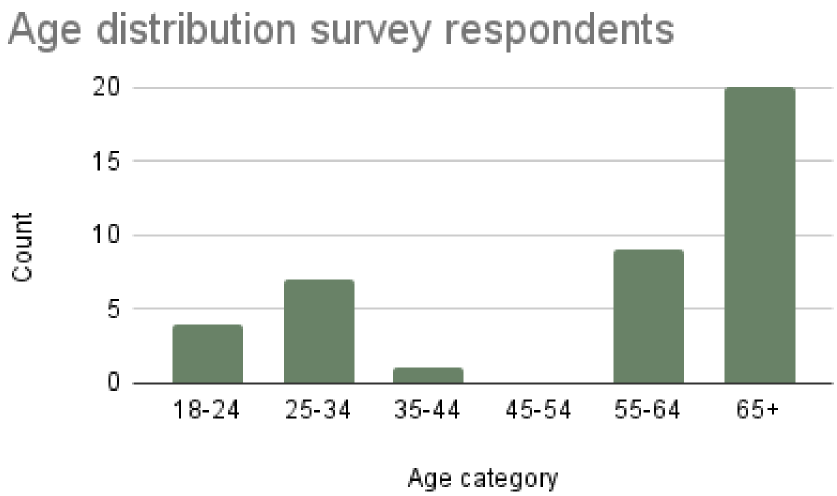

Figure 5.6: Age distribution of the respondents.

When it comes to the analysis of the expected acceptance of the system. 14 respondents had used a cycling computer before. These respondents were presented with questions regarding the vibrating handlebars as a replacement or addition to a cycling computer. Generally these respondents found the cycling computers easy to read, though 4 mentioned it was difficult to read in sunlight, and 5 respondents said it could sometimes distract from traffic. The alternative of the vibrating handlebars was received positively by the majority. 5 mentioned it would generally be a good addition, and another 5 mentioned it would be helpful in certain circumstances. Worries were voiced that an excess of vibrations could be annoying to the cyclist (n=2). One respondent felt that improving cycling safety should not be put upon the cyclists themselves, but instead put upon car drivers giving cyclists more space and being more aware of the cyclists. Another respondent mentioned:

”I think it would be great for places with no dedicated bike paths. I think in the Netherlands traffic coming from behind is not often a problem. However this could be great for use in other countries."

Multiple people mentioned how they valued the additional modality to the existing ones of sound/sight (n=4).

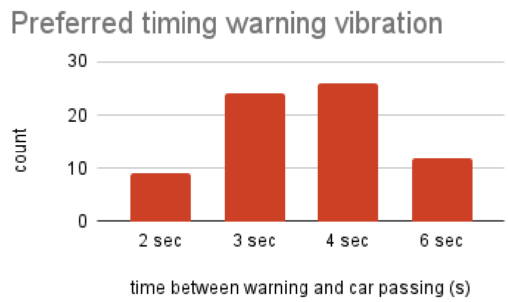

Figure 5.7: Amount of respondents that marked a timing as 'just right' at least once. Respondent could do this for multiple timings, so the total is not equal to number of respondents. n = 41

Regarding the general acceptance of the system, 21 respondents felt that the system could be useful for themselves, and 37 respondents felt that the system could be useful for others. The explanations for why it would be useful was to reduce the startle response in elderly cyclists (n=3). Others mentioned how it would mostly be useful for those with reduced hearing and/or vision (n=9).

6 respondents had hearing and vision issues, and these will be investigated more closely. All of them were older than 55, 3 male and 3 female, and all of them biked at least once a week, four even cycled 4-6 times a week. 5 had hearing problems, and one had vision problems. 5 of the 6 felt that their impairments negatively affected their safety while cycling. 2 of these had rear mirrors on their bicycles, and were satisfied with the solution, though one of them would be open to trying out the new system. Of the other four, one man did not feel that his impairments negatively affected his safety, and therefore did not think the system was useful for him. The other three all thought the system would be useful for themselves. One user said (translated from Dutch):

"I think it would be very nice for the hearing impaired, because now you often startle and make unexpected movements that can lead to a fall."

And another said:

"I am sometimes caught off guard by a passing car, especially if it does not keep enough distance."

5.3 Personal experience prototype

Ideally in this section a user study would be presented. However, due to time constraints as well as the Covid pandemic this was not possible. However, I believe that a description of my personal experiences is still useful as an indication and inspiration for future work.

Figure 5.9: Testing out the prototype

The false positive warnings were quite noticeable, although they were usually quite short, while the ’true’ warnings were of a longer duration. No cars passed without a warning. Usually I could hear the car coming before the warning. The vibration mapping was troublesome. At first, a mapping was made with increasing vibration intensity. However, with the rubber and duct tape on top of the vibration motor, differences in intensity were hard to detect. Pulses were used next, but this made it harder to tell the difference between true warnings and the false positives. Finally I just settled on not using any mapping for the distance, and always output a warning at full vibration. An interesting observation was that even though the vibration was of a moderate intensity, it could still be felt in the other handlebar (which didn’t have any vibration motor embedded) and even in the bicycle seat. I was very surprised such a small vibration motor, attached to the handlebar, resulted in me being able to feel the vibration all the way in the saddle. This could be seen as an undesirable side-effect, but maybe vibrations in the saddle could also be used to provide a higher complexity of feedback.

Chapter 6

Discussion

Although the current system is functional, it is not reliable enough. The warning time is shorter than people prefer, and the false positives, although short, are numerous. Both issues can be traced back to tracker failures. When the tracker makes a wrong association between a tracked and detected vehicle, the direction detection algorithm is likely to give the wrong results. This results in the false positives. The reason that this happens more to parked vehicles is not entirely clear, but most likely it is simply because there much more parked vehicles in the video than moving vehicles, and because they are closer together, often parked headlight to tailpipe. To recover quickly from these errors, the direction detection algorithm is designed to be very responsive, but this also results in having less time to warn the cyclist, as was explained in section 4.1.5.

Improving the tracking algorithm should be the first step taken. The tracking algorithm currently has similarities to the commonly used combination of the Hungarian algorithm and Kalman filter. Using the Hungarian algorithm will probably increase the efficiency of the algorithm, although the tracking consitutes only a small part of the computing time. A Kalman filter could be attempted to improve accuracy, but the disadvantage is that we don't have access to information about the velocity or external forces, meaning multiple assumptions will have to be used. Training the object detection algorithm to differentiate between different orientations of the cars is possible, but since the problem mostly arises with parked cars, this will most likely not solve the issue. A more fruitful approach might be detecting the road - a common computer vision problem for which much work is already done - and filtering out the cars that are not situated on this road, but instead to the side. A similar train of though would be using the width/height ratio to estimate how lateral the car was to us, although varying road-widths will probably not make this reliable. Increasing the reliability threshold of the neural network also seemed to give some small improvements, but not enough time was available to check the quantitative results. More advanced tracking algorithms might also be explored, though this could cause issues with a lower frame rate, and leaves less computational room for new features. Another approach could be making the direction-detection algorithm more robust to tracking failures. One might be able to take a larger window frame when all the data indicates a constant direction/velocity, yet discard past data when sudden or unexpected changes are found. In this way the system might be able to mark approaching cars faster, without losing robustness when dealing with tracker errors or sudden direction changes from the cars. The hindsight paper seem to benefit from a significantly higher fps (around 30), which means there is more data available to determine the increase in size, and makes it easier to track vehicles since there is less change between frames. Increasing the fps could be another way to improve the performance of the system. However, they also report an average warning time of only 1.89 seconds, though with significantly less false positives.

Chapter 7

Conclusions and Future Work

7.1 Conclusions

A survey indicates that a warning system for upcoming traffic on a bicycle interests a large proportion of people, with 51% of respondents indicating such a system would be useful for themselves, and 90% indicating it could be useful for others. The majority of respondents preferred that the warning would be felt around 3 or 4 seconds before the car passes. Most people preferred a type of pulse over constant vibration. No real difference was found in the preference between short or longer pulses.

A functioning prototype has been created. It uses a monocular camera system. However, the prototype has a warning of less than 3 seconds in 44% of the cases. Results from the survey indicated that at minimum a warning should be given 4 seconds before the car passes. This means that the current performance of the system is not sufficient. The biggest choke-point is the inaccuracy of the tracking algorithm. This generates a lot of false positives. To get around this the direction detection algorithm responds very quick by only looking at a small amount of frames in the past. However, the consequence of this is that the subtle size changes in cars far away are not adequately detected.

7.2 Future Work

An important question that is not yet resolved is why choose for a mono-camera approach, instead of a stereo-camera or radar-camera fusion approach. Having access to the velocity information from the radar eliminates the weak point in this project - the tracker. This suggests that for an embedded system in the bicycle - as envisioned in this project - a radar-camera fusion approach could be much more viable. Looking at reviews from the Garmin Varia, as well as conversations with a peer who worked on a similar system using radar, suggests that using radar for this application is very robust and reliable. However, the work done here might still be valuable in an environment that has by default only one camera and no radar - the modern mobile phone. This has as advantage that no additional hardware has to be created or installed, and could massively bring down costs for the consumer, potentially even free. Investigating whether the modern mobile phone has enough computing power and battery life to run such a system would be an interesting new step. Finally, a more hands-on approach, potentially even in virtual reality, in testing the different cyclist behaviours and user experience between auditory, visual and haptic warnings in a cycling environment could provide a lot more information. This also permits for modeling a future traffic situation where the majority of vehicles are electric, as well as representing more complex traffic situations. The personal struggle with mapping distance in the prototype makes clear that future user studies need to use a physical bicycle handle with vibration. Additionally, to have more information from the relevant target group, future user studies should target specifically those with impaired hearing or vision. Finally, the Netherlands also has a unique biking infrastructure which neither academic work nor the existing Garmin Varia takes into account. Biking paths are often separated from vehicles through grass and/or trees. When biking on such paths, nearing cars on the separate road are not very interesting for the biker, but nearing scooters on the bike-path should still be picked up. A user of the Garmin Varia in the Netherlands noticed that from the 27 cars the system picked up, only 10 were actually on the same road as him [30]. Creating a system that detects the difference between separated cycling lanes and shared space with cars could be very useful in adjusting the warning system. I had done some work in this topic too, but due to disappointing results it is not included in this report.

Bibliography

[1] “Ellimans universal embrocation slough 1897 ad,” Jun 2006. [Online]. Available: https://commons.wikimedia.org/wiki/File:Ellimans-Universal-Embrocation-Slough-1897-Ad.png

[2] B. V. Z.-J. Charles, “Impact et potentiel de l’usage du vélo sur l’économie et l’emploi en région de bruxelles-capitale,” 2014.

[3] C. B. Guure, N. A. Ibrahim, M. B. Adam, and S. M. Said, “Impact of physical activity on cognitive decline, dementia, and its subtypes: meta-analysis of prospective studies,” *BioMed research international*, vol. 2017, 2017.

[4] F. B. Schuch, D. Vancampfort, J. Firth, S. Rosenbaum, P. B. Ward, E. S. Silva, M. Hallgren, A. Ponce De Leon, A. L. Dunn, A. C. Deslandes *et al.*, “Physical activity and incident depression: a meta-analysis of prospective cohort studies,” *American Journal of Psychiatry*, vol. 175, no. 7, pp. 631–648, 2018.

[5] I. J. Hendriksen, M. Simons, F. G. Garre, and V. H. Hildebrandt, “The association between commuter cycling and sickness absence,” *Preventive medicine*, vol. 51, no. 2, pp. 132–135, 2010.

[6] “The benefits of cycling - unlocking their potential for europe,” Dec 2018. [Online]. Available: https://www.ecf.com/sites/ecf.com/files/TheBenefitsOfCycling_final-v2.pdf

[7] Dec 2018. [Online]. Available: https://res.cloudinary.com/rampi/raw/upload/v1540861448/Air_quality_in_Europe_-2018_report-compressed_edcg1e.pdf

[8] D. Houthuijs, A. Van Beek, W. Swart, and E. Van Kempen, “Health implication of road, railway and aircraft noise in the european union: Provisional results based on the 2nd round of noise mapping,” 2014.

[9] J. Dekoster, U. Schollaert, and C. Bochu, *Cycling: the way ahead for towns and cities*. Office for Official Publications of the European Communities, 2000.

[10] D. Horton, P. Rosen, and P. Cox, *Cycling and society*. Routledge, 2016.

[11] E. Fishman, S. Washington, and N. Haworth, “Understanding the fear of bicycle riding in australia,” *Journal of the Australasian College of Road Safety*, vol. 23, no. 3, pp. 19–27, 2012.

[12] J. Pucher and R. Buehler, “Cycling for everyone: lessons from europe,” *Transportation research record*, vol. 2074, no. 1, pp. 58–65, 2008.

[13] P. Schepers, D. Twisk, E. Fishman, A. Fyhri, and A. Jensen, “The dutch road to a high level of cycling safety,” *Safety science*, vol. 92, pp. 264–273, 2017.

[14] Jul 2020. [Online]. Available: https://www.cbs.nl/nl-nl/nieuws/2020/31/daling-verkeersdoden-afgelopen-20-jaar-groter-bij-autorijders-dan-fietsers

[15] “Aantal verkeersslachtoffers moet in 2030 gehalveerd zijn,” Jul 2021. [Online]. Available: https://www.anwb.nl/verkeer/nieuws/nederland/2021/juli/aangenomen-motie-halvering-aantal-verkeersslachtoffers-in-2030

[16] H. Parool, “Bijna een vijfde van alle fietsritten is elektrisch,” Oct 2020. [Online]. Available: https://www.parool.nl/nederland/bijna-een-vijfde-van-alle-fietsritten-is-elektrisch~b5d89da7/?referrer=https%3A%2F%2Fwww.google.com%2F

[17] “11 procent 55-plussers in 2018 dagelijks op elektrische fiets,” Nov 2020. [Online]. Available: https://www.cbs.nl/nl-nl/nieuws/2020/47/11-procent-55-plussers-in-2018-dagelijks-op-elektrische-fiets

[18] T. Voermans, “Verkoop elektrische auto’s en fietsen breekt alle records,” Sep 2020. [Online]. Available: https://www.parool.nl/nederland/verkoop-elektrische-auto-s-en-fietsen-breekt-alle-records~b585a512/?utm_source=link&utm_medium=app&utm_campaign=shared+content&utm_content=free

[19] M. van Infrastructuur en Waterstaat, “Mobiliteitsbeeld 2019,” Nov 2019. [Online]. Available: https://www.kimnet.nl/mobiliteitsbeeld/mobiliteitsbeeld-2019#/rapport/1.3.3

[20] “E-bike zakelijk gebruiken als zzp’er: de belastingregels van 2021,” Feb 2021. [Online]. Available: https://zzpservicedesk.nl/e-bike-zakelijk-gebruiken-als-zzper-de-belastingregels-van-2021/

[21] C. Degen, C. Domnik, A. Kürten, M. Meuleners, M. Notz, R. Pohle-Fröhlich, and E. Naroska, “Driver assistance system for pedelecs,” in *2019 20th International Radar Symposium (IRS)*. IEEE, 2019, pp. 1–8.

[22] M. Hägelen, R. Jetten, J. Kassner, and R. Kulke, "Safety and comfort enhancement with radar for a bicycle assistance system," in *2019 20th International Radar Symposium (IRS)*. IEEE, 2019, pp. 1–7.

[23] E. Schoop, J. Smith, and B. Hartmann, "Hindsight: enhancing spatial awareness by sonifying detected objects in real-time 360-degree video," in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, 2018, pp. 1–12.

[24] D. de Waard, K. Edlinger, and K. Brookhuis, "Effects of listening to music, and of using a handheld and handsfree telephone on cycling behaviour," *Transportation research part F: traffic psychology and behaviour*, vol. 14, no. 6, pp. 626–637, 2011.

[25] K. Terzano, "Bicycling safety and distracted behavior in the hague, the netherlands," *Accident Analysis & Prevention*, vol. 57, pp. 87–90, 2013.

[26] "Un regulation on advanced emergency braking systems for cars to significantly reduce crashes," Feb 2019. [Online]. Available: https://unece.org/press/un-regulation-advanced-emergency-braking-systems-cars-significantly-reduce-crashes

[27] "Smart helmet." [Online]. Available: http://livall.com/

[28] "Lumos helmet - a next generation bicycle helmet." [Online]. Available: https://lumoshelmet.co/

[29] "R1 - smart road cycling helmet." [Online]. Available: https://www.sena.com/product/r1

[30] "Garmin varia rtl515 and rvr315 cycling radar in-depth review," Nov 2020. [Online]. Available: https://www.dcrainmaker.com/2020/05/garmin-rtl515-rvr315-cycling-radar-review.html

[31] Www.thingsthat.com, "Laserlight core." [Online]. Available: https://beryl.cc/shop/laserlight-core

[32] A. Dancu, V. Vechev, A. A. Ünlüer, S. Nilson, O. Nygren, S. Eliasson, J.-E. Barjonet, J. Marshall, and M. Fjeld, "Gesture bike: examining projection surfaces and turn signal systems for urban cycling," in *Proceedings of the 2015 international conference on interactive tabletops & surfaces*, 2015, pp. 151–159.

[33] "Skully fenix ar — review," Jan 2019. [Online]. Available: https://www.youtube.com/watch?v=CFF99F3BHsI

[34] H.-Y. Tseng, R.-H. Liang, L. Chan, and B.-Y. Chen, "Lead: Utilizing light movement as peripheral visual guidance for scooter navigation," in *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services*, 2015, pp. 323–326.

[35] F. Pece, J. J. Zarate, V. Vechev, N. Besse, O. Gudozhnik, H. Shea, and O. Hilliges, "Magtics: Flexible and thin form factor magnetic actuators for dynamic and wearable haptic feedback," in *Proceedings of the 30th annual ACM symposium on User interface software and technology*, 2017, pp. 143–154.

[36] M. Pielot, B. Poppinga, W. Heuten, and S. Boll, "Tacticycle: supporting exploratory bicycle trips," in *Proceedings of the 14th international conference on Human-computer interaction with mobile devices and services*, 2012, pp. 369–378.

[37] "World's first connected bike grips." [Online]. Available: https://smrtgrips.com/

[38] C.-Y. Wang, H.-Y. M. Liao, Y.-H. Wu, P.-Y. Chen, J.-W. Hsieh, and I.-H. Yeh, "Cspnet: A new backbone that can enhance learning capability of cnn," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, 2020, pp. 390–391.

[39] J. K. Jung, "Jk jung's blog," May 2019. [Online]. Available: https://jkjung-avt.github.io/setting-up-nano/

[40] A. Bochkovskiy, C.-Y. Wang, and H.-Y. M. Liao, "Yolov4: Optimal speed and accuracy of object detection," Apr 2020. [Online]. Available: https://arxiv.org/abs/2004.10934v1

[41] J. K. Jung, "Jk jung's blog," Jul 2020. [Online]. Available: https://jkjung-avt.github.io/tensorrt-yolov4/

[42] A. Kumar, "How to use gpio pins on jetson nano developer kit: Nvidia jetson," Jun 2020. [Online]. Available: https://maker.pro/nvidia-jetson/tutorial/how-to-use-gpio-pins-on-jetson-nano-developer-kit

[43] *Fietsen van Heerhugowaard naar Alkmaar*. YouTube, Jun 2018. [Online]. Available: https://www.youtube.com/watch?v=9Lqt2294_U4

Appendix A

Links to various resources

The code developed in this thesis can be found at:
`https://git.snt.utwente.nl/smart.cycling/vision.based.warning`

The survey can be viewed at:
Dutch: link
English: link