Developing a Bicycle Lane Assist

ARTHUR HOFMA, University of Twente, The Netherlands

Fig. 1. Bicycle path

Bicycling is a major form of transportation for people in The Netherlands. However, traffic incidents and collisions, including single crashes, have increased. Modern cars are equipped with technology, such as lane assist, to make travel safer. However, bicycles have not caught up. In this paper, we will look into implementing similar lane assist technology, adopted for bicycles. The technology is based on Computer Vision and Image Segmentation. Using self-recorded data sets, the proposed solution achieved a 95% accuracy for detecting road surfaces, and a 68% accuracy for detecting bicycle paths. Using this, the location of the bicycle can be tracked. As a result, the proposed technology can be used as a basis for developing a bicycle lane assist.

Additional Key Words and Phrases: Bicycles, Bicycle path, Virtual bicycle path, Bicycle accidents, Lane assist technology

1 INTRODUCTION

In 2019, cycling was the primary form of transportation for 28% of the population in the Netherlands [CBS 2019]. The reason for the high amount of bicycle-use is the excellent bicycle infrastructure.

However, according to statistics of the last few years, the number of one-sided accidents has greatly increased [CBS 2022]. These accidents can partly be attributed to cyclists colliding with stationary objects, such as trees and (lamp)posts. These objects are often situated on the side of the cycling path, and thus pose a danger to cyclists that veer off the cycling path.

2 PROBLEM STATEMENT

As mentioned in the introduction, accidents occur when cyclists collide with stationary objects located beside the bicycle path. These accidents can be prevented if cyclists do not veer off the bicycle path. This is, of course, easier said than done. Therefore, a solution can be implemented in the form of a lane assist for bicycles. Such technology is already present in modern cars, and can be adjusted to also work for bicycles. for this research, the technology will be based on vision-based road and lane detection.

2.1 Research questions

Based on the problem statement mentioned above, the main research question will be:

To what extent is lane assist technology feasible to implement for bicycles?

To help answer this question, the following sub questions are formulated:

- *How can existing bicycle paths and roads be recognized automatically using computer vision?*
- *How can the location of the bicycle on the bicycle path be measured and tracked?*
- *How can the trajectory of the bicycle be predicted?*

3 RELATED WORK

Lane assist is not a new technology: it has been present in modern cars for a few years. Even though the final implementation for bicycles will differ from the current one in cars, there is a lot of overlap. As the bicycle-adopted version will make use of vision-based lane detection, the focus for related work will be placed on vision-based lane detection.

For lane detection, Xing et al. describe two main methods to identify lanes: feature-based and model-based [Xing et al. 2018]. The former is based on features such as road markings, while the latter is based on modelling the road itself. In their article, they describe pros and cons of both. The most important distinction between the two is how computation-heavy model-based lane identifying is, while feature based is less heavy. As the final system should be portable on a bicycle, the feature-based approach is more suitable in this context.

To further support feature-based lane detection, Haque et al. introduce a system that can efficiently detect lane markings [Haque et al. 2019]. As already discussed, a feature-based lane detection approach is more suitable in this context, which also means that their system can provide more directions in how to better adapt and develop a bicycle lane assist. As mentioned by Haque et al. in the article, the proposed system proves effective in different brightness conditions, which is especially useful for bicycles, as their headlights can differ greatly per model.

TScIT 38, February 3, 2023, Enschede, The Netherlands

However, not all existing bicycle infrastructure has the necessary features, such as road markings, to properly detect the bicycle path. For this reason, image segmentation will be considered as an alternative. Image segmentation is a process where different parts of an image are separated based on features. As the solution should not be computation-heavy, it will be thresholding-based image segmentation. For determining the right threshold, colour extraction is necessary. Therefore, work related to colour extraction will be explored.

For colour extraction, Muzammel et al. describe a method for extracting the average colour of a feature in an image [Muhammad Muzammel 2017]. This is done by selecting multiple patches in the given image, on locations where there is a high probability of that feature being present. Using the mean colour of those patches, the mean color of the feature can be calculated for use in feature detection or segmentation.

4 METHODOLOGY

4.1 Literature review

The first step is to research the different types of bicycle paths in The Netherlands, and what their physical appearance (and differences) are. These features will be used to recognize and classify different bicycle paths later. Attributes that will be taken into consideration are:

- Colour of the bicycle path
- Presence of a white (dashed) line, acting as a divider

These features will dictate what the bicycle path detection algorithm will focus on, which will result in higher accuracy.

4.2 Collecting the data

The data necessary for this research is collected during the research. The data consists of footage, taken from the centre of the bicycle its handle shaft using a smartphone camera. Its perspective will be similar to the one in Fig. 2.

Fig. 2. A two-way isolated bicycle path without markings

4.3 Preparing the data

For data preparation, two steps have to be taken before it can be analysed further.

First, from the collected footage, every 8th frame will be isolated. Based on footage with 24 frames per second, taking a measurement every 8th frame will result in three measurements per second, which is adequate for the average bicycling speed of 20 km/h.

Second, after extracting the relevant frames, the images are pre-processed in order to determine the type of road. This pre-processing will consist of checking for the features mentioned in section 4.1.

4.4 Data analysis

4.4.1 Detecting bicycle paths and roads. For research question 1, detecting bicycle paths and roads, computer vision, as provided in the OpenCV package, is used in order to properly identify the bicycle path in the given image. An algorithm is used in order to achieve this. Below is given a description of this algorithm.

First, the given frame is converted from RGB (red, green, blue) to the HSV (hue, saturation, value) colour space. The relevant difference between these two colour spaces is how a colour is defined. For image processing and isolating colours, HSV is better, therefore that colour space will be used.

Second is the Region Of Interest extraction. This step is based on the research of Muzammel et al. [Muhammad Muzammel 2017]. Four patches of road are chosen based on the high certainty of road surface presence there. These can be found in figure 3 for grey roads and in figure 4 for red roads. The difference between these can be attributed to the width of the region of interest, which is smaller for red bicycle paths. If the same patches were to be used, the algorithm would often categorize the road as bicycle path, resulting in dangerous situations. Therefore, smaller patches are used for greater accuracy and safety. After extracting these four regions, the mean colour of the corresponding colour channel is calculated for the next step. For grey roads, the S and V channel are used. For red roads, all three channels are used.

Fig. 3. Region of interest extraction areas for grey roads

Third is thresholding the image. Thresholding is a method to convert a given image to a binary image, where pixels are either black or white. This is done by defining a lower and upper limit for a pixel in each of the three channels (H, S and V). Then, for each pixel, its value is compared to the lower and upper value. If its value lies between the limits, it is converted to white. Otherwise, it will be converted to black.

Fig. 4. Region of interest extraction areas for red roads

Fourth, closing is applied. As can be seen in figure 2, the road surface does not have a smooth, uniform colour. This is known as noise. Closing is a type of morphology used during image processing to remove noise from the image. The removal of noise makes it easier for the computer to recognise the whole road surface.

Finally, the largest contour (in this case, the road) is located and drawn. The location of this contour is also used later as it provides the coordinates for the next steps. The contour will from now on be referenced as road mask.

4.4.2 Creating a virtual bicycle path. For the next two sections, pixels and centimeters may be converted back and forth. In appendix A, an explanation is given.

First, the portion of the road that is suitable and allowed for bicycling has to be determined. Using the coordinates from the previous step, there are two possibilities for determining the width for the virtual bicycle path. (1) If the left coordinate of the road mask is smaller than or equal to zero, the width is calculated by subtracting the path width from the right coordinate of the road mask. (2) If the left coordinate of the road mask is bigger than zero, then the width is determined by the width of the road mask. The virtual path width is 1.7m, as suggested by Fietsberaad Crow [Crow 2015].

Second, the safe zone is determined. The safe zone is where the cyclist may ride without the need for correcting. It is defined as the middle 50% of the bicycle path. The edges of this area are the soft borders. Next to this area is a buffer zone. This is an area where the cyclist may temporarily be, but should leave and return to the safe zone. The outside edge of this area is a hard border. If that border is crossed, the cyclist should return immediately to the safe zone. in figure 5, this is visualised. The hard and soft border coordinates are calculated based on the virtual bicycle path centre and width.

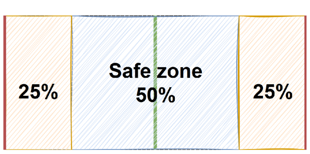

Fig. 5. Virtual bicycle path

Finally, for tracking the current location of the bicycle, the assumption is made that the camera is mounted on the centre of the bicycle its handle shaft. Using this information, the location of the bicycle is always the centre bottom of the image, equal to the coordinates x=(image-width / 2) and y=image-height (note: with images, y starts at the top left corner instead of the bottom left corner). Furthermore, the location of the virtual bicycle path is known, as described in the first step. Using this information, the coordinates of the bicycle can be checked against the coordinates of the virtual bicycle path, and if the coordinates of the bicycle lie outside the virtual bicycle path its boundaries, a warning is given. For these warnings, two types are defined: (1) If the location is outside the soft borders, but inside the hard borders, a warning informing the cyclist that they should soon move back in to the safe zone. (2) If the location is outside the hard borders, a more severe warning is raised, urging the cyclist to immediately move back in to the safe zone.

4.4.3 Bicycle trajectory prediction. For predicting the bicycle trajectory, the current, as well as the last six measurements are used. As already mentioned in section 4.3, every 8 frames the bicycle location is tracked, resulting in three measurements per second. Therefore, to get the past three seconds of the bicycle location, the last 9 measurements are used.

For the prediction itself, it should predict the trajectory for up to three seconds in the future, meaning that there are nine x coordinates, starting at 300, and ending at 3000. This results in approximately three coordinates per second, which is the same rate at which measurements are taken.

The past measurements are plotted, and a linear regression is made using these coordinates. A linear regression is a method to predict the value of a variable based on past measurements. It can predict values by attempting to model the relationship between two variables, by fitting a linear equation to observed data [Yale 1997]. Using this linear regression model and the coordinates described in the paragraph above, the next nine coordinates can be calculated and plotted.

Once the coordinates have been plotted, they are compared against the coordinates of the virtual bicycle path. If the coordinates of one point lie outside the virtual bicycle path its boundaries, a warning is given.

5 RESULTS

5.1 Detecting bicycle paths and roads

The results for research question 1 consist of an algorithm to detect the road or bicycle path.

In figure 6 and 8, the input image can be seen. This is a frame isolated as described in section 4.3, from footage taken as described in section 4.2. On the images, a road with either a suggested bicycle lane or red bicycle lane can be seen. Suggested, in this context, means that cyclists are not required to cycle here.

In figure 7 and 9, the output of the algorithm is visualised, where the detected road surface or bicycle path are marked with light-blue. As can be seen, the edges are not completely straight. This can be caused by either debris on the road or by the dirt being similar in colour to the road surface.

Fig. 6. Grey road with suggested bicycle path

Fig. 7. Road with suggested bicycle path with road surface detected

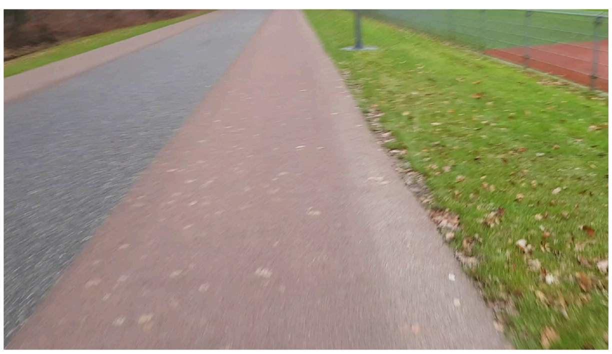

Fig. 8. Road with red bicycle path

In order to test the accuracy of the road surface detection algorithm, for both road types (grey and red), a number of frames are put through the algorithm. They are checked for accurately detecting only the road surface. A surface is marked accurately if both edges are correct, as well as at least the bottom half of the road vertically is recognized. For This purpose, the bottom half is adequate, as the bicycle cannot realistically travel further between two frames.

As can be seen in table 1, the accuracy differs between the grey and red road surface. This is mainly due to the Region Of Interest extraction patch areas for red roads being smaller, which makes them more sensitive to noise from the surroundings.

Fig. 9. Road with red bicycle path detected

Table 1. Accuracy of road surface detection

Road surface	Number of frames	Accuracy
Grey	84	95%
Red	93	68%

5.2 Creating a virtual bicycle path

The results for research question 2 consist of a virtual bicycle path created by an algorithm as described in section 4.4.2, as well as an algorithm that checks the current bicycle location against the virtual bicycle path. If the bicycle is no longer in the safe zone, a warning as described in section 4.4.2 is raised. An example of those messages can be seen in figure 12

The accuracy of the virtual bicycle path fully depends on the accuracy of the detected road mask. Therefore, no additional accuracy testing is performed. Additionally, the location detection also relies on the coordinates following from the detected road mask, and comparing these to the current bicycle location. Therefore, no specific testing for this step is done. The result of this step can be seen in figure 10 and in figure 11. As can be seen in these figure, there are multiple lines visible. These are drawn to visualise the different areas of the virtual bicycle path. The area between the yellow lines is the safe zone, while the red lines mark the edge of the bicycle path. How these are calculated is described in section 4.4.2.

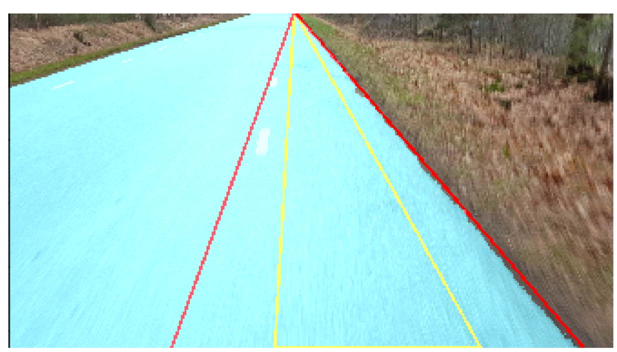

Fig. 10. Road with virtual bicycle path

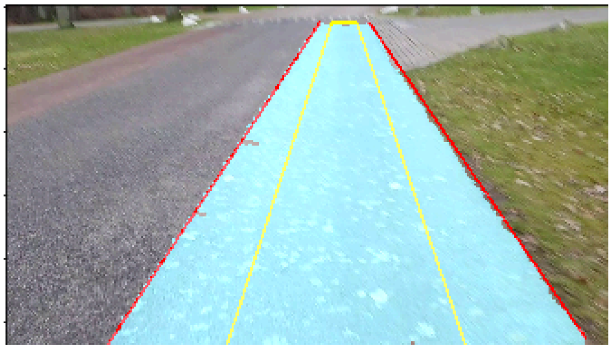

Fig. 11. Bicycle lane with virtual bicycle path

```
No action required
No action required
Warning: turn left soon!
Warning: turn left soon!
Warning: turn left soon!
Warning: turn left now!
```

Fig. 12. Example of given (warning) messages

5.3 Bicycle trajectory prediction

For research question three, the result is a trajectory prediction in the form of a simple prediction algorithm to predict whether the bicycle will stay on course with the current trajectory. In its current form, the prediction algorithm uses a linear regression. An example of this plot can be seen in figure 13. On the x-axis, the time is plotted, and on the y-axis, the location of the bicycle.

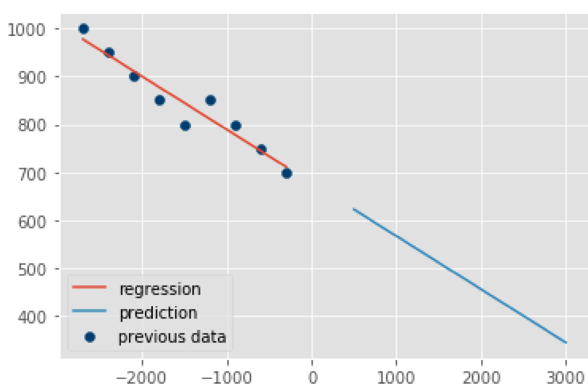

Fig. 13. Predicted trajectory

For testing the prediction model, 150 random bicycle routes were generated with each containing 19 points. For each route, the trajectory was predicted for the next three seconds. After that, the status of each predicted second (1, 2, 3) was compared to the actual route. From this, three confusion matrices (one for each predicted second) can be constructed showing the accuracy of the prediction algorithm. These can be found in respectively table 2, 3 and 4. Following from these matrices, the accuracy can be found in table 5. Accuracy is calculated as:

Table 2. Confusion matrix for results of 1st second trajectory prediction

		Actual	
		Safe	Not safe
Predicted	Safe	41	56
	Not safe	20	33

Table 3. Confusion matrix for results of 2nd second trajectory prediction

		Actual	
		Safe	Not safe
Predicted	Safe	11	60
	Not safe	9	70

Table 4. Confusion matrix for results of 3rd second trajectory prediction

		Actual	
		Safe	Not safe
Predicted	Safe	18	32
	Not safe	31	69

Table 5. Accuracy of trajectory prediction

Second	Accuracy
1	49%
2	54%
3	58%

$$accuracy = \frac{TP + TN}{TP + TN + FN + FP}$$

In table 5 can be seen that the accuracy of the prediction is very low, which can have two reasons: (1) Using a linear regression to predict a trajectory of a bicycle is not accurate, or; (2) The simulated bicycle routes are not realistic. Furthermore, the False Positive rate (where the prediction is safe while it actually is not safe) is very high. This is problematic, as the system is supposed to warn the cyclist for danger, but will not in these cases, as it does not classify the situation as a dangerous trajectory.

6 CONCLUSION

6.1 General conclusion

Using computer vision to detect a road surface can be done with 95% accuracy. Furthermore, it is possible to detect a red bicycle path correctly with 68% accuracy. For both of these, there are some limitations, which are described in section 6.2. Moreover, a virtual bicycle path can be constructed using this detected road surface, given it is detected correctly. Finally, predicting a bicycle trajectory can be done with more or less 50% accuracy using a linear regression, which is not adequate, as dangerous situations can arise from wrong predictions. Therefore, the proposed solution can help cyclists stay on the safe part of the road, but not predict if the current trajectory is safe.

6.2 Limitations

Regarding the limitations of this research, there are a few aspects to be discussed.

First, as the proposed technology is vision-based, it will only function with adequate light, which means it only works during daytime.

Second, this method does not rely on road features such as white lane markings, but rather on the road surface its color. As a result, the algorithm its accuracy is greatly reduced when working with wet, reflective roads.

Third, this method does not differentiate between varying road widths, except for 'whole road' and 'designated bicycle path'. Therefore, the algorithm uses a fixed width for the safe zone when there is no designated bicycle path, not taking into account the other road users and their required share of the road.

6.3 Further work

For further work on this proposed technology, a few aspects could be improved.

First, the detection algorithm could be optimized, as it takes one to two seconds from input to output. This time would probably increase as a bicycle-mounted computer would not be as quick as the computer this algorithm was developed on. As a result, while the algorithm works, it is currently not suitable for use on a regular bicycle.

Second, a different approach could be implemented for determining the Region of Interest for roads with a red surface, as the current implementation is prone to noise. Therefore, different Region of Interest patches could be defined and compared against each other.

Third, as already briefly discussed in section 6.1, using a linear regression for prediction the trajectory is inadequate. Therefore, in a future study, research could be done in what the best bicycle path trajectory prediction model could be.

Finally, as mentioned in section 6.2, further work could be done to better estimate the road width. This in order to create a more balanced virtual bicycle path based on fair road area usage.

REFERENCES

CBS. 2019. Hoeveel fietsen inwoners van Nederland? https://www.cbs.nl/nl-nl/visualisaties/verkeer-en-vervoer/personen/fietsen. Accessed: 2022-11-23.

CBS. 2022. Meer fietsdoden na eenzijdige ongevallen. https://www.cbs.nl/nl-nl/nieuws/2022/37/meer-fietsdoden-na-eenzijdige-ongevallen. Accessed: 2022-11-20.

Fietsberaad Crow. 2015. Aanbevelingen fiets- en kantstroken. https://www.fietsberaad.nl/getmedia/68d3ad2c-88d2-4634-8279-d33349ae5c74/Fietsberaadnotitie-Aanbeveling-fiets-en-kantstroken.pdf.aspx?ext=.pdf. Accessed: 2023-01-21.

Md. Rezwanul Haque, Md. Milon Islam, Kazi Saeed Alam, Hasib Iqbal, and Md. Ebrahim Shaik. 2019. A Computer Vision based Lane Detection Approach. https://www.mecs-press.org/ijigsp/ijigsp-v11-n3/IJIGSP-V11-N3-4.pdf. Accessed: 2022-11-24.

Fabrice Meriaudeau. Muhammad Muzammel, Mohd Zuki Yusoff. 2017. Rear-end vision-based collision detection system for motorcyclists. *Journal of Electronic Im* 26 (2017).

Yang Xing, Chen Lv, Long Chen, Huaji Wang, Hong Wang, Dongpu Cao, Efstathios Velenis, Fei-Yue Wang, and IEEE. 2018. Advances in Vision-Based Lane Detection: Algorithms, Integration, Assessment, and Perspectives on ACP-Based Parallel Vision. *IEEE/CAA Journal of Automatica Sinica* 5, 3 (2018), 645–661.

Yale. 1997. Linear Regression. http://www.stat.yale.edu/Courses/1997-98/101/linreg.htm. Accessed: 2023-02-02.

A APPENDIX A

A.1 Pixels and centimeters

Using the reference image that is pictured in figure 6, pixels and centimeters can be converted, as both the pixel count and the dimensions of certain features in the picture are known. The camera used throughout the research for data collection has been constant. Using a different camera will yield different values! From this follows that $1centimenter : 1.7pixels$ This is only true for the bottom of the image because of perspective.