

6-2005

Collision Warning and Sensor Data Processing in Urban Areas

Christoph Mertz
Carnegie Mellon University

David Duggins
Carnegie Mellon University

Jay Gowdy
Carnegie Mellon University

John Kozar
Carnegie Mellon University

Robert MacLachlan
Carnegie Mellon University

See next page for additional authors

Follow this and additional works at: <http://repository.cmu.edu/robotics>

 Part of the [Robotics Commons](#)

Published In

.

This Conference Proceeding is brought to you for free and open access by the School of Computer Science at Research Showcase @ CMU. It has been accepted for inclusion in Robotics Institute by an authorized administrator of Research Showcase @ CMU. For more information, please contact research-showcase@andrew.cmu.edu.

Authors

Christoph Mertz, David Duggins, Jay Gowdy, John Kozar, Robert MacLachlan, Aaron Steinfeld, Arne Suppe, Charles Thorpe, and Chieh-Chih Wang

Collision Warning and Sensor Data Processing in Urban Areas

Christoph Mertz, David Duggins, Jay Gowdy, John Kozar, Robert MacLachlan, Aaron Steinfeld, Arne Suppé, Charles Thorpe, Chieh-Chih Wang

The Robotics Institute, Carnegie Mellon University, 5000 Forbes Ave. Pittsburgh, PA 15213, USA
Phone : 1-412-2683445, Fax : 1-412-268 7350
E-mail: cmertz@andrew.cmu.edu

Abstract: *Providing drivers with comprehensive assistance systems has long been a goal for the automotive industry. The challenge is on many fronts, from building sensors, analyzing sensor data, automated understanding of traffic situations and appropriate interaction with the driver. These issues are discussed with the example of a collision warning system for transit buses.*

Keywords: Collision warning, sensor processing, perception.

1. INTRODUCTION

For driver assistance systems and autonomous driving functions in vehicles, sensors are needed to provide the necessary information about the surrounding of the vehicle. Unfortunately there are no sensors which can directly measure the relevant quantity like “threat” or “dangerousness”. Instead, they measure distance, speed, color, etc. of objects around the vehicle. From these raw data, one needs to infer in what situation the vehicle is. It is helpful to divide the sensing into three steps. The first is the aforementioned direct *measurement* of physical quantities like distance, speed, color, etc. The second is *perception*, where the data points are segmented into objects, the objects are classified and tracked, and other quantities and qualities are deduced from the raw data. The third is *understanding*. The objects are related to each other, to the environment, to models of object behavior, and to the host vehicle in order to understand the situation.

Since sensing is not sufficiently reliable, one needs to make use of other means to get full or partial machine driving. The first option is infrastructure. Here the problem is simplified by simplifying the environment. An example is autonomous trains often found at airports. Railroad tracks keep the trains on their path and physical barriers ensure that no pedestrians or other objects cross the path of the train while it is moving. The second option is to leave the driver in a supervisory role; he has to make the complex decisions or has to take over when the system is at a loss. An example is ACC, where the driver has to keep the vehicle in the lane and has to take over when the vehicle is approaching another vehicle too fast. The third option is to leave all the actuation to the driver and display only

more or less sophisticated auxiliary information to him. Parking aids fall into this third category. Yet another option is to have every object tagged so that its properties can be read remotely. This would considerably simplify the sensing.

In this paper we will explore the issues mentioned in the preceding two paragraphs with the example of a collision warning system for transit buses [1]. This project arose from the merger of two other projects, one developing a forward collision warning system [2], the other a side collision warning system [3]. We will focus on the side sensing.

2. DRIVING ENVIRONMENT

There is a great variety of situations a driver is exposed to and safety and assistant systems can be quite different for different classes of situations. It is interesting to note that situations at the high and the low end of speeds are easier to handle than the medium speeds. At the high end we find adaptive cruise control and at the low end there are parking aids. While driving at medium to high speeds on interstates one can assume a fairly simple surrounding, only other vehicles are on the street and fixed objects are on the side of the road. While parking the speed is so low, that all the other objects can be assumed to be fixed. The speed range for driving in urban areas is in the middle, the most difficult range. Vehicles, bicyclists, and pedestrian are on the street with various velocities and objects on the side of the road can not be ignored.

In addition, there are several things which point to the specific challenges faced by a transit bus [4]:

1. Many of the most serious accidents involve pedestrians.
2. Only a very small percentage of side collisions are classical lane change or merge accidents.
3. Many of the bus accidents involve objects approaching from the side.
4. The line between safe and unsafe situations is very tight.
5. In a quarter of all pedestrian fatalities, the pedestrian is partially or completely underneath the bus.
6. In many cases the bus driver does not notice that a collision with a pedestrian happened.

- In most cases it is not the bus driver who created the dangerous situation.

One line which separates safe from unsafe situations is the curb. If a pedestrian is on the sidewalk, he or she can be considered much safer than if he or she is on the street, even if in both situations the distance and relative speed to the bus is the same.

3. MEASUREMENT: CHOICE OF SENSORS

From the analysis of the driving environment it became clear that the sensors of the warning system need to be able to detect large and small objects like pedestrians, mailboxes, and vehicles. Location and velocity of the objects need to be determined with good accuracy and the objects need to be classified. Another requirement is that the curb position can be measured. We found that the best sensor for the object detection is a laser scanner. We choose a Sick™ laser scanner. As the 180° field of view allows only using one per side. As we will discuss in the following sections, the laser scanner was sufficient for our project, but it also had some shortcomings.

To determine the location of the curb we developed a triangulation sensor [5]. It consists of a camera and a laser.

Finally, to evaluate the performance of the system we need mounted cameras on the bus, two on each side. Figure 1 shows images and the locations of the various sensors and the computers.




Figure 1: Location of the sensors and the computers on the bus.

In addition to the sensors on the exterior of the bus to observe the environment around the bus, we tapped into the internal data bus to acquire speed and status information (e.g. door open/closed). A gyroscope provided the yaw-rate of the bus.

4. PERCEPTION

In the perception phase the measured raw data are analyzed to extract the desired information about objects and the environment around the bus. We have one perception module for objects and another one to detect the location of the curb.

4.1. Detection, Tracking, and Classification of Objects

The raw data provided by the laser scanner consists of the distances of 181 points at intervals of 1° (see Figure 6). Following operations are performed on this data:

1. Transformation

The data points are transformed into the fixed world coordinate frame. In this frame the apparent movement of the objects is not influenced by the movement of the bus, e.g. fixed objects do not move.

2. Segmentation

The data points are segmented into objects. The criterion of a point belonging to an object is that its distance to the closest point is below the threshold of 0.8m.

3. Line fitting

An attempt is made to fit a corner or a line to the points of an object. An example can be seen in Figure 2 where a corner is fitted to points outlining a car. The end points of the line(s) are the features used for tracking.




Figure 2: A corner fitted to an object.

4. Noise estimation

The lateral error in the position of the feature points is estimated from the quality of the fit, and the longitudinal error is determined max inter-point spacing of the last few points on the line.

5. Data association

The current objects are associated with prior objects based on proximity and similarity. The motion of objects is estimated with the help of a Kalman filter. Inputs to the filter are the positions of the feature points and the estimated noise.

6. Track evaluation

The validity of the dynamic quantities are assessed by checking if they are consistent with the position of the object further in the past.

7. Classification

The shape (corner, line, or neither) and movement of the objects are used to classify them. Vehicles are corners or lines of the right size, pedestrian are small and slow moving.

There are many more relevant details about the detection and tracking algorithm [1]. For example, decisions must be made on which conditions an object is terminated, how to handle inconsistent data, what to do when objects get occluded, etc.

We determined the quality of the velocity estimation by driving the bus past fixed objects like parked cars. The deviation of the velocity from zero is the error in the measurement. The result can be seen in Figure 3.




Figure 3: Error distribution of the velocity estimation.

The error distribution can be fairly well described by a Gaussian with a standard deviation of 0.13 m. However, there are a few outliers. As we will see later, they are cause for some false alarms.

4.2. Curb Detection

The triangulation sensor observes the area directly to the right of the right front corner of the bus. A sample snapshot of what the sensor sees is shown in Figure 4.




Figure 4 Profile of the road and curb observed by the triangulation sensor. Some erroneous readings can be seen above the road.

A fairly robust way to find the curb even in the present of considerable noise is to look at the number of points in a horizontal bin. At the location of the curb there are a large number of points at the same horizontal position. This can be visualized by a histogram (Figure 5).




Figure 5: Histogram of the number of data points versus horizontal distance. The line is the detection threshold.

The position of the curb can now be determined by applying a threshold on the histogram.

Since we also measure the speed and yaw-rate of the bus while it is driving, we are able to map the position of the curb alongside bus. This can be seen in Figure 6.

An in-depth description of this algorithm and a system that uses the data from two more sensors to detect the curb in front of the vehicle can be found in reference [6].

5. SAMPLE DATA SET




Figure 6: Display of the raw and inferred data. The bus is shown from top, the raw laser scanner data are shown as points, the objects are square boxes, and the curb is a blue line. One box is yellow, indicating an alert.

In Figure 6 and Figure 7 the raw and inferred data are visualized. It is a snapshot of our replay and analysis tool. In Figure 6 one can see the various objects and the curb detected by the sensors. The system also knows that the bus is turning left and infers that it is in danger of colliding with the parked vehicle. The probability of the collision is not very high, so an “alert” is issued.

The same situation is displayed in four video images with the data overlaid (Figure 7).



Figure 7: Data overlaid on the video images.

6. UNDERSTANDING

Now that we have all the information about the bus itself and its surrounding, we need to calculate a measure of how dangerous the situation is and then use this understanding to issue appropriate warnings.

In many warning systems the measure is time-to-collision (TTC) or distance-to-collision (DTC) and the warning is issued if the distance in space or time is below a certain threshold. This is a good approach if one considers the 1-dimensional case of one vehicle following another (Figure 8a).




Figure 8: a) A vehicle follows a motorcycle. b) The vehicle and motorcycle are next to each other.

In 2-dimensional cases TTC or DTC approaches have difficulties. The example in Figure 8b where a motorcycle travels next to a vehicle can illustrate the

difficulty. The special distance is very short, but the temporal distance is very long, in fact it is infinite. A small change in the direction of travel of the motorcycle can drastically alter the dangerousness of the situation.

6.1. Probability of collision

We decided to take a different approach and calculate the probability of collision (POC) as a measure of the danger of a situation. For this we take into account the speed and yaw-rate of the bus, the position, velocity and classification of the object, models of bus and objects behavior, and the location of the curb. Reference [7] describes the algorithm in detail, here we illustrate it with an example.

On the left side of Figure 9 one can see a bus making a right turn while an object travels from left to right. On the left side the same situation is transformed into the fixed bus frame. The bus is fixed and the trajectory of the object is the result of the relative motion of the object. To calculate the POC we randomly generate trajectories. The center trajectory is determined by the measured dynamic quantities. The distribution of the trajectories is the result of the uncertainty of the measurements and models of bus and object behavior. For each trajectory we determine if and when a collision happens.




Figure 9: Example of the bus turning right while an object travels from right to left. The situation is shown in two different frames. The point clouds on the right are the distribution of location of the objects for three different times.

The point clouds on the right side of Figure 9 are the distributions of positions of the object for three different times. Red points mean that a collision has happened. The ratio of red points to all points is the POC. The blue line in Figure 10 indicates the POC between 0s and 5s for the example shown above.

6.2. Warning generation

The POC is the measure of danger and the warnings are determined by areas in the POC vs. time graph. As illustrated in Figure 10, the green area is where on warning is given, yellow is the area for “alerts”, and red

is for “imminent warnings”. “Alerts” are warnings which should draw the attention of the driver in a non-intrusive way. “Imminent warnings” are more aggressive and are given for situations where the POC is high. In our example the POC (blue line) reaches into the yellow area and therefore an ‘alert’ is in order.

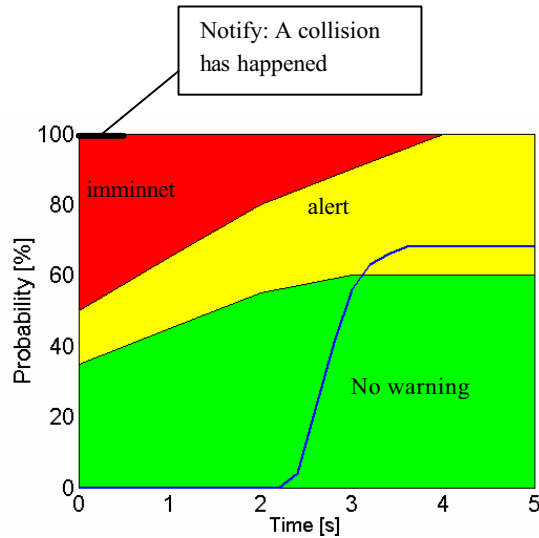


Figure 10: Probability of collision versus time.

As we mentioned in Section 2, the driver sometimes does not notice that a collision has happened. Therefore he needs to be notified if such an event took place. The criterion for issuing a “notify” is that the POC is 100% within the next ½ second

The last category of warning is “under the bus”. The most dangerous situation is when a person slipped under the bus and therefore warrants the highest level of warning. It is even a higher level than “notify”, because if a collision has happened, the driver can obviously not prevent or mitigate it any more. He is notified so he can attend to the accident. The “under the bus” warning does not fall nicely into the described framework of POC. The warning is issued whenever a person is detected underneath the bus.

7. DRIVER VEHICLE INTERFACE

Once the system decided that a warning needs to be issued, it has to be conveyed to the driver in an appropriate way. The design of the driver vehicle interface (DVI) needs to incorporate the four warning levels mentioned above and the warnings issued by the forward part of the warning system. The DIV is a modification of a design developed for snowplows [8] for forward warnings and has been extended to include the side warnings (Figure 11). On each side of the driver is a column of 7 LEDs with two triangles underneath them. The 7 LEDs are for the forward warnings. These grow down from the top as the threat level increases.

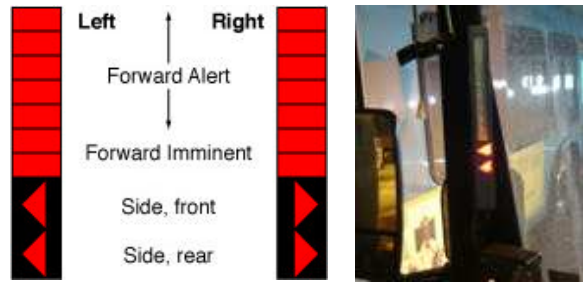


Figure 11: The DVI. On the left is a schematic of the arrangement of the LEDs. On the right is an image of one of the LED bars with only the side warning triangles lit.

For the side, the triangle corresponding to the location of the object (left/right and front/rear side) is lit in following way:

- 1) *Alert*: Yellow.
- 2) *Imminent Warning*: Red.
- 3) *Notify*: The triangles blink yellow.
- 4) *Under the bus*: The triangles blink red.

The DVI does not obstruct the view of the driver, the bars are mounted on the side post (see Figure 11 right) and middle post of the window. Warnings are designed to draw the driver’s attention in the direction of the threat.

8. RESULTS




Figure 12: Density of side warnings around the bus. Two areas are shown, one containing 80% of all warnings and the other 98%.

We have installed the system on two buses and they were used during normal operations for almost a year. We have collected several Tb of data from the hundreds

of hours of operation corresponding to thousands of miles traveled. We used this data to test, refine, and evaluate our system.

Figure 12 shows the density of side warnings (location of the threat) around the bus. The area can be approximated by a half-circle in front of the bus and a rectangle to the left and right of the bus. This image does not document the warnings generated by forward warning component.

We also investigated the false warnings produced by the system and tuned the system in such a way that there are very few false negative warnings (situations where the system did not warn but should have). We found several reasons for false positive warnings (system warned but should not have). The most significant ones are:

- 1) *Vegetation*: Gras, bushes, or branches cause a warning. The system functions correctly but the driver considers the warning a nuisance.
- 2) *Wrong velocity*: The object has the wrong velocity (see Section 4.1).
- 3) *No velocity*: The system needs about ½ second to establish the velocity of an object.
- 4) *Water splashes*: In heavy rain water can splash which is seen by the laser scanner as an object. When the splash is produced by the front tire of the bus, the object appears to be right next to the bus.
- 5) *Ground return*: The laser scanner sees the ground and the system thinks there is a real object.

These five reasons are ordered by severity. In the first one there is an actual object which might collide with the bus, whereas with ground return there is not even a real object. In all five areas improvements are possible through enhancements of the sensor (see [9]) and performance (e.g. [10]).

9. CONCLUSION

Developing the collision warning system for transit buses showed us that building a warning system suitable for driving in urban areas is a great challenge in all areas, from sensing, perceiving, understanding, to interacting with the driver. Other assistance or autonomous system might not face the same specific problems, but they all will have their own challenges in these areas.

10. ACKNOWLEDGEMENT

This work was supported in part by the U.S. Department of Transportation under Grant 250969449000 and PennDOT grants PA-26-7006-02 and PA-26-7006-03.

REFERENCES

- [1] "Integrated Collision Warning System Final Technical Report", FTA-PA-26-7006-04.1, in press.
- [2] Wang, J. Lins, C.-Y. Chan, S. Johnston, K. Zhou, A. Steinfeld, M. Hanson, and W.-B. Zhang, "Development of Requirement Specifications for Transit Frontal Collision Warning System," (UCB-ITS-PRR-2003-29). Richmond, CA: University of California, Partners for Advanced Transit and Highways, November, 2003.
- [3] C. Thorpe, D. Duggins, S. McNeil, and C. Mertz, "Side Collision Warning System (SCWS) Performance Specifications for a Transit Bus," Final report, prepared for the Federal Transit Administration under PennDOT agreement number 62N111, Project TA-34, May, 2002.
- [4] C. Mertz, S. McNeil, and C. Thorpe, "Side collision warning systems for transit buses," IV 2000, IEEE Intelligent Vehicle Symposium, October, 2000.
- [5] C. Mertz, J. Kozar, J.R. Miller, and C. Thorpe, "Eye-safe laser line striper for outside use," IV 2002, IEEE Intelligent Vehicle Symposium, June, 2002.
- [6] R. Aufrère, C. Mertz, and C. Thorpe, "Multiple sensor fusion for detecting location of curbs, walls, and barriers," Proceedings of the IEEE Intelligent Vehicles Symposium (IV2003), June, 2003.
- [7] C. Mertz, "A 2D collision warning framework based on a Monte Carlo approach," Proceedings of ITS America's 14th Annual Meeting and Exposition, April, 2004.
- [8] A. Steinfeld and H.-S. Tan, "Development of a driver assist interface for snowplows using iterative design," Transportation Human Factors, vol. 2, no. 3, pp. 247-264, 2000.
- [9] W. C. Stone, M. Juberts, N. Dagalakis, J. Stone, J. Gorman, "Performance Analysis of Next - Generation LADAR for Manufacturing, Construction, and Mobility." NISTIR 7117; 198 p. May 2004.
- [10] A.E. Broadhurst, S. Baker, and T. Kanade, "Monte Carlo Road Safety Reasoning," IEEE Intelligent Vehicle Symposium (IV2005), IEEE, June, 2005.