

Driving in Traffic: Short-Range Sensing for Urban Collision Avoidance

Chuck Thorpe, Dave Duggins, Jay Gowdy, Rob MacLaughlin, Christoph Mertz, Mel Siegel, Arne Suppé, Bob Wang, Teruko Yata

Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213-3890
Tel.: 01-412-2683612
Fax: 01-412-2687350
thorpe@ri.cmu.edu

Abstract

Intelligent vehicles are beginning to appear on the market, but so far their sensing and warning functions only work on the open road. Functions such as run-off-road warning or adaptive cruise control are designed for the uncluttered environments of open highways.

We are working on the much more difficult problem of sensing and driver interfaces for driving in urban areas. We need to sense cars and pedestrians and curbs and fire plugs and bicycles and lamp posts; we need to predict the paths of our own vehicle and of other moving objects; and we need to decide when to issue alerts or warnings to both the driver of our own vehicle and (potentially) to nearby pedestrians.

No single sensor is currently able to detect and track all relevant objects. We are working with radar, lidar, stereo vision, and a novel light-stripe range sensor. We have installed a subset of these sensors on a city bus, driving through the streets of Pittsburgh on its normal runs. We are using different kinds of data fusion for different subsets of sensors, plus a coordinating framework for mapping objects at an abstract level.

Introduction

Our group at Carnegie Mellon University is developing intelligent vehicles for driving in cluttered urban environments [1]. Our immediate applications are in driver warning systems: sensing objects that are potential causes of collisions, and giving appropriate information to drivers in time to avoid or mitigate a crash. Much of our work is focused on side-looking sensing for transit busses [2]. Our longer-term objectives include fully automated driving, including automation of future military vehicles that will have to work in urban environments during peacekeeping operations.

Current sensing/warning/controlling systems generally work only in relatively simple environments: Applications developed for open highways include Adaptive Cruise Control (ACC), which controls the throttle to

keep a safe gap behind other vehicles; run-off-road collision warning systems, which alert a driver if the vehicle starts to drift out of its lane; and blind-spot sensors on heavy trucks to warn the driver if they start a lane change without seeing a car in the next lane. Some applications are also on the market for slow speed driving: rear-facing sensors as parking aids, for example.

Urban driving raises much more difficult issues than those systems are designed to handle [3]. There are many objects, and many different kinds of objects (vehicles, bicycles, pedestrian, animals, mailboxes, sign posts, ...). There is a great variety of situations (intersections, pedestrians, parked cars, traffic lights...). The vehicle itself and surrounding objects go through a great variety of manoeuvres and actions (vehicle drives backwards into parking spot, person runs to catch a bus, children playing in the street, ...).

Simply applying a system designed for highways will not work. A side-looking sensor designed to cover the blind spot of a heavy truck needs to look for large metal objects (cars) within a range of a few meters (one lane width). The system can have a few spot sensors, spaced along the truck, with gaps in coverage up to nearly the length of a car, and still not miss any important objects. The same system, in an urban environment, will miss small objects (lamp posts, mail boxes, pedestrians) which can be causes of significant collisions. Just as importantly, it can cause numerous false alarms. Transit busses often operate close to parked cars, fixed objects, and their own passengers; if the system triggers an alert every time an object is within a lane width, it will be generating nearly continuous nuisance alarms.

The requirements for a useful collision avoidance and warning system for urban environments, then, include at least the following:

1. *Significant sensing*: The sensor system on an intelligent urban vehicle will need multiple sensors, looking at many different sorts of phenomena. Kinds of sensing include at least the following:

- a. *State of own vehicle*: is it moving, turning, or accelerating; are there indications of changes in state (brake lights, turn signals); is it operating on a fixed route; etc.
 - b. *State of nearby objects*: locations of objects; classification as fixed or moving; speed of moving objects; type of moving objects (car, pedestrian, other).
 - c. *Environment*: cross streets, road lanes, pedestrian crosswalks, traffic signals, visual occlusions, road traction, etc.
2. Knowledge base: The system must have adequate models to interpret sensed data. Examples include:
- a. *Model of the own vehicle and driver*: how fast is the driver's reaction time, how quickly can the vehicle turn or brake, how wide and long is the vehicle, where are blind spots where the driver has poor visibility?
 - b. *Model of other objects*: what are typical motion patterns of vehicles and pedestrians, how are those patterns affected by sidewalks and lanes and crosswalks?
 - c. *Model of environment*: for a transit bus, what is its intended route; where are particularly dangerous corners with occluded vision, where are vehicles likely to approach from with high speed, where are school zones?
 - d. *Model of interaction between all of the above*: how do pedestrians usually behave in the presence of other vehicles, crosswalks, etc.
3. Processing and algorithms: Given the sensor data and the knowledge base, intelligent real-time algorithms must interpret the situation:
- a. *What situation are we in*: what are the nearby objects and their likely future positions, what is the vehicle's state and likely future position?
 - b. *How likely is a collision*: what is the collision probability given the trajectories of all objects in the environment?
 - c. ~~How~~ *How dangerous is the situation*: if a collision is likely, is there adequate time for driver intervention?
 - d. *Is an action needed*: if so, what action, and how likely is the driver to take that action with or without assistance from the system?
4. System response: For a fully automated system, the response will be direct control of vehicle motion to avoid collision. For a driver

assistance system, the response will be through a driver interface. A typical interface will have to have graduated levels, plus ancillary functions, including information for third parties, e.g. pedestrian:

- a. *Aware: Baseline Situational Awareness*. The transit operator and pedestrian see strictly non-intrusive indications be they bumper stickers, running lights, video or the lack of any active alerts, warnings, evasions, or notifications.
- b. *Alert: Potential Obstacles*. Alerts are semi-intrusive information such as enhanced video indicating potential obstacles, lights indicating the close proximity of an obstacle, or a pleasant voice alerting a pedestrian to the presence of a moving bus.
- c. *Warn: High Likelihood of Collision*. Warnings span the spectrum from intrusive information such as voice or melodic sounds to intrusive interference such as shaking the steering wheel and or seat, vibrating the brakes, or a loud buzzer all indicating a high likelihood of collision.
- d. *Evade: Imminent Collision*. Evasive actions include active control of the transit bus such as steering or applying the brakes.
- e. *Notify: Collision has occurred*. Notification involves informing the transit operator through an intrusive light or voice that a collision has occurred and data (either computer and/or video) has been saved.

Overall, this is a difficult task. "Human-level" performance is inadequate; we need to be even better than humans, in order to reduce or eliminate traffic accidents.

In order to make this task tractable, it is important to begin with appropriate sensors. We need multiple sensors to detect the different kinds of objects in the environment. Our group is investigating radar, lidar, stereo vision and a laser stripe range sensor. We want to discuss the laser stripe and its special use in greater depth.

Case study:

A particularly interesting and useful piece of information about the environment is the location of the curb. In most cases fixed objects are on the curb and do not pose a threat to the vehicle. Pedestrians are also usually on the curb and remain on the curb. In the few cases they are not on the curb the system should pay particular attention to them, because they are likely to

want to cross the road. For these reasons we choose curb detection as a case study for short-range sensing in urban areas.

Curb Sensing with laser line striper:

A detailed description of the laser line striper can be found in Ref. [4].

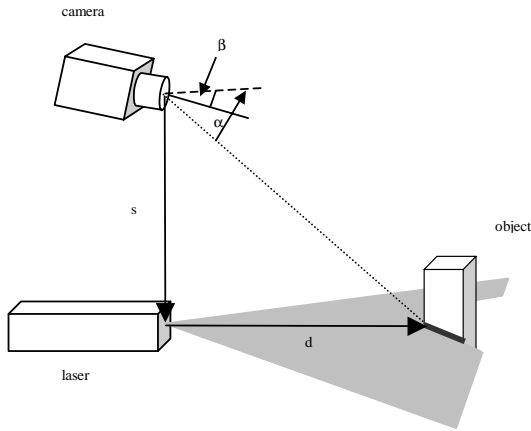


Figure 1 Configuration of laser and camera of the laser line striper.

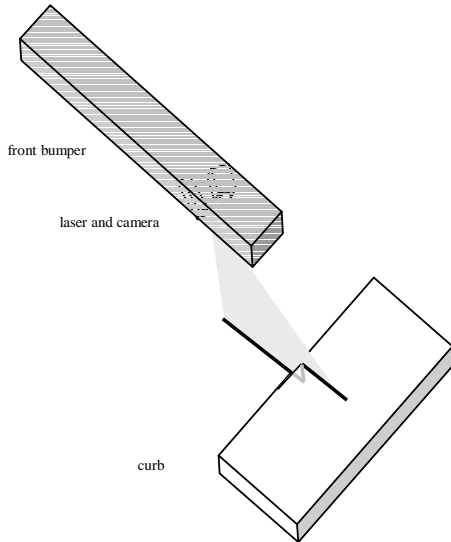


Figure 2 The laser and camera are mounted inside the front bumper of the bus and pointing towards the curb. The profile observed by the line striper is indicated as a thick line.

The working principle is shown in Figure 1. A laser projects a plane of light which intersects the objects in view of the sensor. A camera at a distance from the laser observes the reflection of the laser light and can calculate the distance to the object by triangulation. The amount of power the laser can emit is limited by eye-safety requirements. To make the system work

outside, where the sun produces a strong background of light, we employ two techniques to reduce the background. The camera is equipped with a narrow filter and the laser is pulsed and in sync with the fast shutter of the camera.

The laser and a camera with a field-of-view of 30° were mounted inside the front bumper of the bus and observing straight out the side of the bus (Figure 2).

An example of an observed profile of the curb is shown in Figure 3. The location of the curb was determined in the following way. The points on the left, which form the road, were fitted with a straight line. The first several points which deviated significantly from that line indicate the position of the edge of the curb.




Figure 3 Example of curb profile as seen by the laser line striper. The three arrows point to the road, junk in front of the curb, and the edge of the curb.

Figure 4 shows the bus-curb distance for a 7 second interval during which the bus approaches a bus stop and comes to a halt.




Figure 4 Distance bus-curb for a 7 second interval as measured by the laser line striper. The dashed line shows the uncorrected distances, for the solid line, outliers have been eliminated.

The distances are tracked very smoothly, except for several instances, where the distance is suddenly several centimetres shorter. These outliers are caused by junk (stones, leaves, dirt, etc.) in front of the curb (see Figure 3). We eliminated these outliers by removing all points which are more than 1 cm different than their previous point. The resulting points, joined by a

solid line in Figure 4, have very little variation; the standard deviation of the point-to-point differences is only 1.6 mm.

There are other possibilities to correct the distance measurements. One might develop a curb tracker which determines the most likely curb position from the past measurements and the knowledge of the error function (1.6 mm sigma Gaussian error plus the non-Gaussian error of up to 10 cm outliers). Another way would be to reduce the number of outliers by careful adjustment of the threshold, which determines when a point significantly deviates from the fitted line. The optimal threshold is dependent on the height of the curb. Since the height of the curb is usually uniform along a road, it could be either predetermined or measured on the fly.

So far we have only the distance of the curb to the front corner of the bus. To get the position of the curb alongside the bus the movement of the bus needs to be monitored. Then the position of the curb can be followed while it passes alongside the bus.

Data collection

The data for this discussion was collected on a transit bus in Pittsburgh, PA. The outside of the bus was equipped with a laser scanner which detected objects in a horizontal plane on the side of the bus, a camera monitoring that side, and a laser line striper which was mounted inside the front bumper for curb detection. Several other data were recorded, including speed, heading, GPS position, status of the door and the lights, etc.

In order to establish ground truth, for a time interval of 15 minutes we manually classified all objects at the right side of the bus with the help of the video images. Only objects within about 2 m of the side of the bus were counted. We distinguished between three bus manoeuvres: cruising, turning right, and being stopped at a bus stop. There were no objects close to the right side of the bus while the bus was turning left. The result is tabulated in Table 1.

About half of all the objects encountered are actually on the curb. The great majority of objects on the street are parked cars. Most of the pedestrians on the street are there because they are entering or leaving the bus. Only one pedestrian was on the street while the bus was cruising. The video showed that this person was obviously intending to cross the street. Of everything encountered in these 15 min, only this one person intending to cross the street warranted a safety level higher than the baseline “no danger” level.

Determining the severity of the situations

A proximity sensor and the curb detector together can determine if an object is on the curb or not. If one assumes that the situation is safe as long as things are on the curb or even on the edge of the curb, then already half of the situations in Table 1 can be determined to be safe.

Table 1 Objects at the side of the bus: manual tabulation.

	object \ bus movement	cruising	turning right	stopped
on curb	pedestrian	0	0	10
	mail box, trash can, etc.	1	0	1
	tree, telephone poll, etc.	10	4	0
	sign post, fence post, etc.	38	8	1
	fire hydrant	6	2	0
off curb	parked car	69	2	2
	pedestrian	1	0	8
sum		125	16	22
singular situations (bus was always cruising):				
hedge beyond the curb				
sign on the street				
8 barrels blocking off a bus stop				
pedestrian on the edge of the curb				
tree hanging over the curb				

Fixed objects are no hazard if they are on the side of the bus and the bus is going straight. A sensor which can determine if objects are moving and an internal sensor which measures the heading of the bus provide the data to allow the system to evaluate these situations, in our example the majority of the remaining cases.

Obviously, no warning or alert needs to be given to the driver when the bus is stopped, a situation easy determined by a speedometer.

To evaluate the two parked cars encountered by the bus while the bus was turning right one needs more information than provided by the table. One needs to look at the trajectory of the bus and see, if the sweeping motion of the bus will cause a collision.

The person standing on the street intending to cross the street warrants an ‘alert’. It is necessary to identify the pedestrian as such as well as determining that he is not on the curb. This case emphasizes the importance of determining the location of the curb. But the location of the object with respect to the curb needs to be combined with a pedestrian detection and the knowledge that most people on the street intend to cross it.

Conclusion and Outlook

The analysis of the 15 min of data already showed, that a great variety of sensing, knowledge, and analysis is necessary to correctly identify the situations as safe or give the driver the correct safety level. The curb detector plays an important role to distinguish the safe from the dangerous situations, while at the same time it can only function effectively if it works together with other sensors.

The next steps in our project are to analyse much more of the collected data, improve our sensors and algorithms and develop a more detailed knowledge base.

Acknowledgements

This work was partially sponsored by PennDOT under agreement # 62N111, project TA-34, funded by the USDOT Federal Transit Administration, by the Robert Bosch Corporation, by the Army Research Labs Collaborative Technology Alliance (CTA) Program, contract # D AAD 19-01-2-0012, and by DARPA, contract # MDA972-01-9-0015. The views and conclusions expressed are those of the authors and not necessarily those of the sponsors.

References

- [1] Thorpe, C., Clatz, O., Duggins, D., Gowdy, J., MacLachlan, R., Miller, J.R., Mertz, C., Siegel, M., Wang, C., and Yata, T., Dependable Perception for Robots. Proceedings of International Advanced Robotics Programme IEEE, Robotics and Automation Society, Seoul, Korea, May, 2001.
- [2] Mertz, C., McNeil, S., and Thorpe, C., Side Collision Warning Systems for Transit Buses. IV2000, IEEE Intelligent Vehicle Symposium, October, 2000.
- [3] McNeil, S., Thorpe, C., and Mertz, C., A new Focus for Side Collision Warning Systems for Transit Buses. ITS2000, Intelligent Transportation Society of America's Tenth Annual Meeting and Exposition, May, 2000.
- [4] Mertz, C., Kozar, J., Miller, J.R., and Thorpe, C., Eye-safe Laser Line Striper for Outside Use. Paper submitted to: IV 2002, IEEE Intelligent Vehicle Symposium, June, 2002. www.ri.cmu.edu/pubs/pub_3890.html .