# Safe Robot Driving in Cluttered Environments

Chuck Thorpe,  Justin Carlson, Dave Duggins,  Jay Gowdy,  Rob MacLachlan,
Christoph Mertz, Arne Suppe, Bob Wang,
The Robotics Institute, Carnegie Mellon University, Pittsburgh PA USA

### Abstract

The Navlab group at Carnegie Mellon University has a long history of development of automated vehicles and intelligent systems for driver assistance. The earlier work of the group concentrated on road following, cross-country driving, and obstacle detection. The new focus is on short-range sensing, to look all around the vehicle for safe driving. The current system uses video sensing, laser rangefinders, a novel light-stripe rangefinder, software to process each sensor individually, and a map-based fusion system. The complete system has been demonstrated on the Navlab 11 vehicle for monitoring the environment of a vehicle driving through a cluttered urban environment, detecting and tracking fixed objects, moving objects, pedestrians, curbs, and roads.[*]

## 1    The Need For 360 Degree Safeguarding

Robot driving has concentrated on forward-looking sensing, for road following and obstacle detection. This is an appropriate first step, but real deployment of mobile robots will require additional sensing and reasoning to surround the robot with safeguard sensors and systems. Our group is currently building short-range sensing to surround vehicles to improve the safety of robotic and human-controlled vehicles[1].

In the civilian context, our focus is driver assistance for transit busses. Busses drive at relatively slow speeds, in crowded urban environments. One of the most frequent types of accidents in transit busses is side collision, where the driver does not have adequate awareness of objects near the bus, then turns too sharply and sideswipes a fixed object or (less often) hits a pedestrian. Preventing these accidents requires short-range sensing along the side and front of the bus, detecting fixed objects, detecting and tracking moving objects, predicting the motion of the bus itself, and a suitable driver interface for alerting the driver.

In the military context, our focus is short-range sensing for full automation of scout vehicles. An autonomous vehicle moving through a cluttered environment, such as a forest, may need to move between objects (e.g. trees) with very little clearance on either side of the vehicle. The conventional approach is to sense trees with forward-looking sensors, enter those trees into a map, and estimate the clearance as the vehicle moves forward and the trees move out of the field of view of the forward-looking

sensor. If sensor data is noisy, or if the vehicle slips and slides in mud, the estimated clearance may be incorrect. It is better to directly and continuously sense nearby objects all along the side of the vehicle as it moves through the forest.

Both civilian and military vehicles, and both driver assistance and full automation, need to pay special attention to moving objects, and particularly to humans. People move in unpredictable ways. Seeing a person with a forward-looking sensor, or having the driver note the position of a person in front of a bus, is no guarantee that the person will remain safely out of the vehicle's way. Thus, we pay special attention to detecting moving objects, classifying them if possible as people, continuously tracking their motion, and attempting to predict their future actions.

This work is part of a trend in the intelligent vehicle research community, focusing on better sensing for driver assistance. The Intelligent Vehicles 2002 Conference, for instance, included a number of papers on pedestrian detection [2] [3] [4] [5] [6] and on urban driving [7] [8].

## 2    Testbed Vehicle

The Carnegie Mellon testbed vehicle for this work is the Navlab 11, shown in Figure 1. It is equipped with:



**Figure 1: Navlab 11 testbed**

- Motion sensors (IMU, GPS, differential odometry,  compass, inclinometer, and angular gyro);

- Video sensors (5 video cameras, 3 in front and two looking along the sides);
- Ladars (3 Sick single-axis scanning laser rangefinders mounted in various positions, typically one looking forward and two along the sides);
- A light-stripe rangefinder typically mounted on the bumper looking for the curb on the right;
- Five 500-MHz Pentium computers;
- High-accuracy time synchronization;
- And various other sensors and actuators.

## 3  Perception Modules

We have built, tested, and integrated a number of perception modules, designed specifically for short-range sensing.

*Light-stripe scanner*: We have designed, built, and integrated a light-stripe sensor as a proof-of-principle of a no-moving-parts range sensor, designed for use outdoors at modest ranges[9]. We tested the sensor on board a transit bus operating in the city of Pittsburgh under a full range of weather conditions. The idea of light-stripe sensing is well-known: shine a plane of light (in this case, a vertical plane) on the scene, then look from a different vantage point for the shape that the light makes as it hits the surface. Triangulation gives the 3-D scene geometry. The novel additions we have made to the process enable range sensing under full sun outdoors. First, we filter the camera, so that it only sees light in the same wavelength as the laser illumination. Second, we shutter the camera, in sync with a pulsed laser. That allows much higher instantaneous power output, while still having a total integrated emitted energy low enough to keep the laser eye safe (Figure 2)
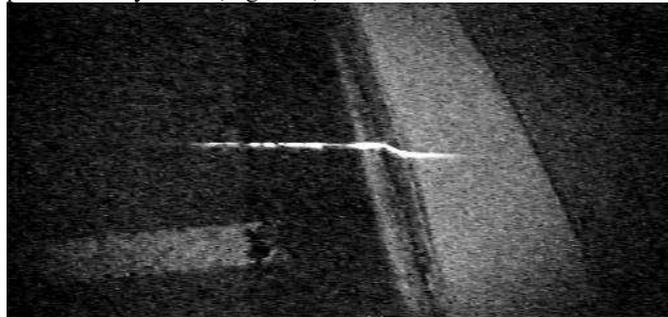


**Figure 2: Curb detection with light stripe sensor. The filtering and shuttering reduce the background illumination while leaving the infrared light stripe clearly visible. The shape of the stripe shows the contour of the street, curb, and sidewalk.**

*Curb detection and tracking*: In use, our side-looking light stripe range sensor measures the vertical profile of the road to the right of the testbed. When there is a curb present, the light stripe sensor can find the curb out to a range of approximately one lane width. We have tuned the software to detect the kinds of curbs commonly encountered in urban environments. We also built a Kalman filter to integrate curb measurements and vehicle motion models over time, to create a track of where the curb has been and an estimate of its current heading.

*Curb and road prediction:* Once the curb has been detected and tracked, we project its location into the video image acquired by the camera mounted on the rear right of the testbed looking along the vehicle. Given the location of the curb in this image, an image processing algorithm generates a template of the curb's appearance and tracks it forward in the image to generate a preview of where the curb and road are headed (Figure 3). Tracking the curb is especially important for pedestrian safety: a pedestrian on the sidewalk is likely to stay safely on the sidewalk, while someone who has stepped off the curb into the street is much less predictable.



**Figure 3: Curb prediction: current detected curb location and future location found by image correlation.**

*SLAM / DTCMO*: We have been using a Sick single-axis ladar for map building, moving object detection, and ego-motion detection. Collectively, these processes are referred to as SLAM (Simultaneous Localization And Mapping) and DTCMO (Detection, Tracking, and Classification of Moving Objects). Since the Sick scans a single line across the scene, a horizontally-mounted Sick sensor maps a single plane of the scene. As long as the vehicle does not pitch or roll violently, and as long as the objects in the scene have adequate vertical extent, we can track those objects from

frame to frame as the vehicle moves. If the data from successive frames matches, the scene is fixed and the shift between frames gives the vehicle's motion (SLAM). If the vehicle is stationary, frame to frame subtraction removes the fixed objects, and the moving objects are easily segmented (DTCMO). Each of these problems has been addressed separately in the literature[10] [11] [12]. Our contribution is to address both problems simultaneously. The combination is both novel and critical: in an urban environment, it is important to first segment out moving objects before building maps and localizing the vehicle; this gives much better maps and ego-location, and the improved maps give a better basis for data subtraction to detect moving objects. We described the details of our approach in[13] [14]. In Figure 4, the solid rectangle is the vehicle itself; segmented moving objects are shown outlined with rectangles, with faint trails showing their motion history
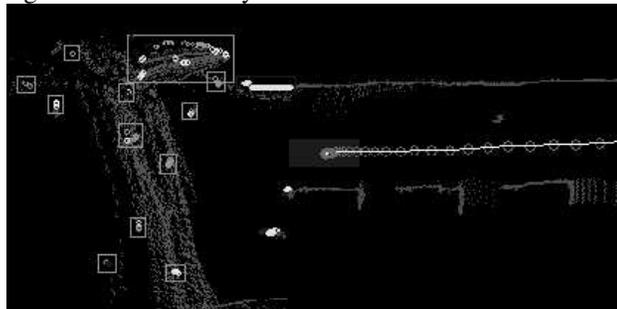


**Figure 4: Sick scanner range data processing, showing detected moving and fixed objects and vehicle trajectory**

*Large-Scale Mapping:* In order for SLAM to work over city-scale environments, an autonomous driving system will need to be able to efficiently store and access detailed maps of regions far larger than memory limitations permit in current systems. Mapping a city the size of Pittsburgh at 1 cm resolution using 8 bytes/cell would consume roughly 10 TB of storage. Using sparse storage techniques, the raw storage requirements can be brought down by a factor of 10 or more. However, even under the most optimistic of assumptions, the amount of storage required will exceed reasonable system memory. We have built a system based on quadtrees and intelligent prefetching. Each cell in the quadtree stores the likelihood that it is occupied, and the number of observations made in that cell. For each ray cast by the laser scanner, the program traverses the quadtree, updating the probabilities of each cell. If the locations of the hits within a particular cell are know accurately, the cell may be subdivided. In Figure 5, the brightness of each cell shows the probability of occupancy. Cells that are relatively far from the vehicle are usually not well mapped, and therefore are larger; cells closer to the vehicle are often subdivided to reflect the more accurate information from shorter-range observations.
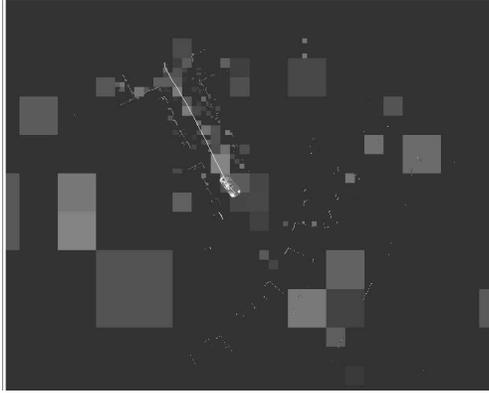
**Figure 5: Quadtree for city-scale map representation**

*Combined data processing*: We have conducted a number of experiments using registered data from video and range sensors. Figure 6 shows a snapshot from the real-time data acquisition and display. The top of the frame shows three-camera video data, aligned by software to provide a 180 degree panoramic view. Overlaid on the image is range data from the forward-looking Sick range scans. The large panel in the lower left is overhead map data; individual objects (including parked cars) are clearly visible. The bottom right shows other data being collected, including vehicle pose and motion. The data can be collected, replayed, and randomly accessed to fetch corresponding data from all sensors for any given time.



**Figure 6: Calibrated data from three forward-looking cameras, forward-looking range sensor, and vehicle state estimator**

## 4   Collision Prediction

Given the location of moving and fixed objects, as collected in the map, and the heading of the vehicle, it is possible to calculate the time of collision for each object. However, since objects may change their motion, and the vehicle may not continue along its current arc, those calculations cannot be exact; instead, they must take into account the probability distribution of likely future states of the testbed vehicle and other objects. This calculation is further complicated by different motion patterns for different kinds of objects: pedestrians on a sidewalk are less likely to step off the curb in front of a moving vehicle, and more likely to continue along the sidewalk. We have developed a probabilistic model of object motion that takes this kind of factor into account and generates a probability of collision as a function of time for each object in the environment. Rather than picking a fixed trajectory, we sample from the trajectories that the object could follow: faster or slower, turning more or less, and (for pedestrians) more or less likely to step off a curb and into the roadway.
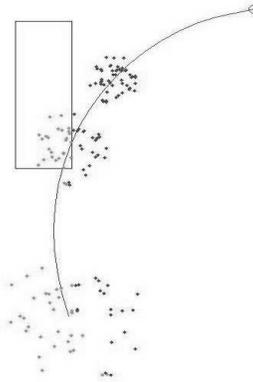


**Figure 7: Probabilistic collision detection.**

Each of these sampled trajectories is then examined to see if it causes a collision with the vehicle, and at what time. Figure 7 shows the process. In vehicle coordinates, the figure shows the position of a pedestrian relative to the bus for randomly sampled trajectories over the next 2, 3, and 5 seconds. Positions that would have resulted in a collision are shown in red. Many more trajectories are actually sampled than can be shown in this figure.

The fraction of trajectories that causes collisions is used as the probability of collision, as a cumulative function of time. High-probability imminent collisions are triggers for either an urgent driver warning (for driver assistance systems) or an immediate cue for the vehicle to take evasive action (for automated vehicles).

# 5    Integrated Warning System



*Replay of the integrated surround sensing system*
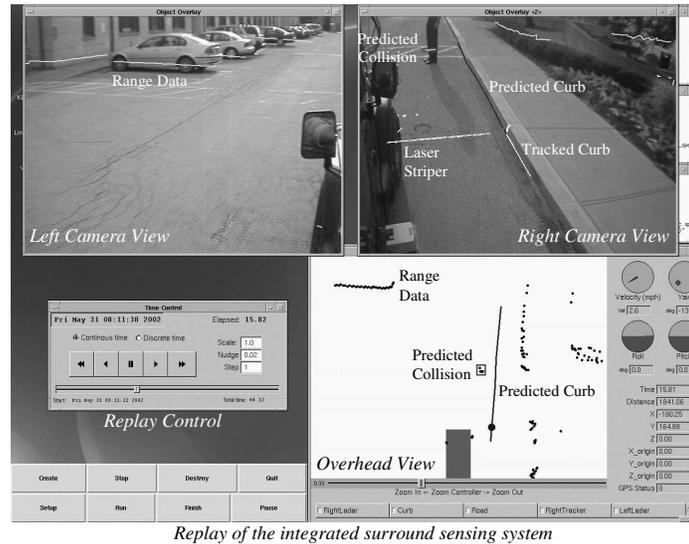
**Figure 8: Integrated Warning system: left camera view with overlaid range data; right camera view with overlayed range data and curb location; and fused map.**

All of these modules have been integrated into a warning system. The system does the following:

- Processes data in real time from two cameras looking along the sides of the vehicle, two laser scanners looking along the sides of the vehicle, the motion estimation sensors, and the side-looking curb tracker.
- Optionally, logs all the data with time tags (sub-millisecond accuracy) for later replay and analysis
- Calibrates all the sensors relative to each other and to the vehicle
- Finds objects in the environment
- Tracks the curb over time
- Overlays the detected objects and the curb on the video data
- Uses video processing to project the curb and road ahead of the vehicle, taking into account occlusions (as noted by 3-D objects projected onto the video image) to terminate tracking where the road edge is not visible
- Determines whether each object is in the roadway or safely on the side
- Estimates the potential of collision for each object

Figure 8 shows the complete system in operation. The dials on the bottom right show vehicle state. The left camera view, on the top left, shows the range data from the left Sick rangefinder and object detection system overlaid on the left camera. The top right shows the right camera view, with the right-hand Sick system data; plus the laser striper's current data, the detected curb, the track of the curb, and its future predicted position. The bottom right shows the combined map information.

The systems as developed to date are good for demonstrations and proof of concept. The next steps are to mature the systems and deploy them on test vehicles for long-duration tests in real weather and real traffic conditions. Over the summer of 2003, the Navlab group is hardening the electronics and improve the algorithms for full automation on robot vehicles, and will install the warning system on transit busses for extended trials in operational service.

## References

1 Driving in Traffic: Short-Range Sensing for Urban Collision Avoidance, C. Thorpe, D. Duggins, J. Gowdy, R. MacLachlan, C. Mertz, M. Siegel, A. Suppe, C. Wang, and T. Yata, in Proc. of SPIE: Unmanned Ground Vehicle Technology IV, Vol. 4715, April, 2002.

2 Vision-based Pedestrian Detection: will Ants Help?
M. Bertozzi, A. Broggi, A. Fascioli, and P. Lombardi, in Proc IEEE Intelligent Vehicle Symposium, June, 2002, Versailles France

3 Shape-based Pedestrian Detection and Tracking
D.M. Gavrila and J. Giebel Proc IEEE Intelligent Vehicle Symposium, June, 2002, Versailles France

4 Probabilistic Template Based Pedestrian Detection in Infrared Videos, Harsh Nanda, Larry Davis Proc IEEE Intelligent Vehicle Symposium, June, 2002, Versailles France

5 Pedestrian Detection and Tracking with Night Vision
Fengliang Xu, Kikuo Fujimura Proc IEEE Intelligent Vehicle Symposium, June, 2002, Versailles France

6 Pedestrian Recognition in Urban Traffic Using a Vehicle Based Multilayer Laserscanner , Kay Ch. Fuerstenberg; Volker Wilhoeft; Klaus C. J Dietmeyer Proc IEEE Intelligent Vehicle Symposium, June, 2002, Versailles France

7 3D-Model-Based-Vision for Innercity Driving Scenes
K. Fleischer, H.-H. Nagel, T. M. Rath, Proc IEEE Intelligent Vehicle Symposium, June, 2002, Versailles France

8 Frontal and Lateral Control for Unmanned Vehicles in Urban Tracks , Ricardo Garcia, Teresa de Pedro, Jose E. Naranjo, Jesus Reviejo, Carlos Gonzalez Proc IEEE Intelligent Vehicle Symposium, June, 2002, Versailles France

9 Eye-safe Laser Line Striper for Outside Use, C. Mertz, J. Kozar, J.R. Miller, and C. Thorpe, in IEEE Intelligent Vehicle Symposium, June, 2002..

[10]A Real-time Algorithm for Mobile Robot Mapping with Applications to Multi-robot and 3D Mapping, S. Thrun, W. Burgard, and D. Fox *, IEEE Int. Conf. on Robotics and Automation*, 2000.

[11] Summer School on SLAM 2002, http://www.cas.kth.se/SLAM/

[12] A solution to the simultaneous localization and map building (SLAM) problem, G. Dissanayake, P. Newman, S. Clark, H. Durrant-Whyte, M. Csorba, IEEE Transactions on Robotics and Automation, Vol 17, No 3, p229-241, June 2001.

13 Simultaneous Localization and Mapping with Detection and Tracking of Moving Objects, C. Wang and C. Thorpe, in Proceedings IEEE International Conference on Robotics and Automation, May, 2002

14 Online Simultaneous Localization And Mapping with Detection And Tracking of Moving Objects: Theory and Results from a Ground Vehicle in Crowded Urban Areas, Chieh-Chih Wang, Charles Thorpe and Sebastian Thrun  In Proceedings of IEEE International Conference on Robotics And Automation, Taipei, 2003