Safe robot driving

Chuck THORPE

Romuald AUFRERE

Justin David CARLSON

David DUGGINS

Terrence W. FONG

See next page for additional authors

Citation

THORPE, Chuck; AUFRERE, Romuald; CARLSON, Justin David; DUGGINS, David; FONG, Terrence W.; GOWDY, Jay; KOZAR, John; MACLACHLAN, Robert; MCCABE, Colin; MERTZ, Christoph; Arne SUPPE; WANG, Chieh Chih; and YATA, Teruko. Safe robot driving. (2002). *Proceedings of 4th International Conference on Machine Automation (ICMA '02), Tampere, Finland, 2002 September 11-13*. 1-8. Available at: https://ink.library.smu.edu.sg/sis_research/8283

Author

Chuck THORPE, Romuald AUFRERE, Justin David CARLSON, David DUGGINS, Terrence W. FONG, Jay
GOWDY, John KOZAR, Robert MACLACHLAN, Colin MCCABE, Christoph MERTZ, Arne SUPPE, Chieh Chih
WANG, and Teruko YATA

Safe Robot Driving

Chuck Thorpe, Romuald Aufrere, Justin Carlson, Dave Duggins, Terry Fong, Jay Gowdy,
John Kozar, Rob MacLaughlin, Colin McCabe, Christoph Mertz, Arne Suppe, Bob Wang,
Teruko Yata
<firstname.lastname>@ri.cmu.edu
The Robotics Institute, Carnegie Mellon University, Pittsburgh PA USA

1 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.

2 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.

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.

3 Testbed Vehicle

Figure 1: Navlab 11 testbed

Our testbed vehicle for this work is the Navlab 11. It is equipped with:
- 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 looking for the curb on the right;
- 5 500-MHz Pentium computers;

- high-accuracy time synchronization; and
- various other sensors and actuators.

4 Perception Modules

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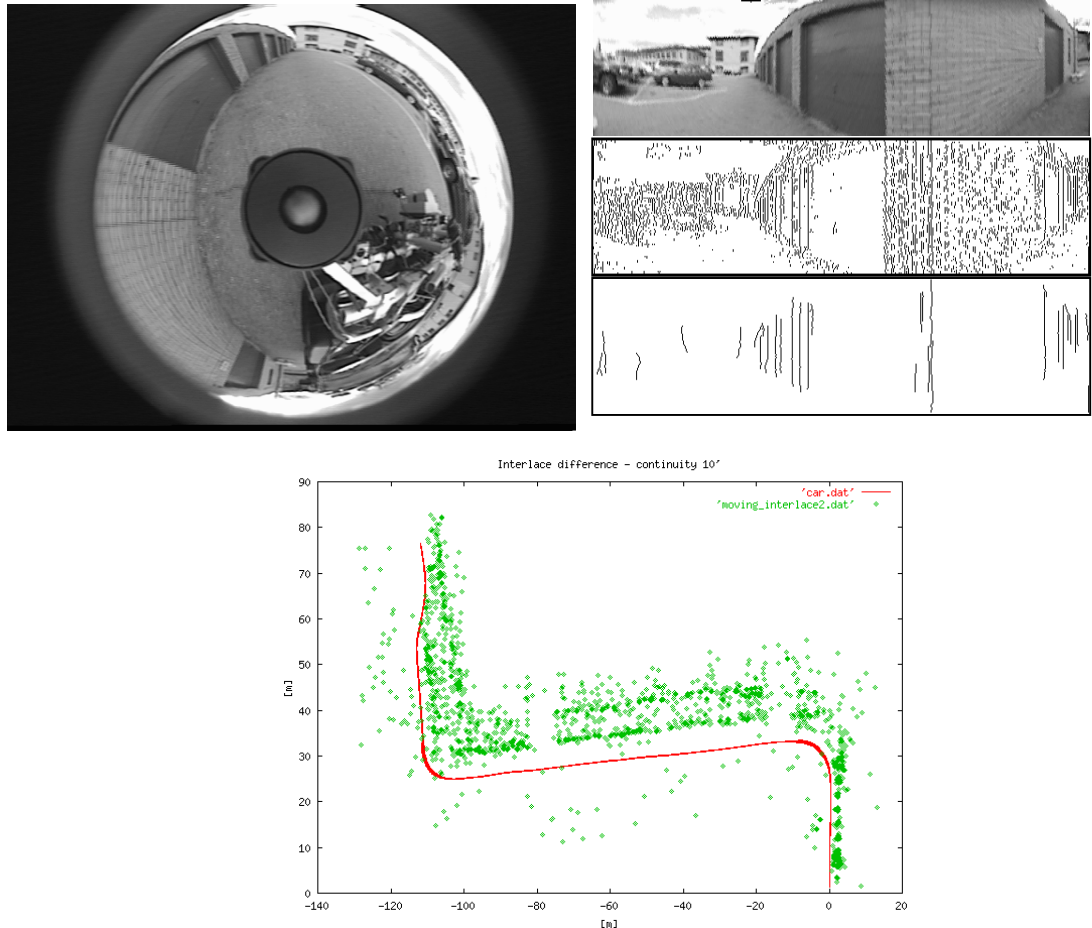

Figure 2: Optical flow tracking: omnicam image, unwarped image with edges, map of tracking result

Optical Flow: Our work with optical flow uses an Omnicam and optical flow tracking to detect nearby objects from a moving vehicle. The omnicam is a video sensor with a hemispherical mirror, which enables one image to capture a 360 degree view of the world. In software, we unwrap the image, detect vertical edges, then track those edges from frame to frame as the vehicle moves. We use a Kalman filter to track each of the detected objects, to update the position estimate of the objects with time. The result is a map of all the detected objects, suitable for obstacle avoidance and map-building. This subsystem has not yet been integrated into the full data fusion path.
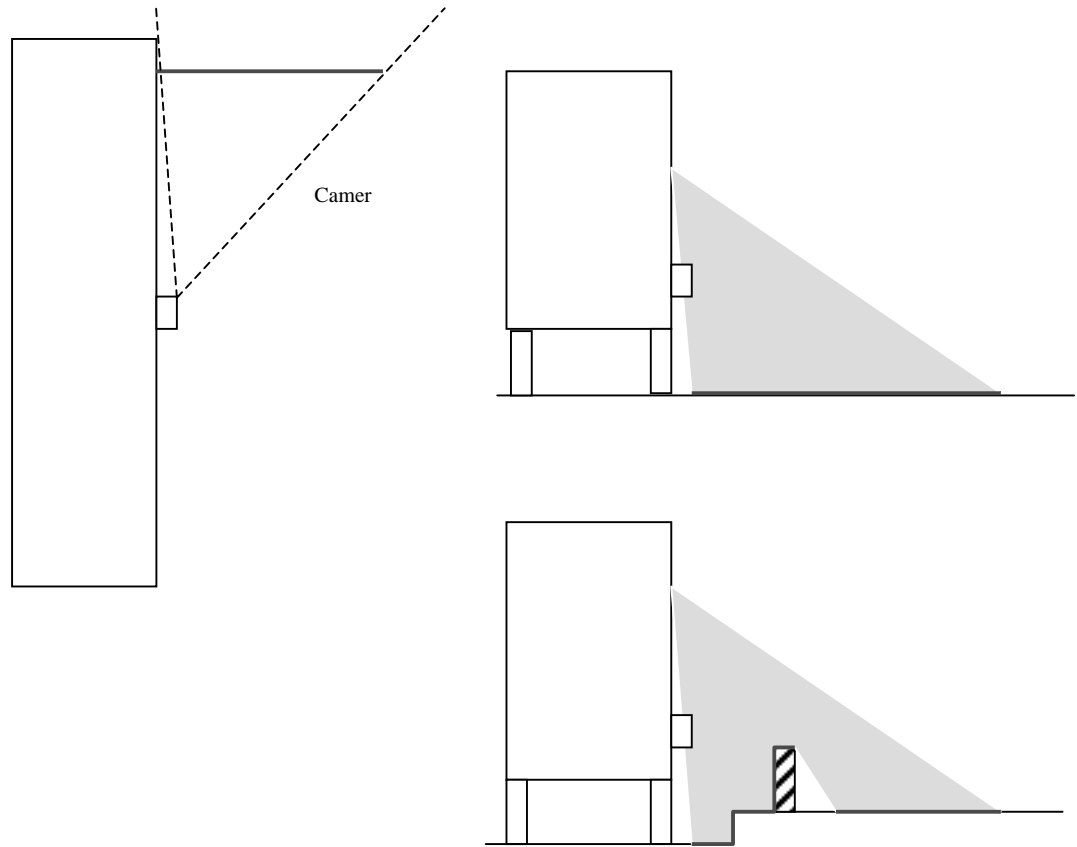
Figure 3: light stripe sensor use on a bus for curb and obstacle detection

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. 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.

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.

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.
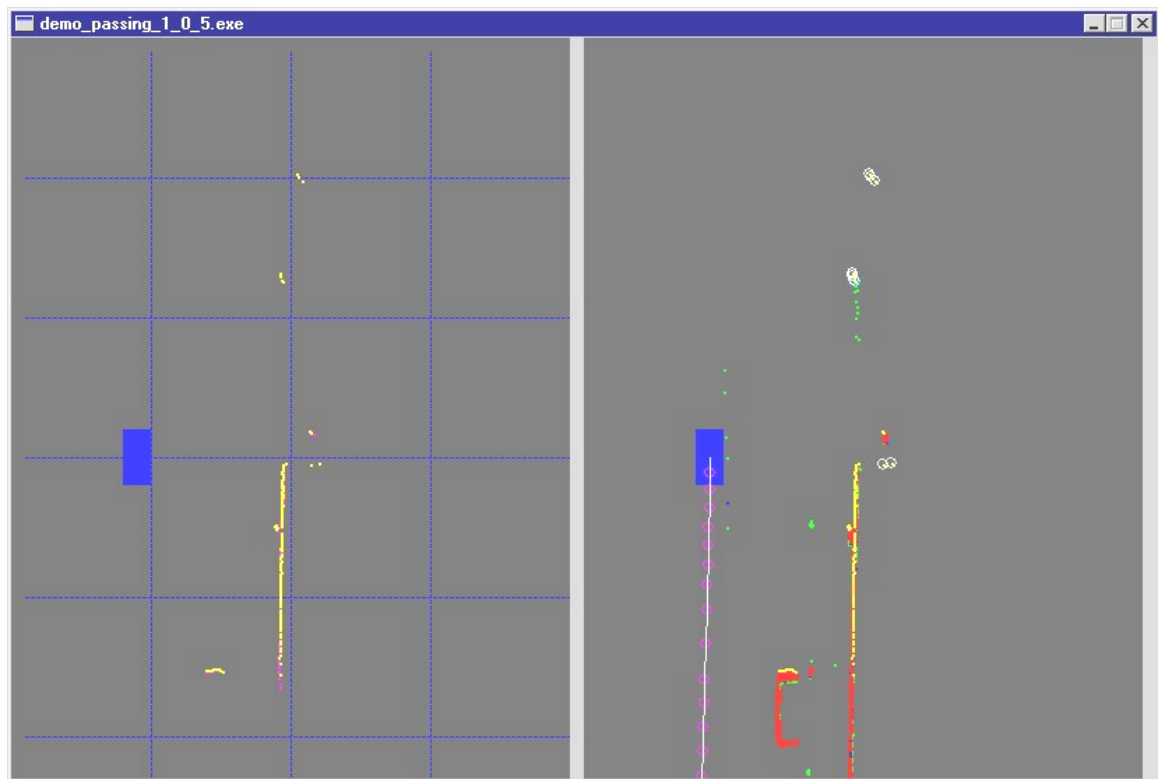

Figure 4: Sick scanner range data processing, showing a single scan (left) and multi-scan matching (right)

Sick processing: We have been using a Sick single-axis ladar for map building , moving object detection, and ego-motion detection. 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. In Figure 4, the left-hand side shows the data from a single rangefinder mounted on the right of the Navlab11. The solid rectangle represents the vehicle itself; the points show range returns. The right-hand side shows the result of data matching. The long vertical line is the wall of a building that has been seen in multiple frames. The partial rectangle near the

bottom is a parked car, that has been observed from the rear, side, and front, so a more complete representation has been built. The circles in the top of the frame show small objects that have been observed multiple times, but with changing locations; these are tentatively classified as pedestrians. The faint chain of circles trailing from the vehicle is the list of positions from which the vehicle made its observations.
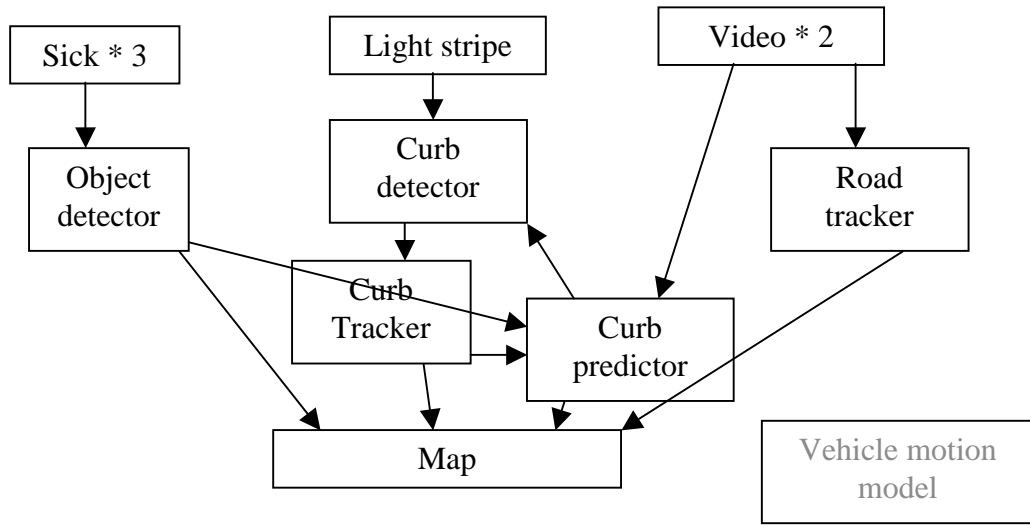

Figure 5: Data Fusion paths

Fusion: The individual perception and interpretation modules all generate output that is fused into a consistent map representation. The data flow is shown in Figure 5. The "vehicle motion model", which keeps track of vehicle pose, is distributed to all modules. The other modules each process a variety of data, culminating in labeled objects flowing into the map.

5   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. Each of these sampled trajectories is then examined to see if it causes a

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

6   Demonstration System
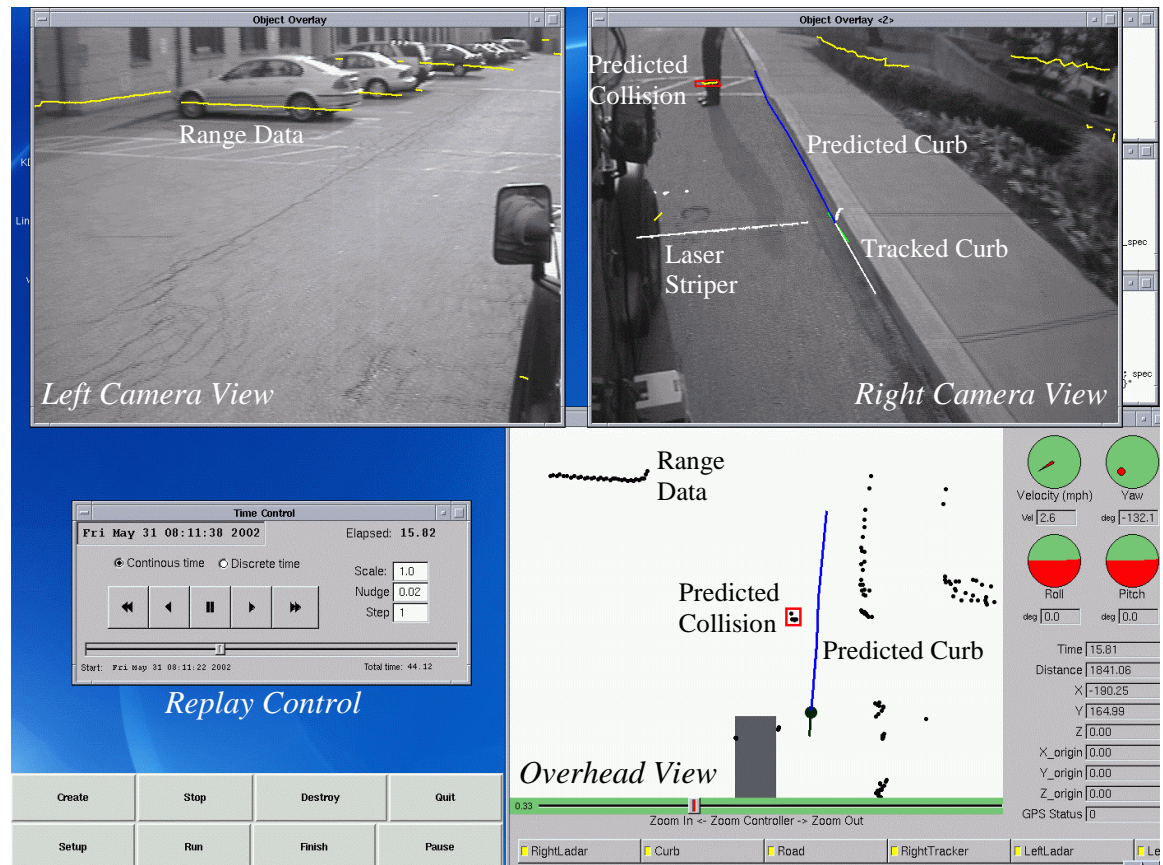
Replay of the integrated surround sensing system

Figure 6: integrated dmonstration 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 demonstration 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 6 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.

7   Status and Ongoing Development

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 next year, the Navlab group will harden the electronics and improve the algorithms for full automation on robot vehicles, and for extended tests on transit busses.

8   Acknowledgements

9   Bibliography

"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.

"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.

"Static Environment Recognition Using Omni-camera from a Moving Vehicle", T. Yata, C. Thorpe, and F. Dellaert, tech. report CMU-RI-TR-02-12, Robotics Institute, Carnegie Mellon University, 2002

"Eye-safe Laser Line Striper for Outside Use", C. Mertz, J. Kozar, J.R. Miller, and C. Thorpe, in IV 2002, IEEE Intelligent Vehicle Symposium, June, 2002., December, 2001.

"Environment Recognition from a Car Using an Omni-camera", T. Yata and C. Thorpe, in Proc. of the 19th Annual Conference of the Robotics Society of Japan, September, 2001