Driver Assistance: Contemporary Road Safety

Andrew Dankers, Luke Fletcher, Lars Petersson, Alexander Zelinsky
Department of Systems Engineering
Research School of Information Sciences & Engineering
The Australian National University
{andrew, luke, larsp, alex}@syseng.anu.edu.au

Abstract

Despite a trend towards increased road safety, fatal and/or serious accidents induce a significant toll each year. We work towards improved road safety through the development of an intelligent driver assistance system that monitors the driver's performance and assists with vehicle guidance. This complex task is decomposed into subsystems including lane-keeping, obstacle avoidance, sign detection, monitoring of driver alertness, situation feedback to the driver, and proactive and preemptive warnings. The set of subsystems combine to produce the desired assistive control and monitoring behaviours. Sensing, feedback and actuation mechanisms are presented. Although refinement and testing is currently in progress, some preliminary results and demonstrations are provided with more thorough experimentation to be conducted shortly.

1 Introduction

Road crashes are a huge cause of human trauma. Since record keeping commenced in 1925, there have been over 169,000 road fatalities in Australia, according to the ATSB [Australian Transport Safety Bureau, 2003] . In addition to the burden of personal suffering, the monetary cost of crashes is in the order of $15 billion per annum (1996 data). From 1970 until 2002 the fatality rate dropped from 30.4 to 8.8 deaths per 100,000 population. This rate reduction has been achieved despite a huge increase in motor vehicle use, and an increasingly aged driving population[1]. From 1970 to 2002, the fatality rate per 10,000 registered vehicles has dropped from 8.0 to 1.4. In terms of 100 million vehicle-kilometres travelled the fatality rate has dropped from 4.4 in 1970 to 1.0 in 2000.

Figure 1 presents the dominant causes of fatal road crashes in Australia between 1988 and 1998 according to the ATSB Fatality Crash Database. These figures suggest that a significant proportion of fatalities can be indirectly attributed to impeded driver concentration and performance.

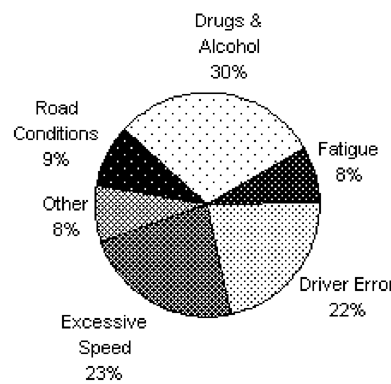

Figure 1: Causes of road fatalities, Australia 1988 - 1998.

Additions of airbags, anti-lock breaks and traction control are just a few examples of factory-fitted technological improvements common on today's vehicles. The significant amount of research and funding resources allocated to developing safer vehicles is evidence of an industry drift towards increased consumer safety. For example, Ford Motor Company's latest concept car is reported to "combine advanced accident avoidance systems and intelligent vehicle technology". Their modified Ford Explorer, dubbed S2RV[2], incorporates "smart" technology such as voice control, reconfigurable displays and switches, a navigation system that uses vehicle-to-vehicle communication and provides incident report information and routing, as well as "safe" features such as

[1] One in every three licensed drivers is over the age of 55 [Wisconsin Govt., 2003]. Peripheral vision, depth perception, night vision and reaction times are shown to deteriorate with age, potentially reducing driving performance.

[2] Smart, Safe Research Vehicle

proximity warning systems, accident avoidance camera systems, and rear collision warning systems.

The contemporary concept of safety through prevention and preparedness is emerging as an automotive industry philosophy. Minimising exposure to potentially risky situations can be achieved through increased driver awareness, and an understanding of the road-scene. Counteracting driver distractions, preventing speeding and lane drift, and reducing reaction times are but a few ways in which such risk can be reduced. Certainly, any efforts consistent with this philosophy of increased attention to safety directly benefits consumers and the automotive industry.

1.1 Direction

The overall goal of our research is to investigate *Driver Assistance Systems*, that is, systems that can relieve the driver during distraction or warn about upcoming situations and possibly take control of the car if an accident is imminent. This approach is different from creating fully autonomous vehicles where the driver becomes a passenger relying purely upon computer control. The driver assistance approach is more likely to be widely accepted by the public since the human driver maintains vehicle control.

To complete a task such as driver assistance, it is beneficial to reduce the broad objectives into simpler, well defined competencies that can be assembled to give the desired outcome. Decomposing a system into well-defined and easily assessable subsystems reduces complexity and helps to define a path to achieving a complete system that behaves in a desired manner. Driving a car for example, can be reduced to a set of sub-tasks such as lane-keeping, obstacle avoidance, velocity control, path planning, etc. This integral approach is adopted in the development of the complete system for monitoring and assisting a driver.

The research performed within this project is unique in that it does not only consider the environment *outside* of the car, but also what is happening *inside*.

There are a number of competencies that are being considered in the project including the following:

- Traffic situation monitoring
- Vehicle state monitoring
- Driver state monitoring
- Communication with the driver
- Vehicle control
- Reasoning system

1.2 Outline

Section 2 introduces the hardware structure of the research platform including sensors, actuators and feedback subsystems. An elaboration of advanced subsystems and concepts is provided in section 3. Subsystem processing and interactions for driver assistance are developed in section 4. Results are discussed in section 5, prior to suggestions for future work (section 6) and conclusion (section 7).

Figure 2: The research platform vehicle.

2 Research Platform

The platform that is used in the project is a 1999 Toyota Landcruiser 4WD. It is equipped with the appropriate sensors, actuators and other hardware to provide an environment in which the desired competencies can be implemented.

An internal Ethernet network connects a number of computers to provide adequate computational resources. Currently, the normal configuration consists of three Pentium computers. Two are dedicated to performing computer vision tasks, the third controls the active vision head and supports other less computationally intensive tasks.

All vehicle sensors and actuators, other than cameras, are monitored and controlled respectively using a STG[3] card.

2.1 Sensing Systems

All sensors gather data that is made available to any processing threads using client/server architectures. Consequently, any threads operating in the vehicle can read data from any device without affecting other processes that may require data from the same sensor. An example is where lane-tracking and object detecting subsystems simultaneously use vision data from the cameras in separate processes. A brief summary of sensors follows.

Road Interaction Sensors

Basic sensors used in the road interactions include numerous encoders, a fibre-optic roll/pitch/yaw rota-

[3]ServoToGo Inc.

tion rate gyro and solid state x/y/z accelerometer, gear/accelerator/break/indicator monitors, and steering angle potentiometer and encoder.

More advanced sensor systems include a lane tracking system (see section 3), obstacle tracking system (section 3), sign detection system (section 3), GPS, SICK laser range scanner, and a Millimeter radar device

Driver Interaction Sensors
Sensors used to monitor driver interactions include steering strain gauges, a TFT touch-screen, Facelab (see section 3), voice recognition, and driver initiated system override mechanisms.

Driver
The driver is the most trusted sensor. The system always obeys driver override, effectively making the driver the primary sensor.

2.2 Feedback and Actuation Systems

As with sensing, feedback and actuation devices use client/server architectures whereby any processes/threads operating in the vehicle can access the actuators. However, in general, the subsystems contribute to an overall desired system actuator state that is effected by higher level processes.

Driver Feedback
Feedback is delivered to the driver via mechanisms including steering-wheel haptic feedback, audio feedback and via the TFT display.

Actuation
The factory-built cruise control actuator unit is used to manipulate vehicle acceleration. DC motors have been installed to actuate the break pedal and steering angle. Gaze control actuation for the vision system is via a high-speed, high accuracy Cable-Drive Active vision Robot (CeDAR) unit [Truong et al., 2000] incorporating common tilt and independent pan axes (see section 3). The driver always has the ability to control/override vehicle actuation through the standard automobile interface.

3 Advanced Subsystems

Brief clarification of the advanced subsystems that have been implemented is now provided (section 3.1). Next, several subsystem concepts currently under research and development are introduced (section 3.2).

3.1 Implementation Details

CeDAR is a high-speed, high precision stereo active vision system [Truong et al., 2000]. The novel mechanism is light-weight, yet capable of motions that exceed human performance. Figure 3 shows the standard CeDAR unit, Figure 4 shows how the unit has been inverted and mounted for use with driver assistance.

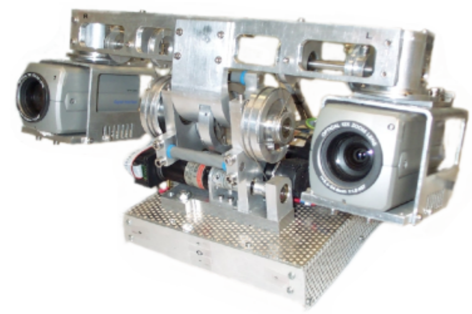

Figure 3: The CeDAR unit.

Figure 4: CeDAR mounted on the research vehicle.

Gaze Stabilisation. Because the vehicle uses an active vision system, visual data quality is improved by stabilising gaze with respect to the road. This provides vertically stabilised images with significantly reduced image shift to any vision-based subsystem. Two mechanisms are adopted to reduce such *retinal shift*. A real-time software algorithm analyses successive images and shifts them such that detected roll and elevations due to bumps on the road are reduced. Detecting slip and shifting of each image takes less than a millisecond in total. The active head gaze is continually moved to reduce the necessary software image shift to zero. The physical head gaze shift is a slightly lagging process due to pipeline effects but the speed of the software algorithm means that images are already visually stabilised. The second mechanism incorporates the use of gyro sensor data to detect bumps and rolls to control the active head such that retinal slip can be reduced by altering the gaze. This is also a slightly lagging process due to pipeline effects. The use of each of these mechanisms, or the combination of these mechanisms for increased robustness, permits that the effective horizon in the images appears stationary in real-time.

Lane Tracking involves developing an understanding of the position of the vehicle with respect to the boundaries of the lane along which it travels. We use a multi-cue, dual phase particle filter system [Loy et al., 2002]. The first particle filter searches for the the road width, the lateral offset of the vehicle from the centerline of the road and the yaw of the vehicle with respect to the centerline of the road. The second particle filter captures the horizontal and vertical road curvature in the mid to far-field ranges using the state information captured by the first particle filter. Lane tracking cues were designed to be simple and efficient while being individually suited to a different set of road scenarios. Used independently, each cue could perform poorly, but when they are combined through the cue fusion process they produce a robust solution. Relevant cues adopted here include a lane marker cue, road edge cue, road colour cue, non-road colour boundary cue, road width cue and an elastic lane cue. See [Apostoloff, Zelinsky, 2002] for further details regarding lane tracking. Figure 5 shows lane detection for separate cues that contribute to the perceived lane position. Figure 6 is an example of lane detection output.

Object Tracking involves the extraction and tracking of objects above and normal to the ground plane, including those on and adjacent to the road. Ob-

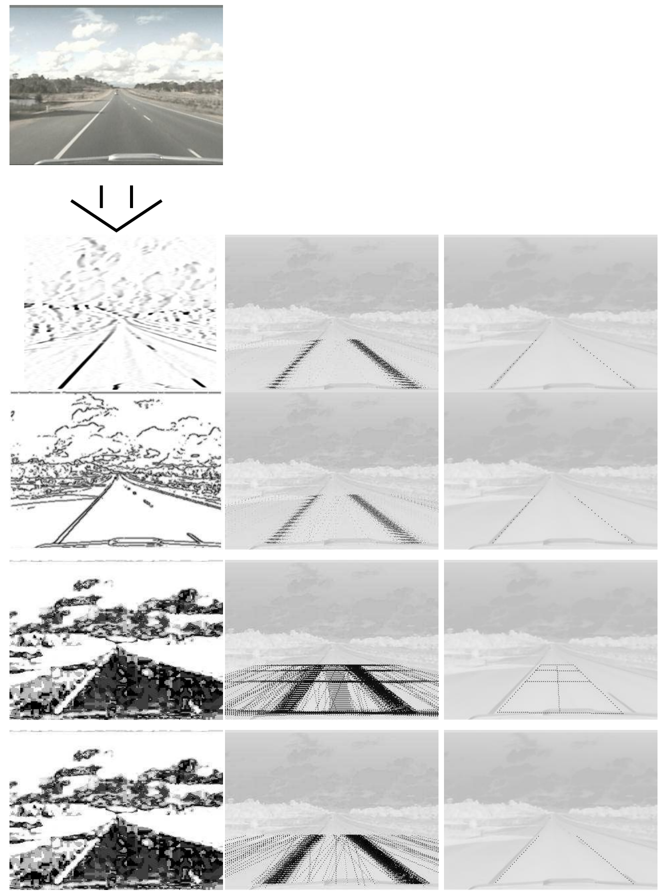

Figure 5: The top left image shows the video input to the lane tracker. The left column below shows preprocessed images which acts as input to the different particle filters. In the middle column, the current set of hypotheses are overlaid in image space. The most likely hypothesis, or particle, for each cue is then shown in the rightmost column. The order of the cues from top to bottom is: Lane Marker Cue, Lane Edge Cue, Road Colour Cue and Non Road Colour Cue.

Figure 6: Lane tracker results. The coloured lines indicate the perceived lane position.

Figure 7: Objects (cars) detected in the road-scene.

jects include vehicles, pedestrians, trees, signposts etc. V-disparity [Labayrade, Aubert, 2003] images are used to extract a segmented ground-plane hypothesis from stereo images of the road scene. Objects above the ground-plane also emerge in the v-disparity image of the scene, and their presence is re-enforced by analysis of a depth map. The most relevant objects, including those on potential collision trajectories are tracked. Extracted objects can be analysed using support vector machines (SVM), to classify them. Objects are separated into two main categories, those commonly seen on the road, and those commonly off the road, so that an inference of the road location can be obtained and matched with the lane tracker. For a broader road-scene awareness, objects on the road can potentially be sub-classified into trucks, cars, motorbikes, pedestrians, etc, and likewise those off the road can potentially be sub-classified into obstacles, signposts, etc. Refining the obstacle detection sub-system is the objective of current work.

Sign Detection is used to gain a perception of road conditions ahead of the vehicle. As the signs we are interested in are are iconic in nature as opposed to textual, the problem remains an image classification problem rather than an optical character recognition (OCR) problem. Common methods of such image classification have been based on off-line parametric analysis or learned models such as artificial neural networks. Recently, SVM technologies have been producing excellent results classifying objects from raw images or from feature vectors.

As speed-signs reliably appear radially symmetric, the circular hough transform can be used to specify the centroid and radius of potential signs. A more recent and efficient technique is the fast symmetry transform (FST) as described by [Loy et al., 2002]. This technique uses image gradients and a voting space to determine the number of pixels contributing to a circle of a specified radius. A subset of radii can be tested and the results combined. The left inset in Figure 8 shows the result of the fast symmetry transform on a typical road scene image. The radius with the maximum FST response is used to crop out the detected sign from the image. Detected signs are then cropped and resized to a standard size before classification.

Once detected, a sign is classified to verify that the detected image is in fact a speed sign and to determine the speed of the sign (using a SVM).Two SVMs are used, the first is a binary classifier for the *SPEEDSIGN / NOT-SPEEDSIGN* distinction. The second is a multi-class classifier for speed differentiation: *50*, *60*, *70*, *80*, *90* and *100*. The split level categorisation allows training of each SVM to emphasise traits specific to the two cases instead of compromising false positive detection at the expense of speed classification and vice versa.

Facelab is an advanced driver monitoring system comercialised by Seeing Machines, based on collaborative research and development work between ANU and Volvo Technological Development Corporation. Incorporating a passive stereo camera pair mounted on the dashboard of the vehicle, it analyses the driver's head pose and gaze direction, providing information on what the driver is likely to be looking at and if the driver is maintaining concentration on the road. Images are monitored in real-time to determine the 3D position of matching features on the

Figure 8: Sign detection. Inset images show (left) most likely sign positions, and (right) signs extracted and sized for SVM use.

drivers face. Features are used to determine the pose of the face as well as eye gaze direction ($\pm$ 3 deg), blink rates, and eye closure to determine the driver's level of alertness.

Figure 9: Facelab cameras mounted to observe driver.

3.2 Advanced Concepts

TFT Telematics. The tft touch-screen device displays the road situation interpretation and all sensor data for driver assessment. Telematic applications that use computers and communications to provide information and support to drivers have a significant and increasing market presence. Predominant applications include navigation, cellular phone interfaces, entertainment, traffic and emergency notification systems. Studies into the use of telemetric driver input reveal that as the interfaces visual demand increases, driving performance decreases [Tsimhoni,

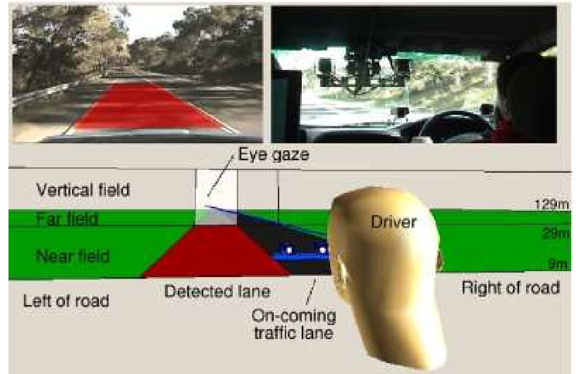

Figure 10: 3D model of how Facelab perceives the driver's head position and gaze. Additionally, top left inset shows results of lane detection, used to build Facelab world model; top right inset shows the scenario.

Green, 2001]. These products must be safe and easy to use and hence standards are being developed for positioning and designing telemetric interfaces for use while driving. With a safety-oriented system such as ours, these recommendations must be taken into account [Green, 2001], [Mayer et al., 2002].

Haptic steering feedback incorporates vibrational steering wheel feedback and steering torque feedback. The steering strain gauges provide information on the significance of the driver's steering guidance. If the force applied to the steering wheel by the driver is large, it suggests the the driver is not in agreement with system assistance, and override occurs. Conversely, the driver can feel the level of guidance through the steering wheel through force-feedback, and the manner in which the system is assisting steering.

A haptic interface used to simulate rumble strips and other road surfaces also immediately suggests to the driver that they review their course. In a haptic warning lane departure study [Suzuki, 2002], it was revealed that steering vibration was effective for lane departure situations. Haptic stimulus transmitted through the steering wheel was found to trigger an effective and beneficial mental response to lane departure and made interpreting road conditions more immediate, even if the driver was not briefed on what the signal was for, it simply stimulates the appropriate response.

Audio feedback. Audio warnings can reduce driver reaction time by incorporating directional feedback to not only alert the driver, but by additionally assisting the driver's search for its source by simulating the direction in which an warned event originates in relation to the vehicle. However, research has shown that such reactions may require some driver training and exposure to the system for directional audio to be effective [Suzuki, 2002]. The study found no significant difference between the use of monaural and directional stereo alerts (where drivers were not briefed that the stereo audio signified the direction of a warned event) until the driver became acquainted with the system.

4 Driver Assistance Processes

Subsystem competencies are combined to produce desired system bahaviours. It is intuitive to categorise high-level behaviours into two main classes: those involved in controlling the vehicle through the perceived road-scene, and those involved in interactions with the driver. Figure 11 sumarises the generalised subsystem interactions.

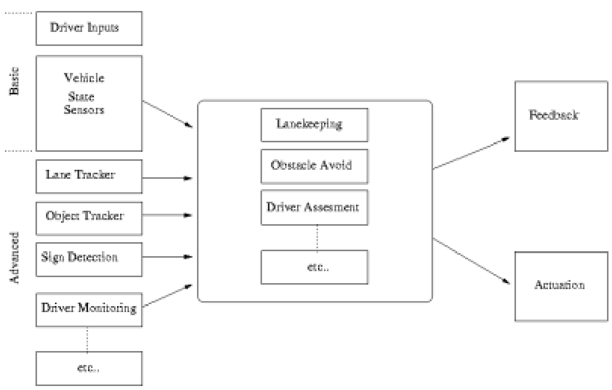

Figure 11: Subsystem interactions: The leftmost column represents system generalised sensors, including basic and advanced. The middle column shows desired behaviours and planning competencies for road or driver interactions. These in turn lead to outputs, the rightmost column.

4.1 Vehicle Control Systems

The objective of the integrated vehicle control assistance subsystems is to determine the desired gaze upon road the road-scene, and to determine a safe vehicle position and path within the road-scene. The lane-tracker sensor provides information on the limits of the lane in which the vehicle is presently situated. The obstacle tracker gives areas to avoid within the road-scene. Sign detection can be used to provide a map-like expectation of what the vehicle is approaching. Reasoning systems combine this information with the basic sensor state data to continually plan the desired position and velocity of the vehicle. The desired present state with respect to the planned path is delivered to the actuators so that driver assistance is effected. Of course, during the interactions, the system takes input from driver actuation in preference to the automated actuation state, effectively using the driver as an additional and preferential sensor/actuator.

4.2 Driver Interaction Systems

The driver interaction systems enable the driver to control the vehicle through a standard vehicle actuation interface, guide the driver assistance systems, assess the system's perceived road-scene structure, be alerted of perhaps unnoticed situations, and be monitored for driving performance and alertness.

This is effected via a summary display on the tft touch screen, showing all system sensor states and any warnings. Warnings may include approaching obstacles, signs that may have been overlooked, speeding, driver alertness, or system concerns. These visual warnings may also be combined with audio and vibrational warnings. For example, if the system detects that the driver is unintentionally departing a lane (where for example, no preemptive lane changing indicators were activated), vibrational haptic steering feedback (simulating lane boundary rumble stripes), a visual warning on the tft screen, and a directional audio warning will combine to alert the driver. In general, the warning signals are combined to most effectively alert the driver as to what sort of warning is occurring and the direction from which the warning originated from, so that reaction times are reduced.

If the driver interaction systems determine that the driver's level of alertness is unacceptable, or that any form of collision is imminent, precautionary actions are executed. For example, where the system detects driver distraction, it is handled by first warning the driver; if the driver does not react, a critical warning reveals the system's intention to go off-line (shut down driver assistance), or pull over - rather than assuming complete autonomous control of the vehicle.

5 Results

Testing has proven the usefulness of all individual vehicle subsystems including all installed sensory devices and all actuators. The advanced subsystems discussed herein have been shown to implement the desired competencies in real-time under real road conditions (see demonstration footage). Lane tracking, gaze stabilisation, obstacle detection, driver gaze detection and sign detection subsystems all produce consistently impressive results.

Presently, experimentation is being developed to help assess and improve the system in its complete form. Results of more thorough testing will be available shortly, in preparation for vehicle demonstration at the Intelligent

Vehicle Conference to be held at Eastern Creek Raceway in early 2004.

Demonstration result footage is available online. Please see section 7 for details.

6 Future Work

In addition to experimentation, planned work includes the addition and refinement of mechanisms including:

- Head-up display.
- Speech recognition input, minimising driver distraction [Tsimhoni et al., 2002].
- GPS navigation assistance.
- Millimeter radar incorporation.
- SICK laser scanner device incorporation for increased road-scene awareness and confirmation of visual data.

Additional driver assistance systems and competencies potentially include:

- Night driving assistance. A driver's ability to adjust to low light conditions and to recover from the glare is affected by oncoming headlights. Additionally, at the age of 55, human eyes require eight times as long to recover from glare as at age 16 [Wisconsin Govt., 2003]. Using a transparent LCD head-up-display overlay on the windscreen, and Facelab, the glare from oncoming vehicle lights can be eliminated by shadowing the relevant parts of the windscreen.
- Blind spot sensing. Presently, our efforts concentrate on what happens in front of the vehicle. A system that also senses regions that are visually occluded for the driver may provide information that would otherwise have been unnoticed.
- Destination detection. Reducing distractions such as map-reading and destination searching as they are approached (an address, building, etc) will not only allow drivers to maintain concentration on the road, but will also give the driver routing assistance in unfamiliar areas.
- Injury to pedestrians is a significant road-safety issue. Driver assistance systems that react to the detection of potentially endangered pedestrians could prevent accidents involving pedestrians.
- Active visual tracking of most significant gaze-points with natural return to stabilised gaze upon the road-scene horizon. Gaze may be trained on objects of interest such as those that may collide with the vehicle, signs, etc.
- Object classification to gain an increased understanding of the road-scene and enable additional behaviours such as following other vehicles, and understanding the behaviours of objects common in the road-scene.

Finally, in order that benefits be well understood, an impact study is planned to assess the effectiveness of the driver assistance system.

7 Conclusion

The significant rate-reduction in serious road accidents over recent years can be attributed to an industry trend towards advancing vehicle safety systems. In accordance with this philosophy, we have developed systems that monitor both the road and the driver, providing feedback and assistance for increased safety. Basic sensors as well as advanced sensing systems such as lane-tracking, obstacle detection, and sign detection allow the vehicle to safe-guard the driver's performance. Driver monitoring systems are used to maintain the drivers level of alertness. Effective feedback mechanisms provide the driver with sensory information with minimal distraction, and warnings that encourage constructive reactions. Intuitive driver interfaces allow easy driver input to the system.

Results from real-world testing has proven the effectiveness of subsystems. Experimentation to be conducted shortly will evaluate the impact of the driver assistance system as a whole on driver safety and performance.

Demonstrations

The system will be demonstrated live in October at the Intelligent Vehicles Conference, Eastern Creek Raceway. We will work towards demonstrating the vehicle during the ACRA conference.

Presently, additional demonstration footage is available online. Demonstration footage includes:

- Lane tracking
- Lane keeping
- Gaze stabilisation
- Object detection
- Object tracking
- Driver alertness monitoring
- Driver gaze detection
- Sign detection
- Introduction to the platform

Footage can be viewed at:

`http://www.syseng.anu.edu.au/rsl`

References

[Australian Transport Safety Bureau, 2003] Australian Transport Safety Bureau (ATSB) *Road Safety Statistics.* http://www.atsb.gov.au.

[Truong et al., 2000] Truong H. Abdallah S. Rougeaux S. Zelinsky A. A novel mechanism for stereo active vision. In *Proceedings of the Australian Conference on Robotics and Automation (2000)*

[Fletcher et al., 2001] Fletcher L. Apostoloff N. Chen J. Zelinsky A. Computer vision for vehicle monitoring and control. In *Australian Conference on Robotics and Automation (2001)*

[Loy et al., 2002] Loy G. Fletcher L. Apostoloff N. Zelinsky A. An adaptive fusion architecture for target tracking. In *Fifth International Conference on Automatic Face and Gesture Recognition (2002)*

[Apostoloff, Zelinsky, 2002] Apostoloff N. Zelinsky A. Vision in and out of vehicles: Integrated driver and road scene monitoring. In *Proceedings of International Symposium on Experimental Robotics (ISER)*

[Labayrade, Aubert, 2003] Labayrade R. Aubert D. In-vehicle obstacles detection and characterization by stereovision. International Workshop on In-Vehicle Cognitive Computer Vision 2003

[Tsimhoni, Green, 2001] Tsimhoni O. Green P. Visual demand of driving and the execution of display-intensive in-vehicle tasks. In *Proceedings of the human factors and ergonomics society 45th annual meeting (2001)*

[Green, 2001] Green P. Synopsis of Driver Interface Standards and Guidelines for telematics as of mid-2001. Technical report, UMTRI (2001)

[Mayer et al., 2002] Green P. HUD feedback to minimize the risk of cellular phone use and number entry while driving. Technical report, UMTRI (2002)

[Tsimhoni et al., 2002] Tsimhoni O. Smith D. Green P. Destination entry while driving: speech recognition versus a touch-screen keyboard. Technical report, UMTRI (2002)

[Suzuki, 2002] Suzuli K. Analysis of driver's steering behavior during auditory or haptic warnings in lane departure situations. Technical report, Japan Automobile Research Institute

[Wisconsin Govt., 2003] http://www.dot.wisconsin.gov/.. drivers/drivers/aging/mature.htm