IMPROVING SITE AND MATERIALS HANDLING OPERATIONS THROUGH AUTONOMOUS VEHICLE-RELATED TECHNOLOGIES

Ashley Tews

Commonwealth Scientific and Industrial Research Organisation (CSIRO)
Queensland Centre for Advanced Technology
Pullenvale, Queensland, Australia
Ashley.Tews@csiro.au

ABSTRACT

In 2005, we automated a forklift-based Hot Metal Carrier to be capable of typical metal transfer operations around a smelter. The project was highly successful and the vehicle has demonstrated hundreds of hours of live autonomous operations to thousands of people from industry and the public. As a result of the exposure and success of the project, we have expanded its focus to demonstrate how various technology components can be utilised to gain benefits in other areas of industrial operations. These include asset tracking, vehicle usage analysis, pedestrian detection and infrastructure profiling. Each of these components was derived from the core autonomous vehicle technology suite and has shown a high potential for improving safety and efficiency of vehicle-related operations, and improved diagnostic processes for measuring deformations of bakes and furnaces. The autonomous vehicle project and its extended technologies are described in this paper.

KEYWORDS

Site automation, vehicle operations, efficiency, site safety

INTRODUCTION

Vehicles are in constant operation around industrial worksites. In many applications in built environments, they may perform repetitive homogeneous tasks such as moving loads between various locations. For example, in the aluminium industry, the primary task of Hot Metal Carriers (HMCs) is to transport molten aluminium from the smelter to the casting shed. Typically, there are a number of static pickup and dropoff zones that the vehicles travel between. These vehicles can weigh 20 tonnes unloaded and may either be converted forklifts (Figure 1) or articulated trucks. The molten aluminium is carried in large metal crucibles. The crucibles weigh over 2 tonnes and they can hold 8 or more tonnes of molten aluminium at 700 degrees Celsius. Therefore, HMC operations are considered heavy, hot, and hazardous, with safety of operation a significant issue.

Figure 1. A Hot Metal Carrier picking up a payload.

Our primary research is focused towards automating the operations of material transport vehicles that operate indoors and out, such as HMCs. There are many challenges in developing these types of vehicles to be autonomous. Indoors, there can be a vast amount of infrastructure, other mobile machines and people. In smelters, there are also large magnetic fields and high temperatures near the pots and furnaces. Outside, the vehicle's paths may be surrounded by infrastructure, fences, and their operation may be effected by the environmental conditions such as rain, fog, snow, and heat. For an autonomous vehicle to demonstrate highly dependable operations, its design needs to take these issues into account. Needless to say, safety, reliability, predictability and repeatability are important factors to facilitate industry and user acceptance.

In 2005 at our worksite, we converted a manned forklift-style HMC into an automated version that can undertake the typical transport operations and movements of a production vehicle. Over the course of the project, the autonomous HMC has demonstrated hundreds of hours or dependable operation in short and long duration testing. The tests have included path and performance repeatability over a short-course pickup and delivery operation, several 5 hour tests of continuous operation and a test of 8 hours of scheduled operations. The HMC is able to operate around other vehicles and people in a safe and considerate manner.

As part of automating the HMC, we developed many hardware and software modules to allow control and monitoring of the vehicle as well as outfitting sensors to allow it to monitor the surrounding area. We have also developed technology for external sensing systems such as using site cameras to expand the sensing range of the vehicle, which allows the HMC to 'see around corners' or blind spots where pedestrians may emerge.

The hardware and software components can be considered as a complete package required to automate an industrial vehicle, but also individually. Examples of component technologies include vehicle and asset tracking systems, infrastructure monitoring and pedestrian detection systems.

The remainder of this paper overviews the main modules in our autonomous vehicle and how some of them can be used as stand-alone units for additional applications. We will finish the paper with relevant conclusions.

VEHICLE ARCHITECTURE

Our HMC has been automated to the level it can carry out all the operations of a conventionally operated vehicle with a driver on-board. However, whereas the driver is responsible for the efficiency, safety and sensing for the operations, the autonomous vehicle has hardware and software systems undertaking these roles.

Figure 2 provides a high-level view of the software and hardware architecture of the autonomous HMC's systems. Low-level components such as throttle, brakes, steering, hook and mast controls are controlled through Programmable Logic Controllers (PLCs). The critical safety components, such as the E-Stop buttons and the watchdog monitor, are controlled through higher grade failsafe PLCs. These PLCs provide redundancy checks of relay connections and continuously monitor the input and output state of hardware connections. All software operates under the in-house developed middleware suite (Corke et al, 2004) that allows communication with the sensors and vehicle control systems, as well as high level control modules.

The H/W Abstraction program converts the internal vehicle state sensors to human-readable signals and manages the vehicle demands in an opposite manner. High Level programs work directly with the external sensors and vehicle state to control the vehicle. Vehicle Level programs control and monitor the vehicle hardware systems. The sensing and safety systems are overviewed next followed by a description of the main High Level modules.

Sensing Systems

The vehicle's sensing systems consist of internal and external sensors. The internal sensors are typical for forklifts and provide information about the pneumatic air pressure, engine parameters, and fuel and oil levels. We have added additional sensors for determining the mast state, brake and throttle positions, odometry, gear position, and steering angle. The external sensors allow the vehicle to sense its surroundings to allow for navigation, payload manipulation and obstacle detection. The primary sensors for the autonomous HMC are SICK scanning laser rangefinders mounted on each corner (Figure 3) and are tilted down and combined to provide 360 degrees of coverage to a distance of approximately 30m. These are used for localization and obstacle detection. Secondary sensors consist of two Pan-Tilt-Zoom (PTZ) webcams attached to the mast for locating the crucible via markers on its handle (Pradalier et al, 2008), and three cameras around the front of the vehicle (Figure 3) for pedestrian detection and localization.

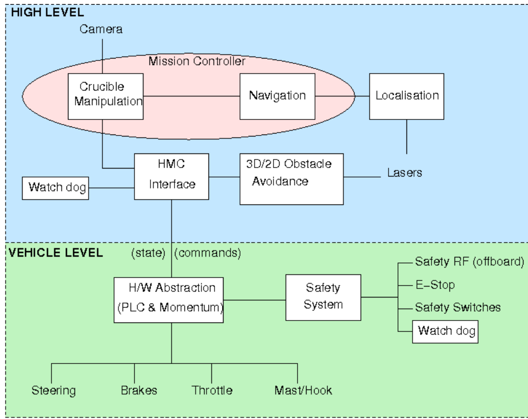

Figure 2. The HMC system architecture. The program blocks are shown in boxes or ellipses with unboxed text representing physical parts of the system.

Since the lasers only provide planar coverage about their horizontal axis approximately 1.2 m from the ground, objects above and below the scan plane will not be detected. However, in the target environments, this will detect a vast majority of objects. We have also mounted an additional laser from the top of the vehicle angled downwards over the engine, which is its normal direction of travel. This laser intersects the ground at approximately 25 m from the vehicle and is used for more advanced object detection.

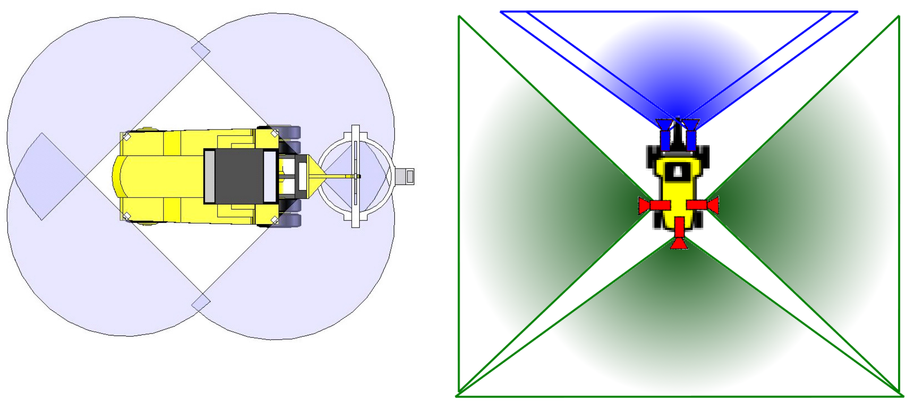

Figure 3. (Left) The lasers are located at each corner of the vehicle and offer overlapping planar coverage out to approximately 30m (indicated by purple sectors). (Right) Camera locations and their fields of view shown on the autonomous HMC. Red cameras are used for pedestrian detection and vision-based localization. Blue PTZ cameras are used for payload manipulation.

Safety Systems

Safety is paramount when considering the size, weight and payload of an industrial vehicle. It covers areas from protecting objects in the environment such as people, vehicles and infrastructure, to ensuring communication paths between software modules and between software and hardware module.

We have designed the autonomous HMC with a number of physical and Radio Frequency (RF) safety interlocks. Physical interlocks consist of door switches, E-Stops and key switches. These allow a physical lockout of the vehicle to prevent inadvertent autonomous operation. The RF interlock consists of a receiver on the vehicle connected to the E-Stop circuit, and a remote portable transmitter. If the receiver does not receive a heartbeat signal within a window of a few milliseconds, it initiates an E-Stop. A remote-triggered

E-Stop can also be initiated from the transmitter. This allows the vehicle to operate without requiring someone onboard. Activating an E-Stop brings the vehicle to a quick halt and shuts down the engine. Hydraulic controls are also locked at their current position.

The software safety systems consist of high-level velocity control when objects are detected close to the vehicle through the peripherally mounted scanning lasers, and low-level watchdog checks between interface level software and the low-level control software. A watchdog timeout initiates an E-Stop since it indicates a communication problem between modules.

Main Modules

Localization

The localization systems are responsible for determining the position of the vehicle in site coordinates, whether it is in sheds or outside.

The choice of localization system depends on the required accuracy for the operations and the sensors available. For autonomous vehicle control, high precision is required in certain parts of the environment such as around infrastructure and people, and during payload manipulation, whereas traversing along wide dedicated roads can allow a less precise method. The main environment sensors on the autonomous HMC are the scanning laser rangefinders as they have been tested over many years in our industrial projects and proven robust and consistent. The camera suite acts as a secondary sensor. We have developed localizers (described below) for each sensor. The laser-based localizers are simpler and more robust than the camera-based since they are not affected by lighting variations or absence and only consist of range and intensity data. The purpose of implementing multiple localizers is to allow for cross-checking and redundancy.

Laser Beacon-based Localization

The primary localization system needs to be reliable and accurate. Artificial landmark approaches allow the user to explicitly effect the accuracy of the system without having to rely on the complexity of detecting key salient features from the environment as is common in other methods. The scanning lasers return both a range profile and the intensity value of the scan points depending on their reflectivity. This feature is exploited in laser beacon-based localization whereby highly reflective marker beacons are placed around the environment at intervals depending on the accuracy required (typically between 5-20m) as shown in Figure 4. A drawback of this approach is the beacon locations need to be defined through manual surveys which can be costly and time-consuming. The beacon locations are stored in a database and cross-referenced during operation to determine the vehicle's location through Bayesian estimation methods.

Figure 4. The main operational area for the HMC. The beacons are evident by the reflective tape on bollards around the area. The area is approximately 50m x 35 m.

Laser-based SLAM

Simultaneous Localization and Mapping (SLAM) is a method that identifies and extracts features from an environment to create maps without any a-priori knowledge. As part of this process, the location of the vehicle is determined using its reference map and new areas added as identified. There are many challenges with implementing a successful SLAM system such as determining key 2D or 3D features of the environment, calculating the vehicle's motion, and recognising when the same area has been mapped which can occur sometime after it was last visited. In our implementation (Bosse & Zlot, 2008), the system is robust to people and vehicles moving around and can assimilate permanent changes to the environment into the reference feature database. This system is useful for tracking, rather than controlling vehicles around a worksite. In industrial environments, SLAM methods are useful for creating maps in all areas without the need for additional manual processes. While they can show high performance in mapping and using these maps for localization, their coordinate system may not be metrically linear due to slight discrepancies in local map matching. Consequently, their coordinate system may not map exactly onto site coordinates which can be problematic for precisely navigating an autonomous vehicle.

Vision-based Localization

The vision-based localization system allows for redundancy as both a localization source, and as a separate physical system to the laser-based localizers. Therefore, if the laser system fails, there is still a localization source available. Our vision-based localizer uses inexpensive off-the-shelf Basler cameras. As with other localizers, it uses a landmark database. The landmarks in this case are the significant edges of major infrastructure such as shed doors and building edges (Figure 5). In our initial approach, these features were derived by manual surveying. More recently, we have developed an automated approach (Borges et al, 2010) using a 3D SLAM system (Bosse & Zlot, 2009). The localizer operates by processing the incoming image stream into its edge-based equivalent, and comparing that to the landmark database. It uses Bayesian methods to estimate the location based on finding the best match of the processed image to the database (Nuske et al, 2008). We have conducted experiments to demonstrate the utility of using this system as a backup to the primary beacon-based localizer with successful results during daytime operations (Roberts et al, 2008).

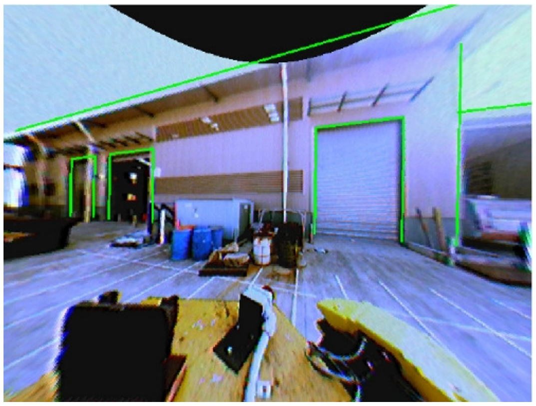

Figure 5. Building edges form the main landmarks for the vision-based localisation system. The raw image is converted into its edge-based representation and fitted to an onboard model (shown in green) to determine the vehicle's location.

Navigation

The navigation system uses waypoints derived automatically by manually driving the required route of operations. Waypoints are recorded after a certain change in distance since the last waypoint or a certain change in vehicle heading. Each waypoint also contains a velocity so ramping speeds can be utilised for smoother navigation. The resulting waypoint list is split into segments with each segment being a homogeneous action such as a forwards traverse (used for normal navigation) or backwards traverse (used for crucible manipulation tasks). Within a segment, the navigation system switches to the next waypoint in the list when it is close to the current waypoint.

Laser-based Object Detection

Object detection consists of determining the locations of obstacles and the payload.

2D Obstacle Detection

Using the lasers mounted around the vehicle, when the proximity of nearby objects such as people or infrastructure reduces, the vehicle slows and will eventually stop if they get too close. This is disabled for the hook-end when the HMC is approaching or carrying the payload.

3D Obstacle Detection

Since the peripheral lasers only detect a 2D plane around the vehicle, objects above and below will not be detected. To alleviate this issue, the top-mounted laser is used. Since it is angled down, as the vehicle moves forwards, any objects in front of the vehicle are scanned and can be used to create a local 3D map. Since the vehicle is autonomous and following a known path, its trajectory can be superimposed on the map to determine if there is the potential for a collision. The size of the object, vehicle wheeltracks and width are taken into account for this calculation. Any object deemed too 'big' will prompt the vehicle's navigation system to determine an appropriate mitigation action. At this point in development, the vehicle will simply stop and notify a supervisor. Path replanning and obstacle avoidance are subjects of future work and are highly dependent on the area of operations. For example, in some areas such as a narrow road, it may be dangerous for the vehicle to manoeuvre around the object and travel off the road where the tires may sink into the ground and render the vehicle immobile.

Payload Detection

The peripheral lasers also serve the function of identifying the payload in the environment. This is a secondary system to the vision-based crucible detection system described in the next subsection. The laser-based system 'learns' what the payload looks like through various laser scans as the vehicle travels around. Scans falling on the payload are saved in a database during a training phase. During the application phase, incoming laser scans are matched to the database and a confidence score incremented in locations where there are a number of unique matches. When the confidence passes a threshold, the associated location is labelled as a payload location (Tews, 2011).

Crucible Operations

The key functionality of a forklift-type vehicle is its ability to handle the payload. Two main operational phases can be distinguished: pickup and drop off. Drop off is an easy manoeuvre from an automation point-of-view once the vehicle is at the dropoff location. No sensing is required and a simple ballistic manoeuvre is sufficient. The pickup manoeuvre is more difficult. It can be divided into two steps: first, an approach where the hook is visually guided toward to the pickup point in the middle of the crucible handle, then the actual pickup. The latter is an easy manoeuvre, again a ballistic movement, similar to a drop off.

The approach part is more complex. It is principally based on the onboard mast-located PTZ webcams detecting the crucible from about 20m. We use specially designed barcodes placed on the bale arm to identify the crucible visually to the HMC (Pradalier et al, 2008). As with most outdoor computer vision applications, it requires proper management of sensitivity to lighting conditions.

Mission Planning (Tasking) and Recovery

The Mission Controller is responsible for switching between tasks and monitoring their performance. A mission is a sequence of tasks with each task returning its status during execution. Once a task has finished, the Mission Controller selects the next task. Contingencies occurring during task execution cause the Mission Controller to select the contingency subtask for that task. For example, a missed crucible pickup will trigger a “missed approach” signal and the HMC will move away from the crucible and try to pick it up again. During one of the 5 hour continuous operation experiments, the only halt in operations occurred after approximately 4 hours when the remote Safety RF unit's battery went flat. This triggered an E-Stop on the vehicle. When the battery was replaced, the vehicle was restarted and as it was about to pickup the crucible at the time of the E-Stop, it executed a missed pickup subtask and successfully continued its operations. In a system fully integrated into a worksite, the missions can be allocated by the schedulers.

APPLICATIONS OF COMPONENT TECHNOLOGIES

The systems discussed in the preceding section form the main technology components for the autonomous HMC. Their joint utility is for creating a reliable and functional autonomous vehicle. However, some provide utility as stand-alone units or have spawned new technologies useful to the overall goal of improving industrial operations through automation. A selection of these are described next.

Anti-collision systems

As discussed in the Object Detection section, the lasers effectively provide proximity detection zone in a relatively horizontal plane around the vehicle. Any object impinging on this plane will result in the vehicle slowing to a stop depending on the range. This simple approach can be implemented into the control system on manned vehicles that can reduce the speed or stop the vehicle in the same manner.

Pedestrian Detection (onboard and offboard)

Situational awareness for industrial vehicle operators is crucial to ensure safety of personnel and equipment. We have implemented camera based pedestrian detection systems on the autonomous HMC, which can reduce the risk of collision between the vehicle and people. Two different approaches have been developed, using statically located offboard cameras and onboard cameras.

Offboard cameras cover areas that are not necessarily visible from the vehicle. While human drivers and onboard sensors are able to detect pedestrians within line-of-sight, in complex environments obscured dynamic objects can unpredictably enter the path of a vehicle. The safety system we have developed (Borges et al, 2010) integrates a vision-based offboard pedestrian tracking subsystem with the onboard localization subsystem discussed earlier in this paper. This combination enables warnings to be communicated to the operator, and effectively extends their field of view to include areas that would otherwise be blind spots. A simple flashing light interface in the vehicle cabin (illustrated in Figure 6) provides a clear and intuitive interface to alert drivers of potential collisions. The pedestrian detection is carried out with efficient computer vision pattern recognition techniques developed by our research group (Borges, 2011; Johnsen & Tews, 2009). Hundreds of hours of tests under real operations have illustrated the applicability of the system, with positive feedback from industrial drivers.

Onboard cameras provide autonomous pedestrian detection from the perspective of the vehicle, where cameras can be mounted facing forward, sideways or to the back, as illustrated in Figure 3. Any pedestrian close to the HMC triggers a warning signal to the driver.

Figure 6. Lights mounted inside the cabin (orange circles). During manual operation, these lights alert the driver of a potential collision route between the vehicle and a nearby pedestrian.

More recently, we have extended the system to perform joint pedestrian detection (Miseikis & Borges, 2012), combining information from on-board and off-board cameras simultaneously. This novel approach is currently being trialled. Preliminary positive results indicate the method can further improve the pedestrian detection system reliability.

Vehicle Utilization

Localization can be used as a standalone component requiring only lasers or cameras connected to an onboard computer. This allows a wide variety of applications and analyses to be undertaken in efforts to reduce maintenance, improve safety and efficiency as described below.

Vehicle Tracking

The data from the localizer consists of a time-stamped position, from which, the vehicle's velocity can be determined. This information can be logged for a holistic analysis of vehicle usage statistics associated with areas on site. E.g, all transit times over a shift from payload pickup to delivery between certain locations can be recorded and automatic reports generated. Live vehicle locations can also be displayed on an online site map, including whether they are carrying a payload which can be useful for the schedulers.

Event Analysis

Adding a simple G-force sensor (Inertial Measurement Unit) to the localization information stream allows for monitoring and reporting high G events such as payload contact, collisions, fast cornering and braking, and even potentially identifying where road surfaces have deteriorated causing bumps.

Automatic Speed Reduction

Since vehicle position and velocity data is available from the localization system, it can be used to impose speed restrictions on the vehicle aligned with those required in different areas. This would require interfacing with the vehicle's control system (brakes, throttle) and storing the different area boundaries with their speed limits in the onboard computer. When the vehicle enters the area as determined by the localizer, the local speed limit can be referenced and enforced.

Asset Tracking

A problem that can arise with manned payload delivery is with wrongful delivery. This can be a costly operation for smelters where the wrong metal put in to a furnace can render it unusable. Our system (Markey & Tews, 2011) uses site cameras to detect and read identification labels on stationary payloads (Figure 7). The cameras can be mounted in key areas where verification checks are required, such as at a weighbridge or payload despatch zone. The system is highly reliable using check digit verification at the end of the label to ensure the label identification is accurate. The information can be used to track the payload and alert drivers if they are delivering to an incorrect location. The system can be extended to provide in-cabin information of routes, and pickup and dropoff locations, based on the asset number.

Initial detection **Zoomed in detection** **Text extraction of ID**

Figure 7. Identifying a crucible label in the environment without initial knowledge of its location.

Infrastructure Profiling

Apart from being used for extracting environment features, our 3D SLAM system can create relatively accurate models of environments or infrastructure. We have applied the technology to mapping caves, underground mines, building interiors, and worksites. The purposes have included visualization, temporal deformation analysis, change detection and determining the dimensions of features/objects. In the smelter-related applications, we have used it to map sites to create 3D edge models for vision-based localization and Carbon Bake Furnace (CBF) pit scans. An example of an interior shed mapping application is shown in Figure 8.

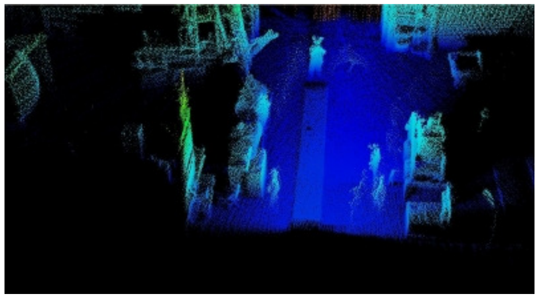

Figure 8. 3D view of the processed data from scanning our robotics bay. The colors are heights from blue (low) to red (high). Note the maintenance pit in the centre and HMC in the upper left.

SUMMARY

This paper has overviewed CSIRO's autonomous Hot Metal Carrier and the application of its main automation components. The components can be considered as part of an automation solution for either retrofitting or developing new material transport vehicles. These components can also be considered independently for their application in other roles such as improving pedestrian safety, vehicle and payload tracking. By considering them from both perspectives, components can be integrated into existing operations on the path towards more autonomous applications and vehicles with a value return along the integration path.

REFERENCES

Corke, P., Sikka, P., Roberts, J. & Duff, E. (2004). *DDX: a distributed software architecture for robotic systems*. Australasian Conference on Robotics and Automation.

Pradalier, C., Tews, A. & Roberts, J. (2008). *Vision-based operations of a large industrial vehicle: Autonomous hot metal carrier.* Journal of Field Robotics, vol. 25, no. 4-5, pp. 243 - 267.

Bosse, M. & Zlot, R. (2008). *Map matching and data association for large-scale 2D laser scan-based SLAM.* International Journal of Robotics Research, vol. 27, no. 6, pp. 667 - 692, June.

Borges, P., Zlot, R., Bosse, M., Nuske, S. & Tews, A. (2010). *Vision-based localization using an edge map extracted from 3d laser range data.* International Conference on Robotics and Automation, May, pp. 4902 - 4909.

Bosse, M. & Zlot, R. (2009). *Continuous 3D scan-matching with a spinning 2D laser.* International Conference on Robotics and Automation, May, pp. 4312 - 4319.

Nuske, S., Roberts, J. & Wyeth, G. (2008). *Visual localisation in outdoor industrial building environments.* International Conference on Robotics and Automation, May, pp. 544 - 550.

Roberts, J., Tews, A. & Nuske, S. (2008). *Redundant sensing for localisation in outdoor industrial environments.* 6th IARP/IEEE-RAS/EURON Workshop on Technical Challenges for Dependable Robots in Human Environments, May.

Tews, A. (2011). *3D payload detection from 2D range scans.* International Conference on Intelligent Robots and Systems, Sep, pp. 834 - 840.

Borges, P., Zlot, R. & Tews, A. (2010). *Pedestrian detection for driver assist and autonomous vehicle operation using off- board and on-board sensing.* Australasian Conference on Robotics and Automation. Dec.

Borges, P. (2011). *Blob motion statistics for pedestrian detection.* 2011 International Conference on Digital Image Computing: Techniques and Applications, Dec.

Johnsen, S. & Tews, A. (2009). *Real-time object tracking and classification using a static camera.* International Conference on Robotics and Automation - Workshop on People Detection and Tracking, May.

Miseikis, J. & Borges, P. (2012). *Joint human detection from on-board and off-board cameras* Workshop on Robots and Sensors Integration, IEEE International Conference Intelligent Robots and Systems, Oct.

Markey, E. & Tews, A. (2011). *A system for reliable text recognition in an industrial environment.* Australasian Conference on Robotics and Automation, Dec.