N Hillier¹, J Ryde¹ and E Widzyk-Capehart²

¹Autonomous Systems Laboratory, CSIRO ICT Centre, Brisbane ²CESRE, Minerals Down Under Flagship, CSIRO, Brisbane.

1 Introduction

Scanning infra-red laser range finders and millimetre-wave radar have seen extensive application in automation and other mapping scenarios in a wide range of research and commercial environments. There has, however, been relatively limited use of these sensing technologies in the mining industry.

This paper presents results obtained through the development of a unique mining terrain mapping system which combines millimetre wave and laser scanners to provide enhanced mining situational awareness capability. This module forms part of an electric rope shovel mining machine (see Fig. 1), a key component implementation of a larger automated swing assist system.

A comparison of the use of laser and radar data for the purpose of creating Digital Terrain Maps (DTMs) and object pose estimation is given which includes, a comparative evaluation of terrain mapping data obtained from prototype millimeter-wave radar and several commercially available 2D scanning lasers mounted on large rotating excavation machinery. The registered laser data is compared with the results of a terrain model obtained through stereo vision, and, subsequently, with the radar data.

The situational awareness component of the control system uses information obtained by scanning infra-red laser range finders, a millimeter-wave radar unit and an inertially aided RTK GPS to create a digital terrain map of the environment. This map is utilised in path planning by a third party trajectory planning system and to provide the location and orientation of the waiting haultruck to the shovel's collision avoidance system. The future automated system will allow the machine control system to autonomously perform swing-dump-return components of the shovel production cycle. The digging, crawler motion and overall operational planning will still be manually controlled by the operator at this stage and computer control can be seamlessly aborted by the operator at any stage.



Fig. 1. The P&H 2100 BLE electric rope shovel

2 Background

Lasers and radars employ significantly different imaging modalities and so have very different operatinal and performance characteristics. An analysis of the sensors' performance under environmental conditions typical of mining scenarios has been previously addressed by the authors in [Ryde and Hillier, 2009].

The use of millimetre wave radar in the mining industry has been limited to, primarily, static monitoring applications [Reeves et al. 2000, Macfarlane and Robertson, 2004, Brooker et al. 2005, Noon et al. 2002] or a few prototype mapping and localisation applications [Widzyk-Capehart et al. 2006; Brooker et al. 2007; Brooker et al. 2007, Scheding et al. 2002, Nienhaus et al. 2007].

Scanning laser range finders operating in the (near) infrared spectrum have found more widespread application in the mining environment, arguably due to the lower sensor costs and maturity of the technology. Examples include those of

Corke et al. 2000, Roberts et al. 2003, Singh 1997; Stentz et al. 1999, Duff et al. 2006, Hall and Keays 1993, Huber and Vandapel 2006, Shaffer et al. 1992, Baker et al. 2004, Nuchter et al. 2004.

During the last two decades, there have been significant improvements in the performance of laser range finding devices, in particular with regards to two areas: sensing capabilities in adverse visibility conditions, such as high suspended dust or water vapour (fog, snow, rain) loadings and direct viewing into the sun. Although, lasers are unable to range transparent objects like glass, it is rare that this is a limiting factor in outdoor environments.

Lasers have much higher range precision and significantly tighter beam widths than radar sensors, allowing for the creation of maps with higher accuracies. Scanning lasers are also usually associated with relatively high scan rates and lower costs than radar. They are considered a mature technology with multiple suppliers and low lead times.

Radar sensors are insensitive to suspended dust and water vapour loadings and the manufacturer of the radar sensor used here claims that a significant amount of debris build-up on the surface of the sensor can be tolerated before any degradation of the signal is noticeable (although this assertion was not specifically tested during the project). The larger beam width and lower operating frequency of radar provides more scope for measurement of multiple downrange targets along a single heading, even when visually obscured by intervening objects. However, the radar as a sensing technology has a significantly wider beam-width and lower range precision than the laser sensors. Some common materials are transparent to the radar (e.g. plastics typically yield low amplitude returns to radar) but, again, it is rare for this to be a limiting factor in outdoor environments.

The higher measurement uncertainties associated with radar often lead to more complex methods being pursued for map generation and data representation, for example volumetric evidence grids [Foessel 2000] by comparison to the simpler 2.5D representations that are usually employed for digital map representation.

Both technologies have gained general acceptance for being safe with the lasers presented here having class I ratings (IEC 60825) and there being no known adverse health risks for exposure to the millimeter-wave radar beam of the instrument used here (IEEE C95.1 1991).

2.1 Sensor performance characteristics

Two scanning laser range-finders, namely the SICK LMS291-S05 and the Riegl LMSQ120, and a 95G Hz scanning millimeter-wave radar (2D HSS) were considered in this study. A brief summary of their key performance characterising is presented in Table 1.

Performance Characteristic	Radar (2D HSS)	SICK (LMS291-S05)	Riegl (LMSQ120)
Min range	1 m ¹	0m	2m
Max range	70m	30-80m	75-150m
Range accuracy	>25mm ²	10mm	5mm
Beam width	1.50	0.7°	0.20
Field of view	360°	180°	80°
Minimum angular resolution	1.2°	0.25°	0.040
Measurement principle	FMCW ³ with CFAR ⁴ peak detection	Single shot time of flight with fog and pixel correction	Single shot time of flight with multiple echo discrimination
Operating wavelength	0.003m (95GHz)	905 nm	"near IR"
Max scan rate	25 Hz ⁵	75 Hz ⁶	100 Hz ⁷

Table 1. Summary of commonly identifiable performance parameters of the tested sensors.

3 Field testing on mining equipment

Four sensors were installed onto a P&H 2100 BLE electric face shovel (Figure 2) for the purpose of generating digital terrain maps and performing vehicle pose and volume estimation tasks. These maps are passed to a third party automation system to allow path planning and collision avoidance.

¹ Radar minimum range is configurable.

² Depends on radar cross section (RCS).

³ FMCW - Frequency Modulated Continuous Wave

⁴ CFAR - Constant False Alarm Rate

⁵ Data presented here was collected with a 3Hz scan rate, which provides the best performance for the radar from a spatial sampling and Doppler smearing perspective.

⁶ Scan rate at best angular resolution is less. Data presented here was collected with a 37.5Hz scan rate and 0.25 degree resolution.

⁷ Scan rate at best angular resolution is less. Data presented here was collected with a 10Hz scan rate and 0.1 degree resolution.

Machine pose information was provided to the sensors via an Applanix Inertially aided RTK GPS (IARTK) system installed on the shovel, which allowed measurement conversion to a common reference frame.

Four laser-scanning units (three SICK LMS291-S05 and a Riegl LMSQ120) and a 2D HSS – 94G Hz scanning millimetre-wave radar provided by the Australian Centre for Field Robotics (ACFR) are were used. The radar has been modified with a pan axis to allow 3D data collection independent of shovel motion. The mounting locations of the sensors on the electric rope shovel are shown in Error: Reference source not found. With the two SICK lasers mounted on either side of the boom pivot and the Riegl laser mounted adjacent to the right-hand SICK, a 3D point cloud data could be obtained during the shovel's motion. A single SICK laser was mounted on the side of the shovel to assist in the localisation of the haul truck. The radar was mounted on the left-hand side of the machine immediately above the lube room.



Fig. 2. Line drawings of the P&H 2100 BLE electric face shovel showing sensor locations.

3.1 Sensors performance characteristics

Fig. 3 and 4 show typical data obtained from the installed sensors during tests in a non-production area of the Bracalba quarry, North of Brisbane, Australia. The scene is centered upon the swing axis of the shovel with the vertical being North-aligned. There is a haul road, which runs from the North-West around to the East of the test area beyond an embankment populated with trees. Access to the shovel is via an access road to the North with light-vehicle parking to the North-West. A low face for digging is to the South-East of the shovel. A previously dug area with elevation below that of the shovel is to the South and West. There is a deeper pit beyond. The shovel crawler base is oriented almost East-West with a parked haul truck to the West of the shovel.

The point cloud plots highlight the sensor characteristics presented in Table 1. A summary of key findings from this data is as follows.

The Riegl has significantly better ranging capability than the other sensors (in practice: to about 140m) showing detail of the haul road to the North and, through gaps in the trees, to the East. Walls of an adjacent pit to the North-West and South-West are also visible. The Riegl also shows a better return rate off glancing surfaces and standing water, evidenced by the flat area to the South-West of Fig. 4(c) (a previously dug region about 5-10m below the shovel, containing some flooding on the Southern end).

The higher angular resolution of the sensor is also apparent by the point cloud density in regions further from the map origin. Some ghosting effects are also apparent in the field data on items of high reflectivity (e.g. vehicle number plates and some signage (Figure 5). Although methods are available to filter these false returns, this is not an ideal approach as it leaves the sensor susceptible to (sometimes significant) blind spots.



Comparison of Scanning Laser Range-finders and Millimetre-wave Radar for Creating a Digital Terrain Map 7

Fig. 3. Point cloud showing sensor coverage from field data from a full rotation of the shovel. Grid lines are at 20m intervals. Red is data from the radar sensor, light blue and green are from two SICKs, dark blue is from the Riegl. Axis re aligned x-y: East-North.





Fig. 4. Point cloud from field data measured using the sensors described in this paper for one full swing of the shovel at moderate speed. Grid lines are at 20m intervals. Points are coloured by height (red is high, blue is low, white areas are occluded or otherwise indicate no return). Image (d) shows an overlay of point cloud data from all four ranging sensors. Axis re aligned x-y: East-North.

The SICK laser scanners provide a larger field of view than the Riegl allowing proprioception of the shovel, including the crawler tracks and rear-mounted cable-reeler, as seen near the map origin. The radar similarly provides such proprioceptive information (with a 360 degree field of view, scanning in the vertical plane) but the proprioceptive information is removed for data clarity. The wider field of view of the SICK and radar also allows sighting of items at higher altitude, including the tops of the trees to the East of the shovel. Manual filtering of the data for mixed pixel effects [Ye and Borenstein 2002] was not conducted due to the large data sets (although automated methods exist for the removal of such points [Tuley et al. 2005].

Data from the radar was originally processed in a first-point-return mode to ease interpretation from possible multiple echoes along the same heading. Analysis has shown that this mode is incorrect for accurate imaging of terrain surfaces as the glancing angle of incidence over the wide beam width reports a nearer target than is actually observed. A similar effect can be observed for similarly processed laser data, but due to the relatively narrow beam width of laser based sensors, such errors are typically negligible.

The radar data required thresholding to remove free-space clutter that occurs if no real targets are present in the beam. Figure 7 presents the distribution of returns against intensity for the radar that was used to select the cut-off value of 75dB. In general, the radar showed expected behaviour with lower point return densities following the filtering process and a higher uncertainty in the range and

angular measurements to the target than the lasers due to the wider beam width and lack of a-priori information on the target's orientation.

There is good correspondence between the filtered point clouds generated by the radar and those by the SICKs, with similar surfaces providing returns (generally those with a high angle of incidence to the beam) and both sensors unable to see the standing water. The radar occasionally showed spurious returns near items with a high-metallic content (e.g. the haul truck shown in Figure 6) when the side-lobes of the radar signal intercepted the metal item. In these cases, the sensor returned a signal with intensity above the free-space clutter filtering threshold, as detected by the beam side-lobe, but with angular offset aligned with centre of the beam.

The significantly lower data acquisition rates and scan-rates for the radar presented the most limiting constraint for the terrain modelling application presented here. The data for the radar presented in Figure 6(a) was collected over a period of approximately 680 seconds whilst the shovel was stationary, using the radar's pan axis actuation to give the 3D point cloud over a 180 degree field of view. This data shows a sparser point distribution than that collected by the laser in the Figure 6(b), which was collected over approximately 17 seconds during a slow shovel slew (for comparison, a 180 degree scan at this rate would take approximately 100 seconds).



Fig. 5. Point cloud data from the Riegl showing the effect of "ghosts" in the data (highlighted in the squares). Points are coloured by return intensity. The false points are due to the high-intensity points immediately below from the vehicle number plates.



Fig. 6. Radar and laser measurements: (a) point cloud data showing spurious returns near items with a high-metallic content (a haul truck) - the radar beam side-lobes are detecting the truck and returning an intensity above the threshold value resulting in points associated with the wrong angle; (b) same scene as in (a) by a SICK laser during a shovel rotation. The points are coloured by height to aid visual clarity.



Fig. 7. Histogram of radar point return intensity versus number of data points at that intensity. The minima around 75dB defined the threshold intensity used to remove free-space data clutter.

3.2 System performance for creating digital model

The total system performance for the task of creating reliable and accurate maps of the surrounding environment is limited by the maximum uncertainties of both the sensor measurements and those introduced by the compound calculations, which are required to register the sensor measurements into a local ground frame. The errors are due to: uncertainty in the estimate of the machine pose in the local ground frame, uncertainty in the time difference between pose and sensor measurements and uncertainty in the estimate of the pose of the sensor with respect to the machine itself (how the sensor is mounted).

To minimise uncertainty in the machine pose estimate, a commercial IARTK pose estimator was employed. System performance typically gave repeatability in position estimates to within approximately 5 cm and heading estimates to about 0.2 degrees at a 50 Hz update rate. Each pose estimate was time-stamped and network time was synchronised using the NTP protocol to be typically within 3 milliseconds between computers. Sensor measurements used linear interpolation between pose updates but did not associate with a pose (and the measurement was discarded) if the pose estimator was unable to achieve a good fix or there was more than 1 missing pose estimate (interpolations were not performed over periods greater than 0.04seconds).

Registration of the sensor location and orientation on the machine was performed by regression analysis between RTK-GPS surveyed target points in the field of view of the sensors and the estimates of these locations given the shovel pose and sensor measurement (similar to the method of US patent #20100034421). This method limits the sensor registration accuracies to be bounded by uncertainty in the measurement of the shovel pose estimator, sensor limitations and the accuracy of the survey points, although a "better" estimate is hopefully achieved using a large data set. We expect the sensor location and orientation errors to be similar to those of the pose estimator for the laser and the radar, although we have no reliable means of confirming this other than measurement repeatability.

To better characterise the accuracies of the constructed digital terrain map, comparison of the face profile at a distance of 15m to 40m from the centre of the shovel generated by the lasers was made to a commercially available 3D measurement system based on an independent sensing technology, the stereo vision solution provided by SiroVision [Williams et al. 2005] as shown in Figures 8 and 9..

A histogram showing the number of points versus distance error between the point clouds generated by a SICK mounted on the shovel and that provided by the SiroVision solution is presented in Figure 10.

Where the datasets intersect, all laser points are within 10 cm of the groundtruth (assumed to be the stereo vision) data. This includes data at a range of between 15 to 40 meters from the centre of the shovel. It should be acknowledged that both the laser and the stereo vision data will each have independent error that will contribute to the 10 cm difference. It is reasonable to state that the separate error of each sensor will be less than 10 cm if compared to an actual "ground truth".

Similar analysis between the point cloud estimates of all three mounted lasers showed less than 0.05m of error.



Fig. 8. 3-D stereo image created by SiroVision as an independent set for comparison of installed sensor system accuracy.



Fig. 9. A plot showing the error between the point cloud data obtained from the front right SICK laser and the stereo vision information presented in Figure 8. Blue points represent regions of low error, red represent high error. Green circles are control points used to register the two data sets.



Comparison of Scanning Laser Range-finders and Millimetre-wave Radar for Creating a Digital Terrain Map 13

Fig. 10. Histogram of the point cloud error presented in Figure 9 with bins at 5 cm intervals. The vast majority of laser data is within 10 cm of the stereo image.



Fig. 11. Histogram of distance error between the point clouds generated by the radar and a SICK when imaging the truck as presented in Figure 6. Bins are at 10 cm intervals. Both sensors mounted on the shovel.

These results are better than expected based on the maximum combined error possible from all system sources, which suggest that the measurement accuracy provided by the lasers would be limited to between 0.1m and 0.2m at a distance of 30m from the centre of the shovel.

Comparison of the point clouds generated by the radar and a SICK laser when imaging the truck as presented in Figure 6 are presented via a similar histogram in Figure 11. This shows that the error is typically in the region of < 0.2m. The authors deduce that this is primarily due to the larger beam width and relatively complex geometry of the item being imaged.

The means by which the data is represented, both internally to situational awareness module and externally to client applications (such as the trajectory planning layer), also directly affects the accuracy of the terrain information. For performance and representation reasons, terrain information is presented as a height-encoded occupancy grid with a spatial resolution of 0.5m in orthogonal axes. Thus, any single point in the terrain can be in error of up to $\pm 0.35m$ in the x and y directions. The height of each cell is typically provided as the median of a rolling window of the most recent data added to that cell, although methods exist to extract the maximum, minimum and variance in cell height. The median value filters out spurious data points to produce a more robust representation of the terrain, with the possible loss of high gradient information. In practice, testing during development has shown that this representation works well, although, in some extreme edge-cases, it may underestimate the height of a sharp obstacle by up to 0.1m.

3.3 System performance for pose estimation of waiting haul truck

Similarly to that of the terrain, the accuracy to which SAM can provide truck information is limited by six constraints:

- 1. The accuracy of the laser and radar sensor data (Ryde and Hillier 2009);
- 2. The accuracy of the shovel pose information;
- 3. The accuracy of the timing information between different data sources;
- 4. The accuracy by which the sensors are able to be registered;
- 5. The suitability of the sensors to the task of providing truck pose information; and
- 6. The means by which the truck information is extracted from the sensor data.

Points 1 to 4 have been covered in the preceding section and shown to be of minimal contribution to possible errors in the accuracy.

Currently, the truck pose is determined by a 2-stage process. Firstly, the truck location and heading in the x-y (horizontal) plane of the shovel is determined using the side laser (seeded by the GPS information from the truck). Secondly, the height of the truck is determined by extracting the height data out of the terrain

information in the region identified by the first step. The sources for error in the truck pose estimate are thus limited by:

- 1. The resolving power of the side lasers to determine the location and orientation of the truck. In particular, the angular resolution of the side laser scanner can introduce an error in truck location of up to 13 cm at a distance of 15m. This error will change depending on the relative location and orientation of the shovel and truck and may appear as a "jittering" of the truck pose estimate;
- 2. The accuracy of the terrain representation; and,
- 3. The particulars of the algorithm used to extract the truck pose.

Our evaluation of the truck pose estimation performance has resulted in truck pose estimates within 0.25m of the surveyed truck position, but the heading error has yet to be assessed.

4 Conclusions

The work presented shows the application of standard off-the-shelf scanning laser range-finder technology and a commercial ready prototype scanning millimeter-wave radar technology to a mine environment for the purposes of generating digital terrain maps.

It has been concluded that the limiting *sensor specific behaviours* for creating a digital model are:

- for the Riegl laser: data ghosts over items with high return intensities, unnecessarily high resolution for this application and (arguably) poorer performance in rain conditions;
- for the SICK laser: marginally poorer performance under heavy dust loading than the Riegl and a lower resolution in the returned intensity data (which makes it more difficult to estimate the sensors' locations and orientations on the machine when using the method employed here which depends strongly on locating control points via intensity information).
- for the radar: high uncertainty in the measurement, erroneous returns due to the side-lobes from metallic objects which would be difficult to identify or filter from the sensor measurement alone and lower data acquisition rates.

However, it has been shown that a scope to combine the output from the sensors exists to create a more robust representation of the surrounding terrain using:

Radar returns to provide a "rough-draft" of the surrounds in which significant obstacles (greater than 2m in size) are clearly identifiable as an obstruction (although not necessarily classifiable) and that is robust to adverse weather and dust;

- Laser information to provide the detail of the surrounds, and other information required for tasks such as volume estimation, segmentation and classification of objects and obstacles;
- Radar returns to determine when information provided by the laser sensors has been degraded by adverse environmental conditions such as rain, mist and dust.

The fusion of data from the disparate systems needs to be further investigated to provide accurate and reliable information for the machine automation.

5 Acknowledgements

The authors would like to thank the Australian Coal Association Research Program (ACARP) for funding project C16031 which provided for the work presented in this paper. We would also like to thank the joint partners in this research project: the Cooperative Research Centre for Mining (CRCMining) and the Australian Centre for Field Robotics (ACFR). In particular, we would like to thank Mike Bosse (CSIRO) for laser sensor calibration and technical advice on the data acquisition and processing and Phil Soole (CSIRO) for providing the SiroVision data, Javier Martinez and Graham Brooker (ACFR) for provision and calibration of the radar and review feedback on a report on which some sections of this paper are based and Ross McAree, Anthony Reid, Kevin Austin and Andrew Denman (CRCMining) for managing the ACARP project, providing access to and operating the test facilities including the research shovel. Thanks also to all those that work behind the scenes to make projects successful.

6 References

- Baker C, Morris A, Ferguson D, Thayer S, Whittaker C, Omohundro Z, Reverte C, Whittaker W, Hhnel D, and Thrun S (2004) A campaign in autonomous mine mapping. Proc. of the 2004 IEEE International Conference on Robotics and Automation.
- Brooker G, Hennesy R, Lobsey C, Bishop M, and Widzyk-Capehart E (2007) Seeing through dust and water vapor: Millimeter wave radar sensors for mining applications. Journal of Field Robotics, vol. 24, no. 10, pp 527-557.
- Brooker G, Hennesy R, Scheding S, and Bishop M (2005) Development and application of millimetre wave radar sensors for underground mining. IEEE Sensors Journal, vol. 5, no. 6, pp 1270-1280.
- Corke P, Roberts J, and Winstanley G (2000) Experiments and experiences in developing a large robot mining system. Experimental Robotics VI, vol. 250/2000, pp 183-192, Berlin Heidelberg. Springer.
- Duff E (2006) Accurate guidance and measurement for excavators using a laser scanner. Technical Report C14043, ACARP.

- Duff E, Usher K, and Ridley P (2006). Swing loader traffic control. Technical Report C13041, ACARP.
- Hall D and Keays J (1993) Dozer track system phase 2. Technical Report C3045/C3056, ACARP.
- Huber D and Vandapel N (2006) Field and Service Robotics: Recent Advances in Research and Applications, vol. 24 of Springer Tracts in Advanced Robotics, chapter Automatic 3D underground mine mapping. Springer.
- Macfarlane D and Robertson D (2004) AVTIS a dual-mode imaging millimetre wave radar/radiometer for volcanological surveying. Geoscience and Remote Sensing Symposium, 2004. IGARSS '04. Proc. 2004 IEEE International, vol. 5, pp 3299-3302.
- Nienhaus K, Winkel R, Mayer W, Gronau A, and Menzel W (2007) An Experimental Study on Using Electronically Scanning Microwave Radar Systems on Surface Mining Machines. IEEE Radar Conference, Boston.
- Noon D, Reeves B, Sticley G, Holmstrom M, Smith G, and Longstaff D (2002) Development of slope stability radar. Technical Report C10036, ACARP.
- Nuchter A, Surmann H, Lingemann K, and Hertzberg J (2004) 6D slam with an application to autonomous mine mapping. Proc. of the 2004 IEEE International Conference on Robotics and Automation.
- Reeves B, Stickley G, Noon D, and Longstaff I (2000) Developments in monitoring mine slope stability using radar interferometry. Geoscience and Remote Sensing Symposium, 2000. Proc. IGARSS 2000, IEEE 2000 International, vol. 5, pp 2325-2327.
- Roberts J, Winstanley G, and Corke P (2003) 3D imaging for a very large excavator. International Journal of Robotics Research, vol. 22, no. 7/8, pp 467-478.
- Ryde J and Hillier N (2009) Comparison of Laser and Radar in Adverse Environmental Conditions. Journal of Field Robotics, vol. 26, no. 9, pp 712-727.
- Shaffer G, Stentz A, Whittaker W, and Fitzpatrick K (1992) Position estimator for underground mine equipment. IEEE Transactions on Industry Applications.
- Scheding S, Brooker G, Hennessy R, Bishop M, and Maclean A (2002) Terrain imaging and perception using millimetre wave radar. Proc. 2002 Australasian Conference on Robotics and Automation, pp 60-65, Auckland.
- Tuley J, Vandapel N, and Hebert M (2005) Analysis and removal of artifacts in 3-D ladar data. Proc. of the 2005 IEEE International Conference on Robotics and Automation, Barcelona, Spain.
- Widzyk-Caperhart E, Brooker G, Hennesy R, Lobsey C, and Scheding S (2006) Application of millimetre wave radar sensor to environment mapping in surface mining. 9th International Conf. Control, Automation, Robotics and Vision, pp 1-6.
- Williams D.J., Stolberg D.J., Poropat G. V., Soole P. K., (2005) Monitoring erosion off unvegetated mine tailings facilities and natural slopes using high-resolution, digital stereo-photogrammetry. Proceedings of the 4th International Conference on Environmental Geotechnics (Brazil 2002)
- Ye C and Borenstein J (2002) Characterization of a 2-D Laser Scanner for Mobile Robot Obstacle Negotioation. Proc. of the 2002 IEEE International Conference on Robotics and Automation, Washington DC.