The 2009 IEEE/RSJ International Conference on
Intelligent Robots and Systems
October 11-15, 2009 St. Louis, USA

Towards Reliable Perception for Unmanned Ground Vehicles in Challenging Conditions

Thierry Peynot, James Underwood and Steven Scheding

Abstract— **This work aims to promote reliability and integrity in autonomous perceptual systems, with a focus on outdoor unmanned ground vehicle (UGV) autonomy. For this purpose, a comprehensive UGV system, comprising many different exteroceptive and proprioceptive sensors has been built. The first contribution of this work is a large, accurately calibrated and synchronised, multi-modal data-set, gathered in controlled environmental conditions, including the presence of dust, smoke and rain. The data have then been used to analyse the effects of such challenging conditions on perception and to identify common perceptual failures. The second contribution is a presentation of methods for mitigating these failures to promote perceptual integrity in adverse environmental conditions.**

I. INTRODUCTION

This work presents a step towards understanding and developing integrity in perceptual systems for Unmanned Ground Vehicles (UGVs). Its purpose is to give a better understanding of what can constitute perceptual failure and how it may be detected and its effects remediated. Such failures would not just include hardware faults (e.g. a broken sensor), but also model failures due to adverse environmental conditions or algorithm failures.

Although researchers still invest significantly in developping sensors that get individually more and more reliable, it is reasonable to state that there is no such thing as one single perfect sensor that will succeed in providing an accurate perception of the world in every situation (even assuming that all systematic[1] errors are properly modeled and taken into account, as in [14]). Therefore, to achieve long-term autonomy with a high level of integrity, autonomous vehicles should be equipped with multiple sensors with different physical properties. This information must then be fused in an *intelligent* way [9], [3], to handle challenging environmental conditions. Some of these sensors may even measure the same quantity by using a different underlying physical process (e.g. measuring the range to obstacles using a laser scanner and a radar scanner) with high benefit for the perception system. Exploiting this *redundancy* allows some failure situations to be identified.

A good example of an advanced perception system for UGVs in real-world outdoor conditions is in the DARPA PerceptOR program, which was undertaken by Carnegie Mellon University [4]. In particular, the extensive experimentation that was conducted in diverse environments provided higly valuable feedback on the remaining problems that researchers in perception need to tackle. Of the issues that resulted in poor performance, particularly in the earlier tests, the vast majority could be described as sensor interpretation errors due to unmodeled environmental conditions. In particular, airborne dust (a common phenomenon for UGVs operating in dry areas) has been noted as an unsolved problem for state-of-the-art systems in the literature (e.g. [7], [6], [8]). The detection of airborne dust particles is used as a recurring theme in this work to illustrate the potential for multi-modal sensor fusion to increase perception robustness. More generally, effects of challenging conditions such as dust, rain or smoke need to be analysed to identify ways of detecting and handling these situations accurately, to ensure perceptual integrity. In this work, an analysis of some of these phenomena is proposed, based on a large amount of data collected with a UGV on real world challenging conditions. Some options to handle perception failures are then discussed.

Most of the work in the literature on detecting and handling perceptual faults focuses on *abrupt* faults, mainly due to hardware problems (e.g. [2], [15]). This is sometimes extended to major software issues [11] or major measurement errors [5]. This work looks into more general problems such as the mis-interpretation of sensor data due to model inaccuracy, although, as will be seen later in the discussion, the methods proposed in this paper to manage model failures should also be able to detect hardware faults such as a broken sensor. Note that in the rest of this paper, *failure* will be used as a general term, including modelling errors. The term *fault* will be associated to abrupt problems such as broken hardware.

This article is structured as follows. Section II presents the multi-sensor UGV that has been used for this work and the data-sets that have been collected and published to study perception reliability in challenging conditions. Section III provides an analysis of the performance of the various sensors, in representative situations. Section IV proposes solutions to deal with some of the perception failures illustrated in the preceding section.

II. MULTI-SENSOR DATA FOR INTEGRITY EVALUATION

To begin to address the issues of sensor data integrity, synchronised data have been gathered from a representative UGV platform using a wide variety of sensing modalities

T. Peynot, J. Underwood and S. Scheding are with the ARC Centre of Excellence for Autonomous Systems, Australian Centre for Field Robotics, The University of Sydney, NSW 2006, Australia tpeynot/j.underwood/scheding @acfr.usyd.edu.au

[1]repeated errors that can be properly modeled, e.g. noise in measurements or odometry errors due to inaccurate wheel radius value

[13]. These modalities were chosen to sample as much of the electromagnetic spectrum as possible, with the limitation that the sensors be feasible and available for use on UGVs.

These large volumes of data are now available to the community to evaluate the performance of perception algorithms, such as for obstacle detection or terrain interpretation. They were gathered in various conditions that are representative of actual situations a UGV is likely to encounter. The authors believe this constitutes a valuable contribution, since they are not aware of any similar, publicly available data-sets that cover this variety of sensors and situations. Note that these data-sets are provided with calibration data-sets (for relative sensor positioning) and calibration parameters. They can be found at the address indicated in [13], with a technical report to describe content and format in detail.

The following sub-sections describe the UGV system and sensors that were used to collect the data. Specific details about the data-sets are also given.

A. System Description

The vehicle that was used to gather the data is an 8 wheel skid-steering platform (Argo, see Fig. 1), with a reliable navigation system composed of a Novatel SPAN[2] System and a Honeywell IMU[3].

Fig. 1. The Argo Vehicle

The vehicle was equipped with the following exteroceptive sensors, mounted on a common frame (see Fig. 2):

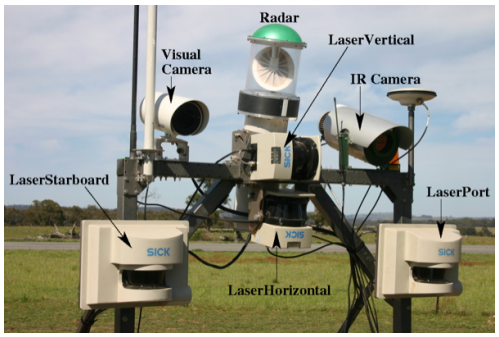

Fig. 2. Sensor Frame on the Argo

[2]Synchronised Position Attitude & Navigation
[3]Inertial Measurement Unit

Fig. 3. The Argo UGV sensing the *static* trial area

- four Sick laser range scanners, with 180 degree angular range and 0.25 degree resolution,
- a $94GHz$ Frequency Modulated Continuous Wave (FMCW) Radar (custom built at ACFR for environment imaging), with a range resolution of $0.2m$,
- a Prosilica mono-CCD megapixel gigabit ethernet colour camera, with an image resolution of 1360×1024 pixels,
- a Raytheon thermal infra-red (IR) camera, with a spectral response range of $7 - 14\mu m$,

B. Datasets

Synchronised data were collected from all of the sensors (in addition to proprioceptive sensors such as wheel encoders) in two different situations: *static* and *dynamic*, refering to the vehicle motion. These data were obtained for various types of environments and conditions. Note that the environments considered in this paper contained no dynamic objects other than moving vegetation (e.g. tree branches).

1) Static tests: First, to properly evaluate the quality of perception and make comparisons of the effects of different conditions, data were collected in a *static* situation: the vehicle was not moving and the observed scene remained the same in all tests. Fig. 3 shows the artificial outdoor environment that was used, surrounded by a large metal frame of $9.3m \times 4.3m$, which supported a sprinkler system to generate artificial rain. It contained artificial features of simple geometry such as boxes or vertical poles, as well as natural features such as a tree branch (see [13] for a complete description of the object sizes and positions). The analysis in this paper is focussed on these *static* data.

2) Dynamic tests: In a second phase, data were collected while the vehicle was moving in three different environments. As the corresponding *dynamic* data-sets are not further analysed within this paper, the reader is invited to refer to [13] for more details.

3) Challenging conditions: The *challenging* environmental conditions (in both static and dynamic cases) included:

- Presence of particulate dust (in various amounts)
- Presence of smoke
- Presence of rain (various flow rates)

In the following sections, *clear* conditions will refer to the nominal situation in the absence of any of these challenging conditions.

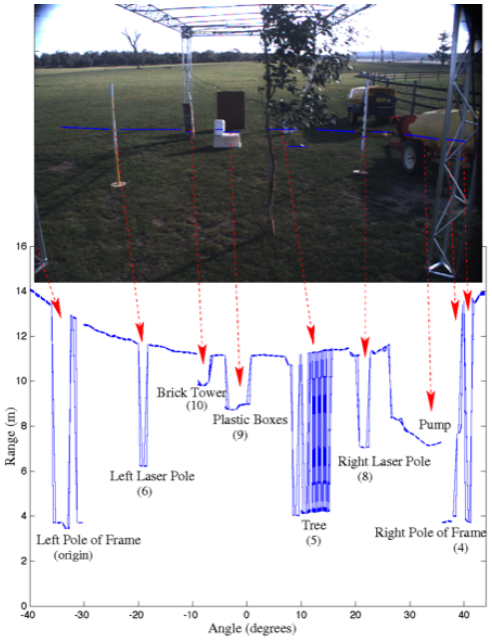

Fig. 4. Colour image of the static scene (above) from the on-board camera and the corresponding laser scan display (below), in *clear* conditions, over the 2 minute complete *dataset 02* (displayed with solid lines). The objects labels refer to [13].

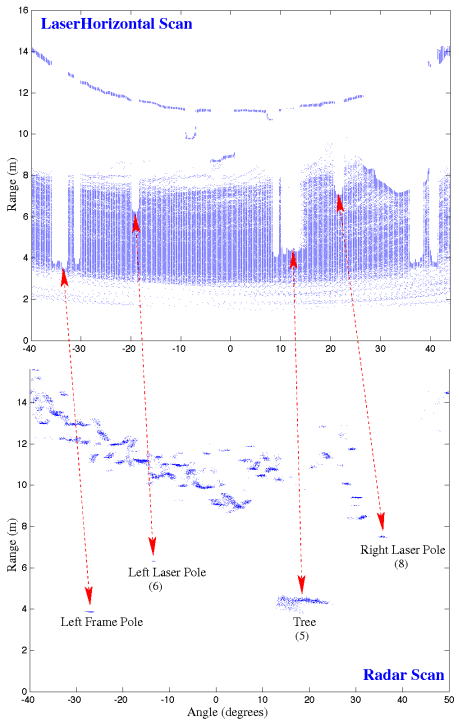

Fig. 5. Laser (above) and Radar (below) scan points over a complete 1 minute data-set, with presence of dust (*dataset 05*). Objects are labelled as in Fig. 4. Airborne dust particles can be seen in the laser scan, whereas the radar scan is almost identical to the one obtained in clear conditions (*dataset 02*).

III. ANALYSIS OF SENSOR DATA IN CHALLENGING CONDITIONS

This section provides an analysis of the effect of the aforementioned challenging conditions on perception when using the sensors described in section II-A. It is based on the *static* data-sets presented in section II-B, which will be referred to using a number corresponding to the labelling in [13] (e.g. static data-sets correspond to numbers from *01* to *21*).

To ease the interpretation of laser scans, Fig. 4 presents the correspondances between objects in the scene as perceived by the on-board colour camera and a single laser scan (displayed as range function of angle), in *clear* conditions. The laser selected for illustration here is the one labelled *LaserHorizontal* in Fig. 2. It was pointing down slightly from horizontal (8 degrees of pitch), with no roll, its beams touching the ground at a distance of about $12m$ in front of the robot. This is a classic configuration for terrain interpretation and obstacle detection on this kind of vehicle. Note that the conclusions that are drawn in this paper for this laser apply similarly to the three other laser devices, since they have the same physical properties. For convenience, all laser and radar scans displayed only show the range of angles corresponding to the perception of the test area.

A. Effect of Dust on Range Sensors (Laser and Radar)

Fig. 5 shows, on the same graph, all laser scans of the static test environment over one minute of data acquisition in the presence of heavy dust. It can be seen that the airborne dust particles are very often detected in mid-air by the sensor, which prevents laser-based perception from detecting the actual obstacles behind the dust cloud. This needs to be taken into account to avoid considering dust as a large obstacle, which would constitute an interpretation error, i.e. a perception failure. On the contrary, the radar range information is not significantly affected (Fig. 5, bottom), apart from a slight attenuation of the reflectivity returned. To evaluate the effect of these conditions on radar reflectivity, the highest peak of average reflectivity in the radar scans is considered. It corresponds to the brick tower seen on the left part of Fig. 4. The attenuation of this peak compared to the *clear* data-sets (namely *datasets 02,11,12*) is: about 5% with presence of heavy dust (*dataset 05*), 3.5% with smoke (*dataset 07*) and 2% in presence of heavy rain (*dataset 08*). This does not significantly affect the range measurement ability of the radar.

Note that according to these data-sets, it appears that the effect of smoke on laser and radar range sensors is very similar to the effect of dust (see [13] for illustration).

Although the radar is not affected by airborne dust or smoke, the range measurements it provides in all conditions are noiser and less accurate than those from the laser. It is therefore desirable to keep the laser information in the perception loop whenever possible, while filtering the inappropriate data (in this case the points returned by the dust cloud for example, see section IV-B).

B. Effect of Dust on Camera Images

To illustrate the effect of dust on the colour camera perception, the average Shannon information is computed on the 8-bit visual images (converted to grey-scale) in the

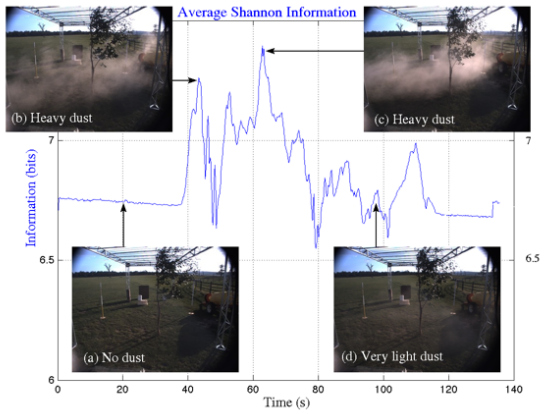

Fig. 6. Shannon Information in the presence of airborne dust (*dataset 05*). The images are snapshots from the colour camera at the corresponding time. It can be seen that the information is quite steady in clear conditions (a) before being significantly affected in the presence of airborne dust (b)(c).

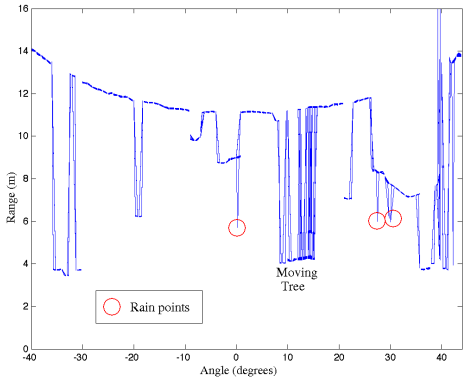

Fig. 7. Laser scans with heavy rain (displayed with connecting lines) over a 2 minute acquisition with constant presence of rain (*dataset 08*). Note the isolated points (circled) out of the actual objects, most probably corresponding to rain drops returns.

presence of dust. Given a discrete set of possible pixel values A_X, the average Shannon information $H(X)$ of an image, in bits, is given by [10, §2.4]:

$$H(X) = \sum_{x \in A_X} P(x) log_2 \frac{1}{P(x)}$$

where the probabilities $P(x)$ were calculated using a normalised histogram. The Shannon information obtained is shown over time in Fig. 6, for the same data-set considered for the laser and radar (i.e. *dataset 05*). It can be seen that in this static environment the presence of dust has a clear impact on the evolution of the level of information, compared to the values in clear conditions, which appear quite steady (since the scene does not change overall). Note that the presence of smoke has a similar effect on the visual images [13]. Airborne dust also affects the information content of the IR images, but to a lesser extent than the colour images. Interestingly, the presence of smoke does not affect the IR Shannon information.

The potential for the use of Shannon information to illustrate and detect the effect of such environmental conditions on camera images will be discussed in section IV-A.

C. Effect of Rain on Range Sensors

In the static tests, artificial rain was generated using sprinklers attached to the top of a metal frame surrounding the test area (cf. Fig. 3), starting at about $3m$ in front of the robot's sensor frame. The flow rate was designed to be constant for the whole area and over time (for a given data-set). Rain appeared not to affect the laser scans significantly, except for a small number of isolated returns from otherwise empty space (cf. Fig. 7). Within the angular region of the laser scan that covered the test area, an average of 0.2% of data appear to be affected by rain (laser returns from specific rain drops). Although this small percentage seems to indicate that the occurrences of such phenomena are quite sporadic, if unmodeled, these can cause highly significant perceptual failures (e.g. up to $3m$ of range error in the example of Fig. 7).

All of the rain-affected points are isolated, spatially and temporally, which would enable simple and efficient filtration. However, this would incur a penalty of reducing the precision of the perception (to about twice the angular resolution).

Further illustration of the effects of the challenging conditions considered in this section are provided in [13]. Methods for mitigating some of the other perceptual failures identified in this analysis are discussed in the following section.

IV. PERCEPTION INTEGRITY IN CHALLENGING CONDITIONS

In this section, methods for promoting perceptual integrity in the presence of adverse environmental conditions are discussed. Firstly, improvements can be made by modeling particular environmental effects on a sensor. This is illustrated by using an information theoretic approach to detect poor visual image contrast, due to phenomena such as airborne dust particles, large occlusions or hardware failure. Secondly, improving the integrity by fusing multi-modal sensor data to form redundancy is discussed. This is illustrated with an example that fuses laser and radar sensors to classify failure conditions that cause a disagreement between the sensors. This can be used to detect hardware faults such as a broken sensor, but also to detect more subtle failures, such as geometric misalignment errors or the interpretive errors caused by the effect of airborne dust clouds on the laser.

A. Improvements to the Perception Model

Failure due to errors in the perception model can be avoided or reduced in severity by increasing the fidelity of the model. Improvements can be made by accounting for nuances of the environment and by including specific failure modes of the perception. By conceiving of possible model deficiencies prior to deployment, or by rigorous testing to determine the causes of failure after they occur, additional complexity can potentially be incorporated into the model so that it better represents the true environment.

As an example, consider a video camera operating in the visible spectrum. There are many phenomena in natural

environments that cause a reduction in the utility of the sensor. These include low (or no) light conditions, near-field occlusions, or atmospheric effects such as dust or fog, all of which can reduce the contrast in the images. If the perception model assumes that the images have usable information in each frame, relating to the surroundings, then the presence of these environmental effects may reduce the robustness of the interpretation of the sensor data. The fidelity of the model can be increased by explicitly incorporating the possibility of such conditions.

Information theoretic measures such as Shannon information or entropy (average Shannon information) [10, §2.4] of the raw sensor data (see section III-B) can be used to determine when there is not enough information in the sensor stream to be useful for processing, with some independence from the algorithm that uses the data. Sensor performance metrics can be difficult to specify for complicated algorithms, whereas the information theoretic approach provides a general purpose measure of the utility of raw sensor information. Although this is generally true, the appropriateness of such metrics depends on the type of sensor and processing method, and it is certainly not guaranteed. For example, if an image processing algorithm is specifically designed to determine whether the images were taken at night or day, then an image with almost no Shannon information (all black) actually contains a large amount of useful information to that algorithm. Likewise, a reasonably small cloud of airborne dust appearing in the perceived environment (as in Fig. 6) can actually constitue an informative feature in the image. In general, however, streams of raw sensor data that contain little or no Shannon information are likely to be of little use to processing algorithms, and may even indicate a sensor or perception failure.

To test the appropriateness of such an information theoretic approach to detect adverse conditions in the field, the following experiment was performed. The Argo vehicle seen in Fig. 1 remained stationary by the edge of a dirt road, as a car was driven past at increasing speeds between approximately 10 and 90 km/h. A volume of dust was stirred into the air, behind the passing car. Anecdotally, the volume of dust is larger when the vehicle passes more rapidly. The Shannon information content of the colour camera is shown for this experiment in Fig. 8. Because the sensing platform is stationary, the information content of the background scene forms the baseline of the graph. There is a significant increase in information content due to the appearances of the car, which are relatively informative compared to the background scene. After the car passes, as the airborne dust covers most of the images, it reduces their contrast, resulting in a reduction of the information content. Interestingly, the magnitude of the reduction in information is correlated to the car's velocity. The Shannon information content can provide a useful metric relating to the clarity of the image, in situations where the environment can have an adverse effect on the contrast of the data. This could potentially be useful for providing a model of uncertainty for complex algorithms that use the data, where low information content signifies higher uncertainty of the results.

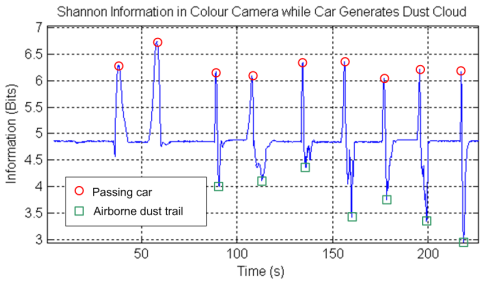

Fig. 8. The entropy or average Shannon information of images in a colour video sequence, in which a vehicle repeatedly drives across the camera viewpoint along a dirt road, stirring dust as it passes.

B. Multi-Modal Sensor Fusion

Perceptual integrity can be improved by incorporating sensors that employ different underlying physical properties. Fusing information from sensors with differing modalities into a perception model provides two main benefits: redundancy is achieved from the similarities between the sensors, and discrimination or classification power is achieved from the differences. Duplicate sensors could be used, however, detrimental environmental conditions will likely affect all of the sensors in the same way, which circumvents the redundancy. Multi-modal sensor combinations are less sensitive to this, because they react differently to their surroundings due to the differing physical properties. Furthermore, these differences can be exploited to discriminate and classify the nuances in the environment that give rise to the measurements. In this section, the discriminative power obtained from the fusion of a laser range scanner and a scanning radar is used to model airborne dust. This is a nuance in the environment that can cause significant misinterpretation of laser sensor data, when it is used in isolation. Although dust is studied in particular here, the model effectively captures the redundancy between the two sensors, to indicate a failure is present when the sensor data is contradictory. This captures other types of failure, such as hardware faults (a broken sensor), or geometric sensor alignment errors (poor calibration).

In Section III, it was shown that in the presence of airborne dust, the laser reports range readings that correspond to the cloud of dust, rather than the solid objects found behind. A naïve interpretation of this data might lead to the assumption that the dust cloud is a solid object. By contrast, it was shown that the dust is practically invisible to the radar. As can be seen in Fig. 2, the *LaserHorizontal* and the *Radar* were physically aligned to have a similar perspective of the environment, so the data from both of these sensors can be combined to detect when the laser has been affected by this phenomenon.

To classify laser returns that are affected by dust, a metric e_{rl} is formed for each matching laser and radar point within a synchronised scan pair. This metric is defined to be the squared error between the laser and radar range for a single bearing angle, in the sensor frame. When the environment

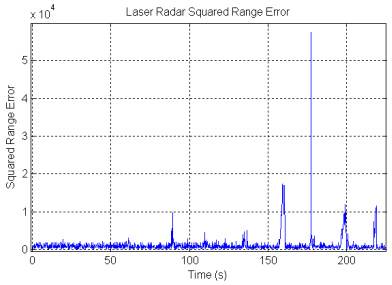

Fig. 9. The squared range error e_{rl} between a 2D laser range scanner and a 2D radar scanner. A car is driven on a dirt road through the viewpoint of the sensors, at increasing velocities. The first two passes at 40 and 59 seconds do not create a large volume of airborne dust, and there is no corresponding increase in the error metric. For all instances where dust is stirred, the error metric is above the noise floor.

affects the sensors differently, e_{rl} is large, and when the sensors agree, e_{rl} is small. For a laser range r_l, a radar range r_r, both at bearing θ, the error value is given by:

$$e_{rl} = (r_l - r_r)^2$$

The metric e_{rl} was calculated for the data-set from Section IV-A, which featured repeated passes of a car at increasing velocity. The result is shown in Fig. 9. For the first two passes of the car at times 40 and 59 seconds, little dust was stirred into the air and e_{rl} remains below the noise floor. Importantly, this shows that the movement of the car through the laser and radar viewpoint does not induce a significant error between the two sensors. This is because they are almost co-located and they are accurately time synchronised. When the car passes more rapidly, stirring more dust into the air behind it, the presence of airborne dust in the viewpoint of the sensors causes a discrepancy, and a corresponding increase in e_{rl}. For all instances where there exists significant airborne dust, e_{rl} is above the noise floor, clearly demonstrating the discrimination power of this combination of sensors for this phenomenon.

A classification data-set of approximately 55,000 laser and radar range pairs was gathered from a more typical operating environment of the Argo UGV. It was driven aggressively (with sharp turns), causing dust to be stirred into the atmosphere. This reduced the possibility of learning specific features of the controlled experiment with the passing car. The error metric e_{rl} was calculated for this data, and is plotted against time in Fig. 10. When the laser is affected by airborne dust, the errors are significantly higher than the noise floor .

Expectation maximisation [1, §9], [12] is used to learn a naïve Bayes binary classification for when the sensors agree or disagree. Disagreement is caused primarily by airborne dust. From the classification data-set of $55,000$ samples, a training subset of 2000 samples is used. The resulting Gaussian likelihood functions are difficult to visualise on the same graph because they vary significantly in scale, instead they are conveyed numerically:

$$\begin{aligned} Class &= [clear, dust] \\ P(e_{rl}|clear) &= \mathcal{N}(e_{rl} - 0.2043, 0.0816) \end{aligned}$$

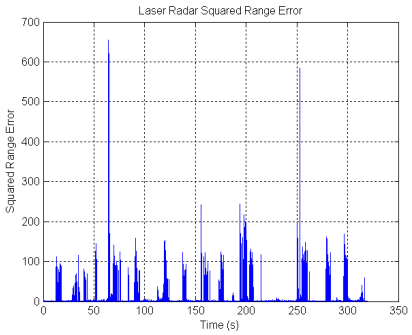

Fig. 10. The squared range error e_{rl} between a 2D laser range scanner and a 2D radar scanner, for a classification training data-set. The UGV was driven manually with fast turns, causing dust to be stirred, which in turn creates an error that is much higher than the noise floor.

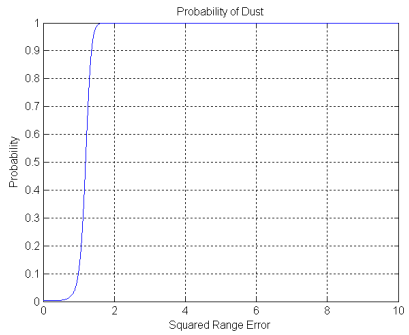

Fig. 11. The learnt probabilistic model, describing the probability that an error in range between the laser and radar sensors is due to airborne dust, as opposed to noise.

$$P(e_{rl}|dust) = \mathcal{N}(e_{rl} - 39.6013, 10970)$$

The difference in scale is due to the fact that e_{rl} is close to zero when the environment is clear, yet the dust induces a very large range of error values (with a standard deviation of $100m^2$) depending on factors such as the distance between the dust cloud and the ground behind.

Assuming a uniform prior, the probability of a sensor range being affected by dust is given by Bayes rule:

$$\begin{aligned} P(dust|e_{rl}) &= \frac{P(e_{rl}|dust)}{P(e_{rl}|dust) + P(e_{rl}|clear)} \\ P(clear|e_{rl}) &= 1 - P(dust|e_{rl}) \end{aligned} \quad (1)$$

The probability of the sensors being affected by dust is shown for different values of e_{rl} in Fig. 11. This figure illustrates the learnt laser-radar dust model. The class boundary point occurs at $e_{rl} = 1.19m^2$, meaning that range errors smaller than this are more likely to be due to noise than due to dust.

The model is then applied to the entire classification data-set to filter laser points that have been affected by airborne dust. In Fig. 12(a), the unfiltered laser data from the entire classification set are mapped to a three dimensional point cloud. In Fig. 12(b), the dust affected laser points have been removed, using the classification model in (1). The

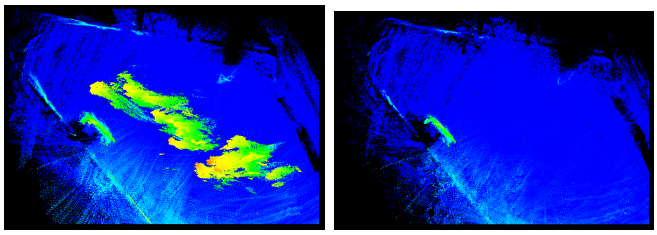

(a) Raw data (b) Filtered data

Fig. 12. A 3D point cloud of the laser points from the dust classification data-set, shaded by elevation. On the left, all of the data is displayed and the area that is affected by airborne dust can be seen in the centre of the scene. On the right, the points that have been classified as dust by the model of (1) have been removed. The dust cloud has been filtered, yet the car to the left of the scene and most of the wall and fence line remain.

figures are shaded by elevation, so that the area affected by dust can easily be seen. The filter has succeeded in removing *every* sample that was affected by airborne dust, while leaving the ground and other objects in the scene. The object to the centre left of both figures is a stationary car, and the linear objects are fence lines and building walls. Some points corresponding to the wall have been classified as dust due to the geometry of the sensors. When a vertical wall is viewed in close proximity by the laser and radar, the imperfect collocation causes an increased range discrepancy. If the laser-radar dust model is considered more generally as a redundancy based failure model, then it has successfully captured the geometric conditions under which the assumption of perfect sensor collocation is invalid. These figures show that the model that was learnt from the training subset has successfully captured the classification boundary, and is generally applicable to the wider data-set. Although the model is very successful at classifying airborne dust in *this* test environment, further testing in different environments would be required before stating that this particular model captures *every* aspect relating to dust discrimination. Other types of failures such as geometric alignment errors should also be considered specifically during training and evaluation, to determine the appropriateness of a general redundancy based failure model. This work provides a clear indication of the power of multi-modal sensor fusion for detecting nuances in the environment, which can be difficult, if not impossible, with a single type of sensor in isolation.

V. CONCLUSION AND FUTURE WORK

The contributions of this work may be summarized as follows. Large, accurately calibrated and synchronised, data-sets, have been gathered in controlled environmental conditions (including the presence of dust, smoke and rain) by a UGV equipped with various types of sensors. These data-sets have been made available to the public. Based on these data, an analysis of the effects of the considered challenging conditions on perception has been proposed, to identify common perceptual failures. Methods for mitigating these failures to promote perceptual integrity in adverse environmental conditions have then been presented and discussed. In particular, an information theoretic approach to improve the perception model for a camera has been proposed, as well as a multi-modal sensor fusion algorithm that uses a laser and a radar to detect when airborne dust particles are impeding the use of the laser.

Future work will make further use of multi-modal redundancy. For example, mutual information redundancy given by a thermal infra-red camera and a colour video camera is expected to provide a substantial contribution to the mitigation of vision failures [14] (e.g. IR camera images tend to be less affected by dust or smoke than colour camera images). Furthermore, although two-sensor redundancy algorithms such as the laser-radar model in section IV-B are able to detect failures of one or both sensor(s), they cannot identify *which* sensor failed. This may require additionnal redundancy, e.g. using a third distinct sensor.

VI. ACKNOWLEDGMENTS

This work is supported by the US Air Force Research Laboratory (Robotics Research Group) and the ARC Centre of Excellence programme, funded by the Australian Research Council (ARC) and the New South Wales State Government.

REFERENCES

[1] C.M. Bishop. *Pattern Recognition and Machine Learning*. Springer, August 2006.

[2] R. Dearden and D. Clancy. Particle filters for real-time fault detection in planetary rovers. In *Proceedings of the 12th International Workshop on Principles of Diagnosis*, 2002.

[3] C.S. Dima, N. Vandapel, and M. Hebert. Classifier fusion for outdoor obstacle detection. In *ICRA*, 2004.

[4] A. Kelly et al. Toward reliable off road autonomous vehicles operating in challenging environments. *International Journal of Robotics Research*, 25:449–483, May/June 2006.

[5] A. Monteriu et al. Model-based sensor fault detection and isolation system for unmanned ground vehicles (parts I & II). In *ICRA*, 2007.

[6] C. Urmson et al. Autonomous driving in urban environments: Boss and the urban challenge. *Journal of Field Robotics*, 25(8):425–466, 2008.

[7] J. Leonard et al. A perception driven autonomous urban vehicle. *Journal of Field Robotics*, 25(10):727–774, September 2008.

[8] Y.-L. Chen et al. Terramax: Team oshkosh urban robot. *Journal of Field Robotics*, 25(10):841–860, 2008.

[9] I. Halatci, C.A. Brooks, and K. Iagnemma. Terrain classification and classifier fusion for planetary exploration rovers. In *IEEE Aerospace Conference Proceedings*, 2007.

[10] D.J.C. MacKay. *Information Theory, Inference, and Learning Algorithms*. Cambridge University Press, 2003.

[11] T. Mikaelian, B.C. Williams, and M. Sachenbacher. Autonomous diagnosis based on software-extended behavior models. In *Proceedings of the 8th International Symposium on Artificial Intelligence, Robotics and Automation in Space - iSAIRAS*, 2005.

[12] K. Murphy. Bayes Net Toolbox for Matlab. bnt.sourceforge.net, 1997–2002.

[13] T. Peynot, S. Terho, and S. Scheding. Sensor data integrity: Multi-sensor perception for unmanned ground vehicles. Technical report, Australian Centre for Field Robotics (ACFR), The University of Sydney, http://www-personal.acfr.usyd.edu.au/tpeynot/SDI/, 2009.

[14] J. Underwood. *Reliable and Safe Autonomy for Ground Vehicles in Unstructured Environments*. PhD thesis, The University of Sydney, 2009.

[15] V. Verma, G. Gordon, R. Simmons, and S. Thrun. Real-time fault diagnosis. In *IEEE Robotics and Automation Magazine, June 2004*.