Robust Curb and Ramp Detection for Safe Parking Using the Canesta TOF Camera

Orazio Gallo[†], Roberto Manduchi[†], and Abbas Rafii[‡]

[†]{orazio, manduchi}@soe.ucsc.edu, University of California, Santa Cruz.
[‡]arafii@canesta.com, Canesta Inc.

Abstract

Range sensors for assisted backup and parking have potential for saving human lives and for facilitating parking in challenging situations. However, important features such as curbs and ramps are difficult to detect using ultrasonic or microwave sensors. TOF imaging range sensors may be used successfully for this purpose. In this paper we present a study concerning the use of the Canesta TOF camera for recognition of curbs and ramps. Our approach is based on the detection of individual planar patches using CC-RANSAC, a modified version of the classic RANSAC robust regression algorithm. Whereas RANSAC uses the whole set of inliers to evaluate the fitness of a candidate plane, CC-RANSAC only considers the largest connected components of inliers. We provide experimental evidence that CC-RANSAC provides a more accurate estimation of the dominant plane than RANSAC with a smaller number of iterations.

1. Introduction

Sensor systems for assisted backup and parking are the subject of increasing interest by the automotive industry [12, 26, 2, 14]. It has been estimated that about 120 persons die annually due to a vehicle backing up, with most of the victims between 1 and 4 years of age [3, 1]. The most common case is that of a light truck or SUV going in reverse [25]. Arguably, the main reason for it is that minivans, pickups and SUV create large blind spots, possibly extending up to more than 20 meters behind the vehicle [27]. Sensor and display systems that can assist the driver when visibility is impaired have great potential for reducing the risk of these insidious types of collision. Moreover, the same systems may help the driver during parallel parking, which typically involves maneuvering the vehicle in revers while staying clear of obstacles yet very close to and possibly aligned with the curb. Another challenging situation where driver assistance could be useful concerns backing up onto a narrow driveway, possibly containing a ramp.

A number of devices for assisted backup and parking have been commercialized over the past few years, either as OEM (as part of a vehicle) or aftermarket (installable on an existing vehicle). Camera-based systems display the scene behind the vehicle, which is presented on a screen on the dashboard. These devices allow the driver to see portion of the scene (and potential hazards) that would be otherwise hidden in the blind spot. However, it does not have an active warning system, and therefore assumes that the driver is constantly looking at the screen during backup. Sensor systems based on ultrasounds or microwave can detect obstacles at a certain height with respect to the car, and thus help the driver avoid collisions while backing up. These sensors can also be used as part of an automated car parking system (as in the new Lexus LS 600 and LS 460 models). However, as noted in [12], the distance at which these sensors can detect small object is often not large enough to allow the driver to stop the vehicle in time.

Our work is motivated by the consideration that imaging sensors have much better potential than traditional backup sensors to recognize obstacles and other important geometric features in the scene. The ability to reason about the scene structure (identifying, for example, planar patches and discontinuities) provides a richer environment awareness than could be obtained by isolated point measurements as with ultrasound or microwave sensors. Even when a full array of such sensors is deployed (which is usually the case due to their narrow field of view), the whole data collected can only be used to infer the presence of a sizable object at a certain distance. In contrast, imaging sensors may enable a much richer classification of the range data, for example identifying different objects at different distances, or detecting a curb line. Additionally, if the range data can be registered to the video taken by a camera, potential hazards could be shown superimposed on a screen.

A system for real-time detection and tracking of moving objects using the Canesta Time-Of-Flight (TOF) cam-

era was presented in [14]. This system is able to detect a child-sized object at a maximum distance of 5.5 meters. Objects at zero rearward distance (right against the bumper) can also be detected, an important feature that is not available in several commercial backup systems [12]. The first step in the object detection algorithm of [14] is the robust estimation of the *ground plane*. Measured surface points at a certain distance above the ground plane, or whose surface normal is significantly different from the vertical, are identified. These points are then grouped into clusters, whose location can be tracked in time.

In this paper we tackle the problem of robust identification and localization of structures such as curbs and ramps (like those of Fig. 1) using data from a TOF camera. The importance of this task is twofold. As mentioned earlier, curbs and ramps are features that need to be considered while parking. Additionally, they represent insidious geometric structures that may impair hazard detection based on the algorithm of [14]. Due to the presence of multiple planar structures at close vicinity and orientation, detection of the ground plane using classical methods (e.g., RANSAC) may fail. An example is shown in Fig. 2 (c): rather than selecting one of the three possible planar patches forming the curb, the algorithm chose a plane intersecting all three. This type of error, which is by no means unusual [29], may impair height measurements of the objects in the scene (since height, as discussed earlier, is measured with reference to the ground plane). This paper presents an improved algorithm for plane fitting, dubbed CC-RANSAC, shown to be more reliable than RANSAC in these situations. Additionally, CC-RANSAC can produce good results with fewer iterations than RANSAC.

This contribution is organized as follows. The TOF imaging sensor used for our work is first described in Sec. 2. Algorithms for range processing are discussed in Sec. 3, where we introduce our novel CC-RANSAC algorithm. The superior performance of CC-RANSAC with respect to classic RANSAC is shown by way of a detailed case study. Sec. 4 presents some experimental results on real data, and Sec. 5 has the conclusions.

2. TOF imaging sensors

Time of flight range sensors are different from other imaging depth sensors in various ways, which make them more suitable for curb or ramp detection. First, unlike structured light methods that require projecting a detectable known pattern on the surface, TOF cameras can work both by day and night, regardless of the lighting conditions in the scene. Second, unlike stereo, TOF ranging is texture independent. Third, unlike either stereo or structured light methods, TOF systems do not require a baseline to recover the depth and do not require the extra image processing for solving the correspondence or pattern matching problem.

(a) (b)

Figure 1. An example of a curb (a) and of a ramp (b). These types of features must be identified for safe parking.

Parallax errors, and therefore occlusion issues, can be completely eliminated in a TOF system by the camera packaging design choices. Fourth, the depth error does not increase quadratically with distance as in stereo and structured light methods.

Canesta's Electronic Perception time-of-flight sensor (EP TOF) [13] module, which is used in this work, measures distances by observing the time delay ΔT between emission and detection of light that travels from a modulated active light source, bounces off a surface in the scene, and returns to the camera. For a light source modulated at frequency f and target at radial distance r, the delay can be expressed as a phase Φ:

$$\Phi = 4\pi r f / c \tag{1}$$

where c is the speed of light. Thus the radial distance can be recovered by measuring the phase through at least two measurements in sequential or parallel order [13]. The light source can be either laser or LED and typically has a wavelength of 780 nm to 870 nm in the infrared (IR) spectrum.

The modulation frequency f determines the unambiguous range of the camera. A lower f produces a longer unambiguous range but at the expense of causing more depth uncertainty (or jitter). As an example, with a modulation frequency of 44 MHz, a radial distance of about 3.4 meters can be detected without range aliasing. Another method to increase the unambiguous range, which is suitable to curb/ramp detection, is to use two measurement frequencies f_1 and f_2 in order to produce two phase delays Φ_1 and Φ_2, respectively. The aggregated phase

$$\Phi_c = |\Phi_1 - \Phi_2| \mod 2\pi \tag{2}$$

yields an unambiguous range corresponding to frequency $f_1 - f_2$. For instance, if $f_1 = 31$ MHz and $f_2 = 44$ MHz , the unambiguous range is of about 11.5 meters.

One important parameter affecting the performance of range processing algorithms is the amount of jitter in the depth measurement. In EP TOF camera the level of measurement uncertainty is relatively independent of distance

and is a function of the active brightness level at the pixel. The following relationship is empirically observed: in the absence of strong ambient light,

$$\text{depth jitter} \approx K \cdot b^{-0.4} \tag{3}$$

where depth jitter (or uncertainty) is defined as the amount of instantaneous 3D point measurement deviation from its time average mean, b is the amount of active light received by the pixel, and K is a normalization constant. Therefore, depth jitter can be kept below an application-dependent value by varying the shutter or adjusting the light source power, which directly affect the value of b.

Of course, an application such as safe parking needs to operate under a wide variety of outdoor lighting conditions. In order to minimize the risk of pixel saturation under strong ambient light, an EP TOF camera uses both optical and electronic methods. The former involves coating the lens or the sensor with a narrow-band filter that blocks a substantial part of visible light spectrum. The latter relies on each pixel flushing out remaining ambient light many times during integration (i.e. while the shutter is open) while maintaining the strength of the active light signal. This allows the camera to operate at conditions close to 100 Klux of ambient light.

3. Algorithms for range processing

3.1. State of the art

A simple categorization of range analysis algorithms may be drawn based on whether local or global descriptors are employed. Local descriptors include local surface normals [18, 24], ridges [10], and discontinuities [32, 34, 4, 5]. Contiguous point sets with similar local descriptors may be clustered in space in order to identify extended regions. For example, chains of points with high curvature may form a curb line, and groups of adjacent points with the same normal may identify a planar patch. Only local analysis of the range data is required for this type of descriptors, which can therefore be computed very quickly. For the same reason, however, local descriptors are susceptible to measurement noise and missing measurements.

"Global" descriptors, on the converse, are parametric representations (typically planar or quadratic) of relatively large surface patches. All measurements in a patch contribute to the estimation of the parameters of the global descriptor. For example, if a set of measurements are known to be part of a plane, then simple linear regression (perhaps using principal component analysis, PCA) can provide the corresponding planar equation.

When the measurements are affected by "outliers" (data points that differ substantially from the standard noise model), robust procedures should be employed [30, 22]. For example, M-estimators find the model parameters that minimize a cumulative "robust" loss function. With respect to the quadratic loss function used for standard linear regression, robust loss functions penalize more those samples that deviate heavily from the model. Possibly the best known robust parametric estimators in Computer Vision are RANSAC [11] and the Hough transform [16]. Both can be seen as particular instances of M-estimators [30]. Another popular robust estimation method is the Least Median of Squares (LMedS) [28] and its variants, which include the Least K–th of Squares (LKS). It can be shown that LMedS and LKS are instances of so-called S–estimators, which are a particular case of M–estimators [9]. Another approach to dealing with outliers is to explicitly model them as uniformly distributed. This assumption is at the basis of the MLESAC algorithm [35].

An important parameter of robust estimators is the "scale", ϵ, at which they operate. Intuitively, those points that are at a distance larger than ϵ from the estimated plane are considered "outliers"; the remaining points are "inliers". Clearly, the scale depends on the variance of the inliers, usually modeled as normally distributed. The choice of scale may critically affect the performance of an estimator. A number of solutions to the scale estimation problem exists, including joint estimation with the model parameters [15], minimum unbiased scale estimation (MUSE [23]), adaptive least K–th order square estimation (ALKS [19]) and modified selective statistical estimation (MSSE [6]). When the variance of the inliers is not constant (heteroscedastic data), then more complex robust algorithms should be used [31].

In general, a given planar patch occupies only a finite portion in the image, with other, competing planar regions present as well. There are three main approaches for the simultaneous segmentation and estimation of planar regions in the same image. The first approach, which we use for the experiments in this paper, is to simply use a robust estimator to extract a "dominant" planar region, by considering all the remaining points (including any other planar regions) as outliers. After finding the planar region and removing the inliers, the operation is repeated on the remaining points, until no more sizable planar structures can be found. This algorithm is simple and intuitive, however, the presence of multiple structures may impair the estimation of individual planar patches, especially if the scale is not estimated correctly. This phenomenon was studied in detail in [29, 9].

The second approach to multiple model estimation is to run a simultaneous, concurrent optimization over all planes visible in the image. This can be obtained using the Expectation-Maximization algorithm (an iterative technique akin to K-means clustering) [20, 36] or the recently developed Generalized PCA algorithm [38]. In this case, each plane is represented explicitly, rather than resorting to the notion of "outlier" with respect to a dominant structure. While intuitively more appealing, this approach requires the

joint estimation of the (unknown) number of planar surface elements in the scene, an operation that often proves challenging [21].

The third family of algorithms is based on region growing [7, 33]. Starting from some "seed" points or regions, homogeneous patches are grown concurrently by adding neighboring points consistent with the model. Regions that have similar models can then be merged together. Both region growing and merging can be performed using robust criteria [8, 17]. Region growing is a simple and fast algorithm, but relies on the selection of good seed points, which may be difficult to obtain, especially when planar patches of interest occupy only a small portion of the image. Note that region growing can also be used as an initial step for subsequent robust parameter estimation [37].

3.2. Regression and CC-RANSAC

Planar regression from a set of 3D point seeks for a plane $\mathcal{P}$ that minimizes some measure of observed "fitness" to the data points. If d_i is the Euclidian distance of the i-th data point to the plane $\mathcal{P}$, different measures of the fitness $o(\mathcal{P})$ can be considered:

$$\begin{array}{ll} o = -\sum_i d_i^2 & \text{(LS)} \\ o = -\text{median}\{d_i^2\} & \text{(LMedS)} \\ o = |I(\mathcal{P})| & \text{(RANSAC)} \end{array} \qquad (4)$$

where $I(\mathcal{P})$ is the set of *inliers* (i.e. data points with $d_i \leq \epsilon$ for a given threshold ϵ) and $|I|$ represents the cardinality of the set I. In the LS case, the plane $\mathcal{P}$ maximizing the fitness o can be found in closed form. In the other two cases, one has to resort to random sampling. In general, when the variance of the noise is known, at least approximately, RANSAC is preferable to LMedS due to its lower computational cost. Both methods are superior to LS in the case of outliers or, as in the scenarios considered here, multiple models. However, as mentioned earlier, even RANSAC (or LMedS) may provide poor results when the scene contains two or more planar patches at short distance from each other [29]. This is not a defect of sampling: rather, the plane maximizing the fitness o may simply not coincide with any of the planes in the scene. This is shown by way of example in Fig. 2 (c), where the "incorrect" plane is the one receiving the highest number of supporting inliers.

In order to mitigate this effect, one may exploit the fact that the inliers corresponding to a given planar patch usually cluster contiguously in space. We propose a simple modification of the RANSAC algorithm, dubbed CC-RANSAC, by defining the following measure of fitness:

$$o = |I_C(\mathcal{P})| \quad \text{(CC-RANSAC)} \qquad (5)$$

where $I_C(\mathcal{P})$ is the largest connected component of inliers, with 8-neighbor topology inherited from the image grid.

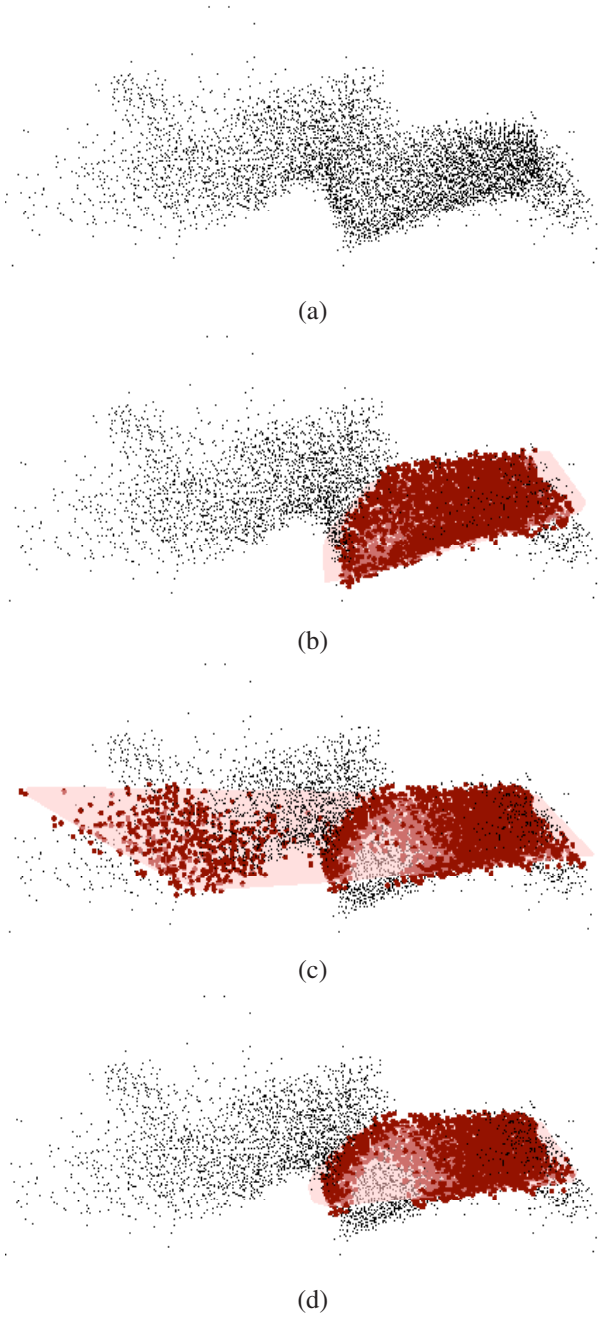

Figure 2. (a): Range data collected in front of a curb. (b): Optimal planar fit to the lower planar patch (inliers with respect to the plane are shown in red). (c) Incorrect planar fit (P) and inliers in $I(\mathcal{P})$. (d) Same plane as in (c) but with inliers in $I_C(\mathcal{P})$.

The idea is that in situations like the one of Fig. 2 (c), different clusters of inliers (corresponding to different patches) are likely to be disconnected. Using I_C for evaluating the fitness of a candidate plane ensures that only the inliers from a single planar patch will contribute to this measure. This is indeed the case for Fig. 2 (d), where the red points represents the inliers belonging to I_C for the same plane as in

Fig. 2 (c).

In the next section we present a simple case study, based on the data of Fig. 2, providing statistical evidence that CC-RANSAC can yield a more accurate fit with a smaller number of iterations than RANSAC.

3.3. A case study

Consider the step shown in Fig. 2 (a). The goal is to find the prominent planar patch, shown in red in Fig. 2 (a). (Although there are two more planar patches visible in the scene, the one shown in red in Fig. 2 (b) has the largest number of data points.) Let $\mathcal{P}_0$ be the plane containing such a patch. This plane was selected manually based on careful observation of the data. Given a candidate plane $\mathcal{P}$, we define its quality $q(\mathcal{P})$ as the number of inliers of $\mathcal{P}$ that are also inliers of $\mathcal{P}_0$, normalized by the number of inliers of $\mathcal{P}_0$:

$$q(\mathcal{P}) = |I(\mathcal{P}) \cap I(\mathcal{P}_0)|/|I(\mathcal{P}_0)| \tag{6}$$

Note that $q(\mathcal{P}) = 0$ when $\mathcal{P}$ is far enough from the planar patch, and $q(\mathcal{P}) = 1$ when $\mathcal{P}$ coincides with $\mathcal{P}_0$. Hence, q seems like an appropriate and simple to compute measure for describing how well a given plane fits the planar patch.

By sampling the space of possible planes, we can estimate the joint probability density function (pdf) of q and o, $f_{q,o}(q, o)$, where for each plane sample $\mathcal{P}$, o is set equal to either $|I(\mathcal{P})|$ or to $|I_C(P)|$ based on the data of Fig. 2 (a). These two joint pdf's, computed using the Parzen window method from a set of 5000 random samples, are shown in Fig. 3.

We show now how $f_{q,o}(q, o)$ can be used to evaluate the expected performance of RANSAC or CC-RANSAC. More precisely, let q_N be the random variable describing the quality of the plane chosen by either algorithm after N iterations (where each iteration corresponds to a randomly selected candidate plane). If $\{o_n\}$ are the measured fitness values of the N candidates planes, then each algorithm chooses the plane $\mathcal{P}$ with $o(\mathcal{P}) = \bar{o}$, where $\bar{o} = \max\{o_i\}$.

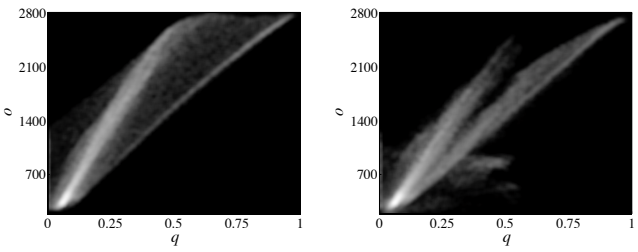

(a) (b)

Figure 3. Graphical representation of the joint densities $f_{q,o}(q, o)$ for the data of Fig. 2 (a) with (a) $o = |I(\mathcal{P})|$ and (b) $o = |I_{CC}(\mathcal{P})|$.

The pdf of q_N can be found as follows:

$$f_{q_N}(q) = \int_{-\infty}^{\infty} f_{q_N|\bar{o}}(q|o) f_{\bar{o}}(o)\, do \tag{7}$$

given that q_N represents the quality of the plane with the highest fitness measure, it is clear that $f_{q_N|\bar{o}}(q|o) = f_{q|o}(q|o)$. The pdf of $\bar{o}$ can be easily derived based on the fact that the samples are drawn independently:

$$f_{\bar{o}}(o) = N f_o(o) F_o^{N-1}(o) \tag{8}$$

where $F_o(o)$ is the cumulative distribution function (cdf) of o:

$$F_o(o) = \int_{-\infty}^{o} f_o(u)\, du \tag{9}$$

All of those quantities are easily computed by numerical integration starting from $f_{q,o}(q, o)$.

Fig. 4 shows the pdf $f_{q_N}(q)$ for RANSAC and CC-RANSAC for different numbers N of iterations. It is interesting to note that RANSAC yields a bimodal distribution: since the two planes of Fig. 2 (a) and (c) both receive good inlier support, the algorithm may choose one or the other with almost the same likelihood (although the incorrect plane receives higher probability mass as the number of iterations increases). RANSAC-CC, instead, yields a unimodal distribution that is peaked around a high quality value, meaning that it almost invariably chooses a plane that is close to the optimal one. Even if only a limited amount of iterations (e.g., 100) is used, the chosen plane is likely to have a good quality value.

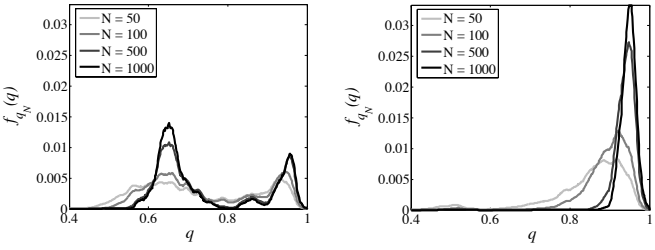

(a) (b)

Figure 4. Plots of the pdf $f_{q_N}(q)$ of the quality of the plane chosen in the case of Fig. 2 using (a) RANSAC and (b) CC-RANSAC with a variable number N of iterations.

4. Experiments

A few experimental results using CC-RANSAC are presented in Fig. 5, in order to highlight the potential of this approach for curb and ramp detection. In this figure, as well as in Fig. 2, inliers are represented with thick points, with color indicating to the plane they are closest to. For each fitting plane, we show the convex hull of its closest inliers, projected onto the planes.

Fig. 5 (a) shows the three best fitting planes to the data of Fig. 2. After the dominant plane has been found, the

corresponding inliers are removed from the data, and the operation is repeated until a maximum number of planes is found, or the highest planar fitness for the remaining point is below a certain threshold.

Fig. 5 (b) shows the result to a similar curb taken from a larger distance. In this case, only two fitting planes were found. Note that the fit is pretty good, in spite of the planar patches being close to each other.

Examples of ramp modeling are shown in Fig. 5 (c) and (d). In particular, Fig. 5 (d) is based on measurements taken of the ramp shown in Fig. 1 (b). The red planar patch corresponds to the descending concrete surface in the ramp; the blue patch represents the asphalt surface at the bottom of the ramp; while the green patch corresponds to the surface covered in soil to the left of the ramp. Note that, contrary to what one would hope, the green and the blue patches do not intersect. This is due to the fact that the measurements supporting the green patch are biased by the presence of a tree stump, visible near the left edge of Fig. 1. Nonetheless, the algorithm is shown to produce very good planar fits to the different elements of the scene, which may enable further reasoning and recognition.

5. Discussion and conclusions

TOF range cameras have great potential for use in backup safety systems. The ability to accurately measure the geometry of the scene enables the sensor system to recognize features that are critical for safe parking. In this paper we concentrate on the detection of curbs and ramps, which are ubiquitous in urban environments. These are particularly challenging features that cannot be reliably detected using conventional ultrasound and microwave sensors.

In order to describe the geometry of a curb or of a ramp, we perform robust fitting to the different visible planar patches. We have shown that the popular RANSAC algorithm may fail in the case of a shallow curb; this result is in agreement with previous work by Stewart [29]. In order to deal with these situations, we propose a new algorithm, CC-RANSAC, that uses only the largest connected component of inliers to evaluate the fitness of a candidate plane. This seemingly minor modification may in fact yield substantially better fits than RANSAC, even at a moderate number of iterations.

A critical analysis of CC-RANSAC brings two main considerations to light. First, the assumption that inliers cluster together into one large connected component, although intuitively correct, needs to be investigated further. It is clear that the size of the largest connected component depends on the distribution of the distances d_i of the data points to the candidate plane as well as on the chosen threshold ϵ. If ϵ is too small, only isolated inlier clusters will form. If ϵ is too large, clusters of inliers corresponding to different planar patches may end up connecting with each other. Secondly, the cost of finding the connected components of inliers contributes significantly to the overall computational weight of the algorithm. In our experience, this additional cost is offset by the fact that fewer iterations are usually required to obtain better results than with RANSAC. However, a more detailed analysis of this trade-off is imperative to better assess the benefits of this algorithm.

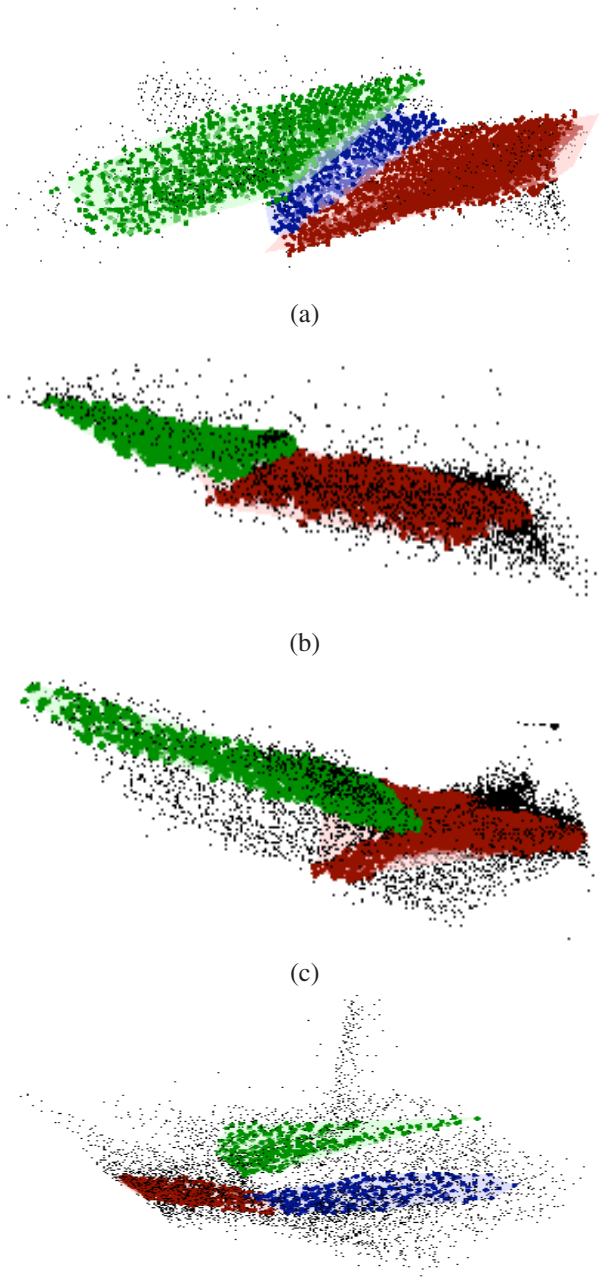

(a) (b) (c) (d)

Figure 5. Some experimental results with curbs ((a) and (b)) and ramps ((c) and (d)).

References

[1] http://www.kidsandcars.org, accessed on 05/02/2008. 1

[2] Backup systems. *Consumer Reports*, 69(10):19–20, October 2004. 1

[3] Data collection study: Deaths and injuries resulting from certain non-traffic and non-crash events. Technical report, U.S. Department of Transportation - National Highway Traffic Safety Administration, 2004. 1

[4] M. Adams and A. Kerstens. Tracking naturally-occurring indoor features in 2-D and 3-D with lidar range amplitude data. *Internation Journal of Robotics Research*, 17(9):907–923, September 1998. 3

[5] M. D. Adams. On-line gradient based surface discontinuity detection for outdoor scanning range sensors. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS '01)*, volume 3, pages 1726–1731, 2001. 3

[6] A. Bab-Hadiashar and D. Suter. Robust range segmentation. In *Proceedings of the Fourteenth International Conference on Pattern Recognition*, volume 2, pages 969–971, 1998. 3

[7] P. J. Besl and R. C. Jain. Segmentation through variable-order surface fitting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10(2):167–192, 1988. 4

[8] K. Boyer, M. Mirza, and G. Ganguly. The robust sequential estimator: A general approach and its application to surface organization in range data. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(10):987–1001, 1994. 4

[9] H. Chen, P. Meer, and D. E. Tyler. Robust regression for data with multiple structures. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2001)*, volume 1, pages 1069–1075, 2001. 3

[10] D. Eberly, R. Gardner, B. Morse, S. Pizer, and C. Scharlach. Ridges for image analysis. *Journal of Mathematical Imaging and Vision*, 4(4):353–373, 1994. 3

[11] M. A. Fischler and R. C. Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM*, 24(6):381–395, 1981. 3

[12] V. Glazduri. An investigation of the potential safety benefits of vehicle backup proximity sensors. In *Proceedings of the International Technical Conference on Enhanced Safety Vehicles*, 2005. 1, 2

[13] S. Gokturk and C. Tomasi. 3D head tracking based on recognition and interpolation using a time-of-flight depth sensor. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR '04)*, 2004. 2

[14] S. Hsu, A. Rafii, and D. Hirvonen. Object detection and tracking using an optical time-of-flight range camera module for vehicle safety and driver assist applications. In *Proceedings of SAE World Congress & Exhibition*, 2007. 1, 2

[15] P. Huber. *Robust Statistics*. Wiley, New York, 1981. 3

[16] J. Illingworth and J. Kittler. A survey of the Hough transform. *Comput. Vision Graph. Image Process.*, 44(1):87–116, 1988. 3

[17] K. Koster and M. Spann. MIR: an approach to robust clustering-application to range image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(5):430–444, 2000. 4

[18] J. F. Lalonde, R. Unnikrishnan, N. Vandapel, and M. Hebert. Scale selection for classification of point-sampled 3D surfaces. In *Fifth International Conference on 3-D Digital Imaging and Modeling (3DIM 2005)*, pages 285–292, 2005. 3

[19] K.-M. Lee, P. Meer, and R.-H. Park. Robust adaptive segmentation of range images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(2):200–205, 1998. 3

[20] Y. Liu, R. Emery, D. Chakrabarti, W. Burgard, and S. Thrun. Using EM to learn 3D models with mobile robots. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2001. 3

[21] G. McLachlan and D. Peel. *Finite Mixture Models*. Wiley-Interscience, October 2000. 4

[22] P. Meer. Robust techniques for computer vision. In G. Medioni and S. B. Kang, editors, *Emerging Topics in Computer Vision*, IMSC Press Multimedia Series, chapter 4. Prentice Hall, 1st edition, July 2004. 3

[23] J. V. Miller and C. V. Stewart. MUSE: robust surface fitting using unbiased scale estimates. In *Proceedings of the IEEE Conference onComputer Vision and Pattern Recognition (CVPR '96)*, pages 300–306, 1996. 3

[24] N. J. Mitra, A. Nguyen, and L. Guibas. Estimating surface normals in noisy point cloud data. *International Journal of Computational Geometry and Applications*, 14(4–5):261–276, 2004. 3

[25] E. P. Nadler, A. P. Courcoulas, M. J. Gardner, and H. R. Ford. Driveway injuries in children: Risk factors, morbidity, and mortality. *Pediatrics*, 108(2):326–8, August 2001. 1

[26] M. Paine and M. Henderson. Devices to assist in reducing the risk to young pedestrians from reversing motor vehicles. Technical Specification 149, Roads and Traffic Authority, New South Wales, Australia, October 2005. 1

[27] M. Paine, A. Macbeth, and M. Henderson. The danger to young pedestrians from reversing motor vehicles. In *Proceedings of the International Technical Conference on Enhanced Safety Vehicles*, 2003. 1

[28] P. Rousseeuw and A. Leroy. *Robust regression and outlier detection*. John Wiley & Sons, Inc., New York, NY, USA, 1987. 3

[29] C. V. Stewart. Bias in robust estimation caused by discontinuities and multiple structures. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(8):818–833, 1997. 2, 3, 4, 6

[30] C. V. Stewart. Robust parameter estimation in computer vision. *SIAM Rev.*, 41(3):513–537, 1999. 3

[31] R. Subbarao and P. Meer. Heteroscedastic projection based m-estimators. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR '05)*, volume 3, pages 38–38, 2005. 3

[32] F. Tang, M. Adams, J. Ibanez-Guzman, and W. S. Wijesoma. Pose invariant, robust feature extraction from data with a modified scale space approach. *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA '04)*, 3:3173–3179, 2004. 3

[33] G. Taubin. Estimation of planar curves, surfaces, and nonplanar space curves defined by implicit equations with applications to edge and range image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(11):1115–1138, 1991. 4

[34] W.-S. Tong, C.-K. Tang, and G. Medioni. First order tensor voting, and application to 3-D scale analysis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR '01)*, volume 1, 2001. 3

[35] P. H. S. Torr and A. Zisserman. MLESAC: a new robust estimator with application to estimating image geometry. *Computer Vision and Image Understanding*, 78(1):138–156, 2000. 3

[36] R. Triebel, W. Burgard, and F. Dellaert. Using hierarchical EM to extract planes from 3D range scans. *Proceedings of the 2005 IEEE International Conference on Robotics and Automation (ICRA 2005)*, pages 4437–4442, 2005. 3

[37] R. Unnikrishnan and M. Hebert. Robust extraction of multiple structures from non-uniformly sampled data. *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003)*, 2:1322–1329 vol.2, 2003. 4

[38] R. Vidal, Y. Ma, and S. Sastry. Generalized principal component analysis (GPCA). *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(12):1945–1959, 2005. 3