Robust Extraction of Multiple Structures from Non-uniformly Sampled Data

Ranjith Unnikrishnan and Martial Hebert

Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213, USA
{ranjith, hebert}@ri.cmu.edu

Abstract— The extraction of multiple coherent structures from point clouds is crucial to the problem of scene modeling. While many statistical methods exist for robust estimation from noisy data, they are inadequate for addressing issues of scale, semi-structured clutter, and large point density variation together with the computational restrictions of autonomous navigation. This paper extends an approach of nonparametric projection-pursuit based regression to compensate for the non-uniform and directional nature of data sampled in outdoor environments. The proposed algorithm is employed for extraction of planar structures and clutter grouping. Results are shown for scene abstraction of 3D range data in large urban scenes. *

I. INTRODUCTION

A basic problem in outdoor scene modeling using data acquired with a laser rangefinder is the reliable identification of multiple structures in large datasets exhibiting a great deal of variability. In this paper we focus on the difficult problem of recovering planar surfaces from large, cluttered 3D point datasets acquired in urban environments. The challenge of delineating multiple instances of a parametric family in the presence of outliers has received considerable interest in the computer vision community. In particular, it has motivated the theory of robust estimation and the adaptation and development of sound statistical methods to vision problems.

A promising approach in this direction using kernel density estimation [1] was introduced for addressing dependence on the estimate of inlier noise scale in a multiple structure scenario. A wide class of robust estimators was expressed as particular cases of M-estimators [2] and revisited in the context of projection pursuit [1] to provide useful non-parametric statistical tools. Under the assumption of higher inlier density, planar structures were recovered from a 3D synthetic dataset contaminated with uniform noise. In this approach, the data is first divided into voxels and each voxel is given a score proportional to local point density. *Samples* are drawn by probabilistic region growing to cover this space, and projection pursuit based M-estimation is performed on each region to obtain

*This work was supported in part by the DARPA MARS2020 program under Grant NBCH1020014

Fig. 1. Aerial view of points sampled by a ground-based laser rangefinder

a robust estimate of plane parameters best fitting the points contained in that region. In a final data fusion step, the parameter estimates for each region are then clustered using a generalization of mean-shift based clustering. The modes found with this clustering procedure having sufficient support are retained as the final plane estimates.

Data collected with active sensors in outdoor settings, however, pose additional challenges due to the nature of sampling and the spatial extent of the observations. As seen in Fig. 1, the data exhibits strong directionality and the appreciable variation in spatial point density helps to violate the assumption of inlier point density inherent to the approach in [1]. Unlike the data analyzed in [1], a global delineation cannot be made in outdoor range data between inliers and outliers based on the inlier density profile, thus making pre-processing steps like nearest-neighbor based filtering unsuitable.

Our implementation closely follows that in [1] with

crucial compensation for variation in spatial point density. We work with datasets covering a very large spatial extent while at the same time regulating computational cost. A robust weighting scheme is used to evaluate local planarity rather than local point density. The initial region sampling is biased toward regions that are likely to contain planes so that only a number of samples smaller than that required to fully cover the space need be used. The sample parameters are clustered in the input space and recover modes with sufficient support. The detected planes are then clustered into walls, and each wall is validated by hypothesis testing with a robust estimate of covariance along the least dominant principal direction. The remaining points constitute scene clutter (trees, bushes, cars etc.) and are grouped in a final step to form an abstraction of the scene.

Popular approaches toward recovering planar structure from outdoor range data have addressed the problem in either the original input domain ($\mathcal{R}^3$) or in the sampling domain as range images. Work in [12] identifies planar regions from raw 3D points to build photorealistic models of buildings. Their approach groups points into neighboring patches, computes the minimum eigenvector for each patch as the normal to the best fitting local plane, and then groups the patches as parts of the same planar surface if they have similar orientations and are close in 3D space. The limitation of this approach lies in its dependence on high point density and low clutter. Recent work [4] in a more challenging scenario segmented buildings from foreground clutter for building textured meshes of building facades. Their approach relied on data being acquired in relative proximity of the surfaces of interest to separate clutter points from points on buildings based on a histogram of distance along-the-ground. Our work attempts to tackle the more general scenario where buildings are observed at varied distances, the manner of data acquisition is not controlled, and a global prior model of noise and clutter distribution is not defensible due to the degree of real-world variability.

Work in [8] uses a variant of the Expectation Maximization (EM) algorithm to fit planar surfaces to indoor laser range data. The biggest drawback to this class of approaches is the prior requirement of the number of model components. In [8], the number of candidate planes is increased incrementally so that a sufficient number of points are explained by surfaces with sufficient support. A post-processing step fuses together planes separated by an angle or distance smaller than a threshold. In statistics, the model selection problem is commonly addressed by scoring functions such as Mallows's C_p statistic and the related AIC (Akaike Information Criterion) which attempt to trade off the fit of the data to model complexity. Analytical extensions to AIC for asymptotic consistency as well as for the more realistic noise model of an underlying t -distribution have been investigated with M-estimators

for surface organization in synthetic range images [9]. However, such scoring criteria are typically biased toward explaining a larger fraction of data points, even though the models themselves may be of low complexity. Hence the approaches mentioned earlier are limited in practice to controlled or low-clutter environments.

We focus attention to the specific problem of recovering multiple planar structures from data acquired using a laser rangefinder. Section II outlines the theory of projection pursuit based regression which we use as the basis of our approach. Section III justifies the modifications made to the basic form of projection-pursuit based regression for handling non-uniformly sampled data. Finally in Section IV, we conclude with a discussion of the potential as well as shortcomings of the proposed method, along with the open issues raised. Results on data acquired in real outdoor scenarios are presented throughout the sections whenever appropriate.

II. BACKGROUND

In this section we summarize the formulation of robust regression introduced in [1], which we employ for recovery of planar structures from point clouds. Let n measurements $y_i \in \mathcal{R}^p$ of dimension p (equal to 3 in our application) be made using a laser rangefinder. Under the linear errors-in-variables (EIV) model, the measurements y_i are corrupted independently by zero mean noise as $y_i = y_{io} + \epsilon_i$, where the subscript o denotes the unknown true value of the measured point. Let the true values obey the linear constraint given by:

$$y_{io}^\top \theta - \alpha = 0 \quad i = 1, \dots, n \quad \|\theta\| = 1 \quad (1)$$

where the normalization constraint is imposed to ensure uniqueness of the parameter values. If the noise is modeled as Gaussian, identically distributed and independent in each variable, the optimal estimator can be proven to be the orthogonal total least squares estimator which minimizes the sum of squared distances between each point y_i and its projection $\hat{y}_i$ on the hyperplane defined by $[\alpha, \theta]$. The corresponding robust M-estimator is given by:

$$[\hat{\theta}, \hat{\alpha}] = \arg \min_{\theta, \alpha} \left[\frac{1}{n} \sum_{i=1}^n \rho \left(\frac{|y_i^\top \theta - \alpha|}{s} \right) \right] \quad (2)$$

where s is a scale parameter, and $\rho(u)$ is a non-negative, symmetric loss function, non-decreasing with $|u|$, scaled to the range $[0, 1]$ and having a unique minimum at $\rho(0) = 0$.

The M-estimator (2) can be rewritten as

$$[\hat{\theta}, \hat{\alpha}] = \arg \max_{\theta, \alpha} \left[\frac{1}{n} \sum_{i=1}^n \kappa \left(\frac{|y_i^\top \theta - \alpha|}{s} \right) \right] \quad (3)$$

where $\kappa(u) \equiv c_p(1 - \rho(u))$ is defined as the M-kernel function in [1] and c_p is the constant of proportionality ensuring that $\kappa(u)$ integrates to 1.

Fig. 2. **Using voxel density for weighting:** Dataset of Fig. 1 with each point colored by weight in proportion to the number of points in its voxel. Colors range from blue (minimum) to dark red (maximum), as per the color-bar on the right.

Consider now the projection of the data points y_i onto a direction in $\mathbb{R}^p$ specified by θ . Let the projected points be denoted $x_i \in \mathbb{R}$ where $x_i = y_i^\top \theta$. The sequence of points x_i can be regarded as being samples from an unknown density $f_\theta(x)$ where the subscript indicates dependence on projection direction. A non-parametric estimator of this density is given by the kernel density estimator [17] which has the general form:

$$\hat{f}(x) = \frac{1}{nh_\theta} \sum_{i=1}^n \kappa\left(\frac{y_i^\top \theta - x}{h_\theta}\right) \quad (4)$$

where $\kappa(u)$ is the kernel function scaled to the bandwidth h_θ .

Equations (2) and (4) have similar form. Replacing the scale parameter s with the kernel bandwidth h_θ , the optimization in M-estimation can be rewritten as

$$[\hat{\theta}, \hat{\alpha}] = \arg \max_{\theta} \left[h_\theta \max_x \hat{f}_\theta(x) \right] \quad (5)$$

As can be seen from the expression, the inner maximization process is that of finding the mode in the density profile of the projected points.

$$\hat{\alpha} = \arg \max_x \hat{f}_{\hat{\theta}}(x)$$

The geometric interpretation of (5) is that of finding the direction θ normal to the candidate hyperplane for which the 1-dimensional projection of the points y_i will possess the maximum mode density. The value of the projected point corresponding to the mode of the density profile is equal to the estimate of plane intercept α . This paradigm

Fig. 3. **Using local planarity for weighting:** Dataset of Fig. 1 with points in each voxel colored by weight in inverse proportion to the minimum eigenvalue of the local robust covariance matrix. Voxels deemed to contain points on the ground have been given minimum weight. Colors range as in Fig. 2

of recovering structure in multidimensional data by looking for trends in lower dimensional (almost always one-dimensional) projections of the data is called projection pursuit [6]. The equivalence given by (5) indicates an approach for robust M-estimation as searching for the θ that maximizes the mode density of the projected density profile [19]. Interested readers are referred to [7] for a discussion of projection pursuit and its many favorable properties and applications.

For computing the objective function, the choice of scale in M-estimation is now replaced by a choice of bandwidth in kernel density estimation. It can be shown that computation of the optimal bandwidth value, that minimizes the asymptotic integrated square error (AMISE) between the reconstructed density profile and the true density, requires knowledge of the underlying density to be estimated [17]. There are several rules of thumb for bandwidth selection based on the extent of approximation of the AMISE. In our experiments, we have found the optimal bandwidth based on an approximated upper bound of the AMISE [17] to give slightly smaller bandwidths than that given by “Scott’s rule” [11], but both have worked well for our application in most cases.

This section revisited the formulation connecting projection pursuit, a tool for non-parametric data analysis, with M-estimators popularly used in the computer vision community for robust regression. The next section describes the modifications made to this basic formulation to constitute our proposed algorithm for manipulating real-world data.

Fig. 4. **Fixed neighborhood size for region growing:** The red colored cuboids are bounding boxes for 70 randomly grown samples. Notice the absence of samples in the sparser region towards the top left.

III. MULTIPLE STRUCTURE EXTRACTION

We now describe our implementation based on the framework of projection-based regression described in the previous section. For the robust recovery of multiple planes, point data is analyzed at two levels. At the local level, candidate points are tested on the degree of planarity exhibited within a small neighborhood. At the global level, candidate planar structures are validated by the degree of support received by several such neighborhoods. The steps of the proposed algorithm are outlined below, with the details of each step relegated to individual subsections:

- A. Discretize input space $\mathcal{R}^p$ into a grid of p -dimensional voxels: This step sets an appropriate scale of local analysis to form sets of spatially proximal points that can be initially considered as statistically coherent groups.
- B. Weight each voxel using an appropriate robust weighting function: A robust estimate of local planarity biases initial analysis of point groups towards those with structure.
- C. Draw s samples from the input space by probabilistic region growing in the voxel grid: Spatially consistent voxels are selectively grouped in proportion to the weights computed in step (2).
- D. Compute projection-pursuit based M-estimates of the best fit plane for each of the s samples: This results in the generation of s samples drawn from the distribution of surface normals.
- E. Fuse the s candidate plane parameters by mean-shift clustering into m modes: Planes with appreciable spatial support form significant modes in the parameter density distribution and are retained as strong

Fig. 5. **Varied neighborhood size for region growing:** The red colored cuboids are bounding boxes for 70 randomly grown samples. Sample clouds in the sparser region in the top left can be extracted due to the increased connectivity of voxels there.

candidates.

- F. Delineate the points associated with each of the m modes into inliers and outliers, and cluster the inliers into walls: This forms physically meaningful group from the extracted points for later testing and validation.
- G. Validate the extracted walls using domain knowledge
- H. Cluster the points in the dataset that were not labeled inliers into groups for further analysis.

A. Discretization of input space

To ease the computational burden of dealing with each point separately, the space of data points is discretized into a regular grid of p -dimensional voxels. Points in each voxel are statistically coherent by choice of an appropriate analysis scale. This is done by setting the dimension of each voxel equal to the AMISE optimal bandwidth estimate computed separately in each dimension. The voxel size is set to have a minimum value of $1m$ corresponding to the scale of structured objects typically found in the operational environments of our application, and also to reduce the overall number of voxels to be considered.

B. Voxel weighting

Due to the directional nature of sampling, the assumption of higher inlier density as made in [1] is no longer valid. This is seen in Fig. 2[†], where the computed voxel weights vary with distance along with the expected variation in spatial point density. To reduce the dependence on

[†]The reader is encouraged to consult the electronic version of this paper for clear color figures.

Fig. 6. **Inferred scene layout:** The dataset of Fig. 1 with beige walls indicate bounding boxes of detected planar points. The clutter groups are colored depending on local planarity, varying from deep green (most uniform distribution) through red (strong planar distribution).

local density and emphasize the local structure of the data, we choose a robust estimator of covariance to compute the weights of each voxel. We opted to choose the *Minimum Volume Ellipsoid (MVE)* estimator introduced in [10]. The multivariate location estimator of mean of a set of n points X is defined as:

$$T(X) = \text{center of minimum volume ellipsoid} \\ \text{containing (at least) } \lceil \frac{n}{2} \rceil \text{ points}$$

The corresponding covariance estimate is given by the ellipsoid itself multiplied by a suitable factor to obtain statistical consistency. This estimator can be shown to possess many desirable properties, including affine equivariance and 50% theoretical breakdown point, and is computationally affordable for moderate values of n .

The weight assigned to each voxel is set to be inversely proportional to the minimum eigenvalue of the robust covariance matrix associated with its neighborhood. The size of the neighborhood considered for determining the covariance matrix is a function of the expected point density at the location of the voxel, increasing with distance from the location of the sensor. As seen from the overhead view in Fig. 3, the voxel weights are higher at locations that exhibit local planarity. Since the ground plane is not of interest in our application, it is computationally advantageous to discard points corresponding to the ground in the steps that follow. One reliable preprocessing step for isolating ground points is cone based filtering [16]. In our experiments, we have found it sufficient in most cases to simply mark the lowest bed of voxels within a specified radius from the scanner as containing points corresponding to the ground, and set their weights to a minimum value.

C. Region growing

For assessing local spatial support for the orientation of a candidate plane, it is necessary to group spatially

Fig. 7. **Scene analysis:** An example where a majority of points in the scene constitute clutter (top). The clutter groups are colored depending on local planarity (bottom) as in Fig. 6. The strong planarity of points on the ground has been correctly identified.

related voxels that exhibit local structure. To do this, first a seed voxel is chosen at random having a weight equal at least to the median weight of all non-empty voxels. We define the boundary of a voxel as the list of voxels in its 4-connected neighborhood. A voxel on the boundary of the seed is chosen at random in proportion to its weight relative to other voxels on the boundary, and then added to the seed to form the current region. The boundary of a region is the union of the boundary of each voxel in the region excluding the voxels that constitute the region. The above step of adding voxels on the boundary is repeated until no more voxels can be added, or a maximum limit on the number of points is reached. This limit is set to be proportional to the point density in the neighborhood of the seed voxel, on the premise that points belonging to a valid plane have similar spatial distribution locally. The final region obtained constitutes a *sample*.

To compensate for the expected variation in point density, the radius of 4-connectivity of each voxel is inversely proportional to its local point density. Without this variation in radii (Fig. 4), voxels containing points at sparse regions would be disconnected from neighboring points. Hence samples generated from seed voxels in those locations would contain too few points to be statistically representative.

This process is repeated to obtain m such samples

Fig. 8. **Inferred scene layout:** The beige walls indicate bounding boxes of detected planar points. Groups of non-planar points (clutter) are shown in different colors. Points in voxels marked as containing ground points have been removed for clarity.

(typically set to 100 in our application). Since the weighting function penalizes deviation from local planarity, the proportion of samples biased toward regions that are more likely to contain the true planes is increased. This facilitates the use of a relatively small number of samples even for the extent of the feature space analyzed. Thus, this step in our proposed approach differs from that in [1] in two ways: the structure-biased sampling using a robust weighting function, and the variation in neighborhood size for spatial density compensation.

D. Projection-pursuit based regression

Robust plane parameters are extracted from each region to obtain candidate plane parameters. This is done by performing the algorithm outlined in Section II on each sample cloud. The objective function is evaluated on the projection of the cloud at the direction corresponding to each simplex vertex using a bandwidth equal to the AMISE estimate. The modes are found for a set of 10 points at equally spaced quantiles using mean-shift [3]. The value of density computed at the most significant mode is then used for computing the simplex objective function. The resulting parameter candidates estimated for each sample are passed to the next stage of data fusion.

Note that in our application, the points recorded by the laser range finder belong to one of two broad categories: a group of noisy samples of points lying on the surface of objects in the scene, and the smaller fraction of false points resulting from sensor noise or edge effects associated with active sensing. As the latter group of points does not typically exhibit coherent local structure, this step is effectively a biased sampling of the space of surface normals in the scene.

E. Data fusion

The density distribution underlying the samples drawn from step III-C can be analyzed to isolate candidate

plane parameters that are statistically significant. This is accomplished by mean-shift based clustering of the parameters obtained from regression. The points of the intersection of the plane normals $z = \alpha\theta$, where $z \in \mathcal{R}^p$, for each of the sample plane parameters are clustered using a Euclidean distance metric. The modes of the underlying density profile, corresponding to the cluster centers, are the plane parameters with appreciable spatial support at the global level of analysis [1]. For each candidate plane, the points that are associated with it and lie in the basin of attraction of the projected density profile are declared inliers.

In urban scenarios, it often occurs that two nearly parallel planes corresponding to two different walls are within a critical margin of separation; i.e. their normal intercepts are close enough in feature space that they are associated with the same mode after fusion in intercept space, but are associated with two different modes in the projected density profile (Fig. 8). It is thus necessary to detect multiple significant modes in the projected density profile (Fig. 9) corresponding to each candidate normal in order to check for the presence of more than one plane. For this, a bandwidth of half the AMISE estimate is used for fine delineation of points between the two planes. To avoid choosing spurious modes in this stage, a simple test of mode validation is conducted by ensuring that the candidate mode in the density profile is persistent over the density profile computed over a narrow range of bandwidths.

F. Clustering of inliers into walls

Points delineated as inliers to a candidate plane may be associated with more than one wall, and hence need to be clustered. The variation of spatial density of points within a wall is addressed by using a distance measure equal to the angular separation between points. This effectively projects the data onto a unit sphere and eliminates the effect of density variation in clustering.

Figures 6–8 illustrate the grouping of inliers at varying distances into semantically meaningful walls. The datasets shown in the figures in this paper were collected using a prototype Zoller+Fröhlich (ZF) laser having a full 360° field of view with $\pm 35^\circ$ of azimuth and a nominal range of $25m$. The data was collected in the lasers native resolution of 0.045° , but the data used in our experiments as well as figures were downsampled by a factor of 30 to simulate a 1.35° resolution while retaining the original volume of view. Successful experiments have also been carried out in similar environments using a panning SICK LMS291 laser at $0.5-1^\circ$ resolution having a nominal range of $70m$. The viewing volume considered varied between $50m \times 50m \times 15m$ on the ZF laser platform and $75m \times 40m \times 18m$ with the SICK laser platform.

Fig. 9. **Multiple mode extraction:** The dataset of Fig. 8 has two walls (marked with red in top figure) that are near-parallel but offset by a small distance. As a result, on projecting the points along the direction giving maximum mode density, the density profile of the projection shows two distinct maxima corresponding to the two walls.

G. Validation of extracted walls

Points in voxels on the surface of dense vegetation have a higher weight than points in the interior due to the directional nature of sampling of points on the vegetation surface. When a large enough number of samples contain such points, they can gather enough support to pass through the post-fusion stage masquerading as planar regions. Such occurrences are weeded out through hypothesis testing.

A robust estimate of the covariance of the n points about the candidate wall is tested against the expected variance σ_0^2 as determined by the choice of sensor. The MVE estimator is used to determine the covariance matrix and the variance (s^2) along the least dominant principal direction is computed the plane is validated by one-sided hypothesis testing. The null hypothesis is defined as the observed standard deviation being less than or equal to the expected value. The test statistic can then be defined as:

$$T = \frac{(n-1)s^2}{\sigma_0^2}$$

with the rejection region for a level α test given by

$$T > \chi_{(n-1),\alpha}^2$$

On rejection of the hypothesis, the candidate wall is rejected; else it is retained.

H. Grouping of clutter

Clustering of points deemed as not associated with a plane, requires coping with the variation in spatial density of points. Well known statistical methods in kernel density estimation [3] have adapted to this difficulty by using varying kernel width for mean-shift based clustering. The local scale of the underlying density at each point is estimated as the bandwidth maximizing the magnitude of the normalized mean shift vector at that point. However this can be burdensome for large datasets.

For computational efficiency, we opt to choose a smaller set of points in the feature space and perform mean-shift based clustering on them. The point selection is done by non-uniformly tessellating the space into spheres, and the center of each sphere of tessellation is chosen as a representative point. The radius of each sphere is varied with the expected density of points at the location of the sphere center, i.e. the sphere radii are increased with distance to give rise to a non-uniform tessellation. The modes found after mean-shift analysis on the representative points are then fused, and the resulting points form the final cluster centers. Each point is then labeled with the mode corresponding to the center of its associated sphere. A cluster is discarded if the number of its supporting points is below a threshold.

Fig. 8 illustrates the different clutter groups formed as a result of this step. Figures 6–7 display the cluster groups in colors ranging from red to brown based on the degree of local planarity in each group. Of note is the delineation of ground points and points associated with vegetation.

IV. DISCUSSION

Recovering coherent structures in a principled manner from unorganized range information of outdoor scenarios is a difficult problem due to the high degree of natural variability. Our work addressed this problem from the viewpoint of robust estimation without rigid association with prior noise models and with necessary compensation for the nature of sensor data sampling. We have presented results on robust extraction of planar surfaces in large and highly cluttered environments, having ready application for use as landmarks for localization, feature-based mapping and autonomous navigation.

While the proposed method has sound statistical basis, work is needed in incorporating sensor specific information, such as reliable noise models and the directionality of the measurements [18]. For generation of 3D models, surface completion of occluded regions [4][14] in a manner tolerant of severe variation in spatial density of points is also required. Inference of more detailed semantic information and classification of sparse range data within classes of man-made and naturally occurring objects is a work in progress.

V. REFERENCES

- [1] H. Chen, P. Meer, "Robust computer vision through kernel density estimation", *7th European Conference on Computer Vision*, Copenhagen, May 2002, vol. I, 236-250
- [2] H. Chen, P. Meer, D.E. Tyler: "Robust regression for data with multiple structures", *IEEE Conference on Computer Vision and Pattern Recognition*, Kauai, Hawaii, December 2001, vol. I, 1069-1075
- [3] D. Comaniciu, V. Ramesh, P. Meer: "The variable bandwidth mean shift and data-driven scale selection", *8th International Conference on Computer Vision*, Vancouver, Canada, July 2001, vol. I, 438-445M.
- [4] J. Davis, S.R. Marschner, M. Garr, M. Levoy, "Filling holes in complex surfaces using volumetric diffusion", *International Symposium on 3D Data Processing, Visualization and Transmission*, Italy, June 2002.
- [5] C. Fruh and A. Zakhor, "Data processing algorithms for generating textures 3D building façade meshes from laser scans and camera images", *3D Data Processing, Visualization and Transmission (3DPVT) 2002*.
- [6] P.J. Huber. "Projection pursuit (with discussion)", *Annals of Statistics*, 13:435-525, 1985.
- [7] M.C. Jones and R. Gibson. "What is projection pursuit? (with discussion)", *J. Royal Stat. Soc. A*, 150:1-37, 1987.
- [8] Y. Liu, R. Emery, D. Chakrabarti, W. Burgard, and S. Thrun. "Using EM to learn 3D models with mobile robots", *Proc. of the International Conference on Machine Learning (ICML)*, 2001
- [9] M.J. Mirza and K.L. Boyer, "Performance evaluation of a class of M-estimators for surface parameter estimation in noisy range data", *IEEE Transactions on Robotics and Automation*, Vol. 9, no 1, Feb 1993, pp.75-85.
- [10] P. J. Rousseeuw and A. M. Leroy. "Robust Regression and Outlier Detection", Wiley, 1987.
- [11] D.W. Scott, "Multivariate Density Estimation: Theory, Practice and Visualization", *Wiley Series in Probability and Mathematical Statistics*.
- [12] I. Stamos and P. Allen, "3D Model Construction Using Range and Image Data", *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2000.
- [13] C.V. Stewart, "Robust parameter estimation in computer vision", *SIAM Reviews*, vol. 41, pp. 513-537, 1999.
- [14] F. Stulp, F. Dell'Acqua, R.B. Fisher, "Reconstruction of Surfaces behind Occlusions in Range Images", *Proc. 3rd Int. Conf. on 3-D Digital Imaging and Modeling (3DIM01)*, Montreal, Canada, June 2001.
- [15] P.H.S. Torr and A. Zisserman. "MLESAC: A new robust estimator with application to estimating image geometry", *Computer Vision and Image Understanding*, 78:138-156, 2000.
- [16] N. Vandapel, R.R. Donamukkala, and M. Hebert, "Experimental Results in Using Aerial LADAR Data for Mobile Robot Navigation", *International Conference on Field and Service Robotics*, 2003.
- [17] M.P. Wand and M.C. Jones. "Kernel Smoothing", *Chapman & Hall*, 1995.
- [18] G.A. Watson, "Some robust methods for fitting parametrically defined curves and surfaces to measured data", *Advanced Mathematical and Computational Tools in Metrology IV, Series on Advances in Mathematics for Applied Sciences, Volume 53, World Scientific, Singapore*, pp. 256-272 (2000).
- [19] R.H. Zamar, "Robust estimation in the errors-in-variables model", *Biometrika* (1989), 76, 1, pp.149-60.