Laser-based Perception for Door and Handle Identification

Radu Bogdan Rusu*, Wim Meeussen†, Sachin Chitta†, Michael Beetz*

† Intelligent Autonomous Systems, Department of Computer Science, Technische Universität München, Germany

{rusu, beetz}@cs.tum.edu

‡ Willow Garage, USA

{meeussen, sachinc}@willowgarage.com

Abstract— In this paper, we present a laser-based approach for door and handle identification. The approach builds on a 3D perception pipeline to annotate doors and their handles solely from sensed laser data, without any a priori model learning. In particular, we segment the parts of interest using robust geometric estimators and statistical methods applied on geometric and intensity distribution variations in the scan. We present experimental results on a mobile manipulation platform (PR2) intended for indoor manipulation tasks. We validate the approach by generating trajectories that position the robot end-effector in front of door handles and grasp the handle. The robustness of our approach is demonstrated by real world experiments conducted on a large set of doors.

I. Introduction

An important challenge for autonomous personal robots is to be able to enter a human living environment and function in it, that is, find ways of navigating and interacting with the world in an effective manner. This means the robot should be capable of going anywhere where it can physically fit, where it would be able to find energy sources and recharge its batteries when their capacity is running low, and in general it must be able to do useful things such as cleaning tables. These behaviors need the support of complex perception routines which can recognize power plugs, certain structures and objects in the world, etc. Some of these environment structures, such as fixtures (handles and knobs) on doors and pieces of furniture, are of key importance for the robot's performance. Any robot that will operate indoors must be able to correctly locate these fixtures and open doors to be able to better carry out different tasks. In this paper we focus on one particular aspect of indoor environments, namely the perception of doors and handles and what is needed to open or close them.

Since the robot may have to function in a wide variety of environments under varying lighting conditions, a robust door detection ability is essential for the robot. In this work we present an approach to door detection using laser-based perception. Our approach to detecting doors is part of a bigger effort to enable mobile robots to function effectively indoors. This effort has resulted in two platforms: a hardware platform called PR2 (Personal Robot 2 - see Figure 1) and a software platform called ROS (Robot Operating System). PR2 is an indoor mobile manipulation platform with two arms, an omni-directional base and an extensive sensor suite including laser rangefinders, cameras and tactile sensors. ROS is a software platform that provides a communication framework and open-source implementations of planning, control and sensing algorithms useful for robotics.

Laser based perception has gained wider acceptance in recent years. Laser range finders are better suited for indoor environments than outdoor environments due to the absence of dust clouds, rain or foliage. Recent work on building semantic maps [1], [2] has exploited the inherent structure of most indoor environments to annotate them using laser-based perception. Such effort benefits from and uses the geometric structure that laser-based perception systems can provide. While stereo-based visual systems can also be used indoors, they are not as suitable in environments where the lack of texture makes it difficult to recover the 3D structure of walls and doors. Sonar provides another approach to indoor sensing but is not as accurate as laser-based systems. Small form-factor laser sensors are now available that can be easily mounted on indoor robots. We utilize such a sensor (Hokuyo UTM-30 - see Figure 3) on our robot. The sensor is mounted on an actuated tilt platform and is used to build a 3D point cloud representation of the environment (see Figure 2).
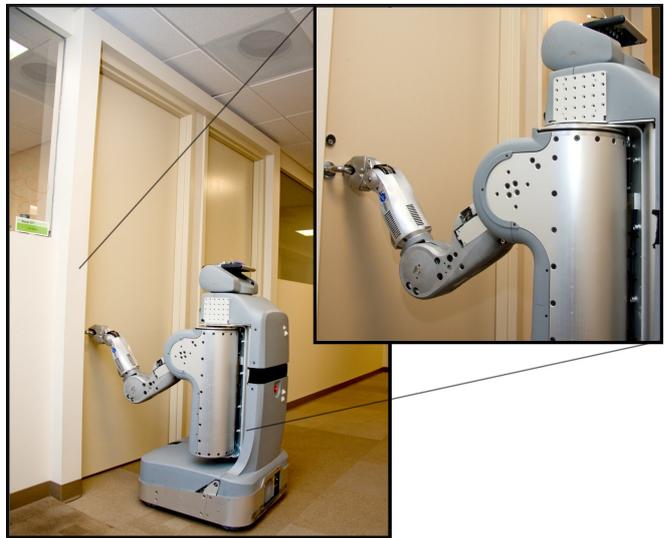

Fig. 1. Snapshot of the PR2 robot used during the experiments.

The space of possible doors and handles is huge. In particular, handles come in a wide variety of shapes and sizes and could be mounted anywhere on a door. This makes the task of searching for doors and handles in a point cloud extremely difficult. In this paper, we restrict our search space

by considering only doors and handles that conform to the American Disability Act (ADA). The PR2 program aims to develop a manipulation platform that is ADA compliant, i.e., the robot can access parts of the environment that a person in a wheelchair should be able to access. ADA compliance places constraints on the geometry and placement of handles, specifying in particular that handles must be above a certain height on the door and should be able to be opened without using a grasping hold. This simplifies the search for handles once a candidate door is found. It also simplifies the process of finding a point on the handle to grab, since the handles tend to be linear with a long lever arm rather than knob-like. ADA compliance also places a restriction on the width of the door since it must be at least wide enough to let a wheelchair through.

Our approach builds on some of these constraints to achieve robust door and handle detection. As input, we make use of 3D point cloud datasets acquired using our tilting laser sensor. The point clouds are sub-sampled and annotated first to find candidate door planes. We then apply several rules based on the constraints arising from ADA rules to prune the search area further and develop a *goodness* score for each candidate door. The candidate with the highest score represents our best choice for a door. We then search within a limited area of the door plane to find handles. Our search is based on multiple criteria including the expected shape of the handle itself and the difference in intensities between the handles and the door plane. Our methods do not need any pre-determined thresholds, but instead automatically cluster and segment the regions of interest using intensity and geometric differences.
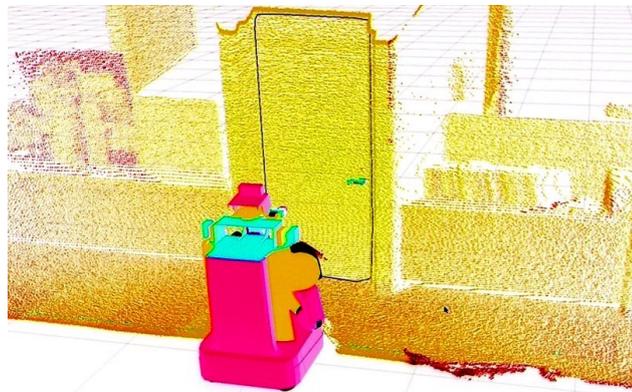

Fig. 2. Door and handle identification example in a 3D point cloud acquired using the tilting Hokuyo laser.

To validate our approach we perform multiple experiments on the PR2 platform using a grasping behavior to confirm the accuracy of the door and handle detection. All software components necessary to perform the tasks discussed in this work are modular so they can be reused individually, and are available Open Source as part of the ROS software project.[1] They can be easily adapted for use with other robots if an appropriate robot description is specified.

[1]http://pr.willowgarage.com/wiki/

The remainder of this paper is organized as follows. Related work is described in Section II. We briefly describe our system architecture in Section III. Sections IV and V present the 3D perception modules for door segmentation and handle identification from the sensed data. Section VI presents the controller used to grasp the door handle. The grasping is carried out to validate the sensing pipeline. We describe experimental results in Section VII, and conclude in Section VIII.

II. RELATED WORK

A 2D implementation for wall and door detection from laser data is presented in [3]. Appearance models are learned using a simple EM algorithm from the geometric data present in 2D laser scans, and an evaluation on 5 doors is presented. The system requires an a priori existing map and heavily relies on precise localization which makes it unsuitable for applications such as ours. An active vision system for door localization and opening is presented in [4], [5]. The door model is extracted using corner based features and edges from images, and is constrained to 2 concentric rectangles parallel to the floor's normal. Then the door handle is identified by training a set of 1500 instances artificially generated for one door handle and extracting Haar-like features. Though the authors give a success rate of 12/15 runs, it's unclear how their method would scale to different doors, such as doors with different frames, same color as the wall, etc, as well as how the handle classification performs on datasets of multiple different handles, rather than a single one. [6] presents an approach for door state identification in images by subsequent filtering and Hough transform search. Unfortunately the authors only present results on a single door, and their system depends on a large number of thresholds (e.g. the door frames must be two or more vertical lines which are separate by more than 20 pixels and at an angle between $85°$ and $95°$, etc).

In [7] a set of behaviors for reliably opening a variety of doors are presented, with over 87.5% success rate. However, the door and handle identification component of the proposed system relies on a user shining a laser pointer at the door handle, and thus the door identification is not autonomously performed. The problem of detecting a small subset of circular handles using Hough Transforms in camera images is presented in [8]. A decision tree like classifier is learned over the RGB space with the assumption that the actual handles are distinct in color from the door, and the results are combined with the Hough circles to obtain the final handle candidates. The results presented are encouraging, but the system implementation depends on many thresholds and is prone to failure for different lighting conditions. A vision-based learning method for handle identification using Haar features and decision trees is presented in [9]. The training dataset consists of approximatively 300 positive and 6000 negative samples of door handles, and the classification results are around 94%. Once a handle is identified, its 3D position is inferred using a logistic classifier for door axis using the information from a stereo camera. However, the

approach assumes that the location of the door is known from a map. It also does not account for cases where the door may already be partially or fully open and only a side-view of the handle may be visible.

Handles for kitchen appliances are modeled in [10] and [1] using 2D line segments extracted either from edges [10] or from 3D point clouds [1]. The environments presented however exhibit mostly classic appliances with a white frontal face, and easily recognizable handles. Unfortunately, both initiatives refrain in presenting the identification success rates. Another initiative for door detection is presented in [11], though the system estimates only the 2D hinge positions on a grid map, and assumes that the handle position is fixed and already known. [12] describes a template-matching technique for handle identification using camera images, but does not present the actual success rate, though the authors acknowledge that the method is susceptible to problems in challenging lighting conditions.

With almost no exception, none of the above research initiatives models the 3D characteristics of the door or the environment itself, thus rendering a reliable 3D map of the world useful for collision detection for arm planning. In addition, the handle identification is based on 2D image features which require stable and usually bright light sources and are thus sensitive to illumination changes (for example, operating in the dark has not yet been shown to be possible with any of the above approaches). Finally, most of the above approaches work only for the cases where the doors are closed and cannot handle the case where the door may be partially open. In contrast, our approach can detect doors and handles even when doors are partially open.

III. System Architecture

Experiments to validate our approach were carried out on a prototype of the PR2 mobile manipulation platform. The PR2 comprises an omni-directional wheeled base, a telescoping spine, and two force-controlled 7-DOF arms equipped with parallel-jaw grippers. The arms use a spring-based counter-balance system for gravity compensation. An EtherCAT link connects all encoders and actuators to a computer running a hard real-time control loop at 1 kHz. Three other dual-core 2.6 GHz computers, connected by a gigabit LAN, provide additional computational power.

Fig. 3. Hokuyo laser mounted on a tilting platform.

The PR2 is equipped with a Videre stereo camera on a pan-tilt stage, and two Hokuyo UTM-30 laser range finders, one mounted on the mobile base, and one on a tilt stage (see Figure 3). This second Hokuyo is tilted up and down continuously, providing a 3D view of the area in front of the robot. The resultant point cloud, which contains both position and intensity information, is the main input to our perception system. The architecture of our system is presented in Figure 4. Each component or subsystem in the ROS architecture is called a *n*ode. The topic of the paper is focused on 3D perception, and therefore we will address its two components (the Door Detector node and the Handle Detector node) below. The remaining subsystems fall outside of the scope of the paper.
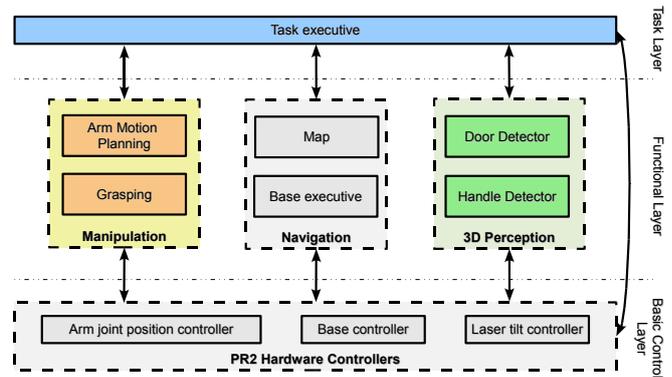

Fig. 4. The architecture of our system using ROS. Boxes with continuous lines define nodes. Boxes with dashed lines delimit groups of nodes with related functionality. Communication between nodes is represented by arrows.

IV. Door Detection

The Door Detection node operates directly on the acquired point cloud data $\mathcal{P}$ from the tilting laser sensor, and makes no use of camera images or any other information sources to identify the door plane and 3D bounds. The main motivation is given by our system's requirements to operate in a variety of situations, and account for any changes in the room ambient light throughout long periods of time. In addition, our perception system natively generates rich 3D annotations for the world map, using the geometric information captured in the point cloud, and thus is an invaluable source of information for the motion planner and grasping system, providing realtime map updates for collision detection and the desired 3D poses for the task goals.

The output of the node is represented by a list of doors described through a set of attributes including the door planes, their bounds, as well as their location with respect to a given world coordinate system. The doors are scored and returned with respect to a given fitness function $\mathcal{F}$, which includes parameters such as the number of point inliers supporting the door model or the current distance of the model to the robot. In practice however, we are mostly interested in the current closest door, and will use the first best candidate in the list for the purpose of the experiments performed in this paper. Any of the subsequent operations applied to a door candidate can uniformly be applied to the rest of the candidates, without any loss of generality.

Because we expect the perception system to run online, we restrict the computational time requirements for the door

detection to the maximum allowable time difference between the acquisition of two subsequent point cloud datasets $\mathcal{P}_i$ and $\mathcal{P}_{i+1}$. This constitutes the only real time constraint of our application. Algorithm 1 presents the main computational steps that are used by the Door Detection node.

The first step takes the set of input points $\mathcal{P}$, and creates a downsampled representation of them $\mathcal{P}_d$ using a fast octree structure. For each point $\boldsymbol{p}_i \in \mathcal{P}_d$, we estimate a surface normal $\boldsymbol{n}_i$ by fitting a least-squares plane to the neighborhood $\mathcal{P}_i^k$ of $\boldsymbol{p}_i$. Due to the fact that we are working with a sparser spatial dataset after downsampling, we select $\mathcal{P}_i^k$ as the set of point neighbors from the original cloud $\mathcal{P}$. Given a viewpoint $\boldsymbol{v}$ from where the point cloud dataset was originally acquired, all resultant point normals $\boldsymbol{n}_i$ must satisfy the equation:

$$\boldsymbol{n}_i \cdot (\boldsymbol{v} - \boldsymbol{p}_i) > 0 \qquad (1)$$

Then, we transform $\mathcal{P}_d$ into a coordinate system defined at the base of the robot, with $\mathcal{Z}$ pointing upwards, and select a subset of all the points $\mathcal{P}_{\mathcal{Z}}$ having their estimated surface normals $\boldsymbol{n}_i$ approximatively perpendicular to the world $\mathcal{Z}$-axis, i.e. $\boldsymbol{n}_i \cdot \mathcal{Z} \approx 0$. The resultant set $\mathcal{P}_{\mathcal{Z}}$ is split into a set of Euclidean clusters using a region growing approach: $\mathcal{P}_{\mathcal{Z}} = \{\mathsf{c}_1 \cdots \mathsf{c}_n\}$. Each cluster c_i represents a potential door candidate in our framework, and we proceed at fitting a plane model to it using a RMSAC (Randomized M-Estimator Sample Consensus) robust estimator [13]. Due to the geometric independence of the clusters, the search is parallelized by dynamically selecting a set of N clusters for concurrent processing, based on the number of CPUs present in the system and their current availability. The inliers of the models found are then projected onto their respective planes and a set of bounding 2D polygonal structures is estimated.

Algorithm 1 Main computational steps for Door Detection

$\mathsf{b}_{\mathrm{ADA}}, \mathsf{b}_r, \boldsymbol{v}$ *// ADA door requirements, robot bounds, $\mathcal{P}$ acquisition viewpoint*
$\mathcal{P} = \{\boldsymbol{p}_1 \cdots \boldsymbol{p}_n\}$ *// set of 3D points*
$\mathcal{P}_d \leftarrow \mathsf{F}(\mathcal{P})$ *// create a downsampled representation $\mathcal{P}_d$*
$\mathcal{P}_d \leftarrow \{\boldsymbol{n}_1 \cdots \boldsymbol{n}_n \mid \boldsymbol{n}_i \cdot (\boldsymbol{v} - \boldsymbol{n}_i) > 0\}$ *// estimate normals at $\boldsymbol{p}_i \in \mathcal{P}_d$*
$\mathcal{P}_{\mathcal{Z}} = \{\boldsymbol{p}_i \mid \boldsymbol{n}_i \cdot \mathcal{Z} \approx 0\}$ *// select set $\mathcal{P}_{\mathcal{Z}}$ with normals perpendicular to $\mathcal{Z}$*
estimate $(\mathcal{C} = \{\mathcal{P}_{\mathcal{Z}}^1 \cdots \mathcal{P}_{\mathcal{Z}}^n\}, \mathcal{P}_{\mathcal{Z}}^i \subset \mathcal{P}_{\mathcal{Z}})$ *// break $\mathcal{P}_{\mathcal{Z}}$ into Euclidean clusters*
for all $\mathsf{c}_i = \mathcal{P}_z^i \in \mathcal{C}$
    // find the best plane fit using sample consensus
    estimate $(\{\boldsymbol{a}, \boldsymbol{b}, \boldsymbol{c}, \boldsymbol{d}\}, \boldsymbol{a} \cdot \boldsymbol{p}_i^x + \boldsymbol{b} \cdot \boldsymbol{p}_i^y + \boldsymbol{c} \cdot \boldsymbol{p}_i^z + \boldsymbol{d} = 0, \boldsymbol{p}_i \in \mathsf{c}_i)$
    estimate $(\mathcal{A} = a_1 \cdots a_n)$ *// estimate geometric attributes for the planar area*
    if $\mathsf{F}(\mathsf{c}_i, \mathcal{A}, \mathsf{b}_{\mathrm{ADA}}, \mathsf{b}_r)$ *// does c_i respect the given constraints?*
        $\mathcal{D} \leftarrow \mathsf{c}_i$ *// add to $\mathcal{D}$, the list of good candidates*

Since an estimated door plane normal is perpendicular to the Z axis, we perform a final robust estimation step to fit the two best vertical lines in each aforementioned 2D polygon, and thus estimate the two major door edges. From these, we estimate a set of geometric attributes $\mathcal{A}$ such as: width, height, minimum and maximum values along an axis, area, number of supporting inliers, and finally the height of the two door edges. These attributes constitute the input to our geometric tests used to select and weight the best door candidates $\mathcal{D}$ for a given dataset $\mathcal{P}$.

The tests refer mainly to the ADA requirements with respect to door dimensions. To retain a degree of flexibility in our solution space, we made the Door Detection node

parameterizable and keep door candidates only if they have a height larger than $\mathsf{h}_{\min}$ (the height of our robot) for example. The width of a candidate has to respect the minimally imposed ADA width for wheelchairs, and using a heuristic assumption we imposed a maximum width of 1.4m. In addition, we impose that each of the two door edges on the side of the door has to have a length of at least $\mathsf{h}_{\min}$. The task executive can change the value of any of these parameters online. Figure 5 presents two examples for door detection in cluttered environments using the previously described algorithm.
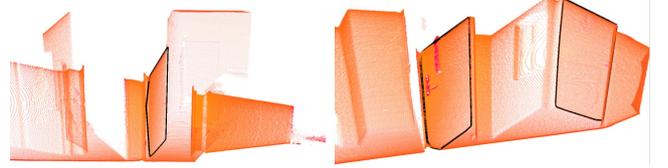

Fig. 5. Two examples of estimated door candidates in point cloud scenes.

V. HANDLE DETECTION

The Handle Detector node is invoked for restricted portions of 3D space, usually in the proximity where a door candidate has already been detected. To retain the same generality level present in the Door Detector node, our handle identification method operates and extracts the handle from the same point cloud data $\mathcal{P}$, without making use of additional data sources such as camera images. However, due to the nature of handles in general, such as their extremely thin geometrical structure and the materials they are made of, the sampled data representing them is extremely sparse and noisy. Taken together, these issues bring additional complexity to the handle identification problem, and lead to situations where it is impossible to perform a pure geometric segmentation.

To solve this, our algorithms combine the sparse geometrical structure with the additional information provided by the intensity (or better said, surface reflectivity) data acquired and present in the laser scan. As shown in the following, this combination increases the robustness of the handle segmentation, and provides solutions which geometry alone would not be able to solve. Figure 6 presents both the intensity and geometry variations for a handle selected from a dataset $\mathcal{P}$. The main computational steps used by the Handle Detection node are presented in Algorithm 2.
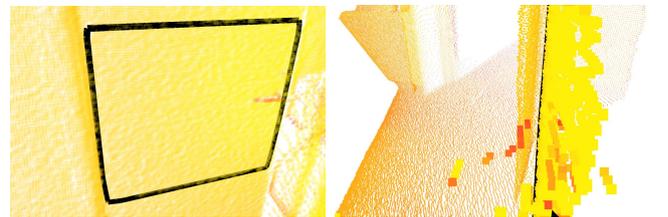

Fig. 6. Intensity and geometry variations for a selected handle. Left: view from the front, right: view from the top.

Our method starts by segmenting a subset of points $\mathcal{P}_s$ from the entire point cloud $\mathcal{P}$, which could potentially contain a handle. The segmentation is performed using the

information provided by the task executive, which states which doors have been detected and offers their geometrical parameters to the Handle Detector node. In particular, for a given door model $d \in \mathcal{D}$ with $\mathcal{D}$ being the set of all detected doors for a dataset $\mathcal{P}$, we select $\mathcal{P}_s = \{\boldsymbol{p}_i \mid \boldsymbol{p}_i \in \mathcal{P}, \boldsymbol{p}_i \subset \mathsf{V}\}$, where V represents the volume of a 3D polygon created from the bounding 2D polygon of d translated along the plane normal with $\pm h_\mathsf{d}$. The parameter h_d is given by the ADA requirements as the maximum distance from a door plane where a handle could be located. A simplification of the above is to get all points $\boldsymbol{p}_i$ whose distance from the plane model of d is smaller than h_d, and check whether their projection on the plane falls inside the bounding polygon of d. Figure 7 presents the segmented $\mathcal{P}_s$ from a larger sequence, together with the ADA vertical bounds on where a door handle must be located.

Fig. 7. The selected $\mathcal{P}_s$ (represented with blue color) that potentially contains the door handle, based on a previously given door model, and the ADA door handle requirements.

For each point $\boldsymbol{p}_i$ in $\mathcal{P}_s$, we take a neighborhood around it $\mathcal{P}_i^k$, and estimate the surface curvature γ_p at $\boldsymbol{p}_i$ from the eigenvalues $\lambda_0 \leq \lambda_1 \leq \lambda_2$ of $\mathcal{P}_i^k$ as follows:

$$\gamma_p = \frac{\lambda_0}{\lambda_0 + \lambda_1 + \lambda_2} \tag{2}$$

Figure 8 presents the distribution of curvatures along the $\mathcal{P}_s$ set of points, with a large percentage of the points having a surface curvature close to 0 (i.e., planar). The points having a spike in curvature space are potentially part of candidate handle clusters.
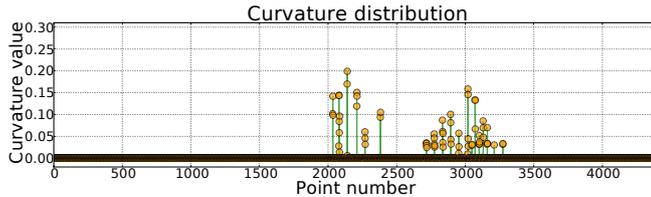

Fig. 8. The distribution of surface curvatures γ_p over $\mathcal{P}_s$. The majority of the points have a curvature value close to 0 which indicates that they lie on a planar surface.

Given the door handle's signature in intensity but also in curvature space, we proceed at analyzing the distribution of values over these spaces for all the points in $\mathcal{P}_s$. Because we are guaranteed that most of the points in $\mathcal{P}_s$ are situated on the door plane and thus part of the door itself, the resultant intensity distribution will have a highly peaked mean μ_h (see Figure 9).

Therefore, in principle we could easily select all the points $\boldsymbol{p}_i \in \mathcal{P}_s$ whose intensity value is outside $\mu_h \pm \alpha_h \cdot \sigma_h$, where σ_h represents the standard deviation of the aforementioned
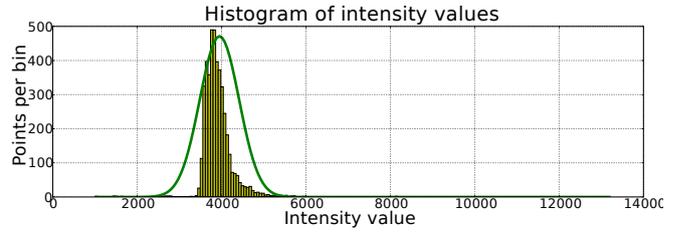

Fig. 9. The distribution of intensity values over $\mathcal{P}_s$. Notice the highly peaked mean μ_h.

intensity distribution, and α_h is a user given parameter. Similarly, another set of points $\boldsymbol{p}_j$ could be selected for the curvature distribution, if the surface curvature value γ_{p_j} of $\boldsymbol{p}_j$ is outside $\mu_c \pm \alpha_c \cdot \sigma_c$, where μ_c and σ_c represent the mean and standard deviation of the curvature distribution, and α_c is a user given parameter respectively.

However, we would like to automatically determine the values α_h and α_c for any door. To do this, we make use of Chebyshev's inequality which states that for any distribution that has a mean and a variance, at least $1 - 1/\alpha^2$ points are within α standard deviations from the mean. Therefore, we select $\alpha_{\max} = 7$ to account for at least 98% of the values, and iterate over the space of standard deviations with a given size step. For each $\alpha_i \in \{0 \cdots \alpha_{\max}\}$, we compute the number of points N_i from $\mathcal{P}_s$ falling outside the interval $\mu \pm \alpha_i \cdot \sigma$, and use the values to create two other distributions (one for intensity values and one for curvature values) over the N_i space.
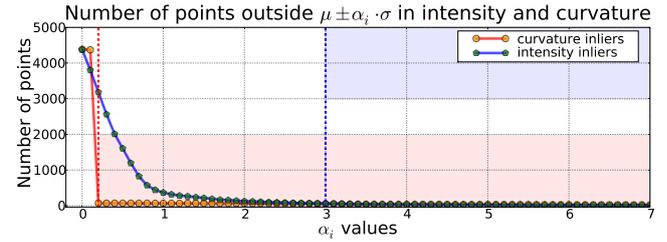

Fig. 10. The number of points $\boldsymbol{p}_i$ falling outside the interval $\mu \pm \alpha_i \cdot \sigma$ for different α_i values, for the two intensity (μ_h, σ_h) and curvature (μ_c, σ_c) distributions. The two estimated cutting points t_h and t_c are represented with dashed vertical lines.

Figure 10 presents the variations of the number of points N_i outside the interval $\mu \pm \alpha_i \cdot \sigma$ for different α_i values in both intensity and curvature spaces. The goal of this analysis is to find a minimal set of points which has extreme values in both spaces. As it can be seen from the two plots, the number of values drops really quickly as α_i grows, especially in the curvature space where almost 99% of the points are around the mean μ_c. To compute a good cutting value, our method makes use of a robust L-estimator, namely the trimean t of each distribution, and selects the best $\alpha_{c_b}, \alpha_{h_b}$ values in both cases as the values closest to the trimeans. If $\mathcal{P}_c$ and $\mathcal{P}_h$ are the point sets outside the intervals $\mu_c \pm \alpha_{c_b} \cdot \sigma_c$ and $\mu_h \pm \alpha_{h_b} \cdot \sigma_h$ respectively, then the final set of point handle candidates $\mathcal{P}_f$ is obtained as:

$$\mathcal{P}_f = \mathcal{P}_c \cap \mathcal{P}_h \tag{3}$$

The actual handle is obtained by first projecting $\mathcal{P}_f$ on the door plane, and then fitting the best horizontal line (largest

number of inliers) along the door axis in it using RMSAC. Figure 11 presents the resultant $\mathcal{P}_f$ handle inliers for the dataset in Figures 6 and 7.

Algorithm 2 Main computational steps for Handle Detection

$h_{ADA}, d, \mathcal{P}$ *// ADA door handle requirements, door candidate, input point cloud*
for all $\boldsymbol{p}_i \in \mathcal{P}$
    $d_i = \text{dist}(\boldsymbol{p}_i, d), p_{r_i} = \text{proj}(\boldsymbol{p}_i, d)$ *// distance and projection of $\boldsymbol{p}_i$ on d*
    if $(d_i \leq h_d) \wedge (p_{r_i}$ inside $d)$ *// check if close to and inside door bounds*
       if $F(\boldsymbol{p}_i, h_{ADA})$ *// does $\boldsymbol{p}_i$ respect the ADA constraints?*
         $\mathcal{P}_s \leftarrow \boldsymbol{p}_i$ *// add $\boldsymbol{p}_i$ to $\mathcal{P}_s$*
$(\mu_h, \sigma_h) \leftarrow F(\mathcal{P}_s)$ *// estimate statistics for intensity distribution*
$(\mu_c, \sigma_c) \leftarrow F(\mathcal{P}_s)$ *// estimate statistics for curvature distribution*
$\alpha_{max} = 7$ *// maximum number of $\pm$ standard deviations to check*
for all $\alpha_i \in \{0 \cdots \alpha_{max}\}$
    $N_{i_h} = F(\mu_h \pm \alpha_i \cdot \sigma_h)$ *// get number of points N_{i_h} outside $\mu_h \pm \alpha_i \cdot \sigma_h$*
    $N_{i_c} = F(\mu_c \pm \alpha_i \cdot \sigma_c)$ *// get number of points N_{i_c} outside $\mu_c \pm \alpha_i \cdot \sigma_c$*
$t_h \leftarrow F(N_{i_h})$ *// compute the trimean for number of points N_{i_h} distribution*
$t_c \leftarrow F(N_{i_c})$ *// compute the trimean for number of points N_{i_c} distribution*
$(\alpha_{b_h}, \mathcal{P}_h) \leftarrow F(t_h, N_{i_h})$ *// estimate the best cutting α_{b_h} value, and $\mathcal{P}_h$*
$(\alpha_{b_c}, \mathcal{P}_c) \leftarrow F(t_c, N_{i_c})$ *// estimate the best cutting α_{b_c} value, and $\mathcal{P}_c$*
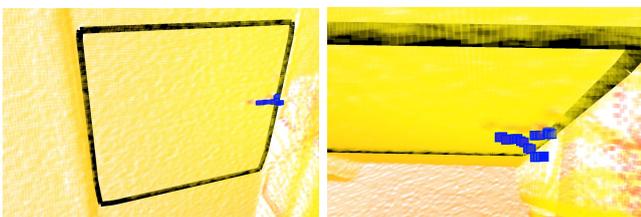$\mathcal{P}_f = \mathcal{P}_h \cap \mathcal{P}_c$

Fig. 11. Resultant handle inliers using dual distribution statistics analysis.

VI. VALIDATION

This section describes our method to validate the effectiveness and accuracy of the door and handle detection nodes in real world situations.

The presented door and handle detection methods compute the 3D pose (position and orientation) of the door and the door handle. To validate the correctness of the handle pose, the detection phase is followed by a validation phase, where the PR2 robot approaches the door and attempts to grasp the door handle. The grasping is performed in "open loop", meaning that no other sensor feedback is used to gather information about the door and handle pose after the initial detection phase. The grasping of the handle consists of two phases. In the first phase the robot base navigates to a pose in front of the door, from where the door handle is within the workspace of the robot arm. The base navigation node moves the base towards the desired pose, based on measurements from the wheel odometry and an inertia measurement unit (IMU). In the second phase, the robot arm moves along a collision free trajectory to position the robot gripper on the door handle. The success rates for the door and handle detection in this paper are based on the robot's ability to achieve a caging grasp around the door handle.

The presented validation method allows for small misalignments of the gripper with respect to the door handle, and still achieves a successful grasp. With the gripper fully opened, the distance between the fingertips is 8cm. For a door handle with a thickness of 2cm, this results in a maximum allowed vertical gripper misalignment of 3cm. The allowed left-right gripper misalignment is bounded by the width of

the door handle to a maximum of 3.5cm. The misalignments observed in real world experiments are typically within 2cm.

VII. EXPERIMENTAL RESULTS

To validate our proposed framework, we have applied our methods to more than 50 different situations, where doors are scanned in different states (open/closed/half open) and from different angles[2]. The datasets have been acquired in varying light conditions, sometimes completely in the dark, thus accounting for a very large variation of situations where 2D image-based segmentation would fail. Both the door and handle detection performed extremely well and succeeded in segmenting and identifying the object components given valid input data which respected the system constraints. Tables I and II present examples of the estimation results obtained on subsets of the above mentioned datasets. The first 4 rows of both tables present valid results, verified using the methods presented in Section VI.

Some of the situations which could not be entirely handled by our methods include examples such as the ones shown in the bottom row of the tables. For example, in Table I, we see examples of (from left to right): i) a wall which respects the geometric constraints of a door being segmented as a door candidate; ii) only one out of two doors being successfully identified in the data, iii) and iv) unsuccessful door segmentation due to no geometric differences between the walls of the room and the door itself. The latter is an issue for our door segmentation approach, but unfortunately it cannot be solved by a revised version of the algorithm, simply because the geometry and intensity levels of the door are indistinguishable from the wall. Instead, we plan to make use of the stereo data available on our PR2 robot in these situations, and attempt to detect the door edges in RGB space, and use that as a starting point for the Door Detection algorithm.

A few failure cases for the handle segmentation method are also presented in the bottom row of Table II. In the first example, the door candidate given by the executive was a refrigerator door, but due to variations in geometry on the door handle, only a partial subset of it was detected and returned. The next two cases are not really failures, as the candidates given are in fact walls, and thus the handle detection algorithm returns an empty list of handles. The last example however presents an undersegmentation of the points belonging to the door handle. The explanation of this error is given in Figure 12, where a closeup of the same dataset is shown. Due to the fact that the sampled geometry of the handle contains only 3 point hits (as presented in the right part of the figure), the handle extraction algorithm rejects the candidate. An immediate solution to this problem is to take a new scan and simply concatenate the previous dataset with the new one to obtain more point hits on the door handle. However we plan to investigate this further to see what other situations could create similar problems.

---

[2]The point cloud datasets of the experiments and a demonstration video explaining the acquisition and processing of point cloud data for the purpose of door and handle identification and validation can be downloaded from http://www.willowgarage.com/icar2009-doorhandle
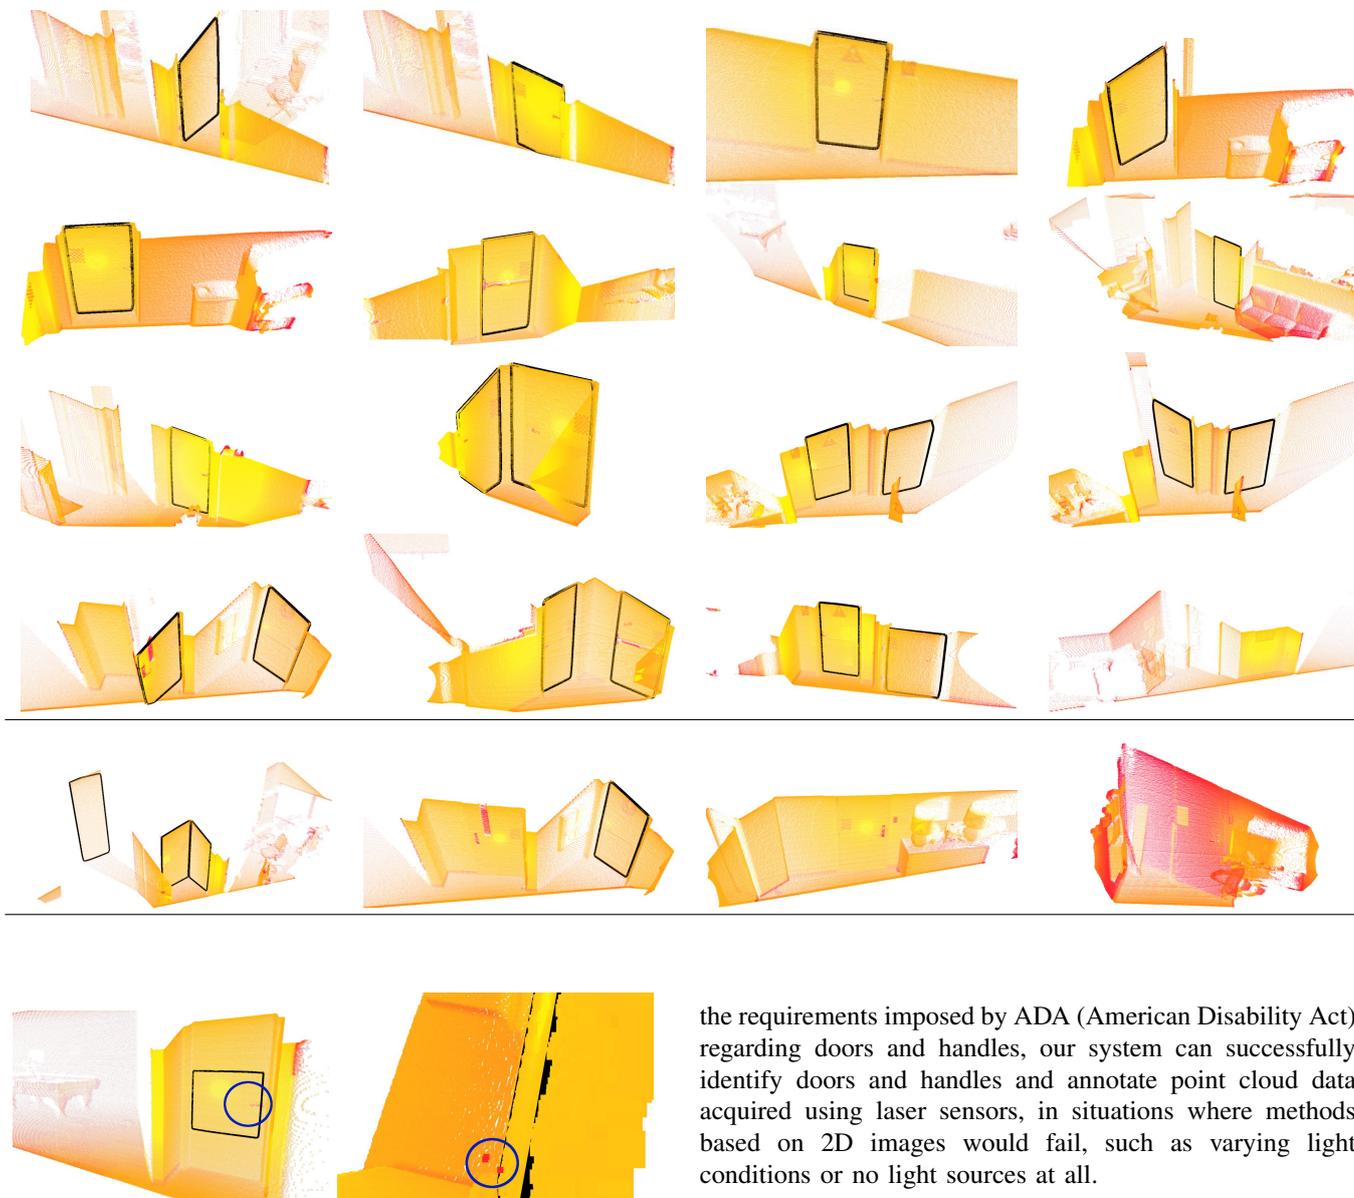
Fig. 12. Handle identification failure due to an extremely low (3) number of point hits on the actual door handle.

As shown in the results presented, our methods are not influenced by the door opening angle or the handle type, as long as some basic constraints are respected, namely: there exists a distinguishable difference in intensity and curvature between the actual handle and the door plane.

VIII. CONCLUSIONS AND FUTURE WORK

In this paper we presented a set of robust methods for the problem of door and handle identification from noisy scanned data in indoor environments. By refraining from using camera 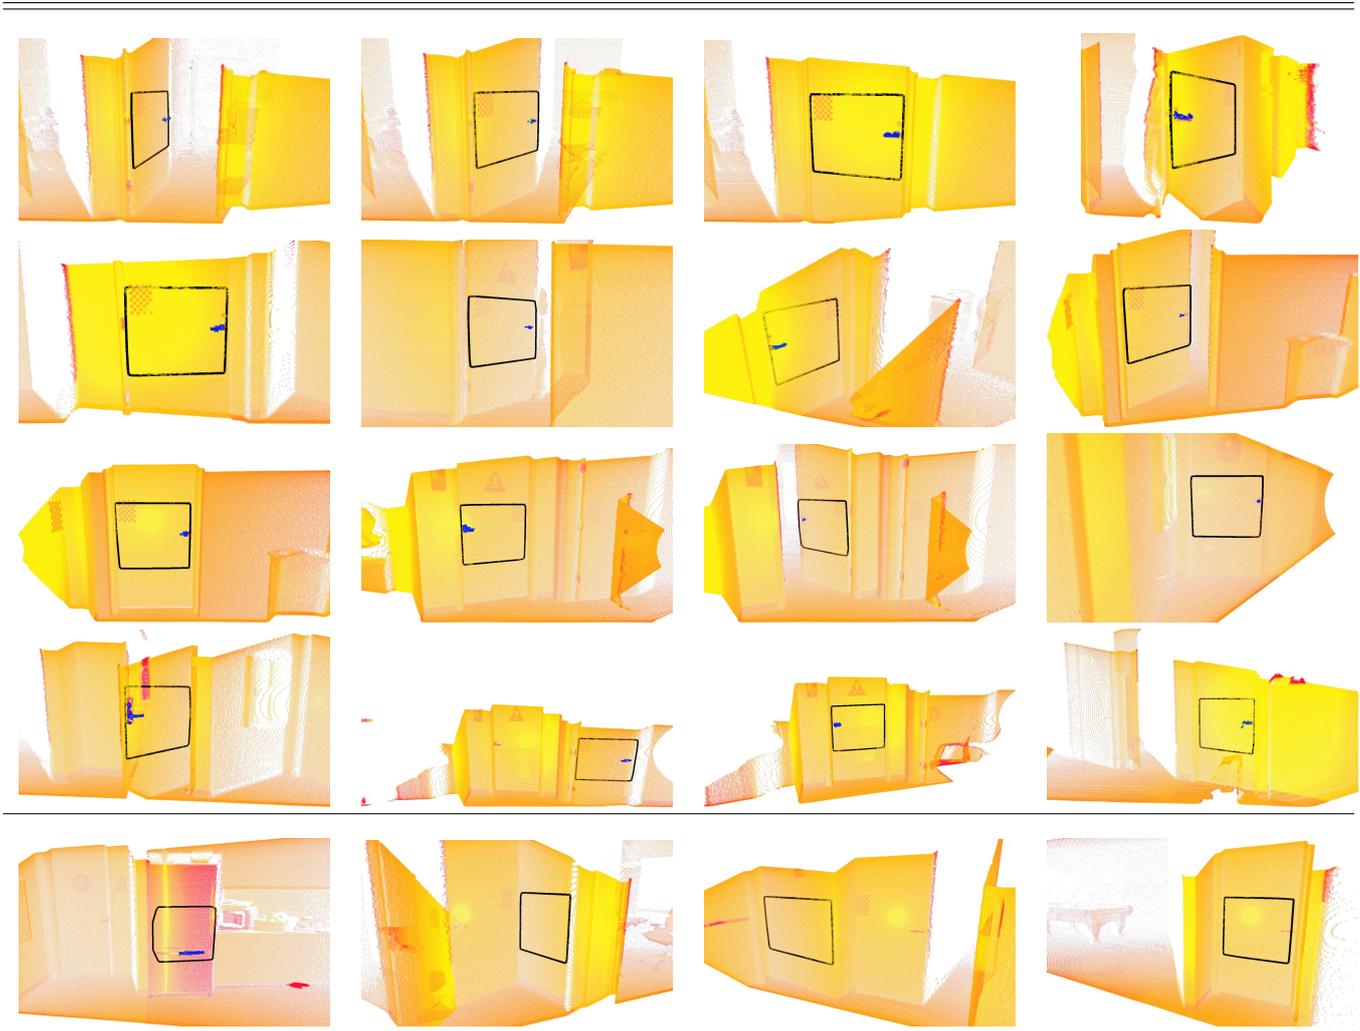images and constraining the search space using the requirements imposed by ADA (American Disability Act) regarding doors and handles, our system can successfully identify doors and handles and annotate point cloud data acquired using laser sensors, in situations where methods based on 2D images would fail, such as varying light conditions or no light sources at all.

The proposed dual intensity-curvature distribution analysis has shown promising results, and we plan to continually improve the results by testing the methods against an even larger set of datasets representing doors. In situations where multiple handle candidates are extracted for a given door, we plan to make use of machine learning classifiers to obtain the best solution by restricting the candidate cluster attributes to a predefined subspace. While this is still work in progress, we already make our software and datasets available as an off-the-shelf component of the ROS project.

TABLE II

A subset of 20 datasets used to test our handle detection algorithm. The first 4 rows show successful identification cases, while the last row presents challenging situations for: a) a refrigerator door; b) and c) testing the handle detection on two walls which resemble door candidates; d) unsegmented handle for a regular door.

REFERENCES

[1] R. B. Rusu, Z. C. Marton, N. Blodow, M. Dolha, and M. Beetz, "Towards 3D Point Cloud Based Object Maps for Household Environments," *Robotics and Autonomous Systems Journal (Special Issue on Semantic Knowledge)*, 2008.

[2] A. Nuechter and J. Hertzberg, "Towards Semantic Maps for Mobile Robots," *Journal of Robotics and Autonomous Systems (JRAS), Special Issue on Semantic Knowledge in Robotics*, pp. 915 – 926, 2008.

[3] D. Anguelov, D. Koller, E. Parker, and S. Thrun, "Detecting and modeling doors with mobile robots," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2004.

[4] J. K. T. Alexander Andreopoulos, "A Framework for Door Localization and Door Opening Using a Robotic Wheelchair for People Living with Mobility Impairments," in *Robotics: Science and Systems, Workshop: Robot Manipulation: Sensing and Adapting to the Real World, Atlanta, June 30*, 2007.

[5] A. Andreopoulos and J. Tsotsos, "Active Vision for Door Localization and Door Opening using Playbot: A Computer Controlled Wheelchair for People with Mobility Impairments," *Computer and Robot Vision, 2008. CRV '08. Canadian Conference on*, pp. 3–10, May 2008.

[6] E. Aude, E. Lopes, C. Aguiar, and M. Martins, "Door Crossing and State Identification Using Robotic Vision," 2006.

[7] A. Jain and C. C. Kemp, "Behaviors for Robust Door Opening and Doorway Traversal with a Force-Sensing Mobile Manipulator," in *RSS Manipulation Workshop: Intelligence in Human Environments*, 2008.

[8] B. S. E. Jauregi, J. M. Martinez-Otzeta and E. Lazkano, "Door Handle Identification: a Three-Stage approach," in *6th International Conference on Intelligent Autonomous Vehicles (IAV)*, 2007.

[9] A. Y. N. Ellen Klingbeil, Ashutosh Saxena, "Learning to Open New Doors," in *Robotics Science and Systems (RSS) workshop on Robot Manipulation*, 2008.

[10] M. I. Mitsuharu Kojima, Kei Okada, "Manipulation and Recognition of Objects Incorporating Joints by a Humanoid Robot for Daily Assistive Tasks," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2008.

[11] A. Petrovskaya and A. Y. Ng, "Probabilistic Mobile Manipulation in Dynamic Environments, with Application to Opening Doors," in *International Joint Conference on Artificial Intelligence (IJCAI), Hyderabad, India*, 2007.

[12] C. Ott, B. Baeuml, C. Borst, and G. Hirzinger, "Autonomous Opening of a Door with a Mobile Manipulator: A Case Study," in *6th IFAC Symposium on Intelligent Autonomous Vehicles (IAV)*, 2007.

[13] P. Torr and A. Zisserman, "MLESAC: A new robust estimator with application to estimating image geometry," *Computer Vision and Image Understanding*, vol. 78, pp. 138–156, 2000.