

3D Pose Estimation, Tracking and Model Learning of Articulated Objects from Dense Depth Video using Projected Texture Stereo

Jürgen Sturm¹

Kurt Konolige²

Cyrill Stachniss¹

Wolfram Burgard¹

Abstract—Service robots deployed in domestic environments generally need the capability to deal with articulated objects such as doors and drawers in order to fulfill certain mobile manipulation tasks. This however, requires, that the robots are able to perceive articulated furniture objects such as cupboards, dishwashers and cabinets. In this paper, we present an approach for detecting, tracking, and learning 3D articulation models for doors and drawers without using artificial markers. Our approach uses a highly efficient and sampling-based approach to rectangle detection in dense depth images obtained from a self-developed projected texture stereo vision system. The robot can use the generative models learned for the articulated objects to estimate their mechanism type, their current configuration, and to predict their opening trajectory. In our experiments we demonstrate that (1) we obtain dense depth images in the workspace of our robot using our camera system, (2) we are able to robustly and reliably detect cabinet fronts from depth images, and (3) are able to learn accurate articulation models for the observed articulated objects. We furthermore provide a detailed error analysis based on ground truth data obtained in a motion capturing studio.

I. INTRODUCTION

Home environments are envisioned as one of the key application areas for service robots. Robots operating in such environments often have to deal with articulated objects such as doors or drawers. In the past, several researchers have addressed the problem of estimating and handling doors and drawers [9], [4], [11], [15]. Most of these approaches, however, are either entirely model-free or assume substantial knowledge about the model and its parameters. Whereas model-free approaches release designers from providing any a-priori model information, the knowledge about objects and their articulation properties may greatly support the state estimation and the simulation, planning, and verification of the actions of the robot. This paper is an extended version of our recent work presented recently [24], with additional experiments on the projected texture system.

We consider the problem of acquiring articulation models of doors and drawers from sequences of depth images acquired with an active stereo camera also presented in this paper. This approach has several advantages. First, it does not rely on artificial markers attached to objects, and second, we do not need to employ expensive range scanners which have the additional disadvantage that they poorly deal

¹ Jürgen Sturm, Cyrill Stachniss and Wolfram Burgard are with the Autonomous Intelligent Systems Lab, Computer Science Department, University of Freiburg, Germany. {sturm, stachnis, burgard}@informatik.uni-freiburg.de

² Kurt Konolige is with Willow Garage Inc., Menlo Park, CA 94025, USA. konolige@willowgarage.com

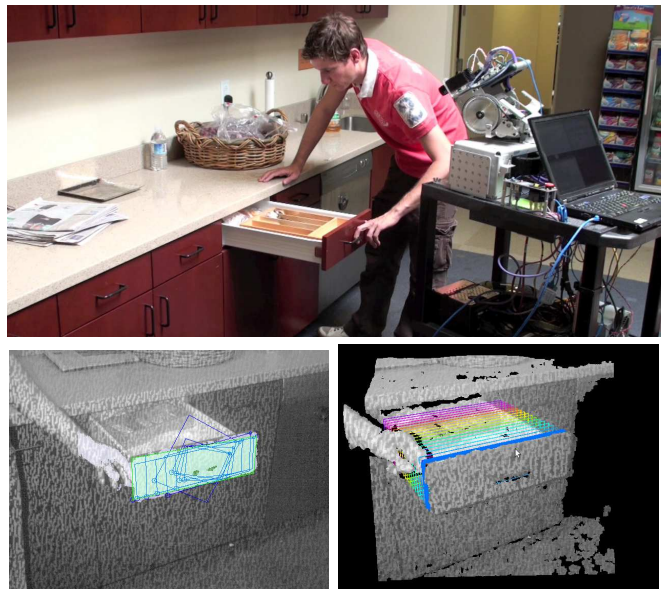


Fig. 1. Top: A drawer is opened and closed and observed with a stereo camera in combination with projected texture. Bottom left: After plane segmentation, we optimize iteratively the pose of a rectangle and evaluate the model fit directly in the disparity image. Bottom right: After combining these detections into a track, we fit an articulation models.

with moving objects, making them inconvenient for learning articulations.

In our concrete scenario, the perception of articulated drawers and doors in a kitchen environment requires the accurate detection of rectangular objects in the depth image sequences, see Fig. 1. We present a highly efficient algorithm that segments the point clouds into planes, and then iteratively fits rectangles to each plane separately. Our perception algorithm can be adapted to the computational capabilities of the robot as it allows to adjust the number of rectangle detections per frame. We furthermore track rectangles over multiple frames. The individual tracks are then used to learn the articulation models. The learning approach instantiates multiple candidate articulation models and selects the one that best explains the data. Once a model has been selected, the robot can use it to predict future configurations of the objects.

Our implementation has been made available within Willow Garage’s open source robotics repository [23].

II. RELATED WORK

A. Projected Texture Stereo

For our application, we require real-time (15 Hz), accurate and dense point clouds of the scene. Flash ladars [2]

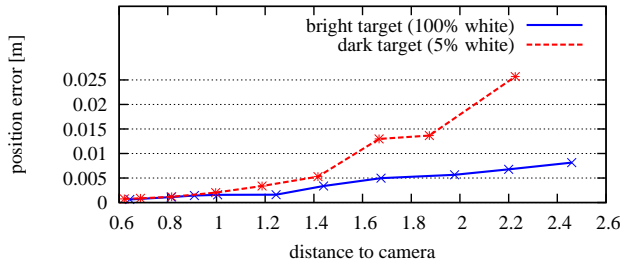


Fig. 2. Positional error of a planar target, observed with our active stereo system. For a white target, the error stays below 2mm until after 1.2m, then goes up to about 1cm at 2.5m. For a very dark target, error is low close up, then becomes larger at distance, when the pattern is difficult to see.

often have poor depth and spatial resolution, and have non-Gaussian error characteristics that are difficult to deal with. Line stripe systems [7], [19] have the requisite resolution but cannot achieve 15 Hz operation, nor deal with moving objects. Monocular structured light systems [22] can achieve reasonable frame rates and can sometimes deal with object motion, but still rely on expensive and high-powered projection systems, while being sensitive to ambient illumination and object reflectance.

Stereo systems that employ matching algorithms to produce dense results [6], [13], [28] can be a suitable sensor for our application. However, passive stereo suffers from the problem of *dropouts*: areas of low texture cannot be matched correctly. An interesting and early technology is the use of stereo with structured light [18], [17]. Unlike structured light systems with single cameras, stereo does not depend on the relative geometry of the light pattern – the pattern just lends texture to the scene. Hence the pattern and projector can be simplified, and standard stereo calibration techniques can be used to obtain accurate and dense 3D measurements.

One variant of this technique, known as Spacetime Stereo (STS) [8], [32], varies the pattern over time and integrates several frames. It produces outstanding results on static scenes and under controlled illumination conditions, but moving objects create difficulties [32], [29], [26].

We have developed a compact projector for active stereo with a fixed, random pattern [14], using a red high performance LED providing 5 optical Watts. It provides a texture for stereo that produces excellent error characteristics at distances up to 3 meters, even for surfaces with low reflectivity, see Fig. 2.

In further experiments, we found that by using the texture projector, we obtain very dense disparity images, see Table I for the number of point correspondences with and without projected texture. Depending on the texture on the objects in the scene, the number of correspondences improved in the order of one to two magnitudes.

B. Model-based Detection

Locating objects from 2D images has a long history in computer vision research [5]. Recent approaches for door detection from camera images include [1] and [3]. For

Dataset	Normal stereo vision	Projected texture stereo	Improvement (textured/normal)
Door	2904.33	194158.27	66.85x
Drawer	6738.98	172426.33	25.59x
Two drawers	12950.42	144916.06	11.19x
Fuse door	7538.08	216917.17	28.77x
Fire extinguisher door	24614.38	211319.51	8.58x

TABLE I
AVERAGE NUMBER OF POINT CORRESPONDENCES WITH AND WITHOUT THE TEXTURE PROJECTOR. USING THE PROJECTOR RESULTS IN A SIGNIFICANT IMPROVEMENT.

3D point clouds, Hough transforms [27], EM-based algorithms [31] and RANSAC-based approaches [21] have been used successfully for line and plane fitting.

For this work, we evaluated several of the above approaches w.r.t. their applicability to the depth data from our active stereo camera system. We found that RANSAC-based plane segmentation combined with iterative fitting yielded the most robust and accurate results in our context.

C. Learning Articulation Models

Yan and Pollefeys [30] present an approach for learning the structure of an articulated object from feature trajectories under affine projections. Other researchers have addressed the problem of identifying different object parts from image data. Ross *et al.* [20] use multi-body structure from motion to extract links from an image sequence and then fit an articulated model to these links using maximum likelihood learning. There exist several approaches where tracking articulated objects is the key motivation and often an a-priori model is assumed. Krainin *et al.* [16], for example, described recently an approach for tracking articulated objects such as a manipulator including its hand, using a depth camera with a texture projector similar to ours. However, they require a geometric model of the manipulator. Katz *et al.* [12] learn planar kinematic models for various articulated objects in 2D using a KLT tracker. Kragic *et al.* [15] describe an integrated navigation system for mobile robots which includes a vision-based system for the detection of door handles that enables the robot to successfully open doors. Anguelov *et al.* [4] model doors as line segments that rotate around a hinge. EM is then used to find the model parameters both from 2D range data and images.

In our previous work [25], we learned articulation models for various objects in full 3D using artificial markers. In this paper, we present an extension of our previous approach, that allows us to observe object parts in 3D directly from depth images and thus learn the models without requiring artificial markers. We regard this as an essential requirement for real-world applications.

III. APPROACH

In this section, we first briefly describe the structured light approach to obtain dense depth images from stereo. We then present our sampling-based rectangle detector for point clouds before we illustrate how the individual observations

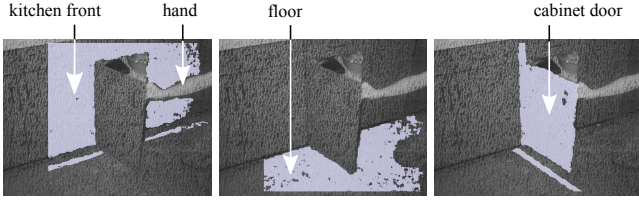


Fig. 3. Finding the three most prominent planes with our RANSAC-based approach.

can be combined into consistent tracks. Finally, we show how articulation models can be learned from such tracks.

A. Dense Depth Images from Stereo and Projected Texture

We consider a projector and a standard calibrated stereo camera configured to be as nearly coincident as possible. For simplicity, the focal length of the projector and camera are similar, so that at any distance the projected pattern appears to be the same size in the camera images. A compact, high-power LED device projects a fixed pattern as a random grid of black and white squares, in sync with the camera exposure. When it is seen by a camera, the pattern produces an image. We use a standard block-matching algorithm to compute the disparity of each pixel [13], that is, the offset between the left and right images. The algorithm runs at 15 Hz for 640x480 at 128 disparities on a single 2 GHz Pentium core.

We tested the device with a 50 deg field of view, using both white and 5% reflectance black planar targets at different distances. The error is taken to be the standard deviation from the best-fit plane. From Fig. 2, the system shows very low error, even out to 2.5 m. For the white target, the error stays below 1 cm throughout this range. Some of the error at the larger distances comes from calibration, as the reconstructed plane will not be perfectly flat. Up to over 1 m, the error is about 2 mm, which is good enough to reconstruct even fine objects. Even with a very dark (5% reflectance) target, the system gives good results up to 2 m, with some increase in error at the larger distance.

From our stereo processing system, we obtain in each frame a disparity image $D \in \mathbb{R}^{640 \times 480}$, that contains for each pixel (u, v) its perceived disparity $D(u, v) \in \mathbb{R}$. The relationship between 2D pixels in the disparity image and 3D world points is defined by the projection matrices of the calibrated stereo camera, and can be calculated by a single matrix multiplication from the pixel coordinates and disparity.

B. Model-based Perception from Depth Images

1) *Sampling planes*: Our RANSAC-based plane fitting algorithm samples three pixels from the depth image, computes from them a plane with coefficients $\mathbf{z}_{\text{plane}} \in \mathbb{R}^4$, and then counts the inliers of that plane. We define the plane to comprise all pixels that are within a certain distance d of the plane, i.e., for which the following holds:

$$\|\mathbf{z}_{\text{plane}}(x \ y \ z \ 1)^T\| \leq d. \quad (1)$$

In general, d depends on the particular noise level of the camera – in our case, we used $d = 0.02$ m. We repeat this

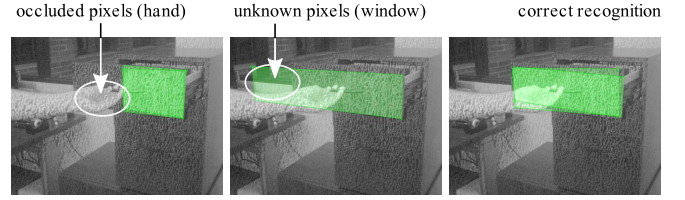


Fig. 4. Effect of the cost parameter for unknown and occluded pixels. Left: cost too high (1.0). Middle: cost too low (0.0). Right: good (0.2).

process of plane candidate generation until we find a plane with a high enough support, or we exceed a given number of iterations. We select the plane with the most inliers and subtract the corresponding inliers from the point cloud. Subsequently, we apply the same strategy to the remaining points in the cloud, until no more points remain.

For each plane, we create an image mask $M \in \{\text{in-plane, free, occluded, unknown}\}^{640 \times 480}$ with labels for the pixels in the depth image, i.e.,

$$M(u, v) = \begin{cases} \text{in-plane} & \text{if } \|\mathbf{z}_{\text{plane}}(x \ y \ z \ 1)^T\| \leq d \\ \text{free} & \text{if } \mathbf{z}_{\text{plane}}(x \ y \ z \ 1)^T > d \\ \text{occluded} & \text{if } \mathbf{z}_{\text{plane}}(x \ y \ z \ 1)^T < -d \\ \text{unknown} & \text{otherwise} \end{cases} \quad (2)$$

Here, “in-plane” indicates that the pixel belongs to the plane for which the mask M is computed. In contrast to that, “free” indicates that the observed pixel lies behind the plane and “occluded” that a pixel in front of the plane has been observed which occludes the plane. “Unknown” means that no depth information is available for that pixel.

In contrast to typical approaches to RANSAC-based plane fitting which always assign pixels to one plane, our masks allow points to belong to several planes at the same time. This is useful, as the infinite planes determined via RANSAC always intersect with the subsequent (less significant) planes, thereby cutting out points that make detection of contiguous rectangles more difficult in the next step of the perception process.

For a visualization of the result, see Fig. 3. In this example, our algorithm automatically segmented three planes from a depth image of a cabinet door.

2) *Sampling rectangles*: The next step is to find rectangles in the segmented planes. A rectangle in 3D space has 8 degrees of freedom: its position, its orientation and its dimensions (3+3+2). After the plane segmentation, we have already fixed 3 DOFs, so that we need to find the remaining 5 DOFs. We apply an iterative fitting approach here. We start with a sampled candidate rectangle and optimize its pose and size iteratively using an objective function g .

For creating an initial rectangle candidate, we sample a random point from the plane, and sample the other DOFs from a prior distribution. The objective function g is based on the average cost of the pixels inside the rectangle $\mathbf{z}_{\text{rect}} \in \mathbb{R}^8$,

$$g(\mathbf{z}_{\text{rect}}) := -\frac{1}{|\text{pixels}(\mathbf{z}_{\text{rect}})|^{1+\alpha}} \sum_{\text{pixels}(\mathbf{z}_{\text{rect}})} \text{cost}(M(u, v)) \quad (3)$$

The parameter α (that we empirically chose around $\alpha =$

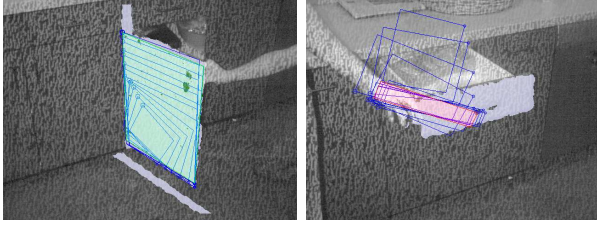


Fig. 5. Left: In each plane, we pick a random starting point from which we optimize iterative the pose and size of the candidate rectangle. Right: As our model fitting procedure is greedy, it can get trapped into local maxima.

0.05) makes g slightly favor larger rectangles over smaller ones.

Finding a good cost metric cost, in particular for occluded and unknown pixels, is non-trivial. If chosen too low, the greedy search converges on too large rectangles, while a too high cost increases the amount of local maxima in g and in turn leads to the detection of partial rectangles in the presence of occlusions (see Fig. 4).

In each iteration, we now individually optimize every DOF of the rectangle. We apply a small set of discrete changes to each DOF, and evaluate the objective function on $\mathbf{z}'_{\text{rect}}$. If $g(\mathbf{z}'_{\text{rect}}) > g(\mathbf{z}_{\text{rect}})$, we continue with the improved parameter set. When this greedy search converges (or we reach the maximum number of iterations), we need to evaluate the quality of the found match. In preliminary experiments, we found that the value of the objective function was not sufficient for discrimination of false and true positives.

Therefore, we decided to evaluate the rectangle candidate \mathbf{z}_{rect} using two measures, that are inspired from statistical classification theory and that have a natural interpretation. First, we evaluate the precision $r_{\text{precision}}$ of the rectangle candidate as the ratio of detected pixels and all pixels in the found rectangle. Second, we evaluate the recall r_{recall} as the ratio of pixels in the found rectangle versus the pixels in the selected plane $\mathbf{z}_{\text{plane}}$. For both measures, we used our cost functions to weight occluded and unknown pixels accordingly.

$$r_{\text{precision}}(\mathbf{z}_{\text{rect}}) := \frac{\sum_{\text{pixels}(\mathbf{z}_{\text{rect}})} 1 - \text{cost}(M(u,v))}{|\text{pixels}(\mathbf{z}_{\text{rect}})|} \quad (4)$$

$$r_{\text{recall}}(\mathbf{z}_{\text{rect}}) := \frac{\sum_{\text{pixels}(\mathbf{z}_{\text{rect}})} 1 - \text{cost}(M(u,v))}{\sum_{\text{pixels}(\mathbf{z}_{\text{plane}})} 1 - \text{cost}(M(u,v))} \quad (5)$$

Empirically, we found that a good condition for thresholding is to require that both ratios are above 0.7, which removes most of false positives.

An example of the iterative pose fitting is given in the left image of Fig. 5: the rectangle candidate started in the lower left of the door, and iteratively converged to the correct pose and size of the door. The candidate is accepted, because both ratios $r_{\text{precision}}$ and r_{recall} have high values. The greedy search however can get stuck in local maxima. In the example depicted in the right image of Fig. 5, the hand is also part of the drawer front plane and the candidate rectangle converged to a rectangle that fits to some extent the hand. Our algorithm then rejects this candidate rectangle because it does not contain the majority of pixels in the plane, i.e., r_{recall} takes a low value.

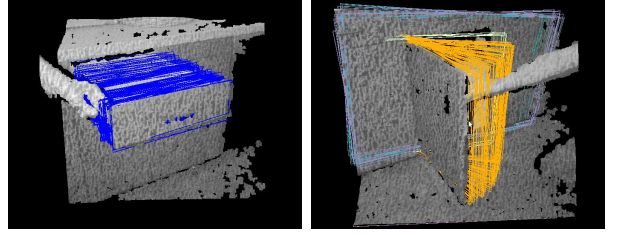


Fig. 6. Observed tracks of a cabinet drawer (left) and a cabinet door (right).

We deal with the problem of local maxima by starting from several rectangle candidates. In this sense, our algorithm is probabilistically complete, as we would find any visible rectangle in the limit with probability 1. In practice, we chose a fixed number m of samples per plane.

C. Tracking

In the remainder of this paper, we drop the subscript in $\mathbf{z}_{\text{rect}} = \mathbf{z}$ to improve readability. The rectangle detector described in the previous section gives us per frame between zero and $n \cdot m$ observations of rectangles (m rectangles in n planes), which need to be integrated into consistent tracks. Checking whether two rectangles \mathbf{z}_i and \mathbf{z}_j are similar requires to take the ambiguity in the representation into account: the same rectangle can be described by eight different parameter vectors (depending on the choice of the corner of origin, and the choice of the front or back side of the rectangle). A track \mathbf{t} is an integrated sequence of k rectangle observations $\mathbf{z}^{1:k} = \mathbf{z}^1, \dots, \mathbf{z}^k$ that were collected until time t representing the trajectory of a single moving part.

D. Learning Models of Articulated Objects

Our approach for learning models of articulated objects aims at estimating the kinematic nature of the observed tracks of objects in the scene and consists of the following parts:

1) *Training Template Models for the Observed Tracks of Object Parts:* Since we have no prior information about the nature of the connection between object parts, we do not aim to fit a single model, but instead fit a set of candidate template models representing different kinds of links. This candidate set consists of parameterized models that occur in various objects including a rotational link ($\mathcal{M}^{\text{rotational}}$), a prismatic link ($\mathcal{M}^{\text{prismatic}}$), and a rigid transformation ($\mathcal{M}^{\text{rigid}}$). All models except $\mathcal{M}^{\text{rigid}}$ have a latent variable q that describes the configuration of the link. For a door, the variable q for example describes the opening angle of the door.

Let us consider a track \mathbf{t} . To train the candidate models for this track, we have a sequence of k noisy observations $\mathbf{z}^{1:k}$ acquired by the tracker. Each candidate template model has its own training or estimation procedure from the track observations $\mathbf{z}^{1:k}$. For example, for a rotational joint model, we need to estimate the rotation axis and the radius. For further details, we refer the reader to [25].

2) *Evaluating a Model:* Besides training each model template, we need to evaluate its performance to subsequently select the model that explains the data best.

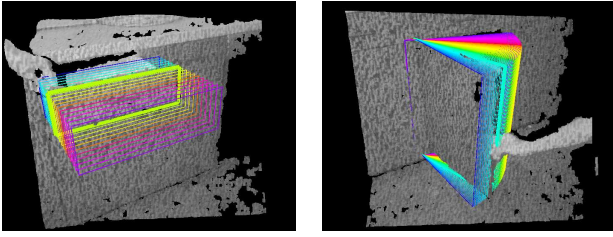


Fig. 7. Left: Articulation model learned from observing a drawer. Right: Same for a door.

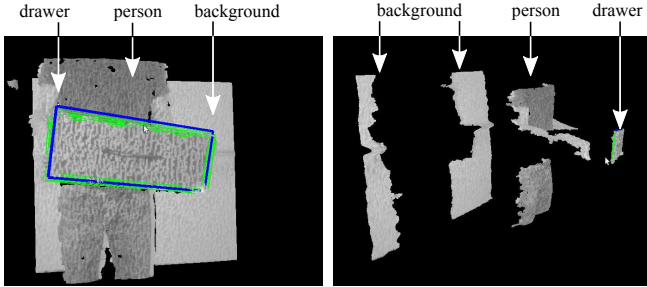


Fig. 8. The blue rectangle shows the ground truth location obtained with a motion capturing studio, while the green rectangles show our estimates.

Let \mathcal{M} be the articulation model describing the observations $\mathcal{D} = \mathbf{z}^{1:k} = \mathbf{z}^1, \dots, \mathbf{z}^k$ of a track \mathbf{t} . To evaluate how well a single observation \mathbf{z} can be explained by a model, we have to determine

$$p(\mathbf{z} | \mathcal{M}) = \int_q p(\mathbf{z} | q, \mathcal{M}) p(q | \mathcal{M}) dq. \quad (6)$$

Under the assumption that no latent configuration state q is more likely than another one, this simplifies to

$$p(\mathbf{z} | \mathcal{M}) = \int p(\mathbf{z} | q, \mathcal{M}) dq. \quad (7)$$

To evaluate $p(\mathbf{z} | q, \mathcal{M})$, that is, a measure for how well model \mathcal{M} parameterized by q explains the observation \mathbf{z} , we first compute the expected transform

$$\hat{\mathbf{z}} = \mathbb{E}_{\mathcal{M}}[\mathbf{z} | q] = f_{\mathcal{M}}(q) \quad (8)$$

using a model-specific transformation function $f_{\mathcal{M}}(\mathbf{z})$ that computes the expected pose of the object given q . The transformation functions for all template models are described in [25]. Under a Gaussian error assumption, the observation likelihood then becomes

$$p(\mathbf{z} | q, \mathcal{M}) \propto \exp(-\|\hat{\mathbf{z}} - \mathbf{z}\|^2 / \sigma^2) \quad (9)$$

and finally, we can compute the marginal data likelihood as

$$p(\mathcal{D} | \mathcal{M}) = \prod_{\mathbf{z} \in \mathcal{D}} p(\mathbf{z} | \mathcal{M}). \quad (10)$$

3) *Model Selection*: With the above mentioned approach, we can estimate for each track a set of actuation models $\mathcal{M}^{\text{rigid}}$, $\mathcal{M}^{\text{prismatic}}$, $\mathcal{M}^{\text{rotational}}$ and the corresponding observation likelihood using Eq. 10.

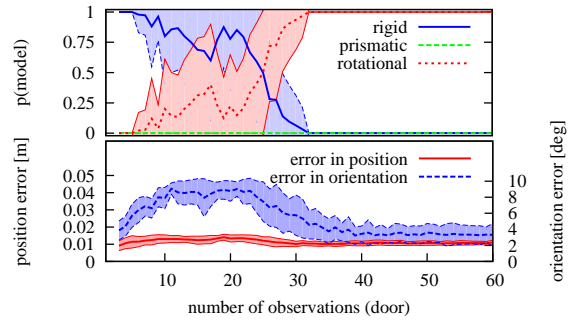


Fig. 9. Evaluation of the articulation models learned for a cabinet door, averaged over 50 runs. The plots at the top show the probability of the articulation model templates, the plots at the bottom show the prediction error of the learned model.

For selecting the model, we assign to each learned articulation model a cost that is equal to the negative expected data log-likelihood plus a complexity penalty of the model:

$$\text{cost}_{\mathcal{M}^{\text{type}}} = -\frac{1}{\|\mathcal{D}\|} \log p(\mathcal{D} | \mathcal{M}^{\text{type}}) + C(\mathcal{M}^{\text{type}}). \quad (11)$$

Then, we select for each track individually the model that has the lowest cost. This articulation model then explains the data of the observed track best while considering at the same time also the model complexity.

IV. EXPERIMENTS

A. Recognition Capabilities

To evaluate the performance of our sampling-based perception approach, we obtained ground truth pose information from a motion capturing studio. Tracking LEDs were added to an unmounted drawer, and a log file containing 19,412 stereo images including pose information was recorded under a large variety of different poses (see Fig. 8).

As a first result, we found that the drawer was correctly detected in more than 75% of the images up to a distance of 2.3 m from the camera. We also found that the number of significant planes identified via RANSAC that need to be searched increases almost linearly with the distance from the camera. This is an expected result since the drawer appears smaller in the depth image the further it is away.

The average position error of the estimator was on average below 1.5 cm. It also was almost independent of the actual distance to the camera. The same holds for the orientation error, that was on average below 3 deg.

In comparison with our previous results [25], the accuracy of our sampling-based perception on active stereo images is approximately five times higher than with the marker-based tracking system [10].

Furthermore, we validated our approach on large number of different doors and drawers in two different kitchens. Also, we successfully tested the detector on a small office pedestal with three drawers of different size, a fuse door and a fire extinguisher door in the wall.

B. Learned Articulation Models

For evaluating the robustness of our articulation model learner, we recorded detailed logfiles of both a door (39.5 cm

$\times 58$ cm) and a drawer (39.5 cm \times 12.5 cm) of a typical kitchen interior that were repeatedly opened and closed. We recorded a total of 1,023 and 5,202 images. From these logs, we sampled uniformly around 100 images for 50 times, and ran our detector and tracker as described in Sec. III on each of these 50 down-sampled logfiles. For the resulting tracks, we trained the three articulation models and evaluated the outcome of the model selection process (see Fig. 9 (top) for the door dataset).

For both datasets, we found that roughly for the first 10 observations, mostly the rigid model was selected, as no substantial motion of the drawer or door was yet detected. The more observations are added to the track, the higher the error between the (rigid) model predictions and the observations becomes. As a result, the prismatic and rotational models are selected more frequently. After 30 observations, model selection has converged in all cases to the true model. For the drawer model we reach a predictive accuracy of 1 cm and 7 deg; for the door we measured a predictive accuracy of 1 cm and 3.5 deg (see Fig. 9 (bottom)). Model fitting and selection takes on average 7 ms, and thus can be easily computed in real-time on a mobile robot.

V. CONCLUSION

In this paper, we presented an approach for learning articulation models for doors and drawers without requiring artificial markers. Instead, our approach detects and tracks doors and drawers in depth images obtained from a self-developed projected texture stereo vision system. It employs a highly efficient approach to detect and track rectangles in sequences of depth images and uses the resulting tracks to learn accurate articulation models for the corresponding objects. We evaluated our algorithm in extensive experiments also including ground truth data.

REFERENCES

- [1] J. J. Guerrero A. C. Murillo, J. Kosecka and C. Sagues. Visual door detection integrating appearance and shape cues. *Robotics and Autonomous Systems*, 56(6):pp. 512–521, 2008.
- [2] D. Anderson, H. Herman, and A. Kelly. Experimental characterization of commercial flash lidar devices. In *Proc. of the Int. Conf. of Sensing and Technology*, Palmerston North, New Zealand, 2005.
- [3] A. Andreopoulos and J. K. Tsotsos. Active vision for door localization and door opening using playbot. In *Proc. of the Canadian Conf. on Computer and Robot Vision (CRV)*, pages 3–10, Washington, DC, USA, 2008.
- [4] D. Anguelov, D. Koller, E. Parker, and S. Thrun. Detecting and modeling doors with mobile robots. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, pages 3777–3784, 2004.
- [5] G. Bradski and A. Kaehler. *Learning OpenCV: Computer Vision with the OpenCV Library*. O’Reilly Media, Inc., 2008.
- [6] T. Brox, B. Rosenhahn, J. Gall, and D. Cremers. Combined region- and motion-based 3D tracking of rigid and articulated objects. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2009.
- [7] B. Curless and M. Levoy. Better optical triangulation through spacetime analysis. In *Proc. of the Int. Conf. on Computer Vision (ICCV)*, 1995.
- [8] J. Davis, D. Nehab, R. Ramamoorthi, and S. Rusinkiewicz. Spacetime stereo: a unifying framework for depth from triangulation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(2), February 2005.
- [9] A.Y. Ng E. Klingbeil, A. Saxena. Learning to open new doors. In *Proc. of the RSS Workshop on Robot Manipulation*, Seattle, WA, USA, 2009.
- [10] M. Fiala. Artag, a fiducial marker system using digital techniques. Technical report, National Research Council Canada, 2004.
- [11] A. Jain and C.C. Kemp. Behavior-based door opening with equilibrium point control. In *Proc. of the RSS Workshop on Mobile Manipulation in Human Environments*, Seattle, WA, USA, 2009.
- [12] D. Katz, Y. Pyuro, and O. Brock. Learning to manipulate articulated objects in unstructured environments using a grounded relational representation. In *Robotics: Science and Systems*, Zurich, Switzerland, 2008.
- [13] K. Konolige. Small vision systems: hardware and implementation. In *Proc. of the Int. Symp. on Robotics Research*, pages 111–116, 1997.
- [14] K. Konolige. Projected texture stereo. In *Proc. of the IEEE International Conference on Robotics & Automation (ICRA)*, Anchorage, Alaska, May 2010.
- [15] D. Kragic, L. Petersson, and H. I. Christensen. Visually guided manipulation tasks. *Robotics and Autonomous Systems*, 40(2-3):193 – 203, 2002.
- [16] M. Krainin, P. Henry, X. Ren, and D. Fox. Manipulator and object tracking for in hand model acquisition. In *Proc. of the Workshop on Best Practice in 3D Perception and Modeling for Mobile Manipulation at the Int. Conf. on Robotics & Automation (ICRA)*, Anchorage, Alaska, May 2010.
- [17] J. Lim. Optimized projection pattern supplementing stereo systems. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2009.
- [18] H. K. Nishihara. Prism: A practical real-time imaging stereo matcher. Technical report, MIT, Cambridge, MA, USA, 1984.
- [19] M. Quigley, S. Batra, S. Gould, E. Klingbeil, Q. Le, A. Wellman, and A.Y. Ng. High-accuracy 3d sensing for mobile manipulation: Improving object detection and door opening. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2009.
- [20] D.A. Ross, D. Tarlow, and R.S. Zemel. Unsupervised learning of skeletons from motion. In *Proc. of the European Conf. on Computer Vision (ECCV)*, 2008.
- [21] 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*, 2008.
- [22] J. Salvi, J. Pages, and J. Batlle. Pattern docification strategies in structured light systems. *Pattern Recognition*, 37(4), 2004.
- [23] J. Sturm. Planar objects package. Willow Garage, Robot Operating System, WG-ROS-PKG, SVN Repository, 2009. Available online at https://code.ros.org/svn/wg-ros-pkg/trunk/sandbox/planar_objects, revision 24017.
- [24] J. Sturm, K. Konolige, C. Stachniss, and W. Burgard. Vision-based detection for learning articulation models of cabinet doors and drawers in household environments. In *Proc. of the IEEE International Conference on Robotics & Automation (ICRA)*, Anchorage, Alaska, May 2010.
- [25] J. Sturm, V. Pradeep, C. Stachniss, C. Plagemann, K. Konolige, and W. Burgard. Learning kinematic models of articulated objects. In *Proc. of the Int. Conf. on Artificial Intelligence (IJCAI)*, 2009.
- [26] F. Tombari and K. Konolige. A practical stereo system based on regularization and texture projection. In *Proc. of the Int. Conf. on Informatics in Control, Automation and Robotics (ICINCO)*, Milan, Italy, 2009.
- [27] G. Vosselman and S. Dijkman. 3D building model reconstruction from point clouds and ground plans. *International Archives Photogrammetry and Remote Sensing (IAPRS)*, 34(3W4):37–43, 2001.
- [28] A. Wedel, C. Rabe, T. Vaudrey, T. Brox, U. Franke, and D. Cremers. Efficient dense scene flow from sparse or dense stereo data. In *Proc. of the European Conf. on Computer Vision (ECCV)*, Marseille, France, 2008.
- [29] O. Williams, M. Isard, and J. MacCormick. Estimating disparity and occlusions in stereo video sequences. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, 2005.
- [30] J. Yan and M. Pollefeys. Automatic kinematic chain building from feature trajectories of articulated objects. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2006.
- [31] Y.L. Yufeng, R. Emery, D. Chakrabarti, and W. Burgard. Using EM to learn 3D models of indoor environments with mobile robots. In *Proc. of the Int. Conf. on Machine Learning (ICML)*, pages 329–336, Williamstown, MA, USA, 2001.
- [32] L. Zhang, B. Curless, and S.M. Seitz. Spacetime stereo: shape recovery for dynamic scenes. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, 2003.