

# Team Applied Robotics: A closer look at our robotic picking system

Wim Abbeeloos<sup>1</sup>, Fabian Gouwens, Simon Jansen, Berend Küpers, Maurice Ramaker, Toon Goedemé<sup>1</sup>

**Abstract**—This paper describes the vision based robotic picking system that was developed by our team, Team Applied Robotics, for the Amazon Picking Challenge 2016. This competition challenged teams to develop a robotic system that is able to pick a large variety of products from a shelf or a tote. We discuss the design considerations and our strategy, the high resolution 3D vision system, the use of a combination of texture and shape-based object detection algorithms, the robot path planning and object manipulators that were developed.

## I. INTRODUCTION

There is a strong interest in applying robotics not only in the typical, well structured environments they normally operate in, but also in less structured, real world situations. While there has been a lot of progress in the fields of computer vision, path planning and robotic grasping, combining them into a reliably working system still proves to be challenging. The APC is a competition in the field of warehouse logistics, in which objects need to be picked either from a shelf, or a tote.

The system we built is based on a Universal Robots UR10 robot (Figure 1). The system was outfitted with a custom vision sensor and vacuum gripper. The software ran on a single laptop and made use of the Robot Operating System (ROS) framework. In the following sections we provide more details on the vision system, object detection, path planning and object manipulation.

## II. PERCEPTION

A 3D vision system was developed to detect and determine the 6D pose of the objects that need to be handled.

### A. Time Multiplexed Structured Light

To be able to recognize objects and determine their pose with sufficient accuracy, while keeping system cost low, a custom 3D scanner (Figure 2) was built from off the shelf components. The system is based on triangulation between binary Gray code pattern sequences projected by a projector, and their image as acquired by a camera [1][2]. The camera and projector are synchronized and the patterns can be projected and acquired at 120 frames per second.

Our choice for this custom sensor over a standard depth sensor such as Time of Flight cameras, the Microsoft Kinect, or the Intel RealSense was mostly based on the higher resolution (1140x912 pixels) and accuracy (0.1mm). The biggest downside is that it requires multiple images (42), and



Fig. 1: The developed picking system consists of a standard six degree of freedom robot equipped with a custom made 3D vision system and vacuum gripper.

some extra processing. This results in a longer acquisition time (about one second for acquisition and processing).

### B. Sensor Calibration

To be able to triangulate points, both the camera and the projector need to be calibrated. The camera is calibrated in the typical way, using a chessboard pattern. Note that while the projector can be modeled as a camera, it is not possible to measure the location of the checkerboard corners in the projector reference frame directly. The method proposed by Moreno et al. [3] was used to estimate the position of the corners in the projected image. This method uses the decoded pattern as observed by the camera and creates a local homography to estimate the code that can be associated to the checkerboard corner. This allows the projector to be calibrated, and to calibrate the camera-projector setup as a stereo camera.

## III. DETECTION

A set of 38 products of varying shape, appearance, material and weight was used in the challenge. To cope with this variety, we chose to employ multiple detection algorithms, and apply the most suitable for every object. One object was searched for at a time, starting with the object that was deemed easiest, according to an ordered list that was created manually (taking into account ease of detection and ease of manipulation). Also, for every object, a preferred object detection method was determined in advance.

For picking from the shelf, three scans were taken from different angles. For picking from the tote, two scans were taken from different angles. After picking it from the tote, the item was placed on a table and an additional scan was taken in order to confirm it was the correct object.

<sup>1</sup>Wim Abbeeloos and Toon Goedemé are with KU Leuven, Department of Electrical Engineering, Leuven, Belgium

The authors thank Alten, KU Leuven and Smart Robotics for their support.

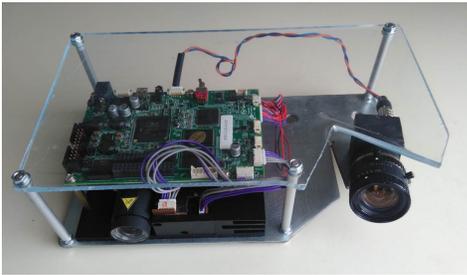


Fig. 2: The 3D vision system consists of a Digital Micromirror Device projector (TI DLP lightcrafter 4500) and a monochrome camera (Pointgrey Flea3 FL3-U3-13Y3) equipped with a c-mount lens (Computar 8mm F1.4 M0814-MP2). The mount was made from a bent steel plate and attaches directly to the robot wrist. A protective acrylic top shields the components. The projector and camera are synchronized to ensure the camera captures the correct projected patterns.

#### A. Pre-processing

As the shelf and tote geometry were known, it is possible to segment the objects from the background. While the exact position of the tote was known, the shelf was slightly moved before the start of the competition. A calibration routine was used at the start of the competition in which the shelves top corners were scanned (Figure 3) and their exact position was measured, allowing to determine the shelf pose.

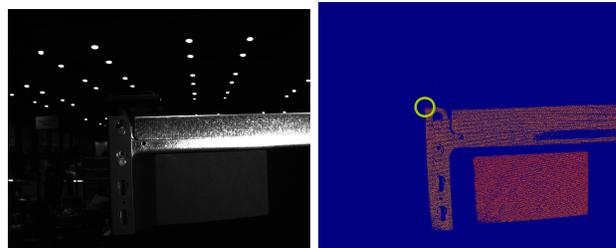
#### B. Detection Algorithms

1) *Point Pair Features*: Some of the objects have a distinctive geometric shape. Point pair features (PPF, Figure 4) can be used to describe this shape [4][5][6]. While a single feature is not very descriptive, the features of all combinations of points on the objects surface typically is.

From a measured scene, a subset of points is selected, and the PPFs of all their combinations are calculated. If a similar PPF is present in the objects model, the PPF votes for a 6D pose of the object. If a pose gets enough votes, it is accepted as a detection. The initial detection is followed by an Iterative Closest Point (ICP) procedure to obtain a refined object pose.

2) *2D Features*: The objects with distinctive texture can be detected using a local features based approach [7][8], in our case, SIFT is used as a feature descriptor. From these matches between the object template and scene (Figure 5) the pose can be estimated using a Perspective-n-Point algorithm, in this case EPnP [9] was used. Note that in our case, the exact 3D location of all features is measured with the structured light scanner. This additional information allows to estimate the pose with much higher accuracy, and allows to eliminate incorrect matches.

3) *Range Image Templates*: The measured pointcloud can also be represented as a range image (Figure 6). If the above algorithms fail, a brute force approach can be used in which templates obtained from the model are compared to this range image [10]. We use 1024 templates per object. Each



(a) Intensity image

(b) Range image

Fig. 3: A scan of the top corners of the shelf is used to measure the pose of the shelf with respect to the robot. The detected corner is shown in the range image with a green circle.

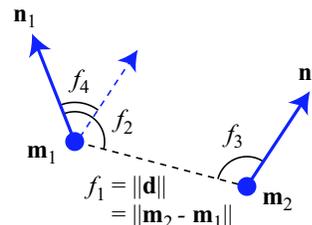


Fig. 4: The four dimensional point pair feature vector  $\mathbf{f}$  of two 3D points on the surface of an object:  $m_i$ , and their normals  $n_i$ . The first element,  $f_1$  is a distance, while  $f_2$ ,  $f_3$  and  $f_4$  are angles.

of the templates is aligned to the scene in an optimization procedure. To initialize the optimization, the scene is first segmented using euclidean clustering, and the closest point of each of the clusters can be used as a starting point, which is initially aligned to the closest point in the range template to be optimized. If a template with sufficiently low cost is found, an extra ICP step is used to refine the object pose.

#### IV. MANIPULATION

Two suction based grippers were custom designed (Figure 7): one for picking from the shelf and one for picking from the tote. They both consist of pieces of standard aluminum tube and parts made using an additive manufacturing process (fused deposition modeling).

Both grippers use high flow vacuum generated by a modified 2000W vacuum cleaner. The air flow is guided through flexible tubing. A small circular piece of soft foam was glued to the grippers to provide better sealing of the gripper to the object.

When grasping an object, the end effector was lowered until a certain force threshold was met. This increased the chance of grasping the object properly.

#### V. PLANNING

Once an object is detected the object needs towards it to grasp it. The grasping locations were manually defined per object. Most of the robot movements were planned in advance, offline, in joint space. Only the final

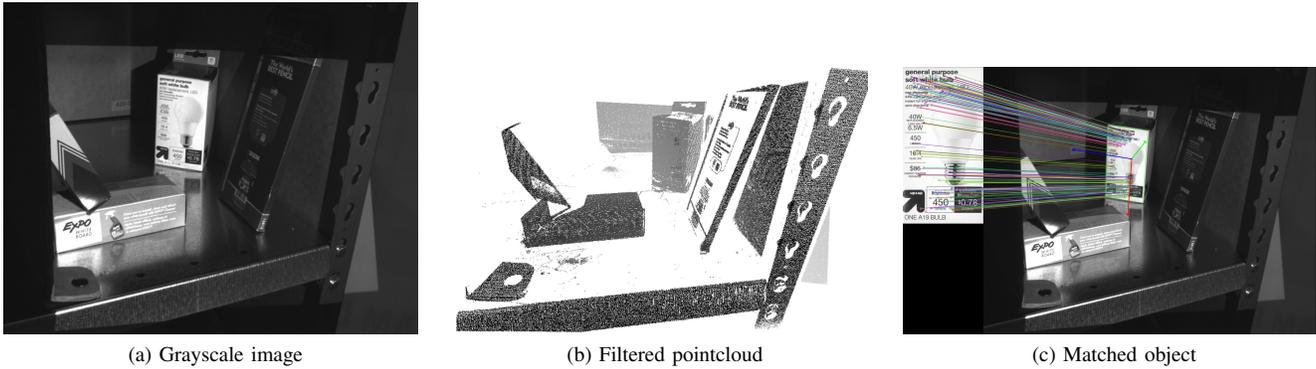


Fig. 5: Local appearance features of the objects are matched to the scene, allowing to estimate the objects pose. Using the 3D location of the feature point, from the pointcloud, results in higher accuracy and allows to eliminate incorrect matches.

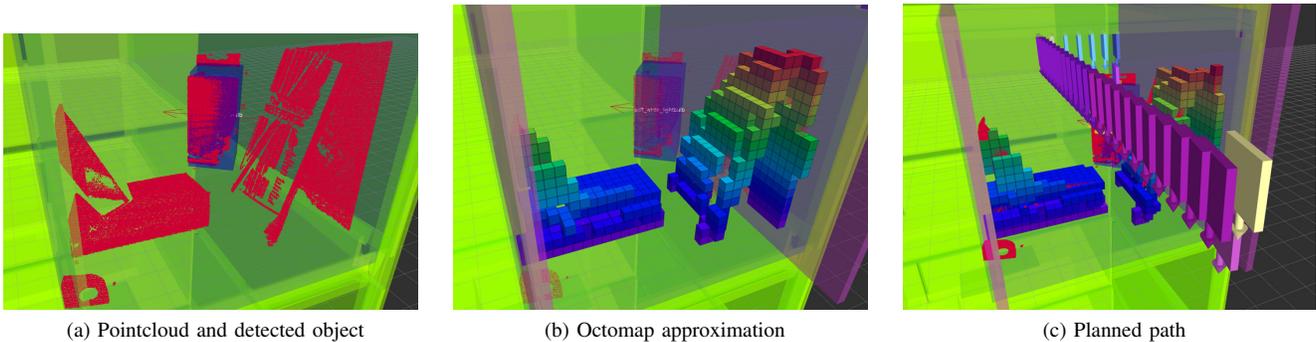


Fig. 8: When an object is detected, a collision map is calculated using an octomap approximation of the pointcloud. Note that the points on the objects are not included. These points are cut out of the pointcloud using a slightly expanded version of the convex hull of the object. A cartesian path planning algorithm is used to search for a collision free path to remove the object from the shelf.

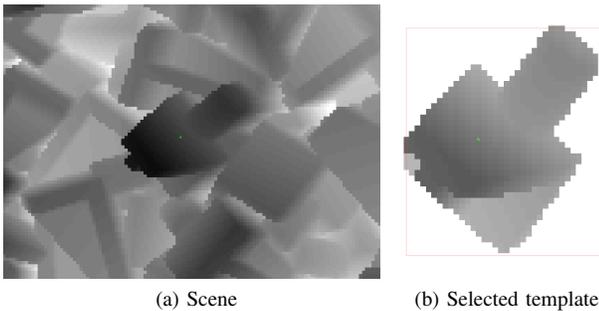


Fig. 6: A set of 1024 templates (range images from different viewpoints) are aligned to the scene using an optimization function. The final pose of the template with the lowest cost is used if it is below a certain threshold.

part of the paths toward the object and out of the shelf need to be planned online. This was done using a cartesian path planning, which creates a tree structure of cartesian movements, with a resolution of 2cm. A greedy search algorithm was used and a check for collisions performed at every node (Figure 8). For removing objects from the tote, no path planning was used and the object was simply lifted vertically.

## VI. CONCLUSIONS

Our approach to the challenges met in the APC was summarized in this paper. The proposed system has some unique features that were not employed by any other teams. Our system also suffered from some drawbacks, one of the major limitations during the challenge was that the 3D scan acquisition speed which had to be lowered due to technical issues. Another limitation was that not all of the objects could be detected reliably. The vacuum detection did not work properly, and could not be used to tell whether an object was successfully picked. Despite these issues, our system was able to pick multiple items correctly in both of the challenges.

## REFERENCES

- [1] S. Inokuchi, K. Sato, and F. Matsuda, "Range imaging system for 3-d object recognition," in *Proceedings of the International Conference on Pattern Recognition*, vol. 48, 1984, pp. 806–808.
- [2] J. L. Posdamer and M. Altschuler, "Surface measurement by space-encoded projected beam systems," *Computer graphics and image processing*, vol. 18, no. 1, pp. 1–17, 1982.
- [3] D. Moreno and G. Taubin, "Simple, accurate, and robust projector-camera calibration," in *Proc. Int'l Conf. 3D Imaging, Modeling, Processing, Visualization and Transmission (3DIMPVT)*, 2012, pp. 464–471.
- [4] W. Abbeels and T. Goedemé, "Point pair feature based object detection for random bin picking," in *Int'l Conf. on Computer and Robot Vision (CRV)*, June 2016, pp. 432–439.

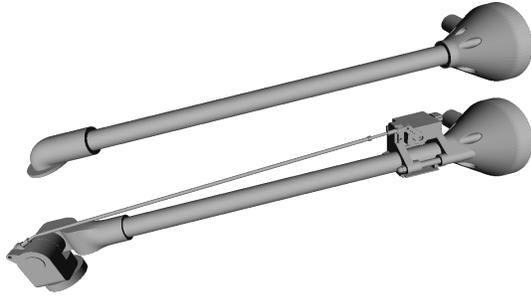


Fig. 7: **Top:** The gripper used for picking from the shelf. **Bottom:** The gripper used for picking from the tote has an extra degree of freedom. The 3D printed joint is designed to have a bit of play, but seals itself when the vacuum is turned on, so that it does not leak air. The joint can be actuated by a servo motor via a bar linkage. The servo motor is controlled by a micro controller with an accelerometer and has three modes: unactuated, fixed at a certain angle, or set to keep the head vertical.

- [5] C. Choi, Y. Taguchi, O. Tuzel, M.-Y. Liu, and S. Ramalingam, "Voting-based pose estimation for robotic assembly using a 3d sensor," in *Proc. IEEE Int'l Conf. Robotics and Automation (ICRA)*. IEEE, 2012, pp. 1724–1731.
- [6] E. Wahl, U. Hillenbrand, and G. Hirzinger, "Surflet-pair-relation histograms: a statistical 3d-shape representation for rapid classification," in *Proc. Int'l Conf. 3-D Digital Imaging and Modeling (3DIM)*. IEEE, 2003, pp. 474–481.
- [7] A. Collet, D. Berenson, S. S. Srinivasa, and D. Ferguson, "Object recognition and full pose registration from a single image for robotic manipulation," in *Proc. IEEE Int'l Conf. Robotics and Automation (ICRA)*. IEEE, 2009, pp. 48–55.
- [8] T. Grundmann, R. Eidenberger, M. Schneider, M. Fiebert, and G. v. Wichert, "Robust high precision 6d pose determination in complex environments for robotic manipulation," in *Proc. Workshop Best Practice in 3D Perception and Modeling for Mobile Manipulation at the Int'l Conf. Robotics and Automation (ICRA)*, 2010, pp. 1–6.
- [9] V. Lepetit, F. Moreno-Noguer, and P. Fua, "Épnp: An accurate o (n) solution to the pnp problem," *Int'l J. Computer Vision*, vol. 81, no. 2, pp. 155–166, 2009.
- [10] M. Germann, M. D. Breitenstein, I. K. Park, and H. Pfister, "Automatic pose estimation for range images on the gpu," in *Proc. Int'l Conf. 3-D Digital Imaging and Modeling (3DIM)*. IEEE, 2007, pp. 81–90.