

3D Robot Perception with Point Cloud Library

1. Introduction

3D perception allows robot to reach levels of autonomy and safety that are not possible with 2D vision alone. Since 2010, 3D depth data became readily available thanks to the introduction of new sensors, such as the Microsoft Kinect, that could densely estimate depth at low cost and couple it with color or infrared data. The robotics community – already used to 3D information estimated with lasers and stereo cameras – adopted the new RGB-D sensors very quickly and promoted a big part of research on exploiting combined depth and color information for improving the performance of many robotic perception tasks. Together with the diffusion of these novel sensors, the need for new tools and software libraries that could help in handling three-dimensional data rapidly emerged. It was also important to standardize the representation of depth data and depth+color data, in order, for example, to easily compare data coming from different sensors.

For this purpose, in 2011, the Point Cloud Library (or PCL) was created. PCL is a large scale, open project for 2D/3D image and point cloud processing. Since 2011, PCL grew with an exceptional trend thanks to the contributions of more than 1,000 contributors from more than 100 institutions all over the world. This growth led PCL to become a reference library for 3D processing, providing numerous state-of-the art algorithms for filtering, feature estimation, surface reconstruction, registration, model fitting and segmentation and drivers for the most diffused 3D sensors. Moreover, complete pipelines needed for solving a number of complex perception problems are also present, such as object recognition and tracking, people detection and large-scale 3D reconstruction. PCL's open BSD license, free for commercial and research use, favored its quick spread among both academia and industry. The tools and algorithms within PCL provide a wide basic infrastructure which allows researchers and companies to better concentrate on their specific areas of expertise and product development, without the need of re-implementing all the parts of a complex system. Since the very beginning, PCL has been directly integrated as part of the

Robot Operating System, being the backbone of most of perception algorithms for robots equipped with RGB-D sensors and thus boosting research and development in this fundamental field.

2. In this issue

This Special Issue on 3D Robot Perception with Point Cloud Library stems from a workshop that was organized in conjunction with the 13th International Conference on Autonomous Intelligent Systems (IAS-13), held in Padova (Italy), on July 15th, 2014. The purpose of this workshop was to bring together researchers presenting novel 3D perception algorithms and 3D systems for mobile robots built on the Point Cloud Library, in order to create a forum to discuss advances in 3D robot perception and the role that the Point Cloud Library had, has and will have, in robotics research and industrial applications. Due to the success of the workshop and the strong interest in the topic of robot perception, we promoted this Special Issue, which was open not only to papers presented at the workshop, but also to contributions not previously published. All papers were then carefully peer-reviewed and a brief introduction to the accepted papers follows.

This special issue starts with a metrological characterization process for time-of-flight (TOF) cameras. In particular, Corti et al. apply this process to the Microsoft Kinect V2, one of the most recently released and promising RGB-D sensors that can be used for robotic perception. This work points out that performances are highly influenced by measuring conditions and environmental parameters of the scene.

When dealing with a mobile robot equipped with such sensors, the huge amount of 3D information that is collected is usually difficult to manage due to the fact that the robot storage system and computing capabilities are insufficient. Therefore, a data compression method is necessary to store and process this information while preserving as much information as possible. Data compression is also useful to efficiently transmit these data to other robots or to a central server for further process-

ing. In the work of Navarrete et al., the authors propose a dataset composed of a set of 3D point clouds with different structure and texture variability to evaluate the results obtained from 3D data compression methods and provide useful tools for comparing different methods.

One fundamental task in 3D robot perception is the ability to model the raw data, for example for the purpose of object reconstruction, recognition and grasping. Mörwald et al. present an algorithm for robustly approximating the boundary of a domain, latent in a planar set of scattered points, by a B-spline curve. Additional to the scientific contribution, they integrated their code into PCL and created a tutorial that guides through the steps of the algorithm. Büttner et al., instead, propose to model objects by means of primitive 3D shapes (cylinders, spheres, cones, etc.), thus creating an estimate of the invisible back sides. Their method detects symmetrical parts of an object, model these symmetrical parts as primitive shapes and fit them to noisy sensor data using sample consensus methods. For the purpose of object grasping, this has the advantage that feasible grasps can be chosen from a precomputed set based on the estimated model, instead of a time-consuming random sampling approach. The authors extended state of the art methods from PCL to include additional relevant shapes (e.g. boxes), constraints (e.g. on size and orientation), and to consider additional information like knowledge about free space or proprioceptive information.

For recognizing an object and its pose, an alternative approach to fitting primitive shapes consists in extracting and matching local feature descriptors. However, the research on 3D local descriptors is still at an early stage. Logoglu et al. describe a new local descriptor, Colored Histograms of Spatial Concentric Surflet-Pairs (CoSPAIR), that encodes both depth and color information and outperforms the state-of-the-art descriptors in both category-level and instance-level recognition task. The proposed descriptor is compared against the state-of-the-art local 3D descriptors that are available in PCL and their object recognition performances are evaluated on several publicly available datasets. We conclude the group of papers on object recognition and pose estimation with the work by Choi and Christensen, who present an object pose estimation approach exploiting both geometric depth and photometric color information available from an RGB-D sensor. In contrast to various efforts relying on object segmentation with a known background structure, their approach does not depend on the segmentation and thus exhibits superior performance in unstructured environments. Inspired by a voting-based approach employing an oriented point pair

feature, they present a voting-based approach which further incorporates color information from the RGB-D sensor and which exploits parallel power of the modern parallel computing architecture.

The availability of dense 3D data allowed to considerably improve also mapping and localization algorithms. Moral et al. present a flexible strategy to register scenes based on their planar structure. The proposed strategy is based on the segmentation of the planar surfaces from the scene, and its representation using a graph which stores the geometric relationships between neighboring planar patches. Quick registration is achieved in indoor structured scenarios, offering advantages like a compact representation, and flexibility to adapt to different environments and sensors. The authors provided also an implementation of their method as a module of PCL.

For the purpose of visual odometry, Gutierrez-Gomez et al. propose to parametrize the geometric error by the inverse depth and switch keyframes based on a visibility criteria between frames. They implemented this method within the scope of PCL as a branch of the code for KinFu Large Scale, an open source implementation of KinectFusion, where the original ICP system for odometry estimation has been completely substituted by their method. Moreover, Correia da Costa et al. describe a modular localization system suitable for a wide range of mobile robots and based on LIDAR / RGB-D data. It demonstrated high accuracy while performing pose tracking with point cloud registration algorithms and high reliability when estimating the initial pose using feature matching techniques. Also this algorithm has been made publicly available.

The semantic localization problem in robotics consists in determining the place where a robot is located by means of semantic categories. In the paper of Martinez et al., the authors propose a framework, implemented within PCL, which provides a set of valuable tools to easily develop and evaluate semantic localization systems. The implementation includes the generation of 3D global descriptors following a Bag-of-Words approach. This allows the generation of fixed-dimensionality descriptors from any type of keypoint detector and feature extractor combinations.

When a user and a robot share the same physical workspace the robot may need to keep an updated 3D representation of the environment and focus on the parts where the user executes manipulation tasks. Monica et al. propose to drive the robot behavior by a next-best view algorithm that computes the most promising range sensor viewpoints to observe the detected salient regions, where potential changes in the environment have occurred. The environment representation is built upon

the PCL KinFu Large Scale project, that has been parallelized to exploit processing on GPUs.

For robots having to interact with humans, in domestic or industrial environments, it is fundamental to be able to detect, track and distinguish people. On top of PCL, OpenPTrack was specifically built for this, as described in the paper of Munaro et al.. OpenPTrack is an open source software for multi-camera calibration and people tracking in RGB-D camera networks. It allows to track people in big volumes at sensor frame rate and currently supports a heterogeneous set of 3D sensors. In the paper of this issue, the authors describe its user-friendly calibration procedure, which consists of simple steps with real-time feedback that allow to obtain accurate results in estimating the camera poses that are then used for tracking people. Moreover, they detail how a cascade of algorithms working on depth point clouds and color, infrared and disparity images is used to perform people detection from different types of sensors and in any indoor light condition.

For what concerns the task of modeling and recognizing people from 3D data, Tombari and Ichim propose an original approach for fitting a pre-defined parametric shape model to depth data by exploiting the 3D body pose tracked through a sequence of range maps. To this goal, they make use of multiple types of constraints and cues embedded into a unique cost function, which is then efficiently minimized. Their algorithm relies on the tools and algorithms provided by PCL, representing a good integration of the functionalities available therein.

We chose to conclude this special issue with a case study on 3D robot perception that heavily relies on the PCL library. The paper of Oliveira et al. describes a 3D object perception and perceptual learning system developed for a complex artificial cognitive agent working in a restaurant scenario. This system, developed within the scope of the European project RACE, integrates detection, tracking, learning and recognition of tabletop objects. This work shows the current importance of PCL in building sophisticated 3D perception systems in robotics and other domains.

We hope that the ideas, tools and open source software presented by the papers collected in this special issue can seed new research ideas and serve as a basis for solving more and more complex 3D robot perception problems. We also hope that the readers of this special issue could consider to contribute back part of their work to PCL, in order to foster the growth of this important open source ecosystem.

Acknowledgements

We wish to thank all the authors and reviewers of the papers in this special issue for their great work in improving, polishing, and commenting the papers you can finally find in the next pages. We also wish to thank Prof. Rüdiger Dillmann, the Editor in Chief of Robotics and Autonomous Systems, for providing us the opportunity to publish this special issue and Christine Brand, Hilda Hu and Mathangi Venkatesan for the help provided during our editorial work.

Matteo Munaro¹

*Intelligent Autonomous Systems Laboratory,
Department of Information Engineering (DEI), School
of Engineering, Universita' degli Studi di Padova, Via
Gradenigo 6/a, I-35131 Padua, Italy
E-mail address: matteo.munaro@dei.unipd.it.*

Radu B. Rusu

*Open Perception, Inc., 68 Willow Road, Menlo Park,
CA 94025, USA
E-mail address: rusu@openperception.org.*

Emanuele Menegatti

*Intelligent Autonomous Systems Laboratory,
Department of Information Engineering (DEI), School
of Engineering, Universita' degli Studi di Padova, Via
Gradenigo 6/a, I-35131 Padua, Italy
E-mail address: emg@dei.unipd.it.*

¹Corresponding editor.