

Preface

This work provides an introduction to the foundations of three-dimensional computer vision and describes recent contributions to the field, which are methodical and application-specific in nature. Each chapter of this work provides an extensive overview of the corresponding state of the art, into which a detailed description of new methods or evaluation results in application-specific systems is embedded.

Triangulation-based approaches to three-dimensional scene reconstruction (cf. Chap. 1) are primarily based on the concept of bundle adjustment, which has been developed in the domain of photogrammetry. The three-dimensional scene structure and the intrinsic and extrinsic camera parameters are determined such that the Euclidean backprojection error in the image plane is minimised, usually relying on a nonlinear optimisation procedure. In the field of computer vision, an alternative framework based on projective geometry has emerged during the last two decades, which allows one to use linear algebra techniques for three-dimensional scene reconstruction and camera calibration purposes. With special emphasis on the problems of stereo image analysis and camera calibration, these fairly different approaches are related to each other in the presented work, and their advantages and drawbacks are stated. In this context, various state-of-the-art camera calibration and self-calibration methods as well as recent contributions towards automated camera calibration systems are described. An overview of classical and new feature-based, correlation-based, dense, and spatio-temporal methods for establishing point correspondences between pairs of stereo images is given. Furthermore, an analysis of traditional and newly introduced methods for the segmentation of point clouds and for the three-dimensional detection and pose estimation of rigid, articulated, and flexible objects in the scene is provided (cf. Chap. 2).

A different class of three-dimensional scene reconstruction methods consists of intensity-based approaches (cf. Chap. 3), which evaluate the pixel grey values in the image to infer the three-dimensional scene structure. Basically, these methods can be divided into shape from shadow, photoclinometry and shape from shading, photometric stereo, and shape from polarisation. As long as sufficient information about the illumination conditions and the surface reflectance properties is available, these methods may provide dense depth maps of object surfaces.

In a third, fundamentally different class of approaches the behaviour of the point spread function of the optical system used for image acquisition is exploited in order to derive depth information about the scene (cf. Chap. 4). Depth from defocus methods determine the position-dependent point spread function, which in turn yields absolute depth values for the scene points. A semi-empirical framework for establishing a relation between the depth of a scene point and the observed width of the point spread function is introduced. Depth from focus methods use as a reference the distance between the camera and the scene at which a minimum width of the point spread function is observed, relying on an appropriate calibration procedure.

These three classes of approaches to three-dimensional scene reconstruction are characterised by complementary properties; thus it is favourable to integrate them into unified frameworks that yield more accurate and robust results than each of the approaches alone (cf. Chap. 5). Bundle adjustment and depth from defocus are combined to determine the absolute scale factor of the scene reconstruction result, which cannot be obtained by bundle adjustment alone if no a priori information is available. Shading and shadow features are integrated into a self-consistent framework to reduce the inherent ambiguity and large-scale inaccuracy of the shape from shading technique by introducing regularisation terms that rely on depth differences inferred from shadow analysis. Another integrated approach combines photometric, polarimetric, and sparse depth information, yielding a three-dimensional reconstruction result which is equally accurate on both large and small scales. An extension of this method provides a framework for stereo image analysis of non-Lambertian surfaces, where traditional stereo methods tend to fail. Furthermore, a method is proposed to integrate photometric information and absolute depth data acquired using an active range scanning device. In the context of monocular three-dimensional pose estimation, the integration of triangulation, photopolarimetric, and defocus cues is demonstrated to behave more robustly and to provide significantly more accurate results than techniques exclusively relying on triangulation-based information.

The developed three-dimensional scene reconstruction methods are examined in different application scenarios. A comparison to state-of-the-art systems is provided where possible. In the context of industrial quality inspection (cf. Chap. 6), the performance of pose estimation is evaluated for rigid objects (e.g. plastic caps and electric plugs) as well as flexible objects (e.g. tubes and cables). The integrated surface reconstruction methods are applied to the inspection of different kinds of metallic surfaces.

The developed techniques for object detection and tracking in three-dimensional point clouds and for pose estimation of articulated objects are evaluated in the context of partially automated industrial production scenarios requiring a safe interaction between humans and industrial robots (cf. Chap. 7). Furthermore, we determine how the developed three-dimensional detection and pose estimation techniques are related to state-of-the-art gesture recognition methods in human-robot interaction scenarios, and typical action recognition results are presented in a realistic industrial scenario.

The third addressed application scenario is completely different and regards remote sensing of the lunar surface by preparing elevation maps (cf. Chap. 8).

While the spatial scales involved differ by many orders of magnitude from those encountered in the industrial quality inspection domain, the underlying physical processes are fairly similar. An introductory outline of state-of-the-art approaches to topographic mapping of solar system bodies is given. The estimation of impact crater depths and shapes is an issue of especially high geological relevance. It is demonstrated for lunar craters that three-dimensional surface reconstruction based on integrated methods yields topographic maps of high resolution, where at several places a comparison with recent topographic data (orbital laser altimetry and stereophotogrammetry) is performed. Another geologically relevant field is the three-dimensional reconstruction of lunar volcanic edifices, especially lunar domes.

Finally (cf. Chap. 9), the main results of the presented work and the most important conclusions are summarised, and possible directions of future research are outlined.

Dortmund, Germany

Christian Wöhler



<http://www.springer.com/978-1-4471-4149-5>

3D Computer Vision
Efficient Methods and Applications
Wöhler, C.
2013, XVIII, 382 p., Hardcover
ISBN: 978-1-4471-4149-5