Preface

The fully automatic processing and analysis of 3D point clouds represents a topic of major interest in the fields of photogrammetry, remote sensing, computer vision, and robotics. Focusing on general applicability and considering a typical end-to-end processing workflow from raw 3D point cloud data to semantic objects in the scene, we introduce a novel and fully automated framework involving a variety of components which cover (i) the filtering of noisy data, (ii) the extraction of appropriate features, (iii) the adequate alignment of several 3D point clouds in a common coordinate frame, (iv) the enrichment of 3D point cloud data with other types of information, and (v) the semantic interpretation of 3D point clouds. For each of these components, we reflect the fundamentals and related work, and we further provide a comprehensive description of our novel approaches with significant advantages. Based on a detailed evaluation of our framework on various benchmark datasets, we provide objective results for different steps in the processing workflow and demonstrate the performance of the proposed approaches in comparison to state-of-the-art techniques.

Target Audience

Although this book has been written as research monograph, the target audience includes a broad community of people who are dealing with 3D point cloud processing, reaching from students at undergraduate or graduate level to lecturers, practitioners and researchers in photogrammetry, remote sensing, computer vision and robotics. Background knowledge on 3D point cloud processing will certainly improve the reader's appreciation but is not mandatory since reviews on fundamentals and related work are provided for the different components of an end-to-end processing workflow from raw 3D point cloud data to semantic objects in the scene.

Difficulty

Beginner to intermediate and expert. This book is intended to address a broad community by reflecting the fundamentals and related work on the different components of an end-to-end processing workflow from raw 3D point cloud data to semantic objects in the scene. Additionally, a comprehensive description of novel approaches with significant advantages is provided, whereby a detailed evaluation of the proposed approaches in comparison to state-of-the-art techniques is carried out on various benchmark datasets. In this regard, beginners may obtain an impression about the performance of different approaches with respect to different evaluation criteria, whereas experts in the related fields may appreciate the novel approaches and be inspired to conduct further investigations.

Organization of the Content

This book addresses the reconstruction and analysis of 3D scenes and thereby focuses on the presentation of an end-to-end processing workflow from raw 3D point cloud data to semantic objects in the scene. To a reasonable extent, all the material about one topic is presented together in a separate chapter, making the book suitable as a reference work as well as a tutorial on different subtopics.

In Chap. 1, we provide an introduction which focuses on explaining the goals of this work, the main challenges when addressing these goals and the scientific contributions developed in this context. We additionally provide a list of our publications, take a glance view on our novel framework for advanced 3D point cloud processing from raw data to semantic objects and finally give an overview on the topics presented in the whole book.

In Chap. 2, we consider fundamentals in the form of general definitions that will be important for a consideration of our work in the global context. Since, in this work, we focus on the use of 3D point cloud data as input for our methodology, we briefly explain how such 3D point cloud data may be acquired indirectly from 2D imagery via passive techniques and how it may be acquired directly by involving active techniques. Particularly the active techniques offer many advantages and we therefore consider respective systems for data acquisition in the scope of this book. An analysis of such acquisition systems reveals that the acquired 3D point cloud data may easily be represented in the form of range and intensity images which, in turn, might facilitate a variety of applications in terms of simplicity and efficiency. In this regard, we consider point quality assessment where the aim is to quantify the quality of the single range measurements captured with the involved acquisition system, and we present and discuss two novel approaches and their potential in comparison to related work.

In Chap. 3, we briefly explain fundamental concepts for extracting features from 2D imagery and 3D point cloud data. In this regard, we derive a definition for the general term of a *feature* and consider different types of features that might occur in 2D imagery or 3D point cloud data. These fundamentals have become the basis and core of a plethora of applications in photogrammetry, remote sensing, computer vision and robotics, and they also provide the basis for the subsequent chapters.

In Chap. 4, we focus on transforming 3D point clouds—where each 3D point cloud has been acquired with respect to the local coordinate frame of the acquisition system—into a common coordinate frame. We consider related work and thereby particularly focus on efficient approaches relying on the use of keypoints. Based on these considerations, we present a novel framework for point cloud registration which exploits the 2D representations of acquired 3D point cloud data in the form of range and intensity images in order to quantify the quality of respective range measurements and subsequently reliably extract corresponding information between different scans. In this regard, we may consider corresponding information in the respective image representations or in 3D space, and we may also introduce a weighting of derived correspondences with respect to different criteria. Thus, the correspondences may serve as input for the registration procedure for which we present two novel approaches. An evaluation on a benchmark dataset clearly reveals the performance of our framework, and we discuss the derived results with respect to different aspects including registration accuracy and computational efficiency.

In Chap. 5, we focus on the enrichment of acquired 3D point clouds by additional information acquired with a thermal camera. We reflect related work and present a novel framework for thermal 3D mapping. Our framework involves a radiometric correction, a geometric calibration, feature extraction and matching, and two approaches for the co-registration of 3D and 2D data. For the example of an indoor scene, we demonstrate the performance of our framework in terms of both accuracy and applicability, and we discuss the derived results with respect to potential use-cases.

In Chap. 6, we focus on 3D scene analysis where the aim is to uniquely assign each 3D point of a given 3D point cloud a respective (semantic) class label. We provide a survey on related work and present a novel framework addressing neighborhood selection, feature extraction, feature selection and classification. Furthermore, we consider extensions of our framework toward large-scale capability and toward the use of contextual information. In order to demonstrate the performance of our framework, we provide an extensive experimental evaluation on publicly available benchmark datasets and discuss the derived results with respect to classification accuracy, computational efficiency and applicability of involved methods.

In Chap. 7, we summarize the contents of the book and provide concluding remarks as well as suggestions for future work.

Abstract

The fully automatic processing and analysis of 3D point clouds represents a topic of major interest in the fields of photogrammetry, remote sensing, computer vision, and robotics. Focusing on general applicability and considering a typical end-to-end processing workflow from raw 3D point cloud data to semantic objects in the scene, we introduce a novel and fully automated framework involving a variety of components which cover (i) the filtering of noisy data, (ii) the extraction of appropriate features, (iii) the adequate alignment of several 3D point clouds in a common coordinate frame, (iv) the enrichment of 3D point cloud data with other types of information, and (v) the semantic interpretation of 3D point clouds. For each of these components, we reflect the fundamentals and related work, and we further provide a comprehensive description of our novel approaches with significant advantages. Based on a detailed evaluation of our framework on various benchmark datasets, we provide objective results for different steps in the processing workflow and demonstrate the performance of the proposed approaches in comparison to state-of-the-art techniques.

In particular, our derived results reveal that (i) our presented point quality measures allow an appropriate filtering of noisy data and have a positive impact on the automated alignment of several 3D point clouds in a common coordinate frame, (ii) the extraction of appropriate features improves the automated alignment of several 3D point clouds with respect to accuracy and efficiency, and even allows a co-registration of 3D and 2D data acquired with different sensor types, (iii) our presented strategies for keypoint-based point cloud registration in terms of either projective scan matching or omnidirectional scan matching allow a highly accurate alignment of several 3D point clouds in a common coordinate frame, (iv) our presented strategies in terms of a RANSAC-based homography estimation and a projective scan matching allow an appropriate co-registration of 3D and 2D data acquired with different sensor types, and (v) our strategy for increasing the distinctiveness of low-level geometric features via the consideration of an optimal neighborhood for each individual 3D point and our strategy for only selecting compact and robust subsets of relevant and informative features have a significantly beneficial impact on the results of 3D scene analysis. Thus, our novel approaches allow an efficient reconstruction and analysis of large 3D environments up to city scale, and they offer a great potential for future research.

Kurzfassung

Die automatische Verarbeitung und Analyse von 3D-Punktwolken stellt in den Bereichen der Photogrammetrie, Fernerkundung, Computer Vision und Robotik ein wichtiges Thema dar. Im Hinblick auf eine allgemeine Anwendbarkeit wird in der vorliegenden Arbeit eine neue und vollautomatisierte Methodik vorgestellt, welche

die wesentlichen Schritte von der Erfassung von 3D-Punktwolken bis hin zur Ableitung von semantischen Objekten in der Szene betrachtet. Diese Methodik umfasst verschiedene Komponenten, welche (i) die Filterung von verrauschten Daten, (ii) die Extraktion von geeigneten Merkmalen, (iii) die angemessene Ausrichtung von mehreren einzelnen 3D-Punktwolken in einem gemeinsamen Koordinatensystem, (iv) die Anreicherung von 3D-Punktwolken mit zusätzlicher Information und (v) die semantische Interpretation von 3D-Punktwolken umfassen. Für jede Komponente werden die Grundlagen sowie der aktuelle Stand der Forschung aufgezeigt. Ferner werden die im Rahmen dieser Arbeit entwickelten Verfahren mit deutlichen Vorteilen gegenüber den bisherigen Verfahren genauer beleuchtet. Basierend auf einer umfassenden Auswertung auf verschiedenen Standard-Datensätzen werden objektive Ergebnisse für verschiedene Schritte in der Datenverarbeitung präsentiert und die Leistungsfähigkeit der entwickelten Methodik im Vergleich zu Standard-Verfahren verdeutlicht.

Im Speziellen zeigen die im Rahmen der vorliegenden Arbeit erzielten Ergebnisse, (i) dass die entwickelten Qualitätsmaße eine angemessene Filterung von verrauschten Daten ermöglichen und sich positiv auf die automatische Ausrichtung von mehreren einzelnen 3D-Punktwolken in einem gemeinsamen Koordinatensystem auswirken, (ii) dass die Extraktion von geeigneten Merkmalen bei der automatischen Ausrichtung von mehreren 3D-Punktwolken sowohl die Genauigkeit als auch die Effizienz der getesteten Verfahren verbessert und sogar eine Ko-Registrierung von 3D- und 2D-Daten, welche mit verschiedenen Sensortypen erfasst wurden, ermöglicht, (iii) dass die vorgestellten Strategien zur Punkt-basierten Registrierung von 3D-Punktwolken über ein projektives Scan Matching und ein omnidirektionales Scan Matching zu einer sehr genauen automatischen Ausrichtung von einzelnen 3D-Punktwolken in einem gemeinsamen Koordinatensystem führen, (iv) dass die vorgestellten Strategien einer RANSAC-basierten Homographie-Schätzung und eines projektiven Scan Matchings für eine angemessene Ko-Registrierung von 3D- und 2D-Daten, welche mit verschiedenen Sensortypen erfasst wurden, geeignet sind und (v) dass die vorgestellte Strategie zur Erhöhung der Einzigartigkeit von einfachen geometrischen Merkmalen über die Betrachtung einer optimalen Nachbarschaft für jeden individuellen 3D-Punkt sowie die vorgestellte Strategie zur Selektion einer kompakten und robusten Untermenge von relevanten und informativen Merkmalen einen signifikanten, positiven Einfluss auf die Ergebnisse einer 3D-Szenenanalyse haben. Auf diese Weise ermöglichen die entwickelten Verfahren eine effiziente Rekonstruktion und Analyse von großen Bereichen bis auf Stadtgröße und bieten damit großes Potential für zukünftige Forschungsarbeiten.

Note

This book has been developed from a Ph.D. thesis written at the Karlsruhe Institute of Technology (KIT) from 2011 to 2015, further details of which are outlined in the German-language title page, and in the English and German abstracts in the front matter.

Karlsruhe Martin Weinmann
October 2015

Springer

http://www.springer.com/978-3-319-29244-1

Reconstruction and Analysis of 3D Scenes
From Irregularly Distributed 3D Points to Object Classes
Weinmann, M.
2016, XXII, 233 p. 81 illus., 69 illus. in color., Hardcover
ISBN: 978-3-319-29244-1