Int J Comput Vis (2013) 102:112–128
DOI 10.1007/s11263-012-0575-y

Online Detection of Repeated Structures in Point Clouds of Urban Scenes for Compression and Registration

Sam Friedman · Ioannis Stamos

Received: 5 November 2011 / Accepted: 11 September 2012 / Published online: 28 September 2012
© Springer Science+Business Media, LLC 2012

Abstract Laser range scans of urban areas have a distinctive geometry dominated by facade and ground planes and repetitive regular fenestration. Detection of these ubiquitous features provides profound insights into the scene. We present a novel method for detecting major planes and repetitive architectural features. Armed with this knowledge we illustrate its application in compression and registration of range scans. What is more our algorithm operates online, processing the scan as it is retrieved by the scanner. This realtime approach opens up new possibilities in range data segmentation, compression and registration.

Keywords Urban range scans · 3D scan registration · 3D scan compression · Regularity detection

1 Introduction

Understanding 3D images is increasingly important in fields as diverse as navigation, architecture, and biology. Each domain is replete with unique patterns that practitioners must uncover as they struggle to make sense of their data. In urban scenes, for example, the topology is characterized by the broad planes of facades, the slightly curving manifold of the street, and the fenestration of regular repeated architectural features like windows, balconies and cornices. Also notable is the fractal noise of vegetation, the Bezier curves of parked cars, and the high frequency spikes of objects in motion faster than the laser can catch them. In this paper we show how a focus on the repeated architectural structures yields significant scene understanding and can be used for compressing and registering range scans. We demonstrate a method to discover these regular structures online, as the scanner scans. We apply knowledge of these structures to facilitate compression and registration.

Processing point clouds quickly is crucial in many scenarios. Analyzing scans column by column allows seamless integration with the 3D camera because the camera is likewise scanning the scene column by column. This synchronizes the execution of the algorithm with gathering the data, accelerating our work to the point of just-in-time scene understanding. While focusing on individual columns gives excellent local knowledge on the scene we do not wish to lose sight of the big picture. After all, each column is just a sliver of a much larger tableau. However, by keeping an evolving data structure with macro features and greedily updating as the scan unfolds we are not forced to choose between speed and global knowledge, in fact our algorithm achieves both.

We can imagine this online processing as a line sweep algorithm. The scanline or column of measurements moves across the scene in discrete steps. At each step we search for periodicity and planes. We aggregate and update this data maintaining best estimates about higher level features in the scene. The processing occurs on the fly so that by the time the scanline reaches the edge of the scene our algorithm has done its work and requires no further manipulations of the data. Integration into the 3D camera hardware is a natural step forward for these types of algorithms.

This work has been supported in part by the following NSF grants: IIS-0915971, CCF-0916452 and MRI CNS-0821384. We would like to thank Tom Flynn for his labeling and registration software. We are also thankful to the anonymous reviewers for their helpful comments.

S. Friedman · I. Stamos (✉)
Hunter College & Graduate Center of CUNY, CUNY, New York City, NY, USA
e-mail: istamos@hunter.cuny.edu

S. Friedman
e-mail: umpteee@yahoo.com