Abstract

This thesis is a contribution towards the methods of 3D reconstruction in computer vision. Most of the 3D reconstruction methods are based on either the principle of triangulation or the principle of photometry. In this thesis, we investigate both of these approaches. Although, fundamentals of these principles are well known in computer vision, the recent availability of new hardware and easy access to the community photo collections have made it possible to investigate some of these methods in a newer context. This is because, the existing methods of 3D reconstruction are required to be modified to cope with the new types of data available by virtue of the current advancements. In this thesis, we propose some novel approaches for 3D reconstruction that improve or combine these methods in accordance with the current advancements.

Broadly speaking, we address two themes in the area of 3D reconstruction: (a) Geometric calibration methods for localizing a set of cameras or scanners and (b) methods for refining 3D shapes obtained from depth scanners.

First, we present a solution to the problem of relative-rotation averaging that occurs in the context of 'structure-from-motion' (SfM). Our approach leverages the geometric Lie-group structure of 3D rotations and incorporates robustness into our estimation method in a principled fashion. Our approach is robust, efficient and can easily handle large-scale problems. The optimality of the method is also analyzed. This approach is the state-of-the-art and is now incorporated into many libraries for SfM. A related problem is one of averaging relative Euclidean motions which occurs in the context of 3D reconstruction using depth scans. We demonstrate the robustness and the efficacy of our approach for 3D reconstruction of real-world models.

Second, we present two approaches to improve the quality of 3D scans obtained using commercial structured-light-stereo scanners, i.e. depth cameras. In the first instance, we carry out a sensitivity analysis of the noise inherent in such scanners. We demonstrate the use of this sensitivity analysis in improving the performance of 3D reconstruction pipelines. In the second instance, we exploit the complementary qualities of 3D estimates obtained from structured-light-stereo scanners and photometric methods. While structured-light-stereo scanners have coarse level fidelity, photometric approaches contain fine-scale 3D details. We develop a robust and accurate method of fusing these two approaches to combine the advantages of both these methods.