Int J Comput Vis (2013) 102:180–197
DOI 10.1007/s11263-012-0563-2

A Performance Evaluation of Volumetric 3D Interest Point Detectors

Tsz-Ho Yu · Oliver J. Woodford · Roberto Cipolla

Received: 4 November 2011 / Accepted: 17 August 2012 / Published online: 13 September 2012

Abstract This paper presents the first performance evaluation of interest points on scalar volumetric data. Such data encodes 3D shape, a fundamental property of objects. The use of another such property, texture (i.e. 2D surface colouration), or appearance, for object detection, recognition and registration has been well studied; 3D shape less so. However, the increasing prevalence of 3D shape acquisition techniques and the diminishing returns to be had from appearance alone have seen a surge in 3D shape-based methods. In this work, we investigate the performance of several state of the art interest points detectors in volumetric data, in terms of repeatability, number and nature of interest points. Such methods form the first step in many shape-based applications. Our detailed comparison, with both quantitative and qualitative measures on synthetic and real 3D data, both point-based and volumetric, aids readers in selecting a method suitable for their application.

Keywords 3D interest points · Volumetric interest points · Feature detection · Performance evaluation

1 Introduction

The applications of object detection, recognition and registration are of great importance in computer vision. Much work has been done in solving these problems using appearance on 2D images, helped by the advent of image descriptors such as SIFT and learning-based classifiers such as SVM, and these methods are now reaching maturity. However, advancing geometry capture techniques, in the form of stereo, structured light, structure-from-motion and sensor technologies such as laser scanners, time-of-flight cameras, MRIs and CAT scans, pave the way for the use of shape in these tasks, either on its own or complementing appearance—whilst an object's appearance is a function not only of its texture, but also its pose and lighting, an object's 3D shape is invariant to all these factors, providing robustness as well as additional discriminative power.

Detection and recognition of 3D objects is not new, e.g. (Fisher 1987), though such applications have seen a recent resurgence. Approaches range from the local to the global. At the global end are those which form a descriptor from an entire object. Such methods generally offer excellent discrimination plus robustness to shape variation, but, since the whole object is required and its extent known, they do not cope well with clutter or occlusion. Also, whilst suitable for recognition, the matching does not provide pose for registration applications. At the other end of the scale are highly local features, such as points. Being completely non-discriminant, such features are usually embedded in a framework that finds geometrical consistency of features across shapes, e.g. RANSAC (Brown and Lowe 2005; Papazov and Burschka 2011). The geometrical consistency framework makes matching, for detection and recognition, costly, but does provide pose for registration applications, and the local nature of features provides robustness to clutter and partial occlusion. Between these two extremes are those methods that describe local features of limited but sufficiently distinctive scope, thereby gaining the discriminability of global methods and the robustness of local methods.

T.-H. Yu (✉) · R. Cipolla
University of Cambridge, Cambridge, UK
e-mail: thy23@cam.ac.uk

R. Cipolla
e-mail: cipolla@eng.cam.ac.uk

O.J. Woodford
Toshiba Research Europe Ltd., Cambridge, UK
e-mail: oliver.woodford@crl.toshiba.co.uk