Computers & Graphics 58 (2016) 31–42

Contents lists available at ScienceDirect

Computers & Graphics

journal homepage: www.elsevier.com/locate/cag

Special Issue on SMI 2016

Color-aware surface registration

Shuai Lin [a], Yu-Kun Lai [b], Ralph R. Martin [b], Shiyao Jin [a], Zhi-Quan Cheng [c,*]

[a] School of Computer, National University of Defense Technology, Changsha City, Hunan Province 410073, China
[b] School of Computer Science and Informatics, Cardiff University, Cardiff, Wales CF24 3AA, UK
[c] Avatar Science Company, Changsha City, Hunan Province 410205, China

ARTICLE INFO

Article history:
Received 7 March 2016
Received in revised form
10 May 2016
Accepted 10 May 2016
Available online 12 May 2016

Keywords:
Surface registration
Textures
Color aware
MeshHOG
Gabor filter

ABSTRACT

Shape registration is fundamental to 3D object acquisition; it is used to fuse scans from multiple views. Existing algorithms mainly utilize geometric information to determine alignment, but this typically results in noticeable misalignment of textures (i.e. surface colors) when using RGB-depth cameras. We address this problem using a novel approach to color-aware registration, which takes both color and geometry into consideration simultaneously. Color information is exploited *throughout* the pipeline to provide more effective sampling, correspondence and alignment, in particular for surfaces with detailed textures. Our method can furthermore tackle both rigid and non-rigid registration problems (arising, for example, due to small changes in the object during scanning, or camera distortions). We demonstrate that our approach produces significantly better results than previous methods.

© 2016 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

1. Introduction

Reconstructing 3D objects from multiple scans taken from different viewpoints is a classical problem. A fundamental step in this task is surface registration, in which geometric data in local coordinate systems are aligned to a global coordinate system. Registration involves two intertwined problems: establishing correct correspondences between points on different surfaces, and finding a suitable spatial transformation or deformation that puts these surfaces into alignment. This is a computationally expensive task; most methods iteratively optimize correspondences and transformations alternately.

Surface registration has made substantial progress, and existing methods perform quite well, even in challenging cases with little overlap, noise and outliers. However, many methods formulate registration as an optimization problem based on *geometric* errors, and only use geometric information for finding correspondences and transformations. Such methods work well with textureless surfaces (in this paper, *texture* refers to surface coloring, not shape detail). However, recently developed low-cost RGB-depth (RGB-D) acquisition devices such as the Microsoft Kinect permit efficient and cheap capture of textured surfaces, leading to many novel applications such as clothed 3D human body reconstruction. Often, the color information has greater detail than the geometric information, so registration based on geometry alone can lead to poor results in which the textures are not well aligned. Worse still, texture misalignments are typically much more noticeable to the eye than geometric misalignments.

To overcome this problem, we present a novel color-aware registration algorithm that produces high-quality registration of textured surfaces. Our method can handle both rigid and non-rigid alignment: even if scans are supposed to be rigidly related, non-rigid alignment may be needed to correct for inaccurately determined camera intrinsic parameters, to allow for lens distortion, or to rectify small geometric changes (e.g. altered wrinkles in clothing when a human subject has moved slightly between successive scans).

Our first contribution is to use color information as well as geometric information to robustly find correspondences, in both rigid and non-rigid cases. We filter out incorrect vertex correspondences by using a combination of color, texture and geometric measures, and further improve pruning by rejecting correspondences which are not mutually consistent. Our second contribution is to also take color into account when using optimization to find the transformation or deformation needed for surface alignment.

We experimentally evaluate our *color-aware registration algorithm* using real scans, demonstrating the effectiveness of our algorithm when the subjects have richly textured clothing. We also compare our algorithm with other state-of-the-art methods using both real scans and a public dataset, showing the superiority of our method. While we have mainly tested the algorithm on our own data, we believe it to be generally useful, as identical problems are likely to arise in other RGB-D capture systems when capturing subjects with rich color textures.

* Corresponding author.
E-mail address: cheng.zhiquan@avatarscience.com (Z.-Q. Cheng).

http://dx.doi.org/10.1016/j.cag.2016.05.007
0097-8493/© 2016 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

2. Related work

Surface registration is an active research area. Various surveys [1–3] have summarized and categorized methods from different perspectives, including performance of descriptors, rigid versus non-rigid registration framework, etc. Here we only summarize the work most relevant to our problem.

2.1. Registration based on geometric closeness

Registration aims to put two (or more) shapes into alignment, so it is natural to use geometric closeness as a criterion to determine the transformations. In the rigid case, the most popular approach is the *iterative closest point* (ICP) method [4,5]. The main ideas are point correspondence selection based on Euclidean distances, and an energy function formulated based on closeness, either *point-to-point* or *point-to-plane*, for transformation optimization. Its many variants address different challenges, including noise, outliers, limited amounts of overlap, variations in initial positions, etc. Recent improvements have been achieved by care in correspondence selection [6], using L^p distances in the energy formulation [7], and use of branch-and-bound [8] or game theory [9] to explore the motion space. The approach has also been generalized largely unchanged to non-rigid registration, using dense reliable correspondences to find optimal local transformations [10–13]. To address the high degrees of freedom, priors are required, such as smoothness and piecewise rigidity.

2.2. Features, salience and texture assisted registration

While ICP approaches are effective, good initialization is required to avoid such iterative methods stopping in a local minimum. As an alternative, feature-based approaches identify correspondences based on intrinsic geometric features. This helps us to identify potential correspondences even when the shapes are far from alignment and reduces the possibility of mismatches. Features used include spin images [14], mean [15] and Gaussian [16] curvature, SHOT signatures [17], and salience measures such as differential properties [15] or multiscale slippage [18]. Color at each vertex can also be used to provide features, either as direct color values [19–21] or by derivation from color values, e.g. texture spin images [22] and orbital point descriptors [23]. Both [24] and [25] use color to assist registration and are thus closely related to our paper. The former uses SURF descriptors to help find correspondences and incorporates color information into the energy representing alignment error. However, their focus is completion and reconstruction of dynamic shapes from real-time data. Thus, their SURF descriptor is augmented with a temporal coherence constraint, which is not relevant in our setting where the multiple views differ substantially both in time and space. The latter focuses on texture correction to improve shape and texture reconstruction. However, their method aims to reconstruct static, rigid objects. Thus, their method starts by reconstructing a textureless model using KinectFusion to establish uniform geometric constraints, and then optimizes texture consistency locally. Although they consider sensor distortion, and apply non-rigid correction to the images, they do not change the geometry provided by KinectFusion. The method thus cannot cope with the data we assume, where as noted, multiple scans can differ substantially from views and minor non-rigid deformations are assumed to be present: KinectFusion reconstruction does not allow for such situations. The strength of our approach is to *combine* geometry and texture information during the whole process, allowing us to find good alignment even in challenging cases.

2.3. Other uses of color and geometric descriptors

Color has also been exploited in addition to geometric information in other applications such as shape retrieval and recognition. Various effective descriptors have been proposed, including heat kernel signatures [26], conformal factors [27], SIFT features [28], wave kernel signatures [29], and MeshHOG (histograms of oriented gradients) [30]. Although intended for applications which only need sparse correspondences, such descriptors have the potential for use in shape registration. In this paper, we use MeshHOG [30] during correspondence search.

Several works [31–33] have considered optimization of textures on a 3D mesh. Multi-view textures can be projected onto a pre-built 3D mesh and rectified to achieve texture consistency via feature correspondences, using e.g. SIFT [31] and optical flow [32], or discrete labeling optimization on each triangle [33]. Although such methods achieve consistent colorization for a 3D mesh, which is one of our goals as well, they assume rather simple geometry which can be reconstructed using existing registration methods, avoiding the difficulty of having to simultaneously optimize both geometry and texture. As Fig. 11 shows later, for challenging situations, independently performing shape registration to initially build a 3D mesh does not work well, whereas our color-aware surface registration exploits both color and geometry uniformly at all stages, producing substantially better results.

2.4. Textured surface fusion

Rapid advances in RGB-D cameras have driven applications based on 3D modeling of highly textured surfaces. The GPU-based capture system proposed in [34] uses a single Kinect camera to incrementally build 3D models. Tong et al. [35] use three Kinect sensors and associated components to produce acceptable 3D human body models. A fixed Kinect can be used with user orientation changes to capture and produce a full-body 3D self-portrait [36]. These systems all use traditional registration techniques to fuse scans. However, data captured by cheap RGB-D cameras has significant geometric distortions due both to camera distortion and errors introduced by depth processing; multiple frames captured even from a perfectly static subject may not align perfectly. If a human subject is captured, misalignment is further exacerbated by small unintentional movements. Our color-aware registration algorithm utilizes color information both when finding correspondences and when improving alignment; the emphasis is on providing a high-quality texture on the recovered surface.

The method in [37] formulates registration in an overall optimization framework which requires solution of a nonlinear system. This is difficult, so local linearization is resorted to. The method is furthermore used in a face tracking problem where strong priors are available. The blendshape model used is a linear combination of coordinates and is inapplicable to general objects such as human bodies. Other recent work also considers real-time non-rigid reconstruction from RGB-D data. The method in [38] is effective but requires a complete static template model to be captured first; this template is deformed to fit a real-time stream of scans. Alternatively, [39] does not require a template model, and demonstrates the effectiveness of using *sparse* color feature points in a tracking scenario (where adjacent frames are very similar) to help improve registration. Both methods require adjacent frames to be highly similar, so are unsuitable for non-rigid registration of general pairs of scans as considered in our problem. Our method also differs in that *dense* color information is used to align two scans to improve the geometry and *texture* alignment. Although having a different purpose, RGB-D SLAM (simultaneous localization and mapping) also considers use of RGB-D data to reconstruct 3D scenes [40,41]. Both color and geometry information are used

Download English Version:

https://daneshyari.com/en/article/6876952

Download Persian Version:

https://daneshyari.com/article/6876952

Daneshyari.com