Image and Vision Computing 32 (2014) 437–451

Contents lists available at ScienceDirect

Image and Vision Computing

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

Global registration of large collections of range images with an improved Optimization-on-a-Manifold approach☆

Francesco Bonarrigo, Alberto Signoroni *

Information Engineering Dept., DII, University of Brescia, Brescia, Italy

ARTICLE INFO

Article history:
Received 29 June 2013
Received in revised form 29 December 2013
Accepted 25 February 2014
Available online 20 March 2014

Keywords:
Global registration
3D scanning
Range images
Correspondence selection
Newton-type optimization
Differential geometry

ABSTRACT

Concurrently obtaining an accurate, robust and fast global registration of multiple 3D scans is still an open issue for modern 3D modeling pipelines, especially when high metric precision as well as easy usage of high-end devices (structured-light or laser scanners) are required. Various solutions have been proposed (either heuristic, iterative and/or closed form solutions) which present some compromise concerning the fulfillment of the above contrasting requirements. Our purpose here, compared to existing reference solutions, is to go a step further in this perspective by presenting a new technique able to provide improved alignment performance, even on large datasets (both in terms of number of views and/or point density) of range images. Relying on the 'Optimization-on-a-Manifold' (OOM) approach, originally proposed by Krishnan et al., we propose a set of methodological and computational upgrades that produce an operative impact on both accuracy, robustness and computational performance compared to the original solution. In particular, always basing on an unconstrained error minimization over the manifold of rotations, instead of relying on a static set of point correspondences, our algorithm updates the optimization iterations with a dynamically modified set of correspondences in a computationally effective way, leading to substantial improvements in terms of registration accuracy and convergence trend. Other proposed improvements are directed to a substantial reduction of the computational load without sacrificing the alignment performance. Stress tests with increasing view misalignment allowed us to appreciate the convergence robustness of the proposed solution. Eventually, we demonstrate that for very large datasets a further computational speedup can be reached by the adoption of a hybrid (local heuristic followed by global optimization) registration approach.

1. Introduction

Objects and scene modeling from 3D scan data requires an accurate alignment of the acquired multiple views. In a conventional alignment pipeline [1], two consecutive alignment steps, usually called coarse (or initial) and fine (or refined), are performed on all the overlapping views. Either these steps are carried out in a pairwise fashion (with possible accumulation of alignment errors) or according to some multi-view heuristics, high quality 3D modeling requires a final optimized global registration of the whole set of views in order to further reduce and evenly distribute the residual alignment errors. Other than guaranteeing the highest alignment *accuracy*, this last step should also provide other desirable properties. First of all *robustness*, because it is not guaranteed that all views are already well aligned, e.g. when loop closure problems are to be solved. Secondly *computational efficiency*, where computation time becomes an issue especially for dense view and large view collections. Moreover, several advanced, specialized and professional application fields (e.g. biomedicine, orthopedics and orthodontia, cultural heritage, industrial design and reverse engineering) tend to generate large data sets (both in terms of number of views, spatial point density, as well as metric accuracy) so that the overall *usability* of the acquisition and processing pipeline emerges as another primary requirement.

Pretty fast heuristic methods have been proposed to handle the global registration problem; the most popular one, and also a recognized reference in terms of performance, is the method proposed by Pulli [2]. However, such kind of methods cannot guarantee (as we also verify in this work) to actually minimize the alignment error, so that the views' position remains suboptimal while, as stated, maximized metric accuracy of the views' alignment is a sine qua non for the aforesaid demanding application fields. For this reason, more connatural methods for approaching the global registration problem are those that adopt some numerical optimization solutions. However, 1) the global nature of the problem can easily generate computational issues, and 2) optimization methods are to be well designed and conditioned in order to reduce the risk to be stuck in local minima or to behave

☆ This paper has been recommended for acceptance by Daniel Rueckert.

* Corresponding author: Information Engineering Dept., DII University of Brescia Via Branze, 38, 25123 Brescia, Italy. Tel.+39 030 3715 432; Fax.+39 030 380 014.
E-mail address: alberto.signoroni@ing.unibs.it (A. Signoroni).

http://dx.doi.org/10.1016/j.imavis.2014.02.012
0262-8856/

incorrectly. Our proposed solution relies on the optimization over the differentiable manifold of 3D rotation matrices (i.e. the orthogonal Lie group SO(3)) [3,4], a functional theoretical framework which, however, in its early proposed global registration implementation [5,6], reveals itself not fully adequate to address the above problems 1) and 2). Our main contribution here consists of a new redesigned algorithmic approach which significantly improves the performance of the original one, demonstrating to be suitable for the most demanding modern applications in terms of convergence, accuracy, and computational performance. A hybrid approach, which combines a first heuristic-based alignment stage followed by the proposed solution, is also tested in order to further reduce the global registration time for the most demanding cases.

Since the proposed solutions are designed to be used in the context of modern and high-quality 3D object modeling chains, a gentle introduction to the problem we tackle is firstly provided in Section 1.1 (which can be possibly skipped by the informed reader). Subsequently, a literature review is presented in Section 1.2. Finally, an outline of the paper concludes this introduction (Section 1.3).

1.1. Problem description

The possibility to create accurate 3D models of real objects is a primal asset of several application domains such as industrial reverse engineering [7] and visual inspection [8], cultural heritage preservation [9], robot localization and navigation [10], and biological and medical imaging [11]. This is increasingly influencing a number of related fields such as computer animation and augmented reality, environmental remote sensing, construction site management, e-commerce, entertainment, fashion industry, and more. When metric accuracy is a priority, the acquisition is usually made by active 3D optical devices such as laser or structured light 3D scanners which today guarantee high quality and dense point sets captured from their field of view (FoV). At the same time, the physical FoV limitation of the scanner and the spatial morphology of the objects to acquire, require the collection of several (dozen or even hundreds) views in order to build a complete 3D model of the object or the scene of interest. Starting from a collection of independent views (lying in their own local reference systems), and under proper view overlapping and object surface coverage hypothesis, the construction of a 3D model is obtained through a pipeline consisting in a series of (geometric) data processing steps, where the main objective is the registration (or alignment) in a common reference system of all views toward a unique digital model, under application-driven requirements that are usually related to accuracy, robustness, computational speed and automatism.

Due to the intrinsic complexity nature of the acquired data (caused by the presence of occlusions, noise, outliers, residual distortions and variable view overlapping area) it is not viable to conceive and formalize the alignment problem according to a unique single-step, closed-form solving approach. A main distinction regards the concepts of 'local' and 'global'. It is clear how the alignment problem would have a global nature as a whole, however it must be recognized that the alignment of a single view is quite a local problem, and that in most cases the acquisition of an object or the navigation of an environment follows a certain path, either decided on-the-fly or planned in advance, and this is usually linked to local decisions as well. Therefore the alignment problem can be conceived as a local-to-global approach, in that each view is registered with respect to its neighboring ones, possibly according to a specific path, with the aim of arriving to a global registration. The role of global registration is twofold: to perform an equal redistribution of residual alignment errors due to data imperfections, and to solve the closure problems that are likely to happen whenever the acquisition path turns around the object, possibly several times and through crossing paths, in order to obtain a complete coverage of the object. The typical strategy employed by well-defined, state-of-the-art approaches is to apply a local coarse alignment technique followed by a local fine alignment step, which is in turn used to initialize a global registration phase. Usually, each phase defines different problems that require distinct solving approaches. In the present work we focus on the final step of the alignment problem, that is, the global registration phase. We initially assume that we can rely on a fairly good initial alignment of the views, as the one that can ordinarily be provided by reference fine alignment techniques, e.g. those based on pairwise Iterative Closest Point solutions, being the fine alignment usually applied in a sequential pairwise order following the object scanning path. We will also strain our system by gradually releasing the strictness of the fine alignment assumption, while concurrently performing an enlarged exploration of the basin of convergence of the proposed optimization solution. Overall, our method will prove to be effective, robust and versatile to guarantee a highly accurate global alignment for a wide class of challenging (large and dense)

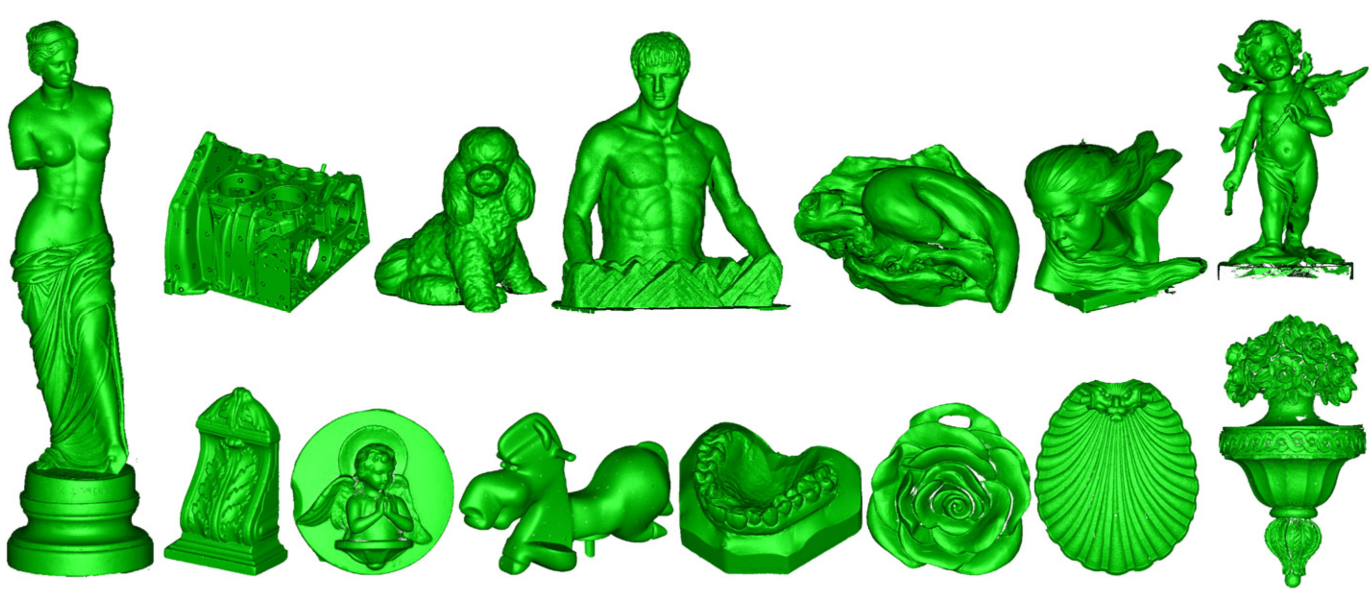

Fig. 1. The datasets employed for global registration assessment. From upper left: Venus, Crankcase, Dog, Neptune, Dolphin, Hurricane, Cupid, Capital, Cherub, Horse, Denture, Rose, Shell and Decoration.