Available online at www.sciencedirect.com
ScienceDirect
Computer Vision and Image Understanding 107 (2007) 195–202
Computer Vision and Image Understanding
www.elsevier.com/locate/cviu

A fast algorithm for ICP-based 3D shape biometrics

Ping Yan *, Kevin W. Bowyer

Department of Computer Science and Engineering, University of Notre Dame, Notre Dame, IN 46556, USA

Received 1 May 2006; accepted 5 November 2006
Available online 13 December 2006

Abstract

In a biometrics scenario, gallery images are enrolled into the database ahead of the matching step, which gives us the opportunity to build related data structures before the probe shape is examined. In this paper, we present a novel approach, called "Pre-computed Voxel Nearest Neighbor", to reduce the computational time for shape matching in a biometrics context. The approach shifts the heavy computation burden to the enrollment stage, which is done offline. Experiments in 3D ear biometrics with 369 subjects and 3D face biometrics with 219 subjects demonstrate the effectiveness of our approach.

Keywords: 3D shape; Shape matching; Iterative closest point; Biometrics

1. Introduction

Since its introduction by Chen and Medioni [1] and Besl and McKay [2], the Iterative Closest Point (ICP) algorithm has been widely used for 3D shape matching [1,3–5]. It has been used in a wide range of application areas, including the integration of range images [6,7] and alignment of CT and MR images [8]. Here, we are specifically interested in 3D shape matching for biometrics [9–13]. The ICP algorithm is known to be computationally expensive. With two clouds of points, source S (probe) and target T (gallery), the complexity of a typical single ICP iteration is $O(N_S \log (N_T))$ using a k-d tree data structure [9] in the expected case, where N_S is the number of points in the source and N_T is the number of points in the target. The ICP algorithm iteratively finds the minimum distance between two surfaces. With N_I iterations, the overall complexity is $O(N_I \times N_S \times \log (N_T))$ [2]. Therefore, matching high-resolution images of both source and target leads to a heavy computational load. A fast ICP implementation is crucial for practical use in biometrics.

Using shapes sensed by a 3D scanner is a major recent trend in biometrics [9–14]. A scan yields a 3D surface that can be used as a representation of the subject. In this paper, we illustrate our approach using both 3D ear and 3D face shapes. There are two types of images in a biometric application, gallery and probe. The gallery images are those that have been enrolled and whose identities are known to the system, while the probe images are those that need to be matched against the images in the gallery. In a recognition scenario, one probe is matched against all the images in the gallery, and the algorithm returns the match with the minimum error distance. In a verification scenario, one probe is matched against just one gallery entry, the one enrolled for the claimed identity. In recognition or verification experiments, enrollment occurs once and is followed by many instances of recognition.

One special characteristic of a biometrics application is that all gallery images are enrolled into the database before the matching takes place. Probe images are introduced into the system for matching. Taking advantage of the fact that the gallery images are enrolled prior to matching, we propose a novel method to accelerate the ICP matching. Our new method is called the "Pre-computed Voxel Nearest Neighbor". The idea is to voxelize a volume which can hold the 3D gallery surface, and for each voxel to

* Corresponding author.
E-mail addresses: pyan@nd.edu (P. Yan), kwb@cse.nd.edu (K.W. Bowyer).

1077-3142/$ - see front matter
doi:10.1016/j.cviu.2006.11.001

pre- compute its distance to the 3D gallery surface and save this for future use.

In Section 2, we review several fast ICP approaches. Then in Section 3 we give details of our approach. Section 4 addresses the applicability of our approach by using the ear and face biometrics, and experimental results are presented and analyzed. Finally, Section 5 discusses further refinements and possible future directions.

2. Literature review

In biometrics applications, 3D shape is used by many researchers in face biometrics [11–13,15–18], has also been used in ear biometrics [9,10], and has also been used in hand biometrics [14].

There have been a number of efforts to speed up ICP matching. One line of work is focused on fast algorithms for computing the nearest neighbor. The use of the k-d tree data structure appears to be the standard method in this area [19,2]. Cleary and co-workers analyzed the "Elias" algorithm for searching nearest neighbor in the n-dimensional Euclidean space [20]. They claimed that by using the "Elias" algorithm, the number of search points is independent of the total number of points on the surface.

In [21], Greenspan et al. proposed a novel nearest neighbor algorithm for small point sets. They report that "Elias" is much faster than a plain k-d tree, and that the "spherical constraint" method improves the speed still further. Zinßer et al. analyzed the performance of the nearest neighbor algorithm for ICP registration [22]. Their work is not limited to range images or triangle meshes, but also can be used with 3D point sets generated by structure-from-motion techniques.

Benjemaa [23] proposed a multi-z-buffer technique to accelerate the ICP algorithm. All points are projected in a z-buffer to perform the local search, and they claimed that this space partition speeds up the search for point-to-projection correspondences. But in order for the multi-z-buffer technique to work properly, the two surfaces need to be sampled with a high and uniform density.

Another line of work in this area looks at different subsample strategies to reduce computation time. One strategy is using multi-resolution approaches; that is, start with a coarse point set and use progressively finer point sets as the algorithm proceeds. The idea of the average distance between points in the current resolution in comparison to the average distance between matched points is the standard way to automate the switching between resolutions [24].

In [3], Gelfand et al. describe the importance of the quality of the point pairs. In the presence of noise or miscalibration in the input data, it is easy to create poor correspondences between pairs of points. Therefore, the least-squares technique might lead to wrong pose, or make it difficult for the algorithm to converge. They propose a technique to decide whether a pair of meshes has good quality by measuring the covariance matrix between two meshes which have been sparsely and uniformly sampled. This technique tries to avoid the unstable movement between two surfaces by sampling the features from the input data which are the best constraint for this kind of movement.

In [25], Rusinkiewicz and Levoy discussed the variants of ICP which affect all phases of the algorithm. They list most of these variants, and evaluate their effects on the speed with which the correct alignment is reached. Also in the paper, they proposed a combination of ICP variants optimized for high speed.

Researchers have also looked at mixing the two lines of work, having some multi-resolution mixed with some constrained search for nearest neighbor. Jose and Hügli proposed a solution that combines a coarse to fine multi-resolution approach with the neighbor search [26]. The multi-resolution approach permits to successively improve the registration using finer levels of representation and the neighbor search algorithm speeds up the closest point search by using a heuristic approach. They claim this technique reduces the time complexity of searching from $O(N \log(n))$ to $O(N)$, while preserving the matching quality [27].

Research related to ICP is also prominent in the graphics community. Leopoldseder et al. used d^2-tree to store a local quadratic approximant of the squared distance function to a surface [28]. Mitra et al. consider a general framework for matching two shapes represented by point clouds, in which the point-to-point and point-to-plane versions of ICP can be considered special cases [29]. Cheng et al. consider a method to fit a subdivision surface to an unorganized point cloud dataset [30]. However, none of these efforts are undertaken in a biometrics context. Also, while Leopoldseder and Mitra use a subdivision of 3D space, they still use a tree search to find the closest point correspondence between two point sets, rather than reducing it to an indexing operation as in this paper.

3. Fast ICP matching for 3D shapes

The most time consuming part of the ICP algorithm is that for each point on the probe surface, the algorithm needs to find the closest point on the gallery surface. By using these pairs of corresponding points, the ICP algorithm iteratively refines the transforms between two surfaces, finding the translation and rotation to minimize the mismatch.

This search for a closest point on the gallery surface is initially done using a k-d tree, as described in [9], and each search takes $O(\log N_G)$, where N_G is the number of the points on the gallery surface. Our goal is to reduce this search time to a constant value. The main idea is that if we can pre-compute the distance from any point in the 3D space to the gallery surface, and use it when needed, then the search time for a closest point is a constant.

Our "Pre-computed Voxel Nearest Neighbor" approach is illustrated on the application of matching 3D surfaces for biometric recognition. At the time of enrollment, the