21st International Conference on Pattern Recognition (ICPR 2012)
November 11-15, 2012. Tsukuba, Japan

A Hierarchical Approach to High-Quality Partial Shape Registration

Ming Zhong Tingbo Hou Qin Hong
Department of Computer Science, Stony Brook University
{*mzhong, thou, qin*}*@cs.stonybrook.edu*

Abstract

This paper presents a feature-driven, hierarchical shape registration algorithm. The central idea is to generate correspondences in multiple levels in a coarse-to-fine manner, with additional features incrementally inserted in each level. The registration starts from the coarsest resolution. Registration results obtained in one level serve as references for the registration in the next level. We adopt the heat kernel coordinates [3] for local shape parameterization, giving rise to a complete solution capable of registering partial shapes undergoing isometric deformation with higher accuracy. Through experiments, we demonstrate the effectiveness of this multi-resolution method and its advantages over the single-resolution method.

1. Introduction

Shape registration is an important problem in shape analysis and computer vision. The goal is to establish valid optimal correspondences between points on two different shapes. Its many applications include shape interpolation, attribute transfer, surface completion, shape matching and retrieval, etc. Comparing with rigid shapes, deformable shapes are more challenging to handle, typically requiring a parameterization that is invariant under certain deformations. Among various types of non-rigid deformations, isometric deformation is of particular value, as numerous real-life deformations can be deemed as approximately isometric, e.g. facial expression and articulated movement.

A common and effective approach to dense correspondence is first matching a small number of preselected feature points, and then using the matched features as references for dense correspondences. Features, encoding important information of shapes, can be used to parameterize the shapes and serve as anchors to bootstrap the matching of the rest points. In general, it is necessary to have a fairly large number of matched

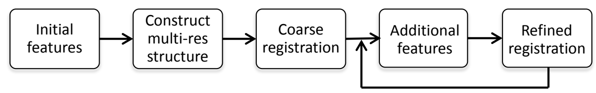

Figure 1. Pipeline overview of our hierarchical registration framework.

features to obtain dense correspondence of good quality. Otherwise, the feature-based parameterization may have difficulty in discriminating nearby elements, especially in parts of shapes that are far away from any features. However, automatically finding and matching a large number of features is very difficult and error-prone. Even in the case of user-assisted feature matching, one would prefer a small set of matched features, since manually corresponding many features is burdensome and time-consuming.

In this paper, we present a feature-driven hierarchical registration method that can achieve high-quality registration results even when the size of initial correspondence set is small. As illustrated in Figure 1, the main steps of our method are: (1) Detect and match features to get a small initial set of feature matches; (2) Construct hierarchical structures of input shapes; (3) Perform registration at the coarsest level using the initial feature set; (4) Select some newly registered points as additional features; (5) Perform registration at the next level using results from the previous level and the expanded set of feature references; (6) Repeat step (4) and (5) until all valid points are registered.

The rationale of our approach is that distinguishing elements that are distant from each other on the surface is much more accurate than nearby elements. Even with a small number of features, we can achieve very good registration on a heavily downsampled version of the original shapes. The registration result of a coarse resolution can serve as seed correspondences when performing registration in a finer level. The large number of available seeds significantly reduce the chances of correspondences being trapped in an incorrect location.

Moreover, the multi-resolution process enables us to pick additional features from already registered points. This greatly enhances the discriminative strength as the meshes become more refined. To summarize, the contributions of this paper are:

- We propose a hierarchical framework for shape registration with incremental reference insertions. The central idea is to establish correspondences in multiple levels in a coarse-to-fine manner.
- We improves the method of heat kernel coordinates (HKC) by developing a complete solution for high-quality registration.
- We articulate the advantages of hierarchical registration over one-level method and demonstrate the performance with experimental results.

2. Related work

In [1], Bronstein et al. proposed the generalized multidimensional scaling (GMDS) framework and compute the least distortion mapping between the two surfaces in the Euclidean embedding space. Conformal mapping is able to flatten non-rigid 3D meshes to a 2D domain; the registration then can be performed in the embedding space[6]. [10] presented a randomized, RANSAC-like feature matching algorithm to match points and optimize the registration by post-processing based on geodesics. The dense correspondence method in [8] is similar to this paper in terms of the coarse-to-fine strategy, where the full correspondences are determined by means of gradually refined patches. Many diffusion-based characteristics, such as heat kernel, are invariant under isometric shape deformation just like geodesics, but tend to be more stable in the presence of holes and noises. Hence, these characteristics are often adopted in isometric registration in place of geodesics. In [7], Ovsjanikov et al. defined heat kernel map (HKM) which parameterizes a surface by computing multi-scale heat kernels from a fixed reference point; given two shapes whose reference points are matched beforehand, the full correspondence can be recovered via a greedy global N-N search. [3] proposed the heat kernel coordinates, which utilize multiple feature points as anchors and globally parameterize the surface via heat kernel. Instead of searching for optimal correspondence in the entire domain as in [7], the search space for each point is restricted to the vicinity of already-registered points, generating results of great geometric compatibility. However, the registration is only performed at one resolution, which is difficult to achieve high-quality results in fine details.

3. Hierarchical registration

Given a source shape S and a target shape T, both represented as triangular meshes, and let $V^S = \{s_i\}$ and $V^T = \{t_i\}$ be their respective vertex sets, the objective of dense registration is to find an optimal mapping $\tau : V^S \to V^T$. In practice, we represent the registration results as a set of correspondences $R = \{v_i^S, v_i^T\}$. When the shapes in question are not complete, some vertices in V^S may not have correspondences in R.

3.1. Initial feature detection and matching

The goal of this step is to obtain a small feature correspondence set C^*. One can employ any good method to find and match features as long as the matched features are stable and representative. In this work, we adopt the heat kernel signature (HKS) [9] to extract multi-scale features and spectral graph matching method [5] to match them. The HKS is defined as the heat kernel from one point to itself: $h_t(x,x)$, with

$$h_t(x,y) = \sum_{k=0}^{\infty} e^{-\lambda_k t}\phi_k(x)\phi_k(y), \quad (1)$$

where λ_k and ϕ_k are the k-th eigenvalue and eigenfunction of the Laplace-Beltrami operator. We strongly refer readers to [3] for more details on this step.

3.2. Multi-resolution structure

Once we obtained features correspondence set C^*, we can use it as reference to propagate the correspondences by searching in the vicinity of already matched vertices, until every source vertex is mapped to a vertex in the target shape [3, 4]. However, when the size of C^* is small, simple propagation approaches often cannot produce satisfactory registration. On one hand, with insufficient features as anchors, it is difficult to distinguish nearby vertices no matter what kind of parameterization scheme we employ. On the other hand, since the sources for propagation are few, wrong correspondences are more likely to accumulate following a mismatch.

To address this issue, instead of computing registration in a single run, we perform it hierarchically in a coarse-to-fine manner. We construct a multi-resolution structure of the original shapes, and in each level we only register vertices that belong to the current resolution. Given a triangular mesh $M_0 = (V_0, F_0)$ and constants $d, m \in \mathbb{Z}$, we downsample M and obtain the mesh hierarchy $\{M_0, M_1, \ldots, M_m\}$. Assume $M_i =$

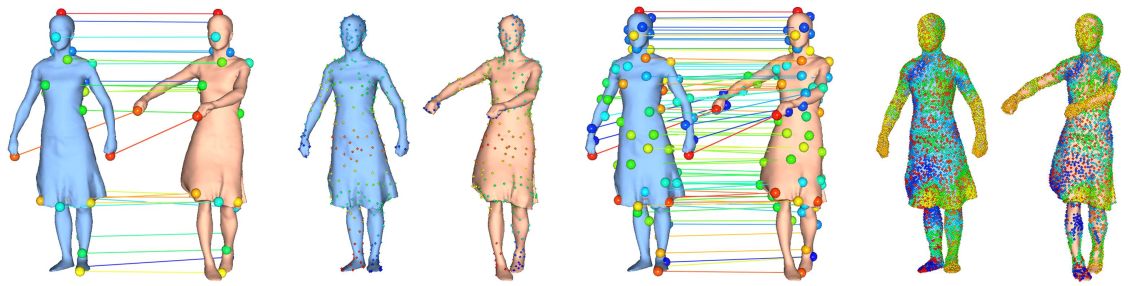

Figure 2. Major steps of our hierarchical registration algorithm. The blue shape is the source and the red one is the target. We use a three level hierarchy in this example. (a) Initial feature correspondences; (b) Coarse registration result (Third level); (c) Expanded feature correspondences (Third level); (d) Final registration result.

(V_i, F_i) and $n_i = |V_i|$, we enforce that $n_{i+1} = n_i/d$. We adopt the method in [2] for mesh downsampling. In our implementation, we select $d = 4$.

3.3. Correspondence propagation and feature expansion

Let both the initial correspondence set R_{m+1} and initial feature set C_{m+1} be C^*. In level l, we input the previous level's registration result R_{l+1} and feature set C_{l+1}. The goal is to find the l-th level correspondence set R_l that registers meshes S_l and T_l, with an augmented feature set C_l.

For each vertex x in S_l and T_l, we compute its heat kernel coordinates

$$\mathrm{HKC}(x) = (h_t(x, c_1), \ldots, h_t(x, c_z)), c_i \in C_{l+1}. \quad (2)$$

Inhering R_{l+1} as the initial correspondence set, we propagate correspondence to match the rest vertices in S_l and T_l. We use a heap to determine the order by which the vertices in S_l are processed, prioritizing on the magnitude of HKC. For an already matched pair (s_j, t_j) and one of s_j's immediate neighbor s_i, we search for s_i's best correspondence $t_i \in V_k^T$ in the neighborhood of t_j, and add (s_i, t_i) into the correspondence set. t_i is selected using the following criterion

$$t_i = \operatorname*{arg\,min}_{t \in n(t_j, T_k)} \|\mathrm{HKC}_S(s_i) - \mathrm{HKC}_T(t)\|_2 \quad (3)$$

where $n(t_j, T_k)$ represent the set of t_j's neighboring vertices in T_k, and HKC_S and HKC_T denote the heat kernel coordinates of points on S and T.

The correspondence propagation continues until all vertices in S_k have been matched and we get the correspondence set $\{(s_i, t_i)\} \subset V_{m-1}^S \times V_{m-1}^T$. Note that for each correspondence (s_i, t_i), the endpoint t_i actually represent a set of vertices $K(t_i)$ in the original mesh T_0. To find the precise correspondence of s_i in the original target mesh, we search $K(t_i)$ and replace t_i with $t_j \in K(t_i)$ if t_j is closer to s_i in the embedding space. The result is the l-th level correspondence set R_k that relates points $s_i \in S_k$ to $t_i \in T_0$. For each correspondence (s_i, t_j), we assign a matching score

$$score(s_i, t_j) = \exp(-\|\mathrm{HKC}_S(s_i) - \mathrm{HKC}_T(t_j)\|_2). \quad (4)$$

We then select from R_l some vertex pairs as new features and insert them into the feature set. These new added feature pairs should be both reliable (having great matching score) and not in the δ-neighborhood of any existing feature points. The expanded feature set C_l enables a more discriminative HKC in the next level. We carry on this process from the coarsest level to the finest level until we obtain the final registration set R_0 between the original meshes S_0 and T_0. Figure 2 shows the major steps of our algorithm.

4. Experimental results

To assess the registration results represented by map $\tau : V^S \to V^T$, we randomly sample $N = 300$ pairs of source vertices $\{(s_1^1, s_2^1), (s_1^2, s_2^2), \ldots, (s_1^N, s_2^N)\}$. We measure the quality of τ in terms of the mean relative error of geodesics:

$$error(\tau) = \frac{1}{N} \sum_{i=1}^{N} \frac{|d_G^S(s_1^i, s_2^i) - d_G^T(\tau(s_1^i), \tau(s_2^i))|}{d_G^S(s_1^i, s_2^i)} \quad (5)$$

where d_G^S and d_G^T are the respective geodesic distance functions on surface S and T.

data	$\lvert V^S\rvert, \lvert V^T\rvert$	#initial features	#final features	#levels	error (multi-level)	error (single-level)
Face	10.5K, 9.8K	11	68	2	0.058	0.108
Horse	8.4K, 8.4K	18	37	3	0.083	0.101
Cat	7.2K, 7.2K	14	65	3	0.123	0.167
Man	10.0K, 10.0K	13	90	3	0.077	0.117
Woman	10.0K, 10.0K	25	106	3	0.048	0.101

Table 1. Evaluation result. Our method has lower errors than the single-level method.

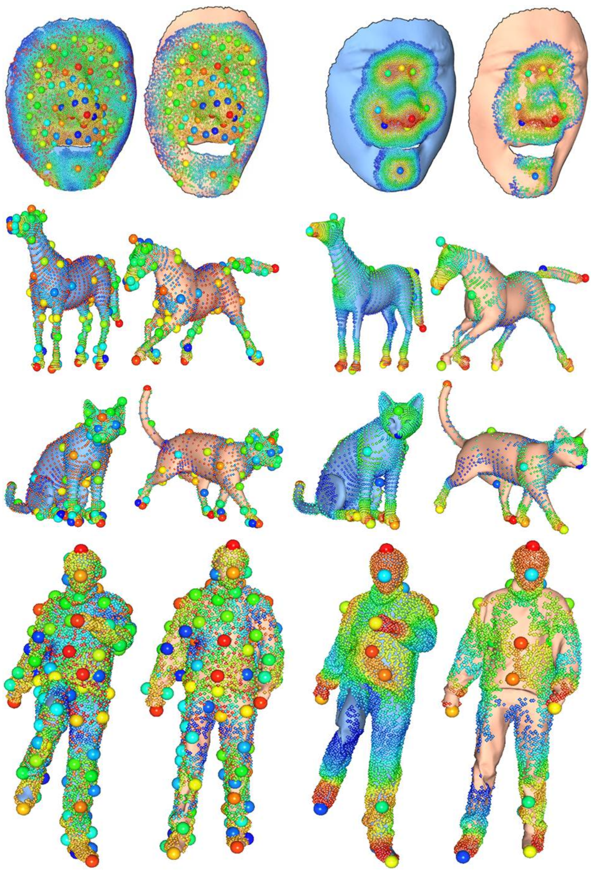

Figure 3. Some registration results by our multi-resolution method (*Left*) and the single-resolution method [3] (*Right*). Large colored dots represent matched features.

We evaluate our algorithm on various models and compare it with the single-resolution method [3]. Table 1 documents the evaluation results. Starting with the same initial feature set, our hierarchical method consistently achieves a better registration than the single-resolution method. In average, the sampled error is only 64% of the single-resolution approach. Figure 3 shows some registration results in our experiments.

5. Conclusion

We have presented a feature-driven, hierarchical framework for high-quality shape registration, in which correspondences are established in multiple levels from the coarse resolution to fine resolution. By integrating heat kernel coordinates into this framework, we developed a complete solution that is capable of registering partial shapes undergoing isometric deformations. We demonstrate via experiments the superiority of our multi-resolution approach over traditional, single-resolution method. As future work, we plan to integrate other parameterization schemes in this framework and explore possible applications.

References

[1] A. M. Bronstein, M. M. Bronstein, and R. Kimmel. Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *PNAS*, 103(5):1168–1172, 2006.

[2] M. Garland. *Quadric-Based Polygonal Surface Simplification*. PhD thesis, Carnegie Mellon University, 1999.

[3] T. Hou and H. Qin. Robust dense registration of partial nonrigid shapes. *TVCG*, 99, 2011.

[4] Q.-X. Huang, B. Adams, M. Wicke, and L. J. Guibas. Non-rigid registration under isometric deformations. *Computer Graphics Forum*, 27(5):1449–1457, 2008.

[5] M. Leordeanu and M. Hebert. A spectral technique for correspondence problems using pairwise constraints. In *ICCV'05*, 2005.

[6] Y. Lipman and T. Funkhouser. Mobius voting for surface correspondence. *ACM Transactions on Graphics*, 28(3), 2009.

[7] M. Ovsjanikov, Q. Mrigot, F. Mmoli, and L. Guibas. One point isometric matching with the heat kernel. *Computer Graphics Forum*, 29(5):1555–1564, 2010.

[8] Y. Sahillioglu and Y. Yemez. Coarse-to-fine combinatorial matching for dense isometric shape correspondence. *Computer Graphics Forum*, 30(5):1461–1470, 2011.

[9] J. Sun, M. Ovsjanikov, and L. Guibas. A concise and provably informative multi-scale signature based on heat diffusion. *Computer Graphics Forum*, 28(5), 2009.

[10] A. Tevs, M. Bokeloh, M. Wand, A. Schilling, and H.-P. Seidel. Isometric registration of ambiguous and partial data. In *CVPR'09*, 2009.