An Efficient Approach to Correspondences between Multiple Non-Rigid Parts

Gary K.L. Tam^{1,2}, Ralph R. Martin², Paul L. Rosin² and Yu-Kun Lai²

¹Department of Computer Science, Swansea University, UK ²School of Computer Science & Informatics, Cardiff University, UK

Abstract

Identifying multiple deformable parts on meshes and establishing dense correspondences between them are tasks of fundamental importance to computer graphics, with applications to e.g. geometric edit propagation and texture transfer. Much research has considered establishing correspondences between non-rigid surfaces, but little work can both identify similar multiple deformable parts and handle partial shape correspondences. This paper addresses two related problems, treating them as a whole: (i) identifying similar deformable parts on a mesh, related by a non-rigid transformation to a given query part, and (ii) establishing dense point correspondences automatically between such parts. We show that simple and efficient techniques can be developed if we make the assumption that these parts locally undergo isometric deformation. Our insight is that similar deformable parts are suggested by large clusters of point correspondences that are isometrically consistent. Once such parts are identified, dense point correspondences can be obtained by an iterative propagation process. Our techniques are applicable to models with arbitrary topology. Various examples demonstrate the effectiveness of our techniques.

1. Introduction

In graphics applications, editing deformable parts in a realistic way by hand is tedious and time-consuming. Instead, a user might like to edit part of a mesh by deforming it (e.g. a finger of a hand), and propagate the results to other similar parts (e.g. other fingers of the same or other hands). A tool which could do this automatically would greatly reduce the time taken and effort. To realise such a tool, two problems must be solved: (i) finding multiple similar deformable parts related to a given query by a *non-rigid* transformation (for conciseness, we refer to them as *non-rigid* parts hereafter) and (ii) establishing pointwise correspondences between them.

Measuring similarity between shapes where parts of them may *non-rigidly* correspond, and using this to determine similar parts, is a special case of *partial matching* discussed in the shape retrieval literature. Techniques for solving this problem focus on developing deformation invariant feature representations and descriptors [TV08]. Since measuring similarity is the main focus, finding correspondences between non-rigid parts is often limited to comparing shape components or parts (e.g. a matching of deformable parts (limbs, body, head) between two human models would in-

© 2014 The Author(s) Computer Graphics Forum © 2014 The Eurographics Association and John Wiley & Sons Ltd. Published by John Wiley & Sons Ltd. dicate a high similarity). These matchings, however, cannot provide a consistent set of point correspondences between the parts to allow geometry or texture transfer. This is largely due to the large range of possible deformation and articulation changes that these techniques need to deal with.

In the *non-rigid* shape registration literature, the focus is on establishing reliable correspondences down to the pointwise level [vKZHCO10]. There are many registration techniques targeting many different types of deformations (e.g. muscle/fat deformation, articulation, isometric deformation). These techniques, however, assume the *non-rigid* shapes are similar and correspond as a whole. Many of them do not support correspondences for *non-rigid parts* due to the assumption of global isometry, or because of topological restrictions. Many cannot handle multiple similar parts.

The goal of this paper is to address the two problems above. We aim to efficiently identify multiple similar nonrigid parts, and establish dense pointwise correspondences between them. Both are challenging problems because the locations of the parts are unknown in advance, and the parts to be found are not an exact match to the query part. They may differ in terms of representation (e.g. differing triangulations of the same shape), minor shape differences (e.g.



Figure 1: Overview. Example query part: right hand of dinopet in reference pose. Target mesh: dinopet in a different pose. (a) 220K initial point correspondences; (b) visualization of cluster of point correspondences; (c) point correspondences after pruning resulting in local clusters; (d) and (e) output point correspondences between right and left hands: corresponding points have the same colour.

front legs and back legs) and pose differences (e.g. leg bent at the knee or straight). What is a match depends on the type of assumed deformation; if a sufficiently general definition of deformation is used, any two parts of a mesh may be put into correspondence.

In this paper, we show that simple and efficient techniques can be developed if we make the assumption that these parts undergo a deformation which is locally isometric. Global isometry has been shown to be a useful and reliable guide in modelling many real kinds of deformation (see e.g. [HAWG08, OMMG10, TBW*11]). In [TMRL14] we recently proposed a point correspondence pruning technique. We build an isometric consistency matrix measuring pairwise consistency of point correspondences that only uses relationships in local neighbourhoods. Diffusion analysis is used to infer global isometric consistency, and remove inconsistent correspondences. The idea is to find the *largest* single cluster of point correspondences through analysis of the isometric consistency matrix. This approach is more general than requiring global isometry, allowing deformations that only respect isometry locally. It prunes inconsistent correspondences between two whole non-rigid surfaces, but does not support multiple non-rigid parts.

Starting from the ideas in [TMRL14], we show how to make further use of the isometric consistency matrix. Embedding the matrix into a low dimensional space using a diffusion map, many small clusters of point correspondences are found. Each of these clusters indicates some point correspondences that are in an isometrically consistent relationship, but they are not globally isometrically consistent between the two surfaces as a whole. (For example, small parts of the front left leg and the back right leg of a horse are consistent similar, but the legs are not consistent as a whole.) These correspondences, however, provide useful hints where similar non-rigid parts may be found. This observation helps to solve our first research problem: given a query part, can we identify and locate similar non-rigid parts on the same or another mesh? Our approach is to extract reliable small clusters of point correspondences through diffusion pruning.

Having extracted these clusters, typically forming a small set of sparse correspondences on the part of the surface concerned, we use these to obtain dense correspondences, solving the second research problem. This is done using an iterative propagation-and-prune strategy.

In summary, our contributions in this paper are two techniques based on the use of an assumption of local isometry: (i) given a query shape, to identify multiple similar *nonrigid* parts of the same or another mesh, and (ii) to establish dense pointwise correspondences between such similar parts. Our techniques are simple, efficient, and can handle arbitrary topology. They are fast enough to allow interactive queries for moderate sized meshes. An example is shown in Figure 1. The query part is the right hand of the dinopet in a reference pose and the target mesh is the whole dinopet in a different pose. Dense pointwise correspondences between the query part and both the left and right hands are established automatically.

2. Related work

Our techniques relate to partial matching and partial shape correspondence. We summarise works most closely related to ours, concentrating on techniques for *non-rigid* shapes.

Local geometry signatures Establishing point or patch correspondences between parts can be done through local geometric signatures. Geometric signatures are descriptive feature vectors that describe the local geometry of a shape. They are usually based on local properties, and therefore are less sensitive to holes. Notable examples include spin images and SHOT signatures [TSDS10]. Such signatures were designed for rigid shape regions, but have also been adopted for non-rigid shape matching. Geometric signatures have also been developed for non-rigid shapes, e.g. heat kernel signatures [SOG09]. However, point correspondences established by geometric signatures are usually sparse, do not cover the whole part, and are often not consistent within the part. Further pruning steps [ZSCO*08] are usually required to obtain consistent point correspondences. In this work, we use SHOT signatures to obtain initial point correspondences. We

138

then find similar non-rigid parts by clustering and pruning, and establish consistent correspondences between parts.

Partial matching and retrieval In the non-rigid shape retrieval literature, the main research focuses are to develop deformation invariant representations, robust descriptors measuring similarity, and fast retrieval through indexing. A few approaches can handle partial matching. A good review of such techniques can be found in [TV08]. Recent work [DLL*10] uses the heat kernel signature function to handle incomplete shape matching and retrieval, while [PBB12] considers partial shape matching in the feature domain to avoid the need to determine point-wise correspondences. In all such techniques, developing a good similarity measure for partial shapes is the main goal. Finding the correspondences between non-rigid parts is often limited to components or parts. Finding correspondences between non-rigid parts down to individual points, which we address in this paper, is often not supported.

Shape registration and correspondences Finding correspondences between non-rigid surfaces has a long research history; a recent survey [vKZHCO10] gives a detailed account. In general, given two input surfaces, these techniques find the reliable point correspondences (or a correspondence map) between them. There are many different types of techniques, depending on the deformation model assumed [TCL*13]. An important category assumes isometric deformation. Under an isometry, the distance between two points on one shape is the same as the distance between the two corresponding points on another. There are various ways to establish correspondences under this assumption. Notable examples include [HAWG08, BBK*10, TBW*11, SNB*12]. Some [ASP*04, LSP08] can work on incomplete surfaces. All such approaches assume that the two input surfaces correspond as a whole (and have substantial overlap). They do not seek to identify similar non-rigid parts, and furthermore, some have topological restrictions (e.g. [LF09, KLF11]).

Symmetry detection Symmetry detection on surfaces is a closely related topic. A good review can be found in [MPWC13]; we focus our discussion on partial intrinsic symmetry. [RBBK10] casts partial symmetry analysis of non-rigid shapes as an optimisation problem to find global and partial symmetries. [MBB10] uses multidimensional scaling to 'unbend' a non-rigid shape into canonical form so that repeating patterns can be detected using rigid transforms. [LTSW09] uses a probabilistic framework for partial intrinsic symmetry, considering the marginal distribution of point correspondences to reduce computational complexity. [BBW*09] uses salient edges for symmetry detection and correspondence establishment. Many of these assume an underlying (near-)isometric deformation using global geodesic distances. Our technique is not a symmetry detection technique. Unlike these works, our approach is based on local isometry which provides a more relaxed assumption about deformation. This also allows us to avoid

the expensive precomputation and storage of global pairwise geodesic distances.

Partial shape point correspondences Our work concerns finding partial correspondences between non-rigid shapes. Some work, including ours, supports finding correspondence between two surfaces where one is significantly smaller than another. [BBCK09] casts partial matching as a sophisticated multi-criterion optimisation problem, maximising similarity and significance. This is computationally expensive. [OMMG10] defines one point correpondence to obtain sparse partial correspondences. However, those sparse correspondences might be inconsistent if the deformation is large, and may require manual initialisation. [SY14] uses coarsely sampled extremities to identify parts by ranking potential partial correspondence maps. The technique assumes near-isometric deformation and may fail if the deformation is large. These technique assumes one single similar part. Our technique identifies multiple parts and establish dense point correspondences between parts. Above all, ours assumes a locally isometric deformation, which is more relaxed than existing works.

3. Background

Our technique is inspired by [HAWG08, TMRL14] which are variants of the pioneering work in [LH05]. To be selfcontained, we briefly describe them here. Both [HAWG08, TMRL14] are pruning techniques which, given a set of candidate point correspondences between two whole surfaces, retain consistent point correspondences and discard inconsistent ones. Both cast the matching problem to retain consistent point correspondences as a quadratic assignment problem, which is NP-hard. [HAWG08] approximates the solution by relaxing binary assignment to a continuous setting and considers the associated eigenproblem. Geodesic distances are used to define pairwise isometric consistency. The elements in the first eigenvector of the eigensolution measure how point correspondences, under isometric deformation, are related to one another. Using spectral graph theory, these point correspondences form the largest cluster in the spectral domain. Instead of using global geodesic distances to define consistency between any pair of point correspondences (no matter how far apart), [TMRL14] considers point correspondences that can be related by *local* geodesic distances only. This reduces the computational overhead; it works well for non-rigid shapes where longer geodesic distances are less reliably invariant under deformation.

Let Q (query) and T (target) be two surfaces for which a set of raw input point correspondences is already available. [TMRL14] begins with the computation of geodesic distances for all points within a small local disc of each point. Let Δ_x^{δ} be a geodesic disc of vertices centred on vertex x: $\Delta_x^{\delta} = \{q | d_g(q, x) \le \delta\}$, where $q \in Q$, and $d_g(\cdot, \cdot)$ is the geodesic distance. δ is a suggested local radius, typically set as $\delta = 0.05D$, where D is the maximum of the diameter of

© 2014 The Author(s)

Computer Graphics Forum © 2014 The Eurographics Association and John Wiley & Sons Ltd.

Q and T. This radius models the size used for local regions within which point correspondences should be isometrically consistent. Similar geodesic discs are computed for T.

Next, the sparse isometric consistency matrix *K* (of size $|C| \times |C|$) is built, where *C* is the set of raw correspondences. The matrix *K* encodes the pairwise isometric consistency between *pairs* of correspondences as follows:

$$K(a,b) = \begin{cases} \left(\frac{k_{ab} - c_0}{1 - c_0}\right)^2 & \text{if } a \neq b \text{ and } (3) \text{ is satisfied} \\ 0 & \text{otherwise} \end{cases}$$
(1)

where $a = (q_i, t_u)$ and $b = (q_j, t_v)$ are two point correspondences: $q_i, q_j \in Q$ and $t_u, t_v \in T$. k_{ab} is a term which measures the isometric consistency of a and b and is based on consistency of their geodesic distances:

$$k_{ab} = \min\left(\frac{d_g(q_i, q_j)}{d_g(t_u, t_v)}, \frac{d_g(t_u, t_v)}{d_g(q_i, q_j)}\right).$$
 (2)

A pair of point correspondences *a* and *b* is only included in the matrix if the following two criteria are satisfied:

$$k_{ab} \ge c_0$$
 and $q_j \in \Delta_{q_i}^{\delta}$ and $t_v \in \Delta_{t_u}^{\delta}$, (3)

where $c_0 \in [0, 1]$ is a threshold that controls the sparsity of *K*. When c_0 is high, *K* is only based on high consistency relationships, while making c_0 smaller relaxes the isometric assumption. $c_0 = 0.7$ is a suggested value [HAWG08]. Eqn. 3 restricts *K* to storing only *local* isometric relationships.

[TMRL14] then computes the row sums of the matrix *K* as the local isometric consistency score for each correspondence $a: \pi(a) = \sum_b K(a,b)$. $\pi(a)$ indicates the confidence that *a* belongs to a cluster that respects isometric consistency, and is similar to the eigenvector computed in [HAWG08]. The higher the value, the more consistent it is with other correspondences and the more central it is within the cluster. A low value indicates that *a* may be an outlier. The algorithm then processes each correspondence *a* in decreasing order of $\pi(a)$, checking whether it is locally isometric consistent with existing ones in a greedy manner. After checking all correspondences, [TMRL14] outputs a set of correspondences *C'* which are globally consistent.

4. Observations and Overview

Given a query part, our goals are to (i) identify similar nonrigid parts, and (ii) establish point correspondences between them. Given the query part Q and the target surface T (|Q| < |T|), we generate some point correspondences C between them using SHOT signatures. Our research then begins by considering the matrix K obtained from [TMRL14]. K, from another viewpoint, encodes a correspondence connectivity graph G(C, K) where point correspondences $a, b \in C$ are vertices and an edge exists between them if K(a,b) > 0 (see Eqn. 1) indicating they have some degree of isometric consistency. This graph can be embedded into a low dimensional space through a diffusion map [CL06]. Our observations show that the graphs have many clusters of vertices in this embedding. Sets of consistent point correspondences between the query part Q and the mesh T are contained in these clusters. One example is shown in Figure 1(b). The horizontal and vertical lines are two clusters of graph vertices (i.e. point correspondences), each corresponding to one of the hand matches in Figure 1(c). This suggests that similar multiple non-rigid parts may be identified in this space. Section 5 discusses a simple algorithm to identify these similar parts. However, as shown in Figure 1(c), the initial set of point correspondences usually only cover part of the query surface. To obtain dense correspondences, we iteratively find new correspondences that satisfy local isometry constraints, and use [TMRL14] to remove any inconsistent ones. This algorithm is considered in Section 6.

5. Identifying similar non-rigid parts

As explained above, it might be expected that simple clustering in the embedding space would easily find all similar parts. However, several problems arise with this approach. Firstly, multiple sets of consistent point correspondences may exist that all respect isometry. For example, for a cylinder, all rotational copies are potential matches. All such copies are isometrically related (K(a,b) > 0), leading to a spread out set of correspondences rather than clusters. Secondly, it may not be easy to identify the intrinsic dimension of the embedding space. Finally, computing an eigendecomposition to obtain the embedding space is slow especially when |C| is large ($|C| \approx 10 - 500K$). We therefore make use of the pruning idea from [TMRL14] to help us identify the most relevant copies efficiently without computing eigen-embeddings.

5.1. Goal, input and output

The goal of the algorithm summarised in Algorithm 1 is, given a query, to quickly identify potentially similar non-rigid parts—our first research problem.

For this step, we assume we already have the following information (i) a query part Q and target mesh T, (ii) their geodesic discs Δ^{δ} of size δ (our experiments used $\delta = 0.25D$ where D is the diameter of Q), (iii) a set of raw input point correspondences C, (iv) the isometric consistency matrix K(Eqn 1), and (v) ψ , the number of similar non-rigid parts to look for. Here, we assume the user specifies ψ , which would be reasonable in a geometry editing scenario like that described in Section 1. However, in Section 9, we describe a possible extension to automatically identify ψ .

To compute *C*, we find the top $r \mathscr{H} |Q|$ most similar point matches to *T* according to SHOT signatures for each vertex $q \in Q$ (i.e., $|C| = r \mathscr{H} |Q|^2$). We do not find point matches from *T* to *Q* because most are mismatches if *Q* is small. Other signatures could be used instead of SHOT signatures, but they give good results and are fast to determine. In all

Algorithm 1 Extract multiple sets of isometric point correthe relative size of geodesic discs a consistency threshold set of input point correspondences isometric consistency matrix the number of similar parts ψ sets of point correspondences - Initialisation ▷ Depth-first search

 $i \leftarrow i+1$ 7: 8: end while 9: ⊳ — - Multiple Point Correspondences 10: $\psi' \leftarrow \min(\psi, i)$ 11: **parfor** $k = 1 \rightarrow \psi'$ **do** ▷ Parallelization 12: $j \leftarrow 0$ $C_{k,j} \leftarrow \mathrm{DP}(C_k^{cc})$ ▷ Diffusion Pruning [TMRL14] 13: 14: $C_k^{cc} \leftarrow C_k^{cc} \setminus C_{k,j}$ ▷ Conflict correspondences $\tilde{C_k^{cc}} \leftarrow \tilde{C_k^{cc}} \setminus R(\tilde{C}_{k,j}, C_k^{cc})$ 15: ▷ Conflict regions 16: $j \leftarrow j + 1$ 17: end parfor - Aggregation 18: ⊳ --19: for $i \leftarrow 1, k = 1 \rightarrow \psi'$ do 20: $j \leftarrow 1$ while $C_{k,j} \neq \emptyset$ do 21: 22: $C'_i \leftarrow C_{k,j}, j \leftarrow j+1$ ▷ Aggregation end while 23: $i \leftarrow i + 1$ 24: 25: end for 26: $\psi \leftarrow \min(i-1, \psi)$ 27: $\operatorname{sort}(C')$ \triangleright in descending order of $|C'_i|$ 28: return $C'_{1...\psi}$

our examples, we used $r \in [1, 10]$ which depends on the descriptiveness of SHOT and the number of identified parts to query. For higher ψ , a larger r is required to allow for potential mistakes in the input correspondences.

The output is a set of clusters of point correspondences C'_1, \ldots, C'_{ψ} which broadly indicate the locations of parts.

5.2. Algorithmic details

spondences $C'_{1...\psi}$

Input:

Output:

2: $i \leftarrow 0$

3: $C_{left} \leftarrow C$

1: ⊳

5:

6:

δ

 c_0

C

K

ψ

 $C'_{1...\psi}$

4: while $|C_{left}| > 0$ do

 $C_i^{cc} \leftarrow \mathrm{DFS}(G(C_{left}, K))$

 $C_{left} \leftarrow C_{left} \setminus C_i^{cc}$

Algorithm 1 has three stages: initialisation, extraction, and aggregation, as detailed below.

1. Initialisation (Line 1-)

Recall that G(C, K) consists of vertices (point correspondences a, b) and edges (when K(a, b) > 0). This G(C, K)often contains disjoint sets of vertices (point correspondences) where point correspondences in one set are not related to another set. For example, in Figure 1(c), there

© 2014 The Author(s)

Computer Graphics Forum © 2014 The Eurographics Association and John Wiley & Sons Ltd

are five groups of point correspondences: head, two legs and two hands. Point correspondences that match a hand to the head are not related to those that match a hand to the legs due to local isometry. These sets of correspondences can be separated by connected component analysis on G(C, K). Thus, we first identify point correspondences that are spatially (linked by local geodesic discs) and geometrically (according to SHOT signature) close, i.e., $C = \bigcup_i C_i^{cc}$ where C_i^{cc} is one set of point correspondences identified by the connected component analysis.

2. Multiple Correspondence Extraction (Line 9-)

As discussed in Section 5, multiple similar copies may exist in each of these components, and simple clustering does not help. We therefore apply the diffusion pruning algorithm [TMRL14] to find the largest cluster of point correspondences $C_{k,i}$ (Line 13) within each C_k^{cc} iteratively. [TMRL14] targets consistent correspondences across the whole surface in a greedy manner. There is a case that new point correspondences cannot be checked for local isometric consistency due to the use of small geodesic discs. Global geodesic distances, though does not respect large deformation well [TMRL14], has to be used. These correspondences may affect subsequent propagation process. To ensure the identified point correspondences C'_i are truly isometrically consistent, we modify the diffusion pruning algorithm such that whenever a point correspondence is accepted, we note the source point correspondences to which it is locally isometrically consistent. When the pruning step finishes, we trace all point correspondences to their sources. Point correspondences with more than one source are discarded, and we return the group having the largest number of point correspondences from a single source.

We also do not want the same region to be extracted twice. E.g., for the cylinder example, we would like to extract the first set of consistent correspondences, but ignore all slightly rotated copies. To do so, our idea is to exclude the corespondences having one end point falls on the extracted part (the conflicting region) from the next pruning step (Lines 14–15). Note $C_{k,j}$, however, covers only part of the query surface and are not sufficient to define such conflicting region. We therefore use a heuristic to define the conflicting region $R(C_{k,i}, C_k^{cc}) =$ $\{(p,q) \in C_k^{cc} \mid N(q,v) \le \tau, (u,v) \in C_{k,j}\} \text{ where } p, u \in Q \text{ and } q, v \in T. \ N(q,v) \text{ counts the distance (in terms of }$ rings of neighbors) between two vertices. We set experimentally $\tau = 1.5 \times \sqrt{\text{Area}(Q)/\text{Area}(S)}$ where S = $\{v|(u,v) \in C_{k,j}\}$ is a set of vertices on the mesh T that correspond well to the query part, and Area(S) defines the surface area associated with S while Area(Q) is the area of the query part. Intuitively, if the extracted region (defined by point correspondences) is an exact copy of Q, then Area(Q) = Area(S). Since we seek similar parts and $Area(S) \leq Area(Q)$, we define the neighbourhood closeness as a function of their area. The ratio is based on an assumption that the extracted $C_{k,j}$ is located at the cen-

Algorithm 2 Establish dense correspondences **Input:** $C'_{1...\psi}$ ψ sets of sparse point correspondences **Output:** $\hat{C}'_{1...\psi}$ ψ sets of dense point correspondences 1: parfor $i = 1 \rightarrow \psi$ do ▷ Parallelization $\hat{C'}_i \leftarrow C'_i$ 2: while $\left| \hat{C'}_i \right|$ is changing do 3: $\hat{C}'_i \leftarrow \hat{C}'_i \cup \text{Propose}(\hat{C}'_i)$ adjust K4: ▷ Propose 5: $\hat{C'_i} \leftarrow \mathrm{DP}(\hat{C'_i}, K)$ 6: ▷ Prune 7. end while 8: end parfor 9: return $\hat{C}'_{1...\psi}$

tre of the whole part, and the part is of disc topology (since we have no prior knowledge of the part). This assumption is highly simplified but works well because as long as some of these close point correspondences are removed, these regions will not provide the highest number of correspondences in the next pruning round. This suffices to avoid finding duplicated regions, and it works well in practice. The whole process can be parallelised by processing each C_k^{cc} independently (Line 11).

3. Aggregation (Line 18–)

Finally, all point correspondences $C_{k,j}$ are collected into a single array C' and sorted in descending order of $|C'_i|$. The first $C'_{1,\psi}$ are returned where ψ is the user defined number of similar parts. $C'_{1,\psi}$ are used in the subsequent process of establishing dense point correspondences.

6. Dense correspondence establishment

Our next goal is to obtain sets of dense point correspondences $\hat{C'}_{1...\psi}$ from correspondences $C'_{1..\psi}$. Note that C'_i is a set of isometrically consistent point correspondences, which often only cover a small part of the query surface. Since local geodesic distance may be more effective in handling deformation [TMRL14], our idea is to grow dense correspondences from $C'_{1..\psi}$ using only the precomputed local geodesic discs. Algorithm 2 iteratively alternates three steps until the correspondence set remains unchanged:

- 1. Find potential point correspondences that respect isometric consistency (Section 6.1).
- 2. Compute a new isometric consistency matrix *K* (Section 6.2).
- 3. Select high quality point correspondences by diffusion pruning [TMRL14], using *K*.

Again Algorithm 2 can be parallelised. We now consider it in detail.

6.1. Proposing Isometric Correspondences

Given an input set of point correspondences C'_i , we find new point correspondences that respect local isometric consis-

tency. For each vertex $v \in Q$, we generate *h* point correspondences as follows (Propose() in Algorithm 2):

- Find a subset a = {q_i,t_u} ∈ A ⊂ C'_i such that v falls in their geodesic discs, i.e., v ∈ Δ^δ_{q_i}.
- Find all k ∈ T such that k also falls in all the geodesic discs related to a, i.e., k ∈ Δ^δ_{tu} ∀a = {q_i,t_u} ∈ A.
- 3. Compute the local isometric score μ for each k: $\mu(v,k) = \sum_{a \in A} |d_g(v,q_i) - d_g(k,t_u)|.$
- Generate *h* new point correspondences {*v*,*k*} using *k* that corresponds to the *h*-smallest μ(*v*,*k*).

These four steps generate a new set of point correspondences $\hat{C}'_i \leftarrow \hat{C}'_i \bigcup$ Propose (\hat{C}'_i) . We use h = 10 in all our examples.

6.2. Isometric Correspondence Selection

 C'_i may contain conflicting point correspondences. We use diffusion pruning [TMRL14] to select isometrically consistent ones. Further, we may use additional information to guide the point correspondences establishment by modifying the kernel *K* (Section 3) before diffusion pruning is applied.

Geometric constraints The diagonal entries of *K* are zeros by construction [TMRL14]. This is useful if the relation between *Q* and *T* comprises mostly isometric deformation where only pairwise isometric relationships $K(a,b), a \neq b$ need to be encoded. On the other hand, if *Q* and *T* are similar in terms of local geometric signatures, we may exploit this information in the diagonal entries of *K* (as has been done for images [LH05]). Let $\mathbf{s}(a)$ be the SHOT signature and $d_s(a = \{q_i, t_u\}) = (\sum (\mathbf{s}(q_i) - \mathbf{s}(t_u))^2)^{-1}$. We redefine

$$K(a,a) = \frac{d_s(a) - \min_{b \in C_j} d_s(b)}{\max_{b \in C_i} d_s(b) - \min_{b \in C_i} d_s(b)}$$

Modifying K(a,a) as above means that the diffusion framework allows diffusion (or random jumps) back to point correspondence *a* with probability proportional to the similarity of geometric signatures between $\mathbf{s}(q_i)$ and $\mathbf{s}(t_u)$. This allows both isometric deformation and geometric information to be considered in a single framework.

7. Results

Here, we evaluate our technique with various examples, with and without ground truth. We use colour coding to demonstrate the quality of the point correspondences. Table 1 gives all timings and resolutions of parts and meshes.

Dinopet Figure 1 demonstrates a simple isometric example. The right hand of the dinopet was used as a query, and we seek the two most similar copies on a *deformed* dinopet. The results show that hands which contain multiple smaller subfeatures (the fingers) can be matched without problems. Despite the SHOT signature indicating many potential matches on the feet and legs, we are able to recover the left and right hands as the two largest sets of point correspondences.



Figure 2: Hand (Olivier) |C| = 192750. There is no exact copy of the query part in this example due to deformation.



Figure 3: Centaur |C| = 34429. There is no exact copy of the query part in this example.

Olivier's fingers Figure 2 demonstrates an example of mostly *isometric* deformation of similar fingers. Olivier's hand is a publicly available model. We took the index finger of the undeformed hand as the query and sought similar parts on the *deformed* hand. In this example, our technique automatically identified four fingers in the deformed hand despite the size differences, and established relatively good point correspondences between them.

Centaur and Geometry Transfer Figure 3 demonstrates another example of querying parts from a whole model. We took the left rear leg of a centaur model and sought similar parts from another *deformed* centaur model (TOSCA dataset [BBK08]). Four legs were identified, including the two front legs. Note that the two front legs differ in bending direction. Figure 4 shows an application in which geometry editing was transferred through reliable correspondences enlargement of the query hoof was propagated to all hooves of the target model.

Bas-Relief and Geometric Constraints In Figures 5 and 6, we attempted to identify copies of a given query region of a mesh representing a scanned periodic bas-relief. These is a challenging problem because all patterns are slightly different due to distortions resulting from the manufacturing process used. The geometric constraints (Section 6.2) were



Figure 4: Simple geometry editing (scaling of hoof) on the query part can be propagated to the other model, since the hooves are all well matched.



Figure 5: The first 5 automatically identified surface patches, given a query. The query established the largest number of point correspondences with itself, whilst other copies differ from the query in terms of small non-rigid deformations due to the manufacturing process.

effective and our technique identifies multiple copies and establishes reliable point correspondences automatically.

Comparison to State-of-the-Art Figure 7 compares our approach to another recent partial correspondence technique. The queried part has many large deformations. [SY14] obtains a map on the front body due to lower isometric distortion score. This example also shows that our technique can support large parts. Note, however, that [SY14] supports scale invariant partial matching whilst our method does not.

Accuracy Analysis Here we use true positive rate to measure accuracy. Centaurs models and ground truth information is provided by the TOSCA dataset [BBK08]. We evaluated our technique on a scene consisting of 6 centaurs (Figure 8). The query part was a cropped back right leg of the fourth centaur (from the left). This experiment was challenging because there were many possible matches, and the topology of the query part was almost cylindrical. We queried the first 12 copies. Our technique identified 6 non-rigid copies of the back right leg as the most similar parts. True positive are colored blue and false positive red. Most point correspondences are correct except at boundaries, which are slightly shifted. The results are to be expected as our technique grows point correspondences from good ones inside the part, and there are fewer isometric con-

© 2014 The Author(s) Computer Graphics Forum (© 2014 The Eurographics Association and John Wiley & Sons Ltd.



Figure 6: The first 6 automatically identified surface patches, given a predefined query. The rightmost patch is a smoothed version of the query in which the geometric surface features are less distinct. Since our technique depends on reliable geometric features, the point correspondences of the rightmost patch are misaligned, and the second rightmost patch is not found—there are more point correspondences for the leftmost patch. Also, the leftmost patch is identified twice as our region growing technique fails to avoid overlapping patches.



Figure 7: Comparison with another leading method.

straints near the boundary. The average percentage of true positives is 93.56%. The next 6 non-rigid parts are coloured with respect to the established point correspondences; all are back left legs with reasonable correspondence.

Timings The running time of our technique depends on several factors: the size of the query parts and queried models, the number of feature point correspondences found, and the number of similar parts to look for. Times for our examples are shown in Table 1. Given a medium sized mesh (say a centaur with around 10000 vertices) and query (a part with 694 vertices), our technique can retrieve all related legs in under 15s. Unlike other techniques, ours does not require down-sampling of the mesh (to 1000-3000 vertices), nor any sophisticated edge detection technique. Excluding preprocessing time (which include SHOT feature, initial correspondences, FLANN and geodesic distance computation time), our algorithm can run in times which are acceptable for interatcive applications for medium sized meshes (taking between 10s and 60s). Our core implementation and timing use MATLAB and its Parallel Computation Toolbox.

8. Discussion

Our technique is built on [TMRL14] which is a pruning technique that can only *remove* incorrect point correspondences. It does not support part correspondences, and makes the strong assumption that all good correspondences form a single cluster linking the two surfaces. Our novel technique here analyses the isometric consistency matrix to look for non-rigid parts. In particular, we focus on establishing *new* correspondences. The strong assumption of connection of all good correspondences is removed by introducing a source clustering step, which is essential to obtain good correspondences during propagation.

Our technique is significant in its flexibility to support various query types: textured patches and components which can be small or large. The assumed locally isometric deformation is more flexible than global isometry or rigid partial matching [GCO06].

In fact, the targeted editing application (see Section 1) may be implemented in two steps, (i) segmenting a mesh into parts [HKG11, SvKK*11, KLM*13] for partial matching/retrieval [FS06, SSS*10, DLL*10] and (ii) establishing a correspondence map between each part using a whole surface correspondence technique. However, such an approach requires answers to questions like: 'Does over-segmentation affect consistency of point correspondences?' 'Which part(s) are the best matches to establish consistent correspondences with?' 'Is the assumed deformation too restricted (general)?' and 'Can the whole surface correspondence technique handles part of arbitrary topology?' This paper suggests that by assuming local isometry, a simple and efficient approach can be developed to provide a reasonable solution. Consistent correspondences from multiple segments (e.g. finger segments) of a part (e.g. a finger) help solve such partial correspondence problems (in term of larger cluster), rather than making it more complicated.

A further advantage over existing techniques is that our method is readily parallelised and is fast, with the benefit of avoiding costly computation of eigenfunctions, canonical forms (e.g. MDS), and global geodesic distances.

9. Limitations

Our technique depends on the ability of the geometric signatures to generate correspondences. If their descriptiveness is not strong, we may not find good parts and point correspondences. An example is shown in Figure 6 where one patch is not identified. Currently we use SHOT signatures because of its reliability and efficiency [TSDS10]. It is possible to replace it with others (e.g. Heat Kernel Signature [SOG09]).



Figure 8: The first 12 similar surface patches (with open cylinder topology) identified from a scene of 6 centaurs. The query part is cropped from the back right leg of the fourth centaur (from the left). Our technique identified 6 non-rigid copies of the back right leg as the most similar parts. True correspondences are blue (true positive) and red (false positive, all are on patch boundaries). The percentages of true positives are: 88.73%, 97.24%, 88.48%, 94.38%, 92.55%, 100% (from the left). The next 6 non-rigid parts are coloured with respect to the established correspondences, and are all back left legs.

Time	Fig 3	Fig 1	Fig 8	Fig 2	Fig 5	Fig 6
#Vertices (part/model)	694/10002	1132/13322	594/60012	1285/12782	5102/38187	4758/39934
#Point correspondences	34700	226400	11880	192750	510200	237900
Pre-processing	4.13	14.11	52.08	9.83	66.81	102.81
Connected components	0.49	6.33	0.08	7.41	30.41	21.83
Clusters identification	3.68	19.92	12.93	38.82	238.19	210.05
Propagation	5.69	14.29	20.58	7.13	201.45	335.87

Table 1: Experiment timing in seconds. Machine configuration: Intel i7 CPU 3.2GHz with 32GB memory.



Figure 9: Number of point correspondences versus cluster id for Figure 8.

We use a (locally) isometric correspondence generator to generate raw point correspondences. This limits our technique to establish point correspondences to those that favour (locally) isometric consistency. If the deformation of the parts substantially deviates from isometry, the approach will not succeed. [TMRL14] is a greedy technique, and our approach can consequently fail due to cascading of errors.

Our technique currently requires several parameter settings, including: c_0 and δ which control the expected level and locality of isometric defomation, r% which controls the number of initial correspondences when looking for clusters, and h which controls the speed and quality of propagated correspondences. All these parameters are determined experimentally. The user must also specify the parameter ψ which is the number parts to be found which match the query. An alternative would be to try to determine ψ by forward search [GMGP05]. For example, a plot of the number of point correspondences versus cluster id for Figure 8 is shown in Figure 9. Two sudden jumps are easily identified after the first 6, and first 12, clusters which correspond to the first 6 matches to right back legs, and the next 6 to left back legs. We hope to further develop our methods to provide a fully parameter-free technique in future.

10. Conclusions

This paper has presented a technique which can automatically identify multiple parts within the same model or other models that match a query shape via a (locally) isometric deformation. For medium sized meshes, the technique takes under one minute. Our technique flexibly allows users to define a query model as a small part (e.g. a finger) or large part (e.g. a hand). The returned results and point correspondences are meaningful. It is suitable for interactive applications such as intelligent geometry editing. We hope to generalise the approach to more general deformations in the future.

Acknowledgements

We thank various reviewers for their constructive comments, Y. Sahillioglu for allowing us to use his code, and Bronstein et al. for the TOSCA dataset. This work was supported by the Engineering and Physical Sciences Research Council [grant number EP/I000100/1].

References

- [ASP*04] ANGUELOV D., SRINIVASAN P., PANG H.-C., KOLLER D., THRUN S., DAVIS J.: The correlated correspondence algorithm for unsupervised registration of nonrigid surfaces. *NIPS 17* (2004), 33–40. 3
- [BBCK09] BRONSTEIN A. M., BRONSTEIN M. M., CARMON Y., KIMMEL R.: Partial similarity of shapes using a statistical significance measure. *IPSJ TCVA 1*, 0 (2009), 105–114. 3
- [BBK08] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: Numerical geometry of non-rigid shapes. Springer Verlag, 2008. 7
- [BBK*10] BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R., MAHMOUDI M., SAPIRO G.: A gromov-hausdorff framework with diffusion geometry for topologically-robust non-rigid shape matching. *IJCV* 89 (2010), 266–286. 3
- [BBW*09] BERNER A., BOKELOH M., WAND M., SCHILLING A., SEIDEL H.-P.: Generalized Intrinsic Symmetry Detection. Tech. rep., Max-Planck-Institut fur Informatik, Tubingen University, 2009. 3
- [CL06] COIFMAN R. R., LAFON S.: Diffusion maps. Applied and Comp. Harmonic Anal. 21, 1 (2006), 5–30. 4
- [DLL*10] DEY T. K., LI K., LUO C., RANJAN P., SAFA I., WANG Y.: Persistent heat signature for pose-oblivious matching of incomplete models. *CGF* 29, 5 (2010), 1545–1554. 3, 8
- [FS06] FUNKHOUSER T., SHILANE P.: Partial matching of 3d shapes with priority-driven search. In Proc. SGP (2006), pp. 131– 142. 8
- [GCO06] GAL R., COHEN-OR D.: Salient geometric features for partial shape matching and similarity. ACM TOG 25, 1 (2006), 130–150. 8
- [GMGP05] GELFAND N., MITRA N. J., GUIBAS L. J., POTTMANN H.: Robust global registration. In Proc. SGP (2005), pp. 197–206. 9
- [HAWG08] HUANG Q.-X., ADAMS B., WICKE M., GUIBAS L. J.: Non-rigid registration under isometric deformations. CGF 27, 5 (2008), 1449–1457. 2, 3, 4
- [HKG11] HUANG Q.-X., KOLTUN V., GUIBAS L.: Joint shape segmentation with linear programming. ACM TOG 30, 6 (2011), 125:1–125:12. 8
- [KLF11] KIM V., LIPMAN Y., FUNKHOUSER T.: Blended intrinsic maps. ACM TOG 30, 4 (2011), 79:1–79:12. 3
- [KLM*13] KIM V. G., LI W., MITRA N. J., CHAUDHURI S., DIVERDI S., FUNKHOUSER T.: Learning part-based templates from large collections of 3d shapes. ACM TOG 32, 4 (2013), 70:1–70:12. 8
- [LF09] LIPMAN Y., FUNKHOUSER T.: Möbius voting for surface correspondence. In ACM TOG (2009), ACM, pp. 1–12. 3
- [LH05] LEORDEANU M., HEBERT M.: A spectral technique for correspondence problems using pairwise constraints. In *Proc. ICCV* (2005), vol. 2, pp. 1482–1489 Vol. 2. 3, 6

- [LSP08] LI H., SUMNER R. W., PAULY M.: Global correspondence optimization for non-rigid registration of depth scans. *CGF* 27, 5 (2008), 1421–1430. 3
- [LTSW09] LASOWSKI R., TEVS A., SEIDEL H.-P., WAND M.: A probabilistic framework for partial intrinsic symmetries in geometric data. In *Proc. ICCV* (2009), pp. 963–970. 3
- [MBB10] MITRA N. J., BRONSTEIN A. M., BRONSTEIN M. M.: Intrinsic regularity detection in 3d geometry. In *Proc.* ECCV (2010), pp. 398–410. 3
- [MPWC13] MITRA N. J., PAULY M., WAND M., CEYLAN D.: Symmetry in 3d geometry: Extraction and applications. *CGF 32*, 6 (2013), 1–23. 3
- [OMMG10] OVSJANIKOV M., MÉRIGOT Q., MÉMOLI F., GUIBAS L. J.: One point isometric matching with the heat kernel. CGF 29, 5 (2010), 1555–1564. 2, 3
- [PBB12] POKRASS J., BRONSTEIN A. M., BRONSTEIN M. M.: Partial shape matching without point-wise correspondence. NM-TMA (2012). 3
- [RBBK10] RAVIV D., BRONSTEIN A. M., BRONSTEIN M. M., KIMMEL R.: Full and partial symmetries of non-rigid shapes. *IJCV 89* (2010), 18–39. 3
- [SNB*12] SOLOMON J., NGUYEN A., BUTSCHER A., BEN-CHEN M., GUIBAS L.: Soft maps between surfaces. CGF 31, 5 (2012), 1617–1626. 3
- [SOG09] SUN J., OVSJANIKOV M., GUIBAS L. J.: A concise and provably informative multi-scale signature based on heat diffusion. In *Proc. SGP* (2009), pp. 1383–1392. 2, 8
- [SSS*10] SHAPIRA L., SHALOM S., SHAMIR A., COHEN-OR D., ZHANG H.: Contextual part analogies in 3d objects. *IJCV* 89, 2-3 (2010), 309–326. 8
- [SvKK*11] SIDI O., VAN KAICK O., KLEIMAN Y., ZHANG H., COHEN-OR D.: Unsupervised co-segmentation of a set of shapes via descriptor-space spectral clustering. ACM TOG 30, 6 (2011), 126:1–126:10. 8
- [SY14] SAHILLIOGLU Y., YEMEZ Y.: Partial 3d correspondence from shape extremities. CGF (2014). 3, 7, 8
- [TBW*11] TEVS A., BERNER A., WAND M., IHRKE I., SEIDEL H.-P.: Intrinsic shape matching by planned landmark sampling. *CGF 30*, 2 (2011), 543–552. 2, 3
- [TCL*13] TAM G. K. L., CHENG Z., LAI Y.-K., LANGBEIN F., LIU Y., MARSHALL D., MARTIN R. R., SUN X., ROSIN P. L.: Registration of 3d point clouds and meshes: A survey from rigid to non-rigid. *IEEE TVCG 19*, 7 (2013), 1199–1217. 3
- [TMRL14] TAM G. K.-L., MARTIN R. R., ROSIN P. L., LAI Y.: Diffusion pruning for rapidly and robustly selecting global correspondences using local isometry. ACM TOG 33, 1 (2014), 4. 2, 3, 4, 5, 6, 8, 9
- [TSDS10] TOMBARI F., SALTI S., DI STEFANO L.: Unique signatures of histograms for local surface description. In *Proc.* ECCV (2010), vol. 6313, pp. 356–369. 2, 8
- [TV08] TANGELDER J. W., VELTKAMP R. C.: A survey of content based 3d shape retrieval methods. *MM Tools Appl. 39*, 3 (2008), 441–471. 1, 3
- [VKZHCO10] VAN KAICK O., ZHANG H., HAMARNEH G., COHEN-OR D.: A survey on shape correspondence. In EG STAR Report (2010). 1, 3
- [ZSCO*08] ZHANG H., SHEFFER A., COHEN-OR D., ZHOU Q., VAN KAICK O., TAGLIASACCHI A.: Deformation-driven shape correspondence. CGF 27, 5 (2008), 1431–1439. 2

© 2014 The Author(s) Computer Graphics Forum © 2014 The Eurographics Association and John Wiley & Sons Ltd.