

Hierarchical Shape Matching for Temporally Consistent 3D Video

Chris Budd

c.budd@surrey.ac.uk

Peng Huang

p.huang@surrey.ac.uk

Adrian Hilton

a.hilton@surrey.ac.uk

University of Surrey
Guildford, Surrey, GU2 7XH, UK

Abstract

In this paper we present a novel approach for temporal alignment of reconstructed mesh sequences with non-rigid surfaces to obtain a consistent representation. We propose a hierarchical scheme for non-sequential matching of frames across the sequence using shape similarity. This gives a tree structure which represents the optimal path for alignment of each frame in the sequence to minimize the change in shape. Non-rigid alignment is performed by recursively traversing the tree to align all frames. Non-sequential alignment reduces problems of drift or tracking failure which occur in previous sequential frame-to-frame techniques. Comparative evaluation on challenging 3D video sequences demonstrates that the proposed approach produces a temporally coherent representation with reduced error in shape and correspondence.

1. Introduction

Capture of 3D video sequences is becoming widely available using multiple camera reconstruction or real-time active sensors. Recent work has focused on the capture and reconstruction of people from multi-view camera setups with visual reconstruction achieving results comparable to laser scanning [12]. These approaches allow the capture of dynamic non-rigid free-form deformation including cloth and hair. Non-rigid shape capture [5, 7, 16, 20] results in an unstructured mesh sequence with changing connectivity and geometry at each frame. Efficient representation, rendering, editing and analysis of captured mesh sequences requires temporal alignment and consistent mesh structure.

Recent research has focused on the problem of reconstructing temporally consistent mesh sequences for people. A number of approaches have been proposed for sequential frame-to-frame non-rigid surface tracking based on correspondence of appearance and geometric features. Starck et al. [14] proposed a continuous surface tracking approach based on geometry images. Aguiar et al [5] use SIFT features and a patch based approach to constrain the deforma-

tion of a high-resolution mesh to match each frame in a sequence. A volumetric Laplacian scheme which minimizes the change in volume is used to regularize the deformation. The use of SIFT appearance features results in a sparse distribution across the surface with uniform regions such as arms and legs commonly having insufficient features to constrain the deformation. Ahmed et al [1] attempt to increase the number of SIFT correspondences using harmonic functions to infer matches from surrounding sparse features. In practice this may still result in insufficient frame-to-frame feature correspondences to constrain the deformation.

Zaharescu et al. [21] introduced MeshHOG which extended SIFT appearance features to the 3D domain, taking into account local geometric detail along with the photometric information. Results demonstrate that this gives an improved distribution of features across the surface providing sparse features matches in areas where there is insufficient variation in appearance for SIFT. Varanasi et al [19] detect the extremities of the limbs for 3D video sequences of people using the maxima of the geodesic integral. Geometric features commonly fail where there is a change in the reconstructed surface topology. This commonly occurs in 3D video of people where the arms are close to the body or self-occlusion from the camera views results in phantom volume protrusions of the reconstructed surface.

Cagniard et al [2, 3] introduce an approach based on iterative closest point ICP registration of rigid surface patches from frame-to-frame. A reconstructed frame from the sequence is used as a reference mesh avoiding the need for a prior high-resolution surface scan [5]. Results on sequences of complex free-form non-rigid surface motion demonstrate that this approach, using purely geometric information, is able to track surfaces undergoing relatively large non-rigid deformations. Cagniard et al. [4] present a probabilistic dense matching approach which introduces improved robustness to errors in the reconstructed surfaces. The use of purely geometric features may fail if deformations are too large for nearest point estimation and can result in drift across the surface for uniform geometric regions. Results of the non-sequential approach proposed in this paper are

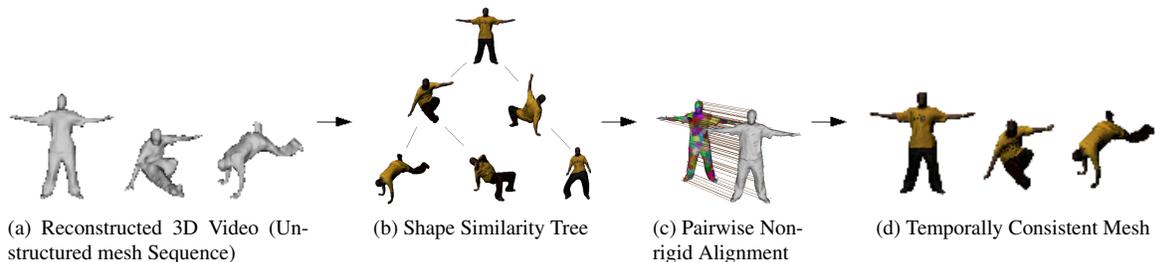


Figure 1: Overview of our approach to non-sequential 3D video alignment

compared to Cagniart et al. [4] and show an improved accuracy of surface fit.

The majority of previous approaches to obtaining consistent mesh structure for 3D video sequences focus on the problem of frame-to-frame non-rigid surface tracking. Under large non-rigid deformations due to fast movement, clothing or errors in surface reconstruction this approach may fail and result in part of the sequence not being tracked. The problem of non-rigid surface matching across large differences in shape has recently been addressed. Starck et al. [15] proposed an approach based on coarse-to-fine matching using both appearance and geometric features with correspondence optimisation performed using belief propagation. This achieves accurate matching between pairs of frames exhibiting large deformations but is relatively inefficient to solve. Tung et al. [18] perform surface matching based on sparse geometric features and a geodesic mapping which ensures a one-to-one matching across the surface.

In this paper we propose an alternative non-sequential approach to alignment of 3D video sequences into a consistent structure. Non-sequential alignment reduces the amount of drift by aligning meshes across the sequence based on their shape similarity rather than temporal ordering and overcomes the problem of sequential tracking failure by aligning frames from all parts of the sequences. The principal novelty of this approach is the extraction of a hierarchical tree representation of a 3D video sequence which represents the shortest path between frames in terms of shape similarity. A Laplacian deformation framework, similar to previous approaches [5, 2], is employed for non-rigid pairwise mesh alignment which incorporates both geometric and photometric feature matching to reduce drift. Other non-rigid alignment techniques could be used within the proposed non-sequential matching framework. An overview of the approach is presented in Figure 1. Results demonstrate that non-sequential matching gives improved accuracy of surface fit and reduces gross errors due to large deformation or poor feature matching.

2. Hierarchical Shape Matching

In this section we introduce a hierarchical tree representation which defines the shortest path between a reference mesh and all other meshes in the sequence. Shortest path is defined in terms of shape similarity between frames. The branches of the tree provide shorter non-rigid alignment paths in shape similarity space which reduces the accumulation of error and localizes gross errors to a single branch.

2.1. Shape Histograms

For the purpose of shape matching between frames of the sequence we use shape histogram comparison which has previously been shown to give good performance for 3D video sequences of people [8, 9]. The 3D space of the mesh is decomposed into a number of bins giving a spacial representation of the frame. The shape histogram is constructed using a spherical coordinate system to partition the space (r, ϕ, θ) around the center of mass. In our implementation we select the number of bins $N = N_r N_\phi N_\theta = 5 \times 10 \times 20 = 1000$ within a fixed size sphere of sufficient radius to encompass all frames of the sequence.

The histogram $H(F_s)$ counts the volumetric occupancy of the mesh F_s for the s^{th} frame in the sequence. To ensure that the similarity measure is invariant to rotation we test for maximum similarity over all rotations in ϕ . This can be achieved efficiently by shifting a fine histogram with 1° bins and re-binning to the required resolution.

For a sequence S of meshes (F_1, \dots, F_n) with corresponding shape histograms (H_1, \dots, H_n) similarity is computed with the L_2 distance. For source mesh F_s and target frame F_t the minimum L_2 distance corresponding to maximum similarity $D(F_s, F_t)$ for each 1° shift $d = (1, \dots, 360)$ is given by:

$$D(F_s, F_t) = \min_d \sum_{i=1}^N \|H_s(i) - H_t(i + d)\|^2 \quad (1)$$

2.2. Shape Similarity Tree

Errors in non-rigid mesh alignment increase as the difference in shape between meshes increases. We therefore propose a tree structure representing all frames in the sequence based on their relative shape similarity according to equation 1. This representation reduces the number and size of the steps in the pairwise non-rigid mesh alignment. Figure 2 illustrates a simple shape similarity tree for a sequence of meshes. First a single frame is selected as a reference mesh from the sequence F_{ref} based on quality of reconstruction. The shape similarity tree is then constructed to represent the shortest path to each frame from the reference mesh such that the pairwise similarity is less than a maximum distance t_{max} (minimum similarity). This constraint avoids large deformations in the pairwise non-rigid alignment. The maximum distance threshold could be computed automatically by evaluating the maximum of the minimum for each row in the shape similarity matrix $D(s, t)$. In practice throughout this paper we have taken a fixed threshold of $t_{max} = 0.001$.

Given a sequence S comprising n meshes (F_1, \dots, F_n) we build a fully connected graph $G(V, E)$ with vertices $V = (F_1, \dots, F_n)$ representing each mesh. Edges (i, j) of the graph are weighted according to the similarity measurement $D(F_i, F_j)$. The minimum spanning tree T for graph G with root F_{ref} subject to the constraint $D(F_i, F_j) < t_{max}$ then represents the minimum path length between the reference mesh and frames in the sequence. Isolated frames with large deformations to all other frames which do not satisfy the maximum distance constraint are added to the tree to give the minimum path length.

The resulting hierarchy defines the path $P_{min} = ((F_{ref}, F_i), \dots, (F_j, F_k))$ as a set of edges connecting the reference mesh F_{ref} and each mesh in the sequence F_k which minimizes the sum of differences in shape:

$$P_{min} = \arg \min_P \left(\sum_{(i,j) \in P} D(F_i, F_j) \right) \quad \text{subject to } D(F_i, F_j) < t_{max} \quad (2)$$

where $P \in \mathbb{P}$ the set of all possible paths. We construct this tree T with the following algorithm:

1. $F_s = F_{ref}$
2. $\forall F_t \notin T$ select children of F_s such that $D(F_s, F_t) < t_{max}$
3. For all children of F_s repeat from step 1
4. $\forall F_t \notin T$ where $D(F_s, F_t) > t_{max} \forall s$ select parent with $\min D(F_s, F_t)$

This tree defines the shortest paths of deformation required to align every frame of the sequence for a given starting frame F_{ref} and matching threshold t_{max} . The mesh at the root of the tree is deformed along all of its branches using a pairwise non-rigid alignment approach described in section 3.

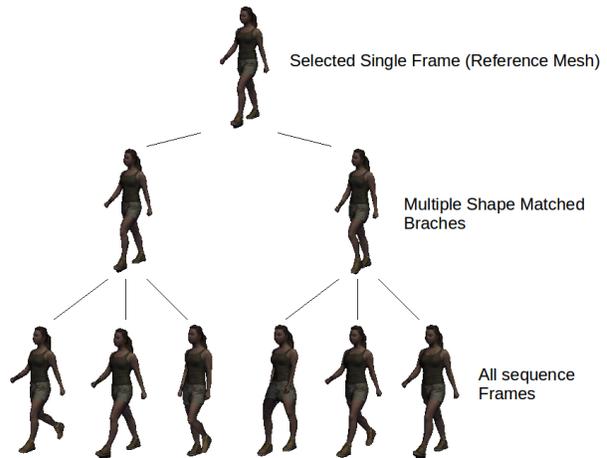


Figure 2: Tree structure of shape matches representing a reconstructed mesh sequence

3. Pairwise Non-rigid Alignment

In this section we introduce a non-rigid pairwise mesh alignment algorithm which uses a coarse-to-fine approach to combine both geometric and photometric matching in a Laplacian mesh deformation framework. This builds on previous work using Laplacian mesh deformation for surface matching [5, 3]. A volumetric Laplacian similar to [5] with soft matching constraints is initially used to allow large deformations with robustness to correspondence errors. A surface Laplacian is then used for refinement to allow a closer fit to the target mesh. Rigidity constraints are also included to give improved robustness to large deformations. In practice, the use of a combination of photometric and geometric features together with volumetric and surface Laplacians has been found to give improved reliability and accuracy of non-rigid alignment.

3.1. Stage 1 - Initial Deformation

An initial alignment between a source F_s and target F_t mesh is estimated using a volumetric deformation technique described in section 4.2. To constrain deformation correspondences are generated between frames using a combination of SIFT [11] and geometric matching [2].

3.1.1 Appearance Feature Matching

SIFT appearance features are computed per camera for each frame of the original video and subsequently projected into 3D. SIFT feature matching is selected over the mesh based variant MeshHOG [21] since it allows for matching from one camera to multiple cameras. This was found to produce a larger number of feature matches. SIFT matches are filtered using a spectral technique[10] based on their relative distance in the source and target frames.

3.1.2 Geometric Feature Matching

Appearance features alone can leave texture-less areas of the mesh or thin parts unconstrained. To combat this a sparse set of geometric features are created evenly distributed on the mesh surface. Geometric features are created using an adaption of Cagniard et al [2] based on ICP fitting of rigid surface patches. Their approach involves separating the source mesh into a number of surface patches and using a variation of the ICP algorithm to fit those patches to the target mesh. The centers of the surface patches are selected as correspondences. The key difference of this approach to traditional ICP is in the selection of target nearest points. The targets for each vertex are selected by first looking at the target mesh and assigning each target vertex a closest compatible point in the source mesh. The final target is then computed as a weighted combination according to the dot product of the motion vector and vertex normal.

3.1.3 Feature Selection

Feature selection is performed to obtain a sparse set of geometric and appearance features which are distributed evenly over the surface. The surface is subdivided into patches of equal size based on geodesic distance and a single feature correspondence is selected for each patch, Figure 3. Each SIFT or geometric feature S_i has associated vertex v_i from the source mesh and a target position t_i . For each patch we select the correspondence whose direction of motion $t_i - v_i$ is most co-linear with the normal of its associated vertex v_i [5]. Patches which have no associated correspondences are constrained to move rigidly based on the location of adjacent patches using the Laplacian deformation framework.

Combination of SIFT and geometric features together with rigidity constraints on unconstrained patches reduce the drift in pairwise matching over successive frames as illustrated in Figure 4. Appearance features allow areas of high texture with little shape information to be accurately aligned, indicated by the waist band which is distorted when only geometric features are used.

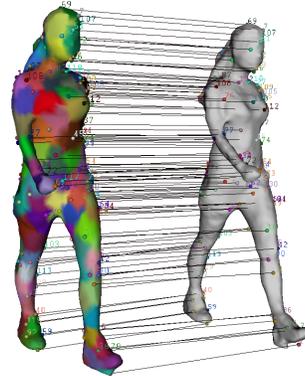


Figure 3: Patches generated on a tetrahedral mesh and selected frame to frame correspondences

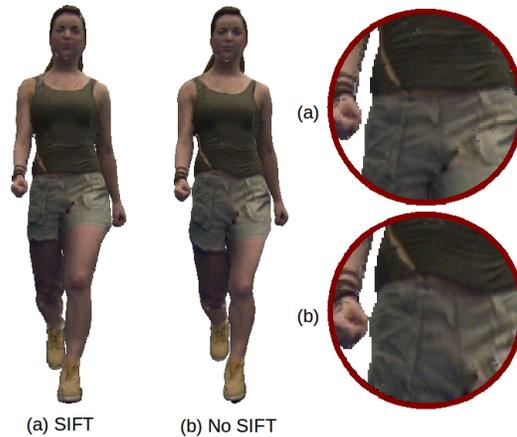


Figure 4: Comparison of SIFT and geometric information (a) with geometric information alone (b)

3.2. Stage 2 - Iterative Surface Fitting

In the second stage of our algorithm we refine the initial surface alignment from stage 1 to accurately match the target mesh. The following steps are iterated:

1. Double the number of geodesic patches
2. Perform the ICP variant on these patches
3. Deform the source mesh with surface Laplacian deformation

With each iteration the size of the patches is decreased allowing a coarse-to-fine matching with increasing density of correspondences. As the patches decrease in size each one contains significantly less shape information. It is therefore essential to get a close initial fit before increasing the number of patches otherwise drift across the surface will occur. Correspondences are selected from the set of geometric

and appearance features as in stage 1. In Stage 2 we use a surface based Laplacian deformation approach which preserves surface area rather than mesh volume. This allows the surface to expand and contract to fit more exactly the shape of the target data. The advantage of this is demonstrated when loose clothing creates a variation in volume between consecutive frames.

3.3. Stage 3 - Dense Fitting

Finally, to achieve the closest possible surface match we attempt to find a target for every vertex of the source frame. Closest point correspondences for every vertex are estimated using the target to source approach taken in the ICP variant. The weighted combination of possible targets provides the target location. This step allows accurate recovery of surface detail.

4. Laplacian Deformation

All the deformation techniques we present in this paper are variations of the Laplacian mesh editing paradigm [13]. Laplacian deformation involves fixing a number of vertex locations and solving for the others by fitting the Laplacian (differential coordinates) of the new geometry to the differential coordinates of the original mesh:

$$Lx = \delta \quad (3)$$

where L is the Laplacian operator matrix, x is a vector of the mesh's vertices stacked $(x_1, \dots, x_n, y_1, \dots, y_n, z_1, \dots, z_n)$ and δ is the differential coordinates of mesh.

4.1. Surface Laplacian

The Laplacian operator matrix is defined on the connectivity of the mesh and is given by:

$$L = G^T DG \quad (4)$$

where D is the diagonal degree matrix which weights the system according to the cotangent weights [6]. G is the gradient operator matrix which contains the gradients of the triangles shape functions ϕ using the faces normal n for orientation.

$$\begin{aligned} G_i &= (\nabla\phi_i, \nabla\phi_j, \nabla\phi_k) \\ &= \begin{pmatrix} (P_1 - P_3)^T \\ (P_2 - P_3)^T \\ n^T \end{pmatrix}^{-1} \begin{pmatrix} 1 & 0 & -1 \\ 0 & 1 & -1 \\ 0 & 0 & 0 \end{pmatrix} \quad (5) \end{aligned}$$

The differential coordinates of a manifold δ are given by multiplying the Laplacian operator L by the vector of the meshes vertices as in 3. Computing deformed vertex locations, x_u , involves solving equation 3 factored according to the constrained vertex locations x_k :

$$x_u = \arg \min_{x_u} \|Lx_u - (\delta + c)\| \quad (6)$$

where $c = Lx_k$ is the result of multiplying the Laplacian operator L by the vector of known constrained vertex locations x_k . This gives a system in which constrained vertices provide hard constraints to deformation. Since our constraints are likely to be subject to some error to maintain smooth deformation we use an energy based formulation to introduce soft constraints [2].

$$x_u = \arg \min_{x_u} \|Lx_u - \delta\|^2 + \|W_c(x_u - x_k)\|^2 \quad (7)$$

Solving this system directly leads to the well known problem with linear interpolation of large rotations. Instead we adopt an iterative approach. Equation 7 is solved repeatedly whilst updating the differential coordinates at each iteration. At each stage rotations are extracted for each facet of the mesh and applied to the corresponding facet in the source mesh. The differential coordinates are then updated accordingly[13].

4.2. Volumetric Laplacian

The Delaunay tetrahedralization of a triangular manifold gives a mesh with volumetric elements. Using this mesh it is possible to define a Laplacian weighted by the volumes of the tetrahedral elements [17]. Here the shape functions from which the gradient operator matrix G is defined are based on the tetrahedral elements of the mesh.

$$\begin{aligned} G_i &= (\nabla\phi_1, \nabla\phi_2, \nabla\phi_3, \nabla\phi_4) \\ &= \begin{pmatrix} (P_1 - P_4)^T \\ (P_2 - P_4)^T \\ (P_3 - P_4)^T \end{pmatrix}^{-1} \begin{pmatrix} 1 & 0 & 0 & -1 \\ 0 & 1 & 0 & -1 \\ 0 & 0 & 1 & -1 \end{pmatrix} \quad (8) \end{aligned}$$

The Laplacian operator is again defined as in 3. Here D is a diagonal weight matrix containing the volumes of the respective tetrahedral elements. The volumetric approach is particularly useful in maintaining the volume and shape of a manifold under large deformation and preventing collapse as used in the first stage of non-rigid alignment.

A section we wish to keep rigid under deformation is defined as a selection of neighboring triangles. On each iteration of the solver the differential coordinates of all members of a rigid section are updated according to the average rotation. In this way a section is unable to bend but can still orientate according to positional constraints.

5. Results and Evaluation

Hierarchical non-rigid matching is applied to a publicly available database of 3D video sequences [16] which includes fast challenging sequences of a brake dancer for

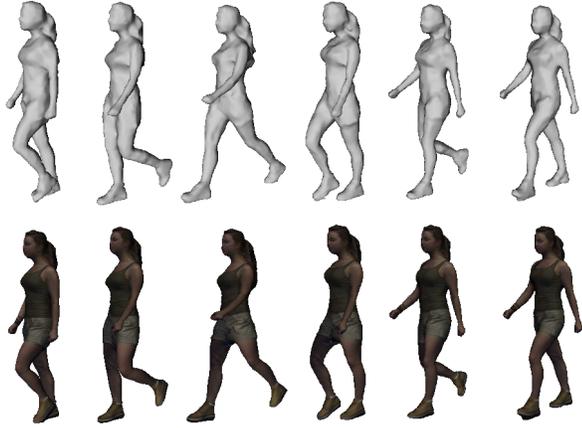


Figure 5: Top: Original reconstructed 3D video sequence, Bottom: Sequence with temporal correspondence. This technique works particularly well on cyclic motions

evaluation. We compare the results on our non-sequential alignment algorithms to results previously reported on this data using a sequential approach [4]. There are two key aspects to testing the quality of a temporally consistent surface alignment technique. The results must accurately represent the target shape and each vertex must maintain accurate correspondence on the surface. The quality of the fit to the target data can be demonstrated quantitatively with the Hausdorff and RMS error measures. The degree to which vertices maintain correspondence can be assessed qualitatively by texturing the first frame of the sequence and transferring the UV map to the remainder of the sequence by exploiting the temporally consistent representation.

5.1. Qualitative

Figure 6 compares temporally consistent data produced using our shape matching approach to the original reconstructed data. The shapes of the meshes produced are visually identical. By texturing the first frame on these sequences and projecting the UV map through the temporally consistent output of our algorithm we observe that the vertices maintain alignment indicated by texture remaining locked to the surface. This is particularly noticeable in cyclic motion such as the walk sequence, Figure 5. In such sequences shapes are commonly repeated at several frames, non-sequential matching is performed across relatively short sequences of frames reducing drift compared to a sequential approach.

5.2. Quantitative

The Hausdorff distance evaluates the maximum deviation of a vertex on the source mesh from the surface of the target data to identify gross errors. Root mean square error

gives an average distance of the vertices to the surface of the target data indicating the accuracy of shape representation. It is important to consider these metrics together with visual information as in cases where reconstruction errors are present in the frame and are correctly avoided by the algorithm large distances can occur.

Figure 7 shows plots of RMS error for multiple 3D video sequences [16]. In all cases we present a very close fit to the target data with the RMS error averaging $\sim 10\text{mm}$ across all sequences. The maximum Hausdorff distance is $\sim 120\text{mm}$ and only occurs in a few frames. Reconstruction error is of the order of $\sim 10\text{mm}$ so this level of rms error indicates a good fit. Frames often contain large errors in the reconstruction itself leading to an exaggerated Hausdorff distance and bias in RMS Error. Figure 7 also compares our results to a state-of-the-art sequential alignment approach [4]. Results indicate comparable or improved fitting accuracy for the proposed non-sequential alignment approach.

5.3. Recovery From Error

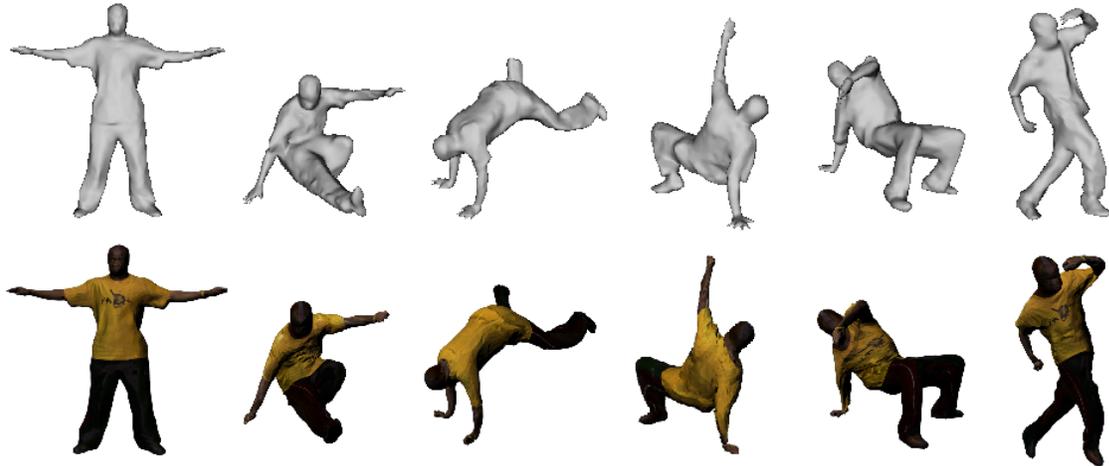
Another important consideration for any sequential tracking technique is recovery from errors. An advantage of the proposed non-sequential tree based approach to non-rigid surface alignment is that errors are only propagated locally along the branches. Misalignment at an intermediate frames will not cause the remainder of the sequence to fail. This is illustrated in Figure 8 where the left hand and collar of the shirt have incorrect alignment in one frame but are correctly aligned in subsequent frames. In practice, alignment failure due to erroneous features or reconstruction errors could be corrected by removing the misaligned frame and interpolating. This correction is not performed in any of the results presented.



Figure 8: Recovery from error. The left hands geometry has degraded and the texture around the collar has started to drift (left) however in the next frame it has recovered (right)

6. Conclusions

A non-sequential approach to non-rigid alignment of mesh sequences has been presented. This approach uses a shape similarity tree to hierarchically match frames across the sequence minimizing the change in shape between



(a) Free Sequence



(b) Head Sequence

Figure 6: Top: Original reconstructed 3D video sequence, Bottom: Sequence with temporal correspondence. UV texture maps are created for the first frame and transferred to subsequent frames exploiting temporal consistency

meshes for pairwise non-rigid alignment. Pairwise alignment is performed using a coarse-to-fine Laplacian deformation approach which combines photometric and geometric features to improve reliability. The non-sequential alignment approach could also be used with other pairwise non-rigid alignment techniques. The principal advantages of the proposed non-sequential approach are to reduce drift inherent in sequential frame-to-frame alignment due to accumulation of errors and to allow alignment across complete sequences even if an intermediate frame is misaligned. There is no visible jitter between non-sequentially reconstructed frames where branches of the tree come back together. Evaluation on a public database of 3D video sequences containing rapid non-rigid movement demonstrates that the non-sequential approach improves the accuracy of

non-rigid alignment and shape representation compared to a state-of-the-art sequential approach. The alignment is sufficiently accurate to transfer a texture map across a sequence maintaining a stable alignment on the surface. Additionally shape matching allows reconstruction across multiple sequences and thus temporal alignment of entire databases.

References

- [1] N. Ahmed, C. Theobalt, C. Rössl, S. Thrun, and H.-P. Seidel. Dense correspondence finding for parametrization-free animation reconstruction from video. In *CVPR*, 2008. 1
- [2] C. Cagniard, E. Boyer, and S. Ilic. Iterative mesh deformation for dense surface tracking. In *ICCV Workshop on 3-D Digital Imaging and Modeling*, 2009. 1, 2, 3, 4, 5

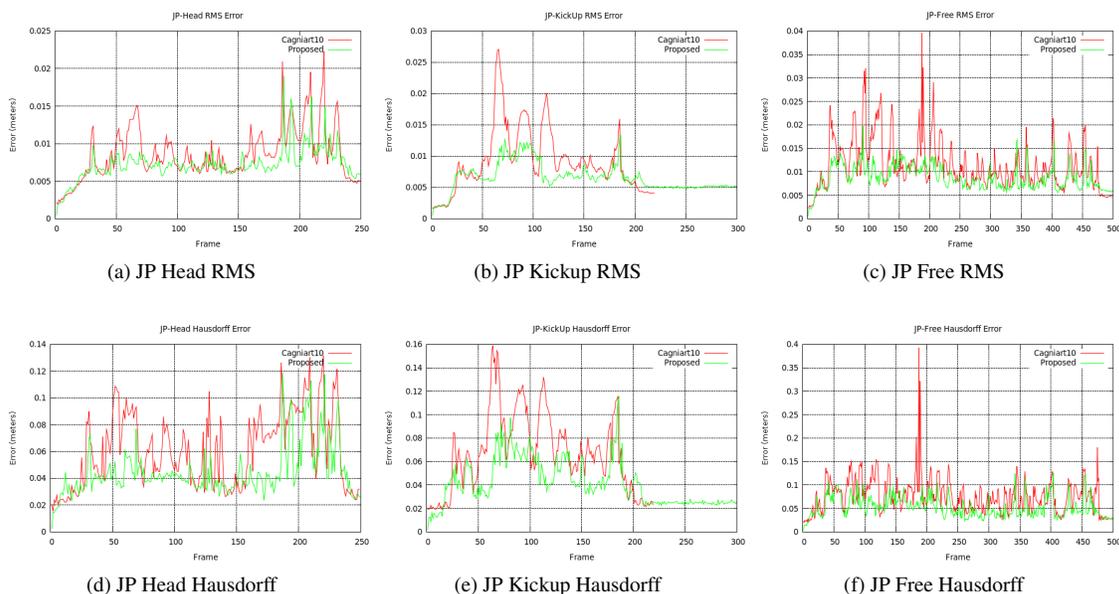


Figure 7: Comparison of RMS and Hausdorff errors for the SurfCap project sequences using our approach and the current state of the art from Cagniat et al [4]

- [3] C. Cagniat, E. Boyer, and S. Ilic. Free-from mesh tracking: a patch-based approach. In *CVPR*, 2010. 1, 3
- [4] C. Cagniat, E. Boyer, and S. Ilic. Probabilistic deformable surface tracking from multiple videos. In *ECCV*, 2010. 1, 2, 6, 8
- [5] E. de Aguiar, C. Stoll, C. Theobalt, N. Ahmed, H.-P. Seidel, and S. Thrun. Performance capture from sparse multi-view video. In *SIGGRAPH '08: ACM SIGGRAPH 2008 papers*, pages 1–10. ACM, 2008. 1, 2, 3, 4
- [6] M. Desbrun, M. Meyer, P. Schröder, and A. H. Barr. Discrete differential-geometry operators for triangulated 2-manifolds. pages 35–57. Springer-Verlag, 2002. 5
- [7] J.-S. Franco, M. Lapierre, and E. Boyer. Visual shapes of silhouette sets. *3D Data Processing Visualization and Transmission, International Symposium on*, 0:397–404, 2006. 1
- [8] P. Huang, A. Hilton, and J. Starck. Shape similarity for 3d video sequences of people. *Int. J. Comput. Vision*, 89:362–381, September 2010. 2
- [9] P. Huang, T. Tung, S. Nobuhara, A. Hilton, and T. Matsuyama. Comparison of skeleton and non-skeleton shape descriptors for 3d video. In *3DPVT*, 2010. 2
- [10] M. Leordeanu and M. Hebert. A spectral technique for correspondence problems using pairwise constraints. In *International Conference of Computer Vision (ICCV)*, volume 2, pages 1482 – 1489, October 2005. 4
- [11] D. G. Lowe. Object recognition from local scale-invariant features. In *ICCV '99*, page 1150, Washington, DC, USA, 1999. IEEE Computer Society. 3
- [12] S. M. Seitz, B. Curless, J. Diebel, D. Scharstein, and R. Szeliski. A comparison and evaluation of multi-view stereo reconstruction algorithms. In *CVPR '06*, pages 519–528, 2006. 1
- [13] O. Sorkine, D. Cohen-Or, Y. Lipman, M. Alexa, C. Rössl, and H.-P. Seidel. Laplacian surface editing. In *Eurographics Symposium on Geometry Processing*, pages 179–188. Eurographics Association, 2004. 5
- [14] J. Starck and A. Hilton. Spherical matching for temporal correspondence of non-rigid surfaces. *ICCV*, pages 1387–1394, 2005. 1
- [15] J. Starck and A. Hilton. Correspondence labelling for wide-timeframe free-form surface matching. *ICCV*, 0:1–8, 2007. 2
- [16] J. Starck and A. Hilton. Surface capture for performance based animation. *IEEE Computer Graphics and Applications*, 27(3):21–31, 2007. 1, 5, 6
- [17] C. Stoll. A volumetric approach to interactive shape editing. Research Report MPI-I-2007-4-004, Max-Planck-Institut für Informatik, Stuhlsatzenhausweg 85, 66123 Saarbrücken, Germany, June 2007. 5
- [18] T. Tung and T. Matsuyama. Dynamic surface matching by geodesic mapping for 3d animation transfer. pages 1402–1409, 2010. 2
- [19] K. Varanasi, A. Zaharescu, E. Boyer, and R. Horaud. Temporal surface tracking using mesh evolution. In *ECCV*, pages 30–43, 2008. 1
- [20] D. Vlastic, I. Baran, W. Matusik, and J. Popović. Articulated mesh animation from multi-view silhouettes. *ACM Transactions on Graphics*, 27(3):97, 2008. 1
- [21] A. Zaharescu, E. Boyer, K. Varanasi, and R. P. Horaud. Surface feature detection and description with applications to mesh matching. In *CVPR*, 2009. 1, 4