

Templateless Quasi-Rigid Shape Modeling with Implicit Loop-Closure

Ming Zeng* Jiaxiang Zheng Xuan Cheng Xinguo Liu†
State Key Lab of CAD&CG, Zhejiang University

Abstract

This paper presents a method for quasi-rigid objects modeling from a sequence of depth scans captured at different time instances. As quasi-rigid objects, such as human bodies, usually have shape motions during the capture procedure, it is difficult to reconstruct their geometries. We represent the shape motion by a deformation graph, and propose a model-to-part method to gradually integrate sampled points of depth scans into the deformation graph. Under an as-rigid-as-possible assumption, the model-to-part method can adjust the deformation graph non-rigidly, so as to avoid error accumulation in alignment, which also implicitly achieves loop-closure. To handle the drift and topological error for the deformation graph, two algorithms are introduced. First, we use a two-stage registration to largely keep the rigid motion part. Second, in the step of graph integration, we topology-adaptively integrate new parts and dynamically control the regularization effect of the deformation graph.

We demonstrate the effectiveness and robustness of our method by several depth sequences of quasi-rigid objects, and an application in human shape modeling.

1. Introduction

3D reconstruction is an important topic in computer vision. With the recent developments of stereoscopic vision and depth acquisition devices, it is becoming easy to obtain the 3D models using images or depth scans captured simultaneously from different views (e.g. [7, 25]). This kind of methods require sufficient amount of spatially arranged cameras. An alternative way is to move a sensor around to capture the object from multiple directions (e.g. [9, 16]). This method requires the captured objects keep static during capture. However, a large number of objects cannot keep static for a reasonable long period. A typical example is the human body modeling, as illustrated in Fig. 1.

Modeling the non-rigid objects with a single camera can be cast as a deformable shape completion problem. De-

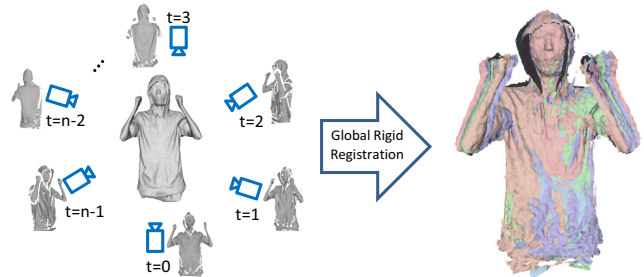


Figure 1. The quasi-rigid object exhibits slight deformation, which can not be reconstructed in a rigid registration fashion.

formable shape completion problem is generally ill-posed, because some shape information is not available due to occlusion. To address the missing information, people take advantage of various priors of the captured objects. The template-based methods (e.g. [10, 24]) confine the shape on un-seen views in the template space, which sacrifices some generality. The images-based methods (e.g. [13]) use the salient feature points in the images to track the correspondences, but they cannot handle textureless parts. The space-time methods (e.g. [15]) assume the locally continuous movements by requiring a sufficiently dense sampling temporally. The methods based on articulated model (e.g. [5]) reduce the degrees of freedom of the deformable objects. There are more general methods (e.g. [26, 23]) which leverage few assumptions on particular objects, but these methods take little care about loop-closure, a crucial issue on object reconstruction.

In this paper, we address the deformable shape completion problem of nearly rigid objects, called quasi-rigid objects. Taking advantage of the as-rigid-as-possible property of the quasi-rigid objects, we can complete the model visual-plausibly without using any template. Furthermore, thanks to reduced degree of the quasi-rigid assumption, our pipeline do not require an explicit global registration. Specifically, we maintain a deformation graph as the representation of the object and develop an incremental integration method to update the graph. For each depth scan, the graph is registered to it using a model-to-part method. The model-to-part method with the non-rigid deformation

*mingzeng85@gmail.com

† Corresponding author. xgliu@cad.zju.edu.cn

model is flexible enough to gradually adjust the graph to fit the new depth scans, and is able to achieve results with loop-closure property. After registration, the depth scans are re-sampled and integrated into the deformation graph. The integration takes care of the topology, and provides topology-adaptive information for regularization on the registration step of succeeding scans. After integrating all depth scans, a global nonrigid warping is adopted to warp the points to their destination positions in the last frame. And finally, we use the Poisson reconstruction method [14] on these warped points to obtain the object’s surface.

The main contribution of this paper is a method to reconstruct the 3D models of quasi-rigid objects, which is templateless and can achieve the loop-closure results. There are several technical contributions:

- First, we propose a model-to-part registration scheme for non-rigid deformation, which distributes the accumulated error in each registration step, and avoids explicit steps for error distribution.
- Second, we introduce a two-stage registration method to make the registration step robust to geometry tracking.
- Third, we propose a topology-adaptive integration and a relaxed regularization for graph update, which improve the robustness of our method.

2. Related Work

Rigid Registration. To obtain the whole 3D model of a static object, scans of different views should be aligned together in a common coordinate by rigid transformations. Iterative closest point (ICP) [2, 6, 18] first sets up some point correspondences by nearest neighbor criterion, and then get a rigid transformation by minimizing the distances of the corresponding point pairs. This procedure is repeated until convergence is reached. The rigid registration is unable to deal with non-rigid shape.

Non-Rigid Registration. For deformable shapes, more degree of freedom is required to represent the deformation of shapes. Based on the embedded deformation introduced by Sumner et al. [20], Li et al. [12] combined the correspondence optimization with deformation optimization together, and improved the robustness of registration. Based on the same deformation model, Huang et al. [8] studied the problem under isometric deformation, and argued that keeping low but necessary deformation freedom will improve the registration reliability. With the similar principle to reduce deformation freedom, Chang and his colleagues introduced an articulated model [3] and a linear blend skinning model [4] for the nonrigid registration. These pairwise registration methods are necessary building blocks for full shape reconstruction.

Shape Completion. The shape completion problem requires merging scans from different views to reconstruct the whole shape. To cope with the under-constrained problem, priors are incorporated to confine the shape in a subspace. Li et al. [10] leverage a coarse template to track the movement of the depth scan. Weiss et al. [27] reconstruct human bodies with a parameterized human model SCAPE[1]. Using this SCAPE model, Tong et al. [24] track the scans of slightly moving human, and impose an explicit global alignment to distribute the accumulated alignment errors. These methods require a template model beforehand. Liao et al. [13] proposed a template-free method, which assumes a locally continuous movement and tracks the correspondences using associated RGB images. The method will fail when the RGB images are noisy or absent. The Space-time methods (e.g. [15, 21]) take the same locally continuous condition by requiring a sufficiently dense sampling on the time axis. These methods optimize the registration problem in a small window of multiple frames to avoid an additional error distribution step. Chang and Zwicker [5] restrict the moving shape to be articulated, and solve a joint labeling/deformation problem with the reduced deformation model. But the method cannot handle details for some deformable objects, e.g. the folds of cloth. Wand et al. [26] and Tevs et al. [23] proposed more general frameworks which impose few assumption on particular objects. These methods are mainly designed for capturing a sequence of moving objects. They take little care about how to match the first and the last scans together, i.e. the loop-closure property.

Our model-to-part method is similar to temporally coherent completion [11] proposed by Li et al. The incremental integration method may incur error accumulation in the rigid case, but the scheme works well in the nonrigid case thanks to the adjustability of the pre-integrated models. Unlike the multi-view scenario in [11], we deal with single-view completion problem, which is prone to fail due to topological issues. In our paper, we study the model-to-part method further than [11] by incorporating the topology information obtained from the integration step into succeeding registration step, which adaptively control the flexibility of the integrated model.

3. Overview

Fig. 2 gives an overview of our method. The input is a sequence of depth scans of a quasi-rigid object captured from different views at different time instances. To represent the motion of the scans, we follow Sumner et al. [20] to represent space deformation by a deformation graph, where each node is a sample point of the object. When a scan comes, the deformation graph is non-rigidly registered to it, and then the new part in this new scan is re-sampled and integrated into the deformation graph. Then the topology of this




Figure 2. Overview. A deformation graph is non-rigidly registered to each new depth scan, which is followed by the integration and update of the deformation graph. After integration of these depth scans, a global non-rigid warping is conducted on each scan according to the deformation records of this graph. Finally, the resulted mesh model is obtained by Poisson reconstruction.

deformation graph is updated accordingly. The registration and the integration/update procedures repeat until the end of the sequence. After all depth scans are processed, the deformation graph records the whole deforming/dynamic behaviors of the scanned object. Using this dynamic information, all depth scans are aligned into one global coordinate, and are registered together compensating slight deformation. Finally, Poisson reconstruction is used to build the mesh model.

In the rest of this paper, we will first introduce the model-to-part registration in Section 4, deformation graph integration and update in Section 5, then we present the implementation details and the experimental results in Section 6, and conclude this paper in Section 7.

4. Model-to-Part Registration

In the rigid alignment scenario, there is inevitable accumulation error without a global registration [17, 19]. A model-to-part scheme [22] incrementally aligns a new scan into a integration model, which improves the alignment of the new scans, but it is unable to adjust the integrated model. Recently, Newcombe et al. proposed KinectFusion [16], which adopts a model-to-part method to incrementally integrate new scans. They argued that the alignment error will be reduced by constant update of the integrated model. Inspired by KinectFusion, we take a similar model-to-part way in our quasi-rigid registration case. In our method, the integrated model is able to be “moved” and “dragged” to match the new scan, and updates its shape iteratively. Furthermore, as is illustrated in Fig. 3, the relatively flexible way inherently keeps loop-closure by adjusting the “head” and the “tail” of integration model to match the last scan together.

For the moving object, the motion can be divided into two parts[12]: global rigid transformation Φ_{global} and local non-rigid deformation Φ_{local} . The global rigid transformation is caused by camera movement or the object’s




Figure 3. Illustration of implicit loop-closure of model-to-part registration. The red is ground truth, the green is deformation graph, and the dark blue is a depth scan. left: the registration suffers from slight drift. middle: the new scan overlaps with the first scan. right: model-to-part registration adjusts the deformation graph as-rigid-as-possible, which exhibits loop-closure implicitly.

overall movement. The local non-rigid deformation comes from slight articulated motion or the small-scale deformation therein (such as folds in the cloth). Therefore, for a point at position p , the transformed position can be computed by

$$\tilde{p} = \Phi_{local} \circ \Phi_{global}(p). \quad (1)$$

The global transformation can be formulated a 3×3 rotation matrix R and a 3×1 translation vector T :

$$\Phi_{global}(p) = Rp + T, \quad (2)$$

To address the local non-rigid deformation, we use the embedded deformation model [20], where a node s^j in the deformation graph determines a local warping field to control the transformation of points near s^j , which can be represented using a 3×3 matrix H^j and a 3×1 translation vector l^j as follows:

$$\Phi_{local}(p) = \sum_{j=1}^k w^j [H^j(p - s^j) + s^j + l^j], \quad (3)$$

where w^j is the normalized weights for p ’s j th-nearest

nodes $s^j, j = 1, 2, \dots, k$. We define w^j by the distance between p and the nodes as follows:

$$w^j = \frac{1 - \|p - s^j\|/d_{max}}{\sum_{m=1}^k 1 - \|p - s^m\|/d_{max}}. \quad (4)$$

with d_{max} is distance between p and its $k + 1$ th nearest node.

4.1. Two-Stage Registration

Let G_t denotes the deformation graph at frame t , and $|G_t|$ denotes the number of the nodes in G_t . To register G_t with a new scan D_{t+1} , we find the best transformation $\Phi = \Phi_{local} \circ \Phi_{global}$ that transforms G_t to fit D_{t+1} .

Following the method proposed by Sumner et al.[20], a fitting term E_{fit} , a rigidity terms E_{rigid} , and a regularization term E_{reg} are defined to form the minimization problem:

$$\min_{R_t, T_t, H_t^i, l_t^i, i=1, 2, \dots, |G_t|} w_{fit} E_{fit} + w_{rigid} E_{rigid} + w_{reg} E_{reg} \quad (5)$$

The fitting term provides a desired position for each graph node. It sums up distances of all m pairs of correspondences:

$$E_{fit} = \sum_{i=1}^m M(\Phi(s_t^i), d_{t+1}^{i*}) \quad (6)$$

where, d_{t+1}^{i*} is node s_t^i 's correspondence in D_{t+1} , and $M(x, y)$ is a metric combining the point-to-point and the point-to-plane distance.

The second rigid term constrains the transformation matrix H_i to be rotational:

$$E_{rigid} = \sum_{i=1}^{|G_t|} Rot(H_t^i) \quad (7)$$

where $Rot(H) = \|H'H - I\|_F^2$.

The regularization term considers the smoothness of the neighboring deformation, which measures the difference of the nearby nodes' transformations:

$$E_{reg} = \sum_i \sum_{j \in N(i)} \|H_t^i(s_t^j - s_t^i) + s_t^i + l_t^i - (s_t^j + l_t^j)\|_2^2. \quad (8)$$

Since Φ_{global} accounts for most of the motion for the quasi-rigid object, we treat the rigid part and the non-rigid part separately, yielding a two-stage registration approach. We first view the graph G_t as rigid, and solve for the global transformation R_t and T_t . Then, we solve the optimization problem in Eq.5 with R_t and T_t fixed.

Rigid Registration Stage: We use the ICP method to estimate rigid transformation between G_t and D_{t+1} . And




Figure 4. Illustration of subset registration. top left: the full-set correspondences introduce wrong correspondences, leading to bad alignment (top right). bottom left: the subset correspondences (green) masks the outdated nodes (grey), leading to expected alignment (bottom right).

we use the fitting energy E_{fit} combining the point-to-point and point-to-plane metrics as follows:

$$E_{fit} = \sum_i \|R_t s_t^i + T_t - d_{t+1}^{i*}\|_2^2 + \rho [n_{t+1}^{i*} \cdot (R_t s_t^i + T_t - d_{t+1}^{i*})]^2 \quad (9)$$

Where, d_{t+1}^{i*} is the position of the correspondence of s_t^i , and n_{t+1}^{i*} is the normal with d_{t+1}^{i*} . The correspondences which are far apart or normal incompatible are pruned [18]. In our experiments, we use the weight $\rho = 0.1$ for point-to-plane distance.

As the scans are integrated into the deformation graph, nodes added in the very beginning of the integration procedure do not overlap with the new scan. These outdated nodes will disturb the registration due to introducing wrong correspondences. Therefore, we adopt a subset registration scheme. After each registration, we record the nodes which are covered by the previous scan, and mark them as active nodes. At the start of the registration of a new scan, we only use the active nodes while mask the outdated nodes in the rigid registration. As is illustrated in Fig. 4, the scheme improves the robustness of the rigid registration, especially for the graph having integrated lots of scans.

Non-Rigid Registration Stage: Given the rigid transformation R_t and T_t from rigid registration stage, we optimize the Eq.5 to obtain the local transformations for each node in the deformation graph. Like rigid registration, we adopt point-to-point and point-to-plane metrics together in E_{fit} , which ensures the non-rigid registration flexible enough to globally “slide” on the target scan. Similar to [20], the optimization Eq.5 can be minimized using Gauss-Newton algorithm.

In the quasi-rigid case, the object’s non-rigid deformation is small, while the deformation graph contains too many nodes to represent the small deformation. We deal with the over-flexibility problem by a simple method which is composed of two points. First, we set a large weight w_{reg} to keep the deformation as-rigid-as-possible. In our




Figure 5. Bijective closest correspondence. Top left: (s^2, p^1) and (s^3, p^1) are wrong, because $Near(p^1)$ is neither s^2 nor s^3 . The wrong correspondences will lead s^2 and s^3 to gather to p^1 . Top right: The wrong “edge point” correspondences are pruned. Bottom: similar to top. But here is “tip point” case. left: failing to prune wrong correspondences leads s^1 , s^2 , and s^3 to gather together. right: the wrong “tip point” correspondences are pruned.

following experiments, we use the weights $w_{reg} = 10^4$, $w_{fit} = 10^2$, and $w_{rigid} = 1$. Second, to prevent the deformation graph from collapsing together, we take a bijective closest correspondence instead of source-to-target closest correspondence. Specifically, for a node s^i with normal n^i belonging to the set of graph nodes S , suppose its nearest point on a depth scan is p^{i*} with normal n^{i*} , we find the id of p^{i*} 's nearest normal-compatible point $Near(p^{i*})$ from S :

$$\begin{aligned}
 Near(p^{i*}) &= \arg_h \min \|p^{i*} - s^h\|_2^2, \\
 s.t. \quad & \arccos(n^h \cdot n^{i*}) < \pi/3 \\
 & \|p^{i*} - s^h\|_2^2 < threshold \\
 & s^h \in S
 \end{aligned} \tag{10}$$

When $Near(p^{i*})$ equals to i , the correspondence (s^i, p^{i*}) is valid, otherwise it is pruned. As shown in Fig. 5, the bijective closest correspondence can prune wrong “edge point” and “tip point” correspondences.

5. Deformation Graph Integration and Update

The deformation graph grows as more depth scans are processed. At the first scan, we directly use down-sampled points from the depth scan to create the nodes of the deformation graph. In our experiments, we set the edge count $k = 6$ for each node. When a new scan is inserted, we register the deformation graph to the scan, and then add sample nodes into graph and update it.

5.1. Feature Preserving Sampling

Our method samples the nodes from the scans by a combination of uniform sampling in position and uniform sampling in normal space [8]. The uniform sampling ensures

the nodes cover all region of the scan, while the normal space sampling makes sure features on the scan are equally sampled.

5.2. Topology-adaptive Integration

For each sample point s_{t+1}^i from depth scan D_{t+1} , if its distance to the nearest point in the deformed graph \tilde{G}_{t+1} exceeds a threshold (we use 2% of the object’s diagonal length), it is added into the graph. The simple method works well on most cases, but fails on some scenarios which need special attentions on topological issues. One example is two sides of a thin cloth. Another case is a re-appearing part which re-appears in a position very close to its other side. We treat the topological issues by considering normal directions. For a sample point which is near the graph nodes, but has a opposite normal (angle > 150), it should be inserted into the graph as well.

5.3. Relaxed Regularization for Re-appearing Parts

After new sample points are added into the deformation graph, the nodes’ nearest neighbors are recomputed and their edges are rebuilt. According to the topology-adaptive integration, two neighboring nodes in the graph are likely to have very different normals. It may result from the two opposite sides of a thin part as mentioned in the previous subsection, or arises when a disappeared part suddenly re-appears in the camera view.




Figure 6. Illustration of relaxed regularization. Left: ground truth of a shape. Middle: the shape suddenly re-appears, and there are a large gap between the current scan (dark blue) and the previous scan (red), the relaxed regularization sets weak influences between them (green dash line). Right: the relaxed regularization allows these two part move apart when the graph are registered to new scans which have large overlapping regions between both red and dark blue parts.

In the latter case, the re-appearing parts are prone to mismatch the previous scans due to little overlapping area between them. Since these two parts belong to the same articulated component and should move together, it is reasonable to link the nodes from these two parts with edges for the registration regularization E_{reg} . However, the influence

between nodes with incompatible normals should be weaker than that of normal-compatible nodes, so as to allow the bad alignment to be adjusted to correct positions. We illustrate this case in Fig. 6. Based on the consideration, we relax the regularization between position-near but normal-incompatible nodes by introducing a relaxed regularization controller. Mathematically, for two sample nodes s_t^i and s_t^j , the relaxed regularization controller $c_{reg}^{i,j}$ between them is defined:

$$c_{reg}^{i,j} = \begin{cases} c_{high}, & n^i \cdot n^j > 0 \\ c_{low}, & n^i \cdot n^j \leq 0 \end{cases} \quad (11)$$

In our experiments, we set $c_{high} = 1$ and $c_{low} = 0.1$.

With the regularization controller, the regularization term E_{reg} is rewritten as:

$$E_{reg} = \sum_i \sum_{j \in N(i)} c_{reg}^{i,j} \cdot \|H_t^i(s_t^j - s_t^i) + s_t^i + l_t^i - (s_t^j + l_t^j)\|_2^2. \quad (12)$$

6. Experiments

We have tested our method on a serial of real or synthetic depth data. For real data, we captured 640×480 depth data by rotating objects in front of a fixed Kinect. The frame rate of the capture is 30 fps. However, since our method does not require dense sampling at the time axis, we uniformly sampled the data and only kept 10 to 15 frames (each frame covers about 24 to 36 degrees) for reconstruction. For synthetic data, we generate depth scans by a virtual camera.

The entire pipeline is implemented with C++, and tested on an desktop computer with Intel Core2 Duo E7400 2.80GH CPU (one core used) and 2GB RAM. In all the experiments, the rigid registration converges within 50 iterations and the non-rigid registration converges within 10 iterations, which costs about 40~60 seconds for a frame. The graph integration and update are very fast, which cost about 1~2 seconds for a frame. The whole pipeline takes 10~15 minutes to finish a model, depending on the number of scans.

6.1. Evaluation of Proposed Techniques

In this subsection we show the effectiveness of our bijective correspondence and relaxed regularization by without/with comparison, respectively.

Bijective Correspondence. Fig. 7 demonstrate the effectiveness of the bijective correspondence. Without the bijective consistence, the graph nodes on the foot shrink to the joints (Fig. 7(a)). While the bijective consistence prevents the gathering phenomena (Fig. 7(b)).

Relaxed Regularization. Fig. 8 shows that the non-relaxed regularization is unable to adjust the poorly integrated graph to correct position (Fig. 8(a)). The nodes on opposite side of the new scan are also “dragged” severely




Figure 7. Comparison between correspondences without/with bijective consistence: (a) vs (b).




Figure 8. Comparison between non-relaxed regularization (a) and relaxed regularization (b).

to the target scan by their neighboring nodes. While our relaxed regularization is more topology-adaptive to deform the graph (Fig. 8(b)), where the opposite nodes are less influenced due to lower regularization controller.

6.2. Real Data

We manually rotated a puppet and a pillow in front of a Kinect to capture their depth data. The hand-held rotation inevitably incurs deformation on the objects. Fig. 9 exhibits effectiveness of our method. In this figure, the first rows are our results, and the second rows are results with rigid global registration[17]. We show and highlight obvious artifacts with red rectangles.




Figure 9. Comparison between our method and global rigid registration on the (a) puppet and (b) pillow.

6.3. Synthetic Data

We use a virtual camera to capture a self-turned human model (1.75m height) with slight but substantial deformation (generated using MAYA). The captured single-view sequence is used as input (Fig. 10(a)), and the model at the last frame is used as the ground truth (Fig. 10(b)). We compute per-vertex nearest distance FROM the ground truth to results of our method (Fig. 10(c)) and global rigid registration[17] (Fig. 10(d)), respectively. The average/max distance of ours is 0.003/0.017, while that of global rigid registration is 0.004/0.037. We also tested KinectFusion on this data, but KinectFusion (Fig. 10(e)) suffers severe misalignment, failing to produce an appropriate result. From this experiment, we show (1) Our result approximates the ground truth globally; (2) Both global rigid registration and KinectFusion are unable to compensate the obvious non-rigid deformation.




Figure 10. (a) Snapshots of partial scans, (b) ground truth, (c) our result, (d) global rigid registration, (e) KinectFusion. Red rectangles highlight regions which deviate the ground truth.

6.4. Application: Human Body Modeling

Our method is suitable for building a human body modeling system. We fix two Kinects on a stand, one up and one down to capture upper and lower parts of a human body, respectively. Views of both Kinects overlap in a small region. To scan a human body, the human is asked to rotate by himself in front of the system, and try to keep his pose. During the process, the system captures depths of the movement. After capture, we merge the upper and lower parts of each frame, then process the merged data using our algorithm, and finally use Poisson reconstruction to get the model. It usually takes about 30~60 seconds to capture the depths, and about 15 minutes to generate the model. Fig. 11 shows some results, where our method preserves necessary details, such as collars, small folds, face, and even the hood with complex topology. These results are pleasing regarding to the low price setup. For more human modeling results, please refer to the supplementary material.

7. Conclusion

We have presented a general method for quasi-rigid shape modeling using depth scans captured at different time instances. In our method, to keep generality we do not require a template to assist the geometry tracking. Without the template, we gradually integrate depth scans into our pipeline and finally obtain a deformation graph representing the whole shape. The method is robust thanks to two key technique contributions. First is our model-to-part scheme to register the deformation graph to match the new scans. To keep the registration robust, we adopt a two-stage registration under the assumption of the as-rigid-as-possible. Second, we handle several topology issues raised with the integration of depth scans. The special cares on topology proposed in this paper in one hand topology-adaptively integrate new parts, and in other hand relax regularization for poorly aligned regions. We demonstrated the proposed method by reconstructing several quasi-rigid objects. As an application, we showed that our method can be used to build a human body modeling system.

Acknowledgements We would like to thank the anonymous reviewers for their valuable comments. This work was partially supported by NSFC (No. 60970074), China 973 Program (No. 2009CB320801), Fok Ying Tung Education Foundation and the Fundamental Research Funds for the Central Universities.

References

- [1] D. Anguelov, P. Srinivasan, D. Koller, S. Thrun, J. Rodgers, and J. Davis. SCAPE: shape completion and animation of people. *ACM Trans. Graph.*, 24(3):408–416, 2005. 2
- [2] P. J. Besl and N. D. McKay. A method for registration of 3-d shapes. *IEEE Trans. Pattern Anal. Mach. Intell.*, 14(2):239–256, 1992. 2
- [3] W. Chang and M. Zwicker. Automatic registration for articulated shapes. *Comput. Graph. Forum*, 27(5):1459–1468, 2008. 2
- [4] W. Chang and M. Zwicker. Range Scan Registration Using Reduced Deformable Models. *Comput. Graph. Forum*, 28(2):447–456, 2009. 2
- [5] W. Chang and M. Zwicker. Global registration of dynamic range scans for articulated model reconstruction. *ACM Trans. Graph.*, 30(3):1–15, 2011. 1, 2
- [6] Y. Chen and G. Medioni. Object modeling by registration of multiple range images. *Image and Vision Computing*, 10(3):145–155, 1992. 2
- [7] E. De Aguiar, C. Stoll, C. Theobalt, N. Ahmed, H.-P. Seidel, and S. Thrun. Performance capture from sparse multi-view video. *ACM Trans. Graph.*, 27(3):98:1–98:10, 2008. 1
- [8] Q.-X. Huang, B. Adams, M. Wicke, and L. J. Guibas. Non-rigid registration under isometric deformations. *Comput. Graph. Forum*, 27(5):1449–1457, 2008. 2, 5
- [9] M. Levoy, K. Pulli, B. Curless, S. Rusinkiewicz, D. Koller, L. Pereira, M. Ginzton, S. Anderson, J. Davis, J. Ginsberg,



Figure 11. Results of human body modeling.

- J. Shade, and D. Fulk. The digital michelangelo project: 3d scanning of large statues. In *Proc. of SIGGRAPH*, pages 131–144, 2000. 1
- [10] H. Li, B. Adams, L. J. Guibas, and M. Pauly. Robust single-view geometry and motion reconstruction. *ACM Trans. Graph.*, 28(5):175:1–175:10, 2009. 1, 2
- [11] H. Li, L. Luo, D. Vlasic, P. Peers, J. Popović, M. Pauly, and S. Rusinkiewicz. Temporally Coherent Completion of Dynamic Shapes. *ACM Trans. Graph.*, 31(1):2:1–2:11, 2012. 2
- [12] H. Li, R. W. Sumner, and M. Pauly. Global correspondence optimization for non-rigid registration of depth scans. *Comput. Graph. Forum*, 27(5):1421–1430, 2008. 2, 3
- [13] M. Liao, Q. Zhang, H. Wang, R. Yang, and M. Gong. Modeling deformable objects from a single depth camera. In *Proc. of ICCV*, pages 167–174, 2009. 1, 2
- [14] K. Michael, B. Matthew, and H. Hugues. Poisson surface reconstruction. In *Proc. of SGP*, pages 61–70, 2006. 2
- [15] N. J. Mitra, S. Flöry, M. Ovsjanikov, N. Gelfand, L. Guibas, and H. Pottmann. Dynamic geometry registration. In *Proc. of SGP*, pages 173–182, 2007. 1, 2
- [16] R. A. Newcombe, S. Izadi, O. Hilliges, D. Molyneaux, D. Kim, A. J. Davison, P. Kohli, J. Shotton, S. Hodges, and A. Fitzgibbon. Kinectfusion: Real-time dense surface mapping and tracking. In *Proc. of ISMAR*, pages 127–136, 2011. 1, 3
- [17] K. Pulli. Multiview registration for large data sets. In *Proc. of 3DIM*, pages 160–168, 1999. 3, 6, 7
- [18] S. Rusinkiewicz and M. Levoy. Efficient variants of the icp algorithm. In *Proc. of 3DIM*, 2001. 2, 4
- [19] G. C. Sharp, S. W. Lee, and D. K. Wehe. Multiview registration of 3d scenes by minimizing error between coordinate frames. *IEEE Trans. Pattern Anal. Mach. Intell.*, 26(8):1037–1050, 2004. 3
- [20] R. W. Sumner, J. Schmid, and M. Pauly. Embedded deformation for shape manipulation. *ACM Trans. Graph.*, 26(3):80:1–80:8, 2007. 2, 3, 4
- [21] J. Süßmuth, M. Winter, and G. Greiner. Reconstructing animated meshes from time-varying point clouds. *Comput. Graph. Forum*, 27(5):1469–1476, 2008. 2
- [22] N. Y. T. Masuda, K. Sakaue. Registration and integration of multiple range images for 3-d model construction. In *Proc. of CVPR*, pages 879–883, 1996. 3
- [23] A. R. T. Tevs, A. Berner, M. Wand, I. V. O. Ihrke, M. Bokeloh, J. Kerber, and H.-p. Seidel. Animation cartography - intrinsic reconstruction of shape and motion. *ACM Trans. Graph.*, 31(2):12:1–12:15, 2012. 1, 2
- [24] J. Tong, J. Zhou, L. Liu, Z. Pan, and H. Yan. Scanning 3d full human bodies using kinects. *IEEE Trans. Vis. Comp. Graph.*, 18(4):643–650, 2012. 1, 2
- [25] D. Vlasic, P. Peers, I. Baran, P. Debevec, J. Popović, S. Rusinkiewicz, and W. Matusik. Dynamic shape capture using multi-view photometric stereo. *ACM Trans. Graph.*, 28(5):174:1–174:11, 2009. 1
- [26] M. Wand, B. Adams, M. Ovsjanikov, A. Berner, M. Bokeloh, P. Jenke, L. Guibas, H.-P. Seidel, and A. Schilling. Efficient reconstruction of nonrigid shape and motion from real-time 3D scanner data. *ACM Trans. Graph.*, 28(2):15:1–15:15, 2009. 1, 2
- [27] A. Weiss, D. Hirshberg, and M. J. Black. Home 3d body scans from noisy image and range data. In *Proc. of ICCV*, pages 1951–1958, 2011. 2