Reliable Fusion of Knee Bone Laser Scans to Establish Ground Truth for Cartilage Thickness Measurement

Ming-Ching Chang[†] Nhon H. Trinh[‡] Braden C. Fleming[*] Benjamin B. Kimia[‡]

[†]GE Global Research Center, Niskayuna NY USA
[‡]LEMS, Division of Engineering, Brown University, Providence RI USA
[*]Dept. of Orthopaedics, Brown Medical School, Rhode Island Hospital, Providence RI USA

ABSTRACT

We are interested in establishing ground truth data for validating morphology measurements of human knee cartilage from MR imaging. One promising approach is to compare the high-accuracy 3D laser scans of dissected cadaver knees before and after the dissolution of their cartilage. This requires an accurate and reliable method to fuse the individual laser scans from multiple views of the cadaver knees. Unfortunately existing methods using Iterative Closest Point (ICP) algorithm from off-the-shell packages often yield unreliable fusion results. We identify two major sources of variation: *(i)* the noise in depth measurements of the laser scans is significantly high and *(ii)* the use of point-to-point correspondence in ICP is not suitable due to sampling variation in the laser scans. We resolve the first problem by performing adaptive Gaussian smoothing on each individual laser scans prior to the fusion. For the second problem, we construct a surface mesh from the point cloud of each scan and adopt a point-to-mesh ICP scheme for pairwise alignment. The complete surface mesh is constructed by fusing all the scans in the order maximizing mutual overlaps. In experiments on 6 repeated scanning trials of a cadaver knee, our approach reduced the alignment error of point-to-point ICP by 30% and reduced coefficient of variation (CV) of cartilage thickness measurements from 5% down to 1.4%, significantly improving the method's repeatability.

Keywords: knee cartilage, morphology measurement, ground truth validation, fusion laser scans, iterative closest point, ICP

1. INTRODUCTION

One major disease causing disability in elderly people is osteoarthritis (OA),[1] which is caused by the degeneration of cartilage in articulating joints. An effective method to monitor its progress and treatment effectiveness is to determine and analyze morphological measurements of the cartilage such as volume, thickness, and surface curvature.[2,3] Obtaining accurate measurements, however, is a very challenging task, partly because knee cartilage is very thin and has complex curvy shapes. Currently the most predominant method is to segment and compute cartilage morphology from magnetic resonance imaging (MRI) of the knees.[4] Despite its superior ability to differentiate soft tissues, MRI images often have low signal-to-noise ratio and may contain artifacts, especially those caused by screws in OA patients under treatment,[2,4] . Typical measuring errors such as the *partial volume effect* can translate to significant relative errors. The voxel size for MRI is typically 0.3-1.0 mm while the average thickness of knee articular cartilage is only about 1.3-2.5 mm.[2] Thus one pixel error could lead to a 25% change in the measured thickness of the cartilage, which is potentially larger than the actual changes caused by the progression of OA and, as a result, prevents its use for monitoring the disease's progression. These difficulties highlight the importance of accurate and reliable measurements. Thus, it is critical that MRI-based methods for measuring cartilage morphology be **validated** for accuracy and reliability before they can be used for clinical purposes.

A key requirement to validate morphology measurements of knee cartilage from MRI is the availability of **ground truth**. Several methods have been described in the literature: *(i)* water displacement of surgically removed cartilage tissue, *(ii)* manual segmentation, *(iii)* microscopic examination of high resolution scans

Corresponding author: Ming-Ching Chang, changm@research.ge.com, Phone: +1 518 387 5242.

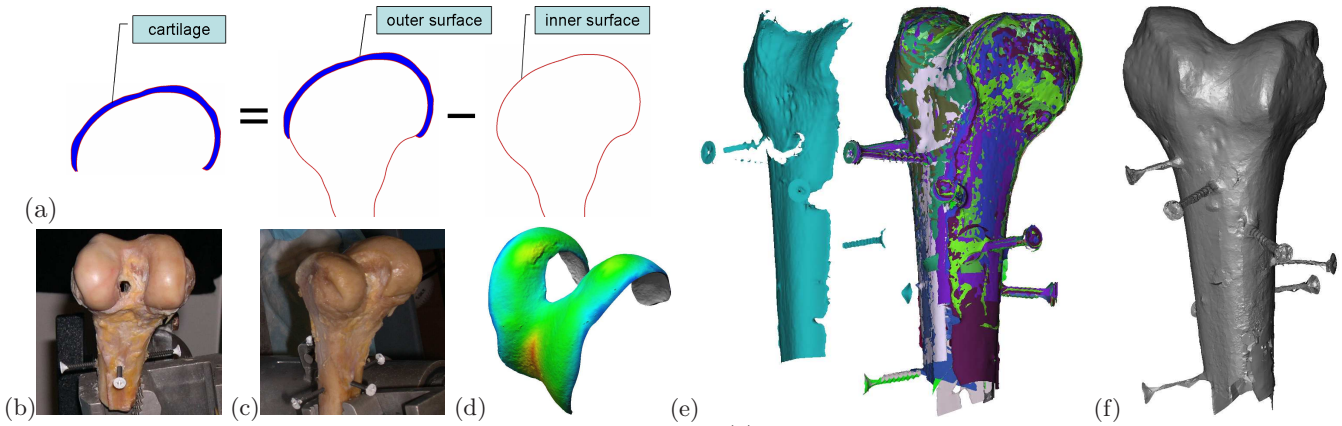

Figure 1. Overview of the laser-scanning validation approach. (a) Knee cartilage surface mesh is constructed as the difference between two bone surfaces: one with the cartilage intact as in (b) and one after it has been removed as in (c). (d) The cartilage surface mesh is constructed from the scans of (b,c) and colored by its thickness.[5] (e) Fusing 20 laser scan surfaces of a dissected femur into a final 3D model shown in (f), which typically consists of several million scan points and triangles. Our approach accurately fuses raw scans and validates the reliability of morphological measurements. The skews serve as fiducials for accurate alignment in fusion.

of anatomical sections obtained by high precision saws, *(iv)* computed tomography (CT) arthrography, *(v)* stereophotogrammetry,[2] and *(vi)* using a high resolution 3D laser scanner to interrogate the surface geometry of the articular cartilage.[5,6] Among them the last method *(vi)* is the most promising because it gives not only thickness and volume measurements but also a 3D model for the cartilage, which allows for further morphological analysis. It computes the cartilage surface as the difference of two bone surfaces, one with the cartilage intact and one after the cartilage has been meticulously removed. Each bone surface is constructed from combining multiple partial scans of the bone. In our previous work,[5] the method has been validated for its accuracy but its repeatability was not fully investigated.

In this paper, we examine the *repeatability* of the laser-scan method. This is a crucial evaluation criteria: if laser scanner measurements are to be used as ground truth to validate MRI measurements, it is necessary that its repeated measurements are more consistent than those of the MRI's. As a preview, we found in experiments with conventional processing of laser scan data, *e.g.* by using off-the-shell software packages such as the PolyWorks, that the repeatability is not sufficient to the extent[5] — the average *coefficient of variation* (CV) of the thickness measurement is 5%, which is about the same level as that of measurements from MR images.[6] Given this insufficiency, we examine the sources of such variations in laser scan measurements and propose an approach to improve the consistency of the fusion of raw scan data. This approach is the main contribution of this paper and is summarized in Figure 1.

In the remaining of the paper, we will first summarize our previous work in measuring cartilage morphology using a 3D laser scanner (§ 2) and identify two sources that cause high variation in the measurements: *sensing* error and *sampling* error. Since the knee bone surface is generally smooth, sensing error is removed by adaptively smoothing the laser scan data (§ 3). Sampling error is removed by replacing point-to-point correspondence in ICP with *point-to-mesh* correspondence (§ 4). After the individual scans are pairwise aligned, we join all the overlapping partial scans by first following the minimal spanning tree of their region adjacency graph and then refining the mesh toward a globally optimal alignment (§ 5). In experiments, we show that our approach reduces the error of pairwise alignment of partial scans by 30% and and improve the repeatability of the thickness measurements significantly —the coefficient of variation is reduced from 5% to 1.4% (§ 6).

2. MORPHOLOGY MEASUREMENT USING 3D LASER SCANS

In this section, we summarize our approach as described in[5] that reconstructs a surface mesh for knee articular cartilage and identify the sources that cause its high measurement variation. In our approach, the bone surface is first scanned with the cartilage intact and then scanned again after the cartilage has been dissolved using

diluted bleach. In all of our experiments, we use the ShapeGrabber® PLM300 Laser Scanner Systems which features a linear motion of 300 mm and a scan head (SG 1000) with a depth field of 250-900 mm and a depth accuracy of 250 μm at the farthest point. Since each scan can only capture a portion of the bone surface, multiple partial scans under different views are necessary to cover the surface entirely. We used a commercial 3D mesh processing software package, PolyWorks, to align the partial scans and produce a mesh surface. Both the interior and exterior cartilage surfaces were extracted, aligned, and combined to measure the cartilage morphology, which includes volume, surface area, and thickness.

Traditionally total cartilage volume has been the major quantification to evaluate cartilage assessment method. However, recent studies have suggested that it is not an accurate metric for the degeneration of osteoarthritis, and that alternative measurements such as *focal volume* and *focal thickness* should be used.[7] In this paper we use focal thickness of the **load bearing regions** of the femoral cartilage, which are the areas in contact with the tibial cartilage as the knee joint moves, as the main quantified assessment for knee cartilage, Figure 2. The focal thickness is obtained as follows. First a cylinder is fit to the bone-cartilage interface of the femoral cartilage model of the tibiofemoral joint, which is essentially the posterior half of the femoral cartilage. The "notch" marking the junction between the tibiofemoral and patellofemoral joints on the lateral condyle is identified on the femur bone surface mesh and marks the 0^o-axis of the cylindrical coordinate system on the femur. Each condyle of the tibiofemoral joint was then divided at 40^o, 70^o, 100^o, and 130^o from the notch point toward the posterior aspect of the condyles to create 6 patches of cartilage (3 medial, 3 lateral), and are denoted (I)-(VI), Figure 2(d). The width of each patch is 20% of the overall medial-lateral width of the femoral cartilage and centered about the centroid of each condyle.

We investigate the repeatability of the laser-scanning method by applying it to a dissected femur of fresh-frozen cadaver knee (ID: SP1133R). The femur surface is scanned repeatedly 6 times with the cartilage intact, each consisting of 20 partial scans to fully cover the bone surface, and 6 times after the cartilage is dissolved. This gives 6 pairs of with-cartilage/no-cartilage bone surfaces, from which we compute the regional thickness of the femoral cartilage, Table 1. The average coefficient of variation (CV) for the thickness measurements of the 6 regions is 5.27%. This is about the same level as that of measurements from MRI images,[6] rendering these measurements not sufficiently reliable to be used as ground truth for MRI.

A major step that affects the resulting measurement significantly is the *alignment* of the multiple partial scans. Pairwise alignment of the partial scans is carried out using the well-known Iterative Closest Point (ICP) algorithm, which minimizes the "distance" between the overlapping regions in an iterative fashion.[8] For accurate and reliable results, it is crucial this distance metric is estimated accurately and reliably. A naive approach is to estimate the distance based on point samples. We define the *point-to-point* distance d_{pt-pt} between two sampled surfaces $A = \{a_i \,|\, i = 1, ..., n\}$ and $B = \{b_j \,|\, j = 1, ..., m\}$, where a_i and b_j are sample points, to be the root mean

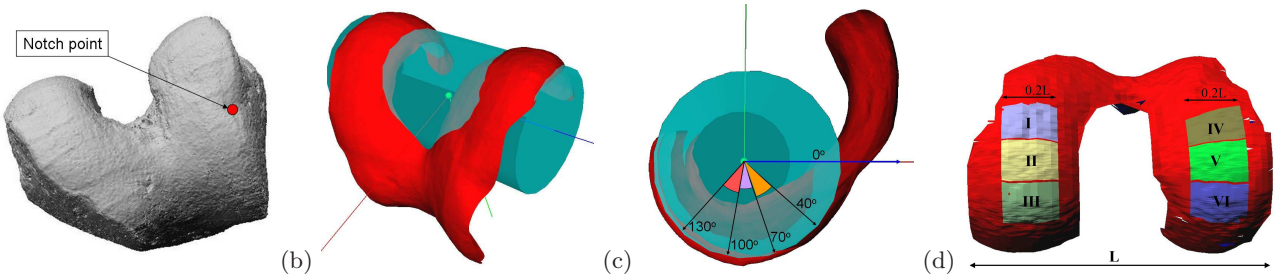

(a) (b) (c) (d)

Figure 2. (a) A reliable landmark referred o as the "notch" point, is marked on the bone surface. It combined with (b) the cylinder fit to femoral cartilage of the tibiofemoral joint to define a coordinate system on the femur. (c) and (d): 6 regions on the load bearing areas are defined using this coordinate system, each of which correspond to a 30^o rotation on the cylinder and has a width equal to 20% with the total width of the fit cylinder.

Region	(I)	(II)	(III)	(IV)	(V)	(VI)
Thickness (mm)	1.43	1.44	1.63	2.30	2.37	2.06
CV (%)	5.49	5.52	7.24	5.68	3.31	4.38

Table 1. Femoral cartilage thickness at the load-bearing regions of the tibiofemoral joint measured using the method of.[5]

square (RMS) of distances from every sample point b_j in B to its closest neighboring sample point $\hat{a}_i$ in A :

$$d_{pt-pt}(A, B, d_{th}) = RMS\left\{d(b_j, \hat{a}_i), \forall b_j \in B,\ \hat{a}_i \in A,\ d < d_{th}\right\}, \tag{1}$$

where $d(b_j, \hat{a}_i)$ denotes the 3D Euclidean distance between b_j and $\hat{a}_i$, and d_th is a constant representing the neighborhood size. This distance definition is robust against sample density changes and extends naturally as the sample density approaches infinity (to the continuous case).

We identified two sources of errors that affect heavily this distance estimation in fusing the laser scans: *sensing error* and *sampling error*. *Sensing error* refers to the error in depth estimation as the laser scanner scans the object surface. *Sampling error* refers to the error distance measurement due to point-wise sampling of the partial scans. We examine each in details and describe solutions to reduce them in the following two sections.

3. SOLUTION TO SENSING ERROR: BILATERAL GAUSSIAN SMOOTHING

Sensing error is caused by high relative error in depth estimation (z-axis) in laser scanning systems, which is typically high in most systems. In our system, the error in depth estimation (z-axis) is approximately 250 μm, while the resolutions along each scan line (y-axis) and between scan lines (x-axis) are 20 μm and 100 μm, respectively. Such high errors can significantly affect the accuracy in alignment, as depicted in Figure 3.

Since the cartilage regions on the dissected bone surfaces are generally glossy, the sensing error can be safely reduced by smoothing out the scanned data without losing important features of the scans. We have examined several smoothing methods, including the median filtering, discrete curve shortening (DCS) along scan lines, and variants of Gaussian smoothing, and find 2D adaptive bilateral Gaussian smoothing to perform the best, with respect to effectiveness and suitability for processing large amount of data. We consider each scan as a two-dimensional signal, $z = f(x, y)$, and apply 2D Gaussian smoothing to each location (x, y) using only neighboring points (kernel size k_σ) whose z-coordinates are not too different from its z-coordinate, $f(x, y)$ (distance threshold t_h). Kernel size k_σ and threshold t_h are estimated automatically from the average sampling density along and across scan lines, Figure 4(b). Similar to *bilateral filtering* in mesh smoothing,[9–11] this method preserves sharp discontinuities of scan boundaries, which is a desirable property. Figure 4 shows an example result, where sensing error is largely reduced.

4. SOLUTION TO SAMPLING ERROR: USE OF POINT-TO-MESH CORRESPONDENCE

Pairwise alignment of the partial scans is performed using the well-known Iterative Closest Point (ICP) algorithm, which is the dominant solution in fusing scans[8] . The ICP algorithm, and its variants, iteratively seeks closest points as corresponding points to register point clouds or meshes and computes a transformation that minimizes a disparity function. *Sampling error* arises because distances between the sample points which determine the

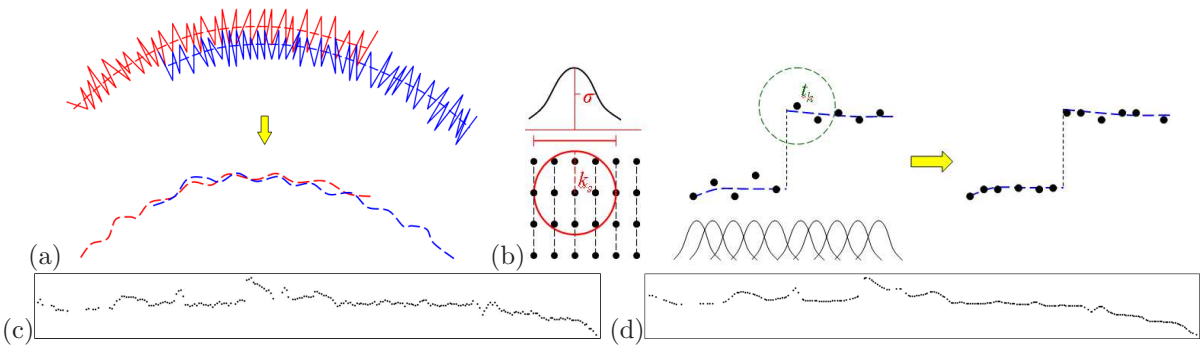

Figure 3. (a) Schematic illustration of the fusion of two 1D profile curves with up to 250% variation in depth relative to the sampling resolution. Scans with such high sensing error can significantly affect the best alignment. We propose to first reduce such sensing error by adaptive smoothing prior to fusion. (b) Illustration of applying the 2D bilateral Gaussian smoothing on the scan grid points. (c,d) An example profile of a raw scan line and after smoothing. Observe how the sharp discontinuities and features are preserved.

alignment between the partial scans do not reflect the true distance between the two scans, and thus distort the distance metric, Figure 5(a). This error can be eliminated by explicitly meshing the point cloud of each scan to approximate its true underlying surface and computing the distances between them using *point-to-mesh* distance measurement $d_{pt-mesh}$:

$$d_{pt-mesh}(A, B, d_{th}) = RMS\left\{d(b_j, \breve{a}_i), \forall b_j \in B,\ \breve{a}_i \in A,\ d < d_{th}\right\}, \tag{2}$$

where $\breve{a}_i$ is the foot point in A when projecting the sample point b_j in B onto the mesh surface of A, Figure 5(c).

We use the algorithm described in[12] to generate a mesh from the point cloud of each scan. This algorithm first computes the *medial scaffold* (medial axis) of the sample points and then uses it to construct a surface mesh by considering triangles that best recover the underlying surface between samples. We choose this algorithm because it is generic in handling unorganized point cloud input and does not enforce strong restriction on the input sampling condition and surface topology.

While *closest* point-to-mesh correspondence is preferred, its exact computation is not trivial.[13] We approximate the exact point-to-mesh distance by using the closest point-to-point computation. Specifically, we find the closest vertex on the other mesh and search in the *one-ring* of triangles of it, Figure 5(c). The chance of finding the true closest point can be increased by exploring the top n closest vertices instead of only the closest vertex.

Figure 5(d) illustrates an example where a combination of smoothing and meshing of raw scans effectively reduce the sensing and sampling error and improve fusion accuracy. In this example one scan (gray color) is fixed while the other (multiple colors) is moved toward the best alignment. The x- and y- scanning resolutions are 0.1 mm and the noise along the z-direction is on average 0.2 mm The multiple colors depict the closet $d_{pt-mesh}$, ranging from 0 mm (green) to 0.3 mm (red). After smoothing and meshing, the two scans aligned well and the final average distance between them was $0.05mm$, which was only 50% of the scanning resolution.

5. FUSION OF MULTIPLE PARTIAL LASER SCANS

The last step in our procedure is to fuse all scans into a final surface model for morphology measurement. Even with all pairwise alignment between overlapping scans computed, this is still a very challenging task[8,14–16] for the following two reasons. First, this is an *over-constrained* problem: Since a scan surface could overlap with many others, locally optimal alignments may conflict with each other. Second, the total amount of scan data could be very large (hundreds of millions of points), rendering a naive brute-force approach not practical.

We propose to fuse the multiple partial scans by first estimating an initial global alignment using a limited number of pairwise overlaps and then continuing refining them using all pairwise overlaps until convergence. The overlaps between individual scans can be represented as a *region adjacency graph* (RAG) , denoted G, where each node is a partial scan and each edge represents an overlap between partial scans. We associate each edge with a *weight* reflecting the number of scan points in the overlapping region, which approximates the area of the overlapping region. We note any *loop* in G introduces redundancy in the alignment. In the inevitable presence of sampling errors, these loops will necessarily cause conflicts between the alignment. We produce an initial alignment of the scans by focusing more on the large overlaps (stronger edges) and ignore the weaker ones.

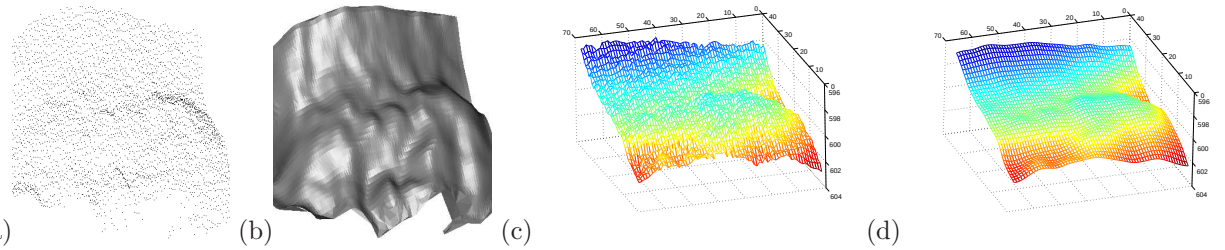

(a) (b) (c) (d)

Figure 4. Example results of smoothing a portion of the knee bone scans: (a) The noisy raw scan point cloud. (b) Result after the adaptive bilateral Gaussian smoothing, where the point cloud are meshed into surfaces for better visualization. (c,d) Respective mesh plots of the raw bone scan and after smoothing. Observe how well the sensing error is reduced.

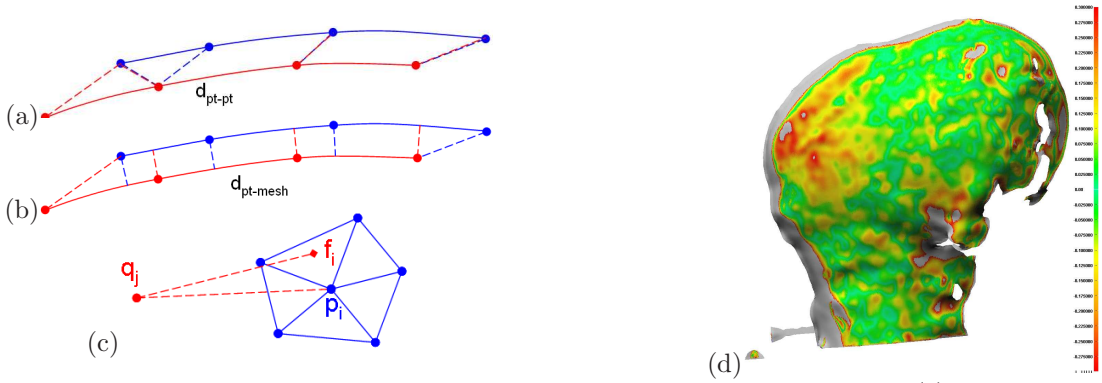

Figure 5. Using point-to-mesh correspondence to reduce the sampling error in alignment. (a) Schematic illustration of fusing two 1D curves using (a) point-to-point and (b) point-to-mesh correspondences, where in the later the error is largely reduced. (c) Approximate the point-to-mesh distance from query point q_i by first finding the closest vertex p_i on the other mesh and search for its *one-ring* adjacent faces $\{f_i\}$. (d) Result of fusing two partial bone scans after smoothing and meshing (see text).

Specifically we follow the *minimum spanning tree* (MST) of G and proceed as follows. Our algorithm mimics Prim's algorithm in computing the MST of a graph:

1. Select the scan with the most overlaps as the root node in the aligned set A.
2. Among the *unvisited* scans which overlap with A, select the scan B whose edge weight reflects the largest overlap with A.
3. Fuse scan B together with scan set A: $A = A \cup B$, and mark B as *visited*.
4. Repeat step 2 until all scans are visited.

Figure 6(a-b) illustrates an example of fusion scans following the MST of the RAG of the scans.

Once an initial global alignment is established, we account for the weak overlaps by distributing the misalignment evenly among the pairs, such that the overall consistency improves.[14] Specifically, we minimize the global error estimation E of all scans of G, defined as the sum of all distances $d_{pt-mesh}$ between all overlapping pairs, in an iterative *best-first* fashion: *(i)* select the edge $e = (v_i, v_j)$ of G with largest error and refine the alignment between scans v_i and v_j to reduce E, and *(ii)* repeat *(i)* until E converges.

The final surface model is computed by *averaging* the aligned overlapping partial scans. First each sample point is moved to its final position by averaging its point-to-mesh correspondences in other overlapping surfaces. Then the combined point cloud is meshed again using[12] to construct the final surface mesh. Due to dense and repeating scans with overlapping regions, the final mesh vertex density is typically several times denser than the original resolution of each individual scan, Figure 6(c-e), typically containing tens of millions of triangles. It can be optionally compressed by applying mesh decimation algorithms such as.[17]

6. EXPERIMENTAL RESULTS

We have implemented our laser scan processing pipeline, including the smoothing, meshing, aligning, and fusing of partial scans into a final surface model. We now report experimental tests of its reliability in morphological measurements. We present two major improvements over previous methods. The first is on the effectiveness of reducing the sensing and sampling errors during pairwise alignment. The second is on the improvement of fusion repeatability for morphological measurements.

First we demonstrate the effectiveness of smoothing and using point-to-mesh correspondence on pairwise scan alignment. We start with two scans A and B which have been aligned using point-to-point ICP alignment. After each scan is individually smoothed – denoting the results as $\widetilde{A}$ and $\widetilde{B}$ – and meshed into a surface (denoted as $\widetilde{A}^*$ and $\widetilde{A}^*$), we apply point-to-mesh ICP to refine the alignment ($\widetilde{A}^*$ is kept fixed and $\widetilde{B}^*$ is moved toward

(a)

vi	vj	overlap pt.
5	19	205956
19	20	201844
1	16	185533
4	19	180072
1	2	178615
5	18	174340
15	16	169685
10	11	158759
8	16	155896
3	4	148243
17	18	140220
6	7	138702
6	11	136569
2	3	129820
6	19	129740
12	13	129123
9	14	126397
13	14	118288
14	15	103990

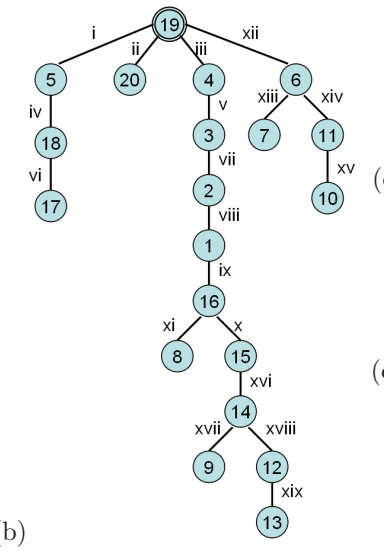

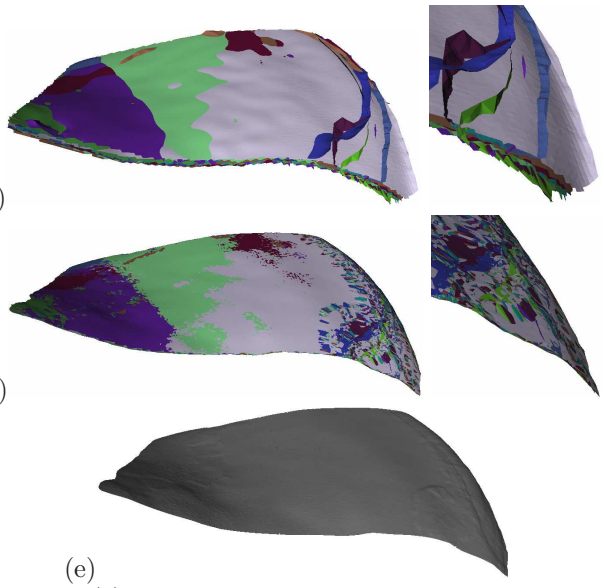

Figure 6. Fusion of 20 scans following the MST of the RAG of the scans. (a) List of MST edges ordered by the number of overlapping points between partial scans of a trial on fusing 20 scans of a knee bone. (b) The MST used to fuse all 20 scans, where each partial scan represented as a vertex here. The Romanic number shows the order of fusion partial scans into the aligned set. (c) Result of a cropped view of the global alignment of 11 overlapping individual scan surfaces following the MST alignment and subsequent refining, each surface scan is show in different color. (d) Result of deforming all overlapping surfaces (locally moving all vertices) to their average position. The estimated 'thickness' (variation of measurement) of the surface is reduced from 0.16 mm from (c) to 0.02 mm in (d), which is far below the scanning resolution (0.1mm). (e) Final surface mesh constructed by meshing all sample points from (d).

it as $\widehat{B}^*$). As shown Figure 7, pairwise alignment error is reduced by 32%. We note that in the case without smoothing (performing the point-to-mesh ICP directly), the error reduction is only 3%, which is much less then the combination of two steps and confirms the crucial role smoothing for accurate subsequent alignment.

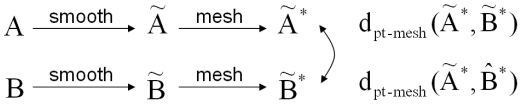

	$d(\widetilde{A}^*, \widetilde{B}^*)$	$d(\widetilde{A}^*, \widehat{B}^*)$	diff. (%)
Raw data	0.1163	0.1129	2.9%
Smoothed	0.0911	0.0791	32.0%

Figure 7. Compare the pairwise point-to-mesh distance measurements between two partially overlapping scans A and B in their initial point-to-point ICP alignment, after smoothing/meshing, and after refinement using point-to-mesh ICP (d_{th} = 0.536 mm, which is twice the grid size). The last column shows the final error/distance that is reduced in the process (up to 32% error reduction).

We repeat the thickness measurement of the load-bearing regions of the SP1133 femur for all 6 trials to examine its repeatability. As summarized in Table 2, the repeatability of the procedure improves significantly over our previous method described in.[5] The mean coefficient of variation (CV) was reduced from 5% to 1.4%. This significantly improves the reliability of morphology measurements using laser scan data and render them sufficiently reliable to be used as ground truth for validating measurements computed from MRI images.

7. CONCLUSIONS

We have proposed a reliable approach to accurately process raw laser scans to produce a consistent ground truth dataset. By reducing the sensing and sample errors in the fusion process, the resulting model has low variation and excellent repeatability, which is useful for validating the MRI measurements of the knee cartilage.

Acknowledgements. This material is based upon work partly supported by the National Science Foundation under Grants NSF/IIS-0413215 and NSF/ITR-0205477.

Region	(I)	(II)	(III)	(IV)	(V)	(VI)
Trial 1	1.73	1.67	1.81	2.45	2.51	2.17
Trial 2	1.72	1.73	1.89	2.42	2.44	2.20
Trial 3	1.72	1.72	1.85	2.45	2.52	2.23
Trial 4	1.75	1.73	1.79	2.41	2.58	2.19
Trial 5	1.75	1.73	1.90	2.44	2.48	2.19
Trial 6	1.71	1.73	1.91	2.41	2.49	2.23
Mean	1.73	1.72	1.86	2.43	2.50	2.20
CV (%)	1.00	1.27	2.48	0.75	1.86	1.10

Table 2. Result of 6 trials of femoral thickness measurements at the load-bearing regions of the tibiofemoral joint of SP1133R. The CV value indicates the repeatability of our method. All thickness measurements are in *mm*.

REFERENCES

[1] Reginster, J. Y., "The prevalence and burden of arthritis," *Rheumatology* **41**, 3–6 (April 2002).

[2] Eckstein, F. and Glaser, C., "Measuring cartilage morphology with quantitative magnetic resonance imaging," *Seminars in Musculoskeletal Radiology* **8**(4), 329–53 (2004).

[3] Ateshian, G. A., Soslowsky, L. J., and Mow, V. C., "Quantitation of articular surface topography and cartilage thickness in knee joints using stereophotogrammetry," *Journal of Biomechanics* **24**(8), 761–776 (1991).

[4] Berquist, T. H., "Imaging of articular pathology: MRI, CT, arthrography," *Clinical Anatomy* **10**(1), 1–13 (1997).

[5] Trinh, N. H., Lester, J., Fleming, B. C., Tung, G., , and Kimia, B. B., "Accurate measurement of cartilage morphology using a 3D laser scanner," in [*Proc. Computer Vision Approaches to Medical Image Analysis (CVAMIA)*], 37–48 (2006).

[6] Koo, S., Gold, G. E., and Andriacchi, T. P., "Considerations in measuring cartilage thickness using MRI: factors influencing reproducibility and accuracy," *Osteoarthritis and Cartilage* **13**, 782–9 (September 2005).

[7] Gandy, S. J., Dieppe, P. A., Keen, M. C., Maciewicz, R. A., Watt, I., and Waterton, J. C., "No loss of cartilage volume over three years in patients with knee osteoarthritis as assessed by magnetic resonance imaging," *Osteoarthritis and Cartilage* **10**(12), 929–937 (2002).

[8] Rusinkiewicz, S. and Levoy, M., "Efficient variants of the icp algorithm," in [*Proc. 3-D Digital Imaging and Modeling*], 145–152 (2001).

[9] Fleishman, S., Drori, I., and Cohen-Or, D., "Bilateral mesh denoising," *ACM Trans. Graph.* **22**(3), 950–953 (2003).

[10] Hou, Q., Bai, L., and Wang, Y., "Mesh smoothing via adaptive bilateral filtering," *Lecture Notes in Computer Science, Workshop on Computer Graphics and Geometric Modeling (CGGM)* **3515**, 273–280 (2005).

[11] Chen, C.-Y. and Cheng, K.-Y., "A sharpness dependent filter for mesh smoothing," *Comput. Aided Geom. Des.* **22**(5), 376–391 (2005).

[12] Chang, M.-C., Leymarie, F. F., and Kimia, B. B., "Surface reconstruction from point clouds by transforming the medial scaffold," *Computer Vision and Image Understanding* **113**, 1130–1146 (April 2009).

[13] Gueziec, A., ""meshsweeper": dynamic point-to-polygonal mesh distance and apps.," *IEEE Transactions on Visualization and Computer Graphics* **7**, 47–61 (2001).

[14] Pulli, K., "Multiview registration for large data sets," in [*Proc. 3-D Digital Imaging and Modeling*], 160–168 (1999).

[15] Curless, B. and Levoy, M., "A volumetric method for building complex models from range images," in [*SIGGRAPH*], 303–312 (1996).

[16] Turk, G. and Levoy, M., "Zippered polygon meshes from range images," in [*SIGGRAPH*], 311–318 (1994).

[17] Hussain, M., "Fast decimation of polygonal models," *Lecture Notes in Computer Science, Advances in Visual Computing* **5358**, 119–128 (2008).