

Arthroial Joint Markerless Cross-Parameterization and Biomechanical Visualization

G. Elisabeta Marai*
Brown University

Cindy M. Grimm†
Washington University

David H. Laidlaw‡
Brown University

ABSTRACT

Orthopedists invest significant amounts of effort and time trying to understand the biomechanics of arthroial (gliding) joints. While new image acquisition and processing methods currently generate richer-than-ever geometry and kinematic datasets that are individual-specific, the computational and visualization tools needed to enable the comparative analysis and exploration of these datasets lag behind.

In this paper, we present a framework that enables the cross-dataset visual exploration and analysis of arthroial joint biomechanics. Central to our approach is a computer-vision inspired markerless method for establishing pairwise correspondences between individual-specific geometry. Manifold models are subsequently defined and deformed from one individual-specific geometry to another such that the markerless correspondences are preserved while minimizing model distortion. The resulted mutually-consistent parameterization and visualization allow the users to explore the similarities and differences between two datasets, and to define meaningful quantitative measures.

We present two applications of this framework to human wrist data: articular cartilage transfer from cadaver data to *in vivo* data, and cross-dataset kinematics analysis. The method allows our users to combine complementary geometry acquired through different modalities, and thus overcome current imaging limitations. Results demonstrate the technique useful in the study of normal and injured anatomy and kinematics of arthroial joints.

In principle, the pairwise cross-parameterization method applies to all spherical topology data from the same class, and should be particularly beneficial in instances where identifying salient object features is a nontrivial task.

CR Categories:

Keywords: pairwise shape correspondence, cross-parameterization, arthroial joints, biomedical visualization, biomechanics

1 INTRODUCTION

Arthroial or gliding joints are the structures that allow us to move; they are formed when two bones come together and move against each other without getting damaged. There are many different types of joints in the human body, and some of them are remarkably complex; for example, the human wrist involves contact among fifteen different bones. Furthermore, the relationships among joint-bones change with motion, age or disease. As new image acquisition and processing methods generate richer-than-ever arthroial joint datasets, the inherent complexity of these data motivates a variety

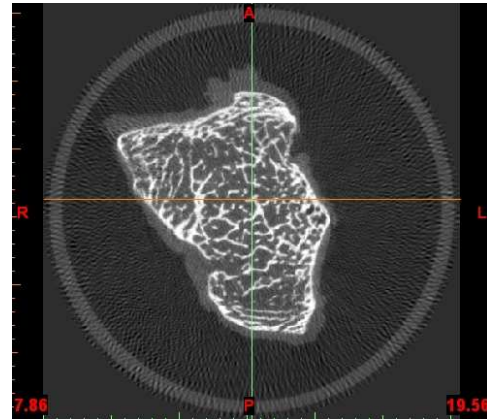


Figure 1: 2D slice through a μ CT-volume image of bone and articular cartilage. The imaged cartilage (bone tissue shown in white, cartilage in soft grey surrounding the bone outline) can be represented as a height-field on the supporting bony surface. The height-field is defined as a collection of base points on the bone surface, the bone surface normal at that point, and the height along the normal at which the normal intersects the cartilage external envelope.

of modeling and visualization techniques designed to assist orthopedics researchers in their analysis.

In the context of arthroial joint data, particular emphasis falls on the comparative analysis and exploration of individual-specific datasets. For example, orthopedic surgeons often compare the injured joint of an individual with the matching uninjured joint of the same individual. In a different application, our collaborators in the Bioengineering Department are studying the effect of *in vivo* motion on articular cartilage in the human wrist. *In vivo* motion can only be measured in live human subjects. On the other hand, highly-detailed, unloaded wrist-cartilage geometry can currently only be measured *in vitro*, by extracting the bones and cartilage from the joint, immersing them in contrast dye for 24 hours, then μ CT-imaging them. The imaged cartilage can be represented as a height-field on the supporting bony surfaces (Fig. 1). Our collaborators would like to be able to transfer this height field, without folding or tearing it, on the corresponding *in vivo* bony surfaces; then apply to the resulted geometry the *in vivo*-measured kinematics in order to estimate how cartilage deforms with motion.

Unfortunately, arthroial bones like the ones in the wrist or ankle are difficult to set in correspondence, due to their round and smooth everywhere (potato-like) shape, with rare clear salient features to aid the matching process. In general, we note that while a number of techniques are available for modeling, visualizing, and animating articular-joint data acquired from a given human individual, the computational and visualization tools needed to comparatively analyze or combine these datasets are limited.

In this paper, we present a framework that enables the cross-dataset visual exploration and analysis of arthroial joint biomechanics. Our key contribution is a markerless method for establish-

*e-mail:gem@cs.brown.edu

†e-mail:cmg@cse.wustl.edu

‡e-mail:dhl@cs.brown.edu

ing pairwise correspondence between individual-specific arthrodistal datasets. The method is markerless in the sense that it circumvents the need to identify corresponding salient geometry features (*markers*). Based on the markerless dataset correspondence, we then propose modeling and visualization techniques for exploring and analyzing cross-dataset variation. We demonstrate our framework on human wrist data, with two applications: transfer of surface-detail (such as soft-tissue insertion sites and cartilage location) from cadaver data to in vivo data; and cross-dataset kinematics analysis.

2 RELATED WORK

A key step in either the comparative analysis or the merging of geometric datasets is establishing pairwise correspondence between shape boundaries or surfaces. In standard biomechanics practice, two similar surfaces are often brought in correspondence by first aligning them through a method like the Iterative Closest Point (ICP) transform [1], and then projecting the source surface on the target surface. Because this approach does not explicitly use shape information, the source surface may either fold or tear during the projection phase. Alternatively, the problem of pairwise correspondence can be posed as that of establishing a common parameterization between the two surfaces. The advantage of cross-parameterization over ICP followed by projection is that: a) it can be constrained to avoid surface folding and tearing; and b) it provides a convenient common frame of reference when quantitatively analyzing the differences between two datasets.

Pairwise correspondence for cross-parameterization. While several cross-parameterization techniques exist, they require in general the existence of an initial pairwise correspondence among several points on the shapes. If the shapes have clear salient features (e.g., extreme curvature points, or characteristic medial axis), the pairwise correspondence process can be automated [2, 3, 4]. In the absence of salient features, pairwise correspondence is often established using user-defined landmarks [5, 6]; however, the landmarking operation is time-consuming, error-prone and subjective even in 2D. In 3D, manual landmarking can become impractical.

Golland et al. [7] avoid 3D manual landmarking by aligning shapes so that the object volume and the center of gravity of a distance transform is the same for all example shapes. This approach does not guarantee anatomical feature correspondence, although the resulting alignment is sufficient for shape classification purposes.

Brechbuhler et al [8] automate the pairwise correspondence process by assigning a correspondence of each individual shape to the sphere (SPHARM). The correspondence of SPHARM is determined by aligning the parameterization so that the ridges of the first order ellipsoid coincide; inherently, the correspondence of objects with rotational symmetry in the first order ellipsoid is ambiguously defined. The method is thus not applicable to potato-like shapes with symmetry about a main axis like the arthrodistal shapes discussed in this paper. In fact, a comparative study [9] on femoral data found that, in case of rotational symmetry in the first order ellipsoid, independent of the higher order terms, the SPHARM correspondence is inappropriate; the correspondence could not be significantly improved using statistical methods like the minimum description length [10] or DetCov [11].

In contrast, the markerless pairwise correspondence method proposed in this paper is automated, works on objects of spherical topology which lack salient features, and is not susceptible to object rotational symmetry in the first order ellipsoid.

Our correspondence approach is inspired by work in computer vision [12, 13], where 2D representations have been used in the context of same-object recognition and object-pose recovery. Global-shape 2D histograms appear often in computer vision literature, but in general they are targeted at same-object recognition.

While in computer vision applications the focus is on selecting just a few strong correspondences that are also geometrically consistent with a rigid-body transform, we introduce new match-selection criteria that allow for non-rigid shape variation. We have briefly described these match-selection criteria before in [14].

Cross-parameterization. A number of statistical methods produce cross-parameterizations across training sets of shapes [5, 7, 10, 11]. In these approaches, a statistical model of shape variation is learned from a training database. The focus of the present paper is, however, the *pairwise* cross-parameterization of shapes, when a training database may not be available.

Previous approaches [15, 16, 17] to consistent pairwise parameterization of meshes use mesh simplification to create a base mesh, align the base mesh with each data set, then re-mesh to establish correspondences. The correspondence depends heavily on the assignment of the base mesh triangles to patches on the surfaces. In our shape-correspondence approach we avoid both the patch creation and re-meshing steps by only pinning a subset of points and using an analytical domain.

Exploratory Visualization of Arthrodistal Joints. Cross-individual arthrodistal joint datasets that combine geometry and kinematic information are relatively new in orthopedics, due to the recent development of technologies for tracking in vivo joint kinematics. Previous attempts at the visual exploration of in vivo joint biomechanics have therefore been aimed mostly at understanding individual-specific kinematics [18, 19], mainly by presenting to the user animations of a specific joint. The lack of salient features makes finding correspondences in arthrodistal joint geometry difficult. In consequence, previous kinematic analysis attempts in general disconnect kinematic data from the geometry, and quantify kinematic differences across datasets by comparing numerically joint-angle series. In contrast, our approach allows the users to explore arthrodistal joint kinematics in the context of joint geometry.

The two scalar data visualization techniques we use, color mapping and iso-contouring, are well known scientific visualization techniques [20].

3 METHODS

Figure 2 summarizes our framework. We use computed-tomography (CT) individual-specific datasets; the data is first processed in order to extract the articulation-joint geometry and the corresponding joint kinematics. Next, given two such datasets, we set their geometry in correspondence using a computer-vision and differential-geometry based approach. We further process the resulting geometry and kinematics in order to explore the similarities and differences between datasets, and to define meaningful quantitative measures.

3.1 Data Acquisition and Preprocessing

Our data is generated by CT-imaging the articulation bones of a volunteer in up to twelve different positions, followed by registration of the bones across all volume-images. Each CT volume-image corresponds to an articulation pose, thus sampling the space of the physiological range of motion.

Through manual segmentation, thresholding, and user interaction, bone surface points are extracted from an arbitrarily-designated reference CT volume image [18]. Each bone is then tracked accurately through the sequence of remaining CT volume images [21]. The tracking procedure reports relative bone-motion from one articulation pose to another. Each resulting dataset is thus individual-specific, and consists of the arthrodistal joint geometry and its sampled kinematics.

Bones in the joint are modeled further both implicitly, as *scalar distance fields*, and parametrically, as *NURBS surfaces*. These two

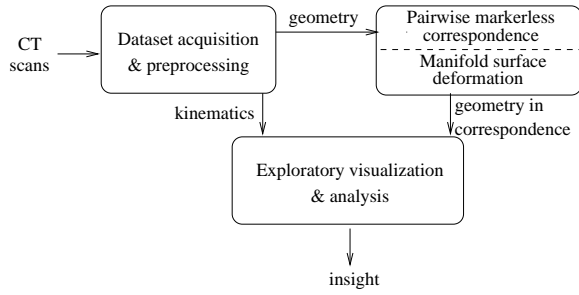


Figure 2: Biomechanics visual analytics framework. Subject-specific CT-datasets are first processed in order to extract the relevant joint geometry and kinematics. Two such datasets are then set automatically in correspondence. We use the correspondence to further explore and analyze dataset differences.

types of representation have complementary strengths for different types of calculations. NURBS surfaces provide an accurate, smooth, and locally controllable representation of the bones. We use the Geomagic software package [22] to generate parametric bone models from the segmented bone surface points. Distance fields on the other hand, have important advantages for geometric computations such as fast distance calculation. A distance field is a scalar field that specifies the signed distance from a point to the bone surface. Numerical sign is used to distinguish the inside from the outside of the bone: negative values are inside the bone, positive values are outside the bone, zero values are on the bone surface. Distance fields surrounding each bone are computed from the parametric representation. These distance fields provide the support for evaluating kinematic changes in the articulation. We use the Closest Point Transform (CPT) level-set software package [23] to generate the distance field representation from the surface representation. A validation study performed on CT-imaged phantom data [24] – in which the distance between two spheres was computed through the CPT software – has shown average errors of under $10^{-3}\%$ of the actual inter-sphere distance.

3.2 Bone Surface Correspondence

Let S and T be the source and target geometry we wish to set in correspondence. We begin by generating a set of pin-points on S and T ; the correspondence is markerless in the sense that a pin-point doesn't necessarily mark a geometrically-salient feature. The corresponding pin-points are generated using global shape information.

Next, we fit a manifold surface to the geometry of S , then deform the manifold to match the geometry of T while preserving the pin-point correspondence. Essentially, the manifold deformation extends the pin-point correspondence to the entire surface. The pin-point correspondence helps prevent manifold folds and self-intersections during the matching process. The manifold correspondence allows us to define consistent parameterizations between datasets – a bijection (i.e., a one-to-one and onto mapping) between the geometry of each surface and the abstract manifold representation.

We describe the markerless-correspondence and manifold-deformation steps below.

3.2.1 Pairwise Markerless Correspondence

Figure 3 summarizes the markerless corresponding procedure. We begin by resampling the bone surfaces of S and T such that they have approximately the same number of vertices. The resampling operation is embedded in the Geomagic preprocessing soft-




Figure 4: Project and bin operation. For each point p on the mesh, we consider the local coordinate system defined by the plane tangent to p and the surface normal N at that point. All other vertices v in the mesh are cylindrically-projected on this local system; we compute the α and β cylindrical coordinates (see text for details) of each vertex. Projected points are finally binned in a 2D histogram.

ware [22]. Corresponding bone surfaces are then normalized with respect to scale and translated so that their centers of mass are aligned. If the two shapes are significantly far apart, the surface-alignment is further refined using an ICP transform [1].

The fundamental shape element we use to generate pin-point correspondences is an oriented point, a 3D vertex with an associated direction [13]. Following Johnson and Hebert, we define an oriented point on the surface of an object using the surface position p and the surface normal n at that position. For each oriented point in a mesh we consider the tangent plane through p and the line parallel to n through p . The tangent plane and the line define a local coordinate system. All the vertices v in the mesh can be mapped on this coordinate system using cylindrical coordinates as follows:

$$\beta(v) = n \cdot (v - p) \quad (1)$$

$$\alpha(v) = \sqrt{(v - p) \cdot (v - p) - \beta(v)^2} \quad (2)$$

For each point p on the surface of a bone instance from dataset S , we compute α and β coordinates for all the vertices in the mesh. Next, we bin the vertices based on their (α, β) coordinates into a 2D histogram, with the bin size equal to the median edge length in the mesh (Fig. 4). We found that 15×15 such bins cover in general all vertices in our models. The result of this step is a collection of 2D histograms, one for each point in the bone mesh. We repeat the procedure for the corresponding bone in dataset T .

Because each 2D histogram encodes the coordinates of *all* the points on the surface of an object with respect to the local (α, β) basis, it is a local description of the global shape of the object. Since each 3D point has associated such a description, we can apply techniques from 2D image matching to the problem of surface correspondence.

Potential pin-point correspondences between the S and T bone instances are established by evaluating the value of all possible matches between points on the S model and points on the T model. We define the value of a match between point s on S and point t on T as a combination of the image-correlation index between the 2D histograms constructed at s and t and the Euclidean distance between the 3D space coordinates of s and t :

$$K_{match}(s, t) = K_{correl}(hist(s), hist(t)) + \frac{1}{\|p(s), p(t)\|} \quad (3)$$

where $K_{correl}(hist(s), hist(t))$ is the 2D correlation coefficient between the histogram at s and the histogram at t . This value function favors matches that have a strong image-correlation index and were generated from points with similar space coordinates.

For each surface point s , we retain the strongest matched point t in the other instance, i.e., the point t that generated the highest

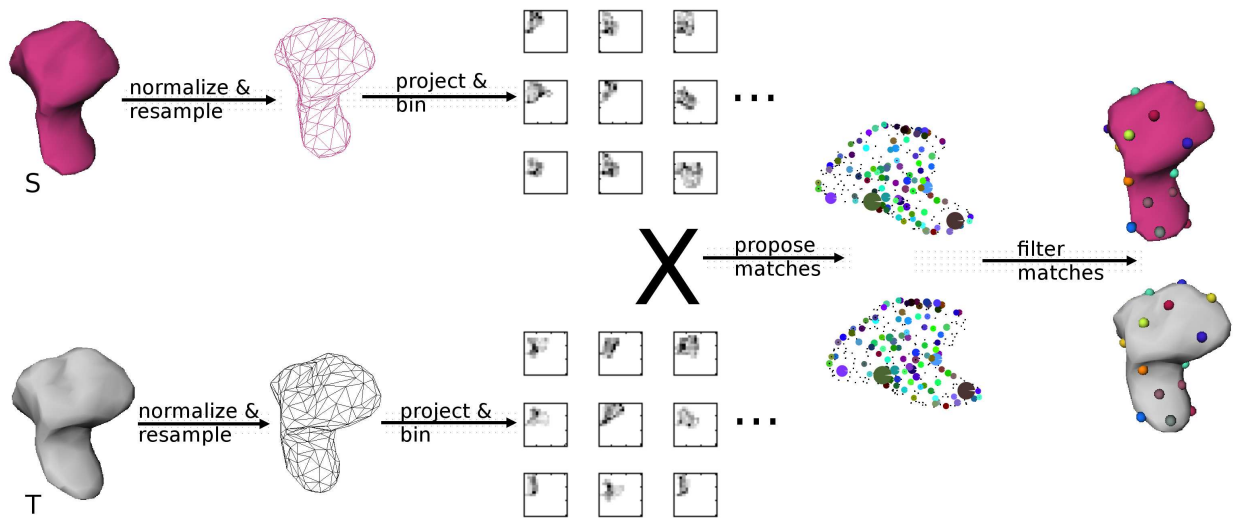


Figure 3: Markerless correspondence pipeline. Corresponding bone instances S and T are first normalized and resampled. For each surface point, we project the bone instance on a local 2D subspace and we generate a 2D histogram (for details on the project and bin operation, see Fig. 4). We evaluate all possible matches between all the points on S and all the points on T , and keep only bipartite matches. Finally, we filter out weaker matches from the neighborhood of stronger matches, generating a Poisson-disk distribution of pin-points on the surface of each mesh.




Figure 5: Two wrist bones belonging to different human subjects and their corresponding pin-points. Pin-points in correspondence are shown in the same color. Right: the 2D-histograms generated by the two pink corresponding pin-points on the left are remarkably similar, indicating good shape correlation at these points. Our markerless method generated this bipartite set of 25 Poisson-disk distributed correspondences in only a few minutes.

$K_{match}(s, t)$ score. The resulting strongest correspondences may not be in a bipartite relationship; i.e., point s_1 's strongest match may be point t_1 , while t_1 's strongest match may not be s_1 , but some other point s_2 . In a first filtering stage, we only keep correspondences that define a bipartite match.

Next, bipartite correspondences are ordered according to their K_{match} value. Beginning with the highest K_{match} value, we filter out weaker correspondences through a greedy programming approach, as follows: for each bipartite correspondence b , in the order of K_{match} , we remove all other bipartite correspondences which are closer to b than $10 \times l$, where l is the median edge length across S and T . The result is a Poisson-disk distribution of pin-points on the surface of each model. Note that the process through which we generate a set of pairwise corresponding pin-points is deterministic.

Corresponding pin-points generated on two wrist bones through this markerless method are shown in Fig. 5. The pairwise markerless correspondence method is implemented in less than 2,000 lines of Matlab code.

3.2.2 Manifold Deformation

To align the surfaces we begin by fitting a default manifold surface to the source geometry S , as described in [25, 26]. Manifolds are smooth, locally parameterized, C^2 continuous surfaces [25]. The overlapped structure of the manifold-surface representation, which is inspired by differential geometry, has several advantages including flexibility in shape adjustments without costly constraints, and smooth transitions and uniformity among patches.

The default manifold we fit to S is roughly the shape of the given bone. Summarizing the [25] reference, the fitting process essentially “shrink wraps” the default manifold around the specific data set. Specifically, we alternate between projecting the S data set points onto the manifold and solving a least-squares problem to adjust the parametric surface control points. Once the manifold surface is fit to S , we project the S pin-points onto the fitted manifold.

Next, we deform the S manifold such that it fits the target geometry T . This time, we introduce additional constraints into the least-squares problem that pull the S manifold pin-points to the 3D pin points of T . Although we cannot formally guarantee no folding or self-intersections during manifold fitting, the use of filtering and of a sufficiently-close starting position as provided by the pin-point correspondence helps us avoid folding problems in practice.

Figure 6 shows a default manifold surface fitted to S , together with its pin-points, the original mesh T with pin-points, and the original mesh T overimposed with the result of deforming the S manifold to fit T while preserving the pin-point correspondence. Note that the resulted T manifold is free of self-intersections/folds. The manifold deformation defines a mutually consistent parameterization between the S surface and the T surface.

3.3 Exploratory Visualization and Analysis

Given two datasets, the markerless correspondence method described earlier allows us to build a mutually-consistent surface parameterization between the dataset geometries. This mutually-consistent parameterization enables the exploration of dataset differences.

In order to contrast the kinematics of two datasets, we trace joint




Figure 6: Manifold deformation: default manifold surface fitted to S with pin-points (left), the original mesh T with corresponding pin-points (middle), and the original mesh T overlaid with the result of deforming the S manifold to fit T while preserving the pin-point correspondence (right); in the right image, the pin-points of T and the deformed manifold are overlapped. The resulted T manifold is free of self-intersections/folds. The manifold defines a mutually-consistent parameterization between the S surface and the T surface.

kinematics on the bone surfaces, as described in [27]. Summarizing this reference, for each bone and kinematic pose, we compute at each bone vertex the distance to neighboring bones. We use the distance field representation to find distances from every vertex in the surface model of one bone to its neighbor.

Because joint kinematics influence how close two bones come together, and where they articulate with each other, we can use the inter-bone joint space as a measure of kinematics. For each bone and space, we define the inter-bone joint space as the cortical surface area on the bone that is less than a prescribed threshold distance (typically 5mm) from the cortical surface of a neighboring bone. As two bones move relatively one to another, tracking through time the location and size of the inter-bone joint space provides insight into the joint kinematics.

We also compute isocontours on the contact area, each contour showing where the inter-bone joint space is equal to a constant distance. The area of the inter-bone joint space is the total area of the surface triangles within the 5mm contour. We characterize the inter-bone joint-space by its area and by the location of its centroid on the cross-parameterized surface. These measures and the common parameterization allow our users to compare quantitatively two datasets.

We visualize inter-bone joint-spaces using color mapping and contouring. Color maps are generated for each bone so that distance values of surface points are mapped to varying color saturations (more saturated colors represent shorter distances). Distances larger than the contact threshold value are neither colored nor contoured and are shown as white surfaces (Fig. 7).

4 RESULTS

4.1 Validation

Pairwise markerless correspondence. We applied our markerless correspondence technique to wrist bone instances collected from 11 individuals (9 hamate bones, and 4 lunate bones). For each pair of bones from the same class, 13 to 22 sets of pairwise pin-points were automatically generated through our method. Because in medical imaging the correspondence ground truth is only known for synthetic and phantom data, the correctness of each pin-point pair was visually examined by an expert user. 99.7% of the more than 700 pin-point correspondences generated were judged to be correct (within human expert accuracy). The bottom orange pin-points




Figure 7: A normal scapholunate joint. Bones are color mapped and contoured. The saturation of red (darker region in black and white) on bone surfaces represents the distance to the nearest point on the opposite bone. Redder regions are closer. The maximum distance visualized is 5 mm. Contour lines are drawn at 0.5 mm intervals. (a) Bones in their correct anatomical context. (b) Bones rotated to show articulated surfaces more clearly.

shown in Fig. 6 are an example of inaccurate correspondence.

For each pair of bone instances (approx. 250 vertices each), sets of pairwise pin-points were generated in under two minutes, a ten-fold speedup compared to the human expert performance. We note however that the focus of our work is reduction of user interaction, and not minimizing running time.

Manifold deformation. Two separate pairs of bone instances (one hamate pair collected from two different individuals, and one lunate pair collected from the same individual, left and right arm) were cross-parameterized using the markerless correspondence followed by manifold surface deformation. The resulted parameterizations were successfully verified for no self-intersections and folds.

For the two surface pairs described above we also compared the curvature of the resulted dense surface correspondence. The similarity-based shape comparison uses the correspondence between the objects being compared, and computes the similarity between corresponding points. The correspondence can be computed from pairs of pin points on the objects. The similarity is measured as the difference between a shape representation called the Curvature Map [28]. The curvature map is based on the surface curvature over a region around a point, and is represented as a 1-D function of the distance to the point. The radius of the region around the point determines the maximum distance over which the curvature

map can be compared, however, more localized comparisons can be generated by taking the difference over a subset of this maximum distance.

The similarity values are associated with a color map in order to plot similarity on the surface of the object (Fig 8). By controlling the range of values associated with the color map, different ranges of values can be emphasized to indicate regions of similarity or the most significant differences between corresponding locations on the shapes. Because the two datasets featured in this experiment were collected from the same volunteer (left and right arm), we expect curvature dissimilarities due to anatomical shape variation to be minimal. Figure 8 shows that the manifold surface maps regions on S to regions on T of similar curvature, a good indication that our method attains anatomical feature correspondence. The small areas of dissimilar curvature appear to correlate with actual anatomical shape variation between the S and T datasets. Manifold models were created and deformed in under 30 minutes per dataset.

Exploratory visualization and analysis. Our visualization and analysis technique was recently applied in a clinical study to 18 forearm datasets [29], normal and matching injured forearm data collected from 9 individuals. In this study, note that the common frame of reference was defined manually. The approach provided unexpected insight into the biomechanics of the forearm: although an earlier numerical joint-angle study had concluded kinematics were not altered significantly by injury, visual exploration of the very first individual highlighted focal changes in the articulation at the distal radioulnar joint and potential soft-tissue constraints. Further numerical analysis of the location and size of inter-bone joint spaces across the 9 individuals showed the focal changes were statistically significant.

In this study, distance fields and inter-bone joint-spaces were pre-computed in under 1 minute per bone. Joint visualizations are displayed on the fly.

4.2 Applications

We demonstrate our framework on two applications: bone surface-detail transfer, and cross-dataset kinematic analysis.

Surface detail transfer. In the first application, our goal is to combine information collected from different individuals. Limitations in current imaging technology enable collection of wrist soft-tissue data only for cadavers. On the other hand, wrist motion is measured in live human subjects. Combining the two types of information – soft tissue and motion – is important when studying wrist biomechanics.

We use our framework to transfer soft-tissue detail like ligament insertion sites and articular cartilage from one individual to another. In Fig. 9, the blue areas represent synthetically-defined ligament insertion sites. The insertion site originally defined on the pink hamate bone instance is automatically transferred to the white hamate bone instance. In Fig. 10, the cartilage originally defined on the left-side bone instance is automatically transferred to the right-side bone instance (lunate bones shown in pink, cartilage in tan). In both examples, the mutually-consistent parameterization between the two bone instances results in no folds or tears during transfer.

Our Bioengineering Department collaborators are quite interested in using this technique to combine μ CT-imaging cadaver soft-tissue data with in vivo bone surfaces and kinematic data.

Exploratory kinematic analysis. In our second application we explore the differences between normal and injured (*scaphoid non-union*) kinematics in the wrists of a human subject. Our data was collected from the injured and uninjured wrists of the same individual. Figure 11 shows a subset of three wrist bones in their correct anatomical context; note the fractured bone in the right image. In such cases, hand surgeons usually reconstruct the fractured bone by inserting a screw through the two fragments. Alternatively, they



Figure 9: Surface detail transfer (ligament insertion site) between the two hamate bone instances shown in Fig. 5. The mutually-consistent parameterization between the two bone instances results in no folds or tears during transfer.




Figure 10: Cartilage transfer between two lunate bones (bones shown in pink, cartilage in tan). The consistent parameterization between the two lunates results in no folds and no tears in the cartilage map during transfer.

may not intervene at all, or they may remove completely the bottom fragment from the joint [30]. For a given individual, we wish to understand which approach is the most appropriate.

We trace the inter-bone joint-spacing area with motion (10 kinematic poses) and compare location and size using the mutually-consistent parameterization (Fig. 12). Visual analysis of the lunate bone with respect to the scapho-lunate articulation shows no significant differences between the uninjured and injured joint (Fig. 13 first two columns). This suggests that, for this individual, despite injury, the injured scaphoid and lunate bones continue articulating correctly. Under these circumstances, reconstruction of the scaphoid bone may appear unnecessary.

However, further analysis of the lunate bone with respect to the radio-lunate articulation highlights differences between the uninjured and injured joint (Fig. 13 last two columns). The lunate-radius distance appears to increase in the injured wrist: note that the innermost iso-contour in the top images disappears in the bottom images. Also note the distal (upwards) shift of concentric iso-contours in the injured wrist compared to the uninjured wrist. Numerical analysis using the common cross-parameterization on the location of the centroid of concentric iso-contours indicates a distal shift of more than 2 mm. These surprising differences indicate that scaphoid injuries may not impact the articulation nearest to the scaphoid (scapho-lunate), but the next articulation (radio-lunate).

Of course, one can not draw sweeping conclusions from the exploratory analysis of a single individual. Our orthopedist collaborators are interested in running this type of analysis on several individuals with the same type of injury and collecting statistical evidence. For the time being, they are interested in the exploratory analysis of all the wrist bones of this injured individual.

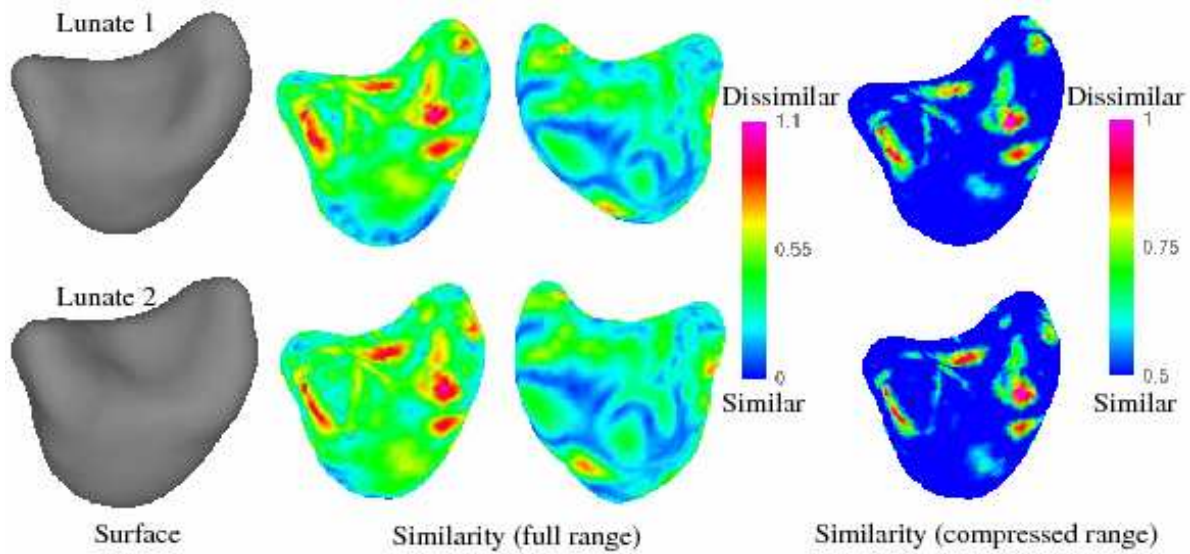


Figure 8: Curvature comparison of source and target manifold surfaces: deforming the manifold surface from Lunate 1 (L1) to Lunate 2 (L2) maps regions on L1 to similar curvature regions on L2. The small areas of dissimilar curvature appear to correlate with actual anatomical shape variation between the L1 and L2 datasets.




Figure 13: Kinematic analysis of the radio-scapho-lunate joint shown in Fig. 11 (top row: lunate bone in the uninjured joint, bottom: lunate bone in the injured joint). The first two columns show the lunate facet articulating with the scaphoid bone; tracing the inter-bone joint-spacing area on the lunate shows similar kinematics between the injured and uninjured scapho-lunate. The last two columns show the lunate facet articulating with the radius bone. This time, note that the innermost contour in the top images disappears in the bottom images; the centroid of the concentric iso-contours also shifts upwards in the bottom images. This indicates that scaphoid injuries may not impact the articulation nearest to the scaphoid, but the next articulation.

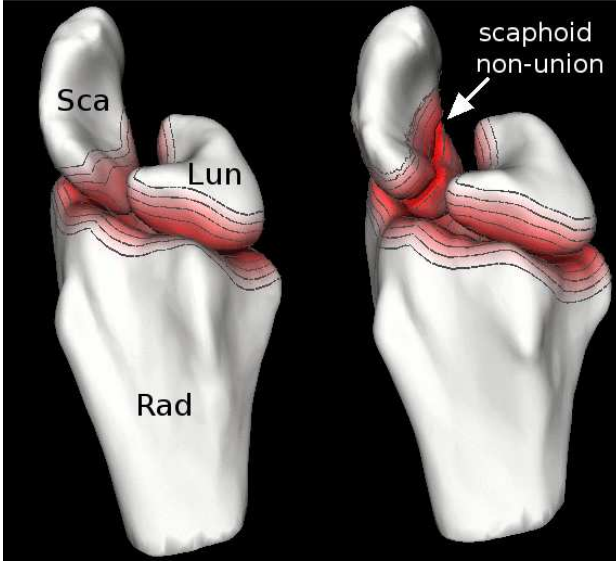


Figure 11: Right-wrist and left-wrist radio-scapho-lunate joints from the same individual. Following injury, the left-wrist scaphoid bone has been fractured in two. Bones are color-mapped and contoured. The color saturation on bone surfaces indicates the distance to the nearest point on the opposite bone or bone-fragment; darker regions are closer.




Figure 12: Pin-point and manifold deformation between the left and right lunate bones in Fig. 11. Left-wrist lunate (injured joint) with pin-points (left), right-wrist lunate with corresponding pin-points (middle), and the original right-wrist lunate overlaid with the result of deforming a manifold from the left-wrist lunate to the right-wrist lunate (right); in the right image, the pin-points on the two surfaces are overlapped. The manifold defines a mutually-consistent parameterization between the S surface and the T surface.

5 DISCUSSION

Our dataset-correspondence method has certain limitations. First, the geometry of the datasets needs to be fairly similar. Second, we need to be able to resample the geometries into similar distributions of points. These requirements are satisfied in the case of arthro-dial-joint bones.

When establishing pairwise pin-point correspondences, we use a correlation coefficient measure. The correlation coefficient is a standard, robust image similarity measure. Although more sophisticated measures have been proposed for special cases where imaging scale, rotation, and perspective distortions are present, it is not clear that the use of such alternative measures would be beneficial in our case.

When further filtering pin-point correspondences, we use a greedy programming approach. Such an approach is not guaranteed to generate an globally optimal set of pin-points, in terms of correlation across the set. The selected set of correspondences could be improved by using a global optimization approach instead of greedy programming. We note that summing the K_{match} values of all the surviving pin-point corresponding for a given S-T pair yields an implicit shape-similarity score between S and T. In our experiments, the higher this score, the more visually similar the shapes of S and T appear to be.

When fitting manifold surfaces we rely on projection to determine the correspondence between the data points and the domain. Problems with folding can arise if projection gives the incorrect correspondence, for example, if the initial surface is poorly aligned with the data. We greatly reduce these problems by using an initial surface which is roughly the right shape and slightly bigger than the data set, and by employing an extra set of smoothing constraints to the control shape when fitting. These smoothing constraints are gradually relaxed as the fit is finalized. The addition of pin points also greatly reduces the chance of folding because the pins pull the surface to the correct area without relying on projection.

6 CONCLUSION

We presented in this paper a framework for the cross-dataset visual exploration and analysis of arthro-dial joint biomechanics. Central to our approach is a markerless method for establishing pairwise correspondences between individual-specific datasets. The resulted correspondence allows the users to combine complementary geometry acquired through different modalities, and thus overcome current imaging limitations. The pairwise correspondence also enables the analysis of kinematic similarities and differences between datasets.

The approach presented is fully automated and works on objects of spherical topology which lack salient features. Unlike previous approaches to pairwise correspondence, the method is not susceptible to object rotational symmetry in the first order ellipsoid. In consequence, our method is applicable to arthro-dial joints like the human wrist or the ankle. Results demonstrate the technique useful in the study of normal and injured anatomy and kinematics of arthro-dial joints.

In principle, the method applies to all spherical topology data from the same class, and should be particularly beneficial in instances where identifying salient object features is a nontrivial task.

Understanding and quantifying differences across groups of human subjects is important in the study of injury mechanisms and prevention, as well as for the design of orthopedic implants. Although many statistical methods for analyzing 3D shape variation exist, they generally require good pairwise correspondence between the different input shape samples. In the case of arthro-dial-joint data, the challenge resides in establishing correspondence between inter-individual bone instances where corresponding features are

difficult to identify without a high level of expertise, due to the smooth, bean-like nature of the arthroidal geometry. Our markerless geometry correspondence method addresses successfully this challenge. The method should be useful as a preprocessing step in the statistical shape analysis of wrist and ankle data.

ACKNOWLEDGMENTS

We wish to thank S. Andrews, C. Demiralp, T. Gatzke, J. Crisco, Dr. Akelman, Jane Casey, Ted Trafton, and J. Hughes for extremely useful conversations and feedback. We also thank NSF CCR-0093238 and 049856 and NIH AR44005 for their support.

REFERENCES

- [1] Besl, P.J. and McKay, N.D., A Method for Registration of 3-D Shapes, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 14(2), pp. 239–256, 1992.
- [2] Thirion, J.P., New Feature Points based on Geometric Invariants for 3D Image Registration, INRIA TR 1901, 1993.
- [3] Wang, Y., Peterson, B.S., Staib, L.H., Shape-based 3D surface correspondence using geodesics and local geometry, *Computer Vision and Pattern Recognition*, Vol 2, pp. 644–651, 2000.
- [4] Styner, M., Gerig, G., Joshi, S.C., Pizer, S.M., Automatic and Robust Computation of 3D Medial Models Incorporating Object Variability, *International Journal of Computer Vision* 55(2-3), pp. 107–122, 2003.
- [5] Cootes, T., Taylor, C., Cooper, D., Graham, J., Active shape models - their training and application. *Comp. Vis. Image Under.* 61, pp. 38–59, 1995.
- [6] Brett, A., Taylor, C., Construction of 3d shape models of femoral articular cartilage using harmonic maps, *Proc MICCAI 2000*, pp. 1205–1214, 2000.
- [7] Golland, P., Grimson, W.E.L., Shenton, M.E. and Kikinis, R., Small Sample Size Learning for Shape Analysis of Anatomical Structures, *Proc. MICCAI 2000*, pp. 72–82, 2000.
- [8] Brechbuhler, C., Gerig, G., Kubler, O., Parametrization of closed surfaces for 3-D shape description, *Comp. Vision, Graphics, and Image Proc.* 61, pp. 154–170, 1995.
- [9] Styner, M., Rajamani, K., Nolte, L., Zsemlye, G., Szekely, G., Taylor, J., Davies, R., Evaluation of 3d correspondence methods for model building, *Information Processing in Medical Imaging (IPMI)*, pp. 63–75, 2003.
- [10] Davies, R., Twining, C., Cootes, T., Waterton, J., Taylor, C., A minimum description length approach to statistical shape modeling, *IEEE Trans. Med. Imaging* 21, pp. 525–537, 2002.
- [11] Kotcheff, A.C.W., Taylor, C.J., Automatic Construction of Eigenshape Models by Direct Optimization, *Med. Image Analysis* 24, 303–314, 1998.
- [12] Belongie, S., Malik, J., and Puzicha, J., Shape Matching and Object Recognition Using Shape Contexts, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 24(4), pp. 509–522, 2002.
- [13] Johnson, A.E., and Hebert, M., Using Spin Images for Efficient Object Recognition in Cluttered 3D Scenes, *IEEE PAMI* 21(5), pp. 433–449, 1999.
- [14] Marai, G.E., Laidlaw, D.H., Markerless inter-subject bone shape matching using 2D projections, *MICCAI 2005 short papers*, 2005.
- [15] Praun, E., Sweldens, W. and Schröder, P., Consistent Mesh Parameterizations, *Proceedings of ACM SIGGRAPH 2001*, pp. 179–184, 2001.
- [16] Kraevoy, V., and Sheffer, A., Cross-Parameterization and Compatible Remeshing of 3D Models, *ACM Trans. Graph.* 23(3), pp. 861–869, 2004.
- [17] Schreiner, J., Prakash, A., Praun, E. and Hoppe, H., Inter-Surface Mapping, *ACM Trans. Graph.* 23(3), pp. 870–877, 2004.
- [18] Crisco, J.J., McGovern, R.D., and Wolfe, S.W., A non-invasive technique for measuring in vivo three-dimensional carpal bone kinematics. *J. Orthopaedic Research*, 17(1), pp. 96–100, 1999.
- [19] Snel, J.G., Venema, H.K., Moojen, T.M., Ritt, M., Grimbergen, C.A. and den Heeten, G.J., Quantitative in vivo Analysis of the Kinematics of Carpal Bones from 3D CT Images Using a Deformable Surface Model and a 3D Matching Technique, *Med. Phys.*, 27(9), pp. 2037–2047, 2000.
- [20] Martin, K., Schroeder, W., Lorensen, B., *The Visualization Toolkit: An Object-Oriented Approach to 3D Graphics*, Prentice Hall, 1997.
- [21] Marai, G.E., Laidlaw, D.H. and Crisco, J.J., Super-Resolution Registration Using Tissue-Classified Distance Fields. *IEEE Trans. on Medical Imaging* 25(2), pp. 177–187, 2006.
- [22] Geomagic Inc., Research Triangle Park, NC, 27709 USA, <http://www.geomagic.com/>
- [23] Mauch, S., A Fast Algorithm for Computing the Closest Point and Distance Transform, <http://www.acm.caltech.edu/seanm/software/cpt/cpt.pdf>.
- [24] Anwar, U., Accuracy of the closest point transform on a 2 sphere phantom CT dataset, Technical Report, Brown University, 2004.
- [25] Grimm, C.M., Crisco, J.J., and Laidlaw, D.H., Fitting locally parametric surfaces to 3D point clouds, *ASME Journal of Biomechanical Engineering* 124(1), pp. 136–140, 2002.
- [26] Grimm, C., Hughes, J., Modeling Surfaces of Arbitrary Topology using Manifolds, *Computer Graphics*, 29(2), July 1995, *Proceedings of SIGGRAPH'95*.
- [27] Marai, G.E., Laidlaw, D.H., Demiralp, C., Andrews, S., Grimm, C.M. and Crisco, J.J., Estimating Joint Contact Areas and Ligament Lengths from Bone Kinematics and Surfaces, *IEEE Transactions on Biomedical Engineering*, 51(5), pp. 790–799, 2003
- [28] Gatzke, T., Zelinka, S., Grimm, C., and Garland, M., Curvature maps for local shape comparison, *In Shape Modeling International*, pp. 244–256, June 2005.
- [29] Crisco, J.J., Moore, D., Marai, G.E., Laidlaw, D.H., Akelman, E., Weiss, A.C., Wolfe, S.W., Effects of Distal Radius Malunion on Distal Radioulnar Joint Mechanics – An In-Vivo Study, *Journal of Orthopedic Research* (in press).
- [30] Akelman, E., personal communication to J. Crisco, 2006.