05/09/2014

An Electric Model for Assessment of Joint Space Morphology in Cone-Beam CT

Computer Integrated Surgery II, Spring 2014

Project Summary

Qian Cao

Mentor: Prof. Jeffrey Siewerdsen

Abstract

High-resolution cone-beam CT (CBCT) of the extremities presents a potentially valuable basis for image-based biomarkers of arthritis, trauma, and risk of injury. We present a new method for 3D joint space analysis that exploits the high isotropic spatial resolution of CBCT and is sensitive to small changes in disease-related morphology. The algorithm yields accurate quantitation of the joint in a manner that is robust against operator and patient setup variation and is a possible image-based biomarker of subtle articular change.

I. Introduction

Normal anatomical variations in knee morphology can function as potent predictors of risk-of-injury when the body is under stress, such as in weight-bearing conditions [1-6]. With the recent development of a dedicated extremity cone-beam computed tomography (CBCT) scanner at Hopkins [7], CT volumes can now be acquired when the subject is standing in a load-bearing state. This combination of high image quality and new configuration provides an opportunity to investigate the relationship between risk-of-injury under load-bearing conditions and features of knee morphology. However, conventional methods of analyzing knee morphology is lacking in consistency and accuracy. In this project, a novel method of characterizing the joint space in the knee will be developed and validated.

Knee morphology has long been used to reveal severity of pathologies such as rheumatoid arthritis, gout and a variety of other diseases. Various methods of assessing the morphology co-exist, each popular within their own modalities and tailored to specific pathologies. In radiographs, a long standing (and evolving) tradition is the Sharp-Larsen Score [8]; in MRI, the gold standard for diagnosing rheumatoid arthritis is the rheumatoid arthritis MRI scoring (RAMRIS) system [9]. Though these methods measure common features such as joint space narrowing and cartilage thinning, they are highly qualitative and/or difficult to reproduce consistently. For example, the Sharp-Larsen score is extremely difficult to standardize, as beam direction, knee position, foot rotation are all variables affecting the final result. Along with RAMRIS, the Sharp-Larsen method also utilizes a simple scoring system based on qualitative descriptions to classify the patient's condition into only a few categories.

In recent years, studies have proposed quantitative methods for assessing the joint space in CT volumes of the knee [10], including projection along a longitudinal axis and the closest point algorithm. In the former method, lines are projected from one surface to the other along a predefined axis (Figure 1), the lengths of these lines are then measured and defined to be the distance between the two surfaces. This method, though straightforward, is difficult to implement in a practical situation, as it is somewhat arbitrary what the predefined axis should be, especially taking into account the complex geometry of the femur and tibia. Moreover, the configuration of the patient (whether

she is standing or sitting) also influences the selection of the predefined axis. Selecting the "right" axis is not just hard but also has significant impact on the result, as shown in figure 1 c) and d), a slight deviation of the axis can produce a drastically different distance map between two surfaces.

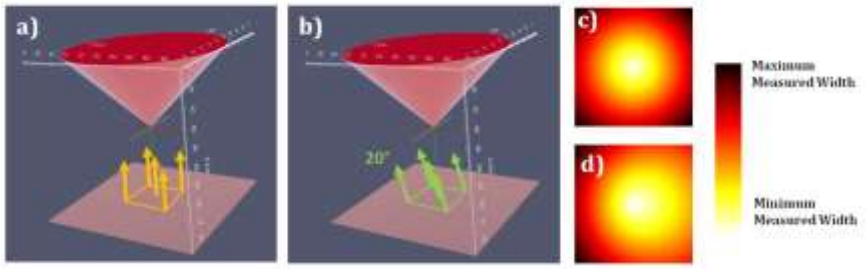

Figure 1 A simple example of measurements with the projection along a longitudinal axis method. In a) the predefined axis is parallel to the z-axis, producing a mapping c) of measured distance between the plane and cone. In b) the predefined axis is tilted by 20 degrees, producing a mapping d).

The latter method evades some of the aforementioned problems. The closest point algorithm calculates for each point on a surface, the distance to its closest point on the other surface. This method is slightly more sophisticated, but not problem-free. For example, figure 2 shows the same cone-plane geometry as depicted in figure 1, when mapping from the cone to the plane (figure 2b), each point on the surface of the cone maps vertically down onto the plane, producing the same distance map as did figure 1a. However, when mapping from the plane to cone (figure 2a), the tip of the cone is the closest point to every point in the region of interest on the plane. Thus the mappings are completely different. This could be a problem when handling sharp protrusions on either surface. As with the case in figure 2a, the full 3D volume information (the majority of the surface area on the cone) is not taken into account.

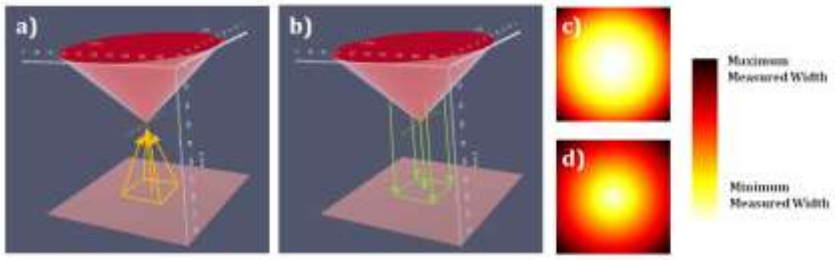

Figure 2 A simple example of measurements with the closest point method. In a) the mapping is constructed from the square region on the plane to the cone. In b) the mapping is constructed from the surface of the cone to the plane.

The two methods described above are based on rather artificially defined distance functions. In this project we derive a more elegant solution (originally proposed in [7]) from the laws of physics, in the form of Maxwell's equations. If we consider the two surfaces to be the terminals of a capacitor differing in electric potential, a gradient of electric potential will exist between them. By tracing the gradient, we obtain the electric field lines from one surface to the other (figure 3). These lines are perfectly bijective, independent of any arbitrary selection of axes, continuously distributed across both surfaces and unique to the geometry of the space.

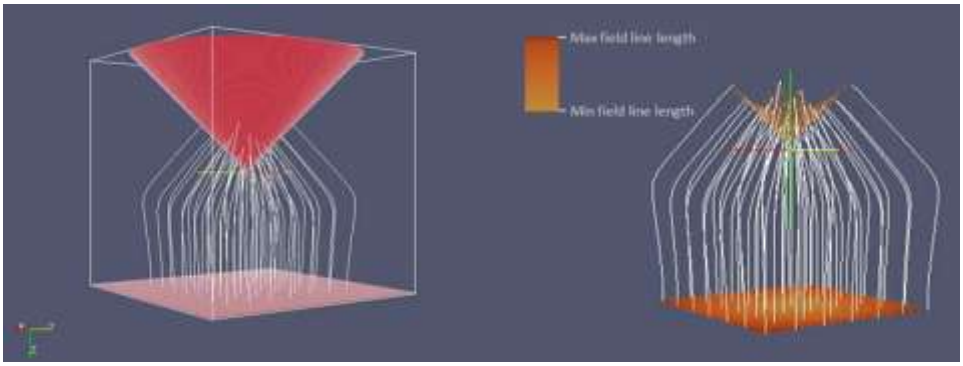

Figure 3 The distance between the cone and plane mapped with electric field lines derived from the capacitor model. In the right subfigure, the region of interest is colormapped so that the smaller distances correspond to lighter color. Figures were created as part of the preliminary findings described in the next section.

The specific aims of the project are to:

1. Implement an algorithm that generates a joint space map (JSM) based on the capacitor model.
2. Validate and test the algorithm on real knee volume data.
3. Provide thorough documentation and analysis of the algorithm for future biomechanical studies.

II. Technical Approach

Theory

Measurement of joint space widths (JSW) is the problem of properly defining and resolving trajectories from one articular surface to the other, with the lengths of each trajectory a measure of JSW at the trajectory origin. As mentioned in the previous section, conventionally, these trajectories were defined as projections along a longitudinal axis of the joint or as beelines directed at the closest point on the opposing surface. In the former method, it is somewhat arbitrary what this longitudinal axis should be. In the latter method, the measurement is often asymmetric and degenerate. We implement a novel method in which the two articular surfaces of the knee joint are conceptualized as surfaces of differential charge densities, resulting in a capacitor. The trajectories are defined as electric field lines from one surface to the other, guided by the electrostatic potential in between. In order to determine the distribution of electrostatic potential $\Phi(\mathrm{x}, \mathrm{y}, \mathrm{z})$, we substitute the definition of electric field:

$$\vec{E} = -\nabla\Phi$$

into Gauss's law in differential form:

$$\nabla \cdot \vec{E} = \frac{\rho}{\epsilon_0}$$

where $\rho(x, y, z)$ is the charge density at a particular point in space and ϵ_0 the electric permittivity of free space. Taking ρ to be zero in the free space between the charges, we arrive at the Laplace's equation:

$$\nabla^2\Phi = 0$$

Literature on the numerical solution to Laplace's equation is extensive. Here we only implement the simplest method. First, we attempt to express this equation in a discrete domain. For a given voxel size h, the Taylor expansion for the forwards and backwards voxel in a given direction (e.g. x) is

$$\Phi(a+h,b,c) = \Phi(a,b,c) + \left.\frac{\partial\Phi}{\partial x}\right|_{(a,b,c)} h + \left.\frac{\partial^2\Phi}{\partial x^2}\right|_{(a,b,c)} \frac{h^2}{2!} + \left.\frac{\partial^3\Phi}{\partial x^3}\right|_{(a,b,c)} \frac{h^3}{3!} + O(h^3)$$

$$\Phi(a-h,b,c) = \Phi(a,b,c) - \left.\frac{\partial\Phi}{\partial x}\right|_{(a,b,c)} h + \left.\frac{\partial^2\Phi}{\partial x^2}\right|_{(a,b,c)} \frac{h^2}{2!} - \left.\frac{\partial^3\Phi}{\partial x^3}\right|_{(a,b,c)} \frac{h^3}{3!} + O(h^3)$$

Summing the two equations above and rearranging the terms, we have the finite difference approximation to the second partial derivative:

$$\left.\frac{\partial^2\Phi}{\partial x^2}\right|_{(a,b,c)} = \frac{\Phi(\mathrm{a}+\mathrm{h},\mathrm{b},\mathrm{c}) - 2\Phi(a,b,c) + \Phi(\mathrm{a}-\mathrm{h},\mathrm{b},\mathrm{c})}{h^2} + \mathrm{O}(\mathrm{h}^3)$$

Substituting this for partials in all directions into the Laplace's equation leads to:

$$\Phi(a,b,c) \approx \frac{1}{6}[\Phi(\mathrm{a}+\mathrm{h},\mathrm{b},\mathrm{c}) + \Phi(\mathrm{a}-\mathrm{h},\mathrm{b},\mathrm{c}) + \Phi(\mathrm{a},\mathrm{b}+\mathrm{h},\mathrm{c}) + \Phi(\mathrm{a},\mathrm{b}-\mathrm{h},\mathrm{c}) + \Phi(\mathrm{a},\mathrm{b},\mathrm{c}+\mathrm{h}) + \Phi(\mathrm{a},\mathrm{b},\mathrm{c}-\mathrm{h})]$$

Thus, Laplace's equation demands that any voxel be the mean of its nearest neighbors. A simple way to satisfy this is to enforce the condition iteratively on all voxels, resetting the boundary conditions (charge densities at the location of articular surfaces) after each iteration. This is equivalent to the Jacobi method. Once $\Phi(\mathrm{x},\mathrm{y},\mathrm{z})$ is determined, the field lines could be found by stepping through the free space, in the interpolated direction of $\nabla\Phi$ or $-\nabla\Phi$.

Implementation

The algorithm is implemented in MATLAB (R2013b), taking advantage of built-in functions and GPU computing capabilities. An overview of the pipeline can be found in figure 1. The initial volumes, in DICOM format with voxel sizes of 0.2169 cm^3, are imported into MATLAB and cropped manually into 256x256x128 subvolumes centered between the medial and lateral tibial plateau.

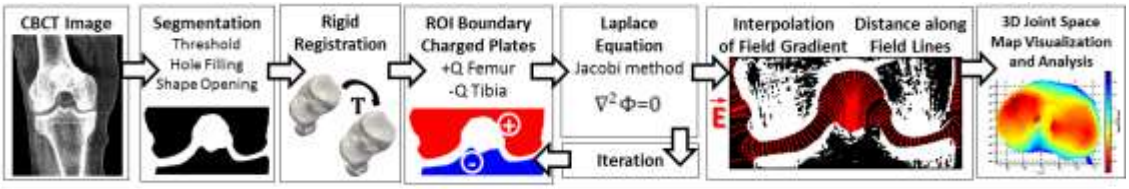

Figure 4. Overview of the joint space analysis pipeline.

The current segmentation method consists of a series of morphological operations designed to separate bones from its surrounding soft tissue. The volumes are first masked to eliminate peripheral artifacts. Then, a grayscale morphological closing with a 6-nearest neighbor structuring element is applied to strengthen the cortical bone. Next, the region with the soft tissue is obtained by slice-wise application of grayscale morphological opening to the resulting image with a disk-shaped structuring element that's half the size of total soft tissue area. This region is then subtracted from the morphologically closed volume, leaving only the bony structures. Next, a hard threshold is applied to segment the bones. A binary morphological operation is applied with the minimum structuring element that's able to fill the holes in cortical bone (determined via trial and error). A fill-hole operation is applied to both the femur and tibia to fill the trabecular structures. Finally, a connected component analysis is carried out and the two largest components are chosen to be the femur and tibia, which are assigned values of +1 and -1, respectively. The rest of the volume is set to zero. In some knee volumes, the distance between the femur and tibia is extremely small, an initial closing operator would connect the two surfaces. In such cases, the initial grayscale closing is not performed.

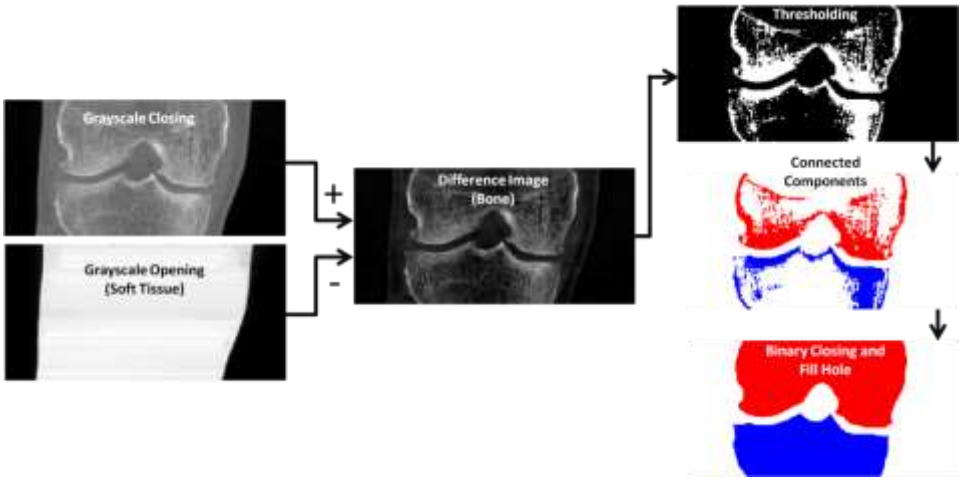

Figure 5. Schematic of important steps in the segmentation pipeline.

In calculating $\Phi(x, y, z)$, iterations are implemented as convolutions of a 7 point stencil, which is zero at the center and $\frac{1}{6}$ at its 6 nearest neighbors, with the segmented and labeled knee volume. After each iteration, the boundary conditions are reinforced by setting the corresponding voxels to +1 or -1. At the volume boundary, the convolution is computed by padding of the nearest edge voxel. Due to the parallel nature of this step, it can be efficiently implemented on the GPU using MATLAB's gpuArray. On our nVidia GTX 470 graphics card, a speed of 17 iterations per second is achieved. Because the region we are interested in measuring is populated by field lines less than 10 cm, complete convergence of every voxel in the volume is unnecessary, instead, 500 iterations were found empirically to yield a converged result. After Φ is calculated, the edge voxels of the tibia is found by subtracting the erosion of the binary image from the original. Electric field lines emerge from these edge voxels and are stepped in the gradient direction of Φ in increments of approximately $\frac{1}{10}$voxel size until it hits either the volume boundary or the femur. Lines that hit the volume boundaries are not saved. Each edge voxel on the tibia is then assigned a numerical value of the length of the emerging electric field line. The final map consists of a scatter plot of these edge points colormapped according to their values.

For the study of joint space narrowing in load-bearing and non-load bearing conditions, it is necessary to register (via intensity) a pair of images for comparison. This is performed immediately after segmentation and labeling before field calculations. This ordering results in a clean map devoid of partial volume effects from interpolation. In

our implementation, the standing tibia is registered to the sitting tibia. The resulting transform is then applied to the entire volume of the standing scan. This aligns both of the tibias and allows differencing between the two maps found by the electrostatic model.

GPU Acceleration

The averaging procedure used in the method of relaxation can be formulated as a multidimensional convolution with a special averaging mask. This is significant because convolution is very easy to parallelize with existing libraries on a GPU. Preliminary benchmarks have revealed that the computation of the field potential can be accelerated up to 50 fold in the size regime of the existing test CT volumes.

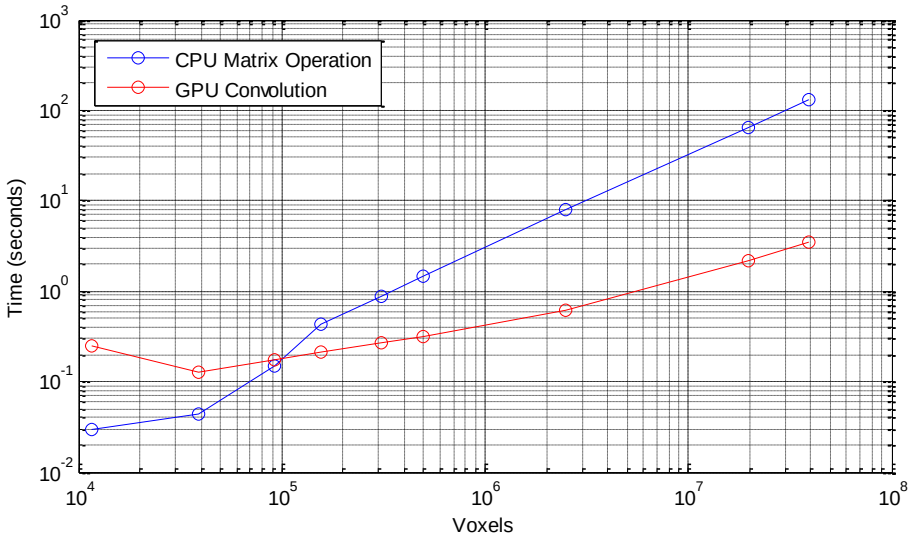

Figure 6. Benchmark comparing the time to complete 10 iterations on the same volume for GPU and CPU implementation

III. Results and Discussion

Axial-displacement Phantom

In order to evaluate the behavior of our algorithm relative to the commonly used closest point distance metric, a phantom was fabricated with a wedge on one fixed end and a

hollowed cone on the other. The hollowed cone is attached to a caliper and could be adjusted towards or away from the wedge. In order for the electrostatic model to return useful information about joint space, a necessary condition is that as the cone is displaced from the wedge, the resulting electric field lines are well-behaved and shift accordingly. The algorithm is applied to the cone-wedge volume. The maps generated are shown in figure 3. The cone was initially brought to contact with the wedge, then displaced about 0.5 mm away. Next, ten scans were acquired, each with the cone displaced in increments of about 0.4 mm.

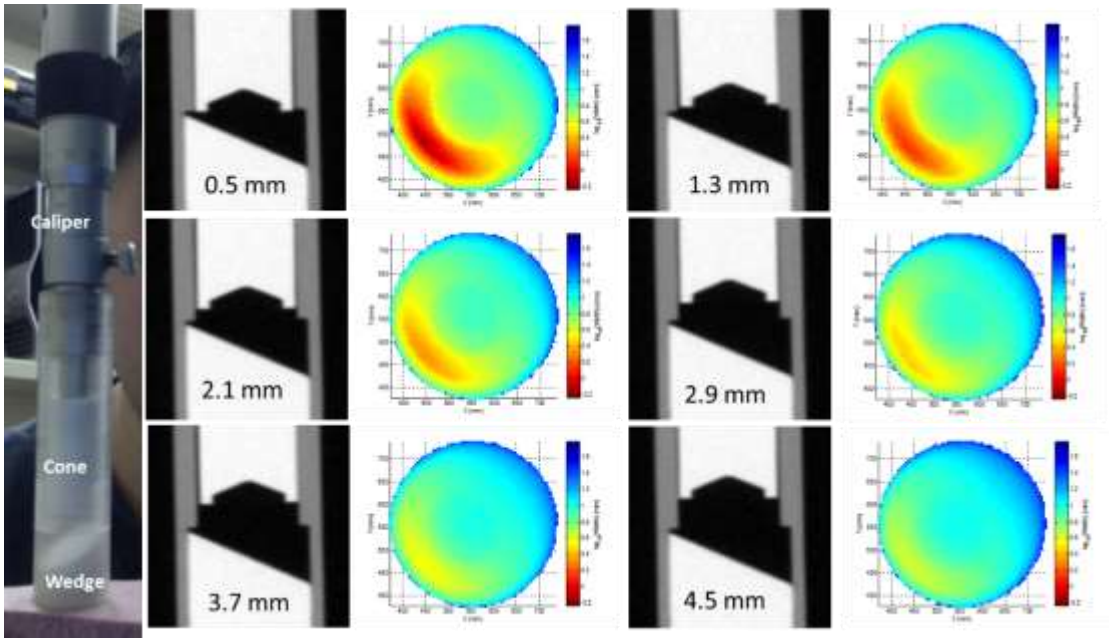

Figure 7. Distance maps on the wedge, as the cone is moving up.

In the electrostatic model, the analogy to the minimum distance between the cone and the wedge is the electric field line with the minimum length. For this specific geometry and displacement along the axial direction, the minimum field line length is roughly linear to the absolute minimum distance. Some offset exists due to the curvature of the field lines.

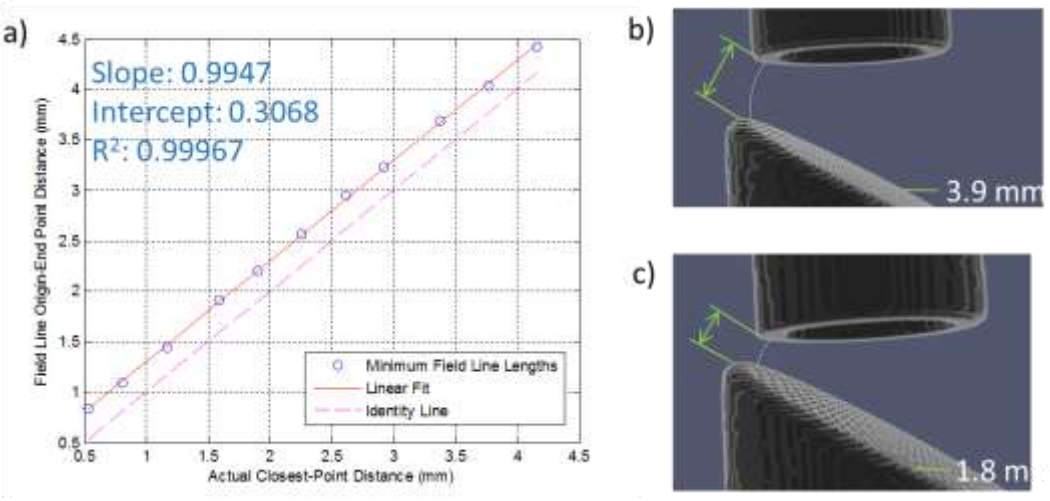

Figure 8. a) Linear fit of the minimum field line lengths to actual closest point distance between the two surfaces. b) and c) shows the closest point distance (in green) and the minimum field line (in white) at two positions of vertical displacement.

While this is not sufficient to show that the electrostatic model produces accurate measurement of axial displacement, the results do reveal some correlation with the conventional intuition of minimum distance.

OA and non-OA Knees

Currently, 31 pairs of knee scans have been made: 18 OA knees and 13 non-OA knees, in sitting and standing positions. Four of these pairs have been mapped by our pipeline and is shown in figure 5. These results are not yet enough to identify patterns of joint space narrowing in OA knees. One possibility is that during acquisition of standing and sitting scans, the degree of knee extension is slightly different. In this case, information contained in the difference maps in figure 5 would be dominated by the complex motions involved in extending the knee as documented in biomechanics literature (e.g. screw-hone mechanism, etc). The next step is to attempt to decouple the joint space narrowing from other motions using a 5-DOF rigid registration of the femurs, somehow preserving the axial displacement information. This might be sufficient to compensate for the difficult-to-control inconsistencies involved in data collection.

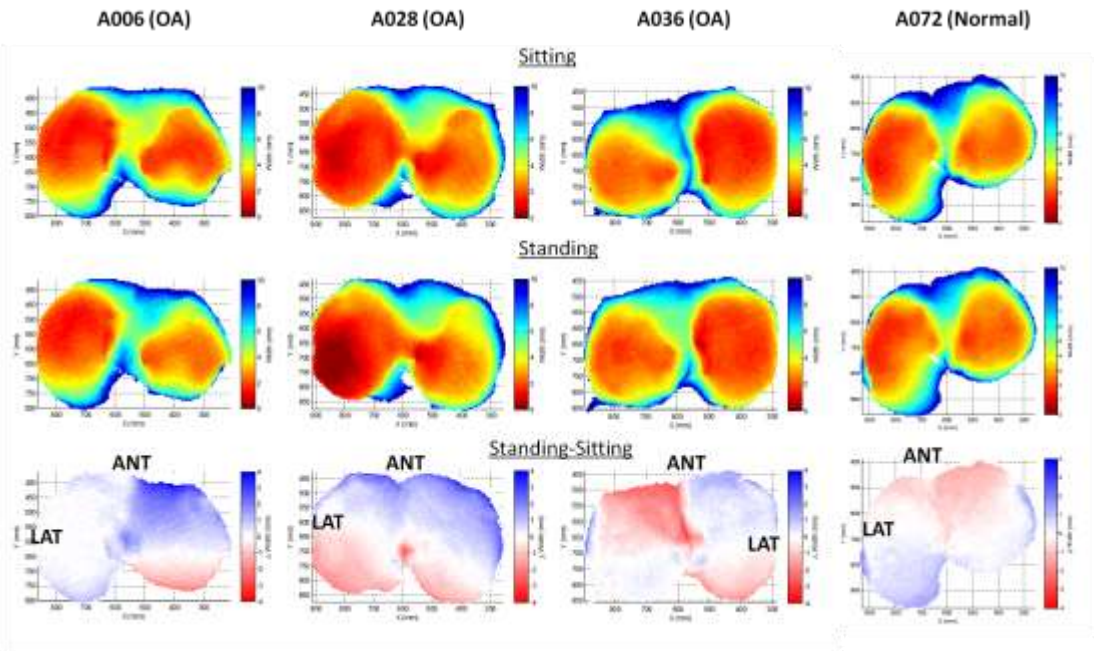

Figure 9. Comparison of 3 OA knees and 1 normal knee in the study. First row are maps calculated from volumes acquired during a sitting configuration; the second row are ones obtained from the standing configuration. The third row is the subtraction sitting from standing.

Extension Angle Correction

The 5DOF correction can be derived from a 6DOF registration between the femurs. Once a 6DOF registration is complete, the position of the moving volume can be shifted back along a defined axial direction to its original position. Implementation-wise, we seek to decompose the 6DOF registration transformation matrix into the axial and the 5DOF component.

$$T_{6DOF}|_{sitting}^{standing} = T_{\mu_{femur}} T_{axial} T_{5DOF} T_{\mu_{femur}}^{-1}$$

Where,

$$T_{\mu_{femur}} = \begin{bmatrix} 1 & 0 & 0 & \mu_x \\ 0 & 1 & 0 & \mu_y \\ 0 & 0 & 1 & \mu_z \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

With the μ_x, μ_y, μ_z denoting the location of the centroid of the moving femur (in this implementation, we are moving the sitting femur to the standing femur). The translation along the axial direction can be formulated as a translation:

$$T_{axial} = \begin{bmatrix} 1 & 0 & 0 & v_x \\ 0 & 1 & 0 & v_y \\ 0 & 0 & 1 & v_z \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

To find the axial displacement components v_x, v_y, v_z, we first solve for the 6DOF transform matrix with the origin translate to the centroid of the moving femur:

$$T_{\mu_{femur}} T_{6DOF} T_{\mu_{femur}}^{-1} = \begin{bmatrix} a & b & c & t_x \\ d & e & f & t_y \\ g & h & i & t_z \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

The rigidity condition is enforced in this matrix, due to the complicated form of the rigid transformation matrix, we abbreviate the entries as $a \ldots \ i. \ t_x, t_y, t_z$ are the translation components. We can then recover the translation along the axial direction with:

$$(v_x, v_y, v_z) = \left[(t_x, t_y, t_z) \cdot \frac{\underline{\mu_{femur}} - \underline{\mu_{tibia}}}{\left| \underline{\mu_{femur}} - \underline{\mu_{tibia}} \right|} \right] \left(\frac{\underline{\mu_{femur}} - \underline{\mu_{tibia}}}{\left| \underline{\mu_{femur}} - \underline{\mu_{tibia}} \right|} \right)$$

Where $\frac{\underline{\mu_{femur}} - \underline{\mu_{tibia}}}{\left| \underline{\mu_{femur}} - \underline{\mu_{tibia}} \right|}$ is the unit vector pointing from the centroid of the tibia to the centroid of the femur. Though we've defined an axial direction, unlike in the projection method, the direction is not arbitrary but is determined by the one-to-one correspondences established by an existing electric field line calculation. Once this is complete, the 5DOF transform can then be recovered:

$$T_{5DOF} = T_{axial}^{-1} T_{\mu_{femur}}^{-1} T_{6DOF} \Big|_{sitting}^{standing} T_{\mu_{femur}}$$

The 6DOF transform is the same as the standard rigid transform, formulated as an optimization problem to minimize the mutual information or correlation. This can rid of the femur/tibia rotation component, preserving the axial displacement information, which is at the heart of clinical diagnosis. There's one hidden caveat to this method. According to the parallel axis theorem in basic dynamics, any rigid body motion can be thought of as the superposition of rotation along some "parallel" axis and the translation of the parallel axis. This means that the axial translation could depend on where we set the origin of the transformation (at the centroid of the femur). It would be a worthwhile endeavor to find and use some other less varying point as the origin.

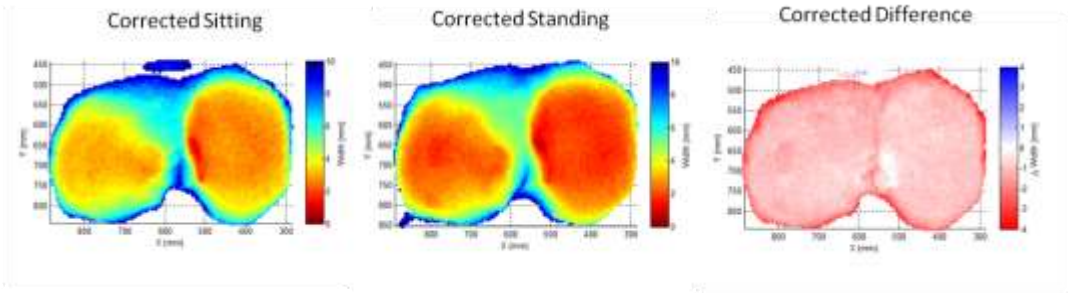

Figure 10. Maps from analysis of subject A036 with geometry correction. Note that the apparent rotation component between the tibia and femur is gone. We are left with a monotonic axial translation between sitting and standing.

IV. Conclusion

Though some work remains to be done, after many trials and error, the main pipeline is implemented and is tested on CT volumes of OA and non-OA knees. Preliminary joint space maps as well as registered difference maps are promising and can be obtained in a reasonable time (~5 min per volume). The main bottleneck is the semi-automatic segmentation component, which requires some manual tuning of parameters in order to obtain a usable labeled volume. Segmentation of the knee is a research topic of its own, with numerous papers describing a host of methods, thus we feel confident that with enough research, the component can readily be swapped with a more robust version in future developments.

Future work also includes processing of all available OA/non-OA knee data and making meaningful clinical interpretations of the resulting maps. This could possibly be done

via an intersubject registration followed by PCA. So far, the current pipeline establishes a good foundation for executing these plans.

Documentation is an important aspect of any algorithm development project. As the code is still evolving with incremental adjustments, most of the documentation exists in-source. These will be transferred and compiled to an official document once the project is near its end for hand-off.

V. Project Management Specifics

Minimum Deliverable (Expected by 03/01/2014)

1. A set of prototyped MATLAB functions for joint space mapping using the capacitor model. ✓
2. A set of prototyped MATLAB functions for segmentation. ✓
3. Documentation of existing code. ✓

Expected Deliverable (Expected by 05/01/2014 → PUSH BACK TO 06/01/2014)

1. A set of validated MATLAB functions for joint space mapping using the capacitor model. ✓
2. A refined MATLAB function for segmentation. ✓
3. Detailed analysis of algorithm performance (convergence characteristics, accuracy, speed etc). ✓
4. MATLAB routines for visualization of the analysis results (volume rendering + GUI) in VTK. ✓
5. Provide relevant documentation for code hand-off. (80% complete)
6. Conduct a phantom study to compare the algorithm with existing closest-point method. ✓
7. Apply the analysis pipeline to the analysis of OA and non-OA knee joints under load-bearing (standing) vs non-load-bearing (sitting) conditions (62 CT Volumes). (? % complete)

Maximum Deliverable (Expected by 07/01/2014)

1. Submit abstract to the American Association of Physicists in Medicine (AAPM) annual meeting. ✓

(Abstract accepted, oral: Q. Cao, G. Thawait, G. J. Gang, W. Zbijewski, T. Riegel, S. Demehri, and J. H. Siewerdsen, "An Electrostatic Model for Assessment of Joint Space Morphology in Cone-Beam CT," The 56th Annual Meeting of the AAPM. July 24th, 2014, Austin TX**)**

2. Submit a technical paper.

Timeline

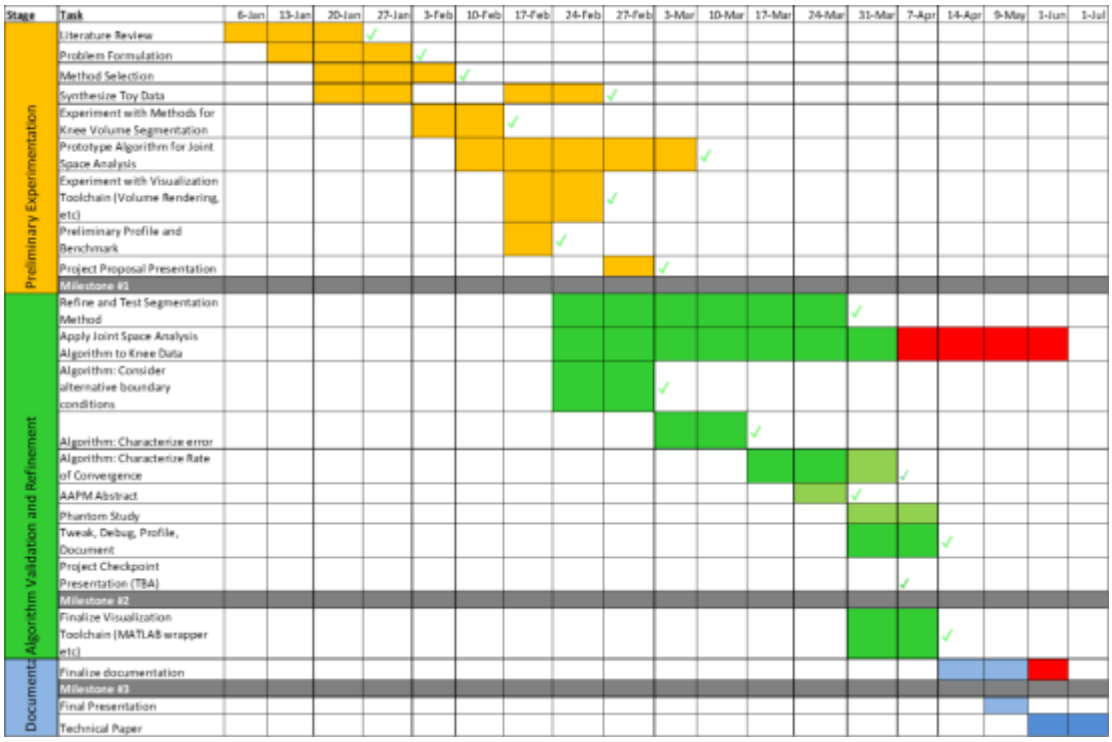

Dependencies (all are met)

1. Bi-weekly meeting with mentor (bi-weekly meeting scheduled with Prof. Siewerdsen).
2. CBCT knee volume test data (Two 61 datasets of OA knees and non-OA knees available for algorithm testing and validation).
3. Phantom study.
    i) Equipment + materials + access to shop (Traylor machine shop).

ii) Access to CBCT scanner (lab bench).

iii) Radiation safety training (radiation badge acquired).

4. Computing resources.
   i) Up-to-date MATLAB w/ image processing and parallel computing toolboxes (R2013b).
   ii) CUDA-enabled graphics card (NVidia GTX470).
   iii) C++ IDE and compiler (Visual Studio 2008).
   iv) Visualization library (VTK).
5. Access to relevant literature (Lab database & JHU Library Website).

VI. References

[1] J. R. Giffin, T. M. Vogrin, T. Zantop, S. Woo, and C. D. Harner, "Effects of Increasing Tibial Slope on the Biomechanics of the Knee," Am. J. Sports Med. 32(2): 376-382 (2004).

[2] J. Hashemi, N. Chandrashekar, B. Gill, B. D. Beynnon, J. R. Slauterbeck, R. C. Schutt Jr., H. Mansouri, and E. Dabezies, "The Geometry of the Tibial Plateau and Its Influence on the Biomechanics of the Tibiofemoral Joint," J. Bone and Joint Surg. 90: 2724-2734 (2008).

[3] J. Hashemi, N. Chandrashekar, H. Mansouri, B. Gill, J. R. Slauterbeck, R. C. Schutt Jr., E. Dabezies, and B. D. Beynnon, "Shallow Medial Tibial Plateau and Steep Medial and Lateral Tibial Slopes : New Risk Factors for Anterior Cruciate Ligament Injuries," Am. J. Sports. Med. 38: 54-52 (2010).

[4] S. G. McLean, S. M. Lucey, S. Rohrer, and C. Brandon, "Knee joint anatomy predicts high-risk in vivo dynamic landing knee biomechanics,"Clin. Biomech. (2010) (in press doi:10.1016/j.clinbiomech.2010.06.002).

[5] K. B. Shelburne,H.-J. Kim, W. I. Sterett, and M. G. Pandy, "Effect of Posterior Tibial Slope on Knee Biomechanics during Functional Activity," Orth. Res. Soc. Feb. 223-231 (2011).

[6] S. J. Shultz and R. J. Schmitz, “Tibial Plateau Geometry Influences Lower Extremity Biomechanics During Landing,” Am. J. Sports Med. 40(9): 2029-2036 (2012).

[7] W. Zbijewski, P. De Jean, P. Prakash, Y. Ding, J. W. Stayman, N. Packard, R. Senn, D. Yang, J. Yorkston, A. Machado, J. A. Carrino, and J. H. Siewerdsen, "A dedicated cone-beam CT system for musculoskeletal extremities imaging: Design, optimization, and initial performance characterization," Med. Phys. 38(8): 4700 - 4713 (2011).

[8] Sokka, T. (2008). Radiographic Scoring in Rheumatoid Arthritis, 66(2), 166–168.

[9] Hodgson, R. J., O’Connor, P., & Moots, R. (2008). MRI of rheumatoid arthritis image quantitation for the assessment of disease activity, progression and response to therapy. Rheumatology (Oxford, England), 47(1), 13–21. doi:10.1093/rheumatology/kem250

[10] Kalinosky, B., Sabol, J. M., Piacsek, K., Heckel, B., & Gilat Schmidt, T. (2011). Quantifying the tibiofemoral joint space using x-ray tomosynthesis. Medical Physics, 38(12), 6672–82. doi:10.1118/1.3662891

[11] Yezzi, A. J. & Prince, J. L. (2003). An Eulerian PDE approach for computing tissue thickness. IEEE Transactions on Medical Imaging, 22(10), 1332–9. doi:10.1109/TMI.2003.817775