

Abstract of “Data-Driven Predictive Modeling of Diarthrodial Joints” by Georgeta-Elisabeta Marai, Ph.D., Brown University, May 2007

This dissertation presents a computational framework for integrating measured data — such as medical images, tracked motion, and anatomy-book knowledge — into the predictive modeling of anatomical joints. The framework is data-driven in the sense that it uses sampled motion data to infer soft-tissue geometry and behavior. The framework allows the generation of adaptable, quantifiable, predictive models and simulations of complex joints, surpassing current measuring limitations.

I instantiate the framework in a collection of tools: 1) a sub-voxel accurate method for tracking bone-motion from sequences of medical images; 2) computational tools for estimating soft-tissue geometry and contact; and 3) a tool for the visual and quantitative exploration of joint biomechanics. The first tool attains accuracy improvements of more than 74% over current tracking methods, when compared to the ground truth computed from marked data; the accuracy improvement enables the analysis of soft-tissue deformation with motion in live individuals. The second tool enables us to overcome current soft-tissue *in vivo* imaging limitations. The third tool facilitates the quantitative and visual analysis of joint models and simulations.

The resulting computational models are somewhat unusual in their hybridization of data representations. Each representation has strengths for various aspects of the modeling and I combine them in unique ways to achieve simple, elegant and accurate estimations of biologically relevant measurements.

I demonstrate the application of this framework to the human wrist and forearm. The results generated through this framework have already impacted orthopedists’ understanding of the many diseases afflicting human joints. With such a better understanding, improvements in treatment for injuries are possible as well as reductions in injuries.

Abstract of “Data-Driven Predictive Modeling of Diarthrodial Joints” by Georgeta-Elisabeta Marai, Ph.D., Brown University, May 2007

This dissertation presents a computational framework for integrating measured data — such as medical images, tracked motion, and anatomy-book knowledge — into the predictive modeling of anatomical joints. The framework is data-driven in the sense that it uses sampled motion data to infer soft-tissue geometry and behavior. The framework allows the generation of adaptable, quantifiable, predictive models and simulations of complex joints, surpassing current measuring limitations.

I instantiate the framework in a collection of tools: 1) a sub-voxel accurate method for tracking bone-motion from sequences of medical images; 2) computational tools for estimating soft-tissue geometry and contact; and 3) a tool for the visual and quantitative exploration of joint biomechanics. The first tool attains accuracy improvements of more than 74% over current tracking methods, when compared to the ground truth computed from marked data; the accuracy improvement enables the analysis of soft-tissue deformation with motion in live individuals. The second tool enables us to overcome current soft-tissue *in vivo* imaging limitations. The third tool facilitates the quantitative and visual analysis of joint models and simulations.

The resulting computational models are somewhat unusual in their hybridization of data representations. Each representation has strengths for various aspects of the modeling and I combine them in unique ways to achieve simple, elegant and accurate estimations of biologically relevant measurements.

I demonstrate the application of this framework to the human wrist and forearm. The results generated through this framework have already impacted orthopedists’ understanding of the many diseases afflicting human joints. With such a better understanding, improvements in treatment for injuries are possible as well as reductions in injuries.

Data-Driven Predictive Modeling of Diarthrodial Joints

by

Georgeta-Elisabeta Marai

B. Sc., Politehnica University of Bucharest, Romania, 1997

M. Sc., Politehnica University of Bucharest, Romania, 1998

Sc. M., Brown University, 2001

A dissertation submitted in partial fulfillment of the
requirements for the Degree of Doctor of Philosophy
in the Department of Computer Science at Brown University

Providence, Rhode Island

May 2007

© Copyright 2007 by Georgeta-Elisabeta Marai

This dissertation by Georgeta-Elisabeta Marai is accepted in its present form by the Department of Computer Science as satisfying the dissertation requirement for the degree of Doctor of Philosophy.

Date _____
David H. Laidlaw, Director

Recommended to the Graduate Council

Date _____
Joseph J. Crisco, Reader

Date _____
John F. Hughes, Reader

Date _____
Nancy S. Pollard, Reader
Carnegie-Mellon University

Approved by the Graduate Council

Date _____
Sheila Bonde
Dean of the Graduate School

Vita

G. Elisabeta Marai was born in 1974 in Bucharest, Romania. In 1985 she played the violin in a chamber orchestra, in 1986 she learned how to code. In 1989 she was thankfully released from communist heaven. In 1997 and 1998 Ms. Marai received a B.S. and an M.S. from the Computer Science Department of the Politehnica University of Bucharest, Romania. She was certified as a high-school teacher and spent another year as a University lab instructor. She won an internship at Philips Research in the Netherlands and then kept walking, on to graduate school at Brown University. Ms. Marai developed an interest in computer graphics and scientific visualization, with a focus on modeling, visualization and automated analysis of medical data. She works closely with researchers at the Brown Medical School and in the Biology and Evolutionary Biology Department. While completing the Ph.D. program at Brown, Ms. Marai has published four journal papers, one book chapter, and numerous conference papers and abstracts.

Acknowledgments

I thank my mentors who saw this dissertation through: my advisor David Laidlaw and co-advisor Joseph Crisco, and my readers John Hughes and Nancy Pollard. Thank you for your support and guidance, and for many inspiring discussions.

I gratefully acknowledge my collaborators Cindy Grimm, Çağatay Demiralp, Stuart Andrews, Douglas Moore, Dr. Edward Akelman, Dr. Arnold-Peter Weiss, Tim Gatzke and Sharon Swartz. This work wouldn't have been possible without the data collected and processed by James Coburn, Sharon Sonnenblum, Anwar Upal, Ted Trafton, Evan Leventhal and Jane Casey, nor without the generous support of NIH, NSF, Pixar and Microsoft.

I wish to thank everyone at Brown and elsewhere who stood by me in times of joy and times of need. Particular thanks go to Anca Ivan, Anne-Marie Bosneag, and Carmen Grigorescu, for picking up the pieces everytime. And thank you, Olga Karpenko, Tomer Moscovich, Peter Sibley, Prabhat, Alex Rasin, Johanna Quattrucci, Song Zhang, Yanif Ahmad, and Aris Anagnostopoulos — no one could wish for more loyal friends.

For saving my wrists and sanity, beyond the call of duty, thank you, Kevin Curtis, Dr. Marsha Miller and Eileen Pappas. I am indebted to Alexis Devine, my kickboxing instructor, for saving many a frustrating day.

I thank my parents for their patience, Maria Urzica-Marai, Lenta Marai and Auta Ionasek for giving me a home, and the Renierises for embracing me as one of their own.

I thank my husband Manos Renieris for believing in me, and for his unwavering support. Above all, thank you for teaching me trust, *κατα του δαίμονα μου*. Words simply won't do.

Contents

List of Tables	x
List of Figures	xi
1 Introduction	1
1.1 Motivation	1
1.2 Anatomy Background	5
1.3 Computational Background and Challenges	7
1.4 State of the Art in Diarthrodial Joint Modeling	8
1.5 A Data-Driven Framework	11
1.5.1 Framework Instantiation and Dissertation Overview	13
1.5.2 Contributions Overview	15
2 Extracting Joint Kinematics from Medical Images	16
2.1 Introduction	16
2.2 Registration Method	17
2.2.1 Overview	17
2.2.2 Object Surface Extraction	18
2.2.3 Localized Distance Fields	18
2.2.4 Tracking Procedure	20
2.2.5 Hierarchical Approach	23
2.3 Validation Method	26
2.3.1 Data Acquisition	26

2.3.2	Experiments	27
2.4	Results	31
2.5	Discussion	35
2.6	Conclusion	42
3	Modeling Ligament Tissue from Bone Surfaces and Motion	43
3.1	Introduction	43
3.2	Related Work	45
3.3	Materials and Methods	46
3.3.1	Data Acquisition	47
3.3.2	Bone Segmentation and Modeling	48
3.3.3	Recovery of Bone Kinematics	49
3.3.4	Inter-Bone Joint Space Area Calculation	50
3.3.5	Ligament Path Estimation	53
3.3.6	Visualization and Analysis of Results	56
3.4	Results and Discussion	57
3.5	Conclusion	62
4	Modeling Articular Cartilage from Bone Surfaces and Motion	66
4.1	Introduction	66
4.2	Related Work	67
4.3	Methods	69
4.3.1	Data Acquisition and Recovery of Kinematics	70
4.3.2	Inter-Bone Joint-Space Modeling	70
4.3.3	Inferring the Cartilage Map Location and Thickness	72
4.3.4	Cartilage Contact Simulation	76
4.4	Validation and Results	78
4.5	Discussion	82
4.6	Conclusion	84

5	Predictive Simulation of Diarthrodial Joints	85
5.1	Introduction	85
5.2	Related Work	87
5.3	Methods	91
5.3.1	Data Acquisition	92
5.3.2	Model Construction	92
5.3.3	Simulation	93
5.4	Results	99
5.5	Discussion	104
5.6	Conclusion	106
6	Diarthrodial Joint Markerless Cross-Parameterization and Biomechanical Vi-	
	sualization	107
6.1	Overview	107
6.2	Introduction	108
6.3	Related Work	110
6.4	Methods	112
6.4.1	Data Acquisition and Preprocessing	112
6.4.2	Bone Surface Correspondence	114
6.4.3	Exploratory Visualization and Analysis	119
6.5	Results	120
6.5.1	Validation	120
6.5.2	Applications	122
6.6	Discussion	127
6.7	Conclusion	128
7	Conclusion	129
A	Orthopedics Terminology	133
	Bibliography	137

★ Parts of this document have been published as [29, 67, 68, 69, 70].

List of Tables

2.1	Datasets used in validation experiments	26
2.2	Validation experiments	30
4.1	Trapezoid cartilage thickness	80
4.2	Scaphoid cartilage thickness	81
5.1	Material properties used when simulating the wrist joint	97
5.2	Wrist ligament fiber lengths across the range of motion	101
5.3	Wrist articular contact size across 7 poses (mm ²)	102

List of Figures

1.1	Diagram of inter-disciplinary research	2
1.2	Diarthrodiial joints	3
1.3	Wrist implant	4
1.4	Modeling diarthrodial joints	4
1.5	Layers in the anatomy of a human joint	6
1.6	Computational modeling diagram	7
1.7	Data-driven framework	12
2.1	Registration method pipeline	17
2.2	Computing distances to material boundaries	19
2.3	Tissue-classified distance-fields	21
2.4	2D illustration of the tracking procedure	22
2.5	The human wrist	24
2.6	Wrist hierarchy	24
2.7	Capture-range expansion	25
2.8	Registration accuracy with image resolution	32
2.9	Registration accuracy with start-point perturbation	33
2.10	Registration error with a 2mm random perturbation in the optimization start point	34
2.11	<i>In vivo</i> registration results	36
2.12	Collision detection	37
2.13	Image resolution impacts voxel information-value	39
3.1	The DRUJ bones	45

3.2	Method pipeline for measurement of inter-bone joint space areas and ligament paths in joints	47
3.3	Manifold surface representation of bones	48
3.4	Distance field representation of bones	50
3.5	Anatomic coordinate system defined on the ulna	51
3.6	2D illustration for obtaining distances to bone	52
3.7	Inter-bone joint space areas in the DRUJ	53
3.8	Shortest path between two points (2D case)	54
3.9	Insertion point location	55
3.10	Shortest paths generated by the ligament model	56
3.11	Proximal and exploded lateral views of an uninjured and an injured radioulnar joint at six rotation positions.	58
3.12	Size of the ulnar inter-bone area	59
3.13	Cylindrical coordinates of the ulnar inter-bone area centroid	60
3.14	Distal radioulnar ligament paths in the injured forearm and in the matching uninjured forearm of the same volunteer	61
3.15	Length and maximum deflection of a dorsal ligament for the injured and uninjured forearms of a volunteer	62
3.16	Length and maximum deflection of a palmar ligament for the injured and uninjured forearm of a volunteer	63
3.17	The effect of insertion point perturbation on the length and maximum deflection of a dorsal ligament	64
3.18	The effect of insertion point perturbation on the length and maximum deflection of a palmar ligament	65
4.1	2D slice through a μ CT-volume image of a scaphoid bone and articular cartilage	67
4.2	Modeling inter-bone joint-spacing from bone surfaces and motion	71
4.3	<i>In vivo</i> scaphoid cartilage map	75
4.4	2D illustration of cartilage deformation process	78
4.5	<i>In vivo</i> trapezoid cartilage map	79

4.6	<i>In vivo</i> carpal contact area	82
5.1	Soft-tissue interacting directly with the scaphoid bone	86
5.2	Ligament fiber interacting with multiple bones	93
5.3	Wrist coordinate system	98
5.4	Ligament fibers in the human wrist	100
5.5	Wrist radial-ulnar deviation and maximum contact pose	102
6.1	2D slice through a μ CT-volume image of bone and articular cartilage	109
6.2	Biomechanics visual analytics framework	113
6.3	Markerless correspondence pipeline	114
6.4	Project and bin operation	116
6.5	Two wrist bones belonging to different human subjects and their corresponding pin-points	117
6.6	Surface mapping through manifold deformation	118
6.7	Visualization of a normal scapholunate joint	120
6.8	Curvature comparison of source and target manifold surfaces	122
6.9	Ligament insertion-site transfer between two hamate bones	123
6.10	Cartilage transfer between two lunate bones	123
6.11	Normal right-wrist and injured left-wrist radioscapholunate joints from the same individual	125
6.12	Pin-point and manifold deformation between a left and a right lunate bone	125
6.13	Kinematic analysis of a radioscapholunate joint	126

Data-Driven Predictive Modeling of Diarthrodial Joints

Georgeta-Elisabeta Marai

May 2007

Chapter 1

Introduction

1.1 Motivation

20% of all computer users damages their wrists due to excessive typing [84]. How do these injuries occur, and why does treatment work only for certain individuals? Subject-specific, computational models of anatomical joints can help answer such questions. However, developing such models poses significant computational challenges — for example, what level of modeling detail is necessary in order to generate biologically significant measurements, while keeping the resulting models efficient to simulate?

Developing such models also requires interdisciplinary collaboration between computer scientists and life scientists. Interdisciplinary research like the one described in this dissertation is a meeting place for experts in different fields. While our overarching goal is gaining insight into how anatomical joints work, the focus of each field is in general on different domains (Fig. 1.1). For example, doctors and biologists target in general applications, bioengineers emphasize data acquisition and validation, while computer scientists focus on developing computational modeling and analysis tools. A research project at the intersection of the data acquisition, computational tools and application domains is, for example, developing an image-based automated system for tracking small animal motion.

This dissertation focuses on modeling anatomical joints. Technically, the correct term is *diarthrodial* joints, joints that move freely — examples of such joints are the knee and

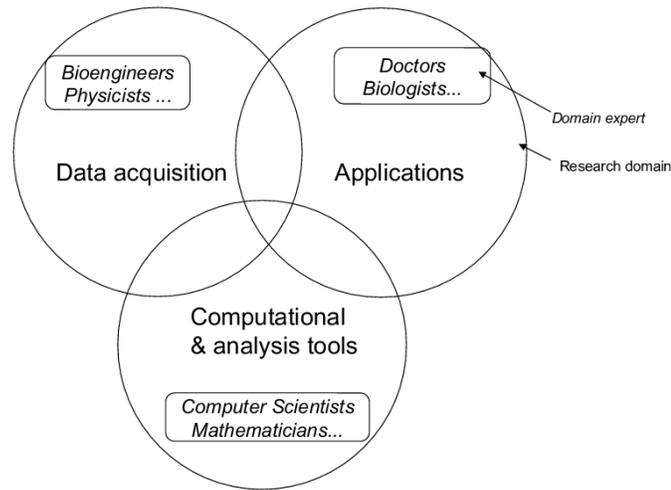


Figure 1.1: Inter-disciplinary research is a meeting place for experts in different fields. The focus of each field, however, is on a different domain: for example, doctors and biologists target in general applications, bioengineers emphasize data acquisition and validation, while computer scientists focus on developing computational modeling and analysis tools.

the elbow (Fig. 1.2). Diarthrodial joints are the structures that allow us to move. They occur wherever two or more bones adjoin and move against each other; surrounding soft-tissues stabilize the joint and protect the bones from motion-related damage.

From the application point of view, our goal is to develop tools that can generate joint models which have: 1) subject-specific capabilities, i.e., the models are adaptable to differences between individuals; 2) quantifiable capabilities, i.e., the models allow users to evaluate not only whether, for example, an injured joint differs from a normal joint, but also how much; and 3) predictive capabilities, in predicting for example the outcome of surgical interventions or therapy.

The impact of obtaining subject-specific, quantifiable, predictive models would be tremendous. For example, such models of joints could predict, for a given individual, how joint motion would be altered after a simulated surgical intervention or after therapy, and therefore help doctors plan their procedures. Second, such models would allow the design of higher-performance robots and orthopedic implants. Fig. 1.3 shows a state of the art wrist implant [6]; attempts at total wrist replacement have historically been fraught with complications, most commonly prosthetic dislocation and loosening [6, 39]. Subject-specific

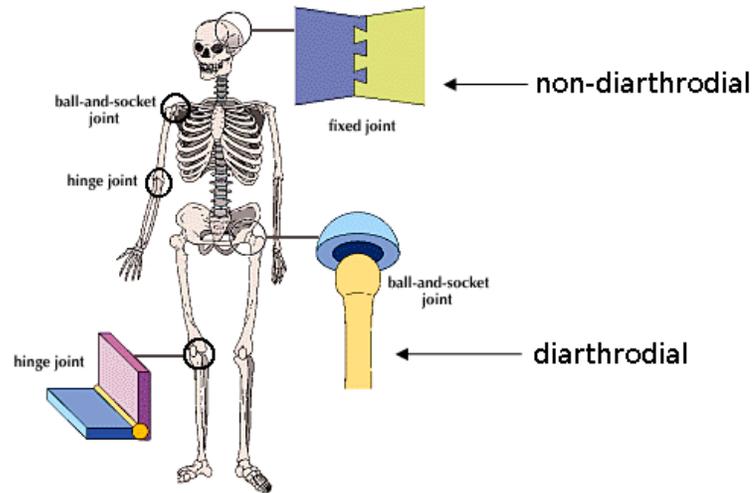


Figure 1.2: Diarthrodial joints are joints that move freely. Such joints are formed wherever two or more bones adjoin and move against each other.

wrist models able to predict contact within the wrist could help us design more performant prosthetics, tailored to specific individuals. Lastly, in computer animation the same models would generate more realistic character motion than current analytical or highly simplified musculoskeletal models. For examples of current animations generated without the help of a large team of skilled artistic animators, see <http://www.theseisgame.com/>; in the released clips, note in particular the unrealistic shoulder motion. Subject-specific, motion predictive joint models would certainly help animators.

When modeling diarthrodial joints, the computer science area of expertise is the development of computational modeling, visualization and analysis tools that take as input individual-specific medical measurements, and generate models and simulations that can provide insight into specific applications. In Fig. 1.4, the left side corresponds to the data-acquisition domain, and the right side to the applications domain; computer scientists contribute primarily to the computational and analysis tools domain. The view shown here is computer science-centric; however, the flow among the three domains is by no means uni-directional. For example, applications generate hypotheses; hypotheses influence the type of data acquired, and thus the development of data acquisition techniques, but also how much we model and at what level of detail.



Figure 1.3: State of the art orthopedic wrist implant; attempts at total wrist replacement have historically been fraught with complications, most commonly prosthetic dislocation and loosening.

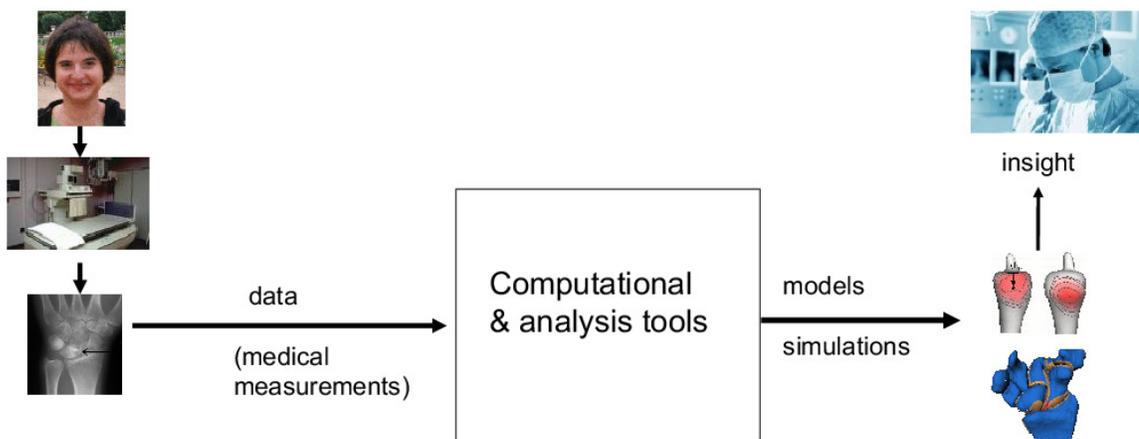


Figure 1.4: When modeling diarthrodial joints, our goal as computer scientists is to develop computational modeling, visualization and analysis tools that can take individual-specific medical measurements and generate models and simulations that can provide insight into specific applications.

1.2 Anatomy Background

The anatomy of a diarthrodial joint comprises several layers (Fig. 1.5). The first and most superficial layer is (1) skin and fat, followed by (2) the neurovascular system layer (shown in Fig. 1.5 as brightly colored threads). Blood vessels nourish the joint tissues, and nerves act as sensors and controls.

The next layer is formed by (3) muscles, active bundles of soft tissue that attach to bones through tough cords called tendons (in Fig. 1.5, muscles are shown in red, and tendons in white-pink). Muscles flex and relax as commanded by nerves; the flexing and relaxation processes modify the length of muscle bundles, and thus the muscles through their tendons apply forces to the joint.

The interesting observation at this point is that, if we remove the top three layers, a joint will still hold together and move appropriately when forces are applied to it. This observation — made thanks to clinical studies on cadaver data — is frequently used as a simplifying assumption in biomechanical modeling [32, 52]. According to this assumption, and depending on the specific application, the top two layers and their influence on the joint can be neglected, and muscles and their actions can be represented as external forces applied to a joint at tendon insertion sites.

The deeper layers that hold a joint together are: (4) ligaments, (5) cartilage, and (6) bones. Ligaments are tough, passive bands of soft-tissue connecting bones. Their role is to stabilize the joint during motion. While the anatomical rendering in Fig. 1.5 shows ligaments (in grey) as separate bands of tissue, in reality ligaments are inter-connected and form a sac; individual ligaments can be described as thickenings of the sac. The shapes and mechanical properties of individual ligaments are in general poorly documented; clinical studies indicate large variation among individuals, and among ligaments of the same individual. The sac itself contains synovial fluid; the fluid's role is to lubricate the joint and thus reduce friction during motion.

Articular cartilage (shown in white in Fig. 1.5) is a complex, living tissue that lines the bony surface of joints. Its function is to provide a low friction surface cushioning the joint bones through the range of motion. In other words, articular cartilage is a very thin shock absorber. It is organized into five distinct layers, with each layer having different structural

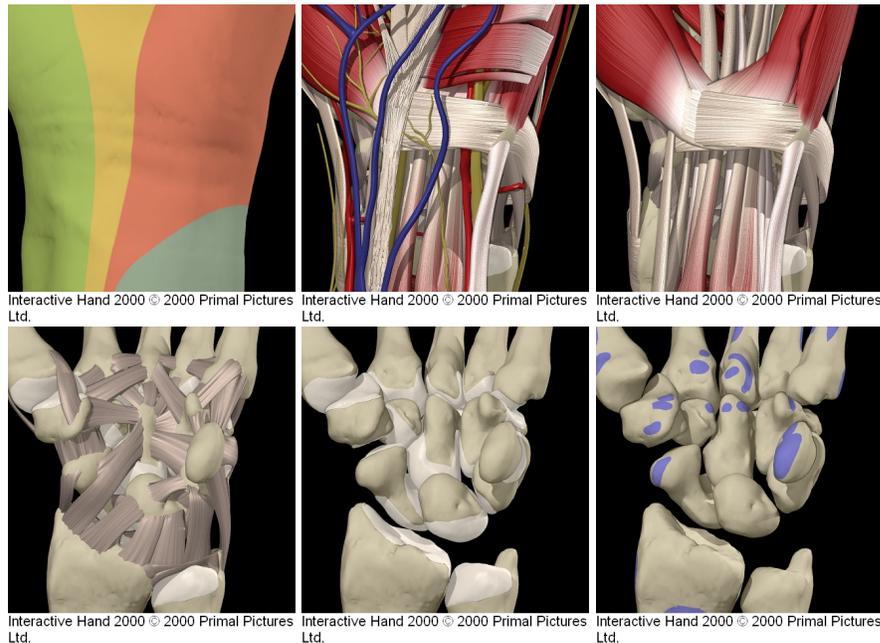


Figure 1.5: Layers in the anatomy of a human wrist — from left to right: skin and fat, neurovascular system, muscles, a subset of wrist ligaments, cartilage, bones with muscle and a few ligament insertion sites.

and biochemical properties. Erosion of this protective layer results in osteoarthritis. Cartilage is extremely slippery — 100 times more slippery than ice; as a result, cartilage contact during motion is practically frictionless.

Bones are organs with a complex internal and external structure that allows them to be lightweight yet strong and hard. The hard outer layer of bones is called compact (or *cortical*) bone tissue due to its minimal gaps or spaces. This tissue gives bones their smooth, white, and solid appearance, and accounts for 80% of the total bone mass. Filling the interior of the bone is a spongy (or *trabecular*) bone tissue which makes the overall bone lighter and allows room for blood vessels and marrow. Spongy bone accounts for the remaining 20% of total bone mass, but has nearly ten times the surface area of compact bone. While bone is essentially brittle, it does have a significant degree of elasticity. However, in the context of motion analysis and considering the large configuration changes normally occurring in a joint during motion, bones in human diarthrodial joints can be considered rigid; in contrast, bones in bat wings are believed to bend during flight [104].

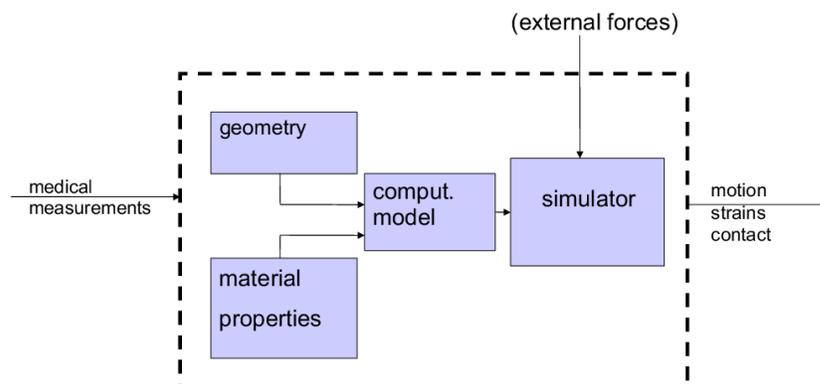


Figure 1.6: Diagram of the computational modeling process

1.3 Computational Background and Challenges

In theory, we could build diarthrodial joint models from the molecular level up to full motor function. Constructing such a model would require accurate, well-defined inputs, including complete digital anatomical models of all joint components, the material properties for all components, and a detailed understanding of the applied forces. Once the model geometry is constructed and the forces defined, algorithms and representations must be implemented to computationally model and then simulate the interactions among joint components and finally the behavior of the joint (Fig. 1.6).

However, building such models from the molecular level up will take a long time, and simulating them will be enormously slow. In practice, slightly less accurate but faster models of anatomy in which we treat bones, for instance, as rigid bodies, and tendons as inelastic bands, could serve to advance life science in the same way that the development of rigid body physics — while failing to take into account molecular forces and relativistic effects — has helped advance physical science and engineering for 150 years. The first challenge here is choosing what we need to model at what level of detail, and developing appropriate representations and approximations so that we obtain biologically significant measurements while keeping the models simulatable.

The second challenge is that many of the inputs we need to build models of diarthrodial joints are not measurable in live individuals. Some *in vivo* measuring restrictions come from current limitations of imaging technology. For example, we are still unable to image

non-invasively small structures (under 0.5mm thickness) such as wrist cartilage or ligaments in live individuals: there just aren't enough imaging capabilities to generate models of such structures. In time, progress in imaging technologies may overcome such limitations. Other measuring limitations are, however, inherently linked to *in vivo* investigation: detailed subject-specific material properties and undeformed geometry cannot be acquired without invasively disrupting the joint and thus altering its kinematics. For example, we may never be able to determine the rest-length and elasticity of a specific ligament of a specific individual without removing the ligament from the individual's joint. Input entities which are not measurable directly become, in fact, latent variables in our models. Latent variables are variables that are not directly observed but are rather inferred from other observable variables. These variables — such as cartilage thickness and location, or ligament rest-lengths — need to be inferred from directly measurable data.

1.4 State of the Art in Diarthrodial Joint Modeling

Approaches to modeling diarthrodial joints can be classified in four categories, according to the type of data they take as input and the representations they use: 1) the 'stick-and-wire', 2) the analytical, 3) the rigid bodies, and 4) the deformable-rigid hybrid approach. The resulting models of diarthrodial joints include only skeletal tissues: bone, ligaments, cartilage, and muscle-tendon units.

In the 'stick-and-wire' approach, researchers build physical models of joints from materials such as wood, nails, wire, or epoxy-resin. The necessary medical measurements are in general acquired by dissecting the joint. For example, Jacob et al. built a physical model of a human wrist from joints dissected *in vitro* [52]. Epoxy-resin casts were generated from silicon-latex molds in which the exact form of the cartilage surfaces was preserved. Ligaments were modeled with stout threads and attached at the locations observed during *in vitro* dissection; material-property differences observed among specimens were ignored.

Historically, the 'stick-and-wire' approach has generated extremely useful anatomical knowledge and a wealth of diarthrodial joint models. Nevertheless, in this approach we destroy the biological joint in order to study it. This limits the applicability of this approach

to cadaver subjects. Furthermore, even in the cadaver case it is difficult to generate and compare subject-specific models. Since in general invasive studies alter joint kinematics and material properties, models generated through the ‘stick-and-wire’ approach are also difficult to validate.

The analytical approach generates simulatable, but greatly simplified and often heuristically defined models of diarthrodial joints. The only inputs used in this approach are the bone lengths and the joint range-of-motion. For example, knee joints are modeled as hinges, and wrists as ball-and-socket devices. The resulting models are often used in computer graphics; such models can roughly replicate the range of motion observed in live individuals. However, the analytical approach has reduced appeal to medical applications and realistic computer animation. Such applications require in general detailed models tailored to subject-specific data.

In the rigid-bodies approach, joints are modeled as collections of interacting rigid bodies. The inputs here are individual-specific medical measurements of bone surfaces, and anatomy-book knowledge. Bones are modeled in general as 3D meshes. If modeled at all, soft tissues are represented as springs or rigid shells surrounding the bones [95, 32, 51, 98]. Some bones are rigidly connected; soft-tissue wrapping and deformation are in general discounted. However, some of the resulting models may be able, on restricted ranges of motion, to correctly predict some bone kinematics. For example, Blankevoort et al. [18] built a rigid-body model of the knee with motion-prediction capabilities. The model was generated from medical images of the bones and anatomy book knowledge such as tendon insertion sites and material properties; bones were constrained to not inter-penetrate during motion. We note that, in general, deformable contact within the joint is not modeled in the rigid-body approach.

The most sophisticated approach to date to modeling diarthrodial joints is the hybrid deformable-rigid approach. The inputs in this case are again non-invasive individual-specific medical measurements and anatomy book knowledge. In this approach, bones are considered to be rigid bodies and the other tissues are considered to be deformable. However, the geometrical representations and properties of deformable tissues vary depending

on the complexity of the model and the medical measurements available as input. For example, the geometry of knee ligaments can be acquired through magnetic resonance imaging, and thus can subsequently be modeled through accurate, sophisticated representations such as finite element methods. The most recent and sophisticated models are capable of predicting strain through entire ligamentous [41] and cartilaginous structures [113, 33, 75] using advanced finite element analysis. In contrast, the geometry of individual-specific ankle ligaments cannot be currently acquired *in vivo*, and thus ankle ligaments are commonly represented as non-wrapping, line spring elements.

With the notable exception of the knee — a relatively large joint of high clinical interest, current rigid-deformable models for most diarthrodial joints are either relatively crude or model only a few components in high detail. They typically model the ligaments as line springs, don't include cartilage, or include only a few bones.

The most sophisticated hybrid model to date of a complex joint has been developed by Carrigan et al. [24]. They created a simplified 3-D finite element model of the carpus, in which hollowed bones were modeled through finite element modeling, then their articulated surfaces were extruded to mimic cartilage. In this model several pairs of bones were fused into single rigid bodies, and ligaments were modeled as non-wrapping line springs. Material-property parameters were specified at the input. However, for unclear reasons, the resulting model was not stable. In the end, bone motion was restricted artificially to certain directions through non-physiological constraints in order to prevent the carpus from collapsing under applied loads. I speculate the instability may have been due to insufficient modeling detail; in particular, this dissertation demonstrates wrapping soft-tissues play an important role in stabilizing the carpal joint.

Important additional limitations of the models generated through the approaches surveyed in this section are the lack of kinematic validation data, and the inability to perform comparisons between subject-specific models.

When we consider the space described by models generated through the four approaches described above, an interesting trend becomes apparent. The more complex a joint, the fewer models for it exist, and the fewer predictive capabilities these models have. For example, knee models (3 bones, 4 ligaments) range from hinge models able to roughly

replicate the range of motion to models able to compute contact and strains in live individuals [118]. In contrast, wrist models (8 bones, dozens of ligaments) are far fewer and have far simpler capabilities; the same holds true for most human joints, from shoulder to ankle joints.

While the knee commands particular clinical interest, this does not quite explain the paucity of models of more complex joints. A first observation is that joint complexity influences simulation efficiency — computing accurate contact among 100 deformable components is certainly more expensive than computing contact among 3 components. This observation ties into the first computational challenge — appropriate representations for efficient simulations — identified in section 1.3. The second, more interesting observation is that, in general, complex joints have smaller components. The smaller the components, the fewer direct medical measurements are available. Unfortunately, we cannot accurately simulate soft-tissue behavior when soft-tissue measurements are not available. In fact, what we see is the impact of measuring limitations and hence paucity of morphological data on the model space. The more complex a joint is, the fewer relevant data are available, and hence the fewer and weaker models we have. This second observation ties into the second computational challenge — measuring limitations — which we also identified in section 1.3.

1.5 A Data-Driven Framework

The thesis of this dissertation is that a data-driven modeling approach, when tightly coupled with visualization and analysis tools, can generate adaptable, quantifiable, predictive models of diarthrodial joints. In the computer science definition, a data-driven approach estimates a hitherto unknown mapping (or dependency) between a system’s inputs and outputs from the available data [76].

I present in this dissertation a data-driven framework for the predictive modeling of diarthrodial joints. The framework allows for the generation of adaptable, quantifiable, predictive models of complex joints, in spite of current measuring limitations. The resulting computational models are somewhat unusual in their hybridization of data representations.

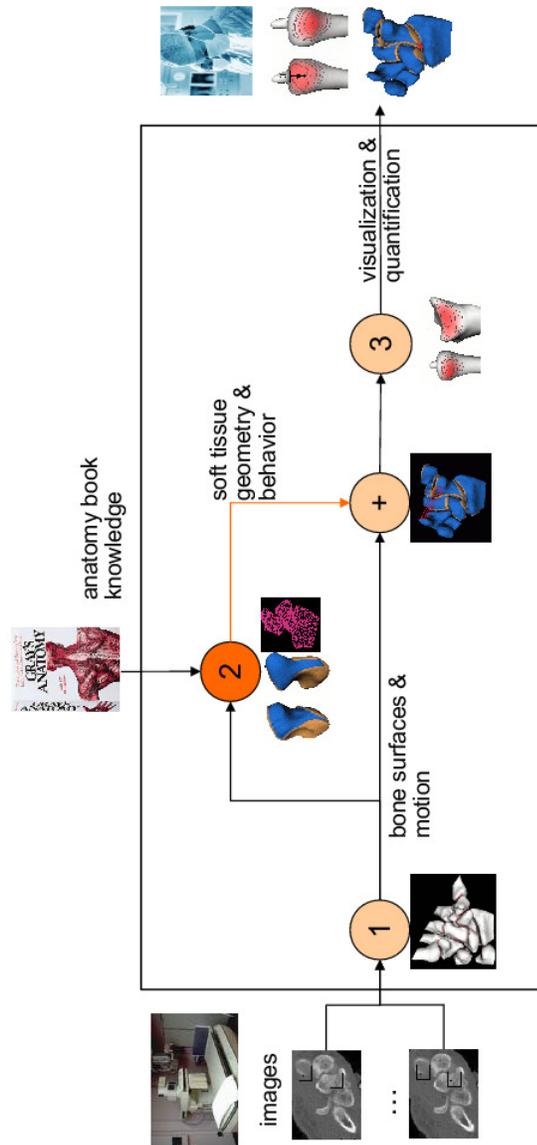


Figure 1.7: Data-driven framework for modeling diarthrodial joints: (1) in a first step, we extract bone surfaces and motion from sequences of medical images (Chapter 2). (2) Next, we use this data and anatomy-book knowledge to infer and model some of the missing data — such as soft-tissue geometry and behavior (Chapter 3 and Chapter 4). (+) We assemble the measured and inferred data into a predictive model of the joint (Chapter 5). (3) Finally, we propose quantitative measures and use them to analyze disease-related joint behavior (Chapter 6). We call the framework ‘data-driven’ because of step (2) above, in which we use sampled data to infer soft-tissue geometry and behavior.

Each representation has strengths for various aspects of the modeling and we combine them in unique ways to achieve simple, elegant and accurate estimations of biologically relevant measurements.

The key idea behind the framework is to use sampled data to infer unknown data. The framework uses as input currently available subject-specific medical measurements and anatomy-book knowledge; but it uses more than one instance of such measurements. In a first step, we augment data acquired through medical imaging — such as bone geometry — with motion information. The idea here is that motion information can provide insight into soft-tissue properties: for example, wider ranges of motion can be associated with laxer soft-tissues, while narrower ranges of motion can be associated with stiffer soft-tissues. Next, we use this augmented data and anatomy-book knowledge to infer and model some of the missing data — such as soft-tissue geometry and behavior. We assemble the measured and inferred data into a model of a joint. Finally, we propose quantitative measures and use them to analyze disease-related joint behavior.

Figure 1.7 shows the flow among the different framework components. The framework is *data-driven* because in step two we infer soft-tissue geometry and behavior from sampled motion data.

In this work we model only the skeletal tissue of a joint: bones, cartilage, and ligaments. Muscles are represented as external forces, provided by the user. As a simplifying assumption, in several instances we approximate tissue behavior; for example, we model articular cartilage as a single layer, not five. Incorporating explicitly the external anatomical layers — muscles, neurovascular system, skin and fat — and modeling tissue behavior in more detail are beyond the scope of this dissertation, and constitute directions of future work.

1.5.1 Framework Instantiation and Dissertation Overview

The framework instantiation described in this dissertation uses as input computed tomography (CT) volume images of a joint. Computed-tomography imaging can be thought of as X-ray imaging in 3D: in the resulting grayscale volume images air shows in black, bony material in bright intensities, and soft-tissues in shades of grey. When imaging joints, we

prefer CT technology over other modalities — such as magnetic resonance imaging — because CT offers superior resolution and spatial accuracy. Bone surfaces are extracted from such a reference volume image.

Next, we recover motion for each joint bone. This itself is a difficult problem: we cannot use sensors on the skin surface to recover bone motion, because there is significant relative motion between skin and the bones underneath. Instead, we first CT-image the joint in a few poses, sampling the space of joint kinematics. We then track the bones across the sequence of volume images through object registration (Chapter 2). By tracking we mean here recovering the rigid transform that takes each bone from one sample pose to another. Tracking accuracy is paramount when analyzing joint kinematics, because even small errors — errors commensurate with the voxel size — can result in false bone inter-penetration during motion. Preserving inter-bone spacing is important because cartilaginous soft-tissue is, in fact, located in this spacing.

Note that the motion tracking tool we describe in Chapter 2 could be replaced by other tracking tools, using perhaps different input data, such as bone surfaces and series of 2D images of the moving joint. The only condition here is that alternative instantiations of this first framework component should generate similarly accurate results. To the best of my knowledge, currently there are no similarly accurate alternative tracking tools.

Next, we use the acquired bone-surfaces, sampled joint-kinematics and anatomy-book knowledge to infer and model ligamentous (Chapter 3) and cartilaginous tissue (Chapter 4). This second component could be modified to incorporate directly measured geometry, if available.

We assemble bones, ligaments, and cartilages and infer some of the model parameters by imposing joint equilibrium at the sampled kinematic poses (Chapter 5).

The last component of our system is an automated tool for the cross-subject analysis and visualization of anatomical joints (Chapter 6). We use this component to explore and measure the influence of injury on joint kinematics. The analysis tool could also be replaced by alternative, for example manually-aided, cross-parameterization techniques.

Applications of this framework instantiation are presented in Chapter 3 (forearm malunion), Chapter 5 (wrist close-pack pose) and Chapter 6 (scaphoid non-union).

Chapter 7 discusses the contributions of this work and proposes directions of future research. Orthopedics terminology is briefly reviewed in Appendix A.

1.5.2 Contributions Overview

This dissertation presents novel representations, computational modeling, visualization and analysis tools that are needed to integrate subject-specific data with the predictive modeling process of diarthrodial joints.

The data-driven framework presented in this dissertation — while having certain limitations as discussed in Chapter 7 — allows for the development of complex, automatically-tuned subject-specific models that have predictive capabilities.

I instantiate the framework in a collection of tools: 1) a sub-voxel accurate method for tracking bone-motion from sequences of medical images; 2) computational tools for estimating soft-tissue geometry and contact; and 3) a tool for the visual and quantitative exploration of joint biomechanics.

The results generated through this framework instantiation have already affected orthopedists' understanding of the many diseases afflicting human joints [29]. With such a better understanding, improvements in treatment for injuries are possible as well as reductions in injuries.

In addition to providing specific insight into joint mechanics, the developed tools and resulting databases should be applicable to the study of pathology and injuries, including arthritis, ligament tears, bone fractures, and surgical reconstructions. The tools and methodologies I demonstrate on forearm and wrist data will be generally useful for the study of bone, cartilage and ligament interactions in other complex multi-articular joints, including the foot and spine, as well as in other joints such as the knee, elbow, and human shoulder. The tools will also be applicable to animal studies, in basic biology research. Ultimately, this work has the potential to create a modeling approach that will more simply and efficiently explain and predict the underlying biomechanics of musculoskeletal systems.

Chapter 2

Extracting Joint Kinematics from Medical Images

2.1 Introduction

As research areas that employ image registration techniques focus on ever-smaller features, they require higher registration accuracy. *In vivo* kinematic analysis of small joints, such as the wrist, exemplifies the need for highly-accurate intra-subject, same-modality registration. A common way to analyze joint kinematics is by CT-imaging the joint bones in several different positions and registering them across all volume images. While early studies have focused on retrieving bone pose and orientation, recent research focuses on measuring how more subtle features like inter-bone spacing change with motion. In the first case, errors on par with the image sampling step-size, like those introduced by existing tracking systems, may be acceptable, while in the latter case errors as small as 0.5 mm can compromise the study by introducing inter-bone collisions. At the same time, decreasing the image sampling-step results in increased imaging cost and time. We need a subvoxel-accurate method for registering features whose size is on par with the image sampling step.

We describe in this chapter an automated intra-subject same-modality registration method that attains subvoxel-accuracy. The method is of interest to any registration applications involving datasets where the image sampling step is larger than features of interest.

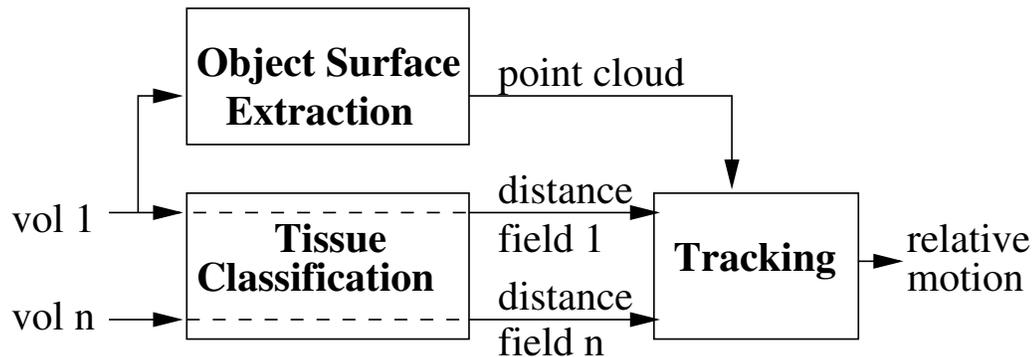


Figure 2.1: Our registration method works in three steps: to register an object across two volume images, we first extract the object surface from one volume; we generate through tissue classification a localized distance field from each volume; we then use the object surface and the distance fields to track the object.

2.2 Registration Method

2.2.1 Overview

The registration method works in three steps on a series of volume images. First, we extract the surface of the object to be registered from an arbitrarily-selected reference image. Next, in order to obtain an accurate localized distance field for registering the object, we classify the tissues in each volume image using a probabilistic approach. Last, we register the object by automatically adjusting its position and orientation, thereby minimizing a distance-field derived cost function (Fig. 2.1).

In the case of multi-object structures (e.g., joints in the human body) we infer from the distance field an object hierarchy that expands the capture range of our procedure beyond the capabilities of previous registration methods. The capture range represents the range of positions from which a registration algorithm can converge to the correct minimum or maximum.

We validate our method using CT data from a cadaver with external markers, an *in vivo* volunteer, and forty subjects participating in a wrist-motion study. We compare the performance of our method against a manually aided segmentation-based method as well as a standard grey-value-registration method.

2.2.2 Object Surface Extraction

Through manual segmentation, thresholding, and user interaction, we extract in this first step an object surface from a reference CT volume image [28]. Summarizing this reference, the contours defining the outer cortical bone surfaces of each object are extracted using thresholding and image algebra processes with a 3-D imaging software package (Analyze AVW 2.5; Biomedical Imaging Resource, Mayo Foundation, Rochester, MN). Each contour is then assigned to the appropriate object using Matlab custom code, which designates contours based on the contiguity of their centroids. Contour lines are output as collections of discrete points, which are distributed densely along each contour and sparsely between different contours.

2.2.3 Localized Distance Fields

In the second step of our method, we classify the tissues in each CT volume image probabilistically in order to generate a localized distance field. Our tissue classifier uses the partial-volume technique described by Laidlaw et al.[63]. This method identifies distances from material boundaries and creates distance fields for individual materials. The technique assumes that, due to partial-volume effects or blurring, voxels can contain more than one material, e.g., both cortical bone and soft tissue. Each voxel is assumed to contain either a pure material or two pure materials separated by a boundary (Fig. 2.2).

We treat each voxel as a region, by subdividing it into 8 subvoxels, and evaluating the image intensity and its derivative at the center of each subvoxel. The intensity is interpolated from the discrete data using a tricubic B-spline basis that approximates a Gaussian. Thus, intensity and derivative evaluations can be made not only at sample locations, but anywhere between samples as well. From this intensity and derivative information we infer a histogram of each voxel, accumulating the contributions from all subvoxels. This gives us a more refined histogram than we would obtain by evaluating only the intensity values at the same number of points. Histograms are next fit by basis functions, each basis function corresponding to either one material or a mixture of two materials.

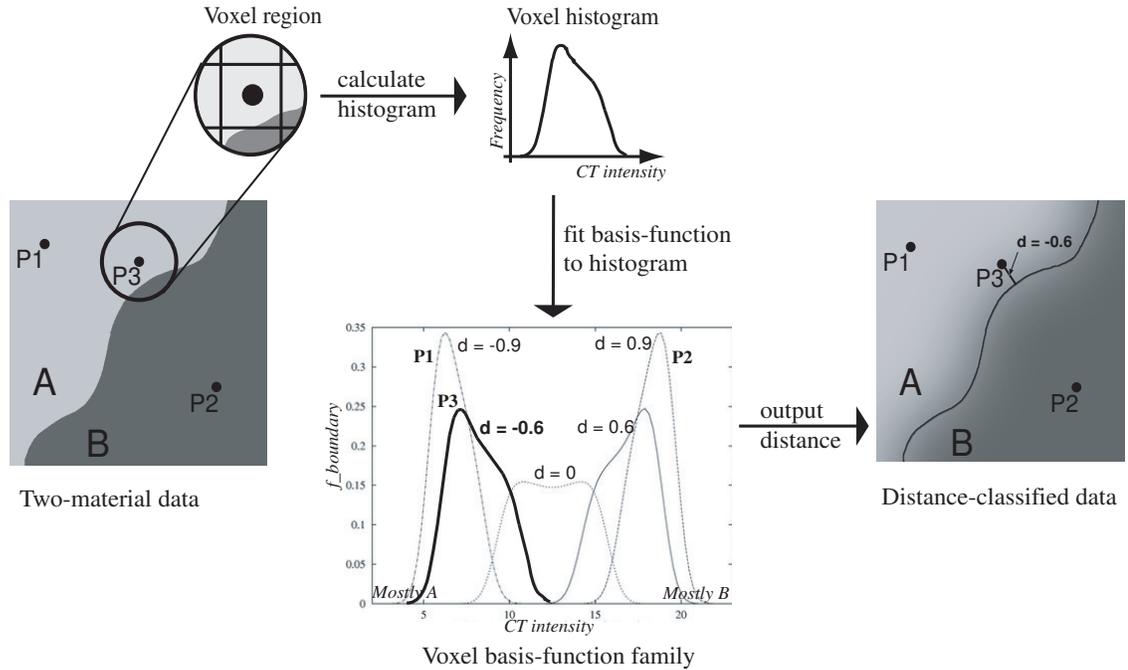


Figure 2.2: The classification algorithm computes distances from sample points to material boundaries. Points P_1 and P_2 lie inside regions of a single material, either A or B. Point P_3 lies near the boundary between A and B. We treat each voxel as a region, by subdividing it into 8 subvoxels, and taking into account information from neighboring voxels. We evaluate the image intensity and its derivative at the center of each subvoxel. The resulting voxel histogram is then fit to a family of basis functions ($f_{boundary}$), whose shapes reflect the estimated distance d to the material boundary. The best-fit basis instance is selected through a maximum likelihood process ($d = -0.6$ fits best the histogram of point P_3). The result is a localized distance field, that specifies, at each point, the signed distance from the point to the material boundary.

Pure material basis-functions are Gaussians whose parameters are the mean CT grey-scale value and standard deviation for that material. Mixture basis functions have an additional parameter, d , describing the distance from the center of the voxel to the boundary between materials. As the distance parameter changes, the shape of the basis function changes (Fig. 2.2). The basis-function shape that best fits each mixture voxel histogram is chosen through a maximum-likelihood process. The derivation of the basis-function formulas and the description of the optimization process are presented in detail in [63]. We repeat the fitting procedure for each material, and select the material basis function that fits each voxel histogram best.

For each tissue type, the sole input required by our tissue classifier is an initial estimate of its CT grey-scale value’s mean and standard deviation. We estimate these measures from sets of approximately one hundred voxel samples, one set per tissue type. We consider three distinct pure materials: air, soft-tissue, and bone. Soft-tissue is present both outside bones and inside bones (as bone marrow). Material samples are collected only once, from the same *in vivo* dataset. We consider two instances of the mixture basis function: one modeling mixtures along air and soft-tissue boundaries, the other modeling mixtures along soft-tissue and bone boundaries. We initialize the basis function parameters to the same values throughout all the datasets, including the *in vitro* datasets.

Through this basis-function tissue-classification process, we generate a localized distance field. The distance field is a scalar 3D grid that specifies at gridpoints the distance to the closest boundary between two materials. The distance field is local in the sense that the distance estimate is specified only as far as gridpoints located within a five voxels band around the material boundary. Distances between gridpoints are approximated through tricubic interpolation.

The classification of a wrist volume image produces one distance field per material type. We use the distance field corresponding to bone material (Fig. 2.3) in the tracking stage of our registration method.

2.2.4 Tracking Procedure

In the third step of our method, we register an object through a sequence of CT volume images classified using the process described in section II.C. For each bone, we recover the rigid body transformation between the reference image that generated the geometrical model and a target image. The rigid body transform is expressed as a rotation around the bone’s center of mass, and translation.

An object’s geometric model is registered with a target image of the object when its signature in the reference distance field, D_R , is most similar with its signature in the target distance field, D_T . We measure this similarity with a sum-of-squared-differences cost function that takes into account the reference and target distance-field values of the vertices in the geometric model. The sum is weighed by the number of vertices that are still inside

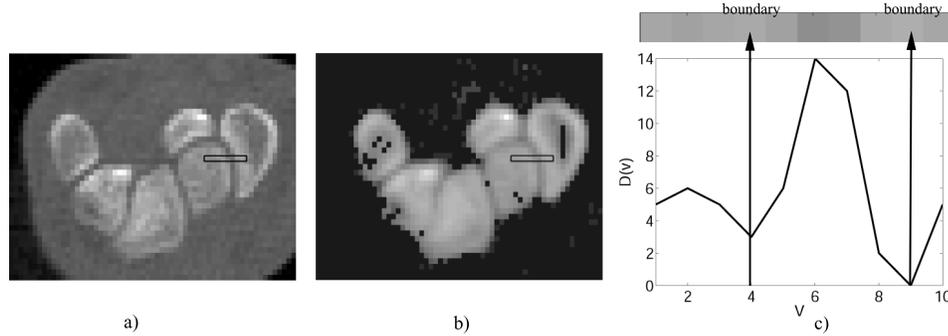


Figure 2.3: Tissue-classified distance-fields quantify the distance from the center of each voxel to the closest boundary. (a) One slice from a low-resolution (0.9 x 0.9 mm) wrist CT volume image. (b) Localized distance field corresponding to bone material. Dark pixels have been classified as either pure soft-tissue, pure air, or soft-tissue and air mixture. The area of interest in the box crosses two bony boundaries and is detailed on the right. Each voxel in the field codifies the distance from the voxel center to the closest bony boundary; the lighter the grey, the closer to a bone boundary the voxel is. (c) Plot of the distance values along the strip on top. Note the two dips in the plot corresponding to the two bone boundaries. In this particular case the bone cortex is very thin (1 voxel wide); consequently there are no samples inside the bone cortex to be associated with negative (‘inside’) distance values; hence the distance function $D(v)$ does not take negative values.

the target distance field after applying the current transform to the model. The cost function is thus:

$$F = \frac{1}{V} \sum_{j=1}^n (D_R(p_j) - D_T(p'_j))^2, \quad (2.1)$$

where p_j are points in the geometric model, p'_j are the 3D points obtained by applying the current translation and rotation to p_j , n is the number of points in the geometric model, and V is the number of points that are still inside D_T after rotation and translation. Whenever p'_j is outside D_T , $D_T(p'_j)$ returns an approximation of the distance from p'_j to the D_T volume, obtained by projecting p'_j on the closest face of the volume. This expands the cost-function gradient outside the volume to register, in order to accommodate partially-scanned bones. Note that, by incorporating D_R in the cost function, we compensate for the small errors in boundary-point location that occur occasionally during segmentation of the geometric model. Because this cost function attempts to match distance-field signatures, geometric model vertices that diverge slightly from the true bone boundary due to segmentation-errors will be off by the same amount in the registered image.



Figure 2.4: 2D illustration of the tracking procedure. In this example we search for the optimal location of the 2D boundary of a bone (shown in white) using a 2D bone and soft-tissue distance field (shown in grey). (Left) In a highly unlikely neighborhood the cost function F has a high value; the bone boundary may become trapped in local minima. (Center) In the neighborhood of the solution the cost function F has a lower value, as some boundary points overlap with lower distance field values. The distance field serves as a local gradient: F decreases smoothly as the location and orientation of the white boundary approaches the correct solution. (Right) At the correct location and orientation the cost function F should be close to zero.

Our tracking procedure searches for the position and orientation of each bone that results in maximal distance-field similarity at registration, i.e., the rotation and translation that minimizes F (Fig. 2.4). We use a quasi-Newton algorithm to solve the optimization problem [1]. The distance volume serves as a smooth local gradient field, which leads to rapid convergence when the search starts from a point where at least a few geometric model vertices are within the capture region of the localized distance field. In practice, we begin by applying to all the bones a rough alignment translation M_{com} . The translation aligns the center of mass of the bony points in the first five slices of the distance field with the center of mass of the five most proximal contours that define the outer cortical bones in the joint (see Section 2.2.2). For example, the alignment transform to pre-register a human wrist would use the first five slices of a wrist distance field and the five most proximal contours of the ulna and radius bones. This approximation suffices as a search start point.

The quasi-Newton method is fast and robust; however, like most optimization procedures, it is susceptible to being confined to sub-optimal local solutions. Consequentially, we use 64 perturbed start positions for each bone and choose the solution that yields the smallest value of the error function. Multiple searches per bone can be performed in parallel. The optimization procedure is stable with respect to perturbations in the space of

possible rotations. This is consistent with the fact that rotations around a spherical object’s center of mass are not likely to change the object’s original capture region. The perturbed start positions were therefore generated by sampling the space of possible initial translations on three concentric spheres of radius 2, 4, and 8 voxels respectively. In our experience, the majority of the repeated optimizations per bone returned the same minimum. The alternative local minima were at least one order of magnitude higher (expressed in squared millimeters).

2.2.5 Hierarchical Approach

The distance field formulation allows us to apply the tracking procedure hierarchically, expanding the capture range of our method. We derive a hierarchy empirically, based on a trial-and-error analysis of the start values of the cost function F on a few separate sequences of volume images. For a complex structure like the human wrist (Fig. 2.5), we use three *in vivo* sequences of volume images. Each sequence consists of ten different wrist poses, each of which corresponds to a different human subject. All possible tree hierarchies starting from the radius and ulna and branching towards the metacarpals were considered; we chose the one which generated best start values of the cost function across all sequences.

We run the optimization procedure on successive layers of the wrist bones, starting with the forearm bones, as shown in Fig. 2.6. We iterate through bones: once we detect the motion of bone b_i through cost function optimization (*optimization* transform), we propagate the motion to all the bones that have b_i as an ancestor in the tree hierarchy (*propagation* transform), then we move on to the next bone. Optimization and propagation transformations are accumulated for each bone.

The hierarchical approach ensures that we always start an optimization step from a reasonable neighborhood, thereby boosting the capture range of the registration procedure from less than 5° rotational pose increments to a full range of wrist motion (about 180°), as shown in Fig. 2.7.



Figure 2.5: The human wrist is a complex structure comprising the distal end of the two forearm bones, and eight small, tightly packed carpal bones. In this X-ray view the five metacarpals are also included. Figure reproduced with permission from [89].

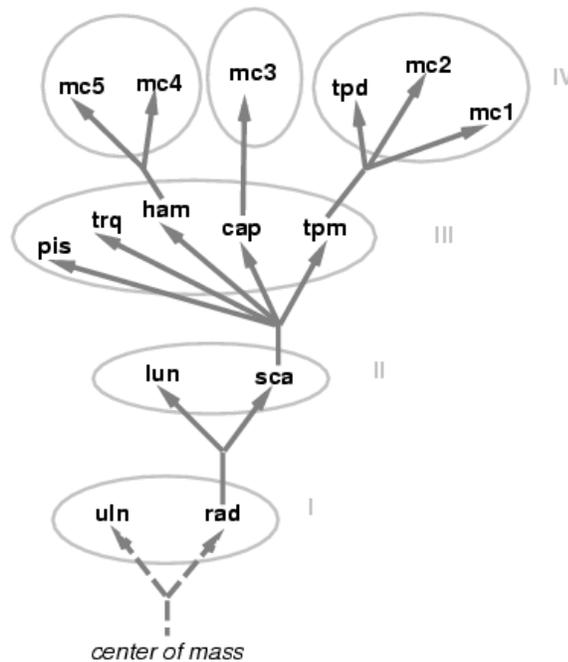


Figure 2.6: Wrist hierarchy induced from distance field information. We consider four layers in ascending order from the forearm: **ulna** and **radius**; **lunate** and **scaphoid**; **pisiform**, **triquetrum**, **hamate**, **capitate** and **trapezium**; **metacarpals** and **trapezoid**. During a propagation step the motion of a bone b_i is propagated to all bones in ascending levels that have b_i as an ancestor. The hierarchy indicates the **radius** and **scaphoid** may be governing the motion of the other bones.

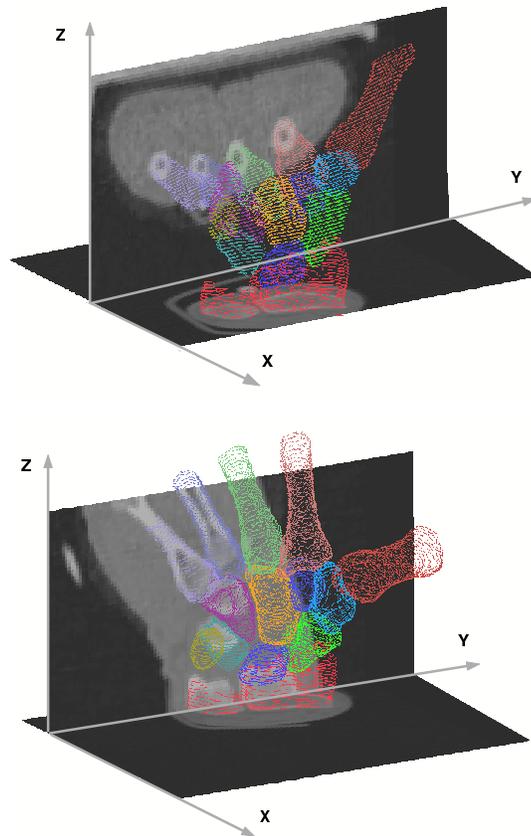


Figure 2.7: Imaged wrist-poses do not necessarily come in small motion increments. The images show the same geometric wrist model, after registration, in two different poses. The orthogonal greyscale planes correspond to vertical and horizontal sections through the CT volume images (darker grey areas correspond to soft-tissue, brighter areas to bones). Note the significant differences in bone posture, orientation, and overall wrist location in the scan volume. As shown here, two subsequent instances of the same wrist can be outside each other's capture region.

Data-set	Type	Number of images (subjects x poses)	Size	Voxel size (mm^3)
A	<i>in vitro</i>	1 x 4	512 x 512 x 141	$0.31^2 \times 1$
B	<i>in vitro</i>	1 x 4	171 x 171 x 141	$0.94^2 \times 1$
C	<i>in vivo</i>	1 x 2	180 x 180 x 60	$0.78^2 \times 1$
D	<i>in vivo</i>	80 x 12	100 x 100 x 80	$0.94^2 \times 1$

Table 2.1: Datasets used in validation experiments

2.3 Validation Method

In this section we describe a series of experiments where the method was used to register wrist intra-subject CT images. In order to compare our method’s performance with earlier results reported in the literature, we begin by evaluating our method’s accuracy on high-resolution, marked *in vitro* data. We then examine our method’s robustness with respect to practical issues such as image-resolution and perturbations in the registration start-point. We take validation one step further by examining our method’s performance on *in vivo*, unmarked data. Finally, we evaluate our method’s robustness with typical *in vivo* factors such as variation in image subject and object pose.

2.3.1 Data Acquisition

Four different datasets (Table I) were used in our experiments. All datasets were acquired using CT technology (Hispeed Advantage, General Electric Medical System, scan parameters: 80kV, 80mA). All images consist of axial slices, with the x axis oriented horizontally right to left, the y axis horizontally front to back, and the z axis vertically up, such that the image resolution is lowest in the z direction. The geometric-model point clouds have between 2000 and 8000 points, depending on the size of the bone.

2.3.2 Experiments

In vitro accuracy and robustness experiment

In this experiment we evaluate *in vitro* accuracy against the ground truth yielded by external marker registration. We further compare our *in vitro* results with those generated by grey-value registration, implemented as described further below.

To enable comparison with earlier results reported in the literature, we use the high-resolution dataset A, consisting of four CT images of a fixed specimen (separated forearm and hand) in different poses. Both components — the forearm and the hand — were encased in plastic resin to prevent relative bone motion. To better reflect the *in vivo* scanning protocol the phantom forearm bones were only partially included in the scan field-of-view for three poses. Seven markers (ceramic spheres of various high-tolerance diameters) were rigidly glued to each specimen component, allowing us to establish the registration ground truth *in vitro*. Marker contours were extracted from each volume image by thresholding at 600 Hounsfield units. The contour images were then processed with a 3D imaging software package (Analyze AVW 2.5). The centroids of the seven spherical markers (one set per specimen component) were used to calculate rigid-body motion by a method of least squares [79].

In both the tissue-classification method and the grey-value method, the optimization procedure is initialized with the ground truth. The resulting registration transforms should deviate from the given true transform due to each method’s translation and rotation error. In both methods we compute for each registered bone the error relative to the true transform. We report the relative error (mean and standard error obtained by cross-registration of the four images in dataset A) as a translation and rotation in *helical axis of motion* (HAM) coordinates [82]; HAM coordinates express rigid-body motion as pairs (θ, t) of rotations around and translations along a unique helical axis.

Next, we examine our method’s robustness with image resolution, since in practice our *in vivo* data’s resolution was limited by the large number of subjects participating in motion studies and the large number of images acquired per subject. To this end, we repeat the accuracy experiment on dataset B. Dataset B, designed to simulate lower resolution

data, was obtained by smoothing and subsampling the images of dataset A.

Last, we examine the impact of initialization on the tracking procedure. We note that true transform data, as yielded by external markers, is usually not available *in vivo*. To simulate this situation *in vitro*, we perform in this experiment a perturbation study, in which the optimization procedure is restarted repeatedly from the ground truth yielded by external-marker registration, plus a small random rigid-body transformation. We perform a set of five trials, with a translational perturbation of 2mm (approx. 2 voxels in image space) in a random direction, followed by a second set of five trials, with a translational perturbation of 5mm in a random direction. Again, we report error relative to the ground truth transform, mean and standard error obtained by cross-registration of the four images in dataset B, for both our method and grey-value registration.

Grey-value registration implementation Grey-value registration is a voxel-property registration method that has been successfully used to track joint motion from sequences of volume images. Given two or more volume images and a surface model of the joint bones, grey-value registration attempts to find the optimum location of each bone across the volume images. The method operates directly on the image grey values, via different paradigms such as cross-correlation or Fourier analysis. For example, Snel et al. [99] use chamfer matching and texture characteristics to track 3D wrist motion across sequences of CT volume images.

Grey-value registration was implemented as in Snel et al. [99], with several modifications to increase accuracy. First, all the points, as opposed to a random 10%, with a greyscale value greater than 600 Hounsfield units of each image were used in the calculation of the root-mean-square cost function; the values in each target image were obtained by tricubic interpolation. We used a high-performance library implementation [2], as opposed to a custom implementation, of the downhill simplex method of Nelder and Mead, with a maximum deviation from the initial transform values of $\Delta t = 6$ voxels per axis and $\Delta \theta = \frac{\pi}{4}$. To further boost this method's ability to deal with partially-scanned bones, the original cost function was also slightly modified to approximate distance to the target volume whenever the model's points were outside the target image during matching (Section 2.2.3).

In vivo accuracy experiment

Because it is technically impossible to know the ground truth *in vivo*, we evaluate our method’s accuracy by comparing results with the mean answer of several manual registration trials (described further below), and with the results generated by grey-value registration. In this experiment we use dataset C, consisting of two low-resolution CT images of the same *in vivo* left wrist, one with the wrist in a neutral pose (targeted by visually aligning the back of the hand with the back of the forearm and the third metacarpal with the long axis of the forearm) and one with the wrist extended.

Note that in this experiment we enhance the grey-value method with the hierarchical approach described in Section 2.2.5. Without the hierarchical enhancement, the capture-range capabilities of the grey-value method are surpassed by the range of joint-motion in dataset C, rendering the method inapplicable. Results from all three methods — tissue-classification, manual, and grey-value are further verified using the following visualization method.

Visual validation is performed by superimposing the registered bone geometric wire-frame models with vertical and horizontal slices of the volume image. Two sliders control the vertical and horizontal slice displayed. The registration results are automatically checked for potential erroneous collisions between objects that coexist in the same image, at the cost of further geometrical processing. To this end, a NURBS surface is fit to each object geometry (Raindrop GeoMagic, Research Triangle Park, NC), a level-set distance field representation is then generated from the NURBS representation [73], and the inter-object distance is evaluated accurately for each vertex of the NURBS surface with respect to all neighboring objects [67]. The generated NURBS surfaces have typically on the order of 10^3 to 10^4 points. Collisions are indicated by negative inter-object distances and reported to the user. When collisions happen, each object surface is further color-mapped and iso-contoured according to the inter-object distance, in order to create an informative visualization (see Section 2.4). Registration results are also evaluated numerically, by examining the final-fit cost function values. Results are visually inspected in cases where fit values were abnormally high, i.e. above 0.01.

Validation experiment	Datasets	Results compared against
<i>in vitro</i> accuracy and robustness (image resol. and start-point perturbation)	A, B	grey-value registration
<i>in vivo</i> accuracy	C	grey-value registration segmentation-based registration visual inspection (collision detection)
<i>in vivo</i> robustness (image subject and object pose)	D	visual inspection numerical analysis

Table 2.2: Validation experiments

Segmentation-based registration Five medical school students, all familiar with the segmentation procedure and the anatomy of the wrist, manually segmented the wrist from each of the two volume images in dataset C. Each segmenter took several runs through the procedure, for a total of twelve runs. Registration of the carpal bones between two volume images was subsequently accomplished with an inertia-matching method [27]. Finally, relative motion to the radius was reported for each bone and run. Statistics on the registration results (rotation and translation mean and standard deviation, per bone) were collected; note that these statistics include intra-observer variation.

In vivo robustness experiment

Finally, we evaluate our method’s *in vivo* robustness with respect to object pose and human subject data (dataset D). Dataset D consists of CT wrist images acquired from forty human subjects. Ten to twelve low-resolution volume images of both wrists were acquired per human subject, spanning a full range of wrist poses. The right and left wrists were subsequently split into separate volumes, for a total of 900 wrist volume images. Registration results are validated both visually and numerically with the method earlier described.

Table II summarizes the datasets and validation methods used for each experiment.

2.4 Results

1) *In vitro* accuracy and robustness experiment

In the *in vitro* accuracy experiment our tissue-classification registration method demonstrated super-resolution accuracy, and generally had smaller translational errors than grey-value registration. For all registrations the mean tissue-classification translational error was less than 0.3mm, compared to a mean grey-value translational error of 0.5mm. The mean rotational error was less than 0.4° in both methods (Fig. 2.8). The grey-value registration results are consistent with those reported by Snel et al. [99]. A one-sample t-test ($\alpha = 0.05$) on the difference between the mean results produced by the two methods confirmed the tissue-classification translational accuracy improvement was statistically significant ($p = 0.007$).

Decreased image resolution affected the accuracy of our method less than the accuracy of the grey-value registration method (Fig. 2.8). Our method introduced a mean translational error of less than 0.4 mm (a 30% translation accuracy decrease when image resolution drops to one third), compared to 0.9mm translation error in the grey-value method (a 80% accuracy decrease when image resolution drops to one third). The mean rotational error increased to 0.6° in the tissue-classification method, and 0.7° in the grey-value method, respectively. A one-sample t-test ($\alpha = 0.05$) on the difference between the mean results produced by the two methods confirmed the tissue-classification translational accuracy improvement was statistically significant ($p = 0.001$).

The last part of the *in vitro* experiment showed that the tissue-classification registration method maintains super-resolution accuracy with perturbations in the optimization start position (Fig. 2.9). The tissue-classification method was practically insensitive to perturbation (less than 0.4mm, 0.6° mean error, 0% degradation with perturbation), while the grey-value method's mean rotational error doubled (0.8mm, 1.5°), as shown in Fig. 2.9, middle column. A one-sample t-test ($\alpha = 0.05$) on the 2mm perturbation results confirmed that the tissue-classification accuracy improvement was statistically significant, with respect to both translation ($p = 0.01$) and rotation ($p = 0.007$). Increasing the amount of perturbation from

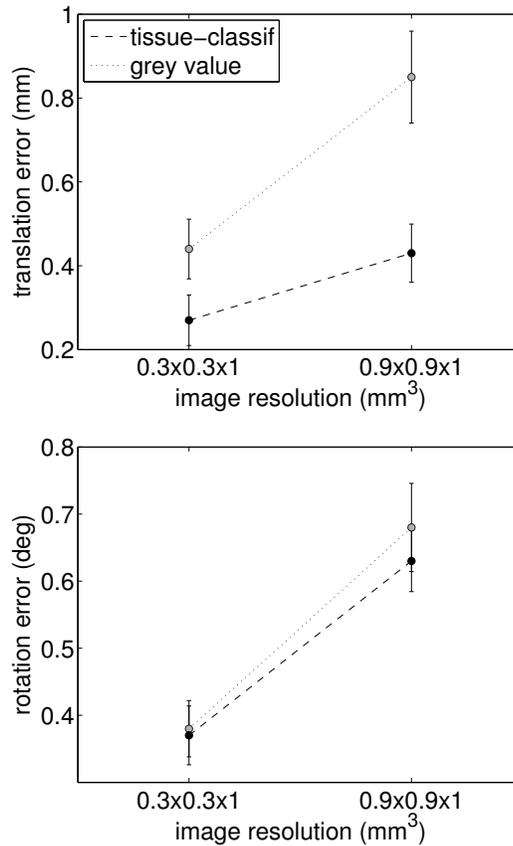


Figure 2.8: The accuracy of tissue-classification and grey-value registration, and the influence of image resolution on both methods. For each resolution, we plot the mean and standard error obtained by registering ten carpal bones across four volume images. The tissue-classification method introduces smaller translational errors than grey-value registration. Tissue-classification accuracy degrades more gracefully than the accuracy of grey-value registration.

2mm to 5mm did not further impact the results of either method (Fig. 2.9). When distortions as small as 2 voxels are present in the optimization start point, the tissue-classification method significantly outperforms grey-value registration, with an average combined (translation and rotation) accuracy improvement of over 74% (Fig. 2.10).

2) In vivo accuracy experiment

The *in vivo* accuracy experiment showed good correlation between the results returned by the three methods tested: tissue-classification, grey-value, and segmentation-based (Fig.

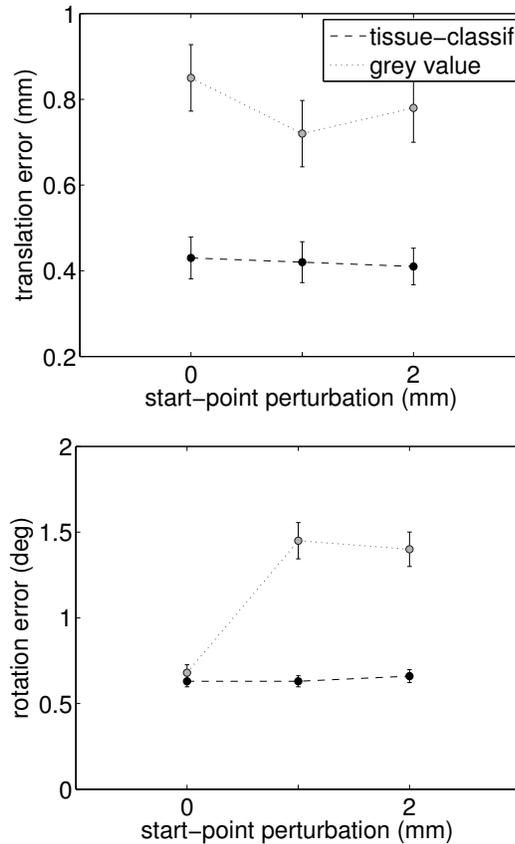


Figure 2.9: The influence of start-point perturbation on tissue-classification and grey-value registration accuracy. The amount of perturbation increases on the horizontal axis from 0mm to 5mm. At each point, we plot the mean and standard error obtained by registering ten carpal bones across four volume images, in five perturbation trials. The tissue classification method is stable with perturbation, while grey-value registration is not. Increasing the amount of perturbation from 2mm to 5mm doesn't impact further the accuracy of either method.

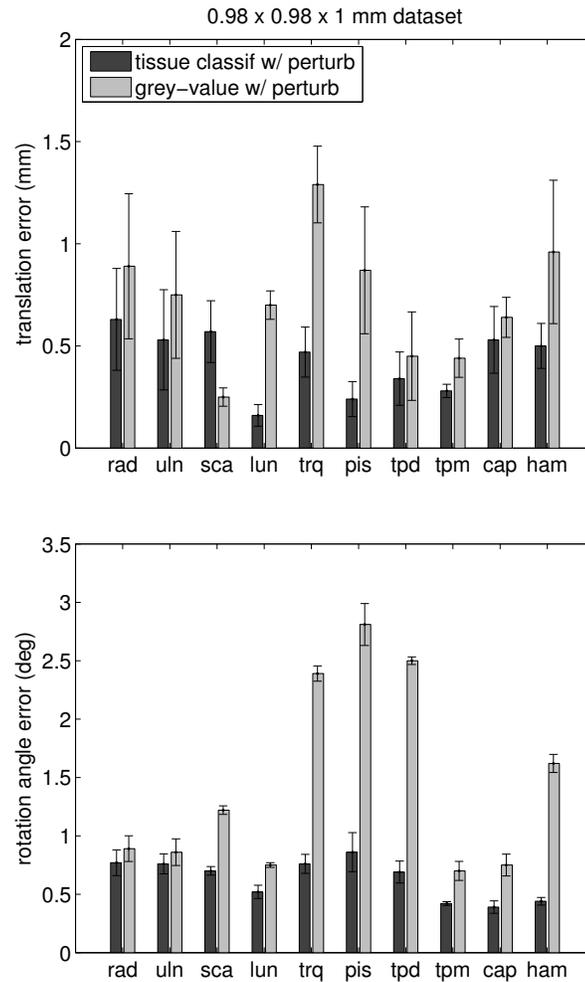


Figure 2.10: *In vitro* tissue-classification registration error and grey-value registration error with a 2mm random perturbation in the optimization start point (top — translation, bottom — rotation). We register each bone across four volume images. Each registration is performed five times, each trial corresponding to a different initial translational perturbation of 2mm in a random direction. For each bone, we plot the mean and standard error thus obtained.

2.11). Since the ground truth is not available *in vivo*, this consistency with the expert segmentation-based results, along with visual inspection, were the best indications of accuracy available. The visual inspection of the grey-value registration results revealed collisions between several pairs of bones. Figure 2.12 illustrates a collision detected between the radius and scaphoid bones. No collisions were detected in the results generated using the tissue-classification method. These results indicate that the accuracy of grey-value registration is insufficient when measuring small features such as inter-bone spacing.

3) In vivo robustness experiment

In the *in vivo* robustness experiment more than 13,500 bones were registered through the tissue-classification method. Visual and numerical validation showed that the method was stable with both object poses and human subjects. In less than 0.1% of cases (8 bone instances) numerical validation indicated suspicious fit values. For each of these cases, further visual inspection revealed that the abnormal fit values resulted from errors in the scanning procedure, wherein the respective bone was only partially included in the target scan. With occasionally as much as half of a bone missing from the volume image, visual verification showed that the bone was still being registered correctly using the information available.

2.5 Discussion

Our results indicate that tissue-classification registration consistently attains subvoxel accuracy. The method maintains subvoxel accuracy despite decreasing image resolution, and is stable with perturbations in the initial optimization start position. Furthermore, visual and numerical validation during clinical application (*in vivo* robustness experiments) shows that the method is robust with varying object poses and subjects.

Tissue-classification should be given credit for our method's accuracy, because the object boundary estimated through tissue classification has super-resolution accuracy. The comparison with grey-value registration shows that without the super-resolution boundary estimate, the matching process is effectively reduced to using voxel-wide estimates, which

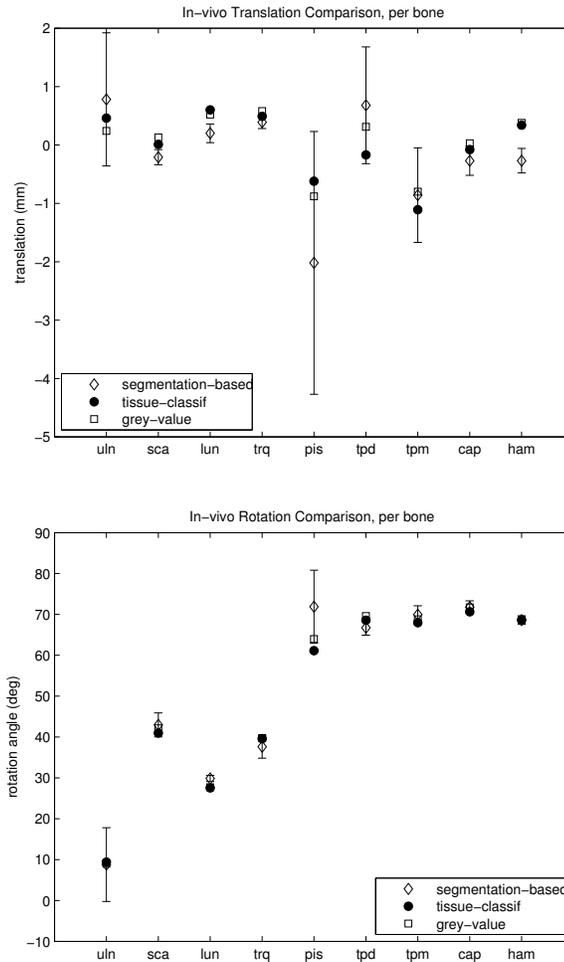


Figure 2.11: Tissue-classification registration results versus segmentation-based registration and grey-value registration results. For each bone, we plot the mean and standard deviation obtained by manual registration in a total of twelve runs, the tissue-classification registration result, and the grey-value registration result. Both tissue-classification and grey-value registration results are generally within one standard deviation of the mean expert segmentation-based registration results.

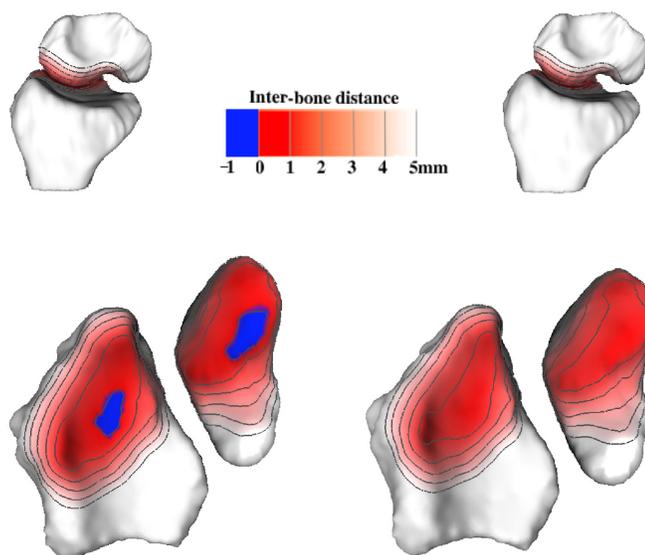


Figure 2.12: Visual inspection shows collisions between bones registered using grey-value registration. (Left) *In vivo* grey-value registration of the radio-scaphoid joint versus (right) tissue-classification registration of the same joint from dataset C. Bones are color mapped and contoured according to the distance between bones post-registration. The saturation of color on bone surfaces represents the distance to the nearest point on the opposite bone. Contour lines are drawn at 1mm intervals. (Top) Bones in their correct anatomical context — note the two registration methods yield similar bone poses. (Bottom) Bones rotated to show articulated surfaces more clearly. Blue corresponds to negative inter-bone distances, indicating collision in the grey-value registration result.

results in lower accuracy. The resulting distance field representation further ensures that registration results converge to the same value, regardless of perturbations in the start position. The distance field acts as a gradient guiding the search to the global minimum. Without this gradient, the search can easily be trapped in local nearby minima, as our *in vitro* perturbation-stability experiments with grey-value registration show.

The comparison with grey-value registration reveals why tissue-classification maintains subvoxel-accuracy with decreasing image resolution. When the sampling step is sufficiently small with respect to the desired features, the raw volume image often offers rich information: object texture-patterns can be implicitly identified and used in the matching process. In this case, *more* raw information can be *better*. However, as the features of interest become smaller, imaging noise effectively blurs the informational content of individual voxels (Fig. 2.13). In this case, *quality* processed information — like accurate boundary-estimates obtained through unsupervised learning — begins to matter. As long as we guarantee super-resolution boundary estimates, registration accuracy stays within the subvoxel range.

In the process of developing our super-resolution accurate registration method, we have proposed new computational methods that are applicable to a broader scope of medical image processing. Our technique uses neighborhood information throughout a volume to generate localized distance fields directly from sampled datasets. No feature points need to be pre-segmented per scan or subject in order to generate the distance fields. Distance fields have been used before to expedite registration processes [110], however no previous work has generated these distance fields with super-resolution accuracy. Recovering material boundaries from sampled datasets, as well as generating distance fields once geometric models have been extracted, is a research topic in several fields [19, 115, 73]. We build on the work of Laidlaw et al.[62], who use Bayesian probability theory to classify accurately tissues in medical volume images. We are not aware of other work in generating super-resolution accurate distance fields directly from sampled datasets. Although we limited the distance field computation in our experiments to a five-voxel band around material boundaries, this distance computation could be either performed or propagated beyond this threshold [73]. However, we showed that the hierarchical approach described

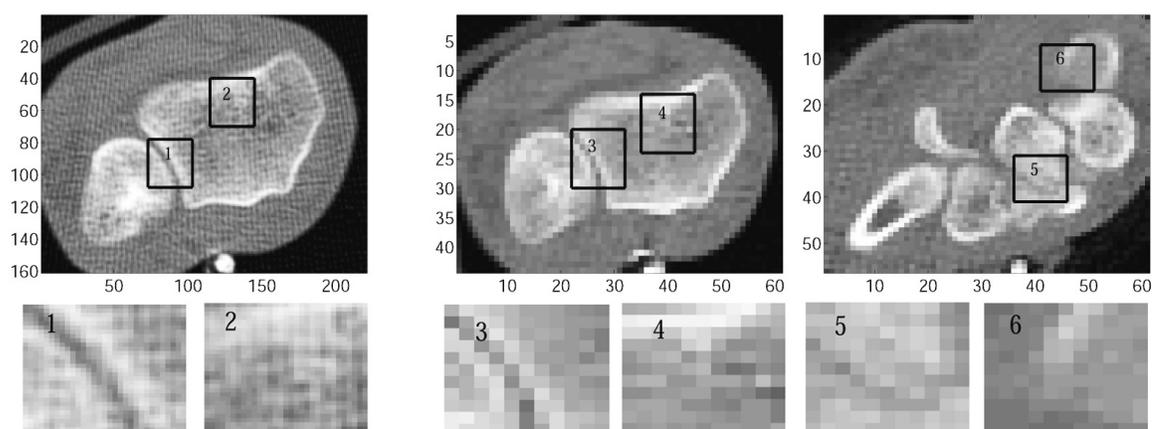


Figure 2.13: Image resolution impacts voxel information-value: as resolution decreases, ‘more’ raw grey-values may deliver less information than ‘fewer’ information-enhanced voxels. From left to right: one slice from a high-resolution CT image (0.3 x 0.3 mm), and two slices from a low-resolution CT image (same wrist, 0.9 x 0.9 mm). Each area of interest is detailed in the bottom row. Inter-bone distances that are narrow in high-resolution images (Box 1) become a mere voxel wide in low-resolution images (Box 3). Boxes 2 and 4, note the differences in bone tissue texture; high-resolution images reveal a distinctive texture pattern (Box 2), while detail is lost in low-resolution images (Box 4). Box 5, note how boundaries are blurred due to partial volume effects, and Box 6, note the soft or diffused bone boundary. As shown in our *in vitro* experiments, accurate boundary estimates based on information throughout the scan contribute more information than such collections of unprocessed grey-values.

in section III.C can overcome potential capture region limitations without the additional cost of extended distance computations.

Our matching procedure employs a new similarity measure that, unlike chamfer matching or iterative closest point measures, incorporates distance field knowledge from both the reference and target images. The measure also accounts for partially-scanned objects. We believe this measure can improve the stability of other registration procedures with respect to errors introduced both by geometric model segmentation and partially scanned objects. While in recent literature a large number of sophisticated similarity measures have been proposed [114], we note these measures were developed mainly for intermodality registration, and likely do not have strong advantages over sums-of-squared-differences or root-mean-square matching when applied to images obtained through the same modality. In our tissue-classification registration approach, sum-of-squared-differences is in fact the natural similarity choice, since the reference and target distance-field intensities corresponding to registered points stand, by construction, in an identity relationship.

The tissue-classified distance field approach helped us identify a motion-directed multi-object hierarchy in the wrist-joint case. The potential physiological implications of this hierarchy are beyond the scope of this chapter but the hierarchy enabled our registration method to trace motion between wrist poses that were not within each other's capture region. A similar approach is likely to boost the capture-range capabilities of other local-search registration methods.

While accurate, our tissue classification method poses computational challenges. Registering a series of twelve wrist volume images (fifteen objects per pose) takes twenty minutes on a multi-node i686 cluster running Linux (AMD *Athlon*TM XP 2700+). We emphasize however that the focus of our work is boosting accuracy, and not minimizing running time. As no user interaction beyond the initial segmentation and potential final visual validation is required, the registration is, after all, run off-line.

While our approach addresses successfully some common problems found in medical image processing, it also relies on several assumptions. Our tissue-classification procedure assumes a simple tissue-structure, in which hand soft-tissue is assimilated with marrow-tissue. While modeling a more complex structure (for example, one that would distinguish

between trabecular and cortical bone) would likely further improve our method’s accuracy, our results indicate the simple model suffices for subvoxel-accurate registration. In the validation phase, we assume implicitly that *in vitro* accuracy is an indicator of *in vivo* accuracy. We note that, in fact, in our *in vitro* experiments the tissue-classifier misclassified several voxels. These misclassifications can be attributed to the differences between the material composition of a fixed specimen encased in resin with all soft tissue removed, and the *in vivo* tissue model we assumed. While the inability to determine the *in vivo* ground truth makes it difficult to compare *in vivo* and *in vitro* accuracy, it is reasonable to assume that errors in the classification process reflected negatively on the registration results. We speculate in this view that our *in vivo* results, in fact, surpass the accuracy of our *in vitro* results. Last, but not least, we note that we interpret collisions or the lack of collisions in our *in vivo* experiments not as a method validation measure, but as an indicator that error size matters.

The accuracy results we obtained challenge opinions widely-held about the superiority of voxel-property-based over segmentation-based registration methods [66]. Segmentation-based registration aligns structures such as feature points, principal axes, moments, curves, or surfaces. Segmentation-based methods may use deformable models, in which case an extracted structure from one image is elastically deformed to fit the second image. They may also use rigid models, in which case the same structures are extracted from any images to be registered. In contrast, voxel-property-based registration methods operate directly on the image grey values, without prior data reduction by the user or segmentation: the full image content is used, via different paradigms such as neighborhood statistics, cross-correlation, Fourier analysis, optic flow, throughout the registration process. Voxel property-based methods integrate segmentation and registration and are thus likely to generate accurate results; these methods have been able to generate more accurate results than segmentation-based methods in several instances [66].

Our tissue-classification registration can be regarded as a hybrid method, combining the strengths of the two approaches — voxel-property and segmentation-based. While our matching procedure evaluates the correlation between source and target volume image

values, without explicitly extracting the target object boundary (like voxel-property methods do), through tissue-classification we nevertheless discard the original volume data that doesn't convey object-boundary information (like most segmentation-based methods do). The recovered object boundary can be thought of as having super-resolution, in that through our approach boundaries are detected with higher accuracy than an image's resolution allows. Nevertheless, due to image noise and modeling assumptions, this boundary is not perfect. Should one be able to recover more accurate object boundary information, generate a distance field from that boundary, then perform distance-field matching as described in Section 2.2, registration accuracy would only increase. This suggests that, contrary to current knowledge, in the long run segmentation-based methods have the potential to surpass the accuracy of voxel-property methods.

2.6 Conclusion

We presented in this chapter a novel intra-subject method for subvoxel-accurate registration of objects from CT volume images. Results show average accuracy improvements of 74% over grey-value registration. The method is of particular interest to applications where collections of tightly packed, small objects need to be registered. To this end, we showed in a wrist data application that earlier registration methods can introduce false inter-object collisions, while the new method does not.

The tissue-classification registration method maintains subvoxel accuracy with decreasing image resolution, and is stable with perturbations in the initial optimization start position. The method is also stable with respect to partially-scanned objects, and with varying object pose and subject. Our approach should be of interest to any registration applications where super-resolution accuracy is desired.

Chapter 3

Modeling Ligament Tissue from Bone Surfaces and Motion

3.1 Introduction

We propose a method for modeling inter-bone joint space areas and ligament paths in articulations. Inter-bone joint space areas define the cortical surface where bones articulate with each other. Modifications in inter-bone joint space areas and ligaments correlate with numerous joint-related post-trauma disabilities and various degenerative diseases, yet little information about the nature of these modifications is currently available. Most articulation and soft tissue studies are performed either *in vitro* or during clinical interventions, and thus reveal little information on potential modifications of soft tissue biomechanics due to injury or disease. *In vitro* specimens illustrating a specific trauma or disease are rarely available; invasive studies alter inevitably joint kinematics and thus introduce false modifications. Although *in vivo* 3-D techniques for studying the structure and kinematics of joint were recently introduced [79, 35, 38, 111], they do not attempt to capture more subtle details such as potential soft-tissue constraints or modifications in articulation. Our method successfully identifies and highlights *in vivo* and non-invasively potential focal (i.e., localized and well-defined) changes and soft-tissue constraints in articulations.

In our approach, the structure and kinematics of an articulation are determined from

segmented CT volume images. Bones in the joint are modeled further both implicitly, as *scalar distance fields*, and parametrically, as *manifold surfaces*. These two types of representation have complementary strengths for different types of calculations. Manifold surfaces provide an accurate, smooth, and locally controllable representation of the bones [47]. Distance fields on the other hand, have important advantages for geometric computations such as fast distance calculation, collision detection, and inside-outside tests [40]. Distance fields computed from the parametric representation provide the support for estimating inter-bone joint space areas. Once inter-bone joint space areas are calculated, focal changes in the articulation are evaluated by comparing the area and location of the bony contact.

We assess potential soft-tissue constraints by calculating the minimum ‘length’ of ligaments as a function of bone kinematics. Ligament paths are also modeled based on the distance field representation. We model ligaments as shortest paths between ligament insertion points — the points at which a ligament is anchored to bones; these paths are constrained to avoid bone penetration, and can be similarly constrained to avoid cartilage penetration. Our model takes into account the ligament fiber orientation, the location of the ligament insertion points, and the locations of adjacent bones. The ligament model reported here is based solely on joint geometry.

We demonstrate our method by applying it to data collected from both forearms of a volunteer diagnosed with a malunited distal radius fracture in one forearm. The distal radioulnar joint (DRUJ), a complex joint involved in forearm rotation, comprises the two forearm bones (radius and ulna — Fig. 3.1) and a number of ligament and cartilaginous complexes. Forearm injuries involving the DRUJ often result in a significantly decreased range of rotational motion, decreased grip strength, and loss of wrist motion. The symptoms can be disabling, especially in physically active individuals or when the pathology affects a work-related activity.

Altered soft tissues and focal changes in the DRUJ articulation may be responsible for the abnormal functioning of the forearm in the absence of evident bone damage, as a recent study suggests [77]. We show that our contact-area and ligament-length model gives unexpected insight into the biomechanics of the forearm and, more importantly, reveals significant differences between uninjured and injured articulations at the DRUJ. Results

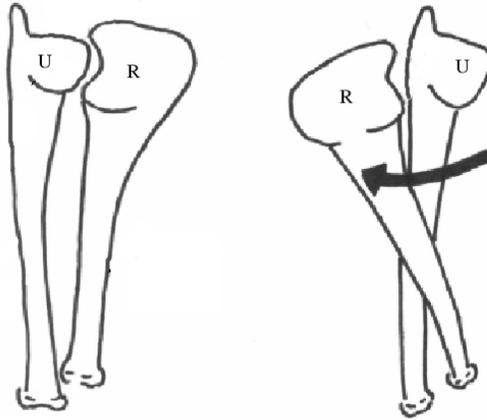


Figure 3.1: The DRUJ comprises the two forearm bones — radius (R) and ulna (U). The wrist is at the upper extremity of the drawing. During forearm rotation the DRUJ goes from supination (left) to pronation (right).

indicate that our method could be useful in the study of the normal anatomy and kinematics of complex joints like the wrist and may also have applications to the study of other joints like the knee or the elbow.

3.2 Related Work

Several approaches to modeling joint surfaces are known; thin-plate splines [20], B-splines [9, 10], and piecewise patches [92] are among them. These methods suffer from problems such as lack of generality, lack of C^2 continuity, and difficulty in enforcing boundary constraints. Our parametric model for bone surfaces is based on manifolds [48].

Distance fields have been used in robotics [55, 65] and computer graphics [40, 44, 21, 87, 49]. Although for the results reported in this chapter we used a brute force approach to generate distance cuboids from the manifold representation, faster techniques such as level set-based methods [81, 96, 97] are available.

Searching for shortest paths in spaces with obstacles is a classical problem in robotics.

A survey of the substantial literature on the shortest-path problem can be found in [117]. Solutions are based on computational geometry methods [91, 83, 25, 11], graph search based algorithms [57], and differential geometry and hybrid techniques [56, 15], depending primarily on the assumed structure of the search space (polyhedral or continuous surfaces). Our technique belongs to the differential geometry and hybrid category. In general, methods in this class generate shortest paths on surfaces. These methods assume a continuous representation of surfaces and are therefore more accurate, although they yield paths that are only locally optimal. Our work extends this approach to 3D spaces with continuous surface obstacles.

The two scalar data visualization techniques we use, color mapping and isocontouring, are well known scientific visualization techniques [72].

Studies of distal radioulnar ligaments are performed in general on cadaver uninjured wrists [36, 5, 94]. A clinical *in vivo* study involving surgery was performed by Kleinman et al. in 1998 [58]. To our knowledge, no *in vivo* noninvasive studies of the distal radioulnar ligaments have been done.

3.3 Materials and Methods

Figure 3.2 depicts our method pipeline. In the first phase, image volumes of the wrists in multiple poses are acquired with a CT scanner (section 3.3.1). From these images bones are manually segmented and further modeled as distance fields and manifold surfaces (section 3.3.2). Kinematic information is recovered via surface registration of the bones (section 3.3.3). Inter-bone joint space areas and ligament paths are estimated using both bone representations (sections 3.3.4 and 3.3.5). We repeat the contact-area and minimum-path computation over all joint poses for a given volunteer. Finally, inter-bone joint space areas and ligaments of the injured and uninjured forearm of the volunteer are compared (section 3.3.6).

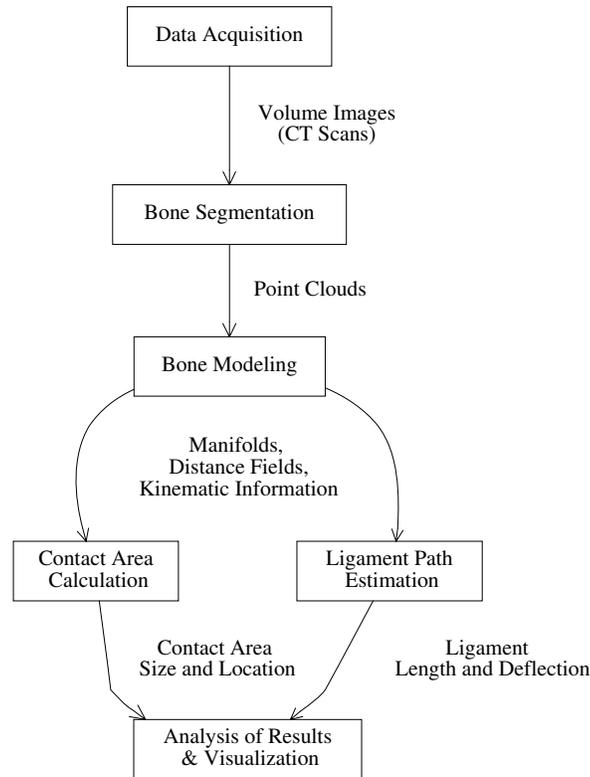


Figure 3.2: Method pipeline for measurement of inter-bone joint space areas and ligament paths in joints. Point clouds corresponding to bone surfaces are segmented from CT volume images. Bones are further modeled as both distance fields and manifold surfaces. From the ligament-path and contact-area models we extract information characterizing the articulation that is further analyzed and presented to the user.

3.3.1 Data Acquisition

CT volume images of both wrists were obtained simultaneously with a GE HiSpeed Advantage CT scanner. Scout and reference scans were performed with the forearm and wrist in the neutral position. Additional scans were performed with the forearm at 30, 60, and 90 degrees of both pronation (i.e., forearm with the palm facing downwards) and supination (i.e., forearm with the palm facing upwards). In the forearm with limited mobility (decreased range of pronosupination), scans were made at 30 degree intervals (above), and then at the maximum rotation that could be comfortably achieved. Approximately 45 1.0 mm CT slices were acquired at each position.

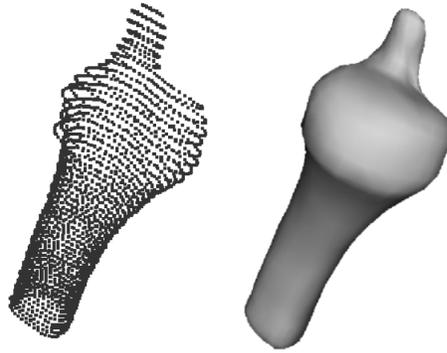


Figure 3.3: Manifold surface representation of bones. Left: segmented point cloud corresponding to the ulna. Right: parametric (manifold) model of the same bone.

3.3.2 Bone Segmentation and Modeling

Points corresponding to the outer bone cortex were manually segmented from each CT slice and grouped to form a separate 3D point cloud for each bone. We reconstruct a bone surface by fitting a manifold surface to the corresponding cloud of 3D points [47] (Fig. 3.3); the result is a smooth, locally parameterized, C^2 continuous surface. The overlapped structure of the manifold-surface representation, which is essentially inspired by differential geometry, has several advantages including flexibility in shape adjustments without costly constraints, and smooth transitions and uniformity among patches.

The manifold model addresses difficulties introduced by the CT scanning process, such as dense sampling along sparse contours and noise [47]. The manifold model is analytic and can therefore be sampled at any resolution to produce smooth distance maps. High-resolution smooth distance maps are necessary in order to build ligament paths, as discussed in Section 3.3.5.

By convention, we reflect left forearm data in order to directly compare it with right forearm data. The mirroring operation is purely mathematical and does not affect the data; it merely allows easier comparisons.

Modeling inter-bone joint space areas and ligament paths requires bone-to-bone distance information (sections 3.3.4 and 3.3.5). The manifold surfaces provide accurate, smooth but computationally expensive distance information. We combine the manifold

representation with interpolated distance fields, which are slightly less accurate but more intuitive and much faster.

Distance fields for each bone are computed using the reconstructed manifold bone models. A distance field is a scalar field that specifies the signed distance from a point to the bone surface (Fig. 3.4). 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.

The distance field is computed from the manifold representation as follows: given a point P in space, the closest point Q on the manifold has the property that the surface normal at Q points in the direction $P-Q$. We find an approximate guess for the point Q by finding the closest point Q on the manifold mesh, then perform a gradient descent to find the Q that meets the above criteria. The inside-outside test simply involves counting the number of intersections with the manifold mesh of any ray from P [8].

In order to increase the speed of lookup operations, the distance fields are sampled on a regular grid. We call the result a *distance cuboid*. The distance cuboid can be regarded as a scalar data set sampled over a regular 3D grid surrounding the bone. Distances to the bone surface are known exactly at grid nodes. Within a grid cell, distances to the bone surface are obtained via tricubic B-spline interpolation of the distance values at grid nodes.

The double bone representation — manifold surfaces and distance cuboids — enables us to perform further joint-related computations, such as calculation of inter-bone joint space areas (section 3.3.4) and estimation of ligament paths (section 3.3.5).

3.3.3 Recovery of Bone Kinematics

Recovering the bone kinematics enables us to analyze our inter-bone joint space area and ligament measurements as functions of wrist motion. Motion of the radius with respect to the ulna was determined for each scanned wrist rotation position. First the ulna bone was registered with respect to its neutral position to account for global changes in forearm positioning. Next, the relative motion of the radius with respect to the ulna was calculated. The forearm data modeled and analysed in these experiments was collected and preprocessed

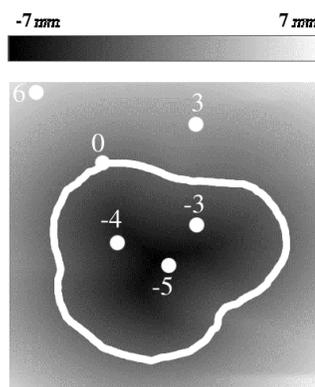


Figure 3.4: Distance field representation of bones: horizontal 2D section through a signed distance field (ulna). The contour corresponds to the boundary of the bone. Sign distinguishes the inside from the outside of the bone: negative values are inside the bone, positive values are outside the bone, zero values are on the bone surface. The dark area is the inside of the bone.

several years before the development of the motion tracking method described in Chapter 2. Therefore in the experiments reported in this chapter registration was accomplished via a manually-aided surface-distance-minimization algorithm [77]. Bone kinematics were reported in a standard anatomic coordinate system defined in the distal ulna as follows: the x -axis was directed proximally along the shaft of the ulna and defined by the centroids of the ulnar bone cross-sections, the z -axis was in a palmar direction and defined to be perpendicular to a plane that passed through the x -axis and the tip of the ulnar styloid, and the y -axis was constructed perpendicular to both the x - and z -axes. The origin of the coordinate system was defined by the intersection of the x -axis with the (ulnocarpal) articular surface of the head of the ulna (Fig. 3.5).

3.3.4 Inter-Bone Joint Space Area Calculation

The inter-bone joint space area is defined as the cortical surface area on the bone that is less than a prescribed threshold distance (typically 5 mm) from the cortical surface of a neighboring bone. Estimating inter-bone joint space areas requires computation of inter-bone distances within the joint.

Once distance cuboids are generated, we calculate the distance from an arbitrary point,

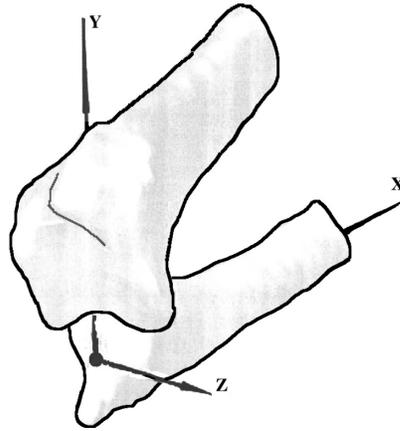


Figure 3.5: Anatomic coordinate system defined on the ulna. The location and orientation of the x -axis were generated from the cross-section centroids of the ulna, while the z -axis was defined to be perpendicular to a plane that passed through the x -axis and the tip of the ulnar styloid. The y -axis was constructed perpendicular to both the x - and z -axes.

p , and a bone surface, b , as follows. Each bone surface has a surrounding distance cuboid f_b . We remind the reader that a distance cuboid can be regarded as a sampled dataset stored over a regular 3D grid; distance values are known exactly at grid nodes, and computed via interpolation inside grid cells.

The point p can be inside or outside the distance cuboid f_b . We make sure that areas of interest (i.e., articulated surfaces) are well within the distance cuboid. For simplicity, Figure 3.6 illustrates the procedure in 2D. We evaluate two cases to find the distance:

p is inside f_b : we look up f_b for p

p is outside f_b : we first find the distance to the nearest point p' on the boundary of f_b . We then add it to the distance value acquired by looking up f_b for p' . Since points outside the distance cuboid are of little interest (i.e., they are far away from articulated surfaces), this distance sum is an acceptable approximation.

With this procedure we find distances from every vertex in the surface model of one bone

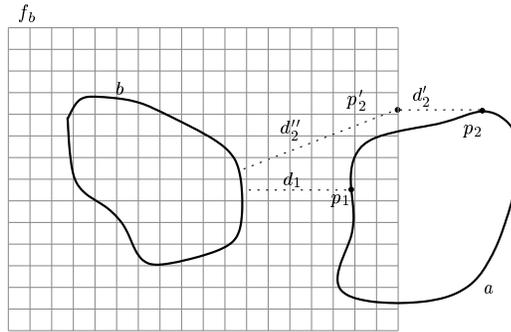


Figure 3.6: 2D illustration for obtaining distances from points p_1 and p_2 to bone b . f_b is the distance cuboid for bone b . Shortest distance values to bone b at the grid intersections are known. We use tricubic interpolation to determine distance values within the grid. Since p_1 is inside the cuboid, the distance from p_1 to b is equal to $f_b(p_1) = d_1$. For p_2 , we first find the distance to the closest point p'_2 in the distance cuboid and then the distance between p_2 and b is approximated as $d'_2 + f_b(p'_2) = d'_2 + d''_2$.

to neighbors of interest.

Using the inter-bone distance we compute isocontours on the inter-bone joint space area, each contour showing where the distance map is equal to a constant distance. For efficient computation, we assume that the distance map is linear over the triangular faces that comprise the surface of the bone and thus the equal distance contours are straight line segments over each triangle. If the distance value of a contour is within the range of the distance values at the vertices, a contour line segment is generated over the triangle.

Figure 3.7 shows typical inter-bone joint space areas in the DRUJ; the joint was exploded to show the articulated surfaces more clearly. The color on bone surfaces codifies the distance to the nearest point on the opposite bone; darker regions are closer.

We characterize the inter-bone joint space area by its size and by the location of its centroid. The size is the area of the surface triangles within the 5 mm contour. The location of the centroid is described in cylindrical coordinates with respect to the same standard coordinate system used to report bone kinematics.

It is important to note that the articular contact calculated here is an estimate of joint contact based upon the distance between cortical bone surfaces. Cartilage thickness, bone and cartilage deformation and stresses in the tissues were not considered in this study.

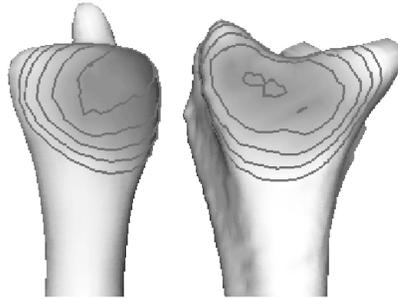


Figure 3.7: inter-bone joint space areas in the DRUJ. Bones are color-mapped and contoured. The color saturation on bone surfaces indicates the distance to the nearest point on the opposite bone; darker regions are closer. The joint is exploded to show the articulated surfaces more clearly. The maximum distance visualized is 5 mm; contour lines are drawn at 1 mm intervals.

3.3.5 Ligament Path Estimation

We can also use the double bone representation to construct ligament paths. We use anatomical landmarks to manually identify the insertion points (the points where the ligament is anchored to the bone) of a given ligament on the bone surface. We generate plausible ligament paths as shortest paths between insertion points, constrained to avoid bone penetration.

We build shortest paths via an optimization approach that exploits the distance field representation of the bones. Unlike graph-based minimum-length path approaches, this technique deals effectively with a large number of bone model vertices without requiring expensive restructuring — in terms of memory and time — of the search space. The resulting paths are also more accurate than those generated, for example, by graph approximation algorithms, as the method allows a large number of path control points and recovers gracefully from obstacle penetration. We begin the description of the algorithm with a simplified 2D example, shown in Fig. 3.8. Here we are required to find a shortest path between two points p_0 and p_n that does not penetrate the 2D obstacle on the right.

We start by attaching a local 2D coordinate system to the obstacle, so that the origin of the system is at p_0 and the x axis is the line defined by p_0 and p_n . We consider $n - 1$ points in addition to p_0 and p_n , equally spaced on the p_0p_n segment. We reformulate our

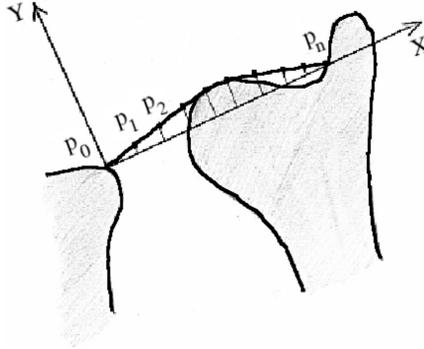


Figure 3.8: Shortest path between two points p_0 and p_n (2D case); the path must not penetrate the 2D obstacle on the right. Following the optimization approach, the points p_1 to p_{n-1} , initially equally distributed on the p_0p_n segment, increase their y coordinate so that the nonpenetration constraint is satisfied.

problem in the following terms: “Find the coordinates of the $n - 1$ points so that the length of the path $p_0p_1p_2\dots p_n$ is minimum and the height of each point with respect to the obstacle surface is nonnegative.” If we fix the x coordinates of the points so that they are initially equally spaced on the p_0p_n segment, our problem amounts to minimizing the Euclidean length of the path over the y_i coordinates of the points:

$$\begin{aligned} \operatorname{argmin}_{y_i} \sum_{i=0}^{n-1} \sqrt{(x_{i+1} - x_i)^2 + (y_{i+1} - y_i)^2} = \\ = \operatorname{argmin}_{y_i} \sum_{i=0}^{n-1} \sqrt{\text{const} + (y_{i+1} - y_i)^2} \end{aligned} \quad (3.1)$$

subject to $f_b(x_i, y_i) > 0, i = 0..n - 1$

where $x_{i+1} - x_i = \text{const}, i = 0..n - 1$

The formulation described above extends to 3D, where we optimize over both the y and z coordinates of the points:

$$\operatorname{argmin}_{y_i, z_i} \sum_{i=0}^{n-1} \sqrt{\text{const} + (y_{i+1} - y_i)^2 + (z_{i+1} - z_i)^2} \quad (3.2)$$

subject to $f_b(x_i, y_i, z_i) > 0, i = 0 : n - 1$

The extension of the algorithm to any number of obstacles is straightforward.

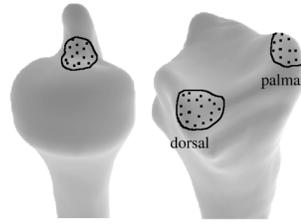


Figure 3.9: Insertion point location: insertion points are chosen manually, based on anatomical information. Points are randomly distributed on the surface of the bones within a circular area with a diameter of 4 mm. Left: insertion site on the ulna. Right: dorsal and palmar insertion sites on the radius.

We use a sequential quadratic programming method [3] to solve the optimization problem. The sequential quadratic programming method is fast and robust and handles both nonlinear objective functions and nonlinear constraints. Although it is a general concern that nonlinear optimizations can become trapped in suboptimal local solutions, in our experience this has not been a problem. We have found that additional iterations of the optimization process with significantly different start positions converge to the same solution. We are currently using three different start solutions:

1. points on the straight p_0p_n line;
2. points on a randomly displaced path;
3. points generated by the procedure in the previous pronosupination position.

The optimization procedure converges to the same solution in all three cases. This outcome is justified by the smooth structure and fine resolution of the search space generated by the distance field representation.

We considered several plausible insertion points for each ligament, as precise information on insertion point location was not available. The insertion points were generated by randomly distributing points around a manually chosen landmark on the surface of the bones, within a circular area with a diameter of 4 mm (Fig. 3.9). The insertions were defined on the ulna at the base of the styloid for both ligaments and on the radius at the

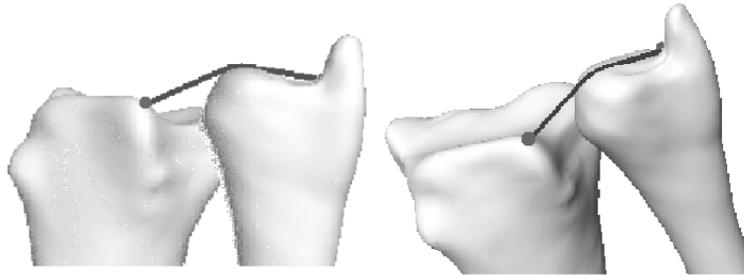


Figure 3.10: Shortest paths (dark gray lines) generated by the ligament model.

dorsal and palmar prominences of the sigmoid notch, respectively. The locations of the insertion sites and the area of insertion were derived from anatomical descriptions in the literature [36, 5, 94, 58]. The results of the insertion point study are presented in section 3.4.

We tried several values for the number of points n . In the DRUJ case, as n approaches 40 the total length of the path converges to a stable value. For this value of n the length of each mini-segment in the path drops below 0.2 mm, which provides sufficient accuracy to detect deflection of the ligament by the bone. Figure 3.10 shows two shortest paths generated with our algorithm.

We characterize the ligament paths by their lengths and their ‘deflection’. Lengths are normalized with respect to the uninjured length in neutral pronosupination. Deflection is defined as the maximum distance across all path points to the straight line defined by the two ligament insertion points.

The ligament-length model reported here is based solely on joint geometry. Structural and material properties of the ligaments were not taken into account in this study. While the paths we generate are not actual ligament paths, they give a useful lower bound on the length of these ligaments and thus help identify potential joint mobility constraints imposed by ligaments.

3.3.6 Visualization and Analysis of Results

The software package we have developed for visualizing the results of our technique consists of C++ and Open Inventor code and runs on the SUN UltraSparc and Windows platforms.

We visualize inter-bone joint space areas 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 (5 mm) are neither colored nor contoured and are shown as white surfaces. Contours and ligament paths are visualized as polylines.

We also analyze the results quantitatively by comparing ligament length, ligament deflection, inter-bone joint space area size, and inter-bone joint space area centroid location between the injured — malunited distal radius fracture — and uninjured forearm of the same volunteer.

3.4 Results and Discussion

Generating inter-bone joint space areas over different forearm rotation positions yields sequences like those in Fig. 3.11. The decreased size and shifted location of the inter-bone joint space area in the injured case is noticeable, especially towards pronation.

Figure 3.12 quantifies the size of the ulnar inter-bone area at a threshold of 5 mm for the volunteer's uninjured and injured forearm. For the uninjured wrist, inter-bone joint space area was positive for a 3 mm threshold as well. For the injured wrist, there were several poses, mostly pronated, in which the 3 mm inter-bone joint space area was absent. Together with the 5 mm inter-bone joint space area changes, this suggests an increased gap between the bones in the injured case.

We measured inter-bone joint space area as a region on the ulnar surface close to the radius; an analogous measure on the surface of the radius can also be defined. We found that the area measure was somewhat larger (10-20%), but followed the same trends as the ulnar inter-bone area. The size difference is consistent with the concave inter-bone joint space area on the radius, which is larger because it is farther from the center of curvature than the corresponding area on the ulna. Measures based on the ulnar area are reported because they reside in the ulnar coordinate system; the ulnar coordinate system was chosen because it is stationary during pronosupination.

Figure 3.13 shows the cylindrical coordinates of the ulnar inter-bone area centroid for

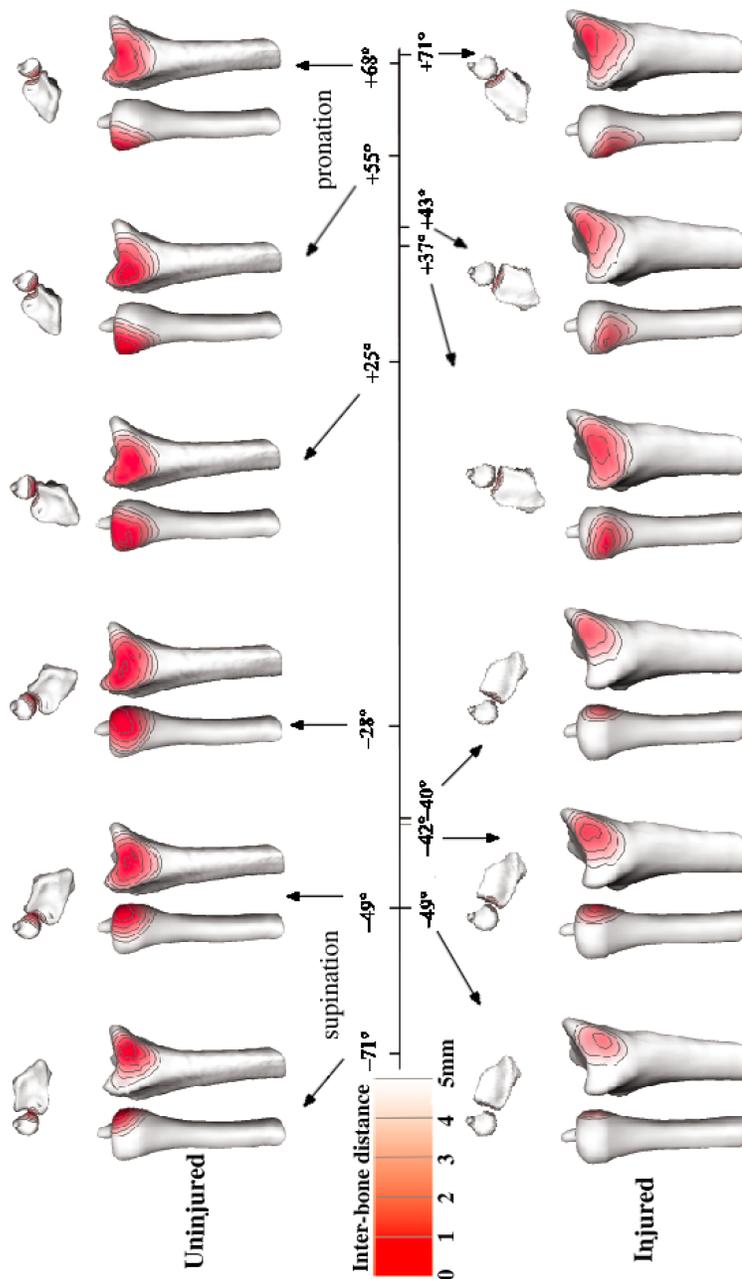


Figure 3.11: Proximal and exploded lateral views of an uninjured and an injured radioulnar joint at six rotation positions. Bones are colored according to the distance between them (the closer they are, the more intense the color). The injured and non-injured views are matched as closely as possible. Note the shift in the location of the inter-bone joint space areas between the uninjured and the injured forearm.

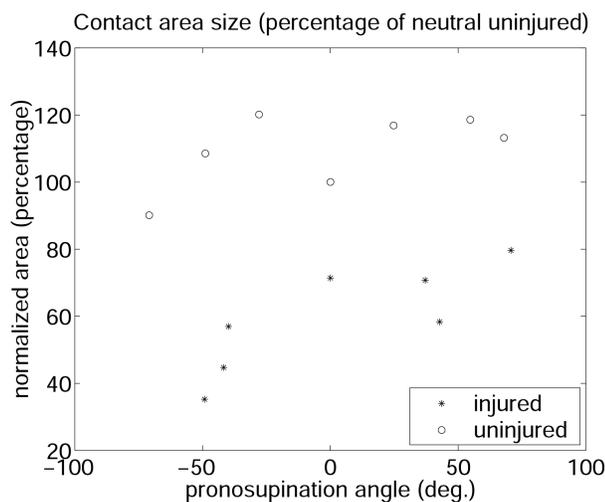


Figure 3.12: Size of the ulnar inter-bone area (5 mm threshold) for both the injured and uninjured forearm of the same volunteer. Areas are normalized by the neutral uninjured area. Pronosupination angles are shown on the x -axis. Note the difference in size between the injured and uninjured forearm.

the uninjured and injured forearm. The increased height coordinate in the injured forearm confirms a shift of the inter-bone joint space area in the proximal direction. The increased distance from the ulnar axis is due to the shift of the inter-bone joint space area on the surface of the ulna to a region of the ulna further from the axis. The angle coordinate plot correlates with the limited range of motion in the injured forearm. The proximal shift in the location of the centroid of the inter-bone joint space area is consistent with the initial diagnosis of radial shortening.

Figure 3.14 shows distal ligament paths generated for the injured and uninjured forearms of the same volunteer. The lengths generated by our approach are similar to those reported in *in vitro* studies; no *in vivo* information is currently available, to the best of our knowledge. Note that the injured forearm presents ligament-bone impingement for both the dorsal and the palmar ligament. No deflection of the ligaments by the bone is present in the uninjured forearm in any of the rotation positions. Figure 3.15 shows the dorsal radioulnar ligament length and deflection corresponding to the entire pronosupination sequence for the injured forearm. We also show the corresponding lengths and deflection computed for

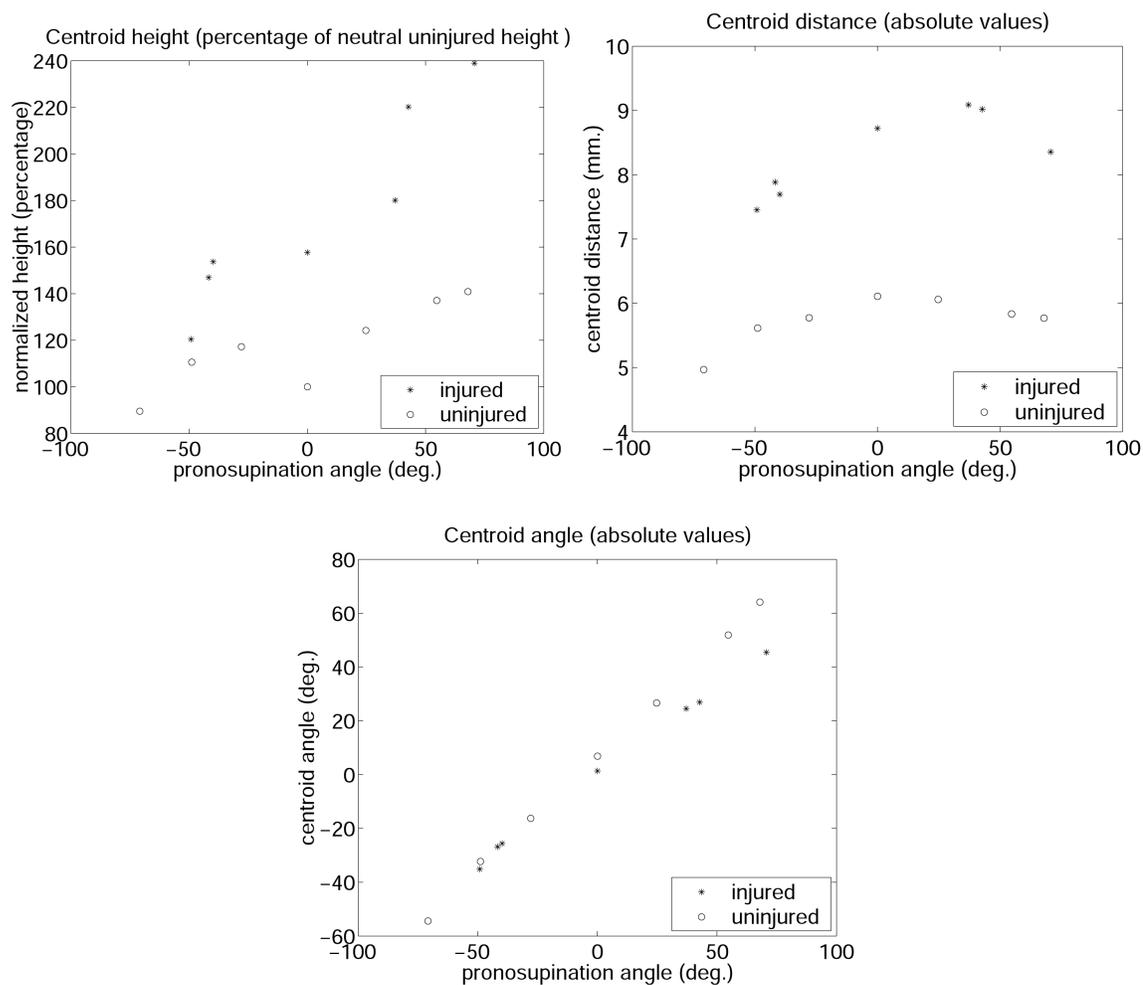


Figure 3.13: Cylindrical coordinates (height, distance, and angle) of the ulnar inter-bone area centroid for the injured and uninjured forearms of the volunteer. Heights are normalized by the neutral uninjured height. Pronosupination angles are shown on the x -axis. Note the difference in height and distance between the injured forearm ligament and the uninjured forearm.

the matching uninjured forearm — note the difference between the two plots. Ligament impingement (measured by the deflection parameter) correlates with ligament path increased length. No ligament deflection is present in the uninjured forearm. The dorsal ligament results generated by displacing the insertion points within the insertion site are plotted in Fig. 3.17. Note that perturbations in the ligament attachment locations do not affect trends

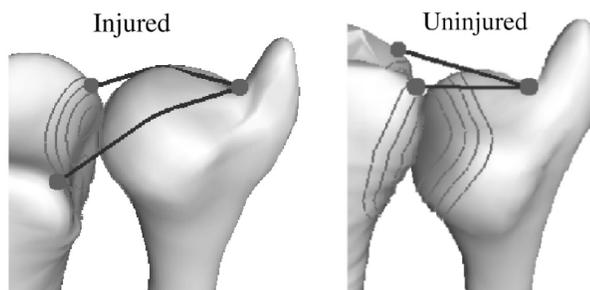


Figure 3.14: Distal radioulnar ligament paths in the injured forearm (left) and in the matching uninjured forearm (right) of the same volunteer. Both forearms are in neutral pronosupination (0° rotation angle). Note the ligament-bone impingement in the injured forearm: both ligaments are deflected by the head of the ulna.

in the comparison measures between the injured and uninjured forearms.

Figure 3.16 shows plots of the palmar radioulnar ligament length and deflection. Although the palmar ligament length plot shows no difference between the injured and uninjured forearm, we note the impingement (deflection) in the injured forearm, lacking in the uninjured case. The palmar ligament results generated by perturbing the insertion points within the 4 mm diameter insertion sites are plotted in Fig. 3.18. Note again that perturbations in the ligament attachment locations do not affect trends in the comparison measures between the injured and uninjured forearms.

The change in the dorsal radioulnar ligament length, but not in the palmar radioulnar ligament length, is consistent with the original malunion (radius tilted dorsally). The change in ligament length and the ligament-bone impingement may be one mechanism for the limitation of forearm mobility.

While a single example cannot distinguish between normal anatomical variation and pathological variation, clinical studies on larger sets of patients may establish or refute a correlation between the differences we found here and the injury. Such studies are beyond the scope of this chapter, but have been published as [29].

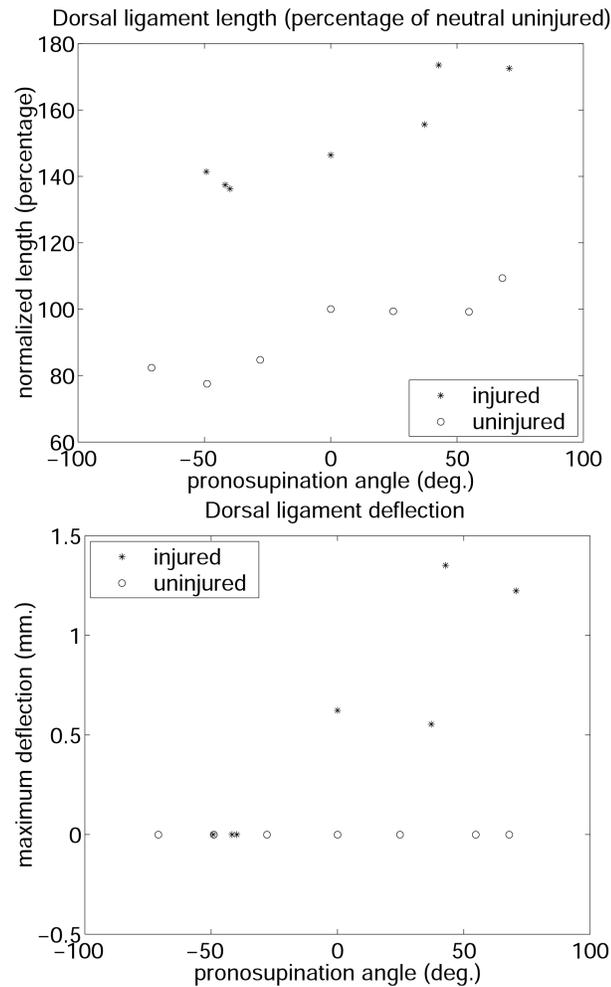


Figure 3.15: Length (top) and maximum deflection (bottom) of a dorsal ligament for the injured and uninjured forearms of a volunteer. Lengths are normalized by the neutral uninjured length. Pronosupination angles are shown on the x -axis. Note the increased ligament length in the injured forearm. Note also that no deflection is present in the uninjured forearm.

3.5 Conclusion

We have demonstrated an *in vivo*, noninvasive technique for modeling the length of ligaments and joint inter-bone joint space areas from bone kinematics and surfaces. Our method uses an implicit model as well as a parametric surface model for each bone. The two types of representation have complementary strengths for different types of calculations.

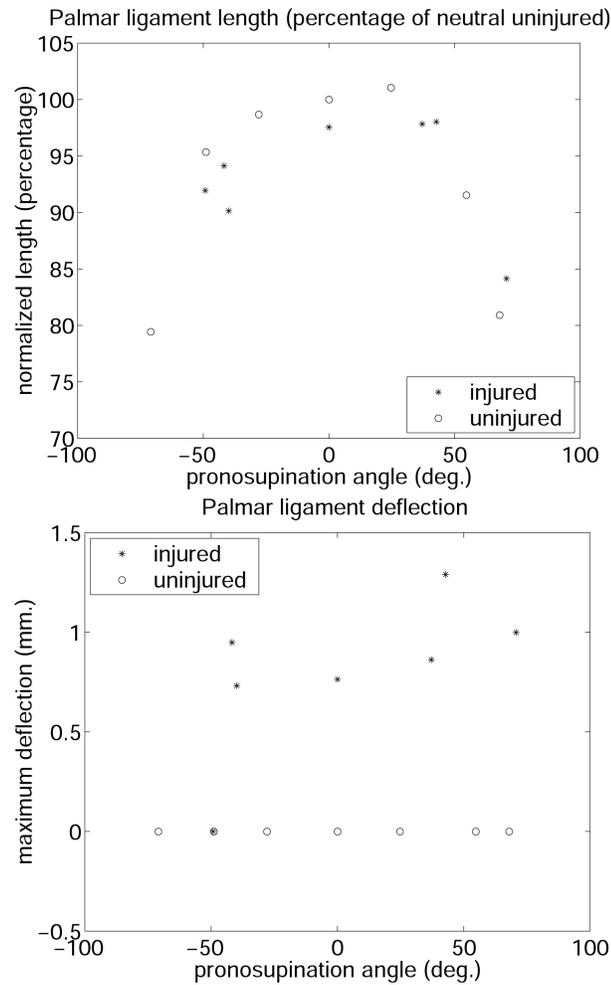


Figure 3.16: Length (top) and maximum deflection (bottom) of a palmar ligament for the injured and uninjured forearms of a volunteer. Lengths are normalized by the neutral uninjured length. Pronosupination angles are shown on the x -axis. Note that no deflection is present in the uninjured forearm.

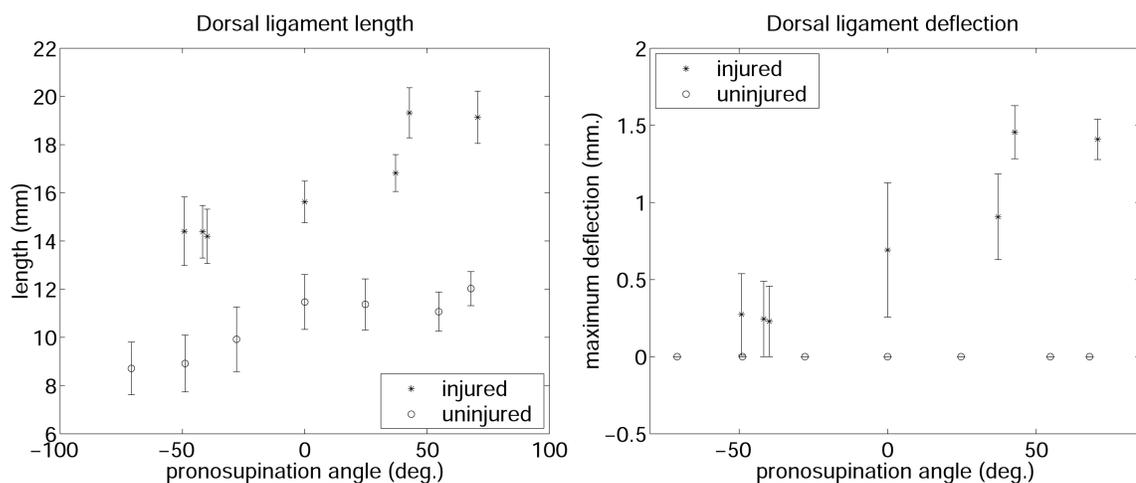


Figure 3.17: The effect of insertion point perturbation on the length (top) and maximum deflection (bottom) of a dorsal ligament for the injured and uninjured forearms of a volunteer (mean and standard deviation calculated over 64 measurements).

The double representation enables us to model secondary types of information from CT data, such as joint inter-bone joint space areas, intra-joint distances, and plausible ligament paths. Our current ligament model could be enriched by considering other intrinsic and extrinsic ligament factors like tissue composition, muscle forces, and joint compression.

In a demonstration on the DRUJ, our approach highlights subtle modifications, otherwise unnoted, in injured wrist kinematics. Although a previous kinematic study [77] on the same data we analyze in this chapter found no significant differences in rigid body kinematics between the injured and uninjured wrist, our method identified potential soft tissue constraints and focal changes in the articulation. The methods presented have the potential to document changes in the joint mechanics that may influence long-term clinical outcome.

Our technique may have applications to the study of wrist disorders such as rheumatoid arthritis, intercarpal ligament tear or attenuation, and carpal-tunnel syndrome. Results suggest that our technique could also be useful in the study of normal anatomy and kinematics of other joints.

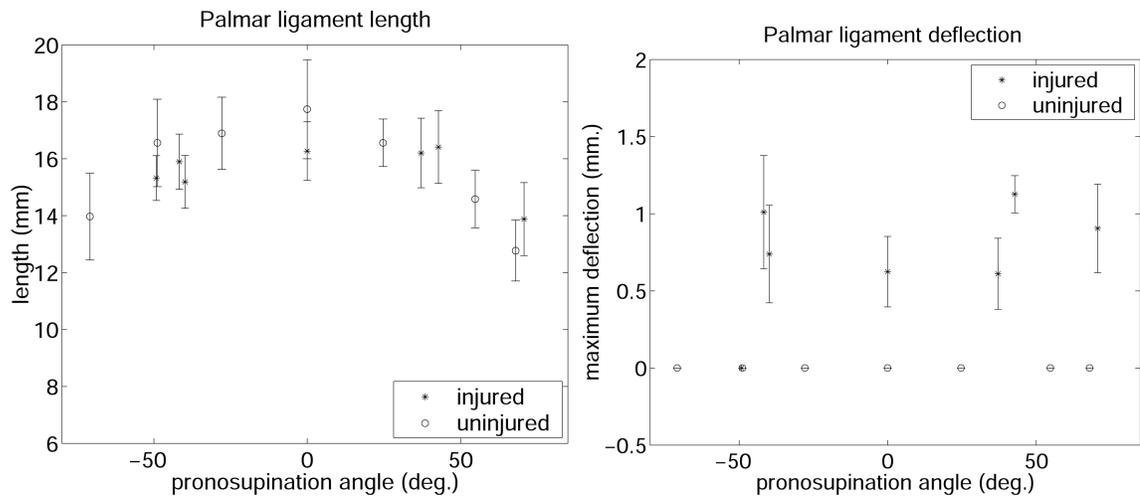


Figure 3.18: The effect of insertion point perturbation on the length (top) and maximum deflection (bottom) of a palmar ligament for both the injured and uninjured forearm of a volunteer (mean and standard deviation calculated over 64 measurements).

Chapter 4

Modeling Articular Cartilage from Bone Surfaces and Motion

4.1 Introduction

Articular cartilage — the living tissue that lines the bony surface of joints — plays an important role in diarthrodial joint motion. Its function is to provide a low friction surface cushioning the joint bones through the range of motion. In other words, articular cartilage is a very thin, slippery shock absorber. Factors such as high impact twisting injuries, abnormal joint anatomy, joint instability, inadequate muscle strength or endurance, and medical or genetic factors can damage the articular cartilage lining. We wish to be able to investigate articular cartilage and to track potential damage progression non-invasively, in live individuals.

However, because it is very thin, articular cartilage may be difficult to image in live individuals and small animals. For example, human wrist cartilage on several of the carpal bones averages 0.5mm in thickness. To generate highly-detailed, unloaded wrist-cartilage geometry, we currently need to extract the bone from the joint, remove all soft-tissue except for cartilage and perhaps ligament insertions, immerse the bone in contrast dye for 24 hours, then μ CT-image the bone. In the resulting μ CT slices, such as the one in Fig. 4.1, cartilage shows as a soft-grey layer surrounding the bone outline.

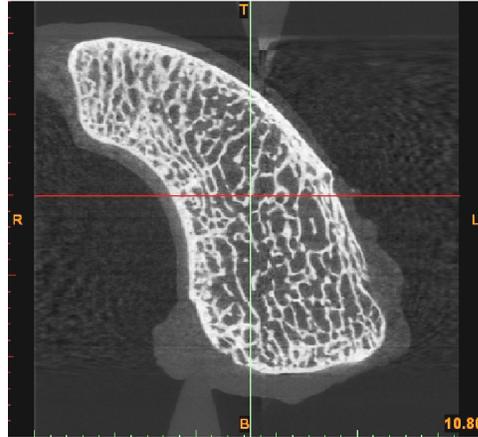


Figure 4.1: 2D slice through a μ CT-volume image of a scaphoid bone and articular cartilage. The imaged cartilage (bone tissue shown in white, cartilage in soft grey surrounding the bone outline) can be represented as a height-field on the supporting bony surface. The height-field is defined as a collection 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.

Because current measuring procedures like the one described above are invasive, their applicability is limited to *in vitro* data. In this chapter I describe a data-driven, subject-specific cartilage modeling approach that allows for the exploration of its functional role non-invasively, *in vivo*. I present a method for estimating subject-specific cartilage maps — location and thickness — directly from *in vivo* kinematic data and computed tomography (CT) volume images. I also introduce a novel algorithm for computing cartilage surface deformations. The resulting cartilage model, a meshless incompressible height-field, captures the physical properties important for estimating the shape, contact area, and deformation magnitude of cartilage at each articulation.

4.2 Related Work

There are two known approaches to generating articular cartilage geometry when direct measurements — in terms of cartilage location and thickness — are not available. Both approaches allow for the estimation of articular contact *in vivo*.

In the first approach (e.g., Thoomukuntla et al. [107]), cartilage is generated as a thin

shell of constant thickness surrounding the bones. Thoomukuntla et al. [107] model carpal cartilage as a uniform 1mm-thick shell wrapping articular surfaces in the human wrist. However, there is no clear clinical evidence that cartilaginous surfaces have in general constant thickness, nor that 1mm would be a good estimate of this thickness in carpal bones, regardless of the individual bone size. Validation against *in vitro* data showed discrepancies between the computed articular contact and contact estimated using pressure-sensitive film.

In the second approach (e.g., Carrigan et al. [24]), articular cartilage geometry is determined from the bone-spacing in one pose. Carrigan et al. [24] grow cartilage by half the inter-bone distance in one reference pose, in the direction of the bone surface normal. However, the cartilage map thus generated is localized to the bony areas that are close one to another in that particular pose. Two adjacent poses may also generate different versions of the map over the same section of the articular surface. Finally, this approach does not account for potential cartilage default deformation in the reference pose. The resulting articular contact was not validated.

Once the cartilage geometry has been either measured directly or generated through one of the approaches above, it can be represented as either a rigid body, a parameterized surface, or a deformable solid. Rigid body and parameterized surface representations [61, 107] allow only for computation of pseudo-deformable contact. In these approaches the nature of the objects being modeled is ignored — for example, volume preservation is not accounted for, and physical contact surfaces are approximated based on the overlap or interpenetration of the two modeled bodies.

In contrast, the deformable solid representation attempts to account for characteristic properties of soft-tissue deformation, such as elasticity or viscosity. Instantiations of this representation are mass-spring systems, generalized particle systems, finite differences and finite element models, and mesh-free models. Early work in physically-based deformable models is surveyed in detail by Gibson and Mirtich [45].

We note that, in general, mass-spring systems can be difficult to fit to physical data and do not preserve volume; they can also be over or under-determined. Generalized particle systems include additional spring forces to prevent shear or bending, and additional non-spring forces to enforce volume preservation. Both mass-spring systems and generalized

particle systems use rough approximations of physical deformation, but are relatively fast to simulate.

In contrast, finite difference and finite element methods (FEM) use a sound derivation of deformation from continuum mechanics. These approaches are appropriate choices for accurately modeling deformations, stresses and strains in complex materials like soft-tissues [108, 41]. However, both finite approaches are computationally expensive. In particular, we note that modeling through FEM simultaneous cartilage contact among many articulated surfaces in order to predict, for example, wrist-joint motion is currently computationally infeasible.

Recent work in deformable solid representations explores quasi-rigid contact through simpler, local, analytical models. The key idea in this approach is that small deformations can be modeled with sufficient accuracy and less computational expense than through finite difference or finite element approaches. For example, Pauly et al. [86] use an approximation to linear elasticity in order to preserve the volume of a mesh-free representation. Their approach is similar to ours in the use of a mesh-free, point cloud surface representation, and the computation of a local response through a surface integral formulation. The advantage of using a mesh-free representation is that contact events can be modeled efficiently and a unique contact surface can be defined without the need for retriangulation. Because Pauly et al. use a linear complementarity formulation to impose volume preservation, their approach can simulate, in principle, effects that are primarily dictated by contact, such as friction. The linear complementarity formulation comes, however, with a higher computational cost. While our approach lacks sophisticated capabilities such as explicit friction integration, it is easier to implement and faster to compute, while it generates sufficiently accurate contact measurements.

4.3 Methods

Our computational approach proceeds in four steps. First, we acquire and segment CT volume images of a given subject’s joint in different joint poses, then recover through registration the bone kinematics. Next, we use the recovered bone surfaces and kinematics

to compute the space between bones during motion. We use this inter-bone space measure to define cartilage map location and thickness. Finally, we model the cartilage map as a deformable height field, and use this model to compute cartilage contact. We describe each step below.

4.3.1 Data Acquisition and Recovery of Kinematics

Our data is generated by CT-imaging a joint in seven different poses. The volume images were acquired with a Hispeed Advantage scanner, General Electric Medical System, scan parameters: 80kV, 80mA, image resolution $0.94 \times 0.94 \times 1 \text{ mm}^3$. Next, bones are registered across all volume-images. Each CT volume-image corresponds to an articulation pose, thus sampling the space of articulation kinematics. An additional higher-resolution scan ($0.31 \times 0.31 \times 1 \text{ mm}^3$) is acquired in a reference neutral pose, allowing us to segment the bone surfaces with higher accuracy.

Through manual segmentation, thresholding, and user interaction, bone surfaces are extracted from the high-resolution reference CT volume image [28]. Bone surfaces are further modeled as NURBS surfaces using the Geomagic software package [43]. The resulting triangular meshes have each approximately 50,000 faces. Next, each bone surface is tracked with sub-voxel accuracy through the sequence of remaining CT volume images, as described in Chapter 2. The tracking procedure reports relative bone-motion from one articulation pose to another. The resulting dataset consists thus of the articular-joint geometry and its sampled kinematics. Fig. 4.2 shows two recovered poses of a volunteer's wrist.

4.3.2 Inter-Bone Joint-Space Modeling

Using the bone geometry and kinematics recovered as described above, we compute the inter-bone joint-space across motion, as described in Chapter 3. The inter-bone joint-spacing defines the cortical surface where bones articulate near each other; it is defined as the cortical surface area on the bone that is less than a proximity threshold p distance from the cortical surface of a neighboring bone. The proximity threshold p is iteratively

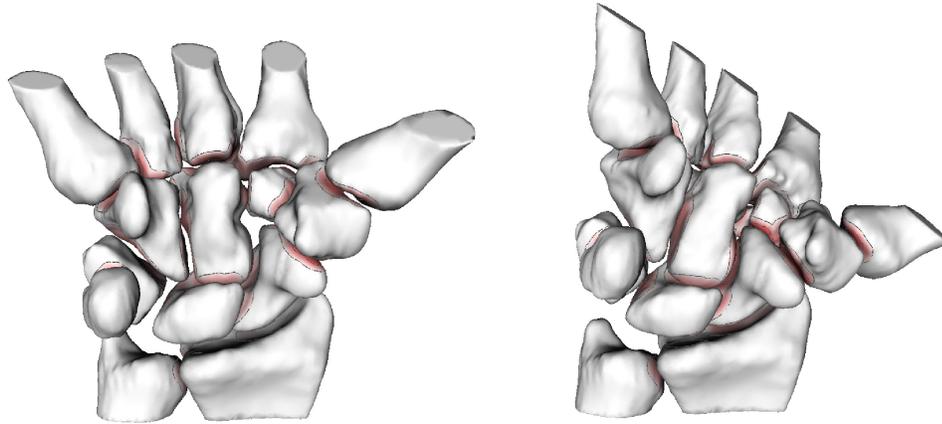


Figure 4.2: Modeling inter-bone joint-spacing from bone surfaces and motion: two different poses recovered through CT-imaging, segmentation and registration of the same wrist joint (eight carpal, two forearm, and five metacarpal bones). Bones are color-mapped and contoured. The color saturation on bone surfaces indicates the distance to the nearest point on the opposite bone; darker regions are closer. Bone surfaces where the inter-bone distances are larger than a proximity threshold ($p=2\text{mm}$ here) are shown in white. We trace these minimum-distance regions across multiple poses in order to estimate cartilage-map location and thickness.

determined, as described in the following section 4.3.3.

The key idea here is that, because cartilage coats bony surfaces wherever the bones articulate with each other during motion, the inter-bone joint-space should correlate with cartilage location. That is, cartilage should be located wherever two bony surfaces are in close proximity during motion. In Fig. 4.2 inter-bone joint-spacing areas are shown 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 proximity threshold value p are neither colored nor contoured and are shown as white surfaces. In Fig. 4.2, note the shift in the location of red areas on the trapezoid bone (Fig. 2.5) between one pose and another. We hypothesize that the cartilage location is given by the union of bony red areas across motion.

Modeling the inter-bone joint spacing requires cortical bone-to-bone distance information for multiple poses. We compute distance fields for each bone from the reconstructed NURBS bone models using a level-set approach. We use the Closest Point Transform

(CPT) level-set software package [73] to generate the distance field representation from the surface representation. We use the distance fields to find, for each pose, distances from every vertex in the surface model of one bone to its neighbors. Using the inter-bone distance we also compute, for each pose, isocontours on the cortical-bone surface. Each contour shows where the distance map is equal to a constant distance. The inter-bone joint-spacing in a given pose is then the area of the surface triangles within the p contour for that pose.

4.3.3 Inferring the Cartilage Map Location and Thickness

We infer the cartilage map location and thickness in a three-step procedure. In the first step, we compute the cartilage map location. To compute the cartilage map location, we need to first select a value for the proximity parameter p . We then compute, across poses, the union of contact areas for the selected p value.

In the last two steps, we compute the cartilage map thickness. We note that at each instant, any two points on the two bones that are approximately on each other's normal vectors must have the property that the cartilage thickness at those points sums up to no more than the distance between them. We hypothesize that for two cartilage maps in contact A and B cartilage thickness is distributed evenly between map A and map B. According to this hypothesis, to find the cartilage map thickness at a cartilage map vertex v we can extrude v by half the minimum distance across poses between A and B, plus a th percentage, in the direction of the bone surface normal at that vertex v .

The th parameter is an extra-thickness measure of the cartilage layer, because we expect that at each pose the cartilage map is already somewhat deformed. Accordingly, we need to first select a value for the th parameter, and then proceed with the extrusion operation. The collection of bone surface points v_k , their oriented normals n_k , and their extrusion heights h_k defines the height-field representation of the cartilage map.

Let $nbones$ be the number of bones in the joint, $nposes$ the number of joint poses available, and RT_{ij} the joint pose (rotation and translation) associated with bone i in pose j . Let $distance(RT_{ij}(v), RT_{lj}(b_l))$ be the minimum distance from vertex v of bone b_i to the neighboring bones b_l in pose j . We use anatomy-book knowledge and an iterative procedure to determine the values of p and th and generate the cartilage maps, as follows:

```

procedure estimate_cartilage
in: bone surfaces  $b_i, i = 1 : nbones$ ;
      sampled bone kinematics  $(RT)_{ij}, j = 1 : nposes$ ;
      initial estimates for  $p$  and  $th$ :  $p_0, th_0$ ;
out: cartilage maps  $C_i = \bigcup_{k=1:|C_i|} (v_{ik}, n_{ik}, h_{ik})$ ;

//1. estimate cartilage location
for  $i = 1$  to  $nbones$ 
   $C_i \leftarrow \emptyset$ 
  repeat
     $p \leftarrow p_0$ 
    foreach vertex  $v \in b_i$ 
      for  $j = 1$  to  $nposes$ 
         $dist(v, j) \leftarrow distance(RT_{ij}(v), RT_{lj}(b_l))$ 
       $d(v) = \min_j dist(v, j)$ 
      if  $d(v) < p$  then
         $h = d(v)/2$ ;
         $n = normal(v)$ ;
         $C_i \leftarrow C_i \cup (v, n, h)$ ;
      endif;
    increase  $p$ 
  until ( $C_i$  overlaps  $C_{i,anatomy-book}$ )
//end step 1.

//2. infer  $th$  parameter
 $th \leftarrow estimate\_th ( C, b, th_0)$ 

//3. estimate cartilage total thickness
for  $i = 1$  to  $nbones$ 
  foreach  $(v_k, n_k, h_k) \in C_i$ 
     $h_k = (1 + th)h_k$ 
//end step 3.

```

In our experiments, p was initialized to 1 mm, and increased in 0.1mm increments.

In the experiments described in this chapter, the overlap between the computed cartilage map and the anatomy-book cartilage map was evaluated visually. For example, for the scaphoid bone shown in Fig. 4.3, increasing the p threshold above 2mm causes certain features like the diagonal palmar strip to disappear, while decreasing the p threshold below 2mm introduces holes in the cartilage map. However, the overlap-evaluation step could be automated using the cross-parameterization technique described in Chapter 6.

Technically, the most complex and delicate step in the **estimate_cartilage** procedure is inferring the extra-thickness parameter th . We assume th has the same value across the collection of joint bones. We calibrate this th value based on anatomy-book knowledge as follows:

```

procedure estimate_th
in: bone surfaces  $b_i, i = 1 : nbones$ ;
      copy of current cartilage map  $C_i$ ;
      initial  $th$  estimate:  $th_0$ 
out:  $th$ ;
 $th \leftarrow th_0$ 
select special-case bones  $b_i^*$ 
repeat
  foreach  $b_i^*$ 
    foreach  $(v, n, h) \in C_i^*$ 
       $h = (1 + th)h$ ;
    resolve_collisions( $C^*$ );
    compute_contact( $C^*$ );
    decrease th;
until (contact w/in special-case anatomy book range);

```

In the experiments reported in this chapter, th was initialized to 10% and gradually

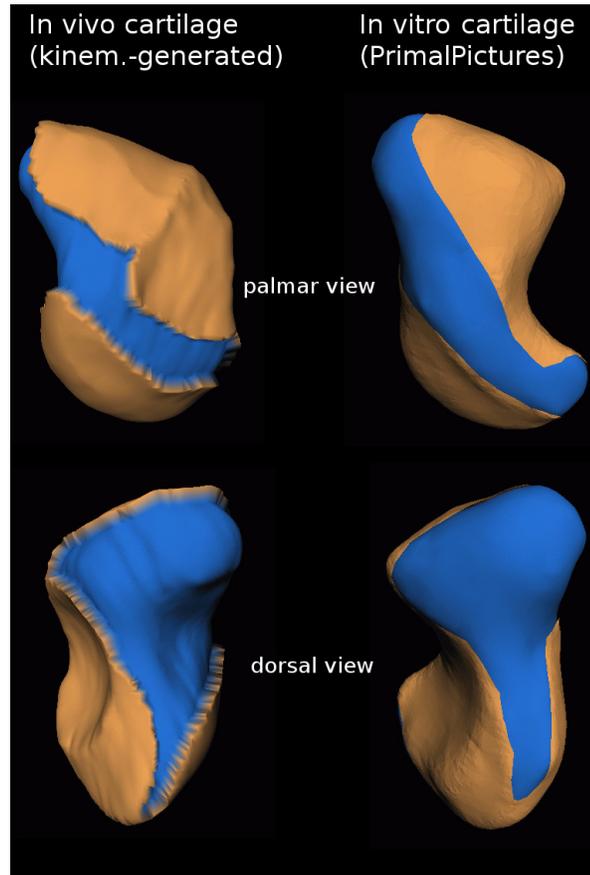


Figure 4.3: *In vivo* scaphoid cartilage map generated through our kinematic modeling approach (left) versus an anatomy-book scaphoid cartilage map estimated through dissection *in vitro* (right). Bone is shown in blue, cartilage in tan (light-grey in grayscale reproductions); top: palmar view, bottom: dorsal view. Note the generally similar location-on-bone of the *in vivo* and *in vitro* cartilage-map. Despite subtle anatomical differences between the *in vivo* and *in vitro* bone-shape, note the diagonal strip cutting through cartilage in both palmar views, and the sliver-cut extending towards the bottom in both dorsal views. *In vitro* data depicted on the right courtesy of Primal Pictures Ltd.; *in vitro* data does not include cartilage-map thickness.

decreased in 1% decrements. The **resolve_collisions** and **compute_contact** steps are presented in section 4.3.4.

In the procedure above, the special-case refers to subsets of bones for which anatomy-book data reporting contact-size results exist. Because anatomy-book studies use pressure-sensitive film to compute contact-size between bones, such studies in general are performed only for joint bones that have at least one flat surface and thus can accommodate the flat pressure-sensitive film. For example, Patterson and Viegas [85] report the contact surface defined by the scaphoid and lunate bones on the radius flat head, when increasing loads are applied to the wrist joint.

Note that, for each pose and pair of articulating bones we keep a list of the height-field points whose original cortical-surface support vertices had a minimum distance below p . These lists are used to speed up the computation of cartilage local deformation.

4.3.4 Cartilage Contact Simulation

Our cartilage model is a meshless, incompressible, deformable height-field whose initial conditions are determined as described in section 4.3.3. While cartilage is a compressible material, the computational requirements for a multi-articular model necessitate some simplifications. Given the small deformations likely to take place in unloaded joints, assuming cartilage is incompressible but deformable is reasonable. Incompressibility is achieved by transferring volume from compressed locations to locations not in contact.

Since we anticipate articular cartilage deformations to be small, our deformation model is based on the physics of linear elasticity. We represent cartilage geometry as unstructured point clouds, i.e., sets of point primitives that sample the position and normal of the underlying surface. For each cartilage surface point we store its height with respect to the underlying bone surface along the cartilage surface normal. Cartilage points are allowed to move along this cartilage normal direction.

We resolve cartilage collisions (**resolve_collisions**) iteratively through repeated:

- collision detection,
- negative displacement in the direction of the normal n_j of colliding points v_j

- volume preservation: positive displacement in the direction of the normal n_k of the non-colliding points $v_k \in$ the collision neighborhood,

for all the objects involved in collisions.

The size of the collision neighborhood correlates with cartilage viscosity and is defined by the proximity threshold p determined in section 4.3.3. The neighborhood consists of vertices which are closer than p to the colliding bone. Cartilage volume is preserved by constraining each height-field response such that the integral of all displacements over the height-field is zero. In our experiments, cartilage contact was usually resolved within under 100 iterations. The height-field points in contact after resolving all collisions defines the cartilage contact area (**compute_contact**).

Figure 4.4 illustrates the cartilage deformation process in 2D. In this example we model the contact between two flexible objects: a computer-generated flexible semicircle of radius 50 units was placed so that it collided with a horizontal flexible line-segment of length 300 units. The maximum collision depth was 10 units at the semicircle South tip. Semicircle and line points were assumed to have variable material thickness: semicircle thickness was maximum at the South tip and tapered off towards the equator; line-segment thickness was also maximum in the middle of the segment and tapered off towards the segment ends. The semicircle and line were iteratively deformed to resolve the collision while attempting to preserve object area by constraining the integral of all displacements over each object surface to be zero. Non-colliding points were deformed proportionally to their material thickness. Although material thickness is only a proxy for cartilage properties, note each object's response is qualitatively appropriate: both objects retract in the collision area and bulge sideways to compensate for the loss of volume thus incurred. Note that the amount of interpenetration modeled in this example — and hence the resulting deformation (approximately 20% of the object thickness) — is an over-estimate of the cartilage map deformations we expect in the human wrist.

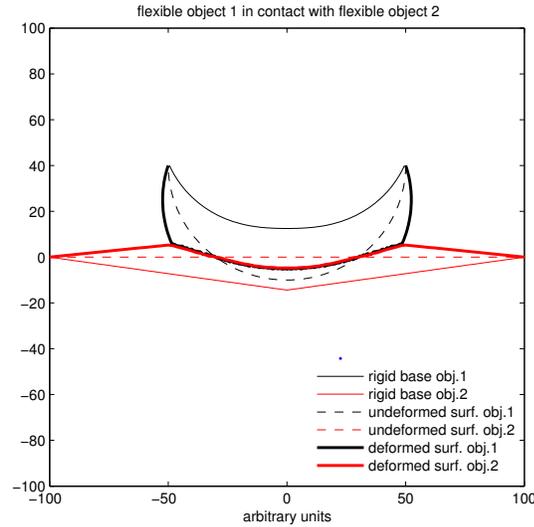


Figure 4.4: 2D illustration of our algorithm for predicting contact deformation of flexible objects with a rigid skeleton (not shown) in 2D. Although material thickness is only a proxy for cartilage properties, note each object’s response is qualitatively appropriate: both objects retract in the collision area and bulge sideways to compensate for the loss of volume thus incurred.

4.4 Validation and Results

We validate against *in vitro* data the location of our kinematically-generated cartilage maps, their thickness, and the computed cartilage contact areas.

Cartilage location

In Fig. 4.3 we compare the location of the *in vivo* cartilage map we generated through our method for the scaphoid bone (live male individual, 25 years old, 7 poses) against a similar cartilage map reported *in vitro* (unknown gender and age; *in vitro* data courtesy of Primal Pictures Ltd., London, UK). In this experiment the inferred value for the proximity threshold p was 2mm. Note the generally similar location-on-bone of the *in vivo* and *in vitro* cartilage-map. Despite subtle anatomical differences between the *in vivo* and *in vitro* bone-shape, note the diagonal strip cutting through cartilage in both palmar views, and the bottom sliver-cut in both dorsal views.

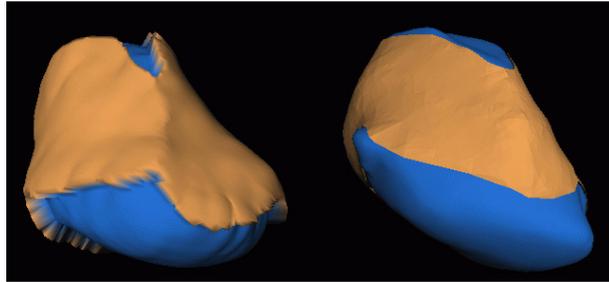


Figure 4.5: *In vivo* trapezoid cartilage map generated through our kinematic modeling approach (left) versus a trapezoid cartilage map estimated through dissection *in vitro* (right). Bone is shown in blue, cartilage in tan (light-grey in grayscale reproductions). Note the similar location-on-bone of the *in vivo* and *in vitro* cartilage-map (tennis-shoe shape), despite slight anatomical differences between the *in vivo* and *in vitro* subject. *In vitro* data depicted on the right courtesy of Primal Pictures Ltd.; *in vitro* data does not include cartilage-map thickness.

Fig. 4.5 shows a similar comparison between the location of the *in vivo* generated cartilage map for the trapezoid bone (live male individual, 25 years old, 7 poses) against a trapezoid cartilage map reported *in vitro* (Primal Pictures Ltd.). The p value inferred through calibration was again 2mm. Again, we note the similar location-on-bone of the *in vivo* and *in vitro* cartilage-map — tennis-shoe shape, despite slight anatomical differences between the *in vivo* and *in vitro* subject.

Cartilage thickness

In a first experiment, to validate the thickness of our cartilage model, we use the trapezoid bone from the left wrist of an unembalmed cadaver upper extremity (female, age 66 years) and the trapezoid bone of a live individual (male, age 25 years). The volar fragment of the cadaver trapezoid was immersed in contrast dye, and a 3-D volume image was made of the contrast-surrounded bone using a desktop μ CT scanner (μ CT 20, Scanco Medical). This technique yielded high-resolution segmentable images where the cartilage was clearly distinguishable from the surrounding contrast dye and from the underlying bone tissue. The thickness of the cartilage in the 1mm section of trapezoid evaluated in this study was relatively uniform, with an average thickness of 0.66 ± 0.14 mm (min. 0.22 mm, max. 1.07 mm).

Table 4.1: Trapezoid cartilage thickness

Trapezoid cartilage thickness	<i>In vitro</i> , invasively (μ CT-imaged)	<i>In vivo</i> , non-invasively (kinem.-generated)
Mean \pm Std.dev.	0.66mm \pm 0.14mm	0.64mm \pm 0.19mm
Min	0.22mm	0.147mm
Max	1.07mm	1.049mm

We used the kinematics-based method presented in section 4.3 to generate the cartilage map for the trapezoid bone of the live individual. The resulting value of the thickness parameter th was 1%. The thickness of the *in vivo* cartilage model was determined by computing the average height of the kinematically-generated trapezoid undeformed height-field. The thickness of our model was also relatively uniform, with an average thickness of 0.64 ± 0.19 mm. The results reported in Table 4.1 show remarkable correlation between the *in vivo* and *in vitro* measurements. However, note that in this first thickness experiment we compare *in vivo* data and *in vitro* data collected from different individuals.

In our second thickness experiment, we use *in vitro* data collected from the same individual. A cadaver wrist was CT-imaged in 12 poses, spanning the wrist range of motion. The wrist bones were segmented and their motion tracked accurately. We used the kinematics-based method to ‘grow’ the cartilage map for the scaphoid bone. The proximity p and thickness th values resulting after calibration were 2mm and 5% respectively. We computed the average, standard deviation, minimum and maximum thickness for the radius facet of the resulting scaphoid cartilage map.

Following CT-imaging, the wrist was dissected. All soft tissue except cartilage and ligament insertion sites were removed from the scaphoid bone. The bone itself was immersed in contrast dye for 24 hours, then μ CT-imaged at a 60 micron resolution. The resulting μ CT-imaged bone surface and cartilage map was reconstructed from the volume image using the Geomagic software package [43]. The μ CT-ed bone surface and cartilage map surrounding it were aligned with the CT bone surface acquired prior to dissection. Finally, the anatomist interactively selected the radius facet of the μ CT scaphoid cartilage map. We computed the average, standard deviation, minimum and maximum thickness for the μ CT

Table 4.2: Scaphoid cartilage thickness

Scaphoid cartilage thickness radial facet	<i>In vitro</i> , invasively (μ CT-imaged)	<i>In vitro</i> , non-invasively (kinem.-generated)
Mean \pm Std.dev.	0.601mm \pm 0.21mm	0.596mm \pm 0.20mm
Min	0.275mm	0.276mm
Max	1.21mm	1.05mm

cartilage facet.

Table 4.2 shows the invasive μ CT thickness measurements and the non-invasive, kinematically generated thickness estimates; note the remarkable correlation between the two columns. The difference between the measured and estimated maximum thickness value (1.21mm measured, 1.05mm estimated) appears to be due to the accidental inclusion of a ligament insertion site vertex in the μ CT cartilage map during the interactive facet selection phase.

Cartilage contact

We compare the cartilage contact areas generated *in vivo* through our method against the *in vitro* pressure-film results reported by Patterson and Viegas [85] for the radius-scaphoid-lunate joint. Although our results are generated for an unloaded joint while Patterson and Viegas report (lightly) loaded joint results, we expect our results to extrapolate the loaded results in the direction of the lowest loads. In Fig. 4.6 we show the cartilage contact generated through our method on the radius distal surface by the lunate and scaphoid bones, in two different poses — neutral and extension. As in [85], the contact areas were localized and accounted for a relatively small fraction of the joint surface (approximately 24%). The scaphoid-radius contact-area shifted from a palmar location to a distal location as the wrist moved from the neutral pose to the extension pose, consistent with the *in vitro* findings of Patterson and Viegas. The *in vivo* radius contact area was 75 mm² in the neutral pose and 49 mm² in the extended pose, again in agreement with the *in vitro* results reported in [85].

In this experiment we used cartilage maps generated with a *th* extra-thickness parameter

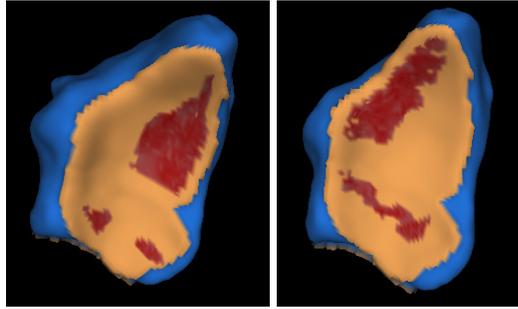


Figure 4.6: *In vivo* contact area between the distal radius bone (shown) and the lunate and scaphoid bones (not shown) in a normal unloaded wrist, neutral pose and an extension pose. Bone is shown in blue, cartilage in tan, cartilage contact in red (largest area depicted corresponds to the scaphoid-radius contact). The contact areas are localized and account for a relatively small fraction of the joint surface. Note the scaphoid-radius contact-area shift from a palmar location to a distal location as the wrist moves from the neutral pose (left) to the extension pose (right).

value of 1%. Note that in the *th* calibration phase described in section 4.3.3, increasing the 1% value generated in general contact areas close to 100 mm²; such numbers would contradict the *in vitro* measurements of Patterson and Viegas [85], who found that for loads under 23 pounds these areas stay under 100 mm².

4.5 Discussion

The validation experiments show good correlation between our *in vivo* kinematically generated cartilage maps and *in vitro*-observed cartilage maps. Interestingly, our cartilage thickness experiment seems to indicate that ‘growing’ cartilage by half the inter-bone distance (as previously done by Carrigan et al. [24]) is a reasonable approach when modeling the wrist joint — larger bones in the wrist like the scaphoid may not grow thicker cartilage. The remarkable thickness correlation between the elderly female dataset and the young male dataset reported in our first thickness experiment indicates joint size may not influence cartilage thickness. In other words, thickness measurements performed *in vitro* could be transferred to *in vivo* data, regardless of differences in the bone size.

Our cartilage map generation method uses two parameters, the proximity threshold p , and the extra-thickness parameter th . The 2mm threshold value resulting in most of our

experiments seems to indicate that modeling carpal cartilage as a uniform 1mm-thick shell wrapping articular surfaces — as previously done by Thoomukuntla et al [107] — is a reasonable approach for at least certain bones like the trapezoid, although our measurements for other carpal bones indicate non-uniform thicknesses.

Similarly, the resulting extra-thickness th value of 1% to 5% may indicate that unloaded cartilage is minimally compressed during motion. Interestingly, this observation ties in with the results of an earlier study by Herberhold et al. [50] on the *in situ* measurement of articular cartilage deformation in intact knee joints under static loading. Herberhold et al. found that patellar cartilage deforms minimally during the first minute of static loading (3% absolute deformation). These findings suggest that, during everyday motion, articular cartilage deforms minimally even in bodyweight-bearing joints like the knee.

Our cartilage-map generation method has certain limitations. Unlike real cartilage, the cartilage maps we generate do not recede towards the boundary of the map. This limitation could be addressed by progressively decreasing the height of the map as we approach the location boundary.

We also require the acquisition of poses spanning the whole space of joint kinematics. However, since our cartilage model deforms locally, missing cartilage areas are not likely to influence the contacts we compute. The correlation between minimum inter-bone distance during motion and cartilage location may have in itself strong biological implications.

When extruding cartilage-location vertices according to the inter-bone distance, our current implementation assumes an implicit correlation between the direction of the surface normal and the inter-bone distance; while this correlation is present in our data, a sounder approach would be an iterative small vertex displacement in the normal direction and reevaluation of the distance until the opposed displaced surface collide.

Another limitation of our model is that we assume cartilage is incompressible; given the small deformations we expect to find in the unloaded wrist during motion, assuming that cartilage is incompressible but deformable is reasonable. This assumption may not hold true in other anatomical joints. We note that for static, loaded poses our model can be coupled with finite element analysis to enable more sophisticated computations of deformation and stress throughout the material.

4.6 Conclusion

We presented in this chapter a method for estimating articular cartilage geometry *in vivo*, directly from CT-imaged bone geometry and joint kinematics. The resulting *in vivo* cartilage maps match well *in vitro* cartilage maps. We augmented the cartilage maps with a novel height-field computational model of cartilage deformation, and showed the resulting contact areas replicate carpal cartilage contact observed *in vitro*.

The results reported in this chapter indicate that articular cartilage deforms minimally during normal motion. The implication is that rigid or almost-rigid, fast-to-calculate representations of cartilage are acceptable building blocks when constructing fully-predictive, dynamic models of diarthrodial joints.

Chapter 5

Predictive Simulation of Diarthrodial Joints

5.1 Introduction

Diarthrodial joints are arguably some of the most complex structures in the human body; they combine in unique and ingenious ways hard-tissue and soft-tissue to allow for everyday motion and athletic activities. A thorough understanding of diarthrodial joint motor function in live individuals is essential in the treatment of injuries and degenerative diseases like osteoarthritis, and also in the design of orthopedic implants. However, currently there is no satisfactory explanation for the motion patterns observed in most complex joints [14, 98, 34], nor for the motion differences observed among different live individuals.

In this chapter we present a medical-image based method for constructing and simulating joint models, using data collected from live individuals. The method uses as input medical volume images of the joint bones, tracked motion and anatomy-book knowledge. The resulting models and simulations are adaptable to individual differences and have predictive capabilities in terms of predicting soft-tissue contact within the joint, and understanding the kinetic response of the joint.

The aim of this chapter is to demonstrate that data obtained *in vivo* through medical

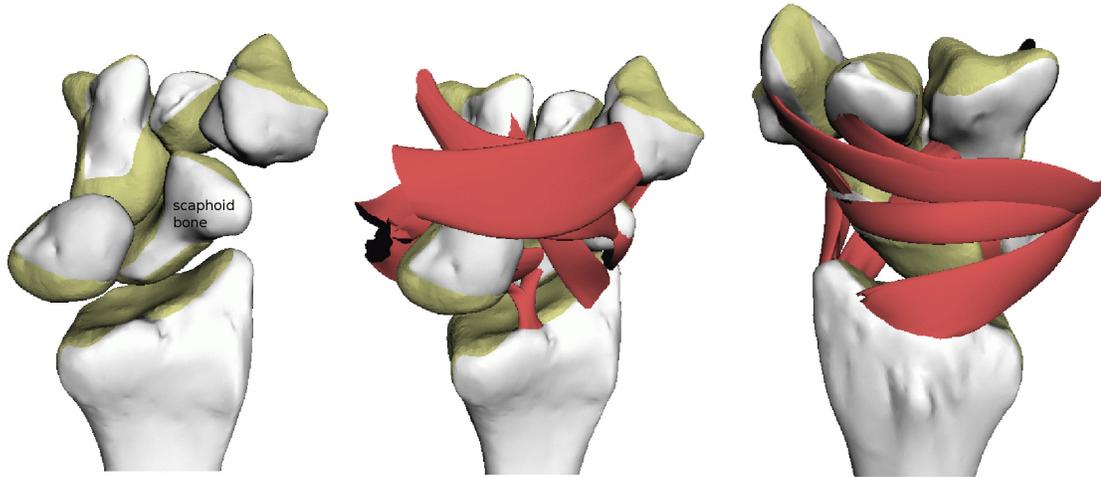


Figure 5.1: Soft-tissues in the human wrist interacting directly with the scaphoid bone (data courtesy of Primal Pictures). Left: a six-bone subset of the human wrist showcasing the bones (white) and articular cartilage (yellow) interacting directly with the scaphoid bone. This six-bone structure includes nine pairwise articulations and fifteen ligaments that may impact scaphoid kinematics. Center: palmar wrist ligaments (red) interacting with the scaphoid bone. Right: dorsal wrist ligaments (red) interacting with the scaphoid bone. The human wrist presents a predictive modeling opportunity: no muscles insert on the scaphoid bone. Scaphoid motion is defined solely by bone shapes, bony articulations, and ligaments.

imaging can be combined with numerical simulation to develop a predictive tool for understanding the role of soft-tissues in the functioning of joints.

We demonstrate our model and simulation on the human wrist. The joint-subset we use throughout our experiments includes nine pairwise articulations and fifteen ligaments that may affect the kinematics of the scaphoid bone (Fig. 5.1). The scaphoid bone and its surrounding articulations are of significant clinical interest: the bone is frequently subject to fractures, and its connecting ligaments are often subject to tears and surgical intervention. Notably, there are no muscle insertions on the scaphoid bone. This lack of muscle insertions provides an excellent modeling opportunity: we can study scaphoid kinematics without having to specify external muscle forces.

5.2 Related Work

Dynamic joint simulations can be divided into two types, based upon the input and output variables [103]: (1) inverse-dynamics simulations, and (2) forward-dynamics simulations. In inverse-dynamics simulations the input variables are the kinematics of the bones, which also define the end-conditions of deformable soft-tissues, and the output variables are the corresponding forces and torques acting on the joint. From a predictive point of view, inverse simulations have been used not only to compute joint moments, but also to estimate muscle forces, contact and stresses in the joint [41, 34, 118]. Such simulations have predictive capabilities in the sense that they can forecast, for example, which component of a joint may break down due to excessive or repetitive stress.

While inverse dynamic simulations are useful for analyzing experimental data, forward dynamics may be used to perform simulated experiments. In forward dynamics the independent variables are the forces and torques acting on the joint, and the output variables are the joint motion. The advantage of forward simulations is — as already stated — that they can be used to perform simulated experiments. For example, given a set of muscle forces and external forces, a forward-dynamics model of a joint could predict how joint motion would be altered after a simulated tear of a ligament or after a bone fracture.

The forces and torques resulting by solving inverse-dynamics problems could be used in a forward-dynamics formulation to yield the complete behavior of a joint, in terms of computed contact, stresses, and motion [13]. However, forward simulations require not only accurate digital models of the geometry of all the objects forming the joint, but also detailed knowledge of the material properties of each deformable object, and of the external forces applied. As discussed in Chapter 1, these inputs are not readily available. Currently, forward-modeled joints use material properties and forces specified by the user, based in general on *in vitro* studies; the values observed *in vitro* may well be different from individual-specific *in vivo* values, due to either anatomical variability or specimen preparation. In addition, accurately modeling deformable contact is often costly [45].

The combined challenges of missing inputs and simulation efficiency lead to a pattern in predictive joint modeling: small-scale, highly detailed *in vivo* subject-specific models have inverse predictive capabilities (contact and stresses), while highly simplified *in vitro*

models may have limited forward predictive capabilities (motion). The motion-prediction capabilities of *in vitro* forward-models are in general difficult to validate, and the models themselves currently discount individual differences.

Current models of diarthrodial joints include only skeletal tissues: bone, ligaments, cartilage, and muscle-tendon units.

We also note that current forward models generate motion, in fact, through quasi-static analysis, and not dynamic analysis. A quasi-static analysis is one where the governing equations are those of statics; however, quasi-static analysis can be applied to joints in motion as long as inertial forces are negligible. The advantage of quasi-static analyses is that they do not require the input of parameters such as mass and inertia properties, or damping properties for which data are rarely available in the analysis of musculoskeletal structures.

In the remainder of this section we briefly review predictive inverse and forward joint models, and current procedures for validating joint predictive capabilities. Work in physically based simulation is surveyed in detail by Pauly et al. [86].

Predictive Inverse Joint Models

Inverse models exist for a variety of joint substructures. Recent work by Gardiner and Weiss [41] can predict *in vivo* strain through one entire knee ligament. Similar work by Warner et al. [113], Donahue et al. [33], or Meakin et al. [75] can predict *in vivo* stress through one cartilaginous structure. In general, these approaches model through finite element methods one deformable component at a time.

A few inverse models also exist for simple joints — joints with a few components. The inverse knee model developed by Zhu et al. [118] uses live individual motion capture and the bone geometry to compute menisci contact and stresses in the joint. Similarly, the temporomandibular joint developed by Donzelli et al. [34] uses tracked motion and the joint geometry to predict contact and stresses in the joint. In both studies ([118] and [34]), neither articular cartilage or soft-tissue wrapping were modeled.

Joints more complex than the knee or the jaw have proven difficult to model due to the number of bodies involved and complexity of the soft tissue interaction [24]. With respect

to modeling complex joints, Carrigan et al. [24] created a simplified 3-D finite element model of the wrist. In this model 3D solid finite element meshes representing the bones were hollowed by eliminating all elements that did not reach the external mesh surface in order to reduce the number of nodes and elements in the mesh, and thereby improve model efficiency. The articulated bony surfaces were extruded to mimic cartilage. Several pairs of bones were fused into single rigid bodies, and ligaments were modeled as line springs. Only one static pose was modeled. In this model bone motion had to be restricted artificially to certain directions through non-physiological constraints in order to prevent the wrist from collapsing under the applied load.

Wrapping soft-tissues may play an important role in stabilizing the joints during motion. However, current inverse models provide almost no insight into the role of soft-tissues in the functioning of joints. In particular, none of the inverse models attempting to predict contact surveyed in this section takes into account soft-tissue wrapping.

The paucity of inverse models for complex joints can be traced back to a general lack of soft-tissue morphological data and to simulation efficiency challenges, as argued in Chapter 1. We show that the data-driven approach introduced in the earlier chapters of this dissertation can overcome such limitations.

Predictive Forward Joint Models

In current forward models, the bones are modeled as rigid bodies and the soft tissues are modeled as tensile and compressive springs, using *in vitro*-determined values. In order to simulate the movement of the joint, forces are applied at tendon insertion sites and the elements of the model move in such a way as to solve the equations of kinetic equilibrium. For example, the *in vitro* forward knee model developed by Blankevoort et al. [18] can replicate knee-motion measured *in vitro*. However, it is essential that the physical properties and dimensions of the structures used in the model are applied correctly in order to obtain meaningful results. Despite repeated attempts, there is currently no single model that can be applied to anticipate the movement patterns of joints more complex than the knee. In the case of the wrist, it has been suggested that such a forward model cannot be determined [37].

Circumventing the lack of wrist soft-tissue data, Sirkett et al. [98] have built a forward model of the wrist that uses as input only the morphology of the bones. The hypothesis behind the model is that soft tissues play a secondary role in governing the movements of the wrist bones, and that the morphology of the bones is what dictates their behavior. The premise is that biological structures are formed using the minimum of materials necessary to provide adequate function. In the case of joints, this is achieved by maximizing contact area through the use of conforming joint surfaces. On a very limited range of motion wrist kinematics were correctly predicted for several bones. However outside this narrow range the model did not match kinematics observed in *in vitro* and *in vivo* studies — e.g., bone rolling during wrist flexion. With respect to the contact-maximization hypothesis, it is worth noting that zoologists believe that adjacent bones can not reach maximum congruence at all poses [90]. The classical zoology interpretation is that the position where all the bones reach maximum contact has particular mechanical significance, most likely that this is the configuration the bones take at maximum load transfer.

Validation of Predictive Capabilities

Validating the predictive capabilities of both inverse and forward models poses significant challenges. In general, contact and strain predictive capabilities can only be validated against *in vitro* data. However, even *in vitro* data has certain limitations. For example, to measure joint contact *in vitro*, researchers typically insert pressure-sensitive film between two articulating bones, apply a load to the joint, then remove the film and estimate the size and location of the resulting print [85, 107]. Because only flat surfaces can accommodate pressure-sensitive film, such measurements have been reported only for bone pairs that include at least one bone with a flat articulating surface. Furthermore, studies performed on different individuals may lead to different measurement values. Even *in vitro* reported values do not always agree; there is large variation in the reported mechanical properties of specimens: for example, for whole scapholunate ligament specimens, Mayfield [74] reports a stiffness of 250N/mm (S.D. 90N/mm) at a displacement rate of 50mm/min, while Johnston et al. [54] reports 66N/mm (S.D. 29N/mm) at 60mm/min. Accordingly, in general,

validation of predictive capabilities with respect to contact and tissue mechanical properties tries to match a range of values reported *in vitro*, and occasionally even extrapolate from available *in vitro* measurements. Oftentimes validation is performed through visual observation, for example, observing whether in a given pose a specific ligament is taut or relaxed [36, 5].

Motion-prediction capabilities are in general validated against sampled data, often *in vitro* motion data [61]. In general, it is essential that the physical properties and dimensions of the structures used in the model — for example, muscle forces — are applied correctly in order to obtain meaningful results. Unfortunately, oftentimes multiple force combinations can result in the same meaningful result.

In this sense, the human wrist offers a unique opportunity in validating motion predictions. The wrist is a uniquely passive joint, consisting of eight small, complexly shaped bones interposed between the distal radius and metacarpal bones. There is only one muscle that has minor insertions on the wrist. Accordingly, motion of the bones is defined solely by the complex, semi-congruent bone shapes, intricate bony articulations, and the network of ligaments that spans the wrist. This frees a potential model from being limited by the accuracy of the estimates of the applied muscle forces. We demonstrate our model and simulation using wrist data.

5.3 Methods

The joint model we build integrates CT-acquired bone geometry, tracked kinematics, and computational models of articular cartilage and ligaments. The model does not include more superficial layers like skin, fat and the neurovascular system, since *in vitro* studies indicate that these layers play a relatively small role in joint motion [18, 52]. We also do not model muscles, which may influence indirectly the joint by compressing the sac of synovial fluid surrounding the joint, nor the synovial sac itself. Again, *in vitro* studies indicate these last two components play a secondary role in the kinematics of the wrist joint [18, 52].

5.3.1 Data Acquisition

Our data was generated by CT-imaging the wrist bones of a male volunteer in seven different poses. The volume images were acquired by using a Hispeed Advantage scanner, General Electric Medical System, scan parameters: 80kV, 80mA, image resolution $0.94 \times 0.94 \times 1 \text{ mm}^3$. Each CT volume-image corresponds to an articulation pose, thus sampling the space of articulation kinematics. An additional higher-resolution scan ($0.31 \times 0.31 \times 1 \text{ mm}^3$) was acquired in a reference neutral pose, allowing us to segment the bone surfaces with higher accuracy.

Through manual segmentation, thresholding, and user interaction, bone surfaces are extracted from the high-resolution reference CT volume image [28]. Bone surfaces are further modeled as NURBS surfaces using the Geomagic software package [43]. Next, each bone surface is tracked with sub-voxel accuracy through the sequence of remaining CT volume images, as described in Chapter 2. The tracking procedure reports relative bone-motion from one articulation pose to another. The motion of each wrist bone is reported in coordinates relative to the fixed forearm.

5.3.2 Model Construction

The model of the wrist includes three-dimensional geometric data of the eight carpal bones, two forearm bones, and five metacarpals. We use the bone geometry and sampled kinematics to generate cartilage maps and ligament fibers.

We generate cartilage maps from the fifteen bone surfaces and their kinematics, as described in Chapter 4 (2mm proximity threshold, 1% compressed thickness).

We manually identify ligament insertion sites based on anatomical expert knowledge [100]. We define three to four equally spaced fibers per ligament bundle to account for the band-like structure of ligaments. For each fiber and joint pose, we automatically generate minimum-length paths constrained to avoid bone penetration, as described in Chapter 3. An example ligament fiber interacting with multiple bones is shown in Fig. 5.2.

For each ligament fiber we compute and report the minimum path length. The resulting fiber paths are visually validated against the anatomy-book data.

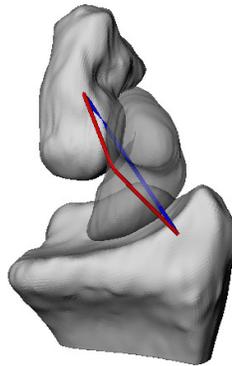


Figure 5.2: Ligament fiber interacting with multiple bones. The computed minimum-length path is shown in red, the straight line between the two insertion sites is shown in blue. Note the significant deflection of the generated path from the straight line.

In the experiments presented in this chapter we do not model ligament-cartilage interactions. Ligament paths are generated over bone, not cartilage surfaces, primarily due to a boundary-related limitation in our cartilage map generation method, as discussed in Section 5.5.

5.3.3 Simulation

Inverse Simulation

In our inverse-dynamics experiments, we assemble bone, cartilage and the ligament fibers resulting for each joint pose. When simulating the joint, we make the following simplifying assumptions:

- bones are rigid; bone deformations are ignored because they are relatively small compared to the displacements in the joint;
- the friction between cartilaginous contact surfaces is ignored; due to the synovial fluid, the friction coefficient between cartilage surfaces is very low (100 times smaller than the coefficient for ice);
- for simplicity, the viscoelastic properties of ligaments and cartilage are approximated;

- bone-mass is considered to be uniformly distributed, again for simplicity.

Cartilage maps are characterized by their location, thickness, and stiffness coefficient. They are modeled and simulated as incompressible, deformable height fields. For each joint pose and pair of articulated bones we compute and report the location and size of the articular cartilage contact. The resulting contact is validated against *in vitro* data as described in Chapter 4.

In this study, the description of deformable articular contact is based on the simplified theory of contact from Blankevoort et al. [18] for thin layers of isotropic, linear-elastic material bonded to a rigid foundation. This implies three assumptions for the cartilage layer. First, the size of the contact area is assumed to be large relative to the cartilage thickness. Second, the cartilage layer is considered to be an isotropic, linear-elastic material. Third, the underlying bone is considered to be rigid. Deformable articular contact is used merely as a first order approximation of the behavior of the articular contact. The simplified contact description is then a first-order approximation of the relation between the normal surface stress σ_n and the surface displacement u_n normalized by the surface thickness t_n :

$$\sigma_n = S(u_n/t_n) \quad (5.1)$$

with

$$S = \frac{(1 - \nu)E}{(1 + \nu)(1 - 2\nu)} \quad (5.2)$$

where E is the elastic modulus and ν is Poisson's ratio. This description of articular contact deformation is strictly linear and will only be valid for small surface displacements. The stiffness parameter S is also known as the confined compression modulus or the aggregate modulus [78].

Following Blankevoort et al. [18], in the present model, where two bodies are in contact, the material properties of the cartilage on both bodies are assumed to be equal. The parameter t_n is then equal to the total thickness of the undeformed cartilage layer. The surface displacement u_n is the relative displacement of the height-field envelope, which is obtained after resolving cartilage penetration as described in Chapter 4.

The compressive cartilage contact force and moment are evaluated by integration of the contact stresses over the surface Ω :

$$f_{cartilage} = - \int_{\Omega} \int_{\Omega} \sigma_n \mathbf{n} \, d\Omega \quad (5.3)$$

$$m_{cartilage} = - \int_{\Omega} \int_{\Omega} \sigma_n \mathbf{c} \times \mathbf{n} \, d\Omega \quad (5.4)$$

where \mathbf{c} is the relative position vector of the cartilage contact with respect to the center of mass, and \mathbf{n} the direction of the normal.

Because we generated the cartilage height-field representation by extruding the bone triangular mesh representation (Chapter 4), we use the bone triangular grid to compute the support area of contact point and evaluate numerically the integrals in the cartilage contact force and moment formulas.

Ligaments are simulated as collections of elastic fibers characterized by their rest length and elasticity coefficient. Invasive studies on the human knee show that ligaments tear when stretched above 10% of their rest length [30]. This indicates that the ligament rest length should be above 90.9% of the fiber length when maximally stretched. To estimate the ligament rest length, we compute the maximum functional length of our fibers across the range of motion. Because the range of motion we use as input is unlikely to stretch ligaments to their maximum length, we estimate the fiber rest length should be around 95% of the computed maximum length. We then evaluate fiber lengthening across poses with respect to the estimated rest length. We consider a ligament to be lax in the poses in which its functional length is less than its estimated rest length.

Ligament contact forces and moments are computed separately for ligament-bone contact due to ligament insertions and ligament-bone contact due to ligament wrapping. Let m be the total number of ligament fibers modeled, and o_j the number of insertions for ligament j . The ligament insertion contact force and moment are computed following Blankevoort et al.[18] as:

$$f_{lig.ins.} = \sum_{j=1}^m \sum_{i=1}^{o_j} K_{lig} \left(\frac{l_j - l_{j0}}{l_{j0}} \right)^2 \mathbf{v}_{ji} \quad (5.5)$$

$$m_{lig.ins.} = \sum_{j=1}^m \sum_{i=1}^{o_j} \mathbf{s}_{ji} \times f_{lig.ins.ji} \quad (5.6)$$

where l_{j0} is the rest length of ligament j , v_{ji} is the ligament direction at the insertion site i , and s_{ji} the relative position vector of the ligament insertion with respect to the center of mass.

Similarly, the ligament wrapping contact force and moment are computed as:

$$f_{lig.wrap} = \sum_{j=1}^m \sum_{i=1}^{n_j} K_{lig} \left(\frac{l_j - l_{j0}}{l_{j0}} \right)^2 \alpha_{ij} p_{ji} \quad (5.7)$$

$$m_{lig.wrap} = \sum_{j=1}^m \sum_{i=1}^{n_j} s_{ji} \times f_{lig.wrap_{ji}} \quad (5.8)$$

where n_j is the number of points at which ligament j contacts the bone surface through wrapping, l_{j0} is the rest length of ligament j , p_{ji} is the sum of the ligament directions left and right of the contact point i , α_{ij} sums the projections of the left and right ligament directions in the p_{ji} direction (i.e., the directional sum of forces passing on either side of the point), and s_{ji} is the relative position vector of the ligament contact point with respect to the center of mass.

We use anatomy-book values for S_{cart} , the cartilage stiffness parameter, K_{lig} , the ligament elasticity coefficient, and the scaphoid mass (Table 5.1) [18, 98, 107]. In the absence of stiffness and elasticity constants for the human wrist, material parameters representative of the human knee were used for the wrist cartilage and ligaments. All ligaments were assumed to have the same elasticity coefficient.

Because each volume image of the joint is acquired in a quasi-static pose, the net force and moment acting on a bone at each pose are:

$$f_{net} = f_{cartilage} + f_{lig.ins.} + f_{lig.wrap} + f_{external} \quad (5.9)$$

$$m_{net} = m_{cartilage} + m_{lig.ins.} + m_{lig.wrap} + m_{external} \quad (5.10)$$

where $f_{external}$ and $m_{external}$ are the external force and momentum.

Because there are no muscle insertions on the scaphoid bone, at each sampled pose the values of $f_{cartilage} + f_{lig.ins.} + f_{lig.wrap}$ and $m_{cartilage} + m_{lig.ins.} + m_{lig.wrap}$ resulting for the scaphoid bone when using the anatomy-book parameters in Table 5.1 should be reasonably small — the same order of magnitude with the gravity on the scaphoid bone. We use this (f_{net}, m_{net}) balance observation to ascertain the plausibility of the anatomy-book parameter values in Table 5.1.

Table 5.1: Material properties used when simulating the wrist joint

parameter	value
S_{cart}	$4MPa$
K_{lig}	$10^6 N/m^2$
<i>Scaphoid mass</i>	$3g$

Forward Simulation

Unlike the load bearing joints of the hip and knee, where the articular cartilage is relatively thick, cartilage depth on the carpal surfaces involved in the present study is typically around 0.7mm. Assuming maximum deformation of the cartilage is 1% to 5% of its thickness, as determined in Chapter 4, then this equates to a maximum deviation in the surface shape of only around 0.007mm to 0.035mm, which is small in comparison to the size of the bones and is likely to be within the bounds of experimental error. For this reason, it was considered a reasonable simplification to employ non-deformable surfaces when modeling cartilage contact in the forward simulation of the joint.

Since the articular surfaces were non-deformable, any surface regions separated by less than 0.01mm during the simulation were deemed to be in contact. The lower bound on the proximity limit was introduced because we found that multiple small intersection regions occurred as a consequence of surface irregularities when two bone surfaces were brought closer together than 0.01mm, causing the bone to bounce repeatedly.

During the simulation, ligaments are allowed to deform up to 10% of their estimated rest-length (see Section 5.3.3). After this threshold, we treat ligament-bone contact as rigid.

In our forward simulations, the start conditions are given by the joint configuration in a given pose. All the bones except for the scaphoid bone are fixed; the scaphoid bone is free to move under the influence of contact forces and gravity. We use the explicit forward Euler method to integrate the scaphoid state through time (0.01s time-step), and the penalty-method (0.005 penalty factor) to resolve inter-penetrations [12].

The goal of the forward experiments described in the remainder of this section is to clarify the role played by ligaments in stabilizing the scaphoid bone. To this end, we first

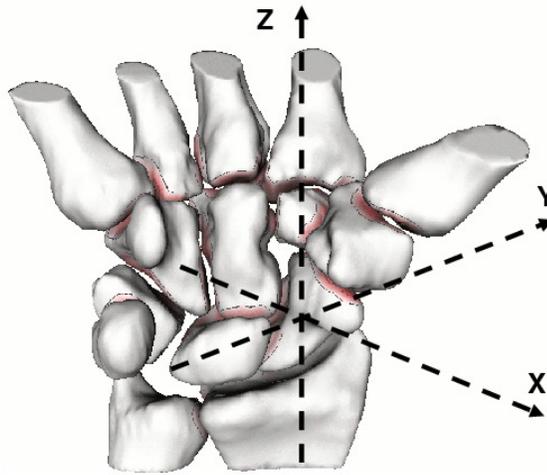


Figure 5.3: Coordinate system defined on the wrist. The origin of the coordinate system is the center of mass of the scaphoid bone; the axes of the coordinate system are aligned with the axes of the CT-scanner coordinate system.

validate the anatomy-book parameter values in Table 5.1 by evaluating (f_{net}, m_{net}) at each sample joint pose. We then simulate the motion of the scaphoid bone under the action of contact forces and gravity, when all the other bones are fixed. The start conditions of each simulation are given by a sample joint pose.

In order to test the stability of the joint, in each simulation below we specify a different direction for the gravitational force. We use the local coordinate system shown in Fig. 5.3 to describe the direction of this force.

We run two sets of forward experiments, each set using a different sample joint pose as start conditions. The first set uses as start conditions the wrist in the neutral pose. In this set of experiments we remove all ligamentous constraints. First, two simulations are run, with the gravitational force acting in the xy plane. In the first simulation the gravitational force pushes the scaphoid bone against the surrounding, fixed bones (direction: $[-1, +1, 0]$). In the second simulation, the gravitational force pushes the scaphoid bone outside the joint (direction: $[+1, -1, 0]$). For each simulation we report the total HAM motion — translation and rotation — of the scaphoid bone.

The second set of experiments uses as start conditions the wrist in an extended pose. We run two simulations with all ligamentous constraints removed, one simulation with the

gravitational force pushing the scaphoid bone towards the joint (direction: $[-1, -1, 0]$), the other with the gravitational force pushing the scaphoid bone outside the joint (direction: $[+1, +1, 0]$). We repeat the last simulation, this time taking into account ligament-related constraints. For each simulation we report the total scaphoid motion.

5.4 Results

The Functional Role of Wrist Ligaments

In this experiment we generated ligament-fiber paths for 15 ligament bundles in the human wrist. We computed the lengths of these paths across the range of wrist motion.

Four computed ligament-fiber paths with their lengths are shown in Fig. 5.4; note the fibers wrapping around the bone at the top insertion site. Also note that the resulting ligament fibers run more or less parallel to each other, justifying the band-like aspect of the ligament bundle.

Table 5.2 shows the resulting minimum and maximum fiber-path length across the range of motion for each of the 15 ligament bundles modeled. The table also shows the number of poses, for each bundle, at which the fiber length was above 95% of the computed maximum length.

Evaluation of the computed fiber lengths across the range of motion (Table 5.2) indicates that most ligament fibers — 10 out of 15 bundles — are active in only a few poses. That is, for most ligaments the computed fiber-length was above 95% of the maximum length in only a few poses (between 1 and 3 poses, out of 7 poses total). In contrast, note that fibers computed for the palmar mid-carpal scaphoid bundle and the dorsal mid-carpal bundle have almost constant length across the range of motion (10.5mm to 10.9mm, 3.8% relative elongation, and respectively 32.0mm to 33.1mm, 3.4% relative elongation).

Wrist Close-Pack Pose

In this experiment we evaluated the size of articular contact across the wrist range of motion. The aim of the experiment was to determine whether there is one pose (or more

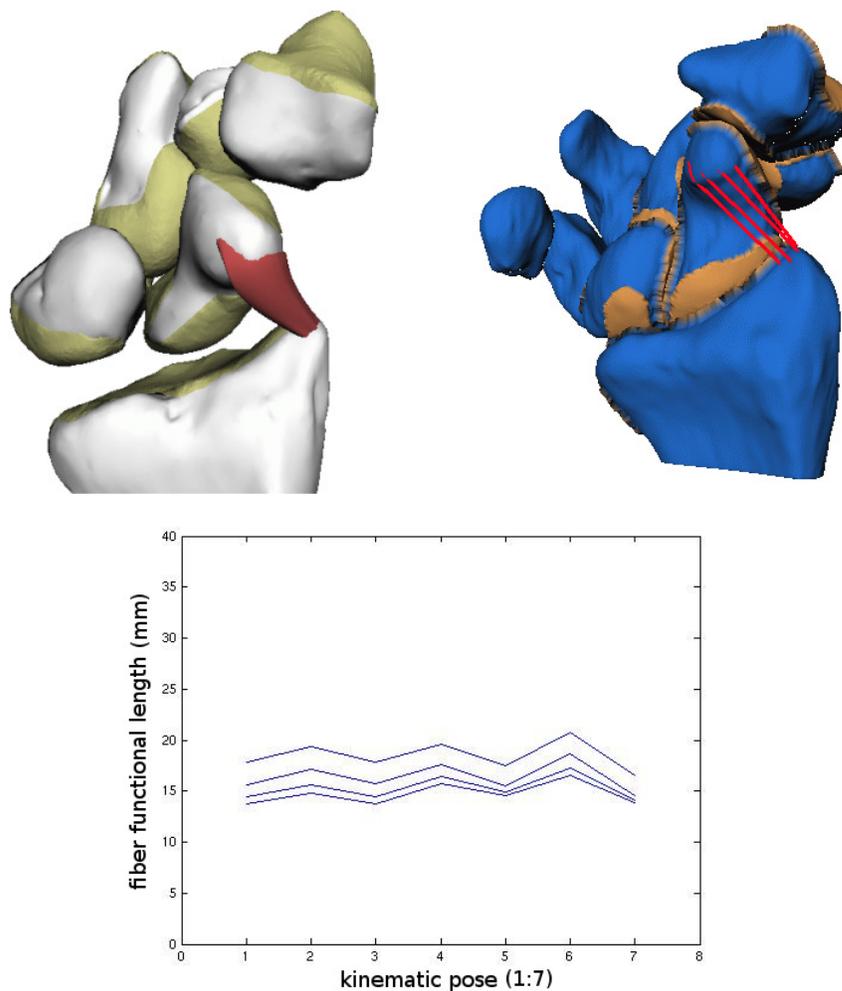


Figure 5.4: Four fibers in the same ligament bundle of the human wrist. Top: an artist's rendering of the palmar radial-scaphoid ligament bundle in a cadaver wrist (left, Primal Pictures; ligament in red, bones in white, cartilage in yellow) versus palmar radial-scaphoid ligament fibers generated computationally using *in vivo* bone surfaces and kinematics (right; ligament in red, bones in blue, cartilage in yellow). The pictured cadaver pose and the *in vivo* pose are slightly different; however, note the general similarity between the ligament geometry estimated invasively (left) and the ligament geometry estimated non-invasively (right). In particular, note the fiber wrapping around the bone at the top insertion site. Bottom: the length of the computational ligament fibers across seven joint poses. The *y*-axis in this plot covers the whole range of fiber lengths across 15 ligament bundles.

Table 5.2: Wrist ligament fiber lengths across the range of motion

Ligament bundle	Minimum length across fibers	Maximum length across fibers	Active poses/ total poses
Dorsal inter-carpal	15.0 ± 0.3mm	16.6 ± 0.6mm	2/7
Scapho-lunate inter-dorsal	2.6 ± 0.4mm	3.0 ± 0.4mm	1/7
Scapho-lunate inter-volar	3.5 ± 0.2mm	4.0 ± 0.2mm	1/7
Scapho-lunate inter-median	2.5 ± 0.2mm	3.3 ± 0.4mm	1/7
Palmar scapho- trapezium	3.2 ± 0.2mm	10.0 ± 0.9mm	1/7
Palmar mid-carpal scaphoid	10.5 ± 0.7mm	10.9 ± 0.8mm	7/7
Radio-carpal colateral	14.7 ± 1.2mm	18.3 ± 1.8mm	1/7
Radio-scapho- lunate	4.7 ± 0.5mm	10.4 ± 0.4mm	1/7
Radio-scapho- capitate	23.1 ± 2.6mm	27.4 ± 2.2mm	3/7
Dorsal mid-carpal	32.0 ± 0.1mm	33.1 ± 0.2mm	7/7
Dorsal mid-carpal trapezoid	32.8 ± 0.3mm	37.6 ± 0.3mm	3/7
Dorsal mid-carpal trapezium	37.4 ± 0.7mm	41.9 ± 1.1mm	4/7
Dorsal radio-carpal	24.5 ± 3.8mm	31.7 ± 2.4mm	1/7
Flexor retinum trapezium	33.6 ± 0.9mm	36.7 ± 1.9mm	4/7
Flexor retinum scaphoid	29.6 ± 1.1mm	32.9 ± 0.7mm	4/7

Table 5.3: Wrist articular contact size across 7 poses (mm^2)

Bone/Pose	1	2	3	4	5	6	7
cap	59.35	152	51.43	37.69	8.36	76.27	126.29
lun	43.46	169.32	18.48	42.62	6.97	105.35	93.23
rad	124.17	123.91	70.28	6.62	0.25	65.03	54.49
sca	120.04	250.3	87.78	30.29	8.11	52.04	284.01
tpd	24.74	77.69	17.98	39.49	68.15	77.98	81.76
tpm	6.04	75.47	0.00	31.05	62.33	24.19	62.43
Total	377.80	848.69	245.95	187.76	154.17	400.85	702.20

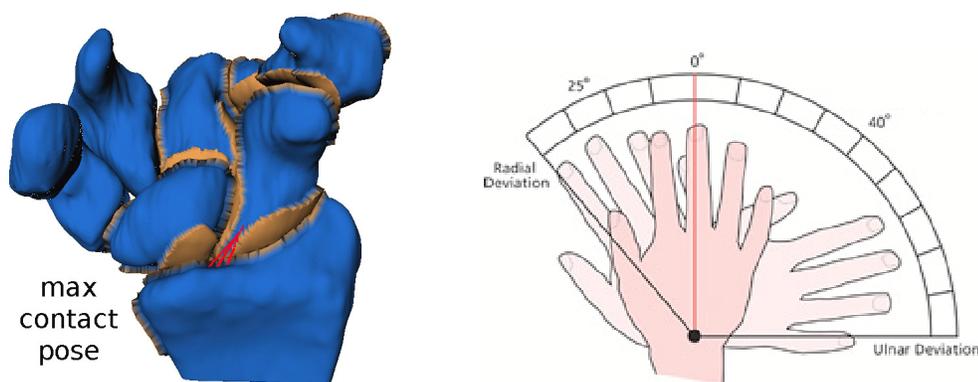


Figure 5.5: Wrist maximum contact pose. The wrist pose that results in maximum contact among bones (left; shown here palm facing) corresponds to a 15° ulnar deviation from the neutral pose (right).

poses) across the range of motion in which the joint bones attain maximum articular contact. Maximum-contact poses are also known as *close-pack* poses — poses in which the joint bones are closely packed.

Table 5.3 shows the size of articular contact across 7 poses. Note that pose 2 maximizes the contact among bones (total contact: 848.69mm^2). Pose 2 roughly corresponds to a 15° ulnar deviation from the neutral pose (Fig. 5.5).

Note that pose 7 also results in a significantly large value of total articular contact (total contact: 702.2mm^2 , 75% larger than the next largest value). Interestingly, both pose 2 and pose 7 fall within the radial-ulnar deviation range for which the contact-maximization approach of Sirkett et al. [98] predicted bone motion correctly.

Scaphoid Dislocation

To verify the plausibility of the parameter values in Table 5.1, we evaluated the values of $f_{cartilage} + f_{lig.ins.} + f_{lig.wrap}$ and $m_{cartilage} + m_{lig.ins.} + m_{lig.wrap}$ for the scaphoid bone, at each one of the sampled joint poses. The resulting values ($0.15\text{N} \pm 0.09\text{N}$ force, $0.74\text{Nm} \pm 0.51\text{Nm}$ torque) were reasonably small, comparable with the gravity on the scaphoid bone (0.03N , 0Nm).

Next, we simulated the motion of the scaphoid bone when the start conditions are given by the neutral pose. All bones except the scaphoid were fixed and ligament-related constraints were not taken into account. In each simulation the scaphoid remained stable, inside the joint. The total scaphoid motion when the gravitational force pushed the scaphoid bone towards the joint was (-0.56mm , 1.41°). The total motion when the gravitational force pushed the scaphoid bone outside the joint was (1.41mm , 1.45°). The simulations were stopped after 6s simulated time. These results indicate that in the neutral pose the scaphoid bone is locked in place by the surrounding bones.

Next, we simulated the motion of the scaphoid bone when the start conditions are given by an extended pose (pose 6 in Table 5.3). All bones except the scaphoid bone were fixed. Initially, ligament-related constraints were not taken into account. The total scaphoid motion when the gravitational force pushed the scaphoid bone towards the joint was (-0.07mm , 3.16°); the bone remained inside the joint. When the gravitational force pushed the bone outside the joint, the scaphoid dislocated from the joint and continued to slide away from the joint, in the direction of the gravitational force (direction: $[+1, +1, 0]$). The simulation was stopped after 15s. The total recorded scaphoid motion was (-5.28mm , -2.45°).

We repeated the extended-pose dislocation experiment, this time taking into account ligament-related constraints. At the start of the simulation, we recorded five active ligament bundles. As the simulation progressed, three of these ligaments were stretched to the maximum length, two became inactive, and one additional ligament was activated. Notably, two of the maximally-stretched ligaments wrap around the scaphoid bone: the radio-scapho-capitate ligament and the palmar mid-carpal scaphoid ligament. The constraints imposed

by ligament-bone contact through ligament insertion points and wrapping points were sufficient to prevent the scaphoid from becoming dislocated from the joint. The simulation was stopped after 6s; the total recorded motion was $(-0.99\text{mm}, -2.45^\circ)$.

Each simulation was computed in under 2 minutes on an AMD Athlon 64 X2 processor, 2 GHz.

5.5 Discussion

This study has successfully coupled data obtained *in vivo* with numerical simulation to develop a predictive tool for understanding the role of soft-tissues in the functioning of the wrist joint. In particular, our results predict soft-tissue contact within the joint, the existence of a close-pack wrist pose, and show that wrist ligaments play an important role in stabilizing the joint.

In our experiments, evaluation of fiber lengthening across the range of motion indicated that ligament fibers are relaxed in most poses. The interesting implication is that, in any given pose, only a few ligaments are active. Subsequent observation of an *in vitro* dissected specimen confirms this finding: wrist ligaments appeared to be lax in most poses. Our collaborators in the Department of Orthopedics expect that the poses in which ligaments become active will indicate the functional role of specific ligaments.

We also found in our experiments that in the close-pack pose contact among the bones was maximized. Because the contact is maximized, loads applied to the wrist in this pose get distributed on a larger contact area. The existence of a unique close-pack pose indicates therefore that there is a joint configuration in which the human wrist can sustain impact most effectively. This observation has interesting potential applications to athletic training and therapy planning.

The close-pack pose finding is difficult to validate otherwise than through visual examination of medical images of the bone configuration at each pose. However, we note that certain primates such as chimpanzees use a similar pose when walking on four limbs [90]. This suggests that the close-pack pose can indeed sustain loads effectively.

Our scaphoid dislocation predictions remain to be validated through invasive studies.

However, we note that our results are consistent with earlier studies on scaphoid motion [24], in which, in the absence of ligament wrapping, the scaphoid motion had to be artificially constrained to a plane to avoid joint collapse.

We used several modeling simplifications in our approach. Our models do not include the top layers of a joint — muscles, the neurovascular system, fat and skin. Ligament interconnectivity is not modeled, nor are the synovial sac and fluid. We approximate the viscoelastic properties of ligaments and cartilage. Bones are approximated as rigid bodies and bone-mass is assumed to be uniformly distributed. The models we generate could be enriched by considering intrinsic and extrinsic factors such as soft tissue composition, bone tissue composition, muscle forces, and joint compression.

When we estimate ligament rest-lengths, we assume that all ligaments are stretched to 95% of their maximum length in some of the sample joint poses. There are no guarantees that this sampling assumption is true in general: some ligaments may be lax in all the sampled poses. In theory, this issue could be addressed by acquiring a larger number of samples. Furthermore, we extrapolate knee anatomy-book knowledge to the wrist when we assume that ligaments tear when stretched above 10% of their rest length; the 10% threshold may not be correct. We also assume that wrist ligaments are characterized by the same elasticity coefficient. When more data about the elastic properties of wrist ligaments becomes available, the models we generate may need to be updated.

Although our contact model is capable of processing ligament-cartilage interaction, in the experiments presented in this chapter this type of interaction is not modeled. Ligament paths are generated over bone, not cartilage surfaces, primarily due to a boundary-related limitation in our cartilage map generation method. Visual analysis of the resulting wrist simulation indicates that during normal motion taut ligaments collide merely with the boundaries of cartilage maps. In anatomical data cartilage maps should recede towards their boundary — an aspect our kinematically-generated cartilage maps do not capture. However, the non-collision observation may not hold in the case of injured kinematics.

Our forward simulation implementation uses the forward Euler integration method and the penalty-method contact model. More intricate techniques exist; such techniques involve subdividing time to find the exact moment of collisions and then applying impulses or

resting forces as needed. Such methods provide in general more accuracy and stability, although simulation stability has not been an issue in our experiments so far.

5.6 Conclusion

We presented in this chapter a medical-image based method for constructing and simulating complex joint models. The method uses as input medical volume images of the joint bones, tracked motion and anatomy-book knowledge.

The resulting models and simulations are adaptable to individual differences: we use as input individual-specific bone surfaces and kinematics, and subsequently customize cartilage maps and ligament fibers according to the input information. The models we generate have predictive capabilities in terms of predicting soft-tissue contact within the joint and understanding the kinetic response of the joint.

The results presented in this chapter prove it is possible to couple *in vivo* data with numerical simulation to predict and understand the role of soft-tissues in the functioning of the wrist joint. Only one bone was allowed to move freely in our forward experiments, while the other bones were fixed. A natural next step is to apply the actual kinematics to some of the bones in the model and concentrate predictive efforts on one bone at a time.

Chapter 6

Diarthrodial Joint Markerless Cross-Parameterization and Biomechanical Visualization

6.1 Overview

Orthopedists invest significant amounts of effort and time trying to understand the biomechanics of diarthrodial 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 chapter, I present a framework that enables the cross-dataset visual exploration and analysis of diarthrodial 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.

I 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 diarthrodial 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.

6.2 Introduction

Diarthrodial 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 diarthrodial joint datasets, the inherent complexity of these data motivates a variety of modeling and visualization techniques designed to assist orthopedics researchers in their analysis.

In the context of diarthrodial 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 invasively, *in vitro*. The invasive procedure involves 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 *in vitro* bony surfaces (Fig. 6.1). Our

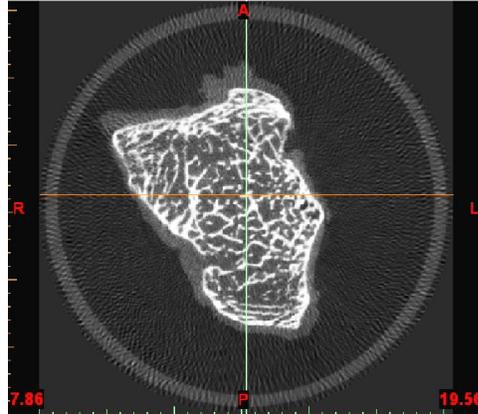


Figure 6.1: 2D slice through a μ CT-volume image of bone and articular cartilage. The imaged cartilage (bone tissue shown in white, cartilage in soft grey surrounding the bone outline) can be represented as a height-field on the supporting bony surface. The height-field is defined as a collection of 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.

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. Both applications — comparative analysis and *in vivo* - *in vitro* data fusion — require establishing a pairwise correspondence between datasets.

Unfortunately, diarthrodial 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 chapter, we present a framework that enables the cross-dataset visual exploration and analysis of diarthrodial joint biomechanics. Our key contribution is a markerless method for establishing pairwise correspondence between individual-specific diarthrodial joint 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.

6.3 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 [17], 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 [106, 112, 102]. In the absence of salient features, pairwise correspondence is often established using user-defined landmarks [26, 23]; however, the landmarking operation is time-consuming, error-prone and subjective even in 2D. In 3D, manual landmarking can become impractical.

Golland et al. [46] 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 [22] 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 diarthrodial shapes discussed in this chapter. In fact, a comparative study [101] 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 [31] or DetCov [59].

In contrast, the markerless pairwise correspondence method proposed in this chapter 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 [16, 53], 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 [71].

Cross-parameterization. A number of statistical methods produce cross parameterizations across training sets of shapes [26, 46, 31, 59]. In these approaches, a statistical model of shape variation is learned from a training database. The focus of the present chapter is, however, the *pairwise* cross-parameterization of shapes, when a training database may not be available.

Previous approaches [88, 60, 93] 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 Diarthrodial Joints. Cross-individual diarthrodial 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 [28, 99], mainly by presenting to the user animations of a specific joint. The lack of salient features makes finding correspondences in diarthrodial 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 diarthrodial 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 [72].

6.4 Methods

Figure 6.2 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.

6.4.1 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.

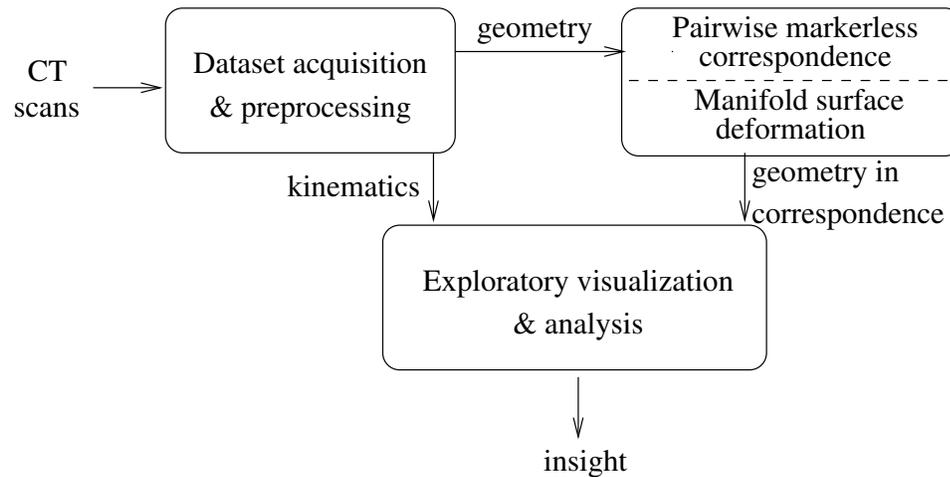


Figure 6.2: Biomechanics visual analytics framework. Subject-specific CT-datasets are first processed in order to extract the relevant joint geometry and kinematics. Two such datasets are then set automatically in correspondence. We use the correspondence to further explore and analyze dataset differences.

Through manual segmentation, thresholding, and user interaction, bone surface points are extracted from an arbitrarily-designated reference CT volume image [28]. Each bone is then tracked accurately through the sequence of remaining CT volume images (Chapter 2). The tracking procedure reports relative bone-motion from one articulation pose to another. Each resulting dataset is thus individual-specific, and consists of the diarthrodial 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 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 [43] 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

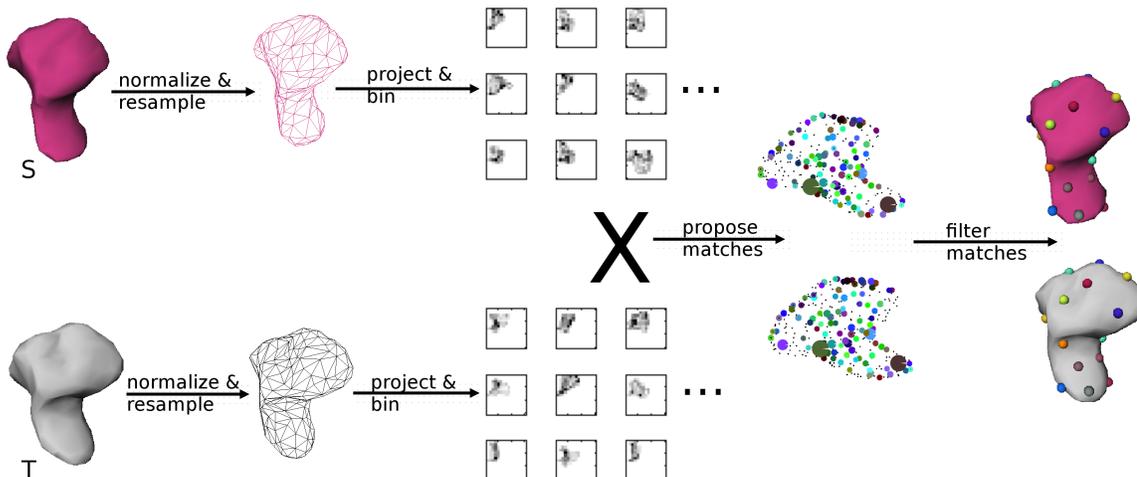


Figure 6.3: Markerless correspondence pipeline. Corresponding bone instances S and T are first normalized and resampled. For each surface point, we project the bone instance on a local 2D subspace and we generate a 2D histogram (for details on the project and bin operation, see Fig. 6.4). We evaluate 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.

the support for evaluating kinematic changes in the articulation. We use the Closest Point Transform (CPT) level-set software package [73] to generate the distance field representation from the surface representation. A validation study performed on CT-imaged phantom data [109] — 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.

6.4.2 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.

Pairwise Markerless Correspondence

Figure 6.3 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 software [43]. 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 [17].

The fundamental shape element we use to generate pin-point correspondences is an oriented point, a 3D vertex with an associated direction [53]. 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:

$$\beta(v) = n \cdot (v - p) \quad (6.1)$$

$$\alpha(v) = \sqrt{(v - p) \cdot (v - p) - \beta(v)^2} \quad (6.2)$$

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. 6.4). 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.

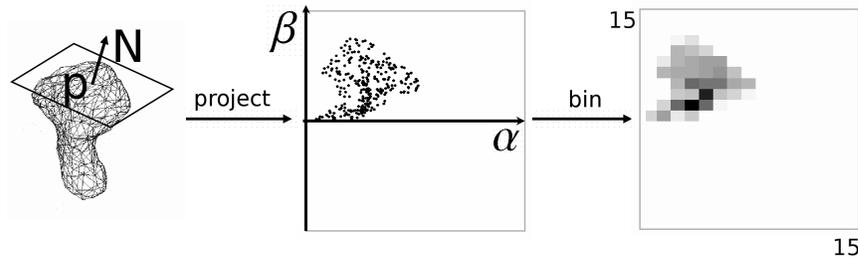


Figure 6.4: Project and bin operation. For each point p on the mesh, we consider the local coordinate system defined by the plane tangent to p and the surface normal N at that point. All other vertices v in the mesh are cylindrically-projected on this local system; we compute the α and β cylindrical coordinates (see text for details) of each vertex. Projected points are finally binned in a 2D histogram.

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 :

$$K_{match}(s, t) = K_{correl}(hist(s), hist(t)) + \frac{1}{\|p(s) - p(t)\|} \quad (6.3)$$

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 $K_{match}(s, t)$ 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.

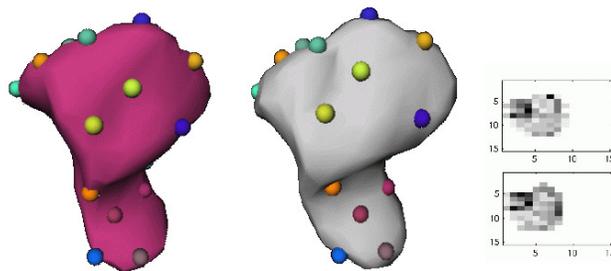


Figure 6.5: Two wrist bones belonging to different human subjects and their corresponding pin-points. Pin-points in correspondence are shown in the same color. Right: the 2D-histograms generated by the two pink corresponding pin-points on the left are remarkably similar, indicating good shape correlation at these points. Our markerless method generated this bipartite set of 25 Poisson-disk distributed correspondences in only a few minutes.

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. 6.5. The pairwise markerless correspondence method is implemented in less than 2,000 lines of Matlab code.

Manifold Deformation

To align the surfaces we begin by fitting a default manifold surface to the source geometry S , as described in [47, 48]. Our manifolds are smooth, locally parameterized, C^2 continuous surfaces [47]. 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 [47] reference, the fitting process essentially “shrink wraps” the default manifold around

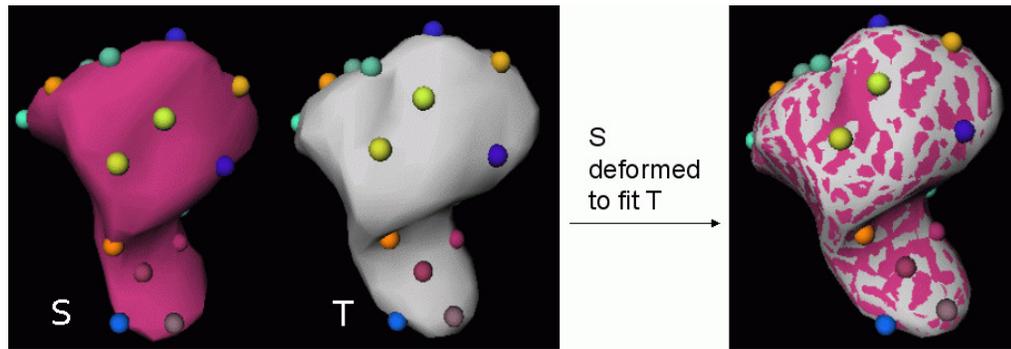


Figure 6.6: Manifold deformation: default manifold surface fitted to S with pin-points (left), the original mesh T with corresponding pin-points (middle), and the original mesh T overlaid with the result of deforming the S manifold to fit T while preserving the pin-point correspondence (right); in the right image, the pin-points of T and the deformed manifold are overlapped. The resulting T manifold is free of self-intersections/folds. The manifold defines a mutually-consistent parameterization between the S surface and the T surface.

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 6.6 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 overlaid with the result of deforming the S manifold to fit T while preserving the pin-point correspondence. Note that the resulting T manifold is free of self-intersections/folds. The manifold deformation defines a mutually consistent parameterization between the S surface and the T surface.

6.4.3 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 kinematics on the bone surfaces, as described in [67]. Summarizing this reference, for each bone and joint 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. 6.7).

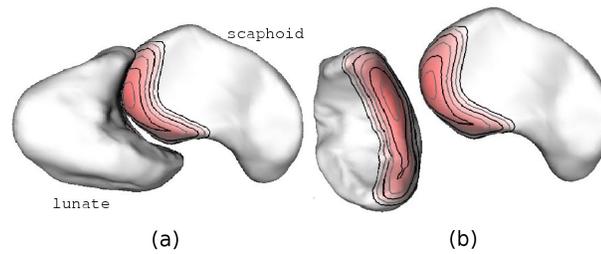


Figure 6.7: A normal scapholunate joint. Bones are color mapped and contoured. The saturation of red (darker region in black and white) on bone surfaces represents the distance to the nearest point on the opposite bone. Redder regions are closer. The maximum distance visualized is 5 mm. Contour lines are drawn at 0.5 mm intervals. (a) Bones in their correct anatomical context. (b) Bones rotated to show articulated surfaces more clearly.

6.5 Results

6.5.1 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 shown in Fig. 6.6 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 tenfold 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 [42]. 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 — described in detail in [42] — are associated with a color map in order to plot similarity on the surface of the object (Fig 6.8). 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 6.8 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 localized, well-defined changes in the articulation at the distal

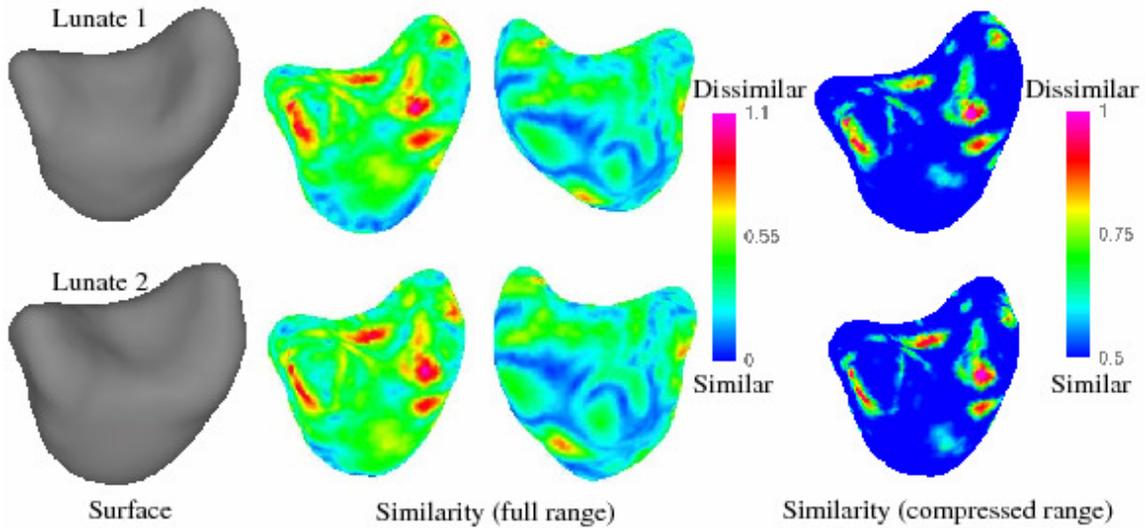


Figure 6.8: Curvature comparison of source and target manifold surfaces: deforming the manifold surface from Lunate 1 (L1) to Lunate 2 (L2) maps regions on L1 to similar curvature regions on L2. The small areas of dissimilar curvature appear to correlate with actual anatomical shape variation between the L1 and L2 datasets.

radioulnar joint and potential soft-tissue constraints (see Chapter 3). Further numerical analysis of the location and size of inter-bone joint spaces across the 9 individuals showed these changes were statistically significant [29].

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

6.5.2 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

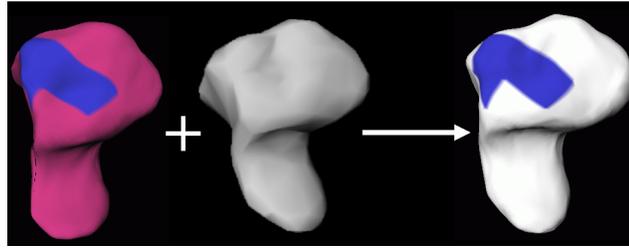


Figure 6.9: Surface detail transfer (ligament insertion site) between the two hamate bone instances shown in Fig. 6.5. The mutually-consistent parameterization between the two bone instances results in no folds or tears during transfer.

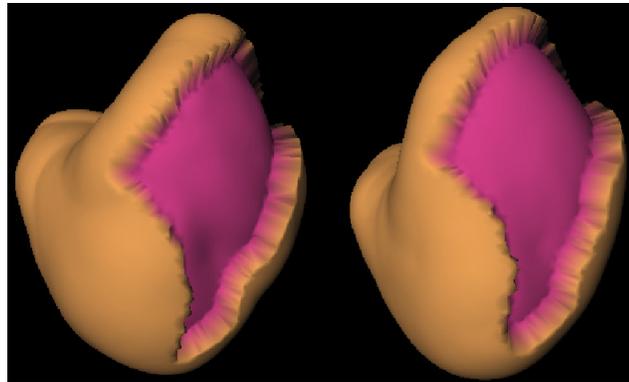


Figure 6.10: Cartilage transfer between two lunate bones (bones shown in pink, cartilage in tan). The consistent parameterization between the two lunates results in no folds and no tears in the cartilage map during transfer.

articular cartilage from one individual to another. In Fig. 6.9, 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. 6.10, 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 6.11 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 may not intervene at all, or they may remove completely the bottom fragment from the joint [7]. 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 joint poses) and compare location and size using the mutually-consistent parameterization (Fig. 6.12). Visual analysis of the lunate bone with respect to the scapho-lunate articulation shows no significant differences between the uninjured and injured joint (Fig. 6.13 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. 6.13 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.

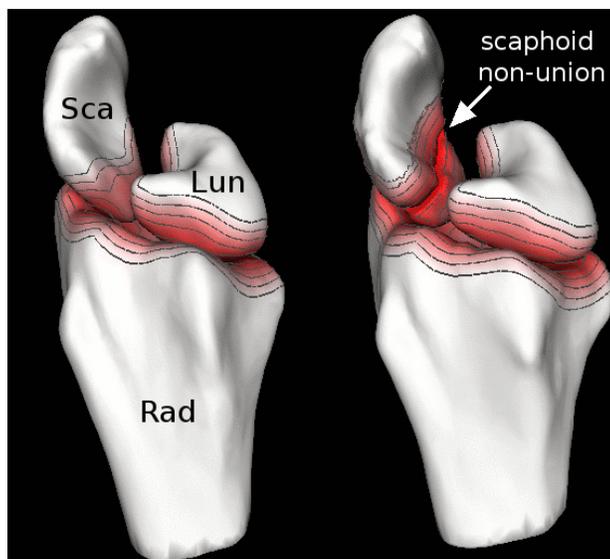


Figure 6.11: Right-wrist and left-wrist radioscapholunate joints from the same individual. Following injury, the left-wrist scaphoid bone has been fractured in two. Bones are color-mapped and contoured. The color saturation on bone surfaces indicates the distance to the nearest point on the opposite bone or bone-fragment; darker regions are closer.

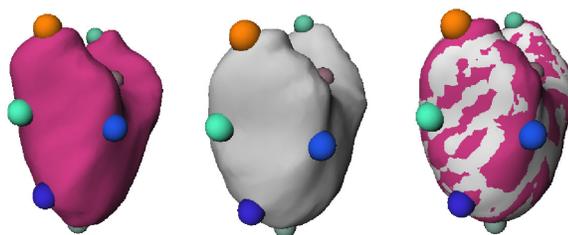


Figure 6.12: Pin-point and manifold deformation between the left and right lunate bones in Fig. 6.11. Left-wrist lunate (injured joint) with pin-points (left), right-wrist lunate with corresponding pin-points (middle), and the original right-wrist lunate overlaid with the result of deforming a manifold 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.

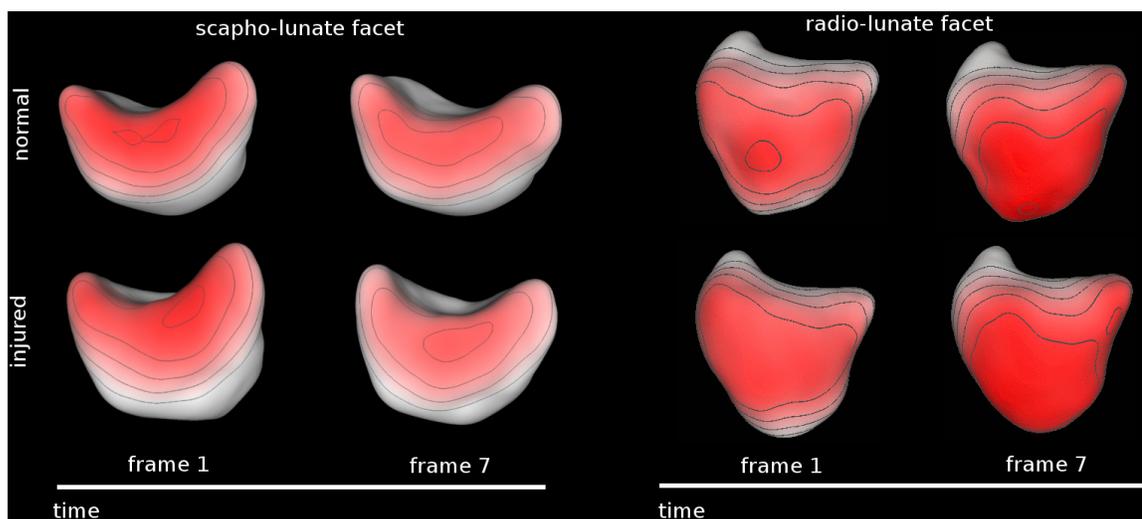


Figure 6.13: Kinematic analysis of the radioscapulohunate joint shown in Fig. 6.11 (top row: lunate bone in the uninjured joint, bottom: lunate bone in the injured joint). The first two columns show the lunate facet articulating with the scaphoid bone; tracing the inter-bone joint-spacing area on the 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.

6.6 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 diarthrodial-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 a 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.

6.7 Conclusion

I presented in this chapter a framework for the cross-dataset visual exploration and analysis of diarthrodial 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 rotational symmetry in the first order ellipsoid. In consequence, our method is applicable to diarthrodial joints like the human wrist or the ankle. Results demonstrate the technique useful in the study of normal and injured anatomy and kinematics of diarthrodial 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 diarthrodial-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 diarthrodial geometry. Our markerless geometry correspondence method successfully addresses this challenge. The method should be useful as a preprocessing step in the statistical shape analysis of wrist and ankle data.

Chapter 7

Conclusion

This dissertation presents a data-driven framework for the predictive modeling of diarthrodial joints. We have shown that the data-driven approach allows us to generate subject-specific, quantifiable, predictive models of complex joints, in spite of current measuring limitations. The resulting computational models use a hybrid data representation, combining volumetric, meshless point-based, and manifold-based components. Each component representation has strengths for various aspects of the modeling and we combine them to achieve relatively simple and sufficiently accurate estimations of biologically relevant measurements.

We instantiated the framework in a system consisting of a collection of computational modeling, analysis and visualization tools. These tools can be used to predictively model and understand the anatomy and dynamics of joints, with a potential positive impact on medical diagnosis and treatment of pathology and injury.

The system we developed takes as input medical images of a joint and generates a model and simulation of the joint. The model is adaptable to individual differences: it uses individual-specific bone shapes extracted from the medical images, and customizes cartilage maps and ligament fibers according to the individual-specific bone geometry and motion. For example, should the individual develop arthritis (erosion of articular cartilage), the disease would impact inter-bone spacing measurements, and therefore the thickness of the cartilage maps we generate and ultimately the simulated motion of the joint. Similarly, should the individual undergo athletic training, training could alter the sampled range of

motion, and be reflected in the length of the soft-tissue fibers we generate.

The resulting models and simulations are quantifiable, in the sense that they quantify the difference between two individuals, or the differences between a normal and an injured joint. The models also have predictive capabilities, in terms of predicting soft-tissue contact within the joint.

Specific contributions of this work include:

- a framework for the data-driven modeling of anatomical joints; the resulting models are adaptable to individual differences, are quantifiable, and have predictive capabilities;
- an instantiation of the framework in a system for the predictive modeling of joints; the models can compute end-conditions for deformable tissues — cartilage contact and ligament function — as a function of joint motion;
- a sub-voxel accurate method for tracking bone-motion from sequences of CT scans. Accuracy improvements of more than 74% over the previous state-of-the-art technique, when compared to the ground-truth motion computed from marked cadaver data, enables the analysis of soft-tissue deformation with motion. The method has also lead to the creation of a wrist motion database of unprecedented detail; the database spans multiple species, including humans and baboons;
- a novel computational model of ligament fibers that models soft-tissue wrapping on bones and the afferent motion constraints. The model requires as sole input bone geometries and kinematics, surpassing current imaging limitations;
- a novel computational volume-preserving height-field model of cartilage maps that only requires CT-imaged bone geometries and joint kinematics; this model allows us to overcome current imaging technology limitations;
- a hybrid representation of diarthrodial joints, combining volumetric, meshless point-based, and manifold-based components. Each representation has strengths for various aspects of the modeling and we combine them in unique ways to achieve estimations of biologically relevant measurements;

- a markerless method for cross-individual dataset analysis. The method also allows users to combine *in vitro* and *in vivo* joint data;
- interactive tools for the exploratory analysis of joint data;
- quantitative metrics for characterizing joint kinematics;
- application to several clinical cases: distal radioulnar joint malunion, wrist close-pack pose, scaphoid non-union.

There are limitations to our work, some due to the data currently available, and some due to our computational approach. From the complex anatomy of a joint, we model only the bones, cartilages, and ligaments. Muscle-tendon units could be easily incorporated in our model. However, the forces these units exercise can not be currently inferred automatically from the data; these external forces would have to be specified as an input, just like in other joint modeling systems. We do not model skin, fat, the neurovascular system, or the sac of synovial fluid surrounding the joints. There are limits to the accuracy and applicability of our tracking procedure, as discussed in Chapter 2; the method could be further improved by modeling a more complex tissue structure — for example, one that would distinguish between trabecular and cortical bone. There are further limitations in our computational modeling of soft-tissues as discussed in Chapters 3 and 4. Someday, more sophisticated and accurate deformation models based for example on finite element modeling techniques may replace our deformation approach, just like progress in imaging technology may supersede our kinematics-based estimates of soft-tissue geometry. We made several further assumptions in our dynamics simulation, as detailed in Chapter 5.

In terms of directions of future work, we note that the ideal, fully-predictive joint model described in Chapter 1 still does not exist. That ideal model would allow doctors to simulate sophisticated surgical procedures and therapies, and animators to generate automatically-tuned, convincing character animations.

A natural next phase is trying to develop bone-cartilage-ligament models that can be simulated dynamically. Our dynamics experiment in Chapter 5 shows it is possible to simulate the motion of one bone under the action of gravity. Next, motion interpolated between two nearby sampled poses could be applied to all the bones but one; the free bone

should also interpolate the sampled poses. Next, we should be able to replicate the effect of wrist injury on bone kinematics. For example, we could virtually fracture the scaphoid in the healthy scapho-lunate joint of Chapter 6 and validate the resulting simulation by comparing it with the sampled motion of the matching injured scapho-lunate joint, also shown in Chapter 6.

The development of fully-predictive joint models opens further research directions. For example, what level of anatomical realism is necessary or useful when we design orthopedic implants and robotic arms, versus animating Pixar-like characters? To perform such comparisons, it would be necessary to first integrate into the joint model the superior layers — muscles, fat and skin. Validation and quantitative evaluation tools will be increasingly necessary. A step further, we investigate biological systems of increasing complexity. Such systems have multiple interacting subsystems — from cellular to the full multi-level virtual human [4]— and are currently developed independently [105, 116, 118]. To interface such subsystems, modeling and simulation standards will have to be proposed and developed.

The collection of tools presented in this dissertation solves significant challenges in the process of integrating measured data into the predictive modeling of diarthrodial joints. The developed tools and applications should be beneficial in the study of pathology and injuries, including arthritis, ligament tears, bone fractures, and surgical reconstructions. The tools and methodologies we demonstrate on forearms and wrists will be generally useful for the study of bone, cartilage and ligament interactions in other complex multi-articular joints, including the foot and spine, as well as in simpler joints such as the knee, elbow, and human shoulder. The tools are also applicable to animal studies, in basic biology research. Ultimately, this work creates a modeling approach that has the potential to more simply and efficiently explain and predict the underlying biomechanics of musculoskeletal systems.

Appendix A

Orthopedics Terminology

orthopedics also *orthopaedics*, the field of medicine and surgery that deals with the injuries and other disorders of the musculoskeletal system;

musculoskeletal of, relating to, or involving both musculature and skeleton;

diarthrodial joints also known as diarthroses or synovial joints (freely movable articulations). This class includes the greater number of the joints in the body; the other two classes are: synarthroses or immovable joints, and amphiarthroses or slightly movable joints. In a diarthrodial joint the contiguous bony surfaces are covered with articular cartilage, and connected by ligaments;

soft tissue tissues that connect, support, or surround other structures and organs of the body. Soft tissue includes muscles, tendons, fibrous tissues, fat, blood vessels, nerves, and synovial tissues. In this dissertation, soft-tissue denotes articular cartilage and ligaments;

articular cartilage hydrated soft tissue covering the surfaces of bones at the joint. Articular cartilage is multi-layered. A thin superficial layer provides a smooth surface for bones to slide against each other. Deeper than that are intermediate layers, which are mechanically designed to absorb shocks and distribute the load efficiently. The deepest layer is highly calcified, and anchors the articular cartilage to the bone;

ligament in its most common use, a band of tough fibrous connective tissue composed mainly of long, stringy collagen fibres. Ligaments connect bones to other bones to form a joint and provide joint stability. (They do not connect muscles to bones; that is the function of tendons.) Some ligaments limit the mobility of articulations, or prevent certain movements altogether. The whole of a diarthrosis is contained by a ligamentous sac;

synovial fluid a thick, stringy fluid found in the cavities of diarthrodial (*synovial*) joints. With its egg-like consistency (*synovial* comes from Latin for "egg"), synovial fluid reduces friction between the articular cartilage and other tissues in joints to lubricate and cushion them during movement;

trabecular tissue also known as spongy bone; a type of osseous tissue with a low density and strength but very high surface area, that fills the inner cavity of bones;

cortical tissue also known as compact bone; one of two main types of osseous tissues. Cortical bone is dense and forms the surface of bones, contributing 80% of the weight of a human skeleton. It is extremely hard, formed of multiple stacked layers with few gaps. The other major type of bone tissue is trabecular tissue;

focal changes localized (as opposed to whole) joint changes. For example, a focal lesion in cartilage is a defect that is well defined, localized, such as a small hole punched through the cartilage. Changes in the location of the contact area between articulating bones may indicate potential focal damage;

in vitro outside the living body and in an artificial environment;

in vivo inside the living body;

carpus wrist; the bones of the wrist;

carpal of or relating to the carpus;

scaphoid one of the eight wrist (*carpal*) bones;

lunate one of the eight wrist (*carpal*) bones;

radius one of the two forearm bones;

radial of or relating to the radius;

ulna one of the two forearm bones;

ulnar of or relating to the ulna;

metacarpals a bone of the part of the hand or forefoot between the carpus and the phalanges that typically contains five more or less elongated bones when all the digits are present;

patellar relating to the *patella*, a thick flat triangular movable bone that forms the anterior point of the knee and protects the front of the joint – called also kneecap;

temporomandibular of, relating to, being, or affecting the joint between the temporal bone and the mandible (jaw) that allows for the movement of the mandible;

meniscus a fibrous cartilage within a joint especially of the knee;

distal situated away from the point of attachment or origin or a central point, especially of the body;

proximal next to or nearest the point of attachment or origin, a central point, or the point of view; especially: located toward the center of the body;

volar relating to the palm of the hand or the sole of the foot; specifically: located on the same side as the palm of the hand;

palmar of, relating to, or involving the palm of the hand;

dorsal of, relating to, or involving the back of the hand;

neutral the default, rest position of a joint;

pronation rotation of the forearm so that the palm faces backwards or downwards;

supination rotation of the forearm and hand so that the palm faces forward or upward;

pronosupination rotation of the forearm and hand, starting with the palm facing backwards or downwards and ending with the palm facing forward or upward;

flexion a bending movement around a joint in a limb (as the knee or elbow) that decreases the angle between the bones of the limb at the joint; also, a forward raising of the arm or leg by a movement at the shoulder or hip joint;

extension an unbending movement around a joint in a limb (as the knee or elbow) that increases the angle between the bones of the limb at the joint;

arthritis inflammation of joints due to infectious, metabolic, or constitutional causes; also : a specific arthritic condition;

osteoarthritis arthritis marked by degeneration of the cartilage and bone of joints.

Bibliography

- [1] NAG Fortran Library Routine Document, E04JYF
<http://www.nag.co.uk/numeric/fl/manual/pdf/e04jyf.pdf>
- [2] NAG Fortran Library Routine Document, E04CCF
<http://www.nag.co.uk/numeric/fl/manual/pdf/E04/e04ccf.pdf>
- [3] NAG Fortran Library Routine Document, E04UCF
<http://www.nag.co.uk/numeric/fl/manual/html/e04ucf.html>
- [4] The Digital Human Project, <http://fas.org/dh/>
- [5] Acosta, R., Hnat, W., Sheker, L.R., Distal radio-ulnar ligament motion during supination and pronation, *The Journal of Hand Surgery (British and European Volume)*, 18B, pages 502–505, 1993.
- [6] Adams, B.D. and Grosland, N.M., Biomechanical Evaluation of a New Prosthesis Design for Use in Total Wrist Arthroplasty, <http://www.uihealthcare.com/depts/med/orthopaedicsurgery/wristreplacement/pdf/evaluation.pdf>
- [7] Akelman, E., personal communication to J. Crisco, 2006.
- [8] Alexandroff, P., *Elementary Concepts of Topology*, Dover Publications Inc., 1961.
- [9] Ateshian, G.A., A b-spline least-squares surface fitting method for articular surfaces of diarthrodial joints, *ASME Journal of Biomechanical Engineering*, 115, pages 366–373, 1993.

- [10] Ateshian, G.A., Generating trimmed b-spline models of articular cartilage layers from unordered 3d surface data points, Proc. ASME Bioengineering Conference, pages 217–218, 1995.
- [11] Baltsan, A. and Scharir M., On shortest paths between Two Convex Polyhedra, JACM, 35(2), pages 267–287, 1988.
- [12] Baraff, D. and Witkin, A., Large Steps in Cloth Simulation, Siggraph Computer Graphics Proceedings, Annual Conference Series, pages 43–54, 1998.
- [13] Barzel, R., Physically-Based Modeling for Computer Graphics: a Structured Approach, Academic Press, 1992.
- [14] Beaulieu, C.F., Hodge, D.K., Bergman, G., Butts, K., Daniel, B.L., Napper, C.L., Darrow, R.D., Dumoulin, C.L., and Herfkens, R.J., Glenohumeral Relationships during Physiologic Shoulder Motion and Stress Testing: Initial Experience with Open MR Imaging and Active Imaging-Plane Registration, Radiology, Vol. 212, pages 699–705, 1999.
- [15] Beck, J.M., Farouki, R.T., and Hinds, J.K., Surface Analysis Methods, IEEE Computer Graphics and Applications, 6, pages 18–37, 1986.
- [16] 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), pages 509–522, 2002.
- [17] 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), pages 239–256, 1992.
- [18] Blankevoort, L., Kuiper, J. H., Huiskes, R., and Grootenboer, H. J., Articular contact in a three-dimensional model of the knee, Journal of Biomechanics, Vol. 24, No. 11, 1991, pages 1019–1031
- [19] Bonnell, K.S., Duchaineau, M.A, Schikore, D.A., Hamann, B., Joy, K.I., Material Interface Reconstruction, IEEE Trans. on Visualization and Computer Graphics, 9(4), pages 500–511, 2003.

- [20] Boyd, S.K., Ronsky, J.L., Lichti, D.D., Salkauskas, D., and Chapman, M.A., Joint surface modeling with thin-plate splines, *Journal of Biomechanical Engineering*, 121, pages 525–532, October 1999.
- [21] Breen, D., Mauch, S. and Whitaker, R., 3D scan conversion of CSG models into distance volumes, *Proc. 1998 IEEE Symposium on Volume Visualization*, pages 7–14, 1998.
- [22] Brechbuhler, C., Gerig, G., Kubler, O., Parametrization of closed surfaces for 3-D shape description, *Comp. Vision, Graphics, and Image Proc.* 61, pages 154–170, 1995.
- [23] Brett, A., Taylor, C., Construction of 3d shape models of femoral articular cartilage using harmonic maps, *Proc MICCAI 2000*, pages 1205–1214, 2000.
- [24] Carrigan, S.D., Whiteside, R.A., Pichora, D.R., and Small, C.F., Development of a Three-Dimensional Finite Element Model for Carpal Load Transmission in a Static Neutral Posture, *Annals of Biomed. Eng.*, Vol 31, pages 718–725, 2003.
- [25] Clarkson, K., Approximation Algorithms for Shortest Path Motion Planning, *Proc. 19th Annual SIGACT Symposium*, pages 56–65, 1987.
- [26] Cootes, T., Taylor, C., Cooper, D., Graham, J., Active shape models - their training and application, *Comp. Vis. Image Under.* 61, pages 38–59, 1995.
- [27] Crisco, J.J. and McGovern, R.D., Efficient calculation of mass moments of inertia for segmented homogeneous 3D objects, *J. of Biomechanics*, 31(1), pages 97–102, 1998.
- [28] Crisco, J.J., McGovern, R.D. Wolfe, S.W., A non-invasive technique for measuring in vivo three-dimensional carpal bone kinematics, *J. Orthopaedic Research*, 17(1), pages 96–100, 1999.
- [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] Crisco, J., personal communication, 2007.

- [31] Davies, R., Twining, C., Cootes, T., Waterton, J., Taylor, C., A minimum description length approach to statistical shape modeling, *IEEE Trans. Med. Imaging* 21, pages 525–537, 2002.
- [32] Delp, S.L. and Loan, J.P., A computational framework for simulating and analyzing human and animal movement, *IEEE Computing in Science and Engineering*, Vol. 2, pages 46–55, 2000.
- [33] Donahue, T.L.H., Hull, M.L., Rashid, M.M. and Jacobs, C.R., How the stiffness of meniscal attachments and meniscal material properties affect tibio-femoral contact pressure computed using a validated finite element model of the human knee joint. *J Biomech* 36, pages 19–34, 2003.
- [34] Donzelli, P.S., Gallo, L.M., Spilker, R.L., Palla, S., Biphasic finite element simulation of the TMJ disc from in vivo kinematic and geometric measurements, *J Biomech* 37, pages 1787–1791, 2004
- [35] Eckstein, F., Gavazzeni, A., Sittek, H., Haubner, M., Losch, A., Milz, S., Englmeier, K.H., Schulte, E., Putz, R., and Reiser, M., Determination of knee joint cartilage thickness using three-dimensional magnetic resonance chondro-crassometry (3d mr-ccm), *Magn. Reson. Med.*, 36, pages 256–265, 1996.
- [36] Ekenstam, F., Anatomy of the Distal Radioulnar Joint, *Clinical Orthopaedics*, 275, pages 14-18, 1992.
- [37] Feipel, V. and Roosen, M., Three-dimensional motion patterns of the carpal bones: an in vivo study using three-dimensional computed tomography and clinical applications, *Surgical and Radiologic Anatomy*, Vol 21 (2), pages 125–131, 1999.
- [38] Feipel, V., Rooze, M., Louryan, S., and Lemort, M. Bi- and three-dimensional CT study of carpal bone motion in lateral deviation, *Surg. Radiol. Anat.*, 14, pages 381–389, 1992.
- [39] Ferlic D.C., Jolly S.N., Clayton M.L., Salvage for failed implant arthroplasty of the wrist, *J Hand Surg*, 17(5), pages 917-23, 1992.

- [40] Frisken, S.F., Perry, R.N., Rockwood, A.P. and Jones, T.R., Adaptively sampled distance fields: A general representation of shape for computer graphics, SIGGRAPH 2000 Conference Proceedings, 2000.
- [41] Gardiner, J. C. and Weiss, J. A., Subject-specific finite element analysis of the human medial collateral ligament during valgus knee loading, *J Orthop Res* 21, pages 1098–1106, 2003.
- [42] Gatzke, T., Zelinka, S., Grimm, C., and Garland, M., Curvature maps for local shape comparison, In *Shape Modeling International*, pages 244– 256, June 2005.
- [43] Geomagic Inc., Research Triangle Park, NC, 27709 USA, <http://www.geomagic.com/>
- [44] Gibson, S., Using Distance Maps for smooth surface representation in sampled volumes, *Proc. 1998 IEEE Volume Visualization Symposium*, pages 23–30, 1998.
- [45] Gibson, S. and Mirtich, B., A Survey of Deformable Modeling in Computer Graphics, Tech. Report No. TR-97-19, Mitsubishi Electric Research Lab., Cambridge, MA, November 1997.
- [46] 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*, pages 72–82, 2000.
- [47] Grimm, C., Crisco, J.J., and Laidlaw, D.H., Fitting locally parametric surfaces to 3D point clouds, *ASME Journal of Biomechanical Engineering*, 124(1), pages 136–140, 2002.
- [48] Grimm, C.M., Hughes, J., Modeling Surfaces of Arbitrary Topology using Manifolds, *Computer Graphics*, 29(2), *Proceedings of SIGGRAPH'95*, 1995.
- [49] Gueziec, A., Meshsweeper: Dynamic point-to-polygonal-mesh distance and applications, *IEEE Transactions on Visualization and Computer Graphics*, 7(1), pages 47–61, 2001.

- [50] Herberhold C., Faber S., Stammberger T., Steinlechner M., Putz R., Englmeier KH, Reiser M., and Eckstein F., In situ measurement of articular cartilage deformation in intact femoropatellar joints under static loading, *J Biomech.*, 32(12), pages 1287-95, 1999.
- [51] Iwasaki, N., Minami, A., Miyazawa, T. and Kaneda, K., 2000. Force distribution through the wrist joint in patients with different stages of Kienbock's disease: using computed tomography osteoabsorptiometry. *J Hand Surg [Am]* 25, 870-876.
- [52] Jacob HA, Kunz C, Sennwald G., Biomechanics of the carpus – functional anatomy and movement analysis of the carpal bones (article in German), *Orthopade*, Feb 21(1), pages 81–87, 1992.
- [53] Johnson, A.E., and Hebert, M., Using Spin Images for Efficient Object Recognition in Cluttered 3D Scenes, *IEEE PAMI* 21(5), pages 433–449, 1999.
- [54] Johnston, J. D., Small, C. F., Bouxsein, M. L. and Pichora, D. R., 2004. Mechanical properties of the scapholunate ligament correlate with bone mineral density measurements of the hand. *J Orthop Res* 22, pages 867–871.
- [55] Kimmel, R., Kiryati, N. and Bruckstein, A., Multi-valued distance maps for motion planning on surfaces with moving obstacles, *IEEE Trans. on Robotics & Automation Symposium*, 14, pages 427–436, 1998.
- [56] Kimmel, R. and Kiryati, N., Finding Shortest Paths on Surfaces by Fast Global Approximation and Precise Local Refinement, *Int. Journal of Pattern Recognition and Artificial Intelligence*, 10(6), pages 643–656, 1996.
- [57] Kiryati, N. and Szekely, G., Estimating Shortest Paths and Minimal Distance on Digitized Three Dimensional Surfaces, *Pattern Recognition*, 26, pages 1623-1637, 1993.
- [58] Kleinman, W., and Graham, T., The distal radioulnar joint capsule: clinical anatomy and role on posttraumatic limitation of forearm rotation, *The Journal of Hand Surgery*, 23A(4), pages 588–599, 1998.

- [59] Kotcheff, A.C.W., Taylor, C.J., Automatic Construction of Eigenshape Models by Direct Optimization, *Med. Image Analysis* 24, 303–314, 1998.
- [60] Kraevoy, V., and Sheffer, A., Cross-Parameterization and Compatible Remeshing of 3D Models, *ACM Trans. Graph.* 23(3), pages 861–869, 2004.
- [61] Kwak, S.D., Blankevoort, L., and Ateshian, G.A., A Mathematical Formulation for 3D Quasi-Static Multibody Models of Diarthrodial Joints, *Computer methods in biomechanics and biomedical engineering*, 3(1), pages 41-64, 2000.
- [62] Laidlaw, D.H., Fleischer, K.W., Barr, A.H., Partial-volume Bayesian Classification of Material mixtures in MR volume data using voxel histograms, *IEEE Trans. Med. Imag.*, 17(1), pages 74–86, 1998.
- [63] Laidlaw, D.H., Geometric Model Extraction from Magnetic Resonance Volume Data, Ph.D. thesis, Caltech, 1995.
- [64] Ledoux, P., Lamblin, D. and Targowski, R., Modifications to the mechanical behavior of the wrist after fracture of the scaphoid, Modeling by finite element analysis. *Acta Orthopaedica Belgica* 67, pages 236-241, 2001.
- [65] Lengyel, J., Reichert, M., Donald, B., and Greenberg, D., Real-time robot motion planning using rasterizing computer graphics hardware, *Proc. SIGGRAPH 1990*, pages 327–335, 1990.
- [66] Maintz, J.B.A., and Viergever, M.A., A survey of medical image registration, *J. Medical Image Analysis*, 2(1), pages 1–36, 1998.
- [67] Marai, G.E., Laidlaw, D.H., Demiralp, C., Andrews, S., Grimm, C.M., Crisco, J.J., Estimating joint contact areas and ligament lengths from bone kinematics and surfaces, *IEEE Trans. on Biomedical Engineering*, 51(5), pages 790–799, 2003.
- [68] Marai, G.E., Grimm, C.M., Laidlaw, D. H., Arthrodial Joint Markerless Cross-Parameterization and Biomechanical Visualization, *IEEE Transactions on Visualization and Computer Graphics* (in press).

- [69] 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), pages 177–187, 2006.
- [70] Marai, G.E., Laidlaw, D.H. and Crisco, J.J., A Kinematics-Based Method for Generating Cartilage Maps and Deformations in the Multi-Articulating Wrist Joint From CT Images, *IEEE 2006 Conference of the Engineering in Medicine and Biology Society (EMBC'06)*, pages 2079-2082, 2006.
- [71] Marai, G.E., Laidlaw, D.H., Markerless inter-subject bone shape matching using 2D projections, *MICCAI 2005 short papers*, 2005.
- [72] Martin, K., Schroeder, W., Lorensen, B., *The Visualization Toolkit: An Object-Oriented Approach to 3D Graphics*, Prentice Hall, 1997.
- [73] Mauch, S., A Fast Algorithm for Computing the Closest Point and Distance Transform, <http://www.acm.caltech.edu/seanm/software/cpt/cpt.pdf>.
- [74] Mayfield, J. K., Patterns of injury to carpal ligaments - a spectrum, *Clinical Orthopaedics* 187, pages 36–42, 1984.
- [75] Meakin, J. R., Shrive, N. G., Frank, C. B. and Hart, D. A., Finite element analysis of the meniscus: the influence of geometry and material properties on its behaviour, *Knee* Vol.10, pages 33–41, 2003.
- [76] Mitchell, T.M., *Machine learning*, McGraw-Hill, 1998.
- [77] Moore, D., Hogan, K., Crisco, J., Akelman, E., DaSilva, M., and Weiss, A., 3-D In Vivo Kinematics of the Distal Radioulnar Joint in Malunited Distal Radius Fractures, *Journal of Hand Surgery*, 27(2), pages 233–242, 2002.
- [78] Mow, V.C., Lai, W.M., and Holmes, M.H., Advanced theoretical and experimental techniques in cartilage research, In *Biomechanics: Principles and Applications*, ed. by Huiskes, R., van Campen, D.H., and de Wijn, J.R., pages 47–74, Martinus Nijhoff, 1982.

- [79] Neu, C.P., McGovern, R.D., Crisco, J.J., Kinematic Accuracy of Three Surface Registration Methods in a Three-Dimensional Wrist Bone Study, *J.Biomech.Eng.*, 122(10), pages 528–533, 2000.
- [80] Oda, M., Hashizume, H., Miyake, T., Inoue, H. and Nagayama, N., A stress distribution analysis of a ceramic lunate replacement for Kienbock's disease, *J Hand Surg [Br]* 25, pages 492–498, 2000.
- [81] Osher, S. and Sethian, J., Fronts propagating with curvature-dependent speed: algorithms based on Hamilton-Jacobi formulation, *J. Computational Physics*, 79, pages 12–49, 1988.
- [82] Panjabi, M.M., Krag, M.H., Goel, V.K., A Technique for Measurement and Description of Three-Dimensional Six Degree-of-Freedom Motion of a Body Joint With an Application to the Human Spine, *J.Biomech.Eng.*, 14, pages 447–460, 1981.
- [83] Papadimitriou, C.H., An Algorithm for Shortest-Path Motion in Three Dimensions, *Information Processing Letters*, 20, pages 259–263, 1985.
- [84] Pascarelli, E. and Quilter, D., *Repetitive Strain Injury: A Computer User's Guide*, Wiley & Sons, 1994.
- [85] Patterson, R.M. and Viegas, S.F., Biomechanics of the Wrist, *J. Hand Therapy*, 8(2), pages 97–105, 1995.
- [86] Pauly, M., Pai, D.K., and Guibas, L., Quasi-Rigid Objects in Contact, *Proceedings of ACM Symposium on Computer Animation*, pages 109–119, 2004.
- [87] Payne, B. and Toga, A., Distance field manipulation of surface models, *IEEE Computer Graphics and Applications*, pages 65–71, 1992.
- [88] Praun, E., Sweldens, W. and Schröder, P., Consistent Mesh Parameterizations, *Proceedings of ACM SIGGRAPH 2001*, pages 179–184, 2001.
- [89] Richardson, M.L., *Radiographic Anatomy of the Skeleton*, <http://www.rad.washington.edu/RadAnat/WristPALabelled.html>.

- [90] Richmond, B.G., Strait, D.S., Evidence that humans evolved from a knuckle-walking ancestor, *Nature* 404, pages 382–385, 2000.
- [91] Scharir, M. and Schorr, A. On shortest paths in polyhedral spaces, *SIAM J. Comput.*, 1(15), pages 193–215, 1986.
- [92] Scherrer, P.K., and Hillberry, B.M., Piece-wise mathematical representations of articular surfaces, *Journal of Biomechanics*, 12, pages 301–311, 1979.
- [93] Schreiner, J., Prakash, A., Praun, E. and Hoppe, H., Inter-Surface Mapping, *ACM Trans. Graph.* 23(3), pages 870–877, 2004.
- [94] Schuind, F., An, K., Berglund, L., Rey, R., Cooney, W., Linscheid, R., and Chao, E., The distal radioulnar ligaments: a biomechanical study, *The Journal of Hand Surgery*, 16A(6), pages 1106–1114, 1991.
- [95] Schuind, F., Cooney, W. P., Linscheid, R. L., An, K. N. and Chao, E. Y., Force and pressure transmission through the normal wrist: A theoretical two-dimensional study in the posteroanterior plane, *J Biomech* 28, 587–601, 1995
- [96] Sethian, J., *Level Set Methods: Evolving Interfaces in Geometry, Fluid Mechanics, Computer Vision, and Material Science*, Cambridge University Press, 1996.
- [97] Sethian, J., A Fast Marching Level Set Method for Monotonically Advancing Fronts, *Proc. Nat. Acad. Sci.*, volume 94, pages 1591–1595, 1996.
- [98] Sirkett, D.M., Mullineux, G., Giddins, G.E.B., Miles, A.W. , A Kinematic Model of the Wrist based upon Maximization of Joint Contact Area, *Proceedings of the Institution of Mechanical Engineers Part H: The Journal of Engineering in Medicine* 218, pages 349-359, 2004.
- [99] 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), pages 2037–2047, 2000.

- [100] Stoller, D.W., Colditz, J.C., McGrouther, D. a., *The Interactive Hand 2000: Radiology Edition* (CD-ROM for Windows and MacIntosh), Primal Pictures Ltd., 2000.
- [101] 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)*, pages 63–75, 2003.
- [102] 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), pages 107–122, 2003.
- [103] Swanstrom, M.D., Zarucco, L., Hubbard, M., Stover, S.M., Hawkins, D.A., Musculoskeletal Modeling and Dynamic Simulation of the Thoroughbred Equine Forelimb During Stance Phase of the Gallop, *J Biomech Eng.*, 127(2), pages 318–28, 2005.
- [104] Swartz, S.M., Allometric patterning in the limb skeleton of bats: implications for the mechanics and energetics of powered flight, *J. Morph.*, 234, pages 277–294, 1997.
- [105] Szekely, G., Brechbuehler, C., Hutter, R. and Rhomberg, A., Modelling of soft tissue deformation for laparoscopic surgery simulation, *Medical Image Analysis*, 4, pages 57–66, 2000.
- [106] Thirion, J.P., New Feature Points based on Geometric Invariants for 3D Image Registration, INRIA TR 1901, 1993.
- [107] Thoomukuntla, B.R., Pillai, R.R., McIff, T.E., Bilgen, M., Ateshian, G.A., and Fischer, K.J., Validation of an MRI-Based Method for In Vivo Joint Contact Mechanics Analyses, American Society of Mechanical Engineers Summer Bioengineering Conference, Vail CO, 2005.
- [108] Un, K. and Spilker, R.L., Finite element simulation of biphasic soft tissue contact with application to the shoulder joint, *Proceedings of the 23rd Annual International Conference of the IEEE*, Vol 2, pages 1504 – 1507, 2001.
- [109] Upal, A., Accuracy of the closest point transform on a 2 sphere phantom CT dataset, Technical Report, Brown University, 2004.

- [110] Van Herk, M. and Kooy, H.M., Automatic three-dimensional correlation of CT-CT, CT-MRI, and CT-SPECT using chamfer matching, *Med. Phys.*, 21, pages 1163–1178, 1994.
- [111] Van Sint Jan, S.L., Clapworthy, G.J., and Rooze, M., Visualization of combined motions in human joints, *IEEE Computer Graphics and Applications*, November/December, pages 10–14, 1998.
- [112] 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, pages 644–651, 2000.
- [113] Warner, M. D., Taylor, W. R. and Clift, S. E., Finite element biphasic indentation of cartilage: a comparison of experimental indenter and physiological contact geometries. *Proc Inst Mech Eng [H]* 215, pages 487–496, 2001.
- [114] West, J., Fitzpatrick, J.M., Wang, M.Y., Maurer, C.R., Kessler, R.M., Maciunas, R.J., Barillot, C., Lemoine, D., Collignon, A., Maes, F., Suetens, P., Vandermeulen, D., van den Elsen, P.A., Napel, S., Sumanaweera, T.S., Harkness, B., Hemler, P.F., Hill, D.L.G., Hawkes, D.J., Studholme, C., Maintz, J.B.A., Viergever, M.A., Malandain, G., Pennec, X., Noz, M.E., Maguire, C.Q., Pollack, M., Pellizari, C.A., Robb, A., Hanson, D., Woods, R.P., Comparison and evaluation of retrospective intermodality image registration methods, *J. Comput. Assist. Tomogr.*, 21, pages 554-566, 1997.
- [115] Whitaker, R., Breen, D, Museth, K., Soni, N., Segmentation of biological volume datasets using a level set framework, *Volume Graphics 2001*, pages 249–263, Springer, Vienna, 2001.
- [116] Winslow, R. L. and J. L. Greenstein, The Ongoing Journey to Understand Heart Function Through Integrative Modeling, *Circ. Res.* 95(12), pages 1135-1136, 2004.
- [117] Yap, C.K., Algorithmic motion planning, *Advance in Robotics*, Volume I, Lawrence Erlbaum Associates, 1985.

- [118] Zhu, Y., 3D Reconstruction, Motion Simulation and Biomechanical Visualization of the Knee Joint, Ph.D. Dissertation, George Mason University, 2001.