

COMBINING DIFFEOMORPHIC MATCHING WITH
IMAGE SEQUENCE INTENSITY REGISTRATION

A Proposal
Presented to
the Faculty of the Department of Mathematics
University of Houston

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy

By
Jeff Freeman
May 2014

COMBINING DIFFEOMORPHIC MATCHING WITH
IMAGE SEQUENCE INTENSITY REGISTRATION

Jeff Freeman

APPROVED:

Robert Azencott, Chairman
Dept. of Mathematics

Ronald Hoppe
Dept. of Mathematics

Jiwen He
Dept. of Mathematics

Stephen H. Little, M.D.
Methodist Hospital

Dean, College of Natural Sciences and Mathematics

COMBINING DFFEOMORPHIC MATCHING WITH IMAGE SEQUENCE INTENSITY REGISTRATION

An Abstract of a Proposal
Presented to
the Faculty of the Department of Mathematics
University of Houston

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy

By
Jeff Freeman
May 2014

Abstract

This thesis presents the research work completed over the past 4 years in the context of a collaborative project between The Methodist Hospital (TMH) and the University of Houston. We have developed and implemented novel algorithmic approaches to: develop patient-specific static mitral valve models from tagged 3D-Echocardiographic image data; incorporate this 3D-Echo data into a new methodology for diffeomorphic valve tracking; and investigate the strain distributions on valve leaflets derived from the deformations captured through valve tracking.

First, we have applied spline techniques in order to generate static models of the mitral valves at discrete instants. Classical smoothing splines are applied to the modeling of valve boundary curves, while tensor product smoothing splines are used to fit surfaces to the mitral valve interior leaflets. Two approaches are presented for this surface modeling: one (lofting) works in all cases but requires more effort to execute, while the other (principal plane) is simpler in its approach but does not work for all mitral valves. These techniques are illustrated by the display of multiple mitral valve models.

Next, we have considered optimal diffeomorphic matching of these mitral valve models by a variational approach based on Hilbert spaces of time dependent vector fields. Since models matched by diffeomorphisms are extracted from sequences of 3D-Echo images, we have proposed an equivalent formulation and solution of this problem, one that involves an iterative scheme that alternates between pure geometric diffeomorphic matching and image intensity registration. Several detailed, concrete examples are presented to validate the performance of our approach.

Finally, we have developed the methods needed to compute, compare, and quantify the distribution of strain values on mitral valve leaflets. We adopted standard finite difference techniques for computing the strain tensor on a leaflet surface and applied the Kolmogorov-Smirnov tests to evaluate how 3 groups of mitral valves (normal, pre-surgery, post-surgery) compare in terms of strain values distributions. These results have been published in a joint paper between UH and TMH. A second joint paper will be submitted in May 2014.

Contents

1	Introduction	1
2	Spline Modeling of Smooth Shapes in \mathbb{R}^3	6
2.1	Splines, B -splines and Smoothing Splines	7
2.2	Fitting Curves and Surfaces in \mathbb{R}^3	10
3	Static Modeling of the Mitral Valve	13
3.1	Mitral Valve Anatomy	14
3.2	Mitral Valve Imaging	14
3.3	Mitral Valve Modeling	17
3.3.1	Mitral Annulus and Coaptation Models	18
3.3.2	Principal Plane Leaflet Models	19
3.3.3	Lofting Leaflet Models	23
3.3.4	Computing Times and Quality of Fit	28
4	Diffeomorphic Shape Matching	31
4.1	Shape Matching Survey	31
4.1.1	Shape Representation in \mathbb{R}^3	31
4.1.2	Shape Deformations	33
4.2	Diffeomorphic Matching of Two Shapes	35
4.2.1	Objective Functional Formulation	36

4.2.2	Diffeomorphic Point Matching of Two Shapes	40
4.2.3	Numerical Algorithm for Diffeomorphic Matching of Two Shapes	43
4.3	Numerical Algorithm for Diffeomorphic Matching with Multiple Target Shapes	48
5	Examples of Diffeomorphic Point Matching	52
5.1	Diffeomorphic Point Matching Implementation	52
5.2	Synthetic Construction of \mathbb{R}^3 vector field flow	55
5.3	Mitral Valve Leaflet (One Target)	58
5.4	Mitral Valve Leaflet (Two Targets)	59
6	Image Registration	66
6.1	Overview	66
6.2	Elements of Image Registration	67
6.2.1	Feature space	67
6.2.2	Deformation Model	68
6.2.3	Similarity metrics	69
6.3	Image Registration for 3D Echocardiography	70
7	Combining Diffeomorphic Shape Matching with Image Registration	73
7.1	Geometric Matching: The need for voxel intensity data	74
7.2	Geometric Matching + Intensity Registration	75
7.2.1	Background	75
7.2.2	Alternative Formulation	78
8	Numerical Methods for Intensity Registration	80
8.1	Intensity Registration Loop	80
8.1.1	Geometric Matching	81
8.1.2	Trajectory Perturbation	82

8.1.3	Velocities, Kernel Inversion, and Diffeomorphism Recovery . . .	85
8.1.4	Performance Monitoring	86
8.2	Examples	89
8.2.1	Example 5.1 – Direct Construction	89
8.2.2	Example 5.2 – Mitral Valve Leaflet	94
9	Exploration of Geometric Matching + IRL	96
9.1	Overview	96
9.2	Performance Impact Analysis for the main parameters of the IRL algorithm	97
9.2.1	Impact of the IRL parameter ϵ	97
9.2.2	Impact of Number of IRL Iterations	98
9.3	Conclusion	101
10	Mitral Valve Strain Analysis	105
10.1	Motivation	105
10.2	Stress/Strain Overview	106
10.3	Strain Tensor and Mitral Leaflets Strain	107
10.4	Statistical Analysis via Kolmogorov-Smirnov	110
10.5	Global Strain Intensity Comparison	112
10.5.1	Normal/Organic Pre-Surgery Comparison	113
10.5.2	Organic Pre-Surgery/Post-Surgery Comparison	113
10.5.3	Normal/Organic Post-Operation Comparison	115
11	Mitral Valve Strain–High Strain Localization Analysis	125
11.1	Motivation	125
11.2	Computerized Generation of Leaflet Scallops	126
11.3	Other Leaflet Zones	126

11.4 High Strain Concentration	127
11.5 Regional Strain Analysis	128
11.6 Anatomical Interpretation and Analysis of Strain Localization Results	129
12 Outline of Future Work: The Shape Space of Human Mitral Valves	133
12.1 Motivation	133
12.2 Geometric Traits as Markers of a Population	134
12.3 Energy of Diffeomorphism as a Metric	135
12.4 Distance between Mitral Valve Models	136
12.5 Example for Normal Population	138
12.6 Further Work	140
Bibliography	146

List of Figures

1.1	Workflow for thesis	3
2.1	Cubic basis (B)-spline $B(x)$ in \mathbb{R} (central “hat” in black) along with the 4 cubic polynomials that piecewise make up the curve (red, green, magenta, and black). The knots are denoted by the blue vertical lines.	8
2.2	Smoothing spline sensitivity to parameter p . Depending on the data (black), small variations in p can have a large impact on the quality of fit for the smoothing spline. Notice the difference between $p = 0.99$ and $p = 1$	9
2.3	Smoothing splines of \mathbb{R}^3 scattered data (black). MATLAB Curve Fitting Toolbox <i>csaps</i> makes these constructions simple and easy to visualize.	11
2.4	Tensor product spline of \mathbb{R}^3 scattered data (black). The smoothing parameter $p = 0.998$ was chosen to display a smooth surface without overfitting. MATLAB Curve Fitting Toolbox <i>csaps</i> allows for quick construction and evaluation of tensor product splines.	12
3.1	Figure from www.heart-valve-surgery.com	15
3.2	Schematic of tagging planes P_i , rotation axis A and anterior horn O_{horn} . This simple example shows three color-coded tagging planes and their corresponding tags.	16
3.3	(left) Interactive tagging in the SLICER 3D environment. (right) A sample tagging plane P_i	17
3.4	Sample models for patient’s mitral annulus (yellow) and coaptation line (blue). The tagged data sites are the black circles.	19

3.5	Patient-specific MV model (a regurgitation case)	20
3.6	Sample principal plane leaflet model. The tagged data sites are the black circles.	21
3.7	Limitations of principal plane modeling. The top figure gives a model where the principal plane (black) allows for a cartesian expression $z = f(x, y)$. The bottom figure gives a model (lofting method, see next section) where no cartesian expression $z = f(x, y)$ is possible. Three candidate plane cross-sections (red) are displayed; each results in double values z_1, z_2 for certain (x, y) pairs.	22
3.8	Schematic illustrating the radial curves C_1 (green), C_2 (blue), and C_3 (red), as well as the angular curves R_1, R_2, R_3, R_4 (black).	24
3.9	Sample lofting leaflet model. The tagged data sites are the black circles.	25
3.10	Human-in-the-loop tools for radial fit. (Top) Interactive spline knot adjustment. Annular tags are yellow circles, anterior leaflet tags are green circles, and posterior leaflet tags are blue circles. Whenever loops or irregularities are detected, knots are re-ordered and re-checked. (Bottom) Once all tagged planes have been processed, neighboring tagged planes are visualized concurrently. The color scheme: previous AL–red, previous PL–magenta; current AL–green, current PL–blue; next AL–cyan, next PL–black. These plots allow for preliminary “continuity” checks before the angular fit is attempted.	26
3.11	Ghost data (yellow) over-layed on sparse PL data. The previous PL plane C_{i-1} (magenta) and next PL plane C_{i+1} (black) do not match the current PL plane C_i (blue). By sampling these curves to M equal values and interpolating corresponding points $x_j \in C_{i-1}, y_j \in C_i, z_j \in C_{i+1}$, a smoother transition is achieved. The current AL plane (green) is shown. It is expected that this valve be closed at systole; this validates the position of the derived ghost data (yellow).	28
3.12	Crease smoothing. (Left) Radial fits in tagging planes converge at anterior horn O_{horn} . Minor changes in azimuthal angle ϕ greatly affect the smoothness of the angular fit in the polar angle θ . This is apparent in the ridges and creases that form near the anterior horn. (Right) Least-squares smoothing of the azimuthal angle ϕ across the range of polar angles θ reduces these ridges. Provided smoothing is restricted to small neighborhood of O_{horn} , the resulting model does not deviate from tagged data.	29

3.13	Patient-specific MVA models. The top row are two models for normal, healthy patients, while the bottom row are models for diseased patients exhibiting larger valves and leaflet regurgitation.	30
4.1	Point grid representation of various smooth shapes in \mathbb{R}^3	33
4.2	The reference F4 Phantom can be identified via overhead satellite imagery through the use of global rigid transformations. Top image: www.eliteday.com . Bottom image: ngm.nationalgeographic.com	50
4.3	A human mitral valve at three different instants. The static models were created using the techniques of Chapter 3.	51
4.4	The deformation of the rectangle (black) is modeled by a flow of time-dependent \mathbb{R}^2 diffeomorphisms F_t , determined by a vector field flow v_t on \mathbb{R}^2 . The vector field flow is shown at three different times (blue, green, red).	51
5.1	The reference shape S_0 (black) and the kernel nodal points N_K (green) shown from two different perspectives.	61
5.2	The reference shape S_0 (left) and target shape S_1 (right).	62
5.3	The bottom unit sphere is evolved under the action of a vector field flow. The reference shape is discretized to a point grid (black) and deformed to match the target shape. The deformed point grid recovered from diffeomorphic matching (red) coincides well with the target surface (blue).	62
5.4	Comparison of recovered deformation trajectories (green) with actual trajectories (black). Only two trajectories shown for ease of viewing. The recovered trajectories, although matching well the beginning and ending points of the actual trajectories, can deviate in different manners at intermediate times. The distances between the actual and recovered trajectories never exceed 0.05 units.	63
5.5	The top surface (black points/blue mesh) is deformed into the bottom surface (blue points/red mesh).	64
5.6	The reference leaflet at begin-systole (red) is deformed throughout systole. At times mid-systole (blue) and end-systole (orange), the deformed trajectories (black) match quite well (< 1 mm disparity).	65

6.1	Side-by-side prestress and poststress long-axis views before (left column) and after (right column) registration. White arrows point to the free wall for the left-ventricle. The marked wall segments in the poststress views are tilted and shifted with respect to the same segments in the prestress views before registration. Note that these segments specifically and the entire left-ventricle in general are better aligned following registration. (Caption and figure from [43])	70
6.2	Illustration of one iteration of the SURE algorithm registering the light gray circle (homologous volume) onto the darker shape (reference). The images (a)-(d) correspond to the four steps described above in the text. (Caption and figure from [27]).	71
7.1	Voxel intensities along 5 specific voxel deformation trajectories for a geometric matching case. Notice the large ranges of intensities, as well as the large oscillations present on some of the trajectories. This is not that surprising, as geometric matching, being a feature-based method of registration, does not consider the voxel intensities information present in the ultrasound movie when the diffeomorphism between two 3D-shapes is computed.	75
8.1	Perturbed trajectories (green) of geometric trajectories (black). The boxes B_ϵ (red) are displayed on a single trajectory to illustrate the procedure.	83
8.2	Sample displays used in performance monitoring of IRL. (Left) A plot of IRQ vs. KEQ for 9 voxel trajectories. (Right) Quantile curve of KEQ for all computed trajectories in IRL. Initial geometric matching quantities are in black, while the cyan values are after 1 iteration of the IRL.	88
8.3	3D time slice of case a). White on black image sequence.	90
8.4	Quantile curves of intensity oscillations (IRQ) and voxel trajectory kinetic energy (KEQ) for case a). The black curve represents the values from geometric matching that are input into the IRL. The blue, cyan, and yellow are successive iterations of the IRL algorithm. Notice the decrease in intensity oscillation and increase in kinetic energy with each iteration. The grey curve is the result of resetting high energy trajectories to more stable, earlier iteration trajectories.	92
8.5	3D time slice of case b). Noisy image sequence.	93

8.6	Performance criteria tracked during successive IRL iterations applied to the dynamic modeling of an anterior leaflet. The black curve corresponds to the values from geometric matching (chapter 5) that are input to the IRL, while the colored trajectories correspond to successive IRL iterations. The grey curve gives the values after trajectories are reset from high-energy, unstable values. The planar display of IRQ vs. KEQ histories demonstrates that approximately 5% of the trajectories began varying wildly in KEQ while demonstrating marginal IRQ gains, seen by the nearly horizontal slopes in the (KEQ,IRQ) plane. However, a large majority of the paths in the (KEQ,IRQ) plane remain bounded in KEQ during the iteration, seen by the large multi-colored “fuzzball” at the lower left of the cloud plot and confirmed by the KEQ quantile plot which shows that about 80% of the trajectories remain bounded above by 5000 KEQ after 5 iterations while beginning the IRL at a 80%-quantile of 3000 KEQ.	95
9.1	The IRQ and KEQ quantile curves monitoring of quality criteria are displayed here for three cases of mitral leaflets dynamic modeling: (top) Normal 42 AL, (middle) Organic 28 PL, (bottom) PostOp 48AL. Notice the marginal gains in IRQ reduction and KEQ increase for higher IRL iterations. This is an indicator of a good solution. For these cases, a small number of IRL iterations (2-4) yields the majority of the benefits with minimal increases in kinetic energy cost.	104
10.1	Two sample strain maps corresponding to a normal patient (top) and a diseased, pre-surgery patient (bottom). The high strain intensities are colored in orange, while low strain intensities are colored blue. . .	116
10.2	Quantiles curves N_{AL} and N_{PL} for the average normal strain intensity (shown above in black) demonstrate strain severity for pre-surgery MR patients; we in fact observe that strain is stochastically higher for pre-surgery cases than for average normal cases. This is quite clear graphically since the 10 pre-surgery strain quantiles curves displayed in green (AL) and blue (PL) are essentially above the black normal strain quantile curves.	117
10.3	Histograms of strain intensities across both leaflets for all 10 normal patients (black) and 10 pre-surgery organic MR patients (red). Notice the clear shift in the red organic data towards higher values of leaflet strain intensity (LSI) compared to the black normal patient data. . .	118

10.4	Leaflets Strain Intensities: color-coded maps at MS for two regurgitation cases. View is top-down, with AL towards the left. All axes have units of mm. Coaptation line is shown in dark green. Notice the reduction in strain after NRD surgery, evident in the increase in blue on the strain maps.	119
10.5	Quantiles curves for two regurgitation patients emphasize a strain reduction due to the NRD surgical repair. The two figures on the left represent the AL strain quantiles, while those in the right represent the PL strain quantiles.	120
10.6	Quantile curves displaying the reduction of strain intensities due to NRD surgery. (Top) For the AL, our figure shows that, after multiplying the pre-surgery strain quantile curve by a shrinking factor of 0.77, the rescaled curve nearly coincides with the post-surgery strain quantile curve. (Bottom) For the PL, our right picture displays a similar result, with a multiplicative strain shrinking factor of 0.60.	121
10.7	Histograms of strain intensities across both leaflets for all 10 organic MR patients, both pre-surgery (red) and post-surgery (green). Notice the clear shift in the green post-op data towards lower values of leaflet strain intensity (LSI) compared to the red pre-surgery patient data. .	122
10.8	For the normal and post-surgery populations, strain quantile curves are nearly identical (as validated by KS-tests with p -values 0.01). . .	123
10.9	Histograms of strain intensities across both leaflets for all 10 normal patients (black) and 10 post-surgery organic MR patients (green). Notice the data for the two populations appear identical, illustrating the reduction of leaflet strain intensity (LSI) back to “normal” levels following the NRD surgery.	124
11.1	Computerized scallop generation process.	131
11.2	Computed geometric zones on one typical posterior leaflet PL. The left figure displays the 3 Posterior Leaflet scallops, while the right figure displays our 5 other geometric PL zones: Z_{coapt} (magenta), Z_{center} (cyan), Z_{ann} (red), and the two commissures Z_{comm} (blue).	132
12.1	Two examples of normal patient mitral valve models.	141
12.2	Two examples of organic pre-surgery mitral valve models	142

12.3	Two examples of post-surgery mitral valve models	143
12.4	Distance between two diffeomorphisms D_1, D_2 given by the average of the kinetic energies $E(F_i)$ of the intermediary diffeomorphisms F_i . . .	144
12.5	The 10 normal anterior leaflets are displayed in 3D Euclidean space via principal component analysis. Neighboring leaflets are connected by colored lines. For instance, the closest leaflet to patient 8 (yellow) is patient 2, and this is marked by yellow connecting line between patient 2 and patient 8.	145

List of Tables

5.1	The runtime for <i>DiffeomorphicPointMatcher.run()</i> with 10 iterations for different values of N_0 . Notice that the computing time is definitely non-linear with N_0 . This is not too surprising given the kernel operations involve $3N_0 \times 3N_0$ matrices.	60
9.1	Results for $\epsilon = 0.2$. For a fixed number of 4 IRL iterations a 20% increase in the kinetic energy of voxel trajectories generates trajectories with at least 30% less intensity fluctuations	99
9.2	Results for $\epsilon = 0.8$. For a fixed number of 4 IRL iterations a 70% increase in the kinetic energy of voxel trajectories is necessary to generate trajectories with at least 40% less intensity fluctuations.	100
9.3	IRL Iteration Trade Results. The successive row difference method given in the text can be used to find the “sweet spot” where we get maximum reductions in IRQ for minimal increases in KEQ. For Normal 42 and Organic 76, we see this occurs at iteration 2. For PostOp48, we see this occurs at iteration 2 or 3.	103
10.1	Group averages of mean strain intensities. Strain intensities are higher for the organic group than the normal group (KS test p value = 0.05) for both leaflets.	112
10.2	Multiplicative strain shrinking factors due to NRD repair surgery for the posterior leaflets and the anterior leaflets, after averaging over all regurgitation cases.	114
11.1	Comparison of high strain concentrations $HSC(Z)$ fo 8 leaflets zones Z	129

12.1 Distances $d(D_1, D_2)$ between normal patient anterior leaflets. Notice patient 8 has the lowest computed mean distance to all other patients. Contrast this to patient 6, which exhibits the largest mean distance to other patients. 138

Chapter 1

Introduction

This thesis presents the research work that we have completed over the past 4 years, in the context of a joint 2008-2014 study between a University of Houston team of mathematicians and computer scientists led by Robert Azencott, and a team of cardiologists at The Methodist Hospital (TMH) led by William Zoghbi, MD. Cardiologists at TMH and elsewhere are interested in detailed evaluation and modelization of human mitral valves, periodic deformations during standard heart cycles, and non-invasive quantification of the effectiveness of mitral valve surgical repair techniques.

One particular mitral valve surgery which has shown promising results over the last 20 years is called Non-Resectional Dynamic repair (NRD), and involves the insertion of an artificial ring along the mitral annulus. NRD surgery has been shown to help restore normal function for mitral valves exhibiting leaflets prolapse and/or mitral regurgitation (see [31], [32] for details).

3D-Echocardiographic image sequences (3D-Echo) are non-invasive and are part of most standard medical examination protocols for heart patients. These 3D-images provide clinicians with direct visual evidence of mitral valve dysfunctions, such as organic regurgitation or mitral leaflets prolapse. Computer analysis of ultrasound images of the mitral valve has already been used to study the areas of regurgitation in pre-surgery patients (see [1], [9], [40]). Following valve repair surgery, patients often undergo post-operative 3D-Echo to help surgeons to evaluate visually the functional impact of surgery. More generally, 3D-Echo image data have been studied through various registration techniques in order to reconstruct and/or modelize cardiac deformations (see [14],[22], [43]).

But to study mitral valve dynamics in the context of elasticity theory, finite elements models have focused on the numerical evaluation of mitral leaflets stress and strain throughout the cardiac cycle. This finite elements modeling (FEM) approach requires complicated elasticity assumptions on leaflet tissue and many patient-specific parameters that are generally impossible to acquire on live human patients. Oftentimes, these FEM approaches require animal-based data for proper evaluation. Specifically, in order to accurately describe how tissue responds to a force (tension, compression), the elastic properties of the tissue must be known (see Chapter 9 for an overview of stress and strain). Since the acquisition of these elastic properties would require active measurements of live patients (an obvious major challenge), most published results of this type replace unknown parameters in their elasticity models by the mitral leaflets tissue properties quantified through in vivo experiments with ovine or porcine animal models (see [29], [30]).

In our study, which relies only on the non-invasive 3D-Echo ultrasound data analysis of human mitral valves, we have deliberately avoided the need to compute or simulate sophisticated elasticity models of the observed mitral valves. We have focused instead on patient-specific computations of mitral leaflet strain distributions, comparisons of strain distributions across patients, as well as strain comparisons before and after mitral valve repair surgery.

The workflow of this project spans three major topics. Figure [1.1] below summarizes the workflow and which chapters correspond to those steps in the thesis.

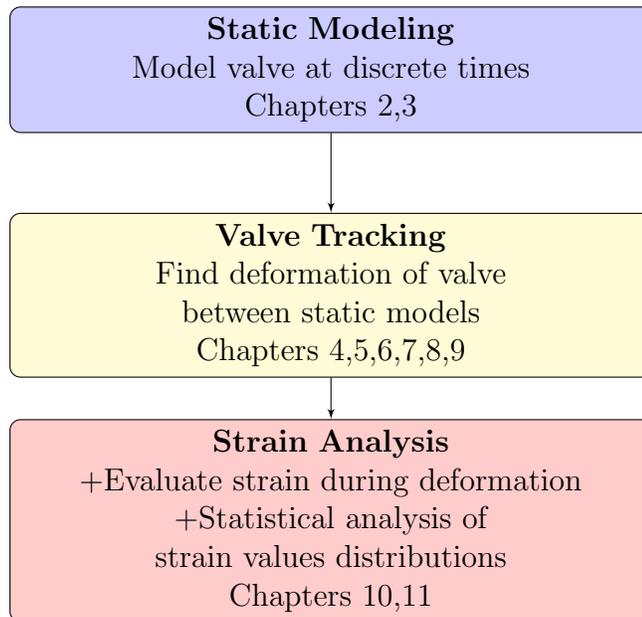


Figure 1.1: Workflow for thesis

Chapter 2 is a brief overview of spline methods used to model static shapes in R^3 . Chapter 3 is the direct application of these methods to compute patient-specific static mitral valve models from tagged 3D-Echo image data. Two approaches are discussed and the applicability of each approach is explored. Chapter 4 begins the

study of mitral valve tracking by first presenting a calculus of variations approach for geometric shape matching. This variational approach is briefly explored and then a numerical algorithm for diffeomorphic matching is presented in detail. Chapter 5 explores these ideas through some synthetic examples. Chapter 6 then introduces some image registration concepts based on image intensities analysis and highlights a few known techniques that have already been used to study 3D-Echo ultrasound image data. Chapters 7 and 8 present the new approach developed in our study to incorporate the 3D-Echo data used in our mitral valves static modeling into an efficient technique for valve tracking. Chapter 9 gives a detailed study of the adequate parameterizations of our diffeomorphic valve tracking algorithm and highlights the performances of this algorithm. With the technique for valve tracking now finely-tuned to 3D-Echo data, Chapter 10 attacks one of the main goals of our research project: strain distribution analysis. The classical notion of strain is presented and we discuss the methods implemented to derive the distributions of leaflet strain values from the dynamic leaflet deformations captured by our valve tracking. After introducing the necessary statistical concepts and tests, strain distributions are computed for each one of the 60 leaflets observed in our 3D-ultrasound data set. Strain distributions are then compared across patient populations, both before and after surgery. Chapter 11 then explores regional strain values for 8 natural leaflet regions; specifically, we study which anatomical regions of the mitral valve exhibit high strain concentrations. Finally, Chapter 12 outlines future research work to analyze shape spaces of mitral valves.

Multiple researchers have contributed to this long-term team project; my own

research is a culmination of the work laid before me. An earlier in-depth mathematical and numerical analysis of diffeomorphic shape matching (see [3]) involved professors Robert Azencott, Roland Glowinski, Jiwen He , and Ronald Hoppe, as well as post-doctoral students Saurabh Jain, Aarti Jajoo, and Andrey Martynenko. In particular, previous work on mitral valves by A. Jajoo and A. Martynenko paved the way for the strain study presented here. Yue Qin and her two advisors, J. He and R. Azencott, explored 2nd-order convergence methods to accelerate 1st-order variational approaches. Essential software tools for the mitral valve tagging interface and static modeling were developed by the post-doctoral fellows Simon Alexander and Saurabh Jain. On the cardiology side, the UH team interacted on a regular basis with Dr. William Zoghbi and Dr. Stephen Little, who provided crucial guidance about the mitral valve anatomy, the problems raised by mitral valve regurgitation and prolapse, and the impacts of mitral valve repair surgery. The clinician Dr. Sagit Ben Zekry worked tirelessly on interpreting the 3D echocardiographic data and provided the team with expertly tagged mitral valve data, a fundamental key to compute accurate patient-specific mitral valve models. Finally, Dr. Gerald Lawrie, a world-renowned surgeon specializing in mitral valve repair, provided invaluable insight into the surgical NRD procedure.

The MATLAB codes that were written to support this thesis can be found on my website:

<http://math.uh.edu/~jwf1>

Chapter 2

Spline Modeling of Smooth Shapes in \mathbb{R}^3

Splines are an important tool in interpolation. Though they have been understood since the late 1940s by Schoenberg (see [41]), their utility as computational tool was not fully realized until the 1970s when stable numerical methods were introduced to evaluate what have become known as B -splines. We begin with a brief overview of splines and B -spline representations. Then we introduce smoothing splines and show how these concepts generalize to arbitrary curves in \mathbb{R}^3 . Then tensor product splines are introduced, which allow for spline representations for arbitrary gridded surfaces in \mathbb{R}^3 .

2.1 Splines, B -splines and Smoothing Splines

Consider a mesh of *knots* $\Delta = \{x_0, x_1, \dots, x_m\} \subset \mathbb{R}$ with $x_0 < x_1 < \dots < x_m$. For some $k \in \mathbb{N}$, the function $s : [x_0, x_m] \rightarrow \mathbb{R}$ such that

- (a) The restriction $s|_{[x_i, x_{i+1}]}$ is a polynomial of degree $d < k$.
- (b) $s \in C^{k-2}[x_0, x_m]$ if $k \geq 2$; otherwise s is piecewise constant

is called a *spline function of order k* for the mesh Δ . Denote by $S_{\Delta, k}$ the set of all spline functions of order k on the mesh Δ . An immediate fact is that $S_{\Delta, k}$ is a vector space of dimension $k + m - 1$. See [10] for details. The cubic B -spline basis introduced by Schoenberg ([41]) allows for compact representation of $s(x)$ in terms of integer shifts of the central “bell” $B(x)$ with

$$B(x) = \begin{cases} 2/3 - (1 - |x|/2)x^2 & \text{if } 0 \leq |x| < 1 \\ (2 - |x|)^3/6 & \text{if } 1 < |x| < 2 \\ 0 & \text{if } |x| \geq 2 \end{cases}$$

Specifically in the case of uniform knot spacing $h = x_{i+1} - x_i$, we have $s(x) = \sum_{k \in K} c(k)B(x/h - k)$ where the indices of the sum $K \subset \mathbb{Z}$ are chosen so that the domain of knots $[x_0, x_m]$ is supported by the shifted B -splines. See Figure [2.1] for a visualization of the prototypical cubic B -spline $B(x)$. The B -spline basis is most useful in the construction of spline functions, while the piecewise polynomial (pp) form implied in (a),(b) of the definition of the spline function is most useful for evaluation and conceptual purposes.

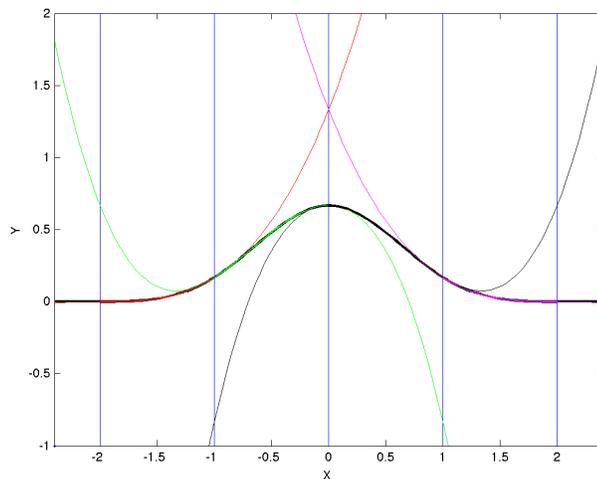


Figure 2.1: Cubic basis (B)-spline $B(x)$ in \mathbb{R} (central “hat” in black) along with the 4 cubic polynomials that piecewise make up the curve (red, green, magenta, and black). The knots are denoted by the blue vertical lines.

Spline functions are most often used in interpolation, and from this point forward it is assumed that all mentioned splines are to be fit to some given data. Denote by y_i a set of arbitrary given values indexed by the knots x_i of the spline grid. We seek a spline function s such that $s(x_i) = y_i$, which will thus smoothly interpolate the data (x_i, y_i) . Moreover, the spline may be chosen so that both the first and second derivatives $s'(x), s''(x)$ are continuous. Doing so requires the solution of an underdetermined linear system; additional boundary constraints may be added to uniquely determine the interpolating cubic spline.

But for interpolants on large datasets, one must be careful to avoid overfitting. There are two general avenues to this end, both of which are described in de Boor’s text (see [12]). The first is to require that the spline interpolant have much fewer knots than available data. This leads to an overdetermined linear system that can be

solved via least squares. The second involves *smoothing splines*. Namely, we consider the minimization problem

$$\min_{s \in \mathcal{S}_{\Delta, k}} p \sum_{i=0}^m [y_i - s(x_i)]^2 + (1 - p) \int_{x_0}^{x_m} [s''(t)]^2 dt$$

The smoothing parameter p offers a balance between data fit and smoothness. For $p = 1$, the solution is the natural cubic spline interpolant of the data (“perfect” fit). For $p = 0$, the solution is the traditional least-squares line fit to the data (“zero curvature”). The transition region for the parameter p is data-dependent and can be quite sensitive. See Figure [2.2].

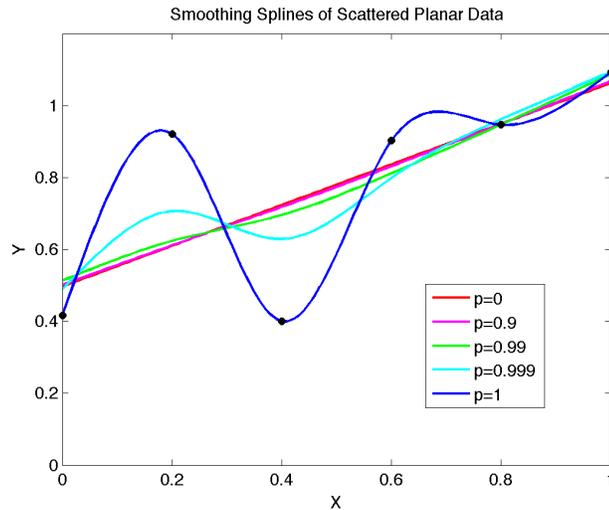


Figure 2.2: Smoothing spline sensitivity to parameter p . Depending on the data (black), small variations in p can have a large impact on the quality of fit for the smoothing spline. Notice the difference between $p = 0.99$ and $p = 1$.

In general, to control the level of fit at different data sites, weights w_i can be introduced to the smoothing spline; namely the smoothing spline is the minimizer of

the variational problem

$$\min_{s \in \mathbf{S}_{\Delta,k}} p \sum_{i=0}^m w_i [y_i - s(x_i)]^2 + (1-p) \int_{x_0}^{x_m} [s''(t)]^2 dt$$

If all data sites are to be treated equally, weights $w_i = 1$ suffice. The MATLAB function *csaps* in the Curve Fitting Toolbox allows for easy computational access to these smoothing splines.

2.2 Fitting Curves and Surfaces in \mathbb{R}^3

The extension of the smoothing spline to general curves in \mathbb{R}^3 is simple. Instead of the knots x_i being coordinates along a real axis, the knots can represent a parametric description of the curve. An obvious example is the arc length of a curve. For ordered sets of discrete \mathbb{R}^3 data, another choice for the knots would simply be the data label or site. With this in mind, the smoothing spline problem becomes

$$\min_{\mathbf{s} \in \mathbf{S}_{\Delta,k}} p \sum_{i=0}^m w_i [\mathbf{y}_i - \mathbf{s}(x_i)]^2 + (1-p) \int_{x_0}^{x_m} |D^2 \mathbf{s}(t)|^2 dt$$

where $\mathbf{s} : [x_0, x_m] \rightarrow \mathbb{R}^3$ is now a \mathbb{R}^3 -valued spline function and $\mathbf{S}_{\Delta,k}$ is the vector space of \mathbb{R}^3 -valued spline functions of order k on the mesh of data sites Δ . Different weights w_i can again be assigned to specific data sites. See Figure [2.3] for an example. Once again, the MATLAB function *csaps* performs efficiently.

Surface fitting in \mathbb{R}^3 with smoothing splines is a slightly different question. When

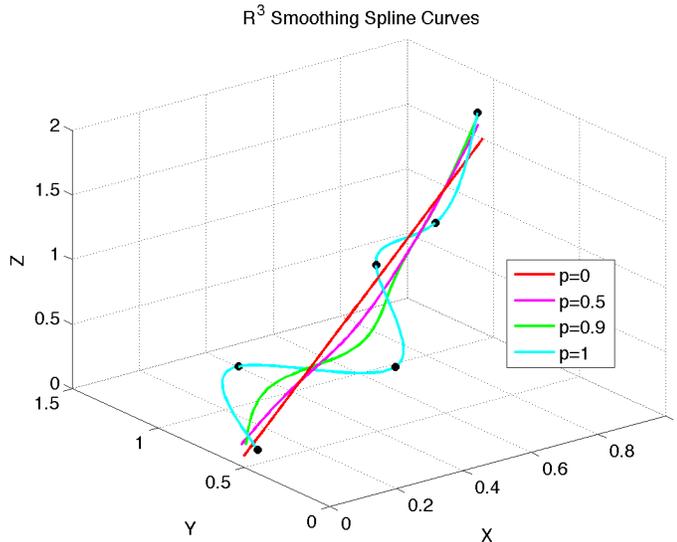


Figure 2.3: Smoothing splines of \mathbb{R}^3 scattered data (black). MATLAB Curve Fitting Toolbox *csaps* makes these constructions simple and easy to visualize.

there is no underlying grid beneath the data $z = f(x, y)$, one can opt for the *thin-plate smoothing spline*, which is an extension of smoothing splines to multidimensional data. The computational complexity increases; one can explore the MATLAB documentation for *tpaps* for details. However, whenever data is given on defined grid, say $z_{ij} = f(x_i, y_j)$, then *tensor product splines* become a viable alternative to the costly thin-plate splines. A tensor product spline $s(x, y)$ is a member of the linear space $S_{\Delta_1, k_1} \otimes S_{\Delta_2, k_2}$. The *B-spline* representation takes the form

$$s(x, y) = \sum_i \sum_j c_{ij} B_{i, \Delta_1}(x) B_{j, \Delta_2}(y)$$

while the pp-form consists of multiple bivariate polynomials. Again, the MATLAB function *csaps* is up to the task, as working with gridded data to produce a tensor

product spline is simple and straightforward. See Figure [2.4]. As we will see in the following chapter, tensor product splines play a fundamental role in the static modeling of the mitral valve.

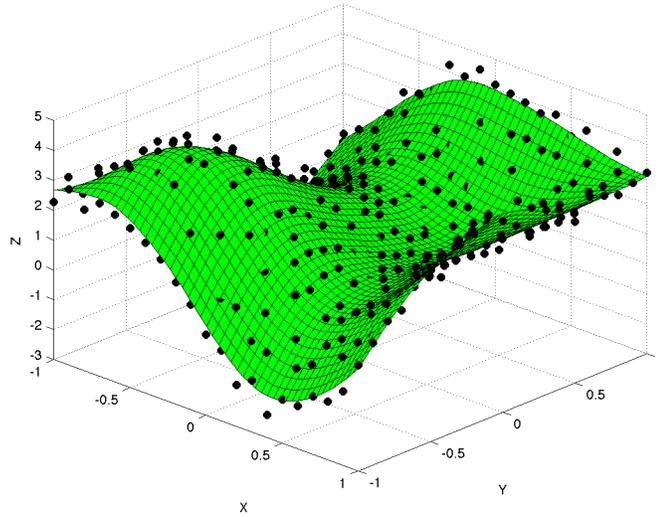


Figure 2.4: Tensor product spline of \mathbb{R}^3 scattered data (black). The smoothing parameter $p = 0.998$ was chosen to display a smooth surface without overfitting. MATLAB Curve Fitting Toolbox *csaps* allows for quick construction and evaluation of tensor product splines.

Chapter 3

Static Modeling of the Mitral Valve

The purpose of this chapter is to highlight the modeling procedures used to make patient-specific static models of the mitral valve. First, background on the mitral valve anatomy and mitral valve imaging is presented. Then, after describing the data provided for the modeling process, the spline methods from Chapter 2 are applied to generate models of the mitral valve. Smoothing splines are applied to the modeling of annulus and coaptation data, while tensor product smoothing splines are used to fit surfaces to the leaflet data. Two methods are presented for this surface modeling; one (lofting) works in all cases but requires more work to execute while the other (principal plane) is simpler in its approach but does not work for all mitral valves. The details of mitral valve static modeling presented here form the first step to the dynamic modeling presented in team's published paper [5].

3.1 Mitral Valve Anatomy

The mitral valve (MV) is a dual-flap valve in the heart that controls the flow of blood. It is known as the left atrioventricular valve because it lies between the left atrium and the left ventricle. A normal-functioning MV opens during diastole to allow blood to flow from the left atrium into the left ventricle. It is opened by contracting papillary muscles, which are pulling open the leaflets during diastole. Following diastole (begin systole), the valve is closed via ventricular pressure to prevent blood flowback (regurgitation) to the atrium.

The mitral valve's central components are the anterior leaflet (AL), the posterior leaflet (PL), the mitral annulus (MA) and the subvalvular apparatus (see Figure [3.1]).

The MA is a ring-like fibrous entity with a 3D shape resembling a saddle. The mitral commissures are the points where both leaflets free edges join—they do not necessarily coincide with the mitral annulus. The subvalvular apparatus consists of the chordae tendinae and the papillary muscles.

3.2 Mitral Valve Imaging

Mitral valve modeling begins with the imaging of a patient's mitral valve via 3D ultrasound echocardiography. Specifically, transesophageal 3D-Echo data sets were acquired for all patients by mid-esophageal full volume echocardiography, with an IE-33 Philips ultrasound system and X7-2t probe. The framerate for the acquisition

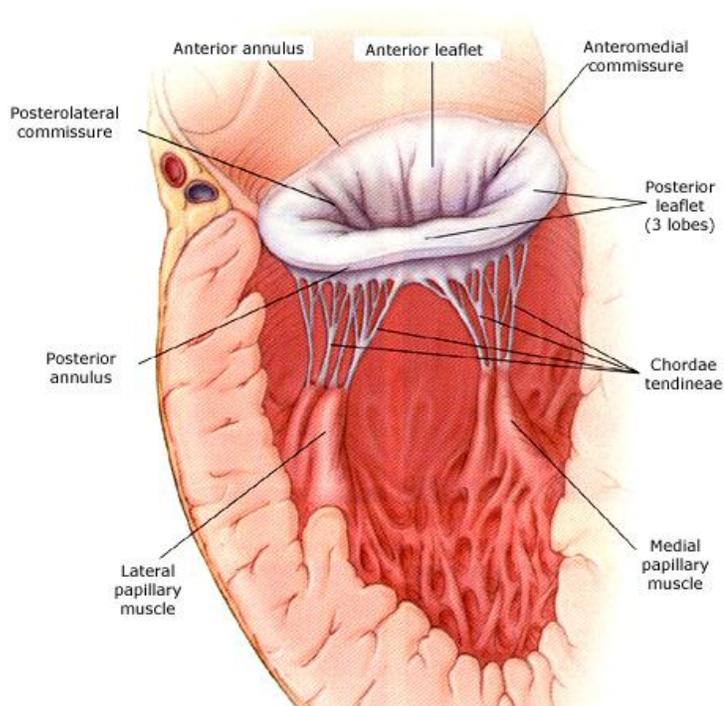


Figure 3.1: Figure from www.heart-valve-surgery.com

was approximately 25 fps, with each frame having resolution $208 \times 224 \times 208$ voxels. All 4D data sets conformed to the standard Digital Imaging and Communications in Medicine (DICOM) format. Each 4D dataset was uploaded to SLICER3D, a freeware 3D-display environment, for interactive tagging by the expert cardiology clinician Dr. Sagit Ben Zekry. This tagging was facilitated through the use of a tagging software module developed by a team at University of Houston. The team is led by Dr. Robert Azencott, and in addition to myself, has included at one time or another Dr. Aarti Jajoo, Dr. Jiwen He, Dr. Ronald Hoppe, Dr. Roland Glowinski, Dr. Yue Qin, Dr. Andrey Martynenko, Dr. Saurabh Jain, Dr. Yipeng Li, and Dr. Simon Alexander. Cardiologists Dr. William Zoghbi, Dr. Stephen Little, and Dr. Gerald Lawrie have also contributed to the team.

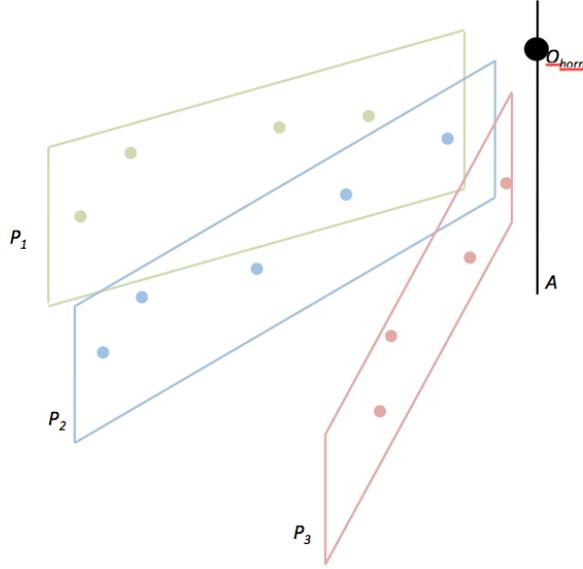


Figure 3.2: Schematic of tagging planes P_i , rotation axis A and anterior horn O_{horn} . This simple example shows three color-coded tagging planes and their corresponding tags.

In order to tag (select) points for a MV model on any 3D image frame, a rotation axis A is first visually identified. This axis is constrained to (a) be “roughly” orthogonal to the main principal components plane associated to the MV and (b) pass through the tip O_{horn} of the MA anterior horn. Next, numerous planes $P_1 \dots P_n$, each containing the axis A , are selected by small successive rotations around A , in order to span the entire leaflets AL and PL (see Figures [3.2-3.3]). Within each plane P_i , the set of tags included the intersection points of P_i with the coaptation line and the annulus, as well as 5 to 10 points on each leaflet. For each tag x in P_i one defines the “polar” coordinates (r, θ) of x as the node label r (determined by distance between O_{horn} and x), and the angle $\theta = \theta_i$ between planes P_1 and P_i . For certain tagging sequences (see lofting section below) this representation required some care.

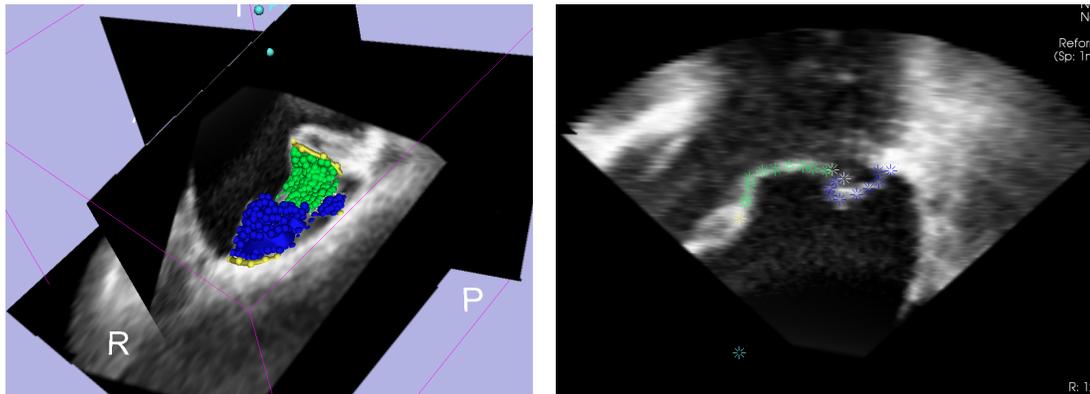


Figure 3.3: (left) Interactive tagging in the SLICER 3D environment. (right) A sample tagging plane P_i .

All tagging is performed at mid-systole (MS) and end-systole (ES) to give two reliable reference positions for each component of the studied MV model. Four components were of primary focus: the AL, the PL, the MA, and the coaptation line (COAPT) along which the leaflets close during systole. It should be pointed out that during these two tagging operations, one does not need to maintain any matching between the tags at time MS and those at time ES; in other words, the models are constructed independent of each other.

3.3 Mitral Valve Modeling

For each tagged MV, the tags recorded at MS and ES are exported to the MATLAB environment in order to compute patient-specific smooth geometric models of the four MV components. Let S_{AL} (resp. S_{PL} , S_{MA} , S_{CO}) denote the set of tags for the AL (resp. PL, MA, COAPT). In all 60 cases studied, the set cardinalities typically

work out to about

$$|S_{AL}| \approx 200, \quad |S_{PL}| \approx 200, \quad |S_{MA}| \approx 60, \quad |S_{CO}| \approx 40$$

We now explain the modeling of each of the MV components. The smoothing spline techniques introduced in Chapter 2 play a significant role.

3.3.1 Mitral Annulus and Coaptation Models

The methods described in Chapter 2 for fitting curves in \mathbb{R}^3 were applied to generate models for the two curves MA and COAPT. Specifically, cubic smoothing splines were obtained for each of the smooth parametric equations

$$x = a(s), \quad y = b(s), \quad z = c(s) \tag{3.1}$$

where s denotes the curve arc-length. See Figure [3.4] for an visualization of these models for a sample valve. Note the annulus is a closed curve, so periodic boundary conditions are required in the spline fitting. Also, for models where the anterior and posterior leaflets do not properly “seal” at systole (this is called regurgitation), the COAPT must be modeled by the union of the two leaflet interior edges ALE and PLE. These curves typically diverge from one another at the regurgitation zone. Thus cubic smoothing splines are needed for each of the curves ALE and PLE, and collection of these models can be designated COAPT. See Figure [3.5] for an example.

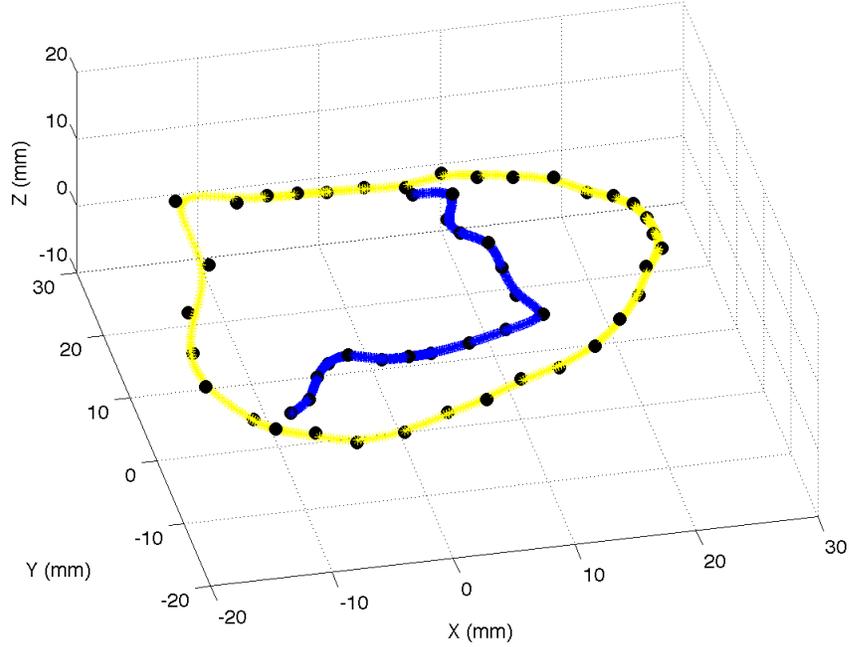


Figure 3.4: Sample models for patient’s mitral annulus (yellow) and coaptation line (blue). The tagged data sites are the black circles.

3.3.2 Principal Plane Leaflet Models

Principal Components Analysis (PCA) of S_{AL} and S_{PL} generates two new orthonormal bases in R^3 , and thus two new coordinate systems (x, y, z) for the AL and (X, Y, Z) for the PL. For notational brevity, we proceed to mention only a single system (x, y, z) for an arbitrary leaflet. For the majority of MV models exhibiting simple geometry, the leaflet surface can be fitted by a single cartesian equation, denoted

$$z = f(x, y) \tag{3.2}$$

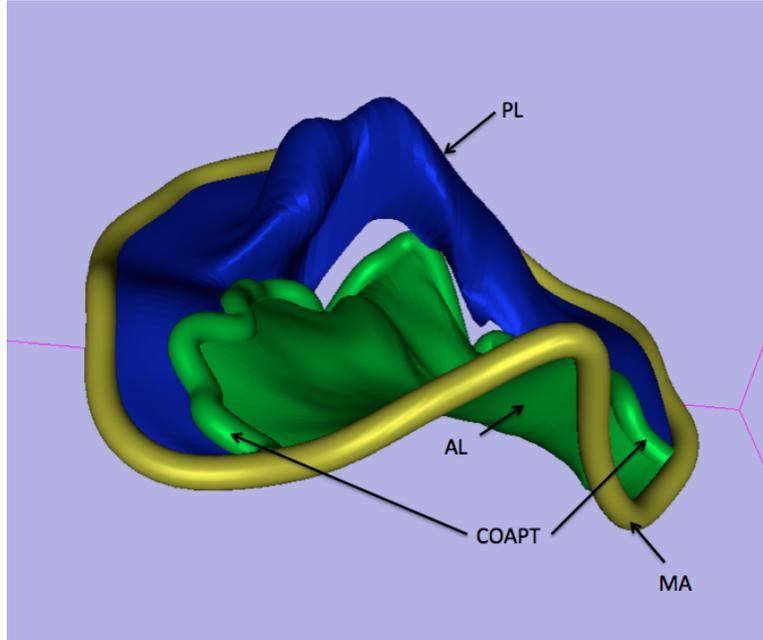


Figure 3.5: Patient-specific MV model (a regurgitation case)

Thus the tagged data has coordinates $(x, y, f(x, y))$ with respect to the principal plane coordinate system. In this system, the tagged data is often scattered and not on a regular grid. By using Delaunay tessellation (see [11]), this scattered data is interpolated onto a regular grid in the (x, y) -plane. On this grid, the function f is represented by a *bi-cubic* (tensor product) *smoothing spline*, which is a linear combination of specific polynomials of x, y of degree 3 with compact support. Recalling from Chapter 2, smoothing spline fitting of the equation $z = f(x, y)$ minimizes the sum of squared errors of fit of this equation to tagged data plus a quadratic term controlling the “smoothness” (curvature) of f . In practice, principal plane leaflet models are computed in two stages—the first is a preliminary smoothing spline construction where equal weighting is applied to each tagged datum. This produces a function $s_1(x, y)$ from which one can compute the errors of fit $\|s_1(x_i, y_i) - f(x_i, y_i)\|$.

Using these errors of fits as weights w_i in a second smoothing spline fit allows for custom weighting of the tagged data and thus fine-tuned adjustments to the model's fit at the tagged locations. See Figure [3.6] for a sample AL model.

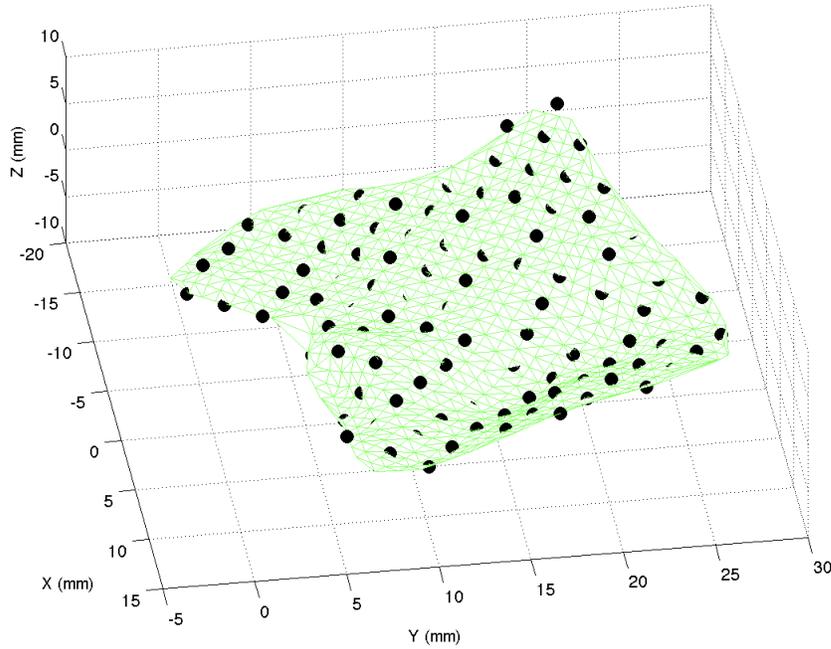


Figure 3.6: Sample principal plane leaflet model. The tagged data sites are the black circles.

Principal plane modeling of leaflets works whenever a component of the tagged data can be expressed as a function of the remaining two components. For select MV models that exhibit large curvatures, exotic loops, folds, cusps, and other complicated shapes, there are geometric obstructions to the preceding cartesian approach. See Figure [3.7]. In the bottom figure, the bulbous prolapsed tissue is curved so that no cartesian representation $z = f(x, y)$ is possible. In layman's terms, that surface does not pass a vertical line test in any cartesian coordinate system.

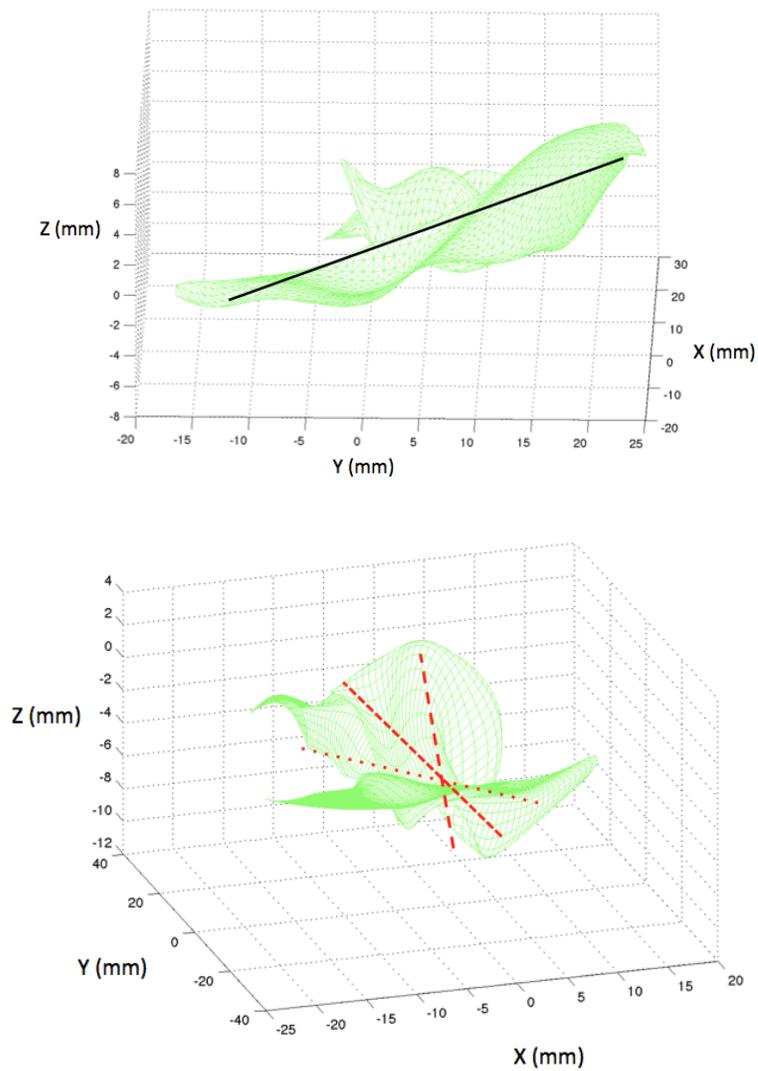


Figure 3.7: Limitations of principal plane modeling. The top figure gives a model where the principal plane (black) allows for a cartesian expression $z = f(x, y)$. The bottom figure gives a model (lofting method, see next section) where no cartesian expression $z = f(x, y)$ is possible. Three candidate plane cross-sections (red) are displayed; each results in double values z_1, z_2 for certain (x, y) pairs.

3.3.3 Lofting Leaflet Models

As the last section showed, whenever one is unable to find a coordinate system (x, y, z) that uniquely expresses a leaflet point by its projected coordinates in the principal planes (that is $z = f(x, y)$), another modeling approach is required. For these cases, we model the leaflet S by 3 parametric equations

$$x = u(r, \theta), \quad y = v(r, \theta), \quad z = w(r, \theta) \quad (3.3)$$

where u, v, w are patient-specific bi-cubic smoothing splines functions of the polar coordinates (r, θ) defined above in the tagging process. The construction of these tensor product splines is as follows. First, for fixed $\theta = \theta_n$, the intersections of S with the plane P_n are fitted by a planar smoothing spline curve C_n . The collection of these planar spline curves C_n will informally be called the *radial fit*. Note that because of the possible ambiguity in assigning knot labels based on Euclidean distance from the anterior horn O_{horn} , some knot re-ordering is often required to produce acceptable spline functions (see below). Once cross-sectional curves C_n have been determined for all θ_n , each spline function C_n is then sampled to produce a fixed N nodes per curve. Then for each node $r = r_j$, with $1 \leq j \leq N$ the collection of all nodes across curves C_n are fitted by an additional smoothing spline curve R_j . The collection of all radial cross-sectional curves R_j is informally called the *angular fit*. See Figure [3.8] for a schematic illustrating this process. This so-called *lofting* technique (see [15]) has proved successful with modeling challenging mitral valve leaflets. See Figure [3.9] for an example lofted leaflet. This modeling method, however, did present certain

difficulties that we had to overcome.

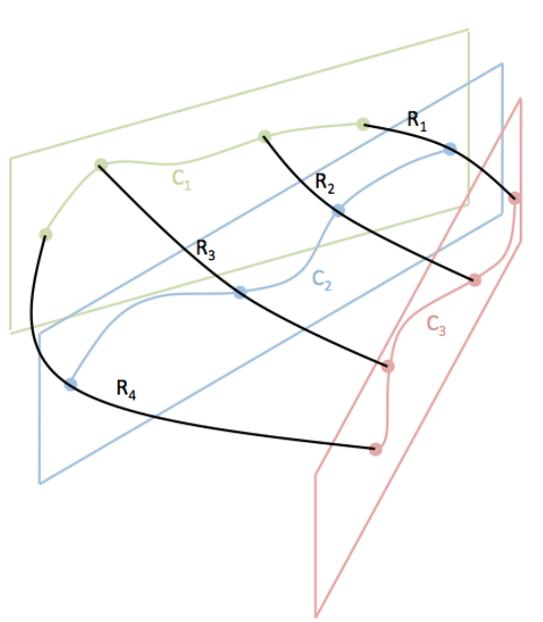


Figure 3.8: Schematic illustrating the radial curves C_1 (green), C_2 (blue), and C_3 (red), as well as the angular curves R_1, R_2, R_3, R_4 (black).

The first difficulty is the previously mentioned re-ordering of the knots required in the radial fit. Most tagged data arrives in a “block” file with no record of the distinct planes P_i used during data acquisition. During data processing, these planes (and thus the angles θ between them) must be recovered. Using a recursive procedure, recovery of these planes is simple. Denoting an arbitrary tag of S by x_1 , the points O_{horn}, x_1 and a fixed point y on the rotation axis A determine a unique tag plane P_1 . Remove all tags in S that are within the plane P_1 . Then proceed to choose another tag x_2 in S and consider the plane formed by O_{horn}, y , and x_2 . Continue until the tag list is exhausted.

Following recovery of the tagging planes, the tagged data within each recovered

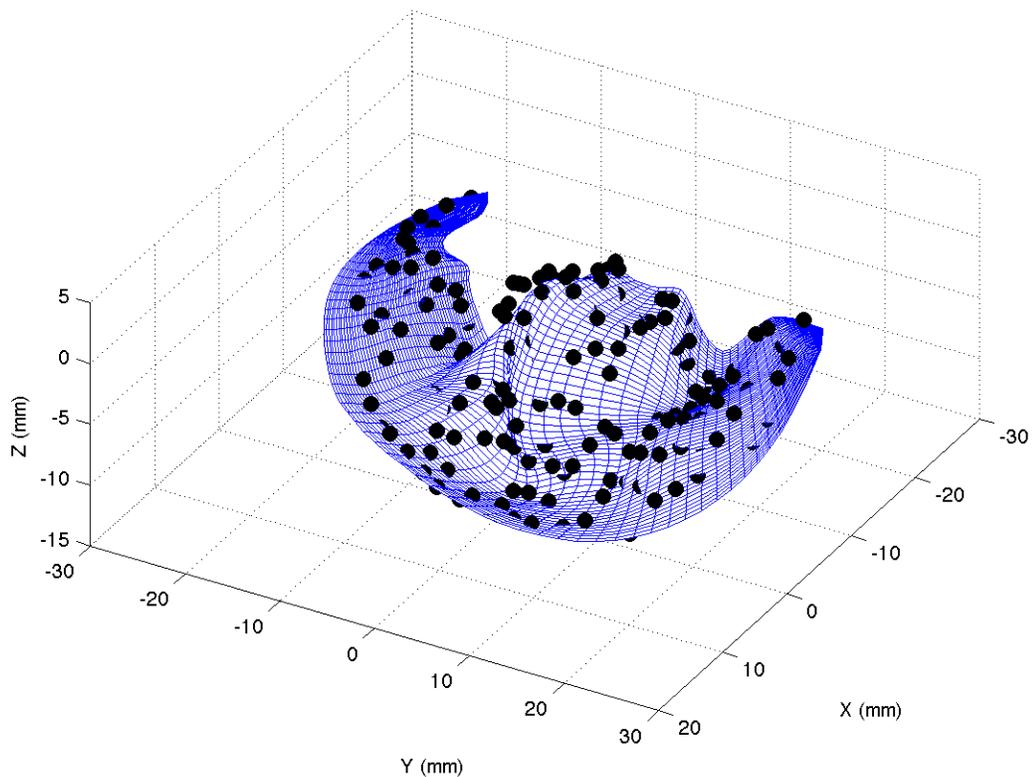


Figure 3.9: Sample lofting leaflet model. The tagged data sites are the black circles.

plane P_i must be ordered so that the fitting spline preserves the shape of the structure that the tagging cardiologist identified. Although a “human-in-the-loop” is required to make these adjustments, we have written several software tools to facilitate and visualize this process. These tools allow for interactive reordering of the tags, visualization of multiple tag planes, and synthetic tagging plane insertion for sparse data sets. See Figure [3.10-3.11] for some snapshots of these tools in action.

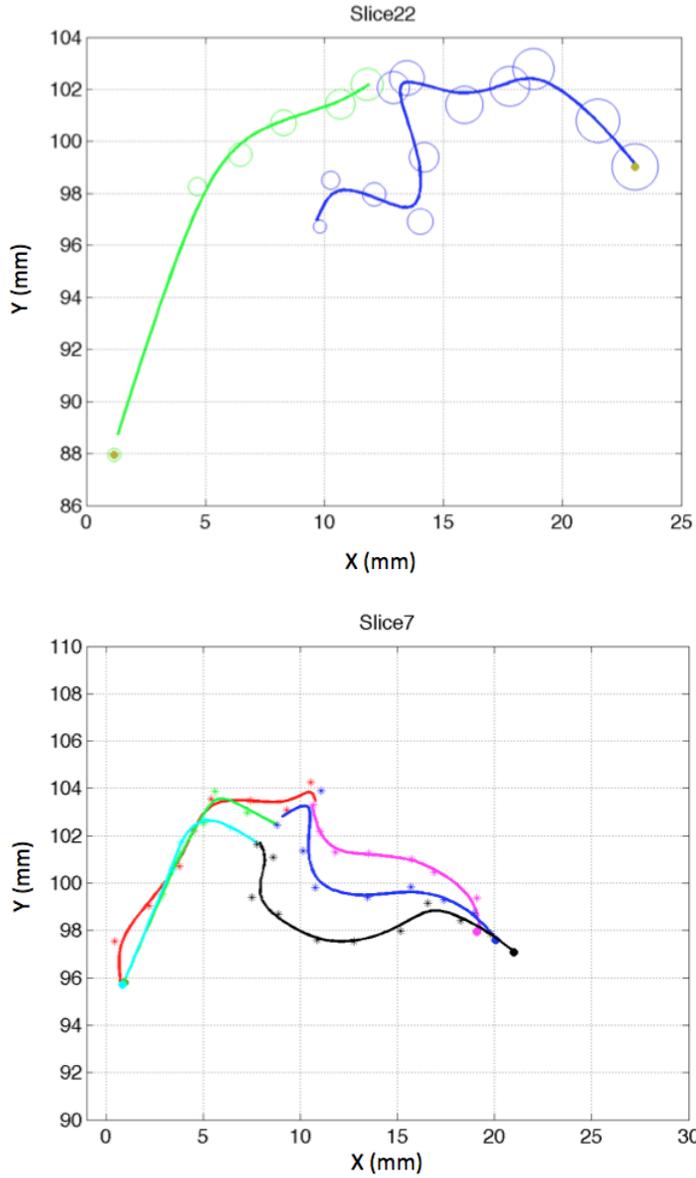


Figure 3.10: Human-in-the-loop tools for radial fit. (Top) Interactive spline knot adjustment. Annular tags are yellow circles, anterior leaflet tags are green circles, and posterior leaflet tags are blue circles. Whenever loops or irregularities are detected, knots are re-ordered and re-checked. (Bottom) Once all tagged planes have been processed, neighboring tagged planes are visualized concurrently. The color scheme: previous AL–red, previous PL–magenta; current AL–green, current PL–blue; next AL–cyan, next PL–black. These plots allow for preliminary “continuity” checks before the angular fit is attempted.

A second issue stems from sparse data during the angular fit. Typical models arrive with approximately 30-50 tagging planes; whenever the number of tag planes is small (< 10), or delivered tags are suspect and deemed inaccurate (even cardiologists make mistakes!), existing tag planes must be “interpolated” to produce synthetic “ghost planes”. These virtual planes ensure that large gaps in θ do not generate ridges or almost planar surface patches to fill the gaps in θ in the resulting model. Experience has shown that recursive insertion of virtual planes is sometimes required; two given accurate radial curves C_1, C_2 can be “interpolated” to find the profile $C_{1.5}$. Specifically, each curve C_1, C_2 is sampled to an equal number of M points. Then corresponding points $x_i \in C_1$ and $y_i \in C_2$ are interpolated. Linear interpolation suffices for this first step. Naturally, the same methodology can be used to compute $C_{1.25}$ and $C_{1.75}$. For small enough $\theta_2 - \theta_1$ these three planes may suffice. Otherwise, the three curves $C_1, C_{1.5}, C_2$ can be sampled and corresponding points can be interpolated with a smoothing spline. This often yields smoother transitions across large gaps $\theta_2 - \theta_1$. But for the valves modeled with lofting, linear interpolation often sufficed. See Figure [3.11] for a snapshot of the “ghost plane” tool in action.

A final problem is the “crease” problem for the angular fit: for knots very close to the rotation axis A (i.e. points along the angular curves R_j close to A), there will be abrupt ridges and visual roughness for the surface fitted to these knots. Moreover reducing the smoothing parameter p in the smoothing spline fit does not help eliminate this issue. One remedy is to pre-smooth the problem knots with a polynomial least squares fit in spherical coordinates (ρ, θ, ϕ) relative to the origin O_{horn} . That is, for numerous fixed ranges $\rho_i \leq \rho \leq \rho_{i+1}$, determine the polynomial

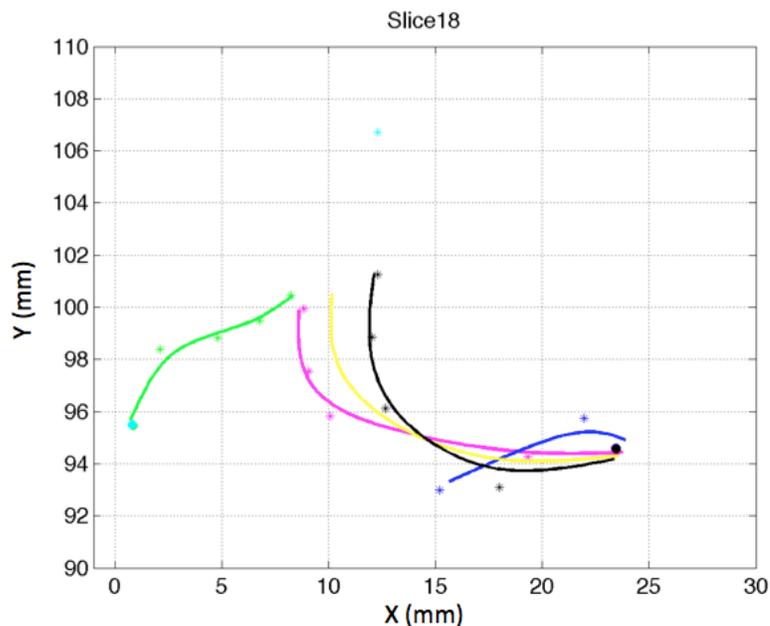


Figure 3.11: Ghost data (yellow) over-layed on sparse PL data. The previous PL plane C_{i-1} (magenta) and next PL plane C_{i+1} (black) do not match the current PL plane C_i (blue). By sampling these curves to M equal values and interpolating corresponding points $x_j \in C_{i-1}, y_j \in C_i, z_j \in C_{i+1}$, a smoother transition is achieved. The current AL plane (green) is shown. It is expected that this valve be closed at systole; this validates the position of the derived ghost data (yellow).

$\phi' = p_i(\theta)$ that minimizes the error $\|p_i(\theta) - \phi\|^2$. Polynomials of degree 2 or 3 have shown to adequately smooth and preserve the desired AL geometry at the anterior horn O_{horn} . See Figure [3.12].

3.3.4 Computing Times and Quality of Fit

These procedures provide smooth, patient-specific smoothing spline models of the four MVA components at tagged instances MS and ES. Moreover, for each model generated, the statistics of the errors of fit are tracked. All models are constrained so

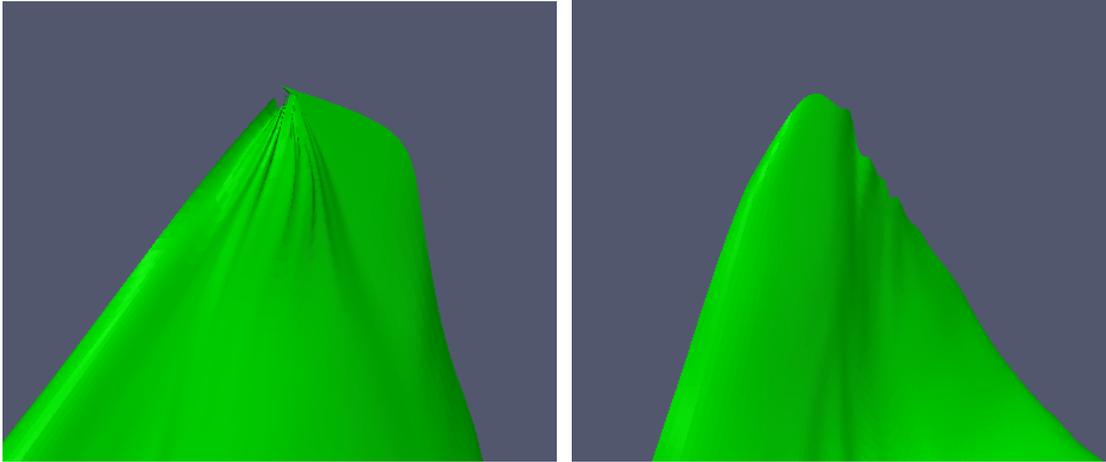


Figure 3.12: Crease smoothing. (Left) Radial fits in tagging planes converge at anterior horn O_{horn} . Minor changes in azimuthal angle ϕ greatly affect the smoothness of the angular fit in the polar angle θ . This is apparent in the ridges and creases that form near the anterior horn. (Right) Least-squares smoothing of the azimuthal angle ϕ across the range of polar angles θ reduces these ridges. Provided smoothing is restricted to small neighborhood of O_{horn} , the resulting model does not deviate from tagged data.

the individual errors of fit are on the order of 1 mm. As mentioned in Chapter 2, the MATLAB curve fitting toolbox function *csaps* makes smoothing spline fits quick and computationally accessible. Typical fits are obtained in the order of a few seconds on a computer with a 2.53 GHz Intel Core i5 processor and 4 GB of RAM. The only slow runtime component of the surface fitting process is the human-in-the-loop knot checking that is required for lofting leaflets. This process takes, on average, the order of 1 hour. See Figure [3.13] for some examples of completed MVA models.

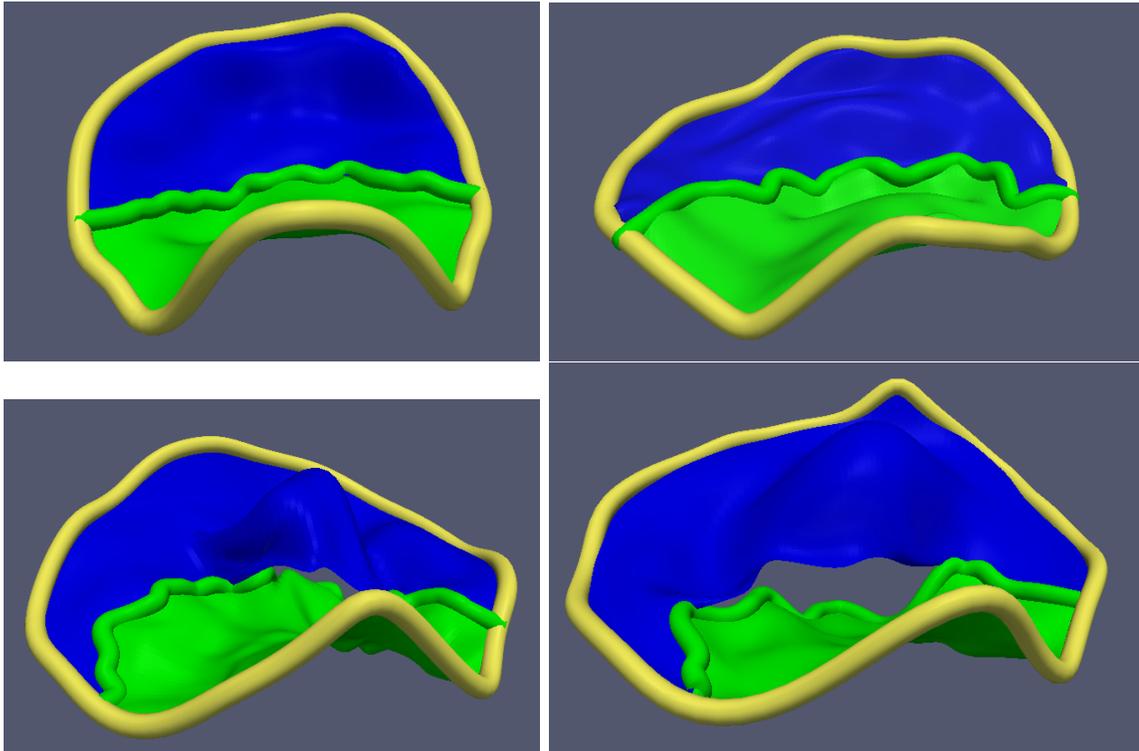


Figure 3.13: Patient-specific MVA models. The top row are two models for normal, healthy patients, while the bottom row are models for diseased patients exhibiting larger valves and leaflet regurgitation.

Chapter 4

Diffeomorphic Shape Matching

4.1 Shape Matching Survey

Shape matching, or shape registration, is fundamental to the concept of image registration, which has many applications ranging from medical imaging to military target recognition. The broad aim of shape tracking is to recover a transformation of source or reference shape into a known target shape. Many methods exist to tackle this problem. They vary in complexity, but can be differentiated in their representations of the shapes and deformations.

4.1.1 Shape Representation in \mathbb{R}^3

We focus on shapes regularly embedded in \mathbb{R}^3 , which are curves, surfaces, or volumes with piecewise boundaries. More precisely, the set SH_d of d -dimensional shapes in

\mathbb{R}^3 will be the set such that any member $S \in SH_d$ satisfies

- (a) S is open and connected
- (b) \bar{S} is compact
- (c) ∂S is a piecewise smooth of class C_r with $r \geq 3$
- (d) S coincides with a whole connected component of $\mathbb{R}^3 - \partial S$
- (e) for each $x \in \partial S$, there is an open neighborhood U of x in \mathbb{R}^3 and a local r -smooth diffeomorphism ϕ of U onto an open ball $\tilde{U} \subset \mathbb{R}^3$ mapping $U \cap \bar{S}$ onto $\tilde{U} \cap H$, where H is the intersection of either one, two, or three closed half-spaces of \mathbb{R}^3 .

This definition is rather technical, but it rigorously specifies the intuitive notion of a “smooth shape”.

Smooth shapes are often represented parametrically. The explicit representation afforded by *spline functions* (see Chapter 2) often provide the computational accessibility and utility that is optimum for numerical work. Such is the case for the concrete models of anatomic shapes studied in this thesis. But a few implicit representations of shapes are also employed in shape matching studies. For instance, a shape can be represented as the set of all points $x \in \mathbb{R}^3$ such that $f(x) > 0$, where f is any smooth function. These implicit representations have proven to be quite useful when coupled with the free form deformation techniques (see [37] and [45]). Note that these representations include shapes that do not belong to SH_d —these types of shapes are not considered henceforth.

In numerical shape studies, one must inevitably discretize a shape into a finite set of points, which we call a point grid, where each point is well-identified by a finite set of its immediate neighbors (see Figure [4.1]). Point grids are often derived from sampling/evaluating a continuous parametric representation (like B-splines) of a shape. We denote a point grid for a curve in SH_1 by x_j , $1 \leq j \leq N$ and a point grid for a surface in SH_2 by x_{ij} , $1 \leq i \leq M$, $1 \leq j \leq N$. Here N (reps. $M \times N$) is the cardinality of the point grid.

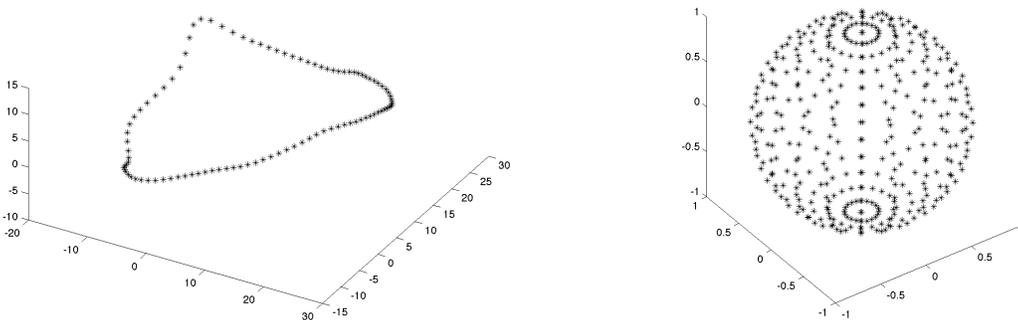


Figure 4.1: Point grid representation of various smooth shapes in \mathbb{R}^3 .

4.1.2 Shape Deformations

Regardless of how a shape is represented, there are multiple options to how to transform it into another shape. Two common circumstances warrant different techniques:

1. Shape A and Shape B represent the physical object at two different instants
2. Shape A and Shape B represent two different physical objects

For case (1), rigid deformations and scalings may suffice if the object is not deformable. Rigid deformations include translations and rotations, and scaling means homothetic scaling. A realistic example is target identification via satellite imagery. A reference image (for instance, military jet) is compared against an extracted target from the surveillance photographs. Thus the identification problem becomes determining the proper rigid transformation and scaling to identify and register the target with the reference image. See Figure [4.2].

Most often in case (1) as well as (2), elastic deformations are needed. These transformations allow for warping and deformation of the shape. Human organs are typically deformable shapes, as is the case of many biological shapes. See Figure [4.3] for 3 discrete instances of human mitral valve model. Here the registration problem becomes determining the proper elastic transformation to deform the top shape (begin-systole) into the middle shape (mid-systole) and then the bottom shape (end-systole), thereby approximating the true motion of the mitral valve during the systolic stage of the cardiac cycle. In general, mathematical methods to recover unknown deformations matching two or more given shapes fall into two categories: free form deformations (FFD) and landmark interpolation methods. In FFD, the deformation of the shape is realized by deforming the enveloping space. Diffeomorphic matching, which is explored below and is the main technique employed for studies in this thesis, falls under this category. FFD techniques often go hand in hand with implicit shape representations, as seen in [23] and [36]. For landmark interpolation methods, sets of landmark points are identified (therefore *known*) on the reference

and target shapes in order to capture the deformation between these sets. The deformation is constrained to map landmarks on the reference shape to corresponding landmarks on the target shape. Interpolants, for instance thin-plate splines or radial basis functions (see [7]), are then used to determine the deformations at other locations on the shapes.

4.2 Diffeomorphic Matching of Two Shapes

For the shape matching studies that follow, variational approaches are used to search for “optimal” diffeomorphic transformations. Consider two shapes S_0, S_1 that members of SH_d for $d \in \{1, 2, 3\}$. The search for an \mathbb{R}^3 diffeomorphism F such that $F(S_0) = S_1$ is an ill-posed problem which requires regularization to be numerically solved by variational methods. These variational methods have been explored by M. Miller, A. Trounev, L. Younes, J. Glaunes et.al. with applications to a quantified comparison of images of human brains (see [4], [13], [20], [25]). In their studies, the regularization is achieved by replacing the constraint $F(S_0) = S_1$ with a “softer” constraint based on the geometric disparity $\text{dis}(F(S_0), S_1)$ between the surfaces. For more detail, see [3].

In our study, we follow this method; specifically, our optimal shape matching transformation is formulated as a the solution to a minimization problem involving actions of \mathbb{R}^3 diffeomorphisms. The objective functional to be minimized consists of two parts: the kinetic energy of the deformations and the geometric matching quality. For numerical computations, point grid representations of the shapes give

rise to *diffeomorphic point matching*. Provided the deformation is sought at discrete times, the search for the optimal transformation reduces to the solution of a finite-dimensional minimization problem. Even then, the dimension of this problem is still high, for two point grids with cardinalities N and a time discretization of cardinality T means complete specification of all deformed trajectories at all times requires N^T data. A matching algorithm based on the first order necessary optimality conditions is presented and supported by some basic examples from geometry and medical imaging. The algorithm presented specifically mentions one target shape, but the methods has been generalized to the matching of multiple target shapes (see [3]).

4.2.1 Objective Functional Formulation

Fundamental to the idea of diffeomorphic matching is the concept of a vector field flow. Let V be any Hilbert space of vector fields on \mathbb{R}^3 and denote by $L^2([t_0, t_1], V)$ the Hilbert space of vector field flows $v : [t_0, t_1] \rightarrow v_t \in V$ with finite kinetic energy $E(v)$

$$E(v) := \frac{1}{2} \|v\|_{L^2([t_0, t_1], V)}^2 := \frac{1}{2} \int_{t_0}^{t_1} \|v_t\|_V^2 dt$$

Define the \mathbb{R}^3 *diffeomorphic flow* F_t as the solution to the dynamics equation

$$\begin{aligned} \frac{\partial F_t}{\partial t} &= v_t \circ F_t, & t \in (t_0, t_1] \\ F_{t_0} &= Id \end{aligned} \tag{4.1}$$

Assuming that our Hilbert space V is continuously embedded in an appropriate Sobolev space of smooth functions, there is a unique solution F_t to the above equation (see [17]).

Denote the solution to (4.1) for a given v by F_t^v and define $\text{Diff}(\mathbb{R}^3)$ to be the collection of \mathbb{R}^3 diffeomorphisms endowed with the topology of uniform convergence on bounded subsets of \mathbb{R}^3 . Furthermore, define $\mathcal{D}(V)$ to be the space of all disparity functionals $\text{Disp} : L^2([t_0, t_1], V) \rightarrow \mathbb{R}^+$ which are of the form $\text{Disp}(v) = \phi(F_{t_0}^v, F_{t_1}^v)$, where $\phi : \text{Diff}(\mathbb{R}^3)^2 \rightarrow \mathbb{R}^+$ is continuous with respect to the product topology on $\text{Diff}(\mathbb{R}^3)^2$. With these notions, we can define the objective functional $J : L^2([t_0, t_1], V) \rightarrow \mathbb{R}$ by

$$J(v) := E(v) + \phi(F_{t_0}^v, F_{t_1}^v)$$

This function ϕ is a necessary regularizing term, as without it the minimization problem is ill-posed ([3]). To enforce a good fit between shapes $S_0, S_1 \in SH_d$ ($d \in \{1, 2, 3\}$), ϕ should return a measure of geometric disparity between the shapes. There are a few choices used in practice. The first is the *pointwise disparity* functional. Identifying reference points x_1, \dots, x_N on reference shape S_0 and corresponding points y_1, \dots, y_N on the target shape S_1 , define

$$\phi(F_{t_0}^v, F_{t_1}^v) = \sum_{j=1}^N \|F_{t_1}^v(x_j) - y_j\|^2$$

Another commonly used disparity functional is the *Hausdorff distance*. Recalling the notation of a distance from point x to a set S by $d(x, S) = \min_{y \in S} d(x, y)$, the

classical Hausdorff disparities $h(S_0, S_1)$ are defined as

$$h(S_0, S_1) = \max_{x \in S_0} d(x, S_1)$$

These disparities determine the Hausdorff distance D_h

$$D_h(S_0, S_1) = \max(h(S_0, S_1), h(S_1, S_0))$$

With these concepts defined, the Hausdorff disparity functional takes the form

$$\phi(F_{t_0}^v, F_{t_1}^v) = D_h(F_{t_1}^v(S_0), S_1)$$

Regardless of which disparity functional is chosen, the following result is useful. For any $\phi \in \mathcal{D}(V)$, the minimization problem

$$\inf_{v \in L^2([t_0, t_1], V)} E(v) + \phi(F_{t_0}^v, F_{t_1}^v)$$

subject to the constraint (4.1) has a solution v^* . For proof, refer to [3].

This variational approach allows for the recovery of a diffeomorphic transformation of shape $S_0 \in SH_d$ at time t_0 into shape $S_1 \in SH_d$ at time t_1 via the warping of the enveloping 3D space. Specifically the deformations of shapes are modeled by a flow of time-dependent \mathbb{R}^3 diffeomorphisms F_t , determined by a vector field flow v_t on \mathbb{R}^3 . See Figure [4.4] for a two-dimensional example.

4.2.1.1 Self-Reproducing Kernel Hilbert Spaces

In applications, it helps to specify more concretely the Hilbert space V . Consider the radial Gaussian kernel

$$K_\sigma(x, y) = \frac{1}{(2\pi)^{3/2}\sigma^3} \exp\left(\frac{-\|x - y\|^2}{2\sigma^2}\right), \quad x, y \in \mathbb{R}^3$$

This choice of K_σ has demonstrated good performance for the applications that follow. The kernel K_σ is smooth, symmetric, bounded, and is known to be positive definite. Recall that a symmetric kernel $K : \mathbb{R}^3 \times \mathbb{R}^3 \rightarrow \mathbb{C}$ is *positive definite* if and only if for any finite set of points $\{x_1, \dots, x_n\} \subset \mathbb{R}^3$ and scalars $\{\lambda_1, \dots, \lambda_n\} \subset \mathbb{C}$

$$\sum_{i=1}^n \sum_{j=1}^n \bar{\lambda}_i K(x_i, x_j) \lambda_j \geq 0$$

Moreover a kernel K is said to be *reproducing kernel* of a Hilbert space H of functions on \mathbb{R}^3 if

(a) For every $x \in \mathbb{R}^3$, the function $K_x : \mathbb{R}^3 \rightarrow \mathbb{C}$ is a member of H , where

$$K_x(y) = K(y, x)$$

(b) For every $x \in \mathbb{R}^3$ and every $f \in H$, there holds $f(x) = (f, K_x)_H$

where $(\cdot, \cdot)_H$ denotes the inner product on H . A Hilbert space H is said to be a *reproducing kernel Hilbert space* (RKHS) when there exists a reproducing kernel K

of H . Notice, for a symmetric reproducing kernel, that

$$K(x, y) = K(y, x) = K_x(y) = (K_x, K_y)_H$$

An important result (see [2]) is that any positive definite kernel $K : \mathbb{R}^3 \times \mathbb{R}^3 \rightarrow \mathbb{C}$ uniquely determines a RKHS H of functions on \mathbb{R}^3 with reproducing kernel K . In particular, since K_σ is known to be positive definite, we may identify the Hilbert space V of vector fields on \mathbb{R}^3 with the RKHS V_{K_σ} of \mathbb{R}^3 vector fields defined by the kernel K_σ . In particular, for $v_t \in V_{K_\sigma}$, $z \in \mathbb{R}^3$, and any fixed finite set of points $\{x_1, \dots, x_n\} \subset \mathbb{R}^3$, we have the representation

$$\begin{aligned} v_t(z) &= (v_t, K_z)_{V_{K_\sigma}} \\ &= \left(\sum_{i=1}^n \alpha_i K_{x_i}, K_z \right)_{V_{K_\sigma}} \\ &= \sum_{i=1}^n \alpha_i (K_{x_i}, K_z)_{V_{K_\sigma}} \\ v_t(z) &= \sum_{i=1}^n \alpha_i K_\sigma(x_i, z) \end{aligned} \tag{4.2}$$

4.2.2 Diffeomorphic Point Matching of Two Shapes

Consider point grid representations X_0, X_1 of shapes $S_0, S_1 \in SH_3$. Then we have $X_0 = \{x_1, \dots, x_{N_0}\}$ and $X_1 = \{y_1, \dots, y_N\}$. Under the action of a diffeomorphism F_t^v , define the point set $\hat{X}_1 = F_{t_1}^v(X_0) = \{F_{t_1}^v(x_0), \dots, F_{t_1}^v(x_{N_0})\}$. This flow of diffeomorphisms actually defines N_0 trajectories $x_i(t) = F_t^v(x_i)$ with $1 \leq i \leq N_0$ and

$t \in [t_0, t_1]$. Moreover, by equation (4.2), the assumed self-reproducing structure of V means these N_0 trajectories completely determine the vector field flow at any point $z \in \mathbb{R}^3$:

$$v_t(z) = \sum_{n=1}^{N_0} K_\sigma(z, x_n(t)) \alpha_n(t), \quad \alpha_n(t) \in \mathbb{R}^3$$

In particular, the reproducing property implies the norm expression

$$\|v_t\|_V^2 = \sum_{i=1}^{N_0} \sum_{j=1}^{N_0} \alpha_i^T(t) K_\sigma(x_i(t), x_j(t)) \alpha_j(t)$$

Before proceeding further, introduce the following notation:

$$\begin{aligned} x^{(0)} &= (x_1, \dots, x_{N_0})^T \in \mathbb{R}^{3N_0} \\ x(t) &= (x_1(t), \dots, x_{N_0}(t))^T \in \mathbb{R}^{3N_0} \\ \alpha(t) &= (\alpha_1(t), \dots, \alpha_{N_0}(t))^T \in \mathbb{R}^{3N_0} \\ A(x(t)) &= [A_{ij}(x(t))]_{i,j=1}^{N_0} \in \mathbb{R}^{3N_0 \times 3N_0} \\ A_{ij}(x(t)) &= K_\sigma(x_i(t), x_j(t)) \mathbf{I}_3 \in \mathbb{R}^{3 \times 3} \\ \mathbf{I}_3 &= \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \end{aligned}$$

With these definitions and SRKHS structure, our minimization problem transforms from a search for vector field flows v_t to the search for the “kernel coefficient trajectories” $\alpha(t)$. Symbolically, we seek α^* (and corresponding trajectory x^*) such

that

$$\begin{aligned} J(\alpha^*) &= \min_{\alpha \in L^2([t_0, t_1], V)} J(\alpha) \\ &= \min_{\alpha \in L^2([t_0, t_1], V)} \frac{1}{2} \int_{t_0}^{t_1} \alpha(t)^T A(x(t)) \alpha(t) dt + \lambda \phi(F_{t_0}^\alpha, F_{t_1}^\alpha) \end{aligned}$$

subject to the constraint

$$\begin{aligned} \frac{dx^*(t)}{dt} &= A(x^*(t)) \alpha^*(t), \quad t \in (t_0, t_1] \\ x^*(0) &= x^{(0)} \end{aligned}$$

Here $\lambda > 0$ is a regularization parameter meant to adjust the balance between a minimum energy solution ($\lambda \rightarrow 0$) and the ideal geometric fit ($\lambda \rightarrow \infty$). The method of Lagrange Multipliers imply the existence of $p^*(t) = (p_1(t), \dots, p_{N_0}(t))^T \in \mathbb{R}^{3N_0}$, called the *adjoint state* such that

$$-\frac{dp^*(t)}{dt} = B(x^*(t), \alpha^*(t))^T \left[p^*(t) + \frac{1}{2} \alpha^*(t) \right], \quad t \in (t_0, t_1)$$

$$\lim_{t \rightarrow t_1^+} p^*(t) = 0$$

$$\lim_{t \rightarrow t_1^-} p^*(t) = \lim_{t \rightarrow t_1^+} p^*(t) + \nabla_{x^*} \phi(F_{t_0}^v, F_{t_1}^v)$$

$$A(x^*(t))(\alpha^*(t) + p^*(t)) = 0, \quad t \in (t_0, t_1]$$

where $B(x^*(t), \alpha^*(t)) = \nabla_x[A(x^*(t))\alpha^*(t)]$. To see this consider the Lagrangian

$$\begin{aligned} L(x, \alpha, p) &= J(\alpha) - \int_{t_0}^{t_1} p(t) \cdot \left(\frac{dx}{dt} - A(x(t))\alpha(t) \right) dt \\ &= - \int_{t_0}^{t_1} p(t) \cdot \frac{dx}{dt} dt + \int_{t_0}^{t_1} \left(p(t) + \frac{1}{2}\alpha(t) \right) \cdot A(x(t))\alpha(t) dt + \lambda\phi(F_{t_0}^v, F_{t_1}^v) \end{aligned}$$

Suppose (x^*, α^*, p^*) is a stationary point of the Lagrangian. Then all partial derivatives of L must vanish, leading to the above equations as first order necessary optimality conditions.

4.2.3 Numerical Algorithm for Diffeomorphic Matching of Two Shapes

Introduce the mesh/partition $\Delta_I = \{\tau^l\}$ of $I = [t_0, t_1]$

$$\Delta_I = \{t_0 = \tau^{L_0} < \tau^{L_0+1} < \dots < \tau^{L_1-1} < \tau^{L_1} = t_1\}$$

Let $L = \text{card}(\Delta_I)$. Introduce the discrete control space $\mathcal{U}^{\Delta_I} = \mathbb{R}^{L \cdot 3N_0}$ equipped with the inner product

$$(\alpha, \beta)_{\Delta_I} = \sum_{l=0}^{L-1} \Delta\tau^l (\alpha^l \cdot \beta^l)_{\mathbb{R}^{3N_0}} = \sum_{l=0}^{L-1} \sum_{j=1}^{N_0} \Delta\tau^l (\alpha_n^l \cdot \beta_n^l)_{\mathbb{R}^3}$$

Introduce $J^{\Delta_I}(\alpha)$, the discrete objective functional, as

$$J^{\Delta_I}(\alpha) = \sum_{l=0}^{L-1} \frac{\Delta\tau^l}{2} (\alpha^l)^T A(x^l)\alpha^l + \lambda\phi(F_{t_0}^\alpha, F_{t_1}^\alpha)$$

In particular, for all subsequent analysis, ϕ is fixed to be the Hausdorff disparity (see below). Discretizing the state equation via the forward Euler method and the adjoint (optimality conditions) equations by the backward Euler method gives

$$\begin{aligned}
\frac{x^{l+1} - x^l}{\Delta\tau^l} &= A(x^l)\alpha^l \\
x^0 &= x^{(0)} \\
\frac{p^{l-1} - p^l}{\Delta\tau^l} &= B(x^l, \alpha^l)^T \left[p^l + \frac{1}{2}\alpha^l \right] \\
p^{L_1+} &= 0 \\
p^{L_1-} &= p^{L_1+} + \lambda \nabla_{x^l} D_h(F_{t_1}(X_0), X_1) \\
A(x^l)(\alpha^l + p^l) &= 0
\end{aligned} \tag{4.3}$$

It turns out these equations are first order optimality conditions for the *finite dimensional* minimization problem

$$\min_{\alpha \in \mathcal{U}^{\Delta I}} J^{\Delta I}(\alpha)$$

subject to the constraint

$$\begin{aligned}
\frac{x^{l+1} - x^l}{\Delta\tau^l} &= A(x^l)\alpha^l \\
x^0 &= x^{(0)}
\end{aligned}$$

In the evaluation of the jump discontinuities for the adjoint p^l , we must evaluate the gradient of the geometric disparity function with respect to the variable x^l . For Hausdorff matching, this presents difficulties since Hausdorff distance is not a smooth

functional of these surfaces. To remedy this, a smooth Hausdorff disparity is adopted. First some notation. For our two point grids \hat{X}_1, X_1 (recall \hat{X}_1 is the deformed point grid X_0), define functions $\phi : \hat{X}_1 \rightarrow X_1$ and $\psi : X_1 \rightarrow \hat{X}_1$ by

$$\begin{aligned}\phi(x) &= \operatorname{argmin}_{y \in X_1} \|x - y\| \\ \psi(y) &= \operatorname{argmin}_{x \in \hat{X}_1} \|y - x\|\end{aligned}$$

Next inductively define the functions $\phi^n : X \rightarrow Y$ and $\psi^n : Y \rightarrow X$ by

$$\begin{aligned}\phi^n(x) &= \operatorname{argmin}_{y \in X_1 - \bigcup_{j=1}^{n-1} \{\phi^j(x)\}} \|x - y\| \\ \psi^n(y) &= \operatorname{argmin}_{x \in \hat{X}_1 - \bigcup_{j=1}^{n-1} \{\psi^j(y)\}} \|y - x\|\end{aligned}$$

With these concepts, the smooth Hausdorff disparity $D_s(F_{t_1}(S_0), S_1)$ between the deformed reference shape $F_{t_1}(S_0)$ and the target shape S_1 is given by

$$D_s(F_{t_1}(S_0), S_1) = \frac{1}{rN_0} \sum_{i=1}^{N_0} \sum_{n=1}^r \|x_i^{L_1} - \phi^n(x_i^{L_1})\|^2 + \frac{1}{rN} \sum_{j=1}^N \sum_{n=1}^r \|y_j - \psi^n(y_j)\|^2$$

Introducing more notation, let $\mathcal{J}_k = \{j : x_k^{L_1} \in \bigcup_{m=1}^r \{\psi^m(y_j)\}\}$ and $\mathcal{N}_{jk} = \{n : x_k^{L_1} = \psi^n(y_j)\}$. From this, we see that the gradient can be computed as

$$\frac{\partial D_s}{\partial x_k^{L_1}} = \frac{2}{rN_0} \sum_{n=1}^r [x_k^{L_1} - \phi^n(x_k^{L_1})] - \frac{2}{rN} \sum_{j \in \mathcal{J}_k} \sum_{n \in \mathcal{N}_{jk}} [y_j - \psi^n(y_j)]$$

It should be mentioned that the parameter r determines how many points are in the computed neighborhoods of the pull-back functions ϕ and ψ . When matching

curves in SH_1 (e.g. boundaries of surfaces in SH_2), the value $r = 3$ has proven successful. For interiors of SH_2 , the value $r = 5$ is taken. Thus to use correctly this smoothed Hausdorff term for general diffeomorphic matching of surfaces in SH_2 , one must systematically impose a “boundary to boundary” and “interior to interior” matching assignment. Denote by ∂X_0 the subset of grid points in X_0 corresponding to the boundary of the shape S_0 . Likewise, let $\text{int}(X_0)$ denote the grid points in X_0 corresponding to the interior of the shape S_0 . Letting similar notations hold for the deformed point grid \hat{X}_1 and target point grid X_1 , we can express this “boundary-to-boundary” and “interior-to-interior” constraint with the cost functional

$$\begin{aligned}
J^{\Delta_I}(\alpha) &= \sum_{l=0}^{L-1} \frac{\Delta\tau^l}{2} (\alpha^l)^T A(x^l) \alpha^l + \dots \\
&\quad \lambda \frac{1}{5|\text{int}(\hat{X}_1)|} \sum_{\{i:x_i^{L_1} \in \text{int}(\hat{X}_1)\}} \sum_{n=1}^5 \|x_i^{L_1} - \phi^n(x_i^{L_1})\|^2 + \dots \\
&\quad \lambda \frac{1}{5|\text{int}(X_1)|} \sum_{\{j:y_j \in \text{int}(X_1)\}} \sum_{n=1}^5 \|y_j - \psi^n(y_j)\|^2 + \dots \\
&\quad \lambda \frac{1}{3|\text{int}(\hat{X}_1)|} \sum_{\{i:x_i^{L_1} \in \partial\hat{X}_1\}} \sum_{n=1}^3 \|x_i^{L_1} - \phi^n(x_i^{L_1})\|^2 + \dots \\
&\quad \lambda \frac{1}{3|\text{int}(X_1)|} \sum_{\{j:y_j \in \partial X_1\}} \sum_{n=1}^3 \|y_j - \psi^n(y_j)\|^2
\end{aligned}$$

Without this restriction, matching quality is poor. Selection of the regularization parameter λ follows the continuation methodology laid out in [3], which we now outline.

For a fixed initial value of λ , a gradient descent in α is performed. In all concrete

implementations considered here, as well as in [3], [8], an Armijo line search is employed to guarantee sufficient convergence with ill-conditioned data. Recall that an Armijo line search ensures that the step-length chosen in the gradient descent results in a decrease of the objective functional by *at least* as much as a fraction β of that guaranteed by the first-order Taylor approximation to the objective functional at the current iterate. Symbolically, if the current iterate is denoted α_k , we select α_{k+1} by

$$\alpha_{k+1} = \alpha_k - \xi \nabla J(\alpha_k)$$

where ξ is chosen (often by bisection) so that

$$J(\alpha_{k+1}) \leq J(\alpha_k) - \xi \beta \|\nabla J(\alpha_k)\|^2$$

Small values of β (< 0.0001) result in slower but more stable decreases in the objective functional. This gradient descent technique proceeds until the norm of the gradient has sufficiently decreased. At this point, the “inner” loop terminates, the value of λ is increased in the “outer” loop, and another gradient descent inner loop is performed. The whole algorithm terminates when sufficient geometric accuracy has been achieved. Oftentimes, this is when some fixed quantile of the geometric errors of fit (say 90%) fall below a predetermined threshold. For our shape matching studies, a threshold on the order of 1 mm often sufficed.

4.3 Numerical Algorithm for Diffeomorphic Matching with Multiple Target Shapes

The above algorithm and underlying theory easily generalizes to sequences of shapes $S_0, S_1, S_2, \dots, S_q$. Specifically, given $q+1$ shape snapshots S_i , each shot being indexed by the time t_i with $t_i < t_{i+1}$, we seek a diffeomorphic deformation flow F_t , fully characterized by its associated vector field flow $v \in L^2([t_0, t_q], V)$, that minimizes

$$J(v) = E(v) + \sum_{i=1}^q \lambda_i D_h(F_{t_i}^v(S_0), S_i)$$

subject to the constraint

$$\begin{aligned} \frac{\partial F_t}{\partial t} &= v_t \circ F_t, & t \in (t_0, t_q] \\ F_{t_0} &= Id \end{aligned}$$

Numerically, the algorithm becomes

$$\begin{aligned} \frac{x^{l+1} - x^l}{\Delta\tau^l} &= A(x^l)\alpha^l \\ x^0 &= x^{(0)} \\ \frac{p^{l-1} - p^l}{\Delta\tau^l} &= B(x^l, \alpha^l)^T \left[p^l + \frac{1}{2}\alpha^l \right] \\ p^{L_q+} &= 0 \\ p^{L_i-} &= p^{L_i+} + \lambda \nabla_{x^l} D_h(F_{t_i}(X_0), S_i) \\ A(x^l)(\alpha^l + p^l) &= 0 \end{aligned}$$

An example of this technique is given in the next chapter. For the majority of the mitral valve dynamic diffeomorphic modeling cases studied in this thesis, however, only two pairs of static surfaces (or curves) are matched by diffeomorphisms, as the echocardiographic tagging was restricted to mitral valves observed at mid and end-systole. Thus the added generality of this approach was only needed in very few of our concrete mitral valve dynamic modeling studies and the methods of the previous section suffice.

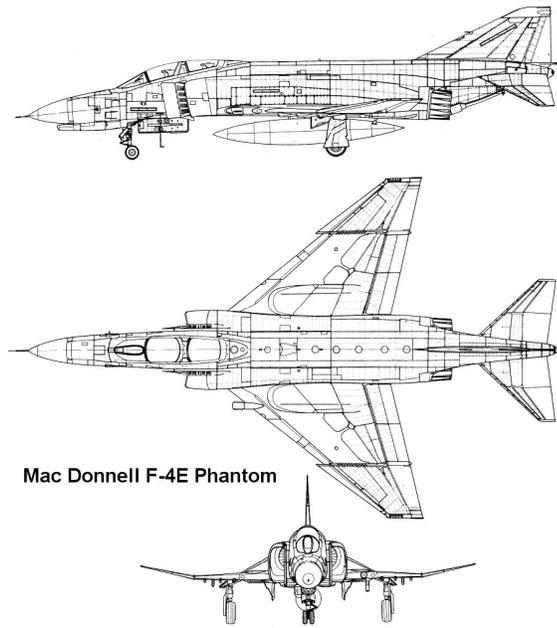


Figure 4.2: The reference F4 Phantom can be identified via overhead satellite imagery through the use of global rigid transformations. Top image: www.eliteday.com. Bottom image: ngm.nationalgeographic.com.

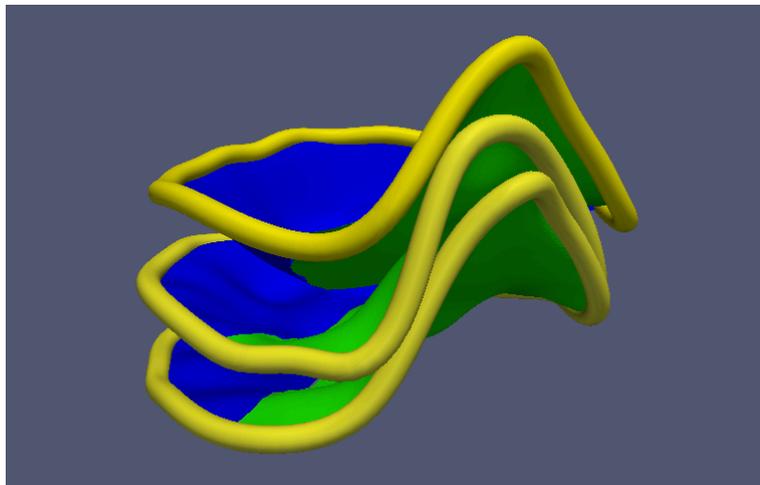


Figure 4.3: A human mitral valve at three different instants. The static models were created using the techniques of Chapter 3.

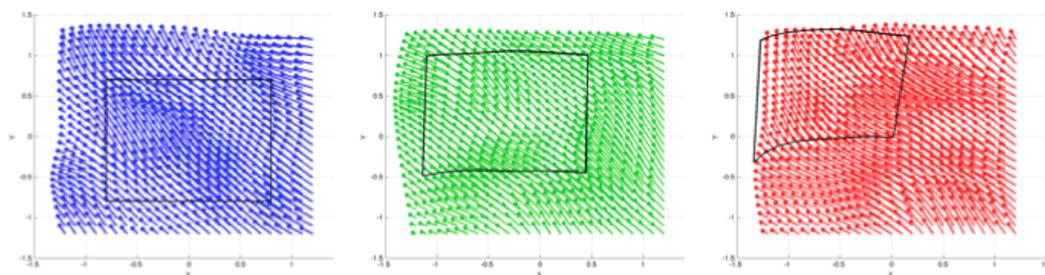


Figure 4.4: The deformation of the rectangle (black) is modeled by a flow of time-dependent \mathbb{R}^2 diffeomorphisms F_t , determined by a vector field flow v_t on \mathbb{R}^2 . The vector field flow is shown at three different times (blue, green, red).

Chapter 5

Examples of Diffeomorphic Point Matching

In this chapter, we briefly discuss the code implementation of the diffeomorphic point matching methods presented in Chapter 4. Then we present some examples to illustrate these techniques. We follow the same notation introduced there. All computations were performed on a Macbook Pro with an 2.53 GHz Intel Core i5 CPU with 4 GB of 1067 MHz DDR3 RAM.

5.1 Diffeomorphic Point Matching Implementation

All numerical modeling was done in the MATLAB environment. The curve fitting toolbox was augmented to the core functionality to access the modeling methods established in Chapter 3 (smoothing spline representation of curves and surfaces).

The prototype class used for all diffeomorphic point matching studies is titled *DiffeomorphicPointMatcher.m*. This is a simulation class; it takes a reference 3d point set and constructs a flow of diffeomorphisms that transforms the point set into another provided target set(s). Its main properties are summarized below. Properties are marked (input) if they are required by the user instantiating an object of the class and (calc) if they are derived from the users input.

- (input) The reference 3d point set is the array 'S0' – N0 points, divided up into S0.int (interior) and S0.bd (boundary).
- (input) The target 3d point sets is the array 'S' – N points each, divided up into S().int (interior) and S().bd (boundary)
- (input) t – Frames at which geometry info is provided
- (input) tscale – Time scale (secs/frame)
- (input) sigma – Smoothing parameter in Gaussian Kernel
- (input) theta – Termination criterion in inner loop
- (input) gamma – Multiplicative criterion in outer loop
- (input) lambda – Array of weights for each disparity term in the cost functional
- (input) THR – Termination criterion in outer loop
- (input) DispG – Function handle to geometry disparity term
- (input) maxnu – Max number of iterations allowed for outer loop

- (input) maxmu – Max number of iterations allowed for inner loop
- (input) gsf – Gradient scale factor in gradient descent (not armijo)
- (input) aalph0 – Armijo initial step length
- (input) armit – Max number of armijo iterations in gradient descent
- (input) armbeta – Armijo parameter
- (calc) N0 – Number of points on reference set
- (calc) N – Number of points on target sets
- (calc) q – Number of times at which geometry is given
- (calc) tau – Total times at which positions of deformation trajectories are sought
- (calc) L – Cardinality of the discretized time tau
- (calc) tauindex – Indexes tau (for sequencing, array calls, etc)
- (calc) tmap – Maps between tau and tauindex for easy spotchecks
- (calc) dtauF – Array of time intervals between each point in tau
- (calc) dtauS – dtauF converted to seconds via tscale (input)
- (calc) bdidx – Boundary index for trajectories
- The history of the transformation is stored in:

- x – trajectories of points
- α – “coefficients” of vector field expansion with RKHS structure
- p – adjoint state, “momentum” vectors
- \ker – kernel $A(x_l)$ for current iteration
- \kergrad – gradient of $A(x_l)\alpha_l$ for current iteration
- KI – value of energy term for current iteration
- $valDG$ – value of geometry disparities for current iteration
- ctr – Iteration counter

The main method of the class (after required properties have been set) is the *run* method. It performs the aforementioned continuation method in Chapter 4. Specifically, it runs “outer” loops that slowly increase λ until a desired geometric fit is achieved. For each outer loop iteration, λ is fixed and an “inner” loop Armijo gradient descent minimizes the cost functional. The method terminates when either acceptable geometric convergence has occurred (driven by THR), if the maximum number of iterations is reached (determined by $maxmu$ and $maxnu$), or if the armijo algorithm fails (determined by $aalph0$, $admit$, and $armbeta$). We now explore a few examples.

5.2 Synthetic Construction of \mathbb{R}^3 vector field flow

Here we opt to directly construct an \mathbb{R}^3 vector field flow. Doing so allows us to check a) how a simple shape embedded in the space deforms with time and b) how well

the recovered diffeomorphic flow matches the known flow.

With $X = [-1, 1]$, define the mesh

$$X_h = \{x_j : x_j = -1 + j/h, 0 \leq j \leq 2h\}$$

and define the kernel nodal points by

$$N_K = X_2 \times X_2 \times \{-0.5\}$$

Adopting the SRKHS mentioned in Chapter 4 (radial Gaussian kernel K_σ with $\sigma = 0.3$), we select our $\alpha_{xy} : [0, 1] \rightarrow \mathbb{R}^3$ by

$$\alpha_{xy}(t) = \begin{pmatrix} 0.05 \\ 0.05 \\ 0.1 \end{pmatrix}$$

We choose the reference surface S_0 to be a point grid sampling of the lower half of the unit sphere $S^{3,-}$

$$S_0 = [X_{10} \times X_{10} \times X_{10}] \cap S^{3,-}$$

When accounting for the boundary of $S^{3,-}$, we have $N_0 = 331$ points $x_i(0)$ on S_0 . See Figure [5.1].

Obviously, any point on the trajectory $x_j(t)$ at time $t \in [0, 1]$ is subjected to a

velocity governed by

$$v_t(x_j(t)) = \sum_{n=1}^{N_0} K_\sigma(x_j(t), x_n(t)) \alpha_n(t)$$

Thus each of the N_0 trajectories can be evolved forward in time with the forward Euler approximation

$$x_j(t_{i+1}) = x_j(t_i) + v_t(x_j(t_i))[t_{i+1} - t_i]$$

Taking the time partition $\{0, 0.25, 0.5, 0.75, 1\}$ of $[0, 1]$, we find at time $t = 1$ the evolved trajectories can be viewed as a point grid representation of a surface S_1 . Anticipating the future smooth Hausdorff matching term in diffeomorphic matching, a much finer grid can be evolved to time $t = 1$, which results in a denser point grid representation of the target surface with $N > N_0$ points. For this example $N = 1354$. See Figure [5.2].

With the reference and target shapes S_0, S_1 defined, *DiffeomorphicPointMatcher.m* can be applied. Taking our time partition $\Delta_I = \{0, 0.25, 0.5, 0.75, 1\}$, we find that with an initial guess of $\lambda = 6000$, the continuation algorithm runs for 3 outer iterations (3-8 inner iterations each) and produces a result such that 90% of the trajectories are within 0.03 of the target S_1 (well within the tolerance of our sampled target surface). The total runtime for this example is 30 seconds. See Figure [5.3]. The choice of λ may seem arbitrary but it is often determined through quick successive simulation attempts. In particular, an experienced human-in-the-loop can easily find “via bisection” a good starting value of λ (i.e. one that will actually get trajectories

to move) in the order of a few minutes; after which, the continuation algorithm is able to take over and proceed to finely increase λ until a satisfactory match is achieved.

At this point, the question remains: how well does the recovered deformation match the actual deformation? The plot in Figure [5.4] gives an insight into the question. We notice that the trajectories recovered via diffeomorphic point matching, although matching the target endpoints, can deviate from the actual vector field flow at intermediary times. This is not all that surprising, as the cost functional J only accounts for vector field energy and geometric disparity at the ending time. Moreover, the numerical scheme terminates once “acceptable” geometric performance has been achieved; further iterations would reduce the extra kinetic energy and trajectories would begin to “straighten out”. However, for general vector field flows v_t , there is *no* reason to expect the diffeomorphic matching solution to match. The latter tends to be smooth and “even” in time (reduce kinetic energy) whereas general flows can be impulsive and possibly nonuniform in time. This discussion will be revisited in subsequent chapters once we add new terms to the cost function J to help reduce this phenomenon.

5.3 Mitral Valve Leaflet (One Target)

An example we briefly mention now, but will explore extensively in Chapter 8, involves the matching of a mitral valve leaflet model (Chapter 3) at two different instants (namely, mid-systole MS and end-systole ES). In this example, we have $N_0 = 259$, $N = 1189$, $\Delta_I = \{1, 2, 3, 4\}$, and an initial $\lambda = 5 \times 10^5$. Again, this choice

of λ comes from experience with working with the tools and models (more will be said in Chapter 8). The kernel parameter σ was chosen to be half of the smallest distance between all points on the sampled reference leaflet. For this choice of N_0 , $\sigma = 3.3$. With this particular choice of parameters, the continuation algorithm runs for two outer iterations (3-4 inner iterations each for a total of 7 iterations) and produces a result such that 90% of the trajectories are within 0.7 mm of the target. This is well within the acceptable bounds given the ≈ 1 mm uncertainties associated with the tagged points from which we generated the models of leaflet surfaces at MS and ES. The total runtime for this example is 8 seconds. See Figure [5.5].

As for the scaling of runtime with the the number of reference points N_0 , the results can be found in Table [5.1]. Note a fixed number of 10 iterations is performed. We see that the compute time is definitely non-linear in the number of points N_0 on the reference shape. This is not too surprising given that operations with the kernel involve $3N_0 \times 3N_0$ matrices. But still ≈ 2 minutes to evaluate a 1000 point-sampled shape match is acceptable. Also, with this method, as long as $N_0 \approx 1000$, the methods can used to match larger shapes with larger mesh sizes or on smaller shapes with finer mesh sizes and one can still expect a runtime of about 2 minutes.

5.4 Mitral Valve Leaflet (Two Targets)

In the data base of tagged mitral valve echocardiographies prepared by our collaborators at by Methodist Hospital Cardiology, only a few patient echocardiographies were tagged at begin-systole (BS) as well. These cases provide perfect opportunities to

N_0	σ	Runtime (s)
97	4.85	4.4
259	3.30	6.1
385	3.00	11.2
540	2.55	28.9
724	2.05	78.2
940	1.95	140.1

Table 5.1: The runtime for *DiffeomorphicPointMatcher.run()* with 10 iterations for different values of N_0 . Notice that the computing time is definitely non-linear with N_0 . This is not too surprising given the kernel operations involve $3N_0 \times 3N_0$ matrices.

test the diffeomorphic point matching algorithm in the case of multiple target shapes (see last section of Chapter 4). Specifically, with $N_0 = 264$, $N_1, N_2 \approx N = 700$, $\Delta_I = \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10\}$, and an initial guess of $\lambda = 5 \times 10^5$, the algorithm runs for 2 outer iterations (10-13 inner iterations each) and produces a result such that 90% of the trajectories are within 0.9 mm of the targets. The total runtime is 2 minutes. See Figure [5.6].

It should be briefly mentioned that with the current framework laid out in Chapter 4, diffeomorphic matching with multiple target snapshots yields trajectories that are not smooth. Particularly, there are corners at the times $\tau^{L_i} = t_i$ where targets are provided. This is not surprising, as the vector field flow was not constrained to be 2nd-order smooth. This problem is avoidable through adequate cost functional adjustment. Aarti Jajoo has done some work in this area with successful results [24]. However, for the remainder of the bio-medical applications to mitral valves considered in this thesis, these corners do not present a problem, as all shape matching consists of one target (valve at end-systole) and one reference (valve at mid-systole).

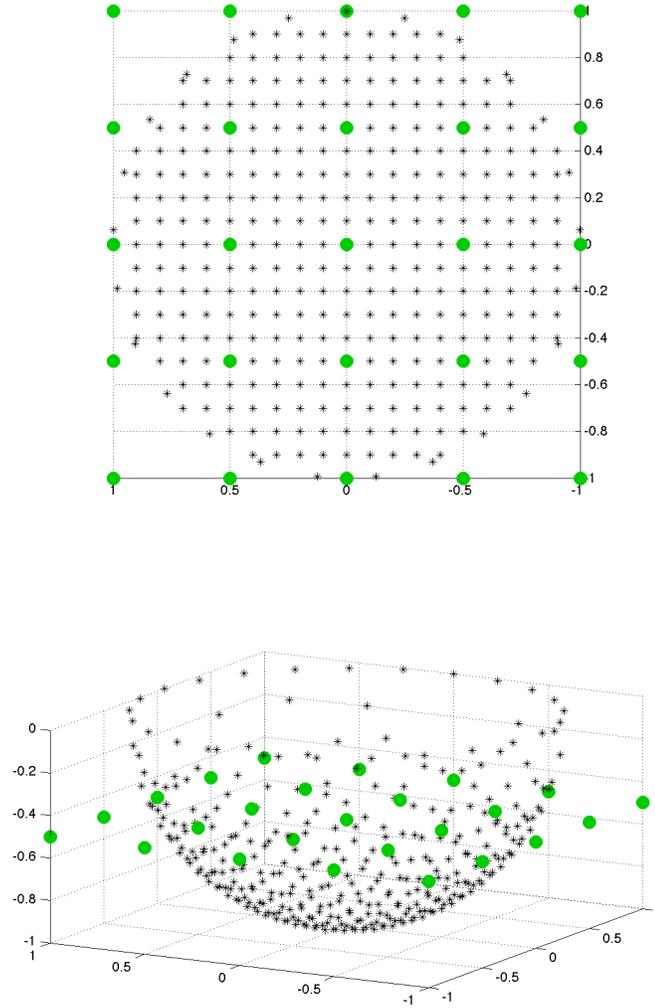


Figure 5.1: The reference shape S_0 (black) and the kernel nodal points N_K (green) shown from two different perspectives.

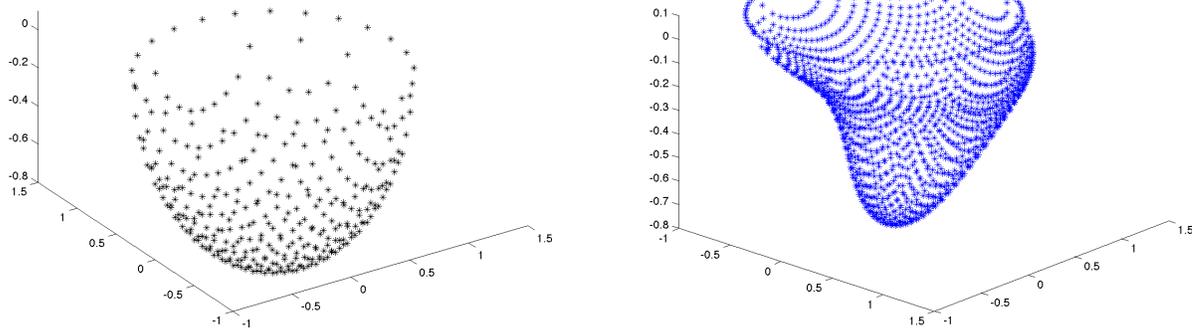


Figure 5.2: The reference shape S_0 (left) and target shape S_1 (right).

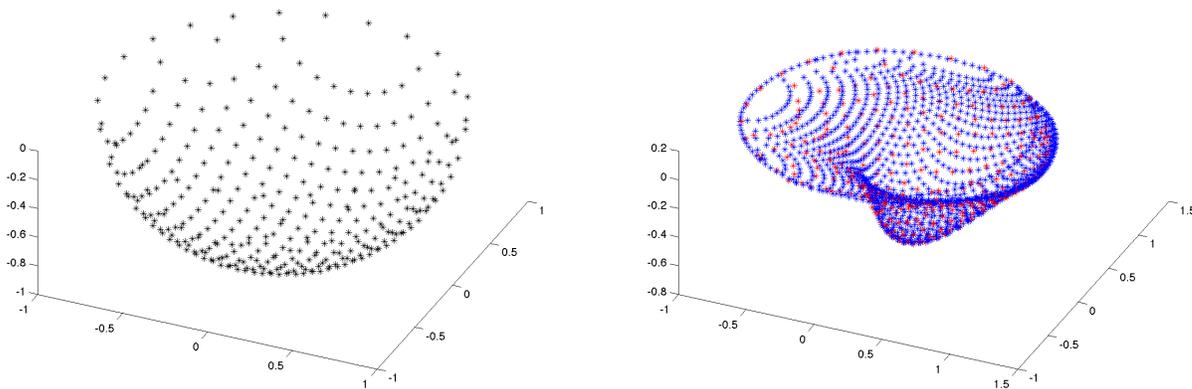


Figure 5.3: The bottom unit sphere is evolved under the action of a vector field flow. The reference shape is discretized to a point grid (black) and deformed to match the target shape. The deformed point grid recovered from diffeomorphic matching (red) coincides well with the target surface (blue).

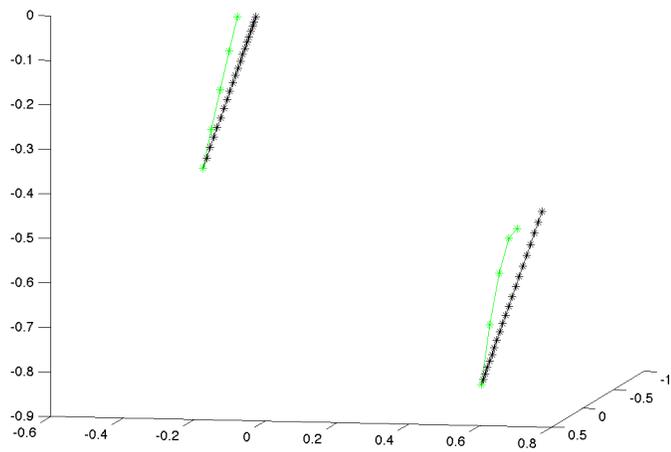


Figure 5.4: Comparison of recovered deformation trajectories (green) with actual trajectories (black). Only two trajectories shown for ease of viewing. The recovered trajectories, although matching well the beginning and ending points of the actual trajectories, can deviate in different manners at intermediate times. The distances between the actual and recovered trajectories never exceed 0.05 units.

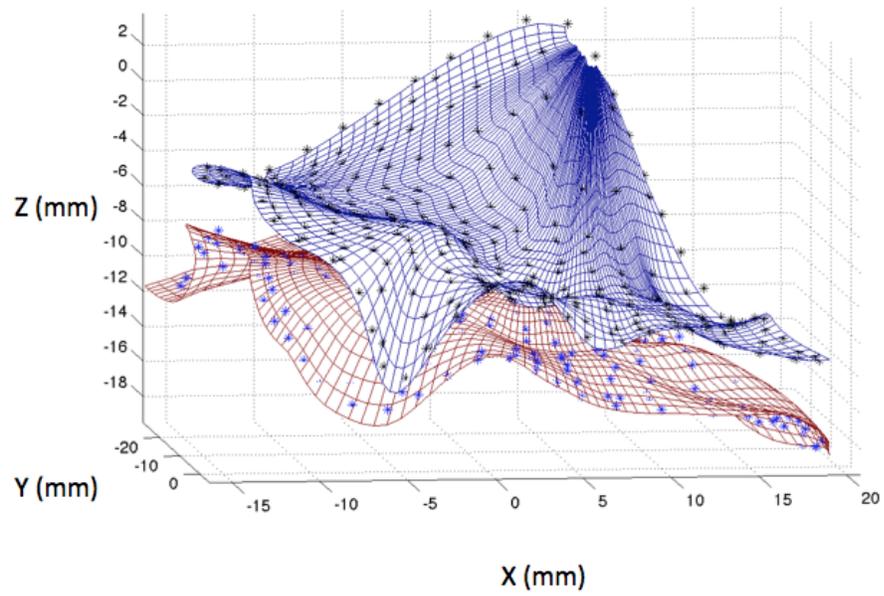


Figure 5.5: The top surface (black points/blue mesh) is deformed into the bottom surface (blue points/red mesh).

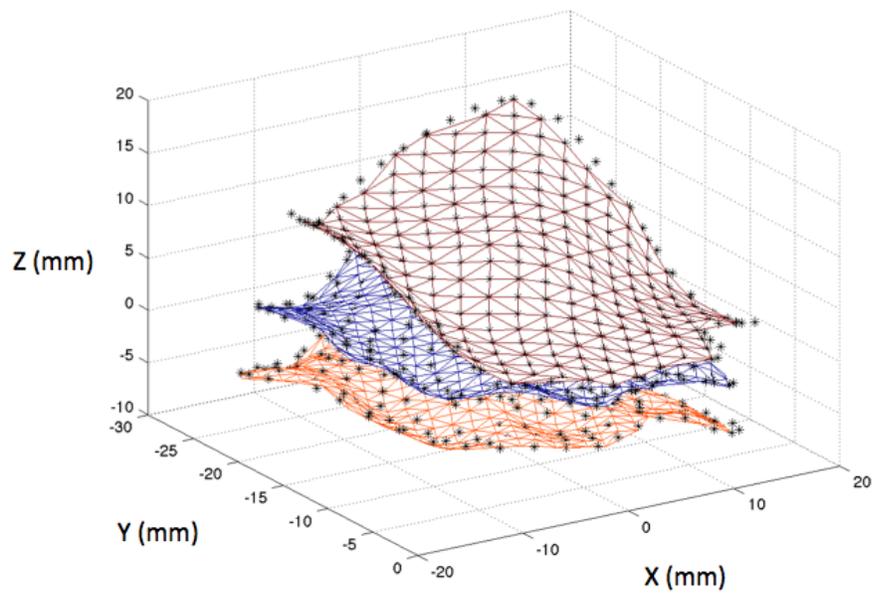


Figure 5.6: The reference leaflet at begin-systole (red) is deformed throughout systole. At times mid-systole (blue) and end-systole (orange), the deformed trajectories (black) match quite well (< 1 mm disparity).

Chapter 6

Image Registration

6.1 Overview

The goal of image registration is to determine an optimal spatial transformation or deformation that will *register* (i.e bring into correspondence) pairs of homologous points in two given images [38]. As with shape tracking discussed in Chapter 4, the mathematical form of the transformation may vary from simple to complex based on a set of physical assumptions. For example, when registering images of a static body, the spatial transformation is rigid and can be constructed by applying successively a translation, a rotation, and a possible dilation. At the other extreme, registration of images of soft organs often requires nonlinear effects that violate rigid body assumptions. Moreover, images acquired from different subjects further complicate matters, as developmental factors including genetics, environment, and random influences all contribute to the complex shape differences inherent to the images. For an excellent

survey of modern image registration techniques, see [16],[34].

6.2 Elements of Image Registration

Existing registration algorithms can be classified by three elements: the feature space, the deformation model and the similarity metric. We briefly explore each of these, noting how each ties in with the established shape tracking methodology given in Chapter 4.

6.2.1 Feature space

Before any registration problem can begin, data must be extracted from the two images to be registered. The extracted image data, whether they are pixel-based or feature-based, belong to a set we call the *feature space*.

Pixel-based data stems directly from the pixel intensity values of the images being registered. Preprocessing is often necessary to suppress the adverse effects of noise and possible differences in acquisition techniques. While it is possible to work directly with pixel values on a discrete coordinate grid (machine limited), it is often advantageous to cast the registration problem into a continuous framework. Specifically, the images are considered as continuous functions of the pixel coordinates. The correspondence between the discrete and continuous versions of the image is established by interpolation. Interpolation methods vary widely in their fidelity and ease of computation, ranging from nearest-neighbor to linear to cubic spline interpolation

(see Chapter 2).

Feature-based data derive from a set of characteristic features extracted from the images. The dimensionality of the features is usually smaller than the dimensionality of the original image data, resulting in a simpler problem description. However, the data extraction process is highly non-linear and often hard to automate. For instance, the static modeling of the mitral valve (see Chapter 3) requires the hand-selection of landmark points from 3D ultrasound images by experienced cardiology experts. The image registration of this type of data by geometric diffeomorphic matching (Chapter 4) accounts for the matching of these features.

6.2.2 Deformation Model

The *deformation model* used in image registration determines how one accomplishes the correspondence between the two images. The methods summarized for shape tracking in Chapter 4 are applicable here. For example, an image registration problem can be cast as a sequence of shape tracking problems, each of which can be approached via diffeomorphic shape tracking outlined in Chapter 4. Alternatively, rigid deformations can be used to track the landmark points which then allows the use of global interpolation methods like thin-plate splines to determine the local elastic deformation. This is discussed in the SURE algorithm for 3D echocardiography below [27].

6.2.3 Similarity metrics

The *similarity metric* determines the quality of registration; in other words, how “good” the correspondence is established by the registration. Consider two images J_{ref}, J_{tar} . Denote the deformed reference image by $F(J_{ref})$. For pixel-based data, a common similarity metric is the pixel difference in l_1 or l_2 norms, that is

$$\|F(J_{ref}) - J_{tar}\|_1 = \sum_{\text{pixels } x} |F(J_{ref}(x)) - J_{tar}(x)|$$

$$\|F(J_{ref}) - J_{tar}\|_2 = \sum_{\text{pixels } x} |F(J_{ref}(x)) - J_{tar}(x)|^2$$

Other similarity metrics are the correlation coefficient

$$r = \frac{\sum_{\text{pixels } x} [F(J_{ref}(x)) - \overline{F(J_{ref}(x))}] [J_{tar}(x) - \overline{J_{tar}(x)}]}{\sqrt{\left(\sum_{\text{pixels } x} [F(J_{ref}(x)) - \overline{F(J_{ref}(x))}]^2\right) \left(\sum_{\text{pixels } x} [J_{tar}(x) - \overline{J_{tar}(x)}]^2\right)}}$$

or the mutual information [27]. For feature-based data, a satisfactory choice is the mean distance between corresponding landmarks in the deformed reference image J_{ref} and target image J_{tar} . When the feature-based data are interpolated onto point grids and the reference points and target points are no longer in a 1-1 correspondence, the Hausdorff matching term (Chapter 4) is used.

6.3 Image Registration for 3D Echocardiography

Image registration for 3D Echocardiography cannot be discussed without mentioning prior studies in that area. Shekhar et al [43] have devised a voxel-based scheme that allows for the registration of 3D ultrasound image volumes provided by 3D echocardiography. The algorithm does so by global rigid transformations (translations, rotations, dilations). This scheme effectively accounts for varying sonography angles during ECG acquisition and facilitates accurate diagnosis during standard echocardiographic stress testing. See Figure [6.1] for an example.

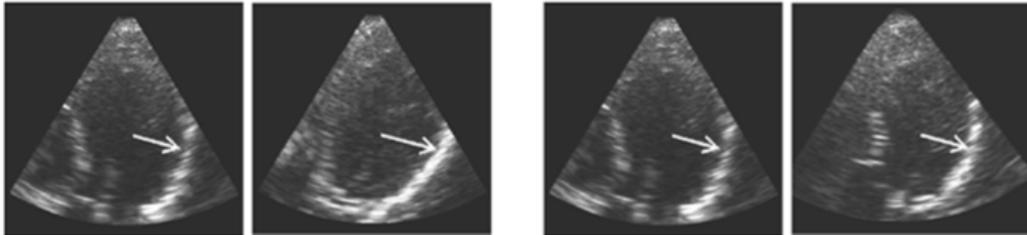


Figure 6.1: Side-by-side prestress and poststress long-axis views before (left column) and after (right column) registration. White arrows point to the free wall for the left-ventricle. The marked wall segments in the poststress views are tilted and shifted with respect to the same segments in the prestress views before registration. Note that these segments specifically and the entire left-ventricle in general are better aligned following registration. (Caption and figure from [43])

A more general elastic approach can be found with algorithm called SURE, which was introduced by Krucker [27]. SURE, which stands for Subvolume-based algorithm for elastic Ultrasound REgistration, is a divide-and-conquer method of registering 3D ultrasound images. It works on the principle that global elastic deformations can be achieved by a superposition of local rigid deformations. It functions in the following four steps:

- (a) Subvolume division: Divide image volume into a grid of subvolumes.
- (b) Subvolume selection: Compute the volumetric overlap, entropy, and pixel similarity measure of each subvolume and discard subvolumes that do not satisfy the minimum criteria for each of these variables. The X in Figure [6.2](b) indicates one such subvolume.
- (c) Subvolume registration: For each remaining subvolume, find the translation vector that maximizes the pixel similarity measure within a search window.
- (d) Global interpolation: The original and translated center coordinates of the remaining subvolumes define two sets of control points, which are used to compute the thin-plate spline transformation from the original to the new homologous volume.

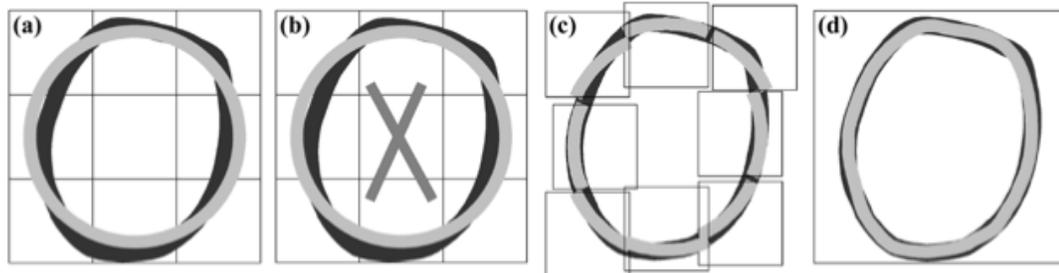


Figure 6.2: Illustration of one iteration of the SURE algorithm registering the light gray circle (homologous volume) onto the darker shape (reference). The images (a)-(d) correspond to the four steps described above in the text. (Caption and figure from [27]).

Both of these methods are for the global (large volume) registration of ultrasound images. But for registration of the mitral valve, which involves a specific, small region of 3D Echocardiography image volume which is difficult to identify accurately

by classical segmentation algorithms, these large-scale techniques prove to be too coarse to be of immediate use. Moreover, as we saw in Chapters 3, cardiologists and mathematicians have already allowed for the mitral valve shape to be described *independently* of the underlying image sequence.

Chapter 7

Combining Diffeomorphic Shape Matching with Image Registration

In many applications, the shapes matched by diffeomorphic deformations are extracted from sequences of images. These images often contain additional information which may be used when registering the shapes. A good voxel-to-voxel transformation of shapes included in successive 3D-images acquired at high frame rates should approximately preserve voxel intensity values to increase confidence that accurate shape deformations are actually recovered. A natural starting point for an implementation of this technique is to add an intensity matching cost to the geometric matching cost. We propose an equivalent formulation and solution of this problem involving an iterative scheme which alternates between geometric matching and image registration. We also study several concrete examples in details to validate the performances of our approach.

7.1 Geometric Matching: The need for voxel intensity data

As we saw in Chapters 4 and 5, geometric matching provides an excellent method for the matching of smooth 3D shapes. Incorporating the language of chapter 6, the geometric matching of static mitral valve models is a feature-based method of image registration of 3D ultrasound images.

Consider the last example given in Chapter 5. Here an anterior leaflet model, after being statically modeled at the three instants BS, MS , ES, is dynamically registered by a diffeomorphic flow matching these three static surfaces. However, the plot in Figure [7.1] gives the intensity vs. time of 5 specific voxel trajectories located throughout the anterior leaflet.

This plot shows that after purely geometric matching, voxel trajectory intensities can vary wildly, exhibiting multiple oscillations and large ranges. For this reason, we seek to linearly combine the geometric matching cost function with a term incorporating voxel intensity data provided by the 3D ultrasound images in order to control these intensity fluctuations during the diffeomorphic tracking.

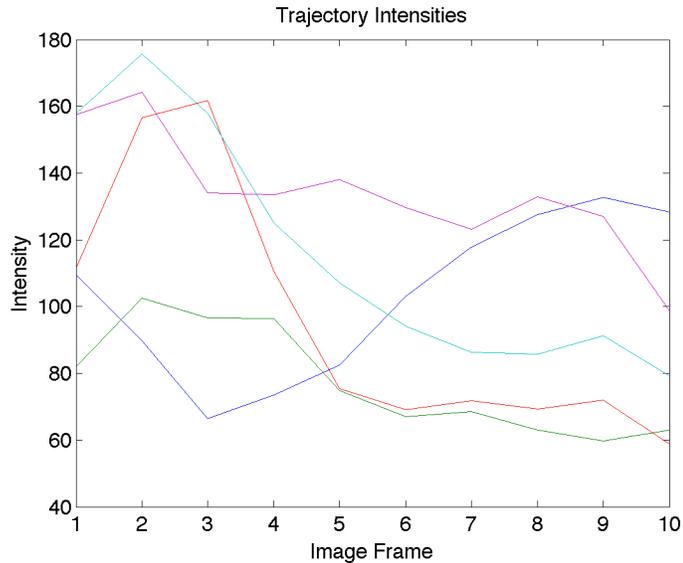


Figure 7.1: Voxel intensities along 5 specific voxel deformation trajectories for a geometric matching case. Notice the large ranges of intensities, as well as the large oscillations present on some of the trajectories. This is not that surprising, as geometric matching, being a feature-based method of registration, does not consider the voxel intensities information present in the ultrasound movie when the diffeomorphism between two 3D-shapes is computed.

7.2 Geometric Matching + Intensity Registration

7.2.1 Background

Recall that in the typical geometric matching of shapes $S_0, S_1 \in SH_d$, the cost functional consists of an “energy term” $E(v)$ proportional to the square of the L^2 norm of an \mathbb{R}^3 vector field flow v_t and a “disparity” term that quantifies the amount of geometric error in the matching (for instance, Hausdorff matching D_h). Specifically,

with the notation already introduced in Chapter 1, we solve the minimization problem

$$\inf_{v \in L^2([t_0, t_1], V)} E(v) + \lambda D_h(F_{t_1}(S_0), S_1), \quad \lambda > 0$$

subject to the constraint

$$\begin{aligned} \frac{\partial F_t}{\partial t} &= v_t \circ F_t, & t \in (t_0, t_1] \\ F_{t_0} &= Id \end{aligned}$$

in order to compute an approximate solution v^* .

But often the shapes $S_0, S_1 \in SH_d$ are derived from extracted from a sequence of 3D images. Thus at intermediary times $t_0 < t < t_1$, there is potential information (namely, the flow of intensities) to exploit, and this information is not used by purely geometric diffeomorphic matching. For this reason, a new disparity term will be introduced below. But first let us define some notations.

Consider a 3D image sequence IS . This image sequence involves successive 3D images J_i of the bounded region $\Omega \subset \mathbb{R}^3$ captured at image frame $i \in \mathbb{N}$. Thus $IS = \{J_0, J_1, \dots\}$. We define the intensity mappings $I_t : \Omega \rightarrow \mathbb{R}$ so that $I_t(x)$ is the intensity of the image J_t at voxel $x \in \Omega$. Typical values of I_t are 8-bit; namely $\text{rng}(I_t) \subset \{0, 1, \dots, 255\}$. Moreover, the definition of I_t generalizes to multi-channel (color) images (unique I_t for each channel). In reality, intensities I_t are not continuous in space. There is always a finite resolution determined by the number of voxels in the image J_t . However, spatial interpolation of the values $I_t(x)$ involves assuming the existence of a piecewise continuous intensity function $I_t \in C(\Omega)$.

Now consider an image sequence $IS = \{J_0, \dots, J_Q\}$ from which two shapes $S_0, S_1 \in SH_k$ ($k \in \{1, 2, 3\}$) are extracted. Let S_0 come from image J_0 and S_1 come from J_Q . Consider a point grid of size N_0 discretizing S_0 and the geometric matching of shapes S_0, S_1 (see Chapter 4). Along each one, the N_0 voxel intensities are going to vary between successive image frames. We define the intensity disparity between frames $i - 1$ and i by the sum of squared differences (SSD)

$$D_{\mathcal{I}}^i = \frac{1}{N_0} \sum_{j=1}^{N_0} [I_i(F_{T_i}(x_j)) - I_{i-1}(F_{T_{i-1}}(x_j))]^2$$

This term seeks to enforce small relative changes in intensity along each trajectory. Moreover, this function is a well defined disparity functional provided $I_i \in C(\Omega)$. Thus we now introduce a new cost function to formulate the *geometric matching + intensity registration* problem

$$\begin{aligned} \inf_{v \in L^2([t_0, t_1], V)} E(v) + \lambda D_h(F_{t_1}(S_0), S_1) + \mu \sum_{i=1}^Q D_{\mathcal{I}}^i, \quad \lambda, \mu > 0 \\ \text{subject to} \\ \frac{\partial F_t}{\partial t} = v_t \circ F_t, \quad t \in (t_0, t_1] \\ F_{t_0} = Id \end{aligned} \tag{7.1}$$

By satisfactory weighting via a parameter μ between the geometry term and this term, we seek to improve the quality of the deformation reconstruction from the image data. However, determining the two regularization parameters λ, μ that provide a good balance between geometric fitting and image optical flow proves

challenging. A generalization of the continuation method has been explored, with little success to this point. We now outline our alternative approach. In the case of geometric matching + image registration, it is far simpler if the registration happens *not concurrently* but *after* the geometric matching. Namely, the trajectories of the computed geometric match are fixed at the times t_0, t_1 but perturbed at intermediary times in order to reduce the intensity disparity. We formalize this idea below.

7.2.2 Alternative Formulation

The geometric matching + intensity registration problem could alternatively be formulated as a two-step procedure. Namely, we introduce the following two vector field flows v^* and v^{**} verifying,

$$\begin{aligned}
 J_{geo}(v^*) &= \inf_{v \in L^2([t_0, t_1], V)} E(v) + \lambda D_h(F_{t_1}(S_0), S_1), & \lambda > 0 \\
 J_{int}(v^{**}) &= \inf_{v \in L^2([t_0, t_1], V)} \mu \sum_{i=1}^Q D_I^i, & \mu > 0 \\
 &\text{subject to the constraints} & (7.2) \\
 E(v^{**}) + \lambda D_h(F_{t_1}^{**}(S_0), S_1) &\leq J_{geo}(v^*) + \epsilon, & \epsilon > 0
 \end{aligned}$$

Solutions v^{**} to (2) are not guaranteed to be the same as solutions v to (1), but both intuitively capture the goal of balancing geometric matching and image registration.

The second formulation has a distinct advantage in the possible approaches that

are immediately viable. Specifically, the $\alpha(t)$ for the vector field v^* are already initialized to “good” values; we expect small perturbations of the $\alpha^*(t)$ can result in a solution v^{**} whose intensity disparities will be lower than that for v^* . Intuitively, perturbing the α^* is a difficult matter. After all, in the language of optimization and optimal control theory, the $\alpha(t)$ are the *controls* and the $x(t)$ (physical trajectories) are the observable *states*. Our understanding of geometric matching involves the states and the derived constraint in (2) is inherently a state space constraint. Moreover, the states $x(t)$ directly determine the intensity disparities $D_{\mathcal{I}}^i$. Intuitively, then, it is ideal to work with perturbations $\hat{x}^*(t)$ of the states $x^*(t)$ that arise from geometric matching. The velocities \hat{v}^* give rise, via inversion of the kernel, to controls $\hat{\alpha}^*(t)$. These controls can then be supplied as inputs to another geometric matching problem. Thus we see an iterative scheme developing, one that alternates between a diffeomorphic point matching problem (geometric matching) and a perturbation problem (intensity registration).

Chapter 8

Numerical Methods for Intensity Registration

8.1 Intensity Registration Loop

Recall the *intensity registration* problem is the sub-problem to (7.2) given by

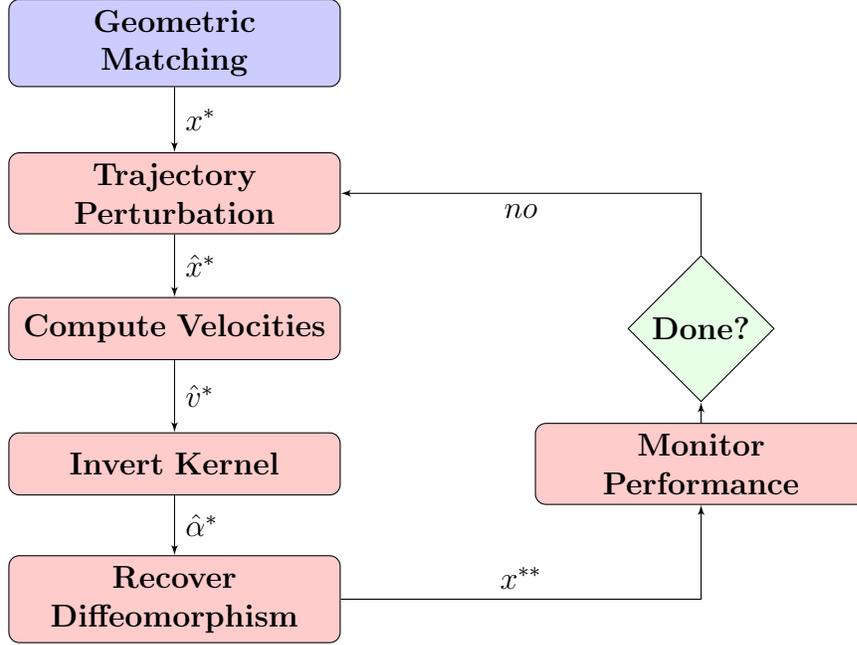
$$J_{int}(v^{**}) = \inf_{v \in L^2([t_0, t_1], V)} \mu \sum_{i=1}^Q D_{\mathcal{I}}^i, \quad \mu > 0$$

subject to

$$E(v^{**}) + \lambda D_h(F_{t_1}^{**}(S_0), S_1) \leq J_{geo}(v^*) + \epsilon, \quad \epsilon > 0$$

Below, we present a computationally efficient algorithm that offers an approximate solution to the problem while preserving the properties of the already-computed

minimum energy geometric matching solution v^* in (7.2). Our algorithm is called the *intensity registration loop* (IRL) and we explore its components in the workflow below.



8.1.1 Geometric Matching

Trajectories $x^*(t)$ of the geometric diffeomorphic matching (Chapter 4) problem are computed for two shapes $S_0, S_1 \in SH_d$ with $d \in \{1, 2, 3\}$. These shapes are assumed to be extracted from an image sequence $\{J_i\}$ consisting of $Q + 1$ 3D image frames. Recall from Chapter 4 that geometric diffeomorphic matching involves a discretization of the reference shape into a $N_0 \in \mathbb{N}$ point grid and in time defined by the time grid $\Delta_I = \{t_0 = \tau^{L_0} < \tau^{L_0+1} < \dots < \tau^{L_1-1} < \tau^{L_1} = t_1\}$. Denote by T_i the times in Δ_I corresponding to the image frames J_i with $0 \leq i \leq Q$. Obviously $T_0 = \tau^{L_0}$,

$T_Q = \tau^{L_1}$, and $\text{card}(\Delta_I) \geq Q + 1$.

8.1.2 Trajectory Perturbation

The perturbation of trajectories $x^*(t)$ is handled in the following manner:

FOR all j with $1 \leq j \leq N_0$

1. Initialize new perturbed trajectory $\hat{x}_j^*(t)$ at *known* starting and ending points.

1a) Set $\hat{x}_j^*(T_0) = x_j^*(T_0)$.

1b) Set $\hat{x}_j^*(T_Q) = x_j^*(T_Q)$.

1c) Calculate $I_{start} = I_0(\hat{x}_j^*(T_0))$ and $I_{end} = I_Q(\hat{x}_j^*(T_Q))$.

2. **FOR** all T_i with $1 \leq i \leq Q - 1$

2a) Compute intensities $I_i = I_i(x_j^*(T_i))$.

2b) Calculate *desired* intensity $\hat{I}_i = f(i, I_{start}, I_{end})$. Here f is a user-defined function of the discrete time index i . The simplest choice for f is the linear interpolation

$$\hat{I}_i = f(i, I_{start}, I_{end}) = I_{start} + (I_{end} - I_{start}) \frac{i}{Q}$$

Notice that $f(0, I_{start}, I_{end}) = I_{start}$ and $f(Q, I_{start}, I_{end}) = I_{end}$.

2c) In the box (of width 2ϵ) B_ϵ around $x_j^*(T_i)$, find the point \tilde{y} whose intensity is closest to \hat{I}_i (that is, $\text{argmin}_{y \in B_\epsilon} |\hat{I}_i - I_i(y)|$). Define $\hat{x}_j^*(T_i) = \tilde{y}$.

3. The new trajectory $\hat{x}_j^*(t)$ has now been defined at the frame times T_i . The positions $\hat{x}_j^*(t)$ at intermediary times between these T_i can then be computed by simple interpolation in \mathbb{R}^3 .

Though the algorithm seems complicated in pseudo-code, it is actually straightforward in principle. The voxel trajectory $\hat{x}_j^*(t)$ within a thin tube around a given voxel trajectory $x_j^*(t)$ (generated by diffeomorphic flow) are constructed so that intensity changes smoothly along $\hat{x}_j^*(t)$, with all perturbations of trajectories $x_j^*(t)$ being implemented at image sequence frame times $t = T_i$. See Figure [8.1].

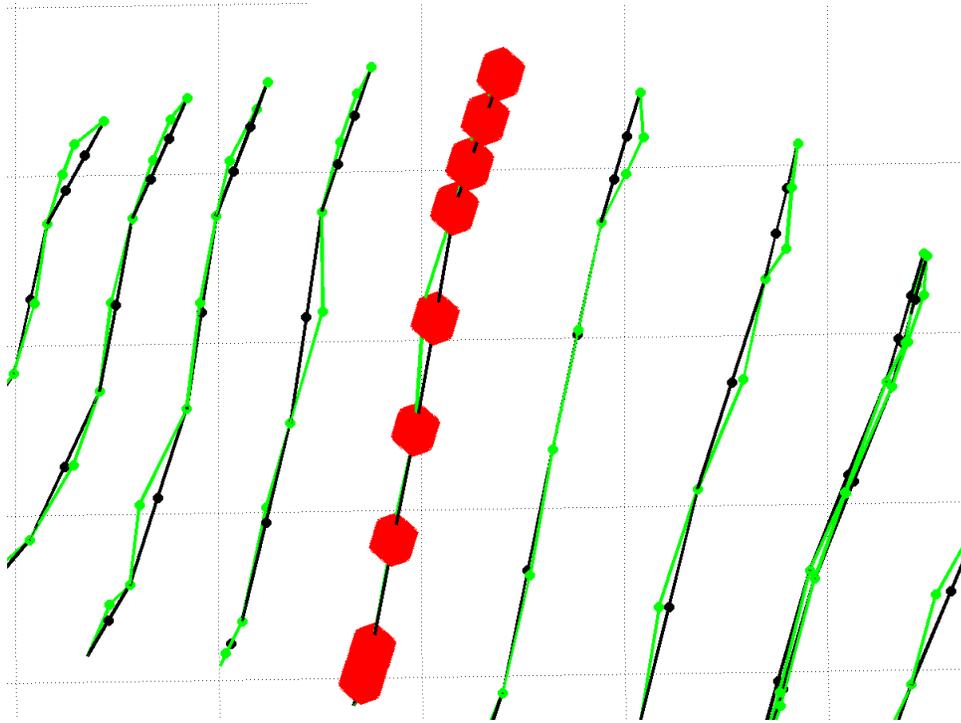


Figure 8.1: Perturbed trajectories (green) of geometric trajectories (black). The boxes B_ϵ (red) are displayed on a single trajectory to illustrate the procedure.

For step 1b, we remark that sometimes an alternative approach is needed. Because intensity registration at intermediary frames T_i is completely dependent upon the intensities recorded at the initial and final times, large mismatches in intensities will demand large corrections and prevent the intensity registration quality (see Performance Monitoring section) from decreasing. A good remedy is to allow the final node the freedom to move, transforming the early stages of trajectory perturbation to a box search problem (analogous to that of step 2c). For instance, the perturbed trajectory at frame index Q is initialized so as to minimize the difference in intensity $|I_Q(\hat{x}_j^*(T_Q)) - I_{start}|$. Whenever the intensity mismatch between initial voxel $x_j^*(T_0)$ and final voxel $x_j^*(T_Q)$ is not an issue, it is simplest to proceed by step 1b.

For step 2b, a simple choice (linear interpolation) is proposed for the intensity interpolating function $f(i, I_{start}, I_{end})$. Any polynomial interpolating function, either in true or least-squares sense, is also a valid option. The linear option for $f(i, a, b)$ is often the simplest approach and gives satisfactory results.

Throughout this algorithm, it should be noted that each perturbation $\hat{x}^*(t)$ of the initial voxel trajectories $x^*(t)$ generally requires more kinetic energy than the kinetic energy of the initial trajectories $x^*(t)$. Moreover, this increase in kinetic energy clearly forces larger values of ϵ in the box search. It becomes important to track the net increase in kinetic energy of the trajectories ΔK to make sure that trajectories do not wildly deviate from their approximately minimum energy state (see Performance Monitoring section).

8.1.3 Velocities, Kernel Inversion, and Diffeomorphism Recovery

With perturbed trajectories $\hat{x}(t)$ now computed at all times $t = \tau^l$, it remains to be seen if these trajectories do indeed result from a new flow of diffeomorphisms. For ease of notation, denote all quantities computed on this time grid Δ_I by superscript l (e.g. $\hat{x}^*(\tau^l) \equiv \hat{x}^{*l}$). Letting \hat{v}^{*l} represent the velocities of the perturbed trajectories (computed by $\hat{v}_j^{*l} = \frac{\Delta \hat{x}_j^{*l}}{\Delta \tau^l}$), we desire $\hat{\alpha}^{*l}$ so that

$$\hat{v}_j^{*l} = \sum_{i=1}^{N_0} K_\sigma(\hat{x}_j^{*l}, \hat{x}_i^{*l}) \hat{\alpha}_i^{*l}$$

or with the block notations introduced in Chapter 4

$$\hat{v}^{*l} = A^l \hat{\alpha}^{*l} \tag{8.1}$$

where $A^l = (A_{ij}^l)$ and $A_{ij}^l = K_\sigma(\hat{x}_j^{*l}, \hat{x}_i^{*l}) \mathbf{I}_3$ and \mathbf{I}_3 is the 3×3 identity matrix. Inverting the matrix A^l in (8.1) gives

$$\hat{\alpha}^{*l} = A_l^{-1} \hat{v}^{*l} \tag{8.2}$$

These computed values of $\hat{\alpha}$ are used to determine trajectories x^{**} that will be *close to* trajectories \hat{x}^* but are indeed generated by a time dependent vector field flow v^{**} of diffeomorphisms described by $\hat{\alpha}^*$.

8.1.4 Performance Monitoring

Once a single iteration of the Intensity Registration Loop (IRL) has been performed, we are in a position to repeat the process. Specifically, the trajectories x^{**} could be further perturbed in order to improve the quality of registration. As previously mentioned, we have to remain conscious of the increasing cost in kinetic energy as well as the possible deviations of our solutions from target shape S_1 . We quantify these concepts below. Assume $N_0 \in \mathbb{N}$ voxel trajectories $x(t)$ have just been computed by applying a finite number of successive IRL iterations.

Definition 8.1.4.1. (*Geometric Matching Quality*) For each voxel trajectory $x(t)$, the geometric matching quality (GMQ) is the distance

$$GMQ(x(t)) = d(x(T_Q), S_1)$$

Here $d(x, Y)$ denotes the standard Euclidean distance from a point x to a set Y .

Definition 8.1.4.2. (*Intensity Registration Quality*) For each voxel trajectory $x(t)$, the intensity registration quality (IRQ) is the mean oscillation in intensity for that trajectory. That is,

$$IRQ(x(t)) = \frac{1}{Q} \sum_{i=0}^{Q-1} |I_{i+1}(x(T_{i+1})) - I_i(x(T_i))|$$

This is precisely the average absolute change in voxel intensity along each trajectory $x(t)$.

Definition 8.1.4.3. (*Kinetic Energy Quality*) For each voxel trajectory $x(t)$, the

kinetic energy quality (KEQ) is the kinetic energy of the trajectory computed on the time grid Δ_I . That is

$$KEQ(x(t)) = \sum_{l=0}^{\text{card}(\Delta_I)} \left[\frac{x^{l+1} - x^l}{\tau^{l+1} - \tau^l} \right]^2$$

Note that this quantity is the approximate kinetic energy of an individual trajectory, while the preceding kinetic energy $E(v)$ concerned a whole vector field flow v .

These three criteria can be monitored for all N_0 trajectories throughout the iterations of the IRL. Statistical plots (quantile curves, histograms) as well as 2D or 3D plots displaying simultaneously history 2 or 3 performance criteria can be used to facilitate performance monitoring for the iterated IRL algorithm. See Figure [8.2] for examples. These monitoring data come from the example given below in section 4.2. The 2D displays allow us to visualize how the initial voxel trajectories generated by the diffeomorphic flow respond to the IRL. Specifically, a history of GMQ, IRQ, and KEQ is stored for each successive IRL iteration. When 2 of the criteria, for instance IRQ and KEQ, are plotted in the (KEQ,IRQ) plane (as in Figure [8.2]), these histories trace out paths for each voxel trajectory processed with the IRL. In the simple example given in Figure [8.2], we see that 8 of the 9 displayed paths in the (KEQ,IRQ) plane behave as expected; these paths correspond to voxel trajectories that exhibit a slight increase in kinetic energy in order to also exhibit smoother image intensity profiles, thus fitting better the ultrasound image sequence data. Visually, this is seen in the (KEQ,IRQ) plane by paths that slope downward to the right.

However, this example also demonstrates that the successive perturbations of a given voxel trajectory generated by iterated IRL may simultaneously increase in kinetic energy and degrade in registration quality (slopes upward to right in (KEQ,IRQ) plane); this can be explained by the kernel inversion in equation (8.2) in the IRL that slightly alters the perturbed trajectories constructed to reduce IRQ.

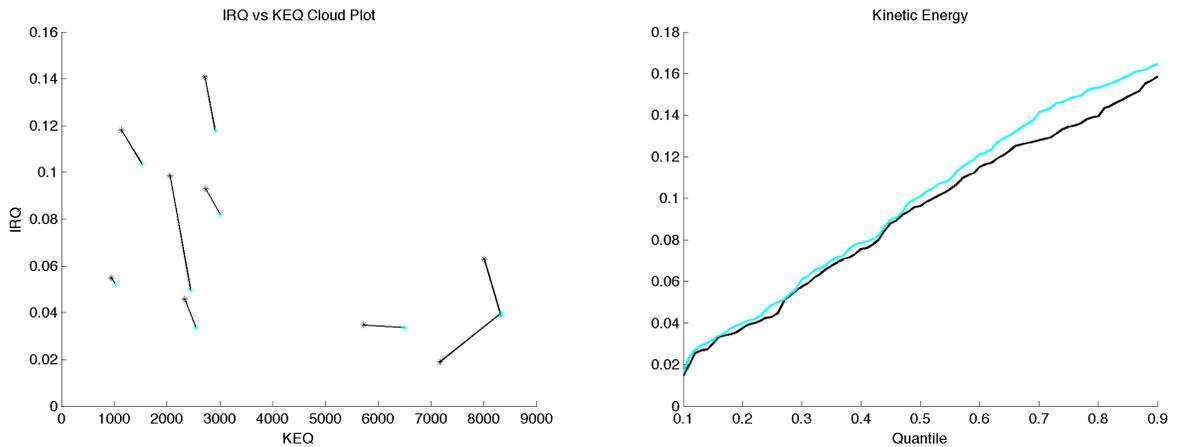


Figure 8.2: Sample displays used in performance monitoring of IRL. (Left) A plot of IRQ vs. KEQ for 9 voxel trajectories. (Right) Quantile curve of KEQ for all computed trajectories in IRL. Initial geometric matching quantities are in black, while the cyan values are after 1 iteration of the IRL.

Deciding how to use these performance monitors to develop an automatic stopping criteria is a complex problem. Using a fixed number of IRL iterations is the simplest strategy. Another termination criterion is if a high percent of the GMQ and KEQ increase above a certain threshold (exact values are problem specific). Also, if enough IRL iterations are performed, the iterated IRQ will eventually bottom out at the noise floor inherent to the ultrasound image sequence under study, yielding yet another possible stopping criterion. As the next section shows, the first few iterations of IRL

might yield trajectories with decreasing IRQ and stable GMQ and KEQ but the next IRL iteration could cause an increase of GMQ and KEQ with no IRQ improvement. To handle these situations, histories of all 3 performance indicators are tracked. If large jumps are detected in GMQ or KEQ, the trajectory can be reset to an earlier iteration.

8.2 Examples

Let us return to a few examples given in Chapter 5.

8.2.1 Example 5.1 – Direct Construction

We come back to the situation in Chapter 5 Example 1. First, we generate two artificial intensity sequences, one with noise and one without, to overlay on the space being deformed by the vector field flow v . We then test the performance of the IRL for each of these image sequences. In both cases, image sequences are 5 frames on a $53 \times 53 \times 53$ voxel grid with voxel size $(0.05)^3$. Physically, this is a partition of the cube $[-1.3, 1.3]^3$ into 148877 equally sized voxels. This grid encompasses the spatial region (deformed unit sphere) studied by geometric diffeomorphic matching.

- a) (*White on Black*) For the first case, we generate an artificial intensity model where the background intensity is given by

$$I_b(x, y, z, t) = \max\{0, \min[260, 100|1 - x^2 - y^2 - z^2| + 10 \sin(x + y)]\}.$$

while the foreground (points interior to deformed lower sphere) intensity is given by

$$I_f(x, y, z, t) = 240 + 5 \sin(t)$$

See Figure [8.3] for a visualization of this image sequence at fixed time. We

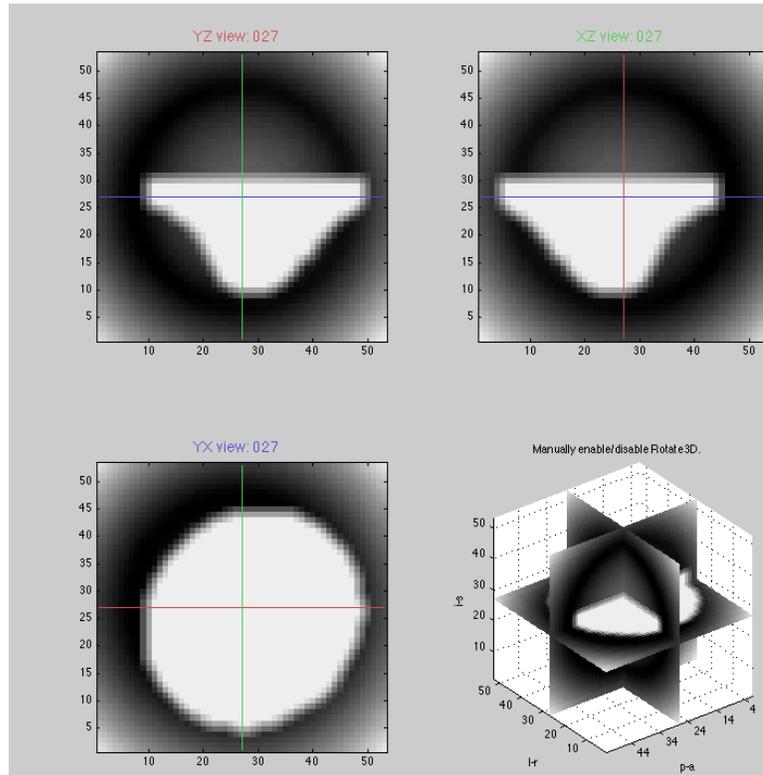


Figure 8.3: 3D time slice of case a). White on black image sequence.

expect that the IRL will perform well because any voxel trajectories that are not aligned with the foreground intensity should be perturbed by IRL within the first iteration. We proceed to run the IRL algorithm for 3 total iterations before terminating it due to high-energy trajectories. Quantile curves for the intensity oscillations along each trajectory (IRQ) as well as the kinetic energy

of each trajectory (KEQ) are displayed in Figure [8.4]. The GMQ is not shown as all iterations remain stable and close to geometric trajectories. Notice that with only 1 iteration of IRL (blue curves), the 90% quantile of IRQ reduces from approximately 0.32 to approximately 0.22, a decrease of almost 50%, while the next two iterations (cyan and yellow) only further reduces the IRQ to approximately 0.2. This validates our hypothesis stated before regarding the quick action of the IRL to white-on-black image sequences. The 90% quantile of KEQ, on the other hand, remains approximately constant at about 0.052 for the first 2 IRL iterations before jumping to a value of about 0.06 on the third IRL iteration. This can be interpreted by voxel trajectories that are very close to the foreground/background interface needing very little energy in order to be properly perturbed. A few trajectories, however, must be further from the interface so that the “step” observed in the KEQ quantile curve corresponds to these trajectories finally having enough energy to transition from the background to the foreground.

In order to find a balance in performance between these two criteria, the high energy trajectories (greater than 85% quantile) of the 3rd IRL iteration are reset to their values from the 2nd IRL iteration. This final solution is represented by the grey curve on both quantile plots and maintains the 90% quantile IRQ reduction to 0.2 while preserving the 90% quantile KEQ of 0.052.

b) (*Spatially Varying + Noise*) In this example the background intensity is given

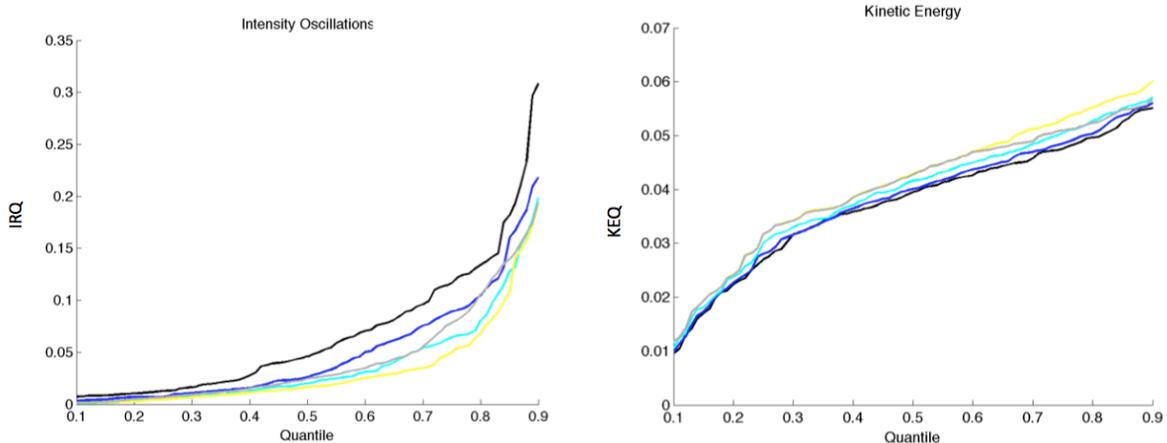


Figure 8.4: Quantile curves of intensity oscillations (IRQ) and voxel trajectory kinetic energy (KEQ) for case a). The black curve represents the values from geometric matching that are input into the IRL. The blue, cyan, and yellow are successive iterations of the IRL algorithm. Notice the decrease in intensity oscillation and increase in kinetic energy with each iteration. The grey curve is the result of resetting high energy trajectories to more stable, earlier iteration trajectories.

by

$$I_b(x, y, z, t) = \max\{0, \min[260, 100|1-x^2-y^2-z^2|+70r_{xyzt}]\}, \quad r_{xyzt} \in N(0, 70).$$

while the foreground (points interior to deformed lower sphere) intensity is given by

$$I_f(x, y, z, t) = 150 + 40 \sin(t) + 70r_{xyzt}, \quad r_{xyzt} \in N(0, 70)$$

Much like case a), the IRL runs for 3 iterations, again terminating due to high energy. Even in the presence of this Gaussian noise, the IRL algorithm maintains the GMQ provided by the input geometric matching trajectories and is

able to adjust the voxel trajectories in this image sequence so that the 90% quantile IRQ reduces by 40% in the first iteration while maintaining an approximately constant 90% quantile KEQ. The quantile curve analysis presents the same conclusions as in a) and is consequently omitted here. See Figure [8.5] for snapshot of the image sequence. Notice the foreground/background interface still appears to be easily identifiable, validating the analysis done in a) applies to this case as well.

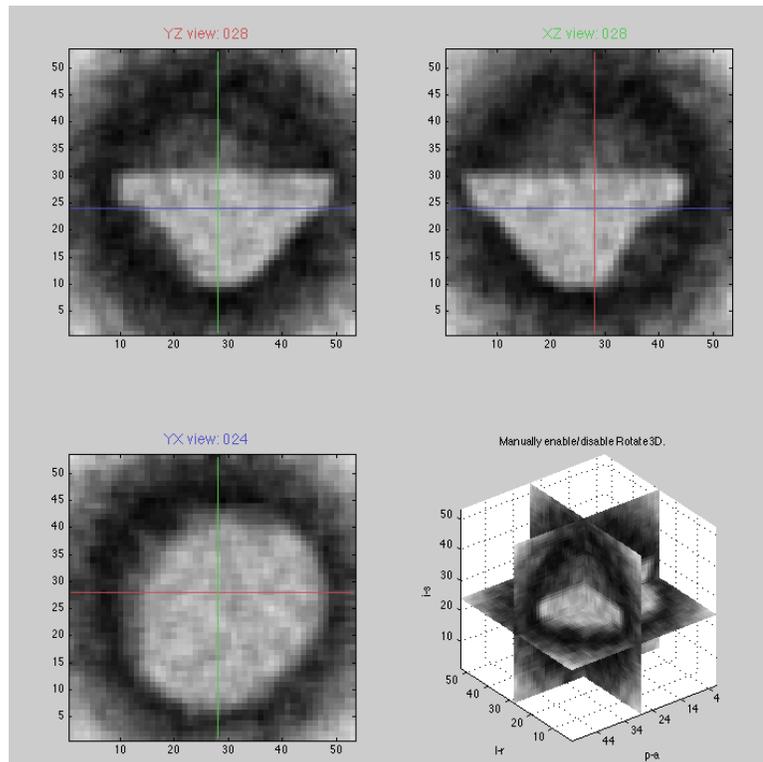


Figure 8.5: 3D time slice of case b). Noisy image sequence.

8.2.2 Example 5.2 – Mitral Valve Leaflet

As explained in Chapters 2-3, all patient-specific static mitral valve models were derived from 3D Echocardiographic image sequences. These image sequences had a resolution of $208 \times 224 \times 208$ voxels with an average voxel size of approximately 0.5 mm^3 . This particular example comes from a sequence of 6 frames, with successive frames separated by real-time $2/27$ seconds. Ultrasound echocardiography is a noisy image modality; for this example the data is smoothed by $3 \times 3 \times 3$ box kernel 3 separate times before the IRL iterations begin. Once the simulation began, the aforementioned performance criteria were monitored. See Figure [8.6]. Here the quantile curve of IRQ shows that the IRQ decreased steadily with each successive IRL iteration (blue to orange). The planar display of IRQ vs. KEQ histories demonstrates that approximately 5% of the trajectories began varying wildly in KEQ while demonstrating marginal IRQ gains. The IRL iterations were terminated after 5 iterations due to these large KEQ increases ($> 100\%$). At this point, trajectories that were consistently exhibiting high energy (top 25%) were reset to earlier, more stable iterations (either geometric matching or IRL iteration 1). The next chapter more carefully explores how the IRL applies to the dynamic modeling study of 30 human mitral valves in vivo which constitute the backbone of our exploration of biomedical applications.

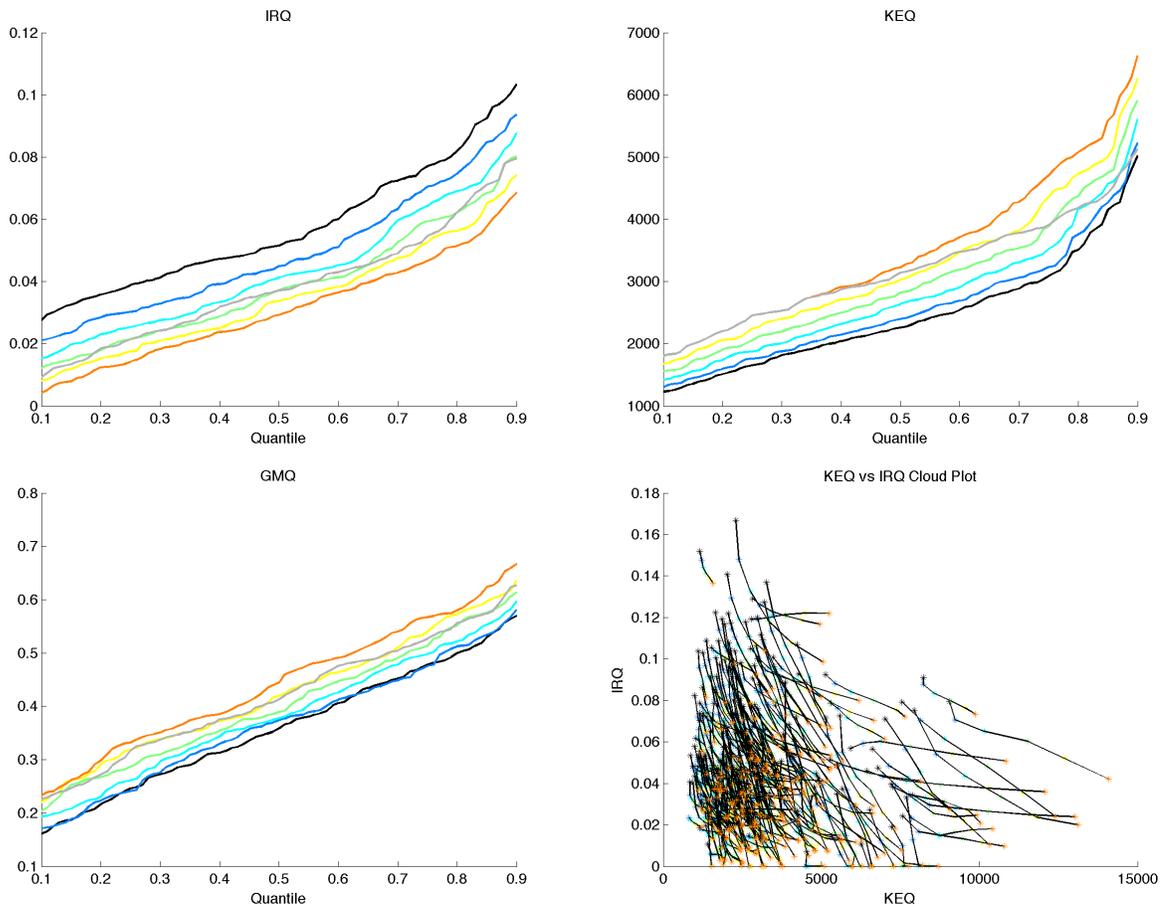


Figure 8.6: Performance criteria tracked during successive IRL iterations applied to the dynamic modeling of an anterior leaflet. The black curve corresponds to the values from geometric matching (chapter 5) that are input to the IRL, while the colored trajectories correspond to successive IRL iterations. The grey curve gives the values after trajectories are reset from high-energy, unstable values. The planar display of IRQ vs. KEQ histories demonstrates that approximately 5% of the trajectories began varying wildly in KEQ while demonstrating marginal IRQ gains, seen by the nearly horizontal slopes in the (KEQ,IRQ) plane. However, a large majority of the paths in the (KEQ,IRQ) plane remain bounded in KEQ during the iteration, seen by the large multi-colored “fuzzball” at the lower left of the cloud plot and confirmed by the KEQ quantile plot which shows that about 80% of the trajectories remain bounded above by 5000 KEQ after 5 iterations while beginning the IRL at a 80%-quantile of 3000 KEQ.

Chapter 9

Exploration of Geometric Matching + IRL

9.1 Overview

The previous chapter highlighted a method of diffeomorphic matching that combines geometric matching with voxel intensity registration. This methodology was applied to the 20 patients study (30 patient specific dynamic models of mitral valve deformations) that formed the backbone of the motivation for work in this thesis. But before proceeding with this study, we evaluated and tested our iterated IRL algorithmics extensively. This chapter highlights some of these validation studies.

9.2 Performance Impact Analysis for the main parameters of the IRL algorithm

9.2.1 Impact of the IRL parameter ϵ

Recall that the parameter ϵ directly determined the width of the box in the box-search algorithm (step 2c) of the IRL trajectory perturbation stage (Chapter 7). This subsection gives the results for 18 dynamic modeling cases (6 normal leaflets + 6 organic pre-surgery leaflets + 6 post-surgery leaflets) when a fixed number of IRL iterations are run for varying values of the ϵ in the voxel trajectory perturbation stage. All solutions are forced to have geometric matching quality (GMQ) on the order of 1 mm (0.5 - 1.5 mm range). We expect that smaller values of ϵ will yield smaller relative changes in kinetic energy quality (KEQ) and intensity registration quality (IRQ) than will larger values of ϵ , as this parameter directly controls the “wobble room” for trajectory perturbation. As we see in Tables [9.1] and [9.2], this trend holds true. Also, for 4 IRL iterations, it appears that $\epsilon = 0.2$ provides a better balance between KEQ gains and IRQ reduction than does $\epsilon = 0.8$. For all normal and organic cases tested, when $\epsilon = 0.2$ we see approximately a 20% increase of voxel trajectories kinetic energy generates voxel trajectories with at least 30% less intensity fluctuations. Contrast this with $\epsilon = 0.8$; here an increase of approximately 70% in trajectory kinetic energy is necessary to generate voxel trajectories with at least 40% less intensity fluctuations. There is a threshold value above which further increases of epsilon yield only small IRQ reduction and only serves to substantially increase

the KEQ. We have also noted rather large kinetic energies for the post-surgery cases. This difference stems from the fact that, for the post-surgery mitral valves at mid and end-systole, the valve remains roughly in the same spatial location and orientation. Thus, the voxel trajectories computed by diffeomorphic matching have small kinetic energy and the intensities computed at mid and end-systole are very similar. When any slight adjustment is made to these trajectories at the IRL perturbation stage, these two points explain both the smaller relative IRQ reduction and higher relative KEQ increase.

On the basis of similar detailed evaluations conducted on 18 mitral leaflets, we have concluded that whenever a small number of IRL iterations (2-4) are to be run, a small value of ϵ (0.2-0.3) should be chosen to yield balanced results that maintain the GMQ at a pre-assigned, physically-acceptable value, reduce the IRQ, and only slightly increase the KEQ.

9.2.2 Impact of Number of IRL Iterations

This subsection gives the results for 3 dynamic modeling cases (1 normal leaflet + 1 organic pre-surgery leaflet + 1 post-surgery leaflet) as the we increase the number of iterations of IRL for a fixed value of ϵ . Keeping with the previous findings, we fix $\epsilon = 0.2$ for our evaluation study presented here. As reported in Table [9.3] and Figure [9.1], this study demonstrates a stabilization of the geometric matching + IRL algorithm; namely, our iterated IRL algorithm converges to a solution where there

Patient Model	# IRL Iterations	$Q_{0.9}[\frac{\Delta(KEQ)}{KEQ}]$	$Q_{0.9}[\frac{\Delta(IRQ)}{IRQ}]$
Normal 1 AL	4	0.27	-0.23
Normal 1 PL	4	0.22	-0.22
Normal 8 AL	4	0.14	-0.24
Normal 8 PL	4	0.19	-0.38
Normal 42 AL	4	0.22	-0.34
Normal 42 PL	4	0.22	-0.36
Organic 28 AL	4	0.33	-0.22
Organic 28 PL	4	0.29	-0.27
Organic 49 AL	4	0.19	-0.29
Organic 49 PL	4	0.20	-0.38
Organic 76 AL	4	0.22	-0.22
Organic 76 PL	4	0.25	-0.30
PostOp 27 AL	4	0.69	-0.13
PostOp 27 PL	4	0.74	-0.12
PostOp 48 AL	4	0.62	-0.11
PostOp 48 PL	4	0.49	-0.13
PostOp 61 AL	4	0.78	-0.17
PostOp 61 PL	4	0.80	-0.09

Table 9.1: Results for $\epsilon = 0.2$. For a fixed number of 4 IRL iterations a 20% increase in the kinetic energy of voxel trajectories generates trajectories with at least 30% less intensity fluctuations

is nearly minimal intensity fluctuations on the trajectories and the constraints imposed by the encompassing 3D ultrasound image sequence data do not allow further decreases in intensity fluctuations without greatly increasing the available energy budget. This behavior is intuitively desirable, because it indirectly validates the fact that the purely geometric diffeomorphic solution approximately matched the true mitral valve position at each image frame of the ultrasound image sequence. Numerical evidence of this stabilization can be seen in Table[9.3] by viewing the ratio of successive row differences. For instance, consider the normal 42 AL data. The

Patient Model	# IRL Iterations	$Q_{0.9}[\frac{\Delta(KEQ)}{KEQ}]$	$Q_{0.9}[\frac{\Delta(IRQ)}{IRQ}]$
Normal 1 AL	4	0.97	-0.38
Normal 1 PL	4	1.21	-0.52
Normal 8 AL	4	0.57	-0.41
Normal 8 PL	4	0.55	-0.60
Normal 42 AL	4	0.72	-0.54
Normal 42 PL	4	0.88	-0.38
Organic 28 AL	4	1.08	-0.40
Organic 28 PL	4	1.18	-0.39
Organic 49 AL	4	0.76	-0.50
Organic 49 PL	4	0.73	-0.32
Organic 76 AL	4	0.69	-0.31
Organic 76 PL	4	0.86	-0.49
PostOp 27 AL	4	2.88	-0.23
PostOp 27 PL	4	2.57	-0.27
PostOp 48 AL	4	2.68	-0.25
PostOp 48 PL	4	2.06	-0.36
PostOp 61 AL	4	2.13	-0.41
PostOp 61 PL	4	2.81	-0.26

Table 9.2: Results for $\epsilon = 0.8$. For a fixed number of 4 IRL iterations a 70% increase in the kinetic energy of voxel trajectories is necessary to generate trajectories with at least 40% less intensity fluctuations.

second iteration requires 4% more energy for a 13% reduction of intensity oscillations, giving a quality ratio $QR_1 = 22/4 = 3.5$. Contrast this with the fifth iteration and tenth iterations, which respectively require 4% and 2% more energy and only yield a 6% and 1% reduction in intensity oscillations, so that $QR_5 = 6/4 = 1.5$ and $QR_{10} = 1/4 = 0.25$. Tracking these quality ratios during IRL offers a good potential stopping criterion for the IRL; once calculated QR values begin to decrease (after at least 2-3 iterations), it is a good time to stop the loop and proceed to the voxel trajectory post-processing to fix any errant trajectories.

9.3 Conclusion

We have just outlined the main results of our detailed evaluation study for the links between the parameters of the IRL algorithm and their impact on our 3 matching quality criteria. These results have directly influenced our IRL parameter choices for the study of our 30 cases of dynamic mitral valve modeling. All geometric matching + IRL were implemented with small ϵ (0.2-0.3) and a small number of fixed IRL iterations (2-4). The exact values depended upon each individual model, but the precise numbers were chosen so that at the end of our IRL iterations, more than 90% of voxel trajectories exhibited a decrease of at least 15% in intensity fluctuations (IRQ), without requiring more than a 20% increase in kinetic energy (KEQ). We also imposed that the geometric matching quality (GMQ) had to remain close to 1mm, which was the known precision of our initial patient-specific static models. As mentioned in Chapter 8, trajectory histories were stored so that any particular trajectory for which one of our 3 criteria (IRQ,KEQ, GMQ) suddenly “blew up” could be reset to a previous stable state. Most often, these “difficult to improve ” trajectories were consistently in the higher percentiles of kinetic energies KEQ, so their status as “difficult to improve” can be traced throughout the IRL iterations history. But the ultimate decision comes down to a human-in-the-loop who monitors the performance of the algorithm and decides if the 2-4 iterations are sufficient and if any trajectories need to be reset.

In summary, the geometric matching + IRL method has proven to be a successful in its application to the UH/Methodist Hospital mitral valve study, as well as to some

synthetic shape matching examples (Chapter 7). But the algorithm is far away from being “automatic”, as each application to a new image registration problem requires a careful study of the underlying image sequence. The good news is that if the UH/Methodist study expands its patient roster, the necessary tools to continue the study are in place for a quicker turnaround.

Patient Model	# IRL Iterations	$Q_{0.9}[\frac{\Delta(KEQ)}{KEQ}]$	$Q_{0.9}[\frac{\Delta(IRQ)}{IRQ}]$
Normal 42 AL	1	0.08	-0.19
Normal 42 AL	2	0.12	-0.32
Normal 42 AL	3	0.17	-0.41
Normal 42 AL	4	0.23	-0.50
Normal 42 AL	5	0.27	-0.56
Normal 42 AL	6	0.32	-0.59
Normal 42 AL	7	0.36	-0.62
Normal 42 AL	8	0.40	-0.65
Normal 42 AL	9	0.44	-0.67
Normal 42 AL	10	0.46	-0.68
Organic 76 PL	1	0.16	-0.11
Organic 76 PL	2	0.14	-0.25
Organic 76 PL	3	0.34	-0.29
Organic 76 PL	4	0.29	-0.36
Organic 76 PL	5	0.50	-0.42
Organic 76 PL	6	0.47	-0.47
Organic 76 PL	7	0.67	-0.49
Organic 76 PL	8	0.66	-0.56
Organic 76 PL	9	0.90	-0.48
Organic 76 PL	10	0.80	-0.53
PostOp 48 AL	1	0.13	-0.05
PostOp 48 AL	2	0.30	-0.08
PostOp 48 AL	3	0.47	-0.11
PostOp 48 AL	4	0.64	-0.13
PostOp 48 AL	5	0.79	-0.16
PostOp 48 AL	6	0.92	-0.18
PostOp 48 AL	7	1.03	-0.20
PostOp 48 AL	8	1.21	-0.21
PostOp 48 AL	9	1.37	-0.23
PostOp 48 AL	10	1.51	-0.25

Table 9.3: IRL Iteration Trade Results. The successive row difference method given in the text can be used to find the “sweet spot” where we get maximum reductions in IRQ for minimal increases in KEQ. For Normal 42 and Organic 76, we see this occurs at iteration 2. For PostOp48, we see this occurs at iteration 2 or 3.

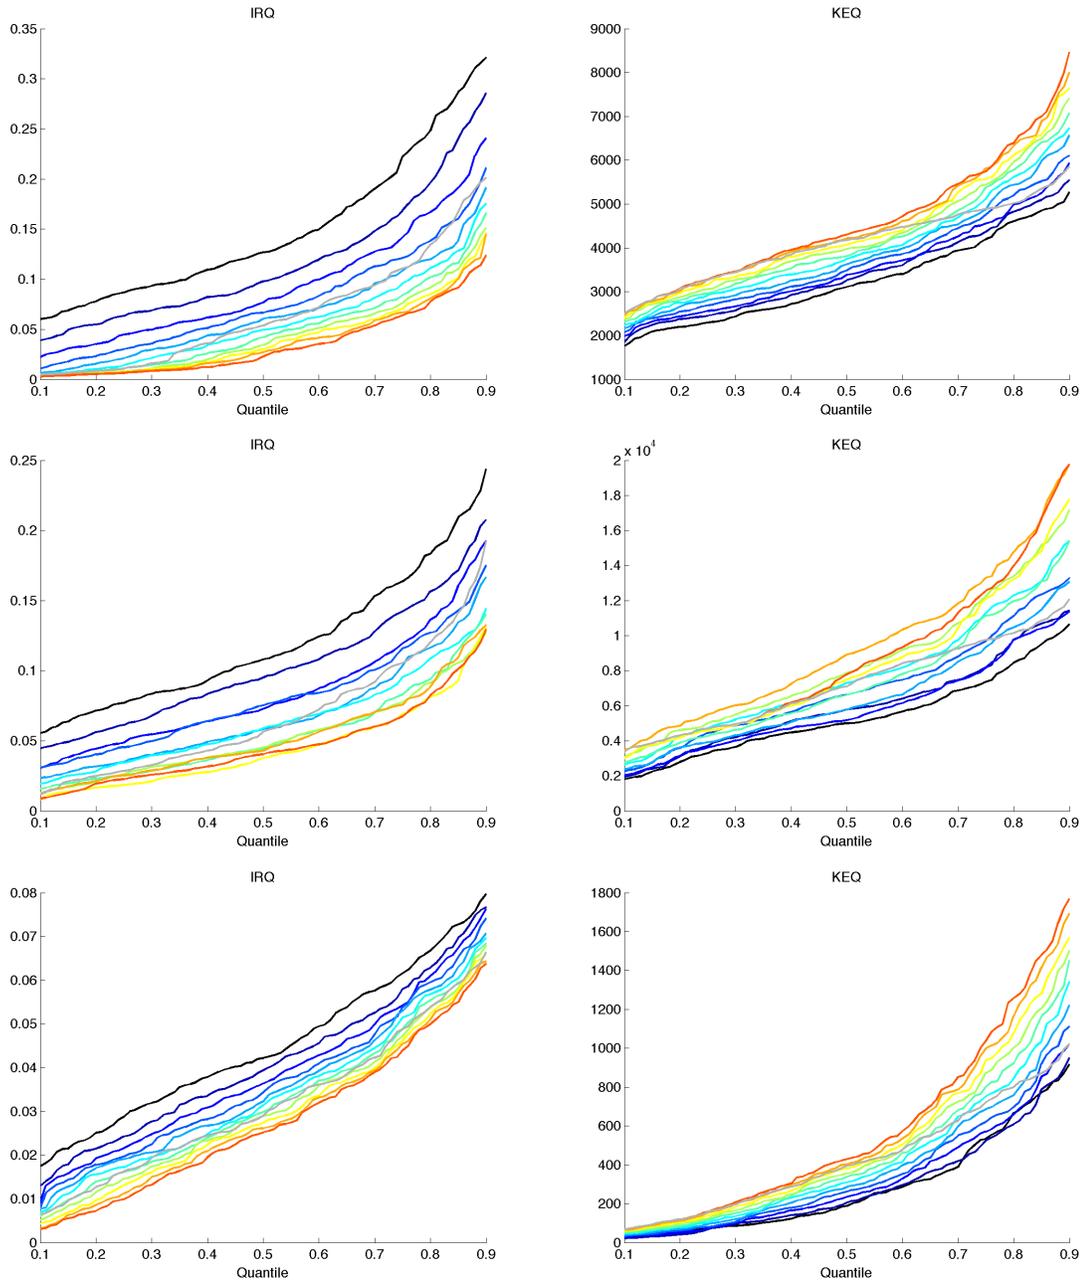


Figure 9.1: The IRQ and KEQ quantile curves monitoring of quality criteria are displayed here for three cases of mitral leaflets dynamic modeling: (top) Normal 42 AL, (middle) Organic 28 PL, (bottom) PostOp 48AL. Notice the marginal gains in IRQ reduction and KEQ increase for higher IRL iterations. This is an indicator of a good solution. For these cases, a small number of IRL iterations (2-4) yields the majority of the benefits with minimal increases in kinetic energy cost.

Chapter 10

Mitral Valve Strain Analysis

10.1 Motivation

The previous chapter showed that our algorithmic combination of geometric matching with iterated intensity registration yielded smooth deformations of the mitral valve between mid and end-systole for the 30 patient specific sequences of 3D-ultrasound images (10 normal, 10 pre-surgery, 10 post-surgery). The next step in the joint UH/Methodist project was to evaluate these deformations in terms of the strain exhibited by the mitral valve apparatus during the deformation occurring at each heart cycle. In this chapter, we first review the concepts of stress and strain. Next, we present the statistical methods needed to compare and quantify the distribution of strain values on mitral leaflets. Finally, we explore how the 3 groups of mitral valves (normal, pre-surgery, post-surgery) compare in terms of strain values distributions. These results have been outlined in our published joint paper [5].

10.2 Stress/Strain Overview

There exist numerous texts on stress and strain in a variety of applied settings; see [42] and [28] for treatments with a respective geoscience and fluid mechanical flavor. At its core, the concept of *strain* quantifies the relative modification of local lengths and surfaces for a somewhat elastic material undergoing a dynamic deformation. This dynamic deformation is linked to the ongoing *stress* of the material (often generated by a load or other external forces, as well as by internal forces generated by local geometric deformations). One of the earliest relations involving these ideas of stress and strain is due to Robert Hooke, whose legacy to introductory mechanics can be summarized with the famous Hooke's law

$$F = -kx$$

where F is an applied force on a ideal spring and x is the displacement of the spring from the neutral position. For more complex settings, this linear relationship between stress and strain can be summarized with the elastic tensor equation

$$\tau_{ij} = c_{ijkl}\epsilon_{kl}$$

where τ is the stress tensor, c is the elastic tensor, and ϵ is the strain tensor. This constitutive relation, along with Newton's Second Law, are the fundamental equations of continuum mechanics. But as we see below, the material strain tensor is a kinematic quantity that can be computed simply once the material displacement field

is known. Stress, on the other hand, depends on the elastic tensor and given values of strain; without well-known elastic parameters, stress computations lose reliable physical interpretation.

10.3 Strain Tensor and Mitral Leaflets Strain

Let us consider how to describe changes in the positions of points within a continuum. The position of a specific particle at time t relative to its position at a reference time t_0 can be expressed as a vector field, namely the displacement field u given by

$$u = u(t, t_0) = r(t) - r(t_0)$$

where $r(t)$ is the position of the point at time t and $r_0 = r(t_0)$ is the reference location of the point. When the final time t is implicitly understood, the displacement field can be written $u(r_0)$. This approach of following the displacements of particles specified by their original positions at some reference time is called the *Lagrangian* description of motion in a continuum (the other is approach is the *Eulerian* description which considers what happens at a fixed location in space). Now at initial time t_0 , consider a reference point x_0 and a nearby point $x_1 = x_0 + \delta$ for some small vector δ . We have by Taylor expansion that at time t

$$u(x_1) = u(x_0) + J\delta$$

where

$$J = \begin{pmatrix} \frac{\partial u_x}{\partial x} & \frac{\partial u_x}{\partial y} & \frac{\partial u_x}{\partial z} \\ \frac{\partial u_y}{\partial x} & \frac{\partial u_y}{\partial y} & \frac{\partial u_y}{\partial z} \\ \frac{\partial u_z}{\partial x} & \frac{\partial u_z}{\partial y} & \frac{\partial u_z}{\partial z} \end{pmatrix}$$

is the spatial Jacobian of the displacement field. This first order approximation neglecting higher order terms is the basis of *infinitesimal strain theory* [42]. The *strain tensor* ϵ is the symmetric part of the Jacobian; that is

$$\epsilon = \begin{pmatrix} \frac{\partial u_x}{\partial x} & \frac{1}{2} \left[\frac{\partial u_x}{\partial y} + \frac{\partial u_y}{\partial x} \right] & \frac{1}{2} \left[\frac{\partial u_x}{\partial z} + \frac{\partial u_z}{\partial x} \right] \\ \frac{1}{2} \left[\frac{\partial u_y}{\partial x} + \frac{\partial u_x}{\partial y} \right] & \frac{\partial u_y}{\partial y} & \frac{1}{2} \left[\frac{\partial u_y}{\partial z} + \frac{\partial u_z}{\partial y} \right] \\ \frac{1}{2} \left[\frac{\partial u_z}{\partial x} + \frac{\partial u_x}{\partial z} \right] & \frac{1}{2} \left[\frac{\partial u_z}{\partial y} + \frac{\partial u_y}{\partial z} \right] & \frac{\partial u_z}{\partial z} \end{pmatrix}$$

or in index form $\epsilon_{ij} = \frac{1}{2} [\partial_i u_j + \partial_j u_i]$. Diagonalization of this tensor gives the *principal strains* (eigenvalues) which are the multiplicative factors modeling local lengths change along the directions specified by the *principal axes* (eigenvectors of ϵ). For the study of strain on a smooth surface S imbedded in \mathbb{R}^3 , the strain tensor expressed in local coordinates X and Y on a plane tangent to S at the point $p \in S$ is the 2 x 2 matrix

$$\eta(p) = \begin{pmatrix} \frac{\partial w_X}{\partial X} & \frac{1}{2} \left[\frac{\partial w_X}{\partial Y} + \frac{\partial w_Y}{\partial X} \right] \\ \frac{1}{2} \left[\frac{\partial w_Y}{\partial X} + \frac{\partial w_X}{\partial Y} \right] & \frac{\partial w_Y}{\partial Y} \end{pmatrix}$$

Here w is any vector expressed in tangent plane coordinates (X, Y) . Numerical evaluation of the strain tensor is straightforward with finite difference techniques.

For instance, consider a densely packed finite point grid $\Lambda \subset \mathbb{R}^3$ representing some smooth surface $S \subset \mathbb{R}^3$. Assume this surface undergoes some deformation u to a surface $u(S)$, meaning our point grid is now $u(\Lambda)$. For some $p \in \Lambda$, consider the intersection $B(p, \gamma)$ of the point grid Λ with the ball of center p and radius γ . A least-squares fit to $B(p, \gamma)$ can determine an approximate tangent plane to S at the point p (if γ is sufficiently small and the grid size $|\Lambda|$ is sufficiently large). This approximate tangent plane is denoted $T(p)$. Within this tangent plane, finite difference techniques can be applied to estimate the partial derivatives in the strain tensor. Higher order finite differences can be used to ensure that more points in the tangent plane $T(p)$ are used in the derivative estimates.

Mitral leaflets tissue elasticity has been studied experimentally by direct observation in open heart experiments performed with ovine and porcine animal subjects, as well as after autopsies of human patients. The elasticity models of mitral leaflets derived from these experimental data are hyper-linear and non isotropic (see [26]). For our live human mitral valve dynamics recorded by 3D-Echo data, in vivo validation of parameterized elasticity models is obviously not an option, which precludes direct parameterization of tissue elasticity models. This has led us to avoid leaflet stress computations, and to focus instead on computing patient-specific leaflets strain, which require no tissue elasticity assumption. Each one of our 30 mitral valve apparatus (MVA) deformation models generates the dynamic trajectories of several thousands of mitral leaflets points x between the times mid-systole (MS) and end-systole (ES). For each initial leaflet point x , and any small tissue patch $P(x)$ around x , the deformation of $P(x)$ between times MS and ES approximately multiplies small lengths

by a factor $g(x) > 0$, called the geometric strain at x . Technically, we compute the classical 2×2 surface strain tensor η at x and the moduli $|\lambda|, |\mu|$ of its two eigenvalues, which are the smallest and largest length dilating (or contracting) factors of deformations around x . Then $g(x)$ is the harmonic mean of these two factors, given by $g(x) = \sqrt{|\lambda||\mu|}$. Tissue deformations around x are roughly dilating if $g(x) > 1$, and contracting if $g(x) < 1$.

The leaflet Strain Intensity $SI(x)$ at x is then defined as the rough percentage of length dilation (or contraction) around x , given by $SI(x) = |g(x) - 1|$. Large values of $SI(x)$ are clearly indicative of high leaflet tissue fatigue around x . Strain intensities $SI(x)$ are computed at several thousand points x of each leaflet to characterize the distribution of strain intensity values on the leaflet surface. These strain maps are then displayed in color code on our patient specific 3D leaflet models (see Figure [10.1]). The statistical distributions of strain values are comparatively studied across patients groups in section 4 below. But first, we outline a quantitative method well-adapted to these specific comparisons.

10.4 Statistical Analysis via Kolmogorov-Smirnov

For each patient-specific mitral valve deformation model computed as described in our preceding chapters, strain intensity was computed for several thousand leaflet points. When one compares two large sets of numerical data $\{X_1, X_2, \dots, X_N\}$ and $\{Y_1, Y_2, \dots, Y_M\}$, with $N, M > 1000$, one classically says (see [33]) that the X values “are stochastically smaller than” the Y values if each quantile $Q_X(z)$ of the X values,

with $0\% < z < 100\%$, is inferior or equal to the corresponding quantile $Q_Y(z)$ of the Y values. This can be verified visually by plotting the two quantile curves $Q_X(z)$ and $Q_Y(z)$ to check if the curve Q_X indeed lies below the curve Q_Y . Reliability of the statement “ X is stochastically smaller than Y ” is evaluated by its p -value, given by one-sided Kolmogorov-Smirnov (KS) test, which is described below.

If the numerical values $\{X_1, X_2, \dots, X_N\}$ are drawn at random from a distribution F and $\{Y_1, Y_2, \dots, Y_M\}$ are drawn at random from a distribution G , then the two-sided KS test is a simple empirical test to the hypothesis $F = G$. It does so through the use of the empirical cdfs F_N and G_M ; specifically, the test statistic

$$D_{MN} = \sqrt{\frac{MN}{M+N}} \sup_z |F_N(z) - G_M(z)|$$

is computed for the data sets and compared to a reference threshold from the Kolmogorov-Smirnov distribution

$$H(z) = 1 - 2 \sum_{i=1}^{\infty} (-1)^{i-1} e^{-2i^2 z}$$

Analogously, the one-sided KS test, used to test the hypothesis $F > G$ proceeds by the computation of the test statistic

$$D'_{MN} = \sqrt{\frac{MN}{M+N}} \sup_z [F_N(z) - G_M(z)]$$

before comparison to $H(z)$. Both procedures works because a) the statistics D_{MN} and D'_{MN} are independent of the underlying distributions F, G and b) one has the

convergence

$$\mathbb{P}[\sqrt{n} \sup_z |F_n(z) - F(z)| \leq z] \rightarrow H(z)$$

in distribution. See [35] for details.

Through the use of quantile curves and the one-sided KS test, we are able to assign precise statistical meaning to such statements “sample A exhibits higher values than sample B”. These comparisons are a valuable metric when comparing the strain intensities derived from the patient-specific mitral valve model deformations from Chapter 9.

10.5 Global Strain Intensity Comparison

In this section, results are presented that compare the 3 populations: normal, organic pre-surgery (pre-op), and organic post-surgery (post-op).

Leaflet	Normal	Organic Pre-Op	Organic Post-Op
Anterior	0.10 ± 0.02	0.13 ± 0.02	0.09 ± 0.02
Posterior	0.13 ± 0.03	0.19 ± 0.05	0.10 ± 0.03

Table 10.1: Group averages of mean strain intensities. Strain intensities are higher for the organic group than the normal group (KS test p value = 0.05) for both leaflets.

10.5.1 Normal/Organic Pre-Surgery Comparison

Table [10.1] indicates that the mean strain intensities for the anterior leaflets AL have a higher group average for pre-op patients than for normal patients, and that this effect is even stronger for the posterior leaflets PL. This *strain severity* observed in the pre-op group is confirmed by the KS test as significant (p value = 0.05). A good visualization of this is given by Figures [10.2] and [10.3], which demonstrates higher strains for pre-surgery MR patients than for normal patients by respective quantile curve and histogram comparisons.

10.5.2 Organic Pre-Surgery/Post-Surgery Comparison

For each regurgitation case, graphic displays of color-coded leaflets stain intensities on our patient specific 3D-models of the MVA enable an immediate visual comparison of these color strain maps before and after non-resectional dynamic (NRD) repair surgery. Recall that this surgery involves the insertion of an artificial annular ring to the valve, thereby tightening the tissue and restoring elasticity to the valve. As seen on two typical such cases (Figure [10.4] below), a striking impact of NRD surgery is the clear reduction of the high leaflets strain intensities initially present in all regurgitation cases.

To validate the strain reduction achieved by NRD surgery, for each mitral valve regurgitation case, we have computed the pre-surgery strain quantile curves $Q_{AL.pre}$ and $Q_{PL.pre}$, and compared them to the corresponding two post-surgery strain quantile curves $Q_{AL.post}$ and $Q_{PL.post}$. Figure [10.5] displays these quantile curves for two

regurgitation cases, and clearly shows that $Q_{AL.post} \leq Q_{AL.pre}$ and $Q_{PL.post} \leq Q_{PL.pre}$. This is visually confirmed by similar graphic displays for our 10 regurgitation cases. To quantify the validity of statistical strain reduction achieved by NRD surgery, we have implemented 20 KS tests (2 per regurgitation case), which showed that stochastic strain reduction was valid for 9 of 10 posterior leaflets, and for 8 of 10 anterior leaflets. We thus conclude that in our 10 regurgitation cases, for 17 of the 20 mitral leaflets (2 leaflets per case), post-surgery leaflets strain intensities were stochastically smaller than pre-surgery strain intensities. This strain reduction effect of NRD surgery was validated by KS tests p -values inferior to 0.05 for each leaflet.

For each fixed percentile q and each regurgitation case, we have computed the reduction ratio $r(q) = Q_{AL.post}(q)/Q_{AL.pre}(q)$ between post-surgery and pre-surgery strain quantiles. To quantify the strain reduction achieved by NRD surgery on anterior leaflets, we compute for each percentile q , the average $StrRedAL(q)$ of reduction ratios $r(q)$ over our 10 regurgitation cases. A similar operation provides the analogous strain reduction coefficients $StrRedPL(q)$ for the posterior leaflets of regurgitation cases. These strain reduction coefficients are displayed in Table [10.2].

Percentile q	10%	20%	30%	40%	50%	60%	70%	80%	90%
$StrRedAL$	0.74	0.76	0.75	0.76	0.77	0.78	0.80	0.80	0.77
$StrRedPL$	0.62	0.60	0.60	0.60	0.60	0.59	0.59	0.60	0.61

Table 10.2: Multiplicative strain shrinking factors due to NRD repair surgery for the posterior leaflets and the anterior leaflets, after averaging over all regurgitation cases.

These strain quantiles reduction factors depend remarkably little on the percentile value q , which shows that NRD surgery roughly multiplies strain intensities by 0.60

for the posterior leaflet and by 0.77 for the anterior leaflet. We further validate this result in Figure [10.6], which displays the pooled post-surgery strain quantile curve $Q_{AL.post}$, the pooled pre-surgery strain quantile curve $Q_{AL.pre}$, and the rescaled quantile curve $0.77 \times Q_{PL.pre}$. It is visually clear that $Q_{AL.post}$ practically coincides with the rescaled $Q_{AL.pre}$ obtained by multiplying all observed pre-surgery PL strain intensities by 0.77. This is confirmed by a KS-test with p -value 0.01. A similar display for NRD surgery reduction of PL intensities by a factor of 0.60 is also presented in Figure [10.6] (and validated with p -value 0.01). For a histogram comparison, see Figure [10.7].

10.5.3 Normal/Organic Post-Operation Comparison

Due to the strong leaflets strain reduction effect achieved by NRD surgery, the mitral valve leaflets of post-surgery patients overwhelmingly exhibit strain intensities distributions which are either stochastically similar or slightly smaller than those of typical normal patients. This result can be visualized in Figure [10.8] below, which displays two nearly superimposed AL strain quantile curves: one strain quantile curve $Q_{AL.post}$ for our pool of 10 post-surgery cases, and one strain quantile $Q_{AL.norm}$ for our pool of 10 normal patients. The KS-test confirms with p -value 0.01 what we see visually: the post-surgery strain is less than or equal to typical normal patient strain. A similar comparison for normal and post-surgery PL strain validates the same conclusion for posterior leaflets. For a histogram comparison, see Figure [10.9].

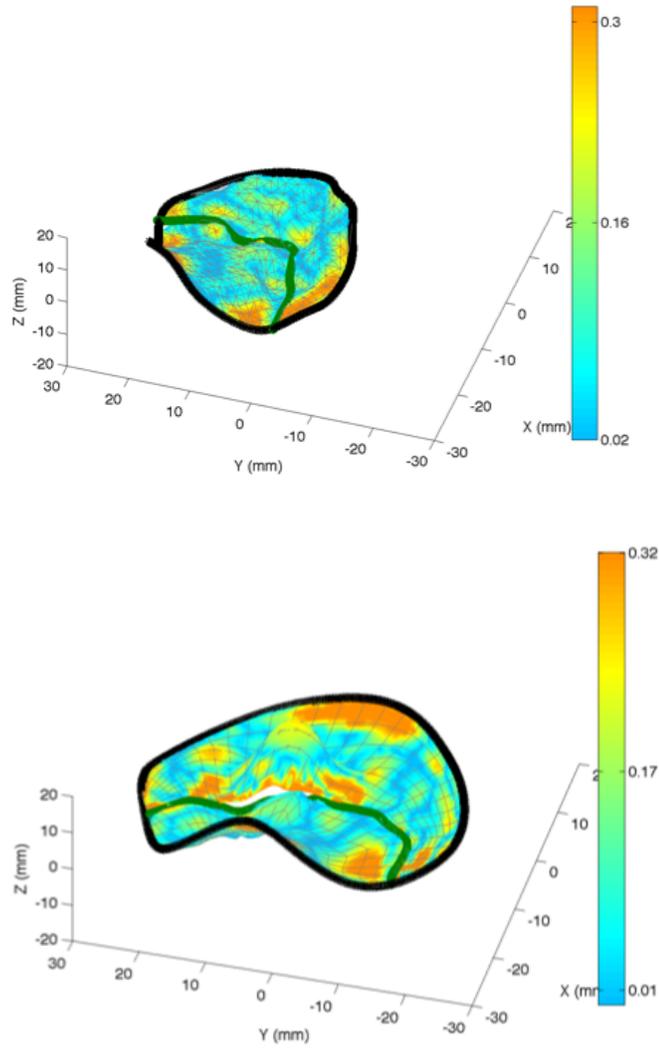


Figure 10.1: Two sample strain maps corresponding to a normal patient (top) and a diseased, pre-surgery patient (bottom). The high strain intensities are colored in orange, while low strain intensities are colored blue.

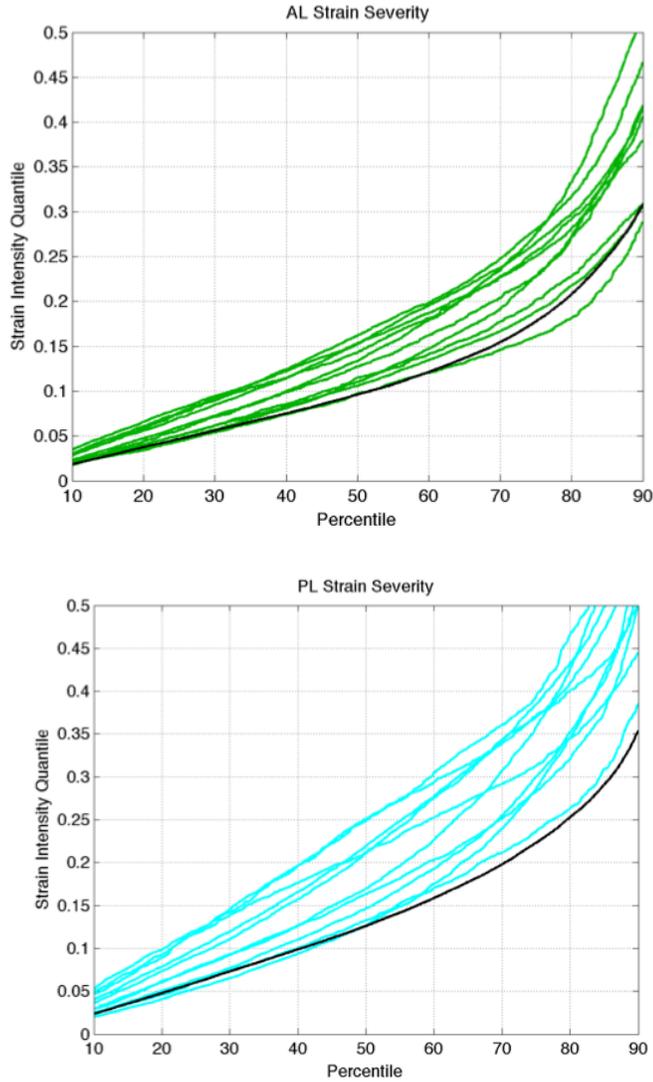


Figure 10.2: Quantiles curves N_{AL} and N_{PL} for the average normal strain intensity (shown above in black) demonstrate strain severity for pre-surgery MR patients; we in fact observe that strain is stochastically higher for pre-surgery cases than for average normal cases. This is quite clear graphically since the 10 pre-surgery strain quantiles curves displayed in green (AL) and blue (PL) are essentially above the black normal strain quantile curves.

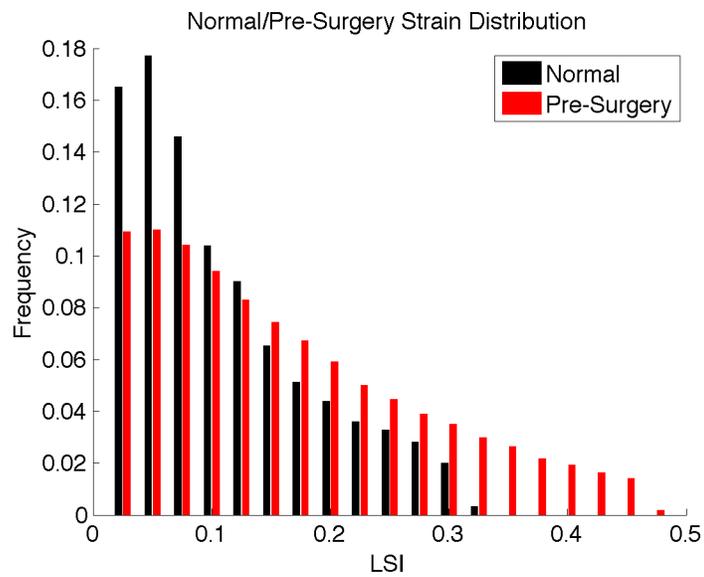


Figure 10.3: Histograms of strain intensities across both leaflets for all 10 normal patients (black) and 10 pre-surgery organic MR patients (red). Notice the clear shift in the red organic data towards higher values of leaflet strain intensity (LSI) compared to the black normal patient data.

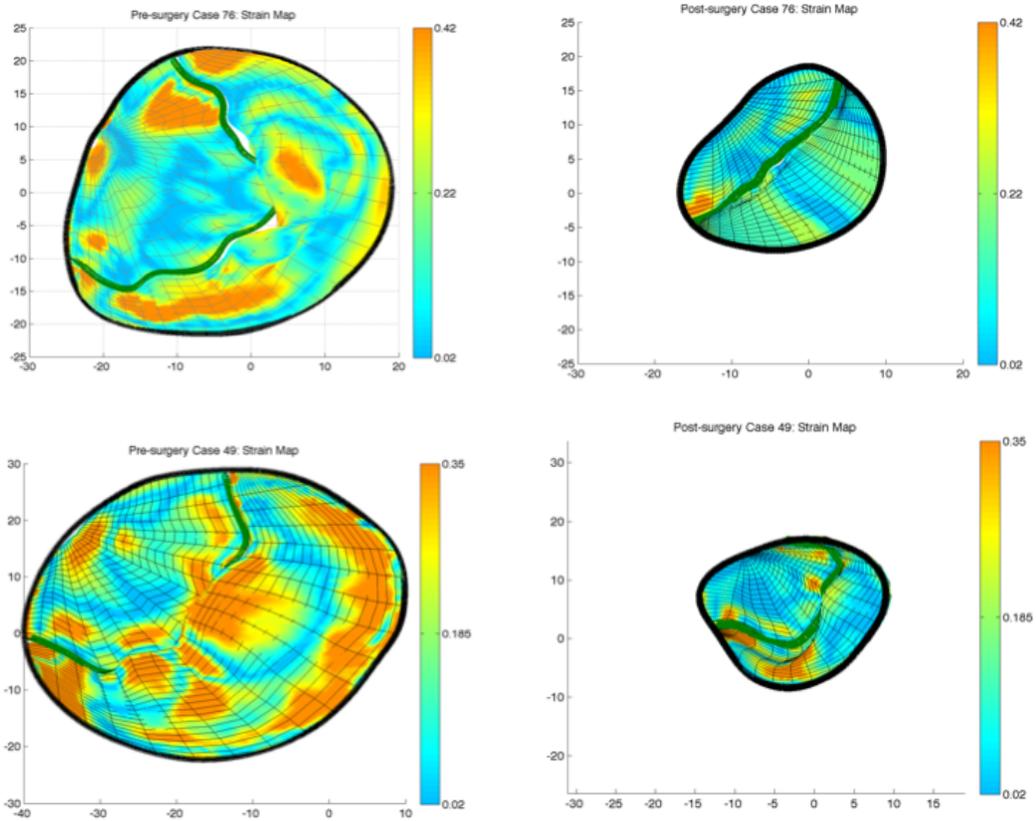


Figure 10.4: Leaflets Strain Intensities: color-coded maps at MS for two regurgitation cases. View is top-down, with AL towards the left. All axes have units of mm. Coaptation line is shown in dark green. Notice the reduction in strain after NRD surgery, evident in the increase in blue on the strain maps.

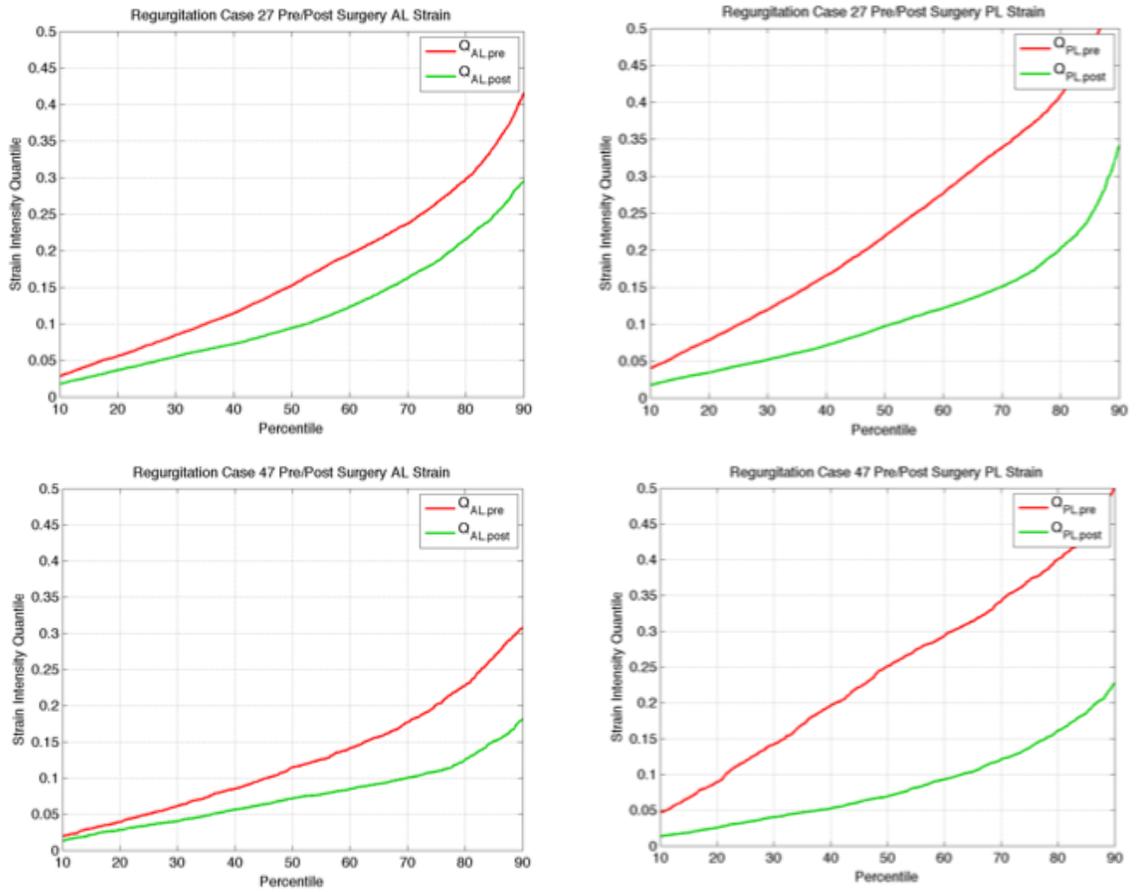


Figure 10.5: Quantiles curves for two regurgitation patients emphasize a strain reduction due to the NRD surgical repair. The two figures on the left represent the AL strain quantiles, while those in the right represent the PL strain quantiles.

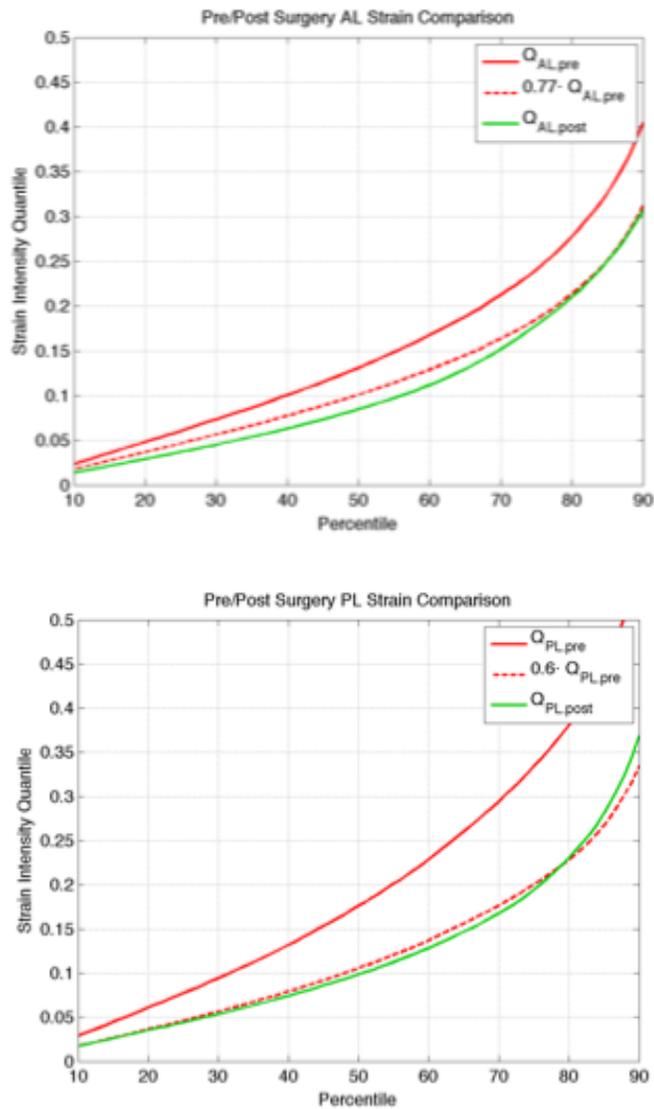


Figure 10.6: Quantile curves displaying the reduction of strain intensities due to NRD surgery. (Top) For the AL, our figure shows that, after multiplying the pre-surgery strain quantile curve by a shrinking factor of 0.77, the rescaled curve nearly coincides with the post-surgery strain quantile curve. (Bottom) For the PL, our right picture displays a similar result, with a multiplicative strain shrinking factor of 0.60.

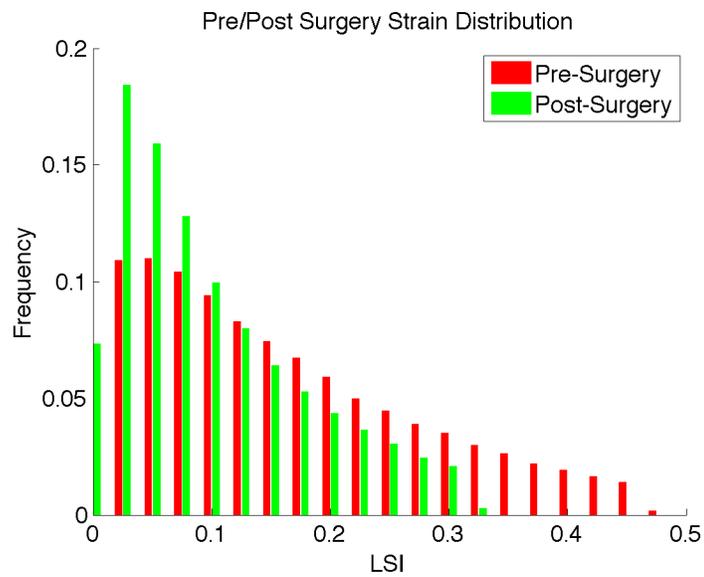


Figure 10.7: Histograms of strain intensities across both leaflets for all 10 organic MR patients, both pre-surgery (red) and post-surgery (green). Notice the clear shift in the green post-op data towards lower values of leaflet strain intensity (LSI) compared to the red pre-surgery patient data.

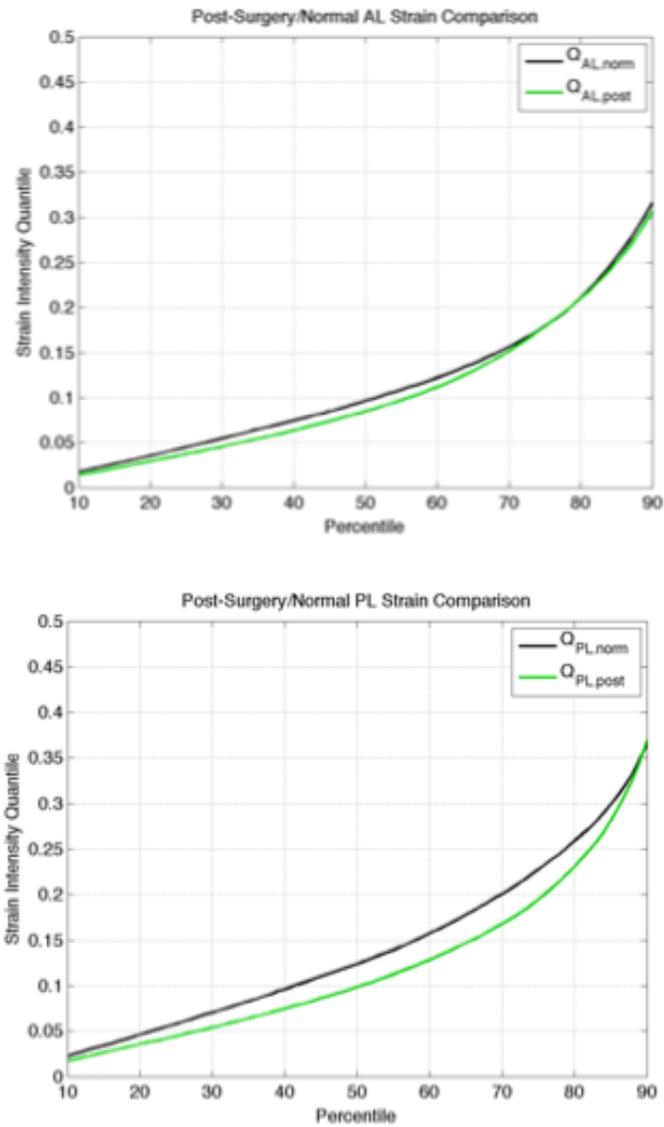


Figure 10.8: For the normal and post-surgery populations, strain quantile curves are nearly identical (as validated by KS-tests with p -values 0.01).

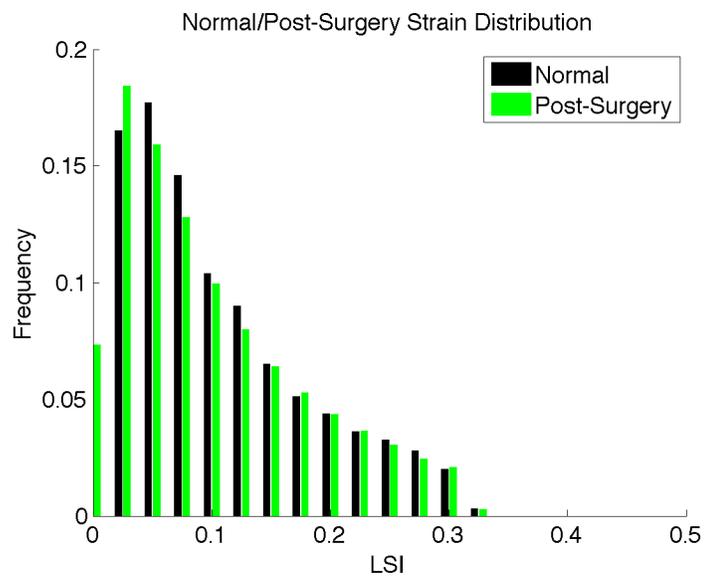


Figure 10.9: Histograms of strain intensities across both leaflets for all 10 normal patients (black) and 10 post-surgery organic MR patients (green). Notice the data for the two populations appear identical, illustrating the reduction of leaflet strain intensity (LSI) back to “normal” levels following the NRD surgery.

Chapter 11

Mitral Valve Strain–High Strain Localization Analysis

Most of the material presented in this chapter is extracted from a joint paper we are currently finalizing for submission to the journal *Cardiovascular Engineering and Technology* [6].

11.1 Motivation

The previous chapter explored the global distribution of strain across all 3 model populations. There it was showed that NRD mitral valve repair surgery reduced the high strain seen in organic in mitral valves exhibiting mitral regurgitation and/or prolapse back to or below strain values seen in normal mitral valves. These results are promising. Methodist Hospital cardiologists collaborating with our research team

also raised the question of characterizing how strain was distributed across *individual* mitral leaflets. The purpose of this chapter is to outline our methodologies for the study of high strain localization, and to present our main results in this direction.

11.2 Computerized Generation of Leaflet Scallops

Anatomical markers were interactively tagged by cardiologists for each one of our MVA models in order to identify the three “scallop” A1, A2, and A3 of the anterior leaflet AL and scallops P1, P2 and P3 of the posterior leaflet PL. For a general leaflet, the three scallops are denoted $S_{c_1}, S_{c_2}, S_{c_3}$ below. Specifically, clinicians studied each patient model and selected two sets S_1, S_2 of tags. Each set S_i consisted of a point O_i near the anterior horn of the mitral annulus MA, a point Y_i on the coaptation line COAPT, and a point X_i on the posterior side of the MA. Then, planes Q_i were fit to the three points in $S_i = \{O_i, Y_i, X_i\}$. The two intersections $Q_i \cap \text{MVA}$ of the planes Q_i with the mitral leaflets determined the boundaries of the classical anatomical subdivision of each leaflet into three scallops. See Figure [11.1] for an example.

11.3 Other Leaflet Zones

To study the distribution of strain on both the anterior and posterior leaflets, we defined 8 regions of interest (see Figure [11.2] which displays the 8 zones for the PL). The three scallops $S_{c_1}, S_{c_2}, S_{c_3}$ were identified by the computerized procedure just described above. This procedure automatically generated scallops $S_{c_1}, S_{c_2}, S_{c_3}$ with

areas respectively equal to $25\% \pm 2\%$, $50\% \pm 2\%$, $25\% \pm 2\%$ of the whole leaflet area, for all the 60 mitral leaflets we had modeled in our patient groups.

The remaining 5 zones were constructed as follows. The annulus band Z_{ann} and the coaptation band Z_{coapt} are resp. the sets of all leaflet points within a distance r of the annulus (resp. a distance s of the coaptation line) so that the areas of Z_{ann} and Z_{coapt} were both equal to 20% of total leaflet area. The boundary zone Z_{bound} is then the union of Z_{ann} and Z_{coapt} . The center zone Z_{center} is the complement of Z_{bound} within the leaflet. Finally, define two ‘‘commissure points’’ cp_1 and cp_2 as the virtual intersections of the MA and the coaptation line. The set of all leaflet points which are within a small distance u of either cp_1 or cp_2 defines the commissure zone Z_{comm} . The distance u was selected so the area of Z_{comm} was equal to 8% of total leaflet area. All geometric zones were computed at mid-systole.

11.4 High Strain Concentration

The high strain zone Z_{HS} is defined as the set of all leaflet points where strain intensity is larger than the patient specific 75% quantile of strain intensities. The area of Z_{HS} is then always equal to 25% of total leaflet area. For each leaflet and each zone Z , the high strain concentration in zone Z is the ratio

$$HSC(Z) = \text{Area}(Z_{HS} \cap Z) / \text{Area}(Z_{HS})$$

Values of $HSC(Z)$ markedly close to 1 do indicate stronger concentrations of high strain points within Z ; the highest possible value of $HSC(Z)$ is 1 and can only be reached when Z contains Z_{HS} .

11.5 Regional Strain Analysis

Analysis of regional strain (see Table [11.1]) revealed high strain concentration on the commissures Z_{comm} (with the highest concentration) and also on boundary zones Z_{ann} and Z_{coapt} . However, there were no distinguishable differences between high strain concentrations on Z_{ann} and Z_{coapt} . The union Z_{bound} of Z_{ann} and Z_{coapt} is thus a fairly homogeneous zone where strain concentration is higher than on Z_{center} . These high strain concentration inequalities

$$HSC(Z_{comm}) \geq [HSC(Z_{bound}) \sim HSC(Z_{ann}) \sim HSC(Z_{coapt})] \geq HSC(Z_{center})$$

hold for both the organic MR and normal group, and for both leaflets AL and PL. These inequalities are confirmed by the KS test as significant (p -value = 0.04).

As for the scallops, scallop Sc_1 exhibited the highest strain concentration, and the high strain inequalities

$$HSC(Sc_1) \geq HSC(Sc_3) \geq HSC(Sc_2)$$

hold for both groups and both leaflets (p -value = 0.05).

	Normal Group HSC		Organic MR Group HSC	
Zone	AL	PL	AL	PL
Z_{comm}	0.44 ± 0.20	0.37 ± 0.20	0.44 ± 0.16	0.37 ± 0.15
Z_{ann}	0.36 ± 0.06	0.31 ± 0.08	0.31 ± 0.15	0.35 ± 0.17
Z_{coapt}	0.35 ± 0.10	0.32 ± 0.05	0.36 ± 0.10	0.37 ± 0.12
Z_{bound}	0.35 ± 0.03	0.31 ± 0.05	0.32 ± 0.09	0.35 ± 0.10
Z_{center}	0.19 ± 0.02	0.21 ± 0.03	0.21 ± 0.05	0.19 ± 0.07
Sc_1	0.42 ± 0.11	0.35 ± 0.14	0.34 ± 0.12	0.32 ± 0.10
Sc_2	0.15 ± 0.07	0.24 ± 0.07	0.16 ± 0.06	0.24 ± 0.06
Sc_3	0.29 ± 0.07	0.18 ± 0.09	0.33 ± 0.11	0.20 ± 0.06

Table 11.1: Comparison of high strain concentrations $HSC(Z)$ fo 8 leaflets zones Z .

11.6 Anatomical Interpretation and Analysis of Strain Localization Results

Regional analysis of High Strain Concentration (HSC) revealed interesting results. The highest mean HSC was noted near the commissures Z_{comm} and in the scallop Sc_1 , followed by the boundary zone Z_{bound} . The lowest values of HSC were seen in the center zone Z_{center} . As could be expected, regional analysis in the organic PL has shown elevated mean HSC in all zones but the center zone. In general, any leaflets edge presented higher deformation, and hence higher strain values. This is consistent with previous report by Stevanella et al [44] who studied anterior mitral leaflet on sheep hearts showing high strain values in leaflets edge. Logically, annulus deformation influences the annulus zone. The coaptation zone is influenced by annulus dynamics, the pressure gradient across the valve (left ventricle and left atrium pressure), as well as papillary muscle relaxation and reduced chordae stress.

In organic patients, tissue properties and size is crucial; since coaptation is not complete, any tissue that is floppy and weak as well as reduced chordae forces, allow for more deformation. This result agrees with those reported from simulations of the central zone alone ([21], [39]), where the authors demonstrate more irregularity in PL collagen fiber, and suggest that collagen fibers are designed to allow leaflets coaptation but guard against any further deformation and possible regurgitation. Higher strain observed near the commissures requires special interest; these two small zones overlap with both the annulus and coaptation zone (combined boundary zone). In both normal and organic patients, annular enlargement, especially at the anterior posterior dimension, accompanied by reduction in leaflet coaptation, allows for more deformation in those zones. A reduced HSC was noted in the Sc_2 scallop on AL while an elevated HSC was noted in the Sc_2 scallop on PL for both organic and normal models. The AL middle scallop Sc_2 is attached to the aortic mitral continuity and is characterized by fibrotic tissue that probably doesnt allow for much deformation.

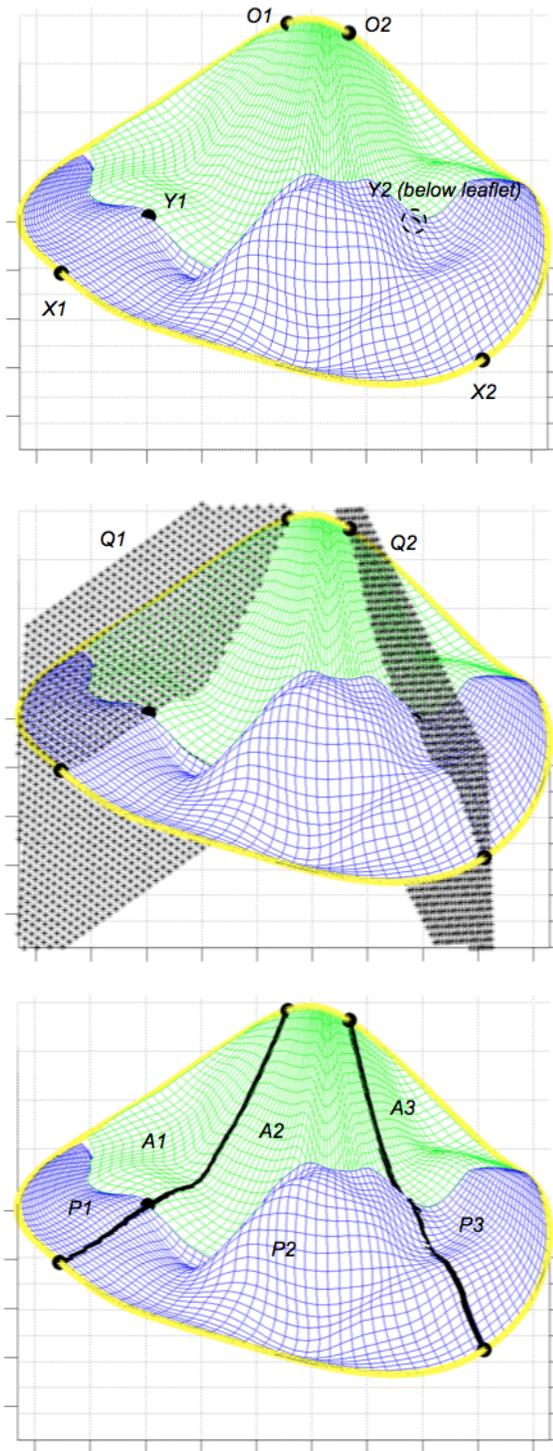


Figure 11.1: Computerized scallop generation process.

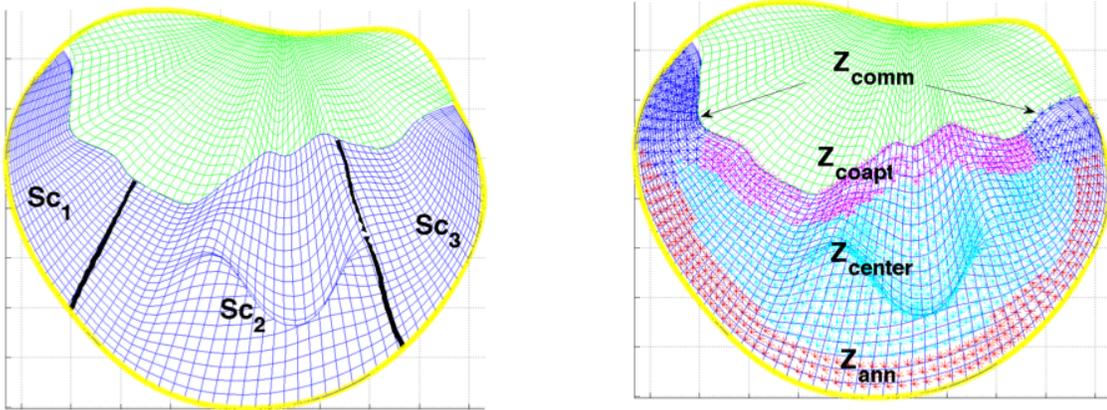


Figure 11.2: Computed geometric zones on one typical posterior leaflet PL. The left figure displays the 3 Posterior Leaflet scallops, while the right figure displays our 5 other geometric PL zones: Z_{coapt} (magenta), Z_{center} (cyan), Z_{ann} (red), and the two commissures Z_{comm} (blue).

Chapter 12

Outline of Future Work: The Shape Space of Human Mitral Valves

12.1 Motivation

The human mitral valves deformations and strain study and the analysis of how NRD mitral valve surgery impacts mitral leaflets strain was undertaken in 2008-2014 by a joint team gathering UH mathematicians led by Prof. R. Azencott and Methodist Hospital cardiologists led by Dr W. Zoghbi, MD. This has provided the scientific context and main motivation for the research work presented in my PhD thesis. The purpose of this chapter is to introduce the notion of metric shapes spaces in anatomy that has been intensively explored for the past 15 years by the research

teams of L. Younes and M. Miller at John Hopkins University and A. Trounev at Ecole Normale Supérieure, Cachan, France. We briefly review the notions of metric distances between shapes in the spirit of the techniques used for instance at Johns Hopkins for the study of human brains, and we explore how these methods might be applied to our set of 60 dynamic models of mitral leaflets.

12.2 Geometric Traits as Markers of a Population

The 30 patient-specific models varied greatly in their size, shape, and features. Compared to the standards set by the normal models, the pre-operation organic models tend to exhibit larger areas, bulbous prolapsed tissues, and exotic folds and creases associated with leaflet flail. The post-operation models are much smaller in area (due to insertion of the artificial mitral ring) and the mitral annulus horn is flat. Even within populations, each mitral valve model is distinct. Certain normal models, for instance, have a high curvature mitral annulus horn while others may have a flatter, wider horn. See Figures [12.1]-[12.3] for some various patient-specific MVA models.

The previous chapters characterized the distributions of strain values on mitral leaflets for three different populations of mitral valves (normal, organic MR pre-surgery, organic post-surgery). There we saw that strain was statistically higher for the organic MR population than for the normal population. Moreover, we saw that NRD surgery reduced strain on the MVA back to levels comparable to that of the normal population. But a natural question follows: are there specific shape and deformations characteristics on the mitral valve that could indicate whether a patient

is at risk for future mitral valve disease? If so, are these characteristics linked at all to the visible geometry of the mitral valve? Similar questions have been asked by a team of computer scientists, mathematicians, and brain specialists at Johns Hopkins in their study of large sets of 3D-images of human brains and early detection of neurological disease [4], [13], [18], [19], [46]. The key idea is to widen the scope from merely looking at static landmarks on the mitral valve and instead compute “how far” mitral valves are from one another in the sense of the diffeomorphic matching (Chapter 4) required to map one valve to another valve.

12.3 Energy of Diffeomorphism as a Metric

The definition of adequate “distances” between diffeomorphisms of smooth shapes has a robust mathematical underpinning [46]. The methodology can be described heuristically using Figure [12.4] below. Suppose that D_1 is a diffeomorphism reconstructed by optimized geometric matching that deforms the left blue shape into the right blue shape. An intermediary shape snapshot (middle blue shape) is also shown. Suppose that D_2 is a separate diffeomorphism reconstructed by optimized geometric matching of the left red shape to the right red shape (the middle red shape is at the some intermediary time). Then a geometric matching can be performed between blue and red snapshots at each pair of corresponding instants. This yields diffeomorphisms F_1 that matches the left blue snapshot to the left red snapshot, F_2 that matches the middle blue snapshot to the middle red snapshot, and F_3 that matches the right blue snapshot to the right red snapshot. We could then define a distance

between these diffeomorphisms as

$$d(D_1, D_2) = \frac{1}{3} \sum_{i=1}^3 E(F_i)$$

where $E(F_i) = \int_0^1 \frac{1}{2} \|v(t)\|^2 dt$ is the kinetic energy of the diffeomorphism F_i (recall that the diffeomorphism F_t is determined by a vector field flow v_t of \mathbb{R}^3 as explained in Chapter 4). Technically, to be precise, the distance $d(D_1, D_2)$ would have to be in the limit of averages of kinetic energies $E(F_{t_k})$ as the number of intermediary instants t_k tends to infinity. However, for practical numerical shape matching studies, a small finite number of intermediary instants t_k is sufficient.

12.4 Distance between Mitral Valve Models

We apply these ideas to our patient-specific mitral valve dynamic deformation models. For simplicity, only geometric matching is optimized when the intermediary diffeomorphisms F_i are constructed. Consider a mitral valve model M_1 with deformation diffeomorphisms D_1 that deforms the model M_1 between mid-systole (MS) and end-systole (ES). Analogously, consider another mitral valve model M_2 with deformation diffeomorphisms D_2 that deforms the model M_2 between MS and ES. We seek diffeomorphisms F_1 and F_2 that respectively map M_1 to M_2 at MS and M_1 to M_2 at ES. But before any matching is done, these valve models M_1, M_2 must be aligned to a common center-point and rotationally aligned. For the studies here, the geometric centroids of the models were aligned. Next the orientations were fixed by

aligning first the principal planes of each model and then the vectors from the centroids to the anterior horns. Moreover, a constant value of $\sigma = 9$ and reference point grids of cardinality $N_0 = 200$ to $N_0 = 300$ were used in determining each F_1 and F_2 . The constant value of σ ensured that a consistent version of the kinetic energy (thus distance) was used for each pair of models M_1, M_2 . Once F_1 and F_2 were determined, we defined the distance between the two diffeomorphic deformations D_1 and D_2 by $d(D_1, D_2) = \frac{1}{2}(E(F_1) + E(F_2))$.

It is important to point out that this $d(D_1, D_2)$ is *not* simply the distance between two static mitral valves M_1, M_2 at a single instant. For instance, at MS, the distance given by $E(F_1)$ is the energy of the diffeomorphism F_1 , and it accounts for specific geometric differences in the shapes of the models M_1, M_2 at MS. The distance $d(D_1, D_2)$ accounts for differences between M_1, M_2 at *both* MS and ES. Thus the actual systolic phase of the cardiac cycle affects the distance $d(D_1, D_2)$ between valve models M_1, M_2 . For instance, two valves could have a very similar size and shape at MS but through any number of biological factors exhibit far different geometric similarity at ES. These valves, when viewed at the snapshot MS, would appear visually similar and the corresponding distance $E(F_1)$ would be relatively small. However, the distance $d(D_1, D_2)$ would account for the dissimilarity at ES, meaning that we would have $E(F_1) \leq d(D_1, D_2)$, which might be unintuitive if one solely looks at static snapshots of the valve. The *dynamic deformation* of the valve must be examined when interpreting the distance results given for mitral valve diffeomorphic deformation models.

12.5 Example for Normal Population

These techniques were applied to the 10 anterior leaflets in the normal population. Specifically, the distance $d(D_1, D_2)$ between each pair of anterior leaflet diffeomorphisms D_1, D_2 was computed for the entire 10 patient population. This results in 10×10 array given in Table [12.1]. Here we see that Patient 8 has the lowest mean distance to the other 9 patients. We designate this model as the *center* of the population.

Patient	1	2	3	4	5	6	7	8	9	10
1	0.00	4.74	3.51	5.92	6.56	8.25	4.64	3.69	5.75	7.07
2	4.74	0.00	5.40	5.36	4.70	4.29	2.71	2.13	3.59	1.00
3	3.51	5.40	0.00	4.68	4.35	7.54	2.80	2.91	3.10	5.02
4	5.92	5.36	4.68	0.00	7.30	5.57	3.91	3.00	4.03	3.07
5	6.56	4.70	4.35	7.30	0.00	9.87	6.15	2.97	7.95	9.17
6	8.25	4.29	7.54	5.57	9.87	0.00	7.61	4.32	6.52	5.32
7	4.64	2.71	2.80	3.91	6.15	7.61	0.00	3.93	2.18	2.02
8	3.69	2.13	2.91	3.00	2.97	4.32	3.93	0.00	5.17	4.41
9	5.75	3.59	3.10	4.03	7.95	6.52	2.18	5.17	0.00	4.69
10	7.07	1.00	5.02	3.07	9.17	5.32	2.02	4.41	4.69	0.00
Mean μ	5.01	3.39	3.93	4.28	5.90	5.93	3.59	3.25	4.30	4.18

Table 12.1: Distances $d(D_1, D_2)$ between normal patient anterior leaflets. Notice patient 8 has the lowest computed mean distance to all other patients. Contrast this to patient 6, which exhibits the largest mean distance to other patients.

To attempt to view these 10 leaflets in 3D Euclidean space, we adopt a principal component analysis (PCA) of the data presented in Table [12.1]. Specifically, denote the top section of data in Table [12.1] by the 10×10 matrix M . This data is *centered* by subtracting the mean value (row μ in Table) of each column; that is

$$U = M - \mu^T \otimes [1 \ 1 \ 1 \ 1 \ 1 \ 1 \ 1 \ 1 \ 1 \ 1]$$

where \otimes denotes the Kronecker product. The *covariance* K of this centered data is computed at $K = U^T U$ and PCA proceeds by finding the spectrum of eigenvalues and eigenvectors of K . Denote the 3 eigenvalues largest magnitude by $\lambda_1, \lambda_2, \lambda_3$ with corresponding eigenvectors $v_1, v_2, v_3 \in \mathbb{R}^{10}$. Finally, the columns k_j of the matrix K are projected onto the vectors v_i ; this yields a 3×10 matrix S with entries $s_{ij} = k_j^T v_i$. The plots of these 10 columns as points in \mathbb{R}^3 is given in Figure [12.5]. Note the quality of this projection of 10 dimensional data onto the 3 principal components can be quantified by the ratio

$$Q = \frac{\lambda_1 + \lambda_2 + \lambda_3}{\sum_{i=1}^{10} \lambda_i}$$

In the case of these 10 normal anterior leaflets, we have $Q = 0.8$. Also, neighboring points in Figure [12.5] can be connected by colored lines. For instance, the closest leaflet to patient 8 (yellow) is patient 2, and this is marked by yellow connecting line between patient 2 and patient 8. Proceeding in this manner reveals that there are two apparent clusters of normal leaflets: the first consists of patients 1,3,5,7,9 and the second consists of patients 2,4,6,8,10. Notice also that the center (patient 8) appears central to the point cloud; this is not surprising as it is our computed center of the population in Table [12.1].

12.6 Further Work

A similar analysis of the 10 diseased pre-surgery anterior leaflets will be completed. This involves computation of another 10×10 table of squared deformation distances between each pre-surgery anterior leaflet, as well as the above procedure to display these leaflets in \mathbb{R}^3 . Finally, we will attempt to generate a kernel-based non-linear separator to automatically distinguish between normal and abnormal anterior leaflets.

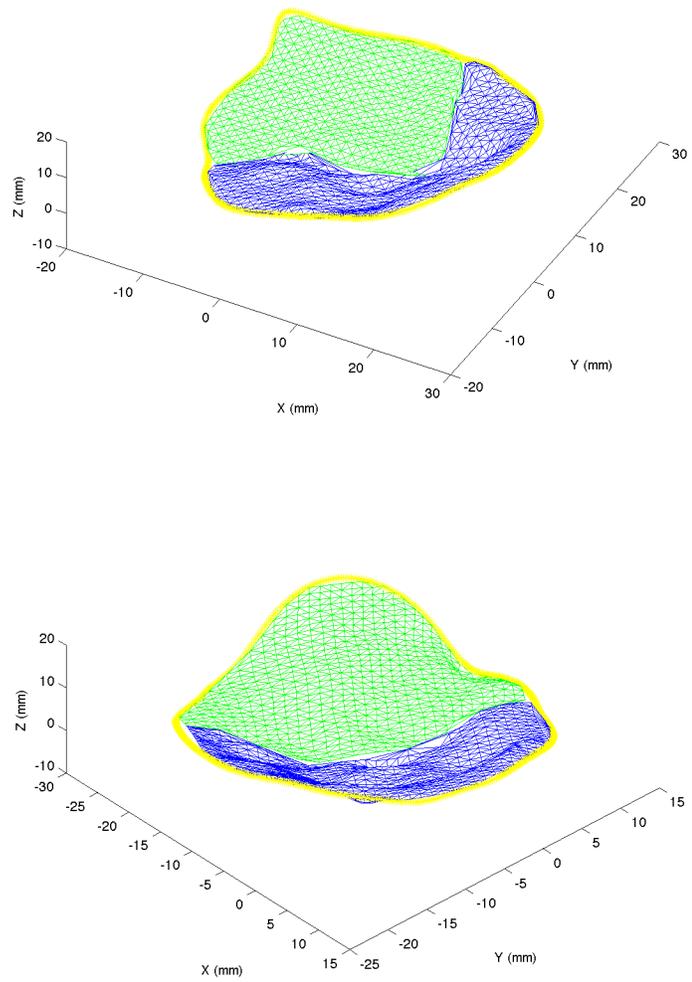


Figure 12.1: Two examples of normal patient mitral valve models.

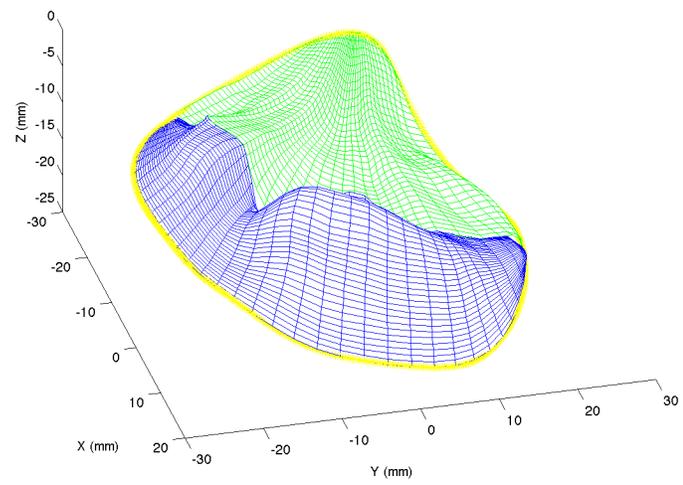
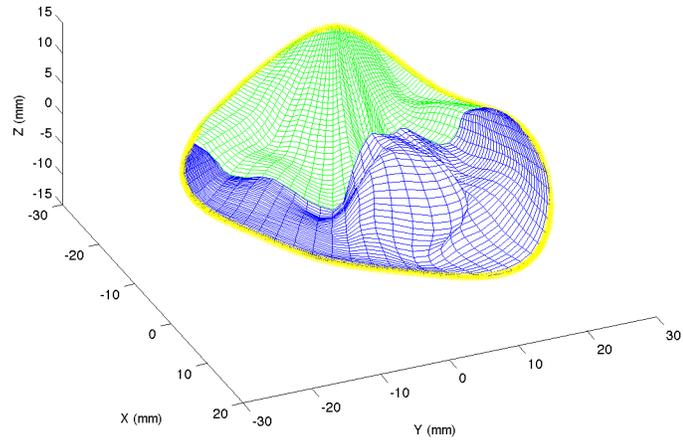


Figure 12.2: Two examples of organic pre-surgery mitral valve models

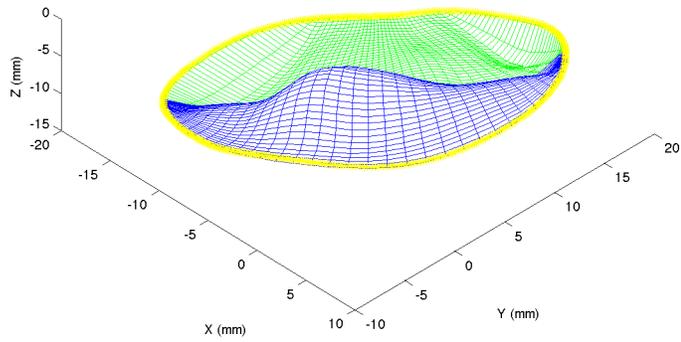
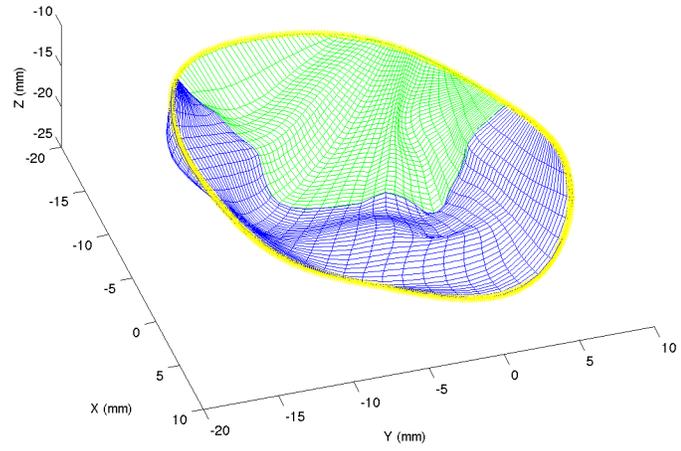


Figure 12.3: Two examples of post-surgery mitral valve models

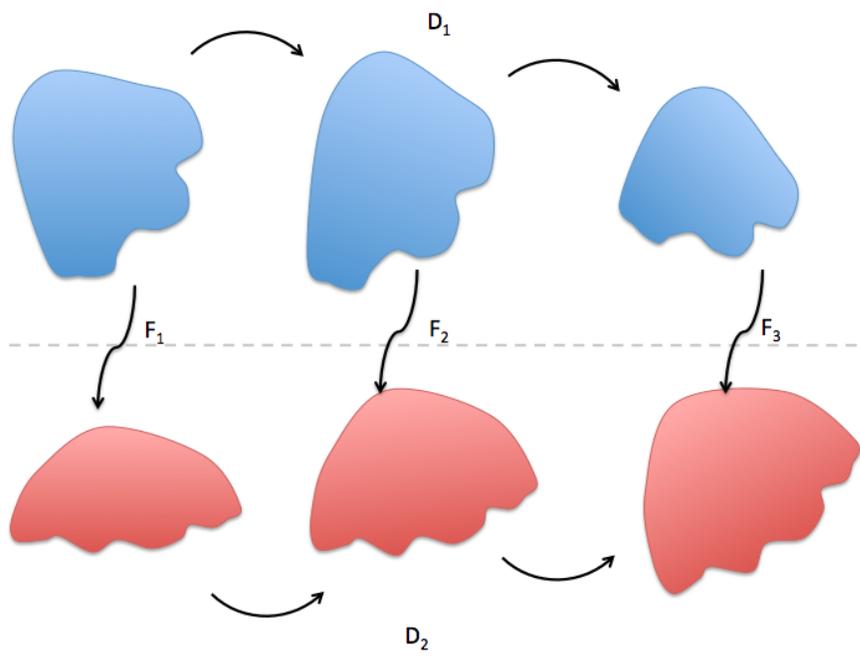


Figure 12.4: Distance between two diffeomorphisms D_1, D_2 given by the average of the kinetic energies $E(F_i)$ of the intermediary diffeomorphisms F_i .

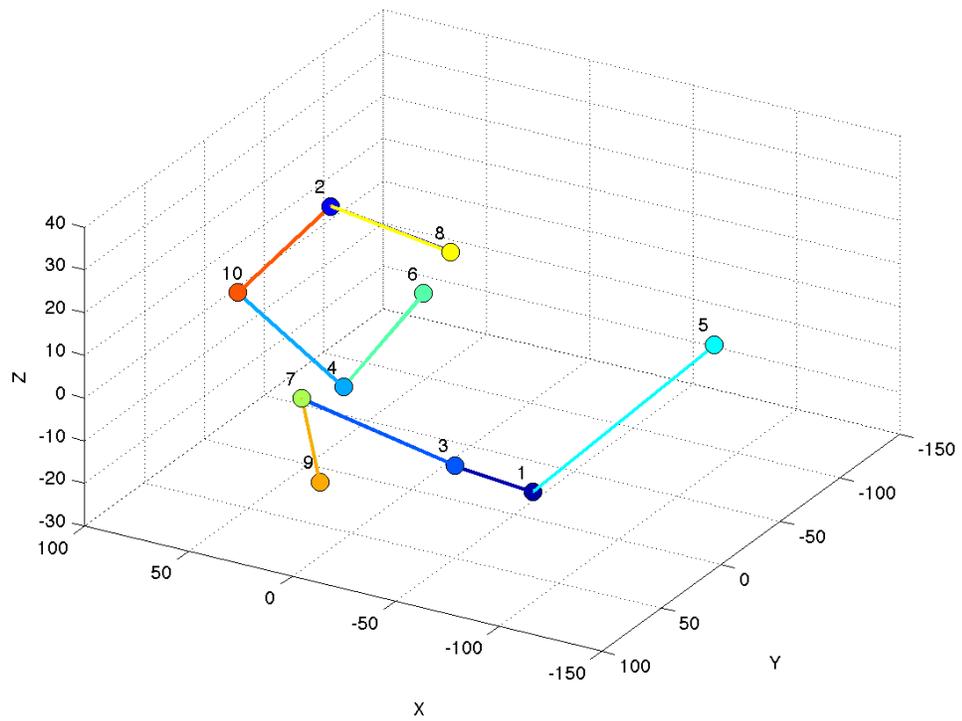


Figure 12.5: The 10 normal anterior leaflets are displayed in 3D Euclidean space via principal component analysis. Neighboring leaflets are connected by colored lines. For instance, the closest leaflet to patient 8 (yellow) is patient 2, and this is marked by yellow connecting line between patient 2 and patient 8.

Bibliography

- [1] E. Altiok MD, S. Hamada MD, S. van Hall, M. Hanenberg MD, G. Dohmen MD, M. Almalla MD, E. Grabskaya MD, M. Becker MD, N. Marx MD, and R.Hoffmann MD, “Comparison of Direct Planimetry of Mitral Valve Regurgitation Orifice Area by 3D Transesophageal Echocardiography to Effective Regurgitant Orifice Area Obtained by Proximal Flow Convergence Method”, *American Journal of Cardiology*, vol. 107 (3), pp 452-458, 2011
- [2] N. Aronszajn, “Theory of Reproducing Kernels,” *Trans. Amer. Math. Soc.*, **68**, pp 337-404, 1950
- [3] R. Azencott, R. Glowinski, J. He, R. Hoppe, A. Jajoo, Y. Li, A. Martynenko, R. Hoppe, S. Ben Zekry, S. Little and W. Zoghbi, “Diffeomorphic Matching and Dynamic Deformable Surfaces in 3D Medical Imaging,” *Computational Methods in Applied Mathematics*, Vol 10 Iss 3, pp. 235–274, 2010
- [4] M.F. Beg, M.I. Miller, A. Trounev, and L. Younes, “Computing large deformations metric mappings via geodesic flows of diffeomorphisms,” *Int. J. Comp. Vis.*, Vol. 61, pp. 139-157, 2005.
- [5] S. Ben Zekry, G. Lawrie, S. Little, W. Zoghbi, J. Freeman, A. Jajoo, S. Jain, J. He, A. Martynenko, R. Azencott, “Comparative Evaluation of Mitral Valve Strain by Deformation Tracking in 3D-Echocardiography”. *Cardiovascular Engineering and Technology*, vol 3, iss 4, pp 402-412, 2012
- [6] S. Ben Zekry, S. Little, W. Zoghbi, J. Freeman, R. Azencott, “Mitral Valve Strain Analysis Derived from 3-Dimensional Echocardiography Imaging: Normal Valves vs. Organic Mitral Regurgitation”. To be submitted to *Cardiovascular Engineering and Technology*.
- [7] M. Buhmann. *Radial Basis Functions: Theory and Implementations*, Cambridge University Press, 2003.

- [8] Y. Cao, M. Miller, S. Mori, R. Winslow, L. Younes, "Diffeomorphic matching of diffusion tensor images", *Computer Vision and Pattern Recognition Workshop*, 2006:67
- [9] S. Chandra, I.S. Salgo, L. Sugeng, L. Weinert, S.H. Settlemier, V. Mor-Avi, and R.M. Lang, "A three-dimensional insight into the complexity of flow convergence in mitral regurgitation: adjunctive benefit of anatomic regurgitant orifice area," *Am J Physiol Heart Circ Physiol*, vol. 301(3), pp. H1015-24, 2011
- [10] G. Dahlquist and A. Bjorck. *Numerical Methods in Scientific Computing, Volume 1*. SIAM, Philadelphia, PA, 2008.
- [11] M. de Berg, O. Cheong, M. Kreveld, M. Overmars. *Computational Geometry: Algorithms and Applications*. Springer-Verlag, 2008
- [12] C. de Boor. *A Practical Guide to Splines*. Springer-Verlag, New York, NY, 2001.
- [13] P. Dupuis, U. Grenander, and M. Miller, "Variational problems on flows of diffeomorphisms for image matching", *Quart. Appl. Math*, **56**, pp 587-600, 1998
- [14] A. Elen, H. F. Choi, D. Loeckx, H. Gao, P. Claus, P. Suetens, F. Maes, and J. D'hooge, "Three-dimensional cardiac strain estimation using spatio-temporal elastic Registration of Ultrasound Images", *IEEE Trans Med Imaging*, **27**(11):1580-1591, 2008
- [15] D. Filip, T. Ball, Procedurally representing lofted surfaces. *IEEE Comput. Graphics and Applications*, pp 27-33, 1989.
- [16] A. Gholipour, N. Kehtarnavaz, R. Briggs, M. Devous, and K. Gopinath, Brain Functional Localization: A Survey of Image Registration Techniques, *IEEE Transactions on Medical Imaging*, vol.26, pp: 427-451, 2007.
- [17] J. Glaunes, A. Qie, M. Miller, L. Younes, "Large deformation diffeomorphic metric curve mapping" *Int. J. Comp. Vision*, **80**, pp 317-336, 2008
- [18] J. Glaunes, A. Trounev, and L. Younes, "Diffeomorphic matching of distributions: A new approach for unlabelled point-sets and sub-manifolds matching", *2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, Vol 2, pp 712-718, 2004
- [19] J. Glaunes, M. Vaillant, and M. Miller, "Landmark matching via large deformation diffeomorphisms on the sphere", *J. Math. Imaging Vis.*, **20** pp 179-200, 2003

- [20] H. Guo, A. Rangarajan, S. Joshi, and L. Younes, "Non-rigid registration of shapes via diffeomorphic point matching," *Proc. Int. Symp. on Biomed. Imaging: From Nano to Macro, Washington, D.C.*, pp. 924-927, 2004.
- [21] Z. He et al. "In vitro dynamic strain behavior of the mitral valve posterior leaflet". *J Biomech Eng.* 2005;127(3):504-511.
- [22] B. Heyde, S. Cygan, H. F. Choi, B. Lesniak-Plewinska, D. Barbosa, A. Elen, P. Claus, D. Loeckx, K. Kaluzynski, J. D'hooge, "Three-dimensional cardiac motion and strain estimation", *IEEE International Ultrasonics Symposium Proceedings*, pp 1534-1537, 2010
- [23] X. Huang, "Shape Registration in Implicit Spaces using Information Theory and Free Form Deformations"
- [24] A. Jajoo, "Control of the acceleration field for the diffeomorphic matching of multiple 2D shapes in R³", Johns Hopkins Shape Meeting, 2013.
- [25] S. Joshi and M. Miller, "Landmark matching via large deformation diffeomorphisms," *IEEE Trans. Image Processing*, Vol. 9, pp. 1357-1370, 2000.
- [26] G. Krishnamurthy, D.B. Ennis, A. Itoh, W. Bothe, J. C. Swanson, M. Karlsson, E. Kuhl, D.C. Miller and N.B. Ingels, Jr., *Material properties of the ovine mitral valve anterior leaflet in vivo from inverse finite element analysis*, *Am J Physiol Heart Circ Physiol*, 295:1141-1149, 2008
- [27] J. Krucker et al. "Rapid Elastic Image Registration for 3D ultrasound", *IEEE Transactions on Medical Imaging*, vol.21, pp:1384-1394, 2002.
- [28] P. Kundu and I. Cohen, *Fluid Mechanics*, 4th ed. Academic Press. 2010.
- [29] K.S. Kunzleman and R.P. Cochran, "Stress/Strain Characteristic of Porcine Mitral Valve Tissue: Parallel Versus Perpendicular Collagen Orientation", *J Card Surg*, 7(1):71-78, 1992
- [30] K.S. Kunzleman, M.S. Reimink and R.P. Cochran, "Annular dilatation increases stress in the mitral valve and delays coaptation: a finite element computer model", *Cardiovasc Surg.* 5:427-34, 1997.
- [31] G. Lawrie, "Mitral valve: toward complete repairability", *Surg Technol Int*, 15:189-197, 2006
- [32] G. Lawrie, "Mitral valve repair vs replacement", *Cardiol Clin*, 16(3):437-448, 1998

- [33] E. Lehmann, H. D. Abrera, *Nonparametrics: Statistical Methods Based on Ranks*
- [34] J. Maintz and M. Viergever, A survey of medical image registration, *Medical Image Analysis*, vol. 2, pp: 1-36, 1998.
- [35] D. Panchenko. 18.443 Statistics for Applications, Fall 2006. (MIT OpenCourseWare: Massachusetts Institute of Technology), <http://ocw.mit.edu/courses/mathematics/18-443-statistics-for-applications-fall-2006> (Accessed 13 Feb, 2014). License: Creative Commons BY-NC-SA
- [36] N. Paragios, "On the Representation of Shapes Using Implicit Functions,"
- [37] B. Payne and A. Toga, "Distance Field Manipulation of Surface Models," *IEEE Computer Graphics and Applications*, Vol. 12, No. 1, pp. 65-71, 1992
- [38] D. Rueckert and P. Aljabar, Nonrigid Registration of Medical Images: Theory, Methods, and Applications, *IEEE Signal Processing Magazine*, vol. 27, pp: 113-119, 2010.
- [39] M. Sacks et al. "Surface strains in the anterior leaflet of the functioning mitral valve". *Ann Biomed Eng.* 2002;30(10):1281-1290.
- [40] M. Enriquez-Sarano MD, L. J. Sinak MD, A. J. Tajik MD, K. R. Bailey, and J. B. Seward MD, "Changes in Effective Regurgitant Orifice Throughout Systole in Patients With Mitral Valve Prolapse", *Circulation*, vol. 92, pp. 2951-2958, 1995
- [41] I. J. Schoenberg. Contributions to the problem of approximation of equidistant data by analytic functions. *Quart. Appl. Math.*, 4:45-99 and 112-141, 1946.
- [42] P. Shearer, *Introduction to Seismology*, 2nd ed. Cambridge University Press. 2009
- [43] R. Shekhar, V. Zagrodsky, M. Garcia and J. Thomas, "Registration of Real-Time 3D Ultrasound Images of the Heart for Novel 3D Stress Echocardiography", *IEEE Transactions on Medical Imaging*, vol. 23, pp:1141-1149, 2004.
- [44] M. Stevanella et al. "Mitral leaflet modeling: Importance of in vivo shape and material properties". *J Biomech.*44(12):2229-2235.
- [45] G. Turk and J. O'Brien, "Shape Transformation Using Variational Implicit Functions," *SIGGRAPH '05 ACM SIGGRAPH 2005 Courses*, Article No. 13, 2005
- [46] L. Younes, "Computable elastic distances between shapes", *SIAM J. Appl. Math.*, **58**, pp 565-586, 1998