

DEFORMATION TRACKING AND Diffeomorphic MATCHING IN MEDICAL IMAGING

A Dissertation

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

Yipeng Li

December 2010

DEFORMATION TRACKING AND DIFFEOMORPHIC MATCHING IN MEDICAL IMAGING

Yipeng Li

APPROVED:

Prof. Robert Azencott, Chairman
Dept. of Mathematics

Prof. Bernhard Bodmann
Dept. of Mathematics

Prof. Manos Papadakis
Dept. of Mathematics

Prof. Ioannis Pavlidis
Dept. of Computer Science

Dean, College of Natural Sciences and Mathematics

DEFORMATION TRACKING AND DIFFEOMORPHIC MATCHING IN MEDICAL IMAGING

An Abstract of a Dissertation
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
Yipeng Li
December 2010

Abstract

This thesis focuses on deformation tracking and diffeomorphic matching in 2D or 3D medical imaging. We have developed and implemented novel algorithmic approaches to track both the myocardial motion in 2D IntraCardiac Echocardiography images and the Mitral Valve deformation in 3D ultrasound images.

We studied the myocardial deformation through the registration of IntraCardiac Echocardiography images, by utilizing the statistical properties of the speckle noise. Within a framework of parametric elastic registration, the deformation was estimated by optimizing the energy function computed by applying the maximum likelihood approach to the speckle noise model. And, the optimization procedure is accelerated by applying the multiresolution method and local masking, which also improve the robustness of our algorithm. Numerical results are verified perfectly by the medical data from Texas Medical Center.

We then consider optimal matching of submanifolds such as curves and surfaces by a variational approach based on Hilbert spaces of diffeomorphic transformations. In an abstract setting, the optimal matching is formulated as a minimization problem, in which the objective functional consists of two parts measuring the elastic energy of the dynamically deformed surfaces and the quality of the matching. The performance of our diffeomorphic tracking in medical 3D-image movies is illustrated by numerical results for the dynamic modeling of the human mitral valve annulus by computer analysis of 3D-echocardiographic image sequences.

Contents

1	Introduction	1
2	Ultrasound imaging	6
2.1	Basic Principles of Ultrasound	6
2.2	B-mode Ultrasound Imaging	8
2.3	Ultrasound Image Artifacts	9
2.3.1	The statistics of speckle	10
3	Historic review	11
3.1	Image registration	11
3.2	Elements of registration	12
3.2.1	Feature space	12
3.2.2	Search space	14
3.2.3	Similarity metrics	15
3.3	Registration of 2D-ICE images	17
3.3.1	Medical Goal	18
3.3.2	B-Splines approximation	19
3.3.3	Heart Anatomy and Mechanism	24
3.4	Diffeomorphic matching of deformable shapes in 3D ultrasound imaging	29
4	Multiscale Speckle Tracking in Intracardiac Echocardiography	33

4.1	Methodology	34
4.1.1	Stochastic disparity cost functional	34
4.1.2	Regularization cost functional	35
4.1.3	Estimation of the speckle noise	35
4.1.4	Image specific speckle noise modelling	36
4.1.5	Minimizing image specific cost functionals	36
4.1.6	Multiresolution and local Masking	37
4.1.7	Validation against experimental heart strain data	37
4.2	Mathematical definition of the cost functionals	38
4.2.1	Cost Functionals	38
4.2.2	B-mode ultrasound display and true Tissue Echogenicity . . .	38
4.2.3	Maximum Likelihood approach	40
4.2.4	Regularization term K_{reg} and final cost functional K	43
4.3	Multi-resolution optimization of the cost functionals	44
4.3.1	B-spline model for the displacement vector field u	44
4.3.2	Multi-resolution optimization of cost functionals	45
4.3.3	Detailed gradient descent at fixed resolution	47
4.3.4	Image sequence analysis and Masking	48
4.3.5	Estimation of the log-speckle noise and of the re-scaling parameter	50
4.4	Detailed analysis of the disparity term K_{disp} derived from log-likelihood	51
5	Numerical performance of speckle tracking	59
5.1	Numerical performances on simulated noise corrupted and warped images	59
5.1.1	Three simulated pairs of images	59
5.1.2	Performances of deformation estimates on simulated cases . .	61
5.1.3	The case of affine image intensities	64

5.2	Estimation of speckle noise characteristics on real ICE image sequences	65
5.3	Time modulations of log-speckle and ICE image intensities	69
5.4	Estimated strains versus sonomicrometric ground truth	70
5.5	Computing time optimization	71
6	Speckle tracking in 3D ultrasound images	80
6.1	Mitral Valve	81
6.2	B-spline model for the displacement vector field u	83
6.3	Multi-resolution optimization of cost functionals	84
6.4	Detailed gradient descent at fixed resolution	85
6.4.1	Image sequence analysis and Masking	87
6.5	Numerical performance on 3D ultrasound images	88
7	Diffeomorphic Matching and Dynamic Deformable Surfaces in 3D Medical Imaging	97
7.1	Introduction	98
7.2	Diffeomorphic Shape Matching	99
7.2.1	Diffeomorphic Matching of Two Shapes in R^3	99
7.2.2	Variational Approach	101
7.2.3	Multiple Snapshots of Dynamic Deformable Shapes	103
7.3	Optimal Matching of Intermediary Snapshots	104
7.3.1	Time Dependent Vector Fields With Finite Kinetic Energy . .	104
7.3.2	Dynamic System of Diffeomorphic Deformations	105
7.3.3	Self-Reproducing Hilbert Spaces	105
7.3.4	Matching Quality	106
7.4	Diffeomorphic point matching	110
7.5	Numerical solution of diffeomorphic point matching	113
7.5.1	The basic matching algorithm	114

7.5.2	Smoothing of the Hausdorff matching term	117
7.5.3	Choice of scale parameters	118
7.5.4	Continuation in the regularization parameter	119
7.6	Numerical Results for the Human Mitral Valve Annulus	120
7.6.1	Diffeomorphic Matching for 5 Annulus Snapshots	121
7.6.2	Diffeomorphic Matching for All Annulus Snapshots in One Cardiac Cycle	124
7.6.3	Diffeomorphic Matching of the Anterior Leaflet	125
7.6.4	Diffeomorphic Matching for Multiple Snapshots of the Poste- rior Leaflet	126
8	Conclusion	139
	Bibliography	143

List of Figures

3.1	The centered spline of degree 3: cubic B-splines	20
3.2	Cardiac cycle of left ventricle	26
3.3	Example of Intracardiac Echocardiography (ICE)	28
5.1	Differently smoothed test images (Top: case 1, bottom: case 3)	62
5.2	Warping field Example	63
5.3	Speckle noise std distribution(top) and mean intensity(bottom) within the myocardium and heart chamber: at the beginning of systole . . .	74
5.4	Speckle noise std distribution(top) and mean intensity(bottom) within the myocardium and heart chamber: at the end of systole	75
5.5	Different regions to compute speckle noise variance	76
5.6	Empirical pdf v.s. theoretical pdf on different regions: region 8 (top) and region 11 (bottom)	77
5.7	Examples of a small mask on real ICE image	78
5.8	Examples of deformation vector field	79
6.1	Example: 2D slices of 3D images	90
6.2	Mitral valve illustration	91
6.3	Mitral valve: the middle line is the coaptation line separating the two surfaces. The left surface in the figure is the anterior leaflet and the right surface is the posterior leaflet. The black thick curve is the annulus	92
6.4	Examples of mask in 3D image	93

6.5	Evolution of cost function in optimization	94
6.6	Histogram of distances between reference and target annulus	95
6.7	Histogram of distances between deformed and target annulus	96
7.1	The five closed curves are five instances of the annulus. The small dotted curve is the reference curve. From bottom to top, the next three curves are the intermediaries and the last one is the final target. The vertical ‘-’ lines are deformations of selected points on the reference curve.	123
7.2	Diffeomorphic matching for multiple mitral annulus snapshots, using smoothed Hausdorff distances: Convergence history for the geometric accuracy indicators and for the smoothed Hausdorff disparity components $Disp_j^G$	124
7.3	Diffeomorphic matching for multiple annulus snapshots, using smoothed Hausdorff distances: Pareto frontiers for the geometric accuracy indicators, for the Hausdorff disparities, and for the global Hausdorff disparity	128
7.4	Reference(solid) and deformed (dash) annulus intensity curve by arc length: top($\mu = 0$), bottom(best μ), data set D_1	129
7.5	Reference(solid) and deformed (dash) annulus intensity curve by arc length: top($\mu = 0$), bottom(best μ), data set D_2	130
7.6	Reference(solid) and deformed (dash) annulus intensity curve by arc length: top($\mu = 0$), bottom(best μ), data set D_3	131
7.7	Reference(solid) and deformed (dash) annulus intensity curve by arc length: top($\mu = 0$), bottom(best μ), data set D_4	132
7.8	Reference(solid) and deformed (dash) annulus intensity curve by arc length: top($\mu = 0$), bottom(best μ), data set D_5	133
7.9	Reference(solid) and deformed (dash) annulus intensity curve by arc length: top($\mu = 0$), bottom(best μ), data set D_6	134
7.10	Reference(solid) and deformed (dash) annulus intensity curve by arc length: top($\mu = 0$), bottom(best μ), data set D_7	135
7.11	Reference(solid) and deformed (dash) annulus intensity curve by arc length: top($\mu = 0$), bottom(best μ), data set D_8	136

7.12	Diffeomorphic matching for four anterior leaflet snapshots: Pareto frontiers for the separate Hausdorff disparities to snapshots (left) and for the global Hausdorff disparity (right)	137
7.13	Pareto frontiers for the maximum distances to snapshots (left) and for their 90th percentiles (right)	137
7.14	Diffeomorphic matching for four posterior leaflet snapshots: Geometric accuracy indicators and Pareto frontiers for strictly equal regularization weights λ_j	138
7.15	Diffeomorphic matching for four posterior leaflet snapshots: Geometric accuracy indicators and Pareto frontiers for dynamically adjusted regularization weights λ_j	138

List of Tables

5.1	Angle error (in degrees)	63
5.2	Relative magnitude error (in percents)	63
5.3	Errors with different α	65
5.4	Statistical properties of different regions	67
5.5	Correlation Coefficient	68
5.6	Computing time per frame(in minutes)	73
7.1	$mean(IME)$ when $GME \approx 1$.	126

Chapter 1

Introduction

The goal of image registration is to determine a spatial transformation that will bring into correspondence (i.e. register) pairs of homologous points in 2 given images. In the simplest cases, the mathematical form of the desired spatial transformation can be restricted by simple physical principles. For example, when registering images acquired from the same subject, it is often possible to assume that the body part being imaged can be treated as a rigid body, which leads to a highly constrained spatial transformation model. Unfortunately, physical processes involved in the acquisition and reconstruction of medical images can cause artifacts and lead to violations of the rigid body model, even when the object being imaged adheres strictly to rigid body constraints. Potential sources of such distortions are prevalent in magnetic resonance (MR) and positron emission tomography (PET) images, also distortions of soft tissues can also lead to nonlinear effects that violate rigid body assumptions. Registration of images acquired from different subjects represents the extreme end of

the spectrum, where developmental factors including genetics, environment, and random influences all contribute to the complex differences between subjects. Chapter 3 gives a historic review of the image registration methods.

2D-ICE is a primary echocardiography modality to acquire visual information on live cardiac anatomy [43][44][45][46][47][48]. However, ICE-images are always perturbed by a strong “speckle” noise, due to diffuse scattering of ultrasound pulses with randomly dispersed small scatterers, at sound wavelength scales. Speckle noise is non-gaussian with a high noise-to-signal ratio. Speckle decorrelation from frame to frame is a standard assumption in this context, as well as for different pixels within a single frame.

Speckle tracking of heart motion focuses on estimation of the non-linear and often large deformations f_t of the heart muscle between a reference image frame J_r and any other image frame J_t in the same heart cycle, in the presence of speckle noise. Several “speckle tracking” algorithms [62][63][64][65][66] have been applied to heart motion recovery in cardiac echocardiography, including optical flow, block matching, and elastic registration.

Classical optical flow approaches [49] [50] [51] have fairly low performance. Indeed they assume pixel intensities to remain constant between consecutive image frames, and hence are highly sensitive to speckle noise. For the block matching techniques [52] [53] [54] [55] [56], capture ability varies with block size, and it is quite difficult to select an optimum block-size that captures both large and small deformations simultaneously. Elastic registration algorithms [57] [58] have performed better for myocardial motion detection, resulting in good spatial capture ability for non-linear

deformations.

Most published cardiac motion detection methods still do not attempt to fully use the statistical features of speckle noise. However a few recent studies [53] [54] [59], have explored maximum likelihood techniques to extract cardiac motion, incorporating a well documented theoretical stochastic model for ultrasound speckle noise. In [59] an elastic registration approach was applied to ICE image data, to compute cardiac motion by minimization of a cost function derived from maximum likelihood principles, and validated the technique by comparison with experimental cardiac strain data acquired by recording displacements of microcrystals surgically inserted in the hearts of animal subjects by Dr D. Khoury’s team at Methodist Hospital.

In chapter four, we have undertaken to generalize and deepen the method implemented in [59], the main principles of our approach are: stochastic disparity cost functional, estimation of the speckle noise, image specific speckle noise modeling and multiresolution and masking.

Clinical diagnosis and therapy planning are increasingly often supported by 3D-imaging modalities, such as MRS (Magnetic Resonance Spectroscopy), PET (Positron Emission Tomography), SPECT (Single Photon Emission Computed Tomography) for functional information, and CT (Computed Tomography), MRI (Magnetic Resonance Imaging), Ultrasound Echography , X-Ray, for anatomical visualization. Thus clinicians and medical researchers become natural users for automated 3D-image registration providing voxel to voxel matching of two 3D-images of the same anatomical object obtained by different imaging modalities, at different times, or

from different perspectives. The search for a good voxel to voxel correspondence between reference and target images J_{ref} and J_{tar} , is guided by one or several matching quality criteria. Image matching is generally achieved by an R^3 -diffeomorphism F matching two given bounded subdomains of the 3D-voxel grid, and can thus be assigned an elastic energy $EE(F)$ measuring the amount of spatial deformation implemented by F . In 2D or 3D-image registration, typical matching quality criteria involve the differences $difint(z, z') = |J_{tar}(z') - J_{ref}(z)|$ in image intensities at all pairs (z, z') of matched voxels. The Intensity Matching Cost $IMC(F)$ is often defined by fixing some exponent $a > 0$ and summing $difint^a(z, z')$ over all voxels z belonging to the domain of interest in J_{ref} . The search for an optimal registration then becomes a variational problem where one seeks a deformation F minimizing a linear combination of $EE(F)$ and $IMC(F)$. Image registration methods were initially designed for 2D-images, for instance to align tomographic slices of different recordings, but in the last decade, 3D-image registration based on volumetric data sets has become the main technical challenge, and involves much heavier computing resources. Surveys of image registration algorithms can be found in [76, 89, 100].

The mitral valve is a dual-flap valve in the heart that lies between the left atrium (LA) and the left ventricle (LV). The mitral valve and the tricuspid valve are known collectively as the atrioventricular valves because they lie between the atria and the ventricles of the heart and control the blood flow. A normally-functioning mitral valve opens due to increased blood pressure from the left atrium as it fills with blood. As the pressure increases above that of the left ventricle, the valve opens allowing blood to flow into the left ventricle during diastole (early rapid filling and atrial

contraction), and closes at the end of atrial contraction to prevent blood flowing back.

Disorders of the mitral valve are the second most frequent heart problems, cumulating 14 percent of total number of deaths caused by Valvular Heart Disease each year in the United States and require elaborate clinical management. Visual and quantitative evaluation of the valve is an important step in the clinical workflow according to experts as knowledge about mitral morphology and dynamics is crucial for interventional planning [30][31]. In chapter 6, we apply variational techniques based on Hilbert spaces of diffeomorphic transformations to numerically construct diffeomorphic flows modeling the dynamic deformations between multiple 3D-snapshots of the human mitral valve apparatus. In an abstract setting, the optimal matching is formulated as a minimization problem involving actions of diffeomorphisms on regular Borel measures considered as support measures of the reference and the target submanifolds. The objective functional consists of two parts quantifying the elastic energy of the dynamically deformed surfaces and the quality of the geometric surface matching. To make the problem computationally accessible, we use reproducing kernel Hilbert spaces with radial kernels and weighted sums of Dirac measures which gives rise to diffeomorphic point matching and amounts to the solution of a finite problem in very high dimensions.

Chapter 2

Ultrasound imaging

2.1 Basic Principles of Ultrasound

Diagnostic ultrasound employs pulsed, high frequency sound waves to image tissue structures and their motion [37]. An ultrasound wave is a form of mechanical energy that propagates through a medium by compression and rarefaction. Ultrasound wavelength determines the spatial resolution achievable along the direction of the beam. A high-frequency ultrasound beam (small wave-length) provides superior resolution and image detail compared with a low-frequency beam.

The emitted ultrasound pulse is the impulse function of the system. When it represents the output of the ultrasound system during interrogation of an ideal point target, the echo pulse is also known as the system's point spread function (PSF). The term PSF is often used to refer to two-dimensional representations of the system

response in pressure amplitude versus space, with the implicit understanding that the actual response has three-dimensional extent.

In ultrasound system, the axial, lateral, and elevational (slice thickness) dimensions determine the spatial resolution and visibility of the system.

Axial resolution defines the ability of the ultrasound pulse to differentiate between two closely spaced objects that lie along the axis of an ultrasound beam. Lateral (azimuthal) resolution defines the ability to resolve adjacent objects perpendicular to the beam direction and is determined by the beam width (diameter). Elevational resolution is dependent on the transducer element height, and is perpendicular to the image plane.

Ultrasound interactions with matter are determined by the acoustic properties of matter. As ultrasonic waves propagate through a medium, some effects that are observed, including reflection, refraction, scattering and attenuation [37].

Sound reflection occurs at tissue boundaries on a scale much larger than the wavelength, and with differences in acoustic impedance. Scattering refers to the interaction of the ultrasound wave with microstructures that are much smaller than its wavelength. There are two types of scattering in human bodies, arising from the spatial arrangement of the scatterers in the ultrasound resolution cell. If the scatterers have a periodic arrangement, a coherent scattering is introduced, producing periodicity in the echo spectrum. If the arrangement of these scatterers is spatially random, the resulting diffuse scattering gives rise to speckle [38]. The texture of the observed speckle pattern does not correspond to underlying structure. However,

the local brightness of the speckle pattern does reflect the local echogenicity of the underlying scatterers. Attenuation is primarily due to the inner friction or viscosity of the tissue that transforms sound energy into other energy forms such as heat. Signal attenuation depends highly on the carrier frequency. Higher frequencies allow a better spatial resolution, but are more attenuated than lower ones and thus have less penetrating ability.

2.2 B-mode Ultrasound Imaging

The basic principle of ultrasound imaging is to emit pulses, and to collect reflected echoes. In the imaging system, the strength or amplitude of each reflected wave is represented by a dot. These dots are combined to form a complete image. The brightness of the dot represents the strength of the returning echo. The position of the dot represents the depth from which the returning echo was received.

The detected echoes may be displayed in one-dimensional formats such as amplitude mode (A-mode), brightness mode (B-mode) or motion mode (M-mode) formats. B-mode is the electronic conversion of the A-mode. In this mode, a line of brightness-modulated dots is displayed, and the line represents the orientation of the transducer.

An ultrasonic scanner generally operates in B-mode, and presents a gray-scale image that represents a spatial map of echo amplitude. In the B-mode image, white dots represent strong reflections, e.g., the reflection caused by diaphragm, gallstones and bones; grey dots denote weaker reflections, e.g., solid organs and thick fluid; and

black dots indicate no reflection, e.g., fluid within a cyst.

2.3 Ultrasound Image Artifacts

Ultrasound image artifacts arise from an incorrect display of anatomy or noise during imaging. Incorrect anatomical imaging can cause shadowing, reverberation, and speed displacement artifacts. However, in this study, we are mainly concerned with the system noise artifact, called speckle. Speckle has a textured appearance that results from the juxtaposition of small, closely-spaced structures that are too small to be resolved by the PSF. Speckle therefore is the result of diffuse scattering, and it can be considered as an inherent property of the ultrasound image. Speckle generally does not reflect the structure of the underlying tissue. The regional mean brightness of texture pattern, however, reflects the regional echogenicity of tissue. Therefore, speckle can be considered as noise, since it may obscure structures in the medium under observation.

Speckle is present in both RF data and envelope-detected data, and the texture of the observed speckle pattern does not correspond to underlying structure. While the mean speckle brightness at each region of the image reflects the original echogenicity map, the speckle noise itself does not reflect the underlying tissue structure, and provides practically no information on the detailed echogenicity map or the corresponding scattering function.

2.3.1 The statistics of speckle

Given the stochastic nature of speckle noise, we must describe this noise pattern statistically to draw general conclusions about imaging systems. The statistics used here to describe ultrasound speckle are drawn from the literature of laser optics. Each one of the diffuse scatterers in the isochronous volume contributes a small component to the echo signal, which is then modeled as the position of a random walk in the complex plane. The steps of this random walk are considered independent random variables, and one can apply the Central Limit Theorem to their sum. Therefore, in fully developed speckle, this complex radio-frequency echo signal from diffuse scatterers alone has a zero mean, two-dimensional Gaussian probability density function (PDF) in the complex plane.

Envelope detection removes the phase component, creating a signal with a Rayleigh amplitude PDF:

$$p_A(s) = \frac{s}{\sigma^2} \exp\left(-\frac{s^2}{2\sigma^2}\right), \quad a \geq 0 \quad (2.1)$$

Speckle brightness is greater if there are fewer, longer steps in the random walk than if there are many shorter steps. This could be accomplished by improving the spatial resolution of the system. On the other hand, if the scatterer density is doubled, a $\sqrt{2}$ increase in brightness results.

Chapter 3

Historic review

3.1 Image registration

The goal of image registration is to determine a spatial transformation that will bring into correspondence (i.e. register) pairs of homologous points in 2 given images [18][19][20][21][22][23][24]. In the simplest cases, the mathematical form of the desired spatial transformation can be restricted by simple physical principles. For example, when registering images acquired from the same subject, it is often possible to assume that the body part being imaged can be treated as a rigid body, which leads to a highly constrained spatial transformation model. Unfortunately, physical processes involved in the acquisition and reconstruction of medical images can cause artifacts and lead to violations of the rigid body model, even when the object being imaged adheres strictly to rigid body constraints. Potential sources of such distortions are prevalent in magnetic resonance (MR) and positron emission

tomography (PET) images, also natural deformations of soft organs can also lead to nonlinear effects that violate rigid body assumptions. Registration of images acquired from different subjects represents the extreme end of the shape deformation spectrum, where developmental factors including genetics, environment, and random influences all contribute to the complex shape differences between subjects.

3.2 Elements of registration

Existing registration algorithms can be classified by three elements: feature space, search space and similarity metrics.

3.2.1 Feature space

In a preliminary step, some intermediate data are extracted from the two images being registered. These data live in a feature space, which can be pixel-based, transform-based, or feature-based. Pixel-based algorithms work directly with the pixel values of the images being registered. Preprocessing is often used to suppress the adverse effects of noise and differences in acquisition [10], or to increase or uniformize pixel resolution [11]. It is possible to work directly with the pixel values on the discrete coordinate grid. However, to get a subpixel resolution, the registration problem is often cast into the continuous framework. The images are considered as functions of real arguments: the pixels (or voxels) coordinates. The correspondence

between the discrete and continuous versions of the image is established using interpolation. The crudest interpolation method is the nearest-neighbor, and the most often used one is linear (resp. bi- or trilinear) interpolation. Among the high-end methods, spline interpolation [60][61][9] provides the best tradeoff between accuracy and computational cost [1][2]. Transform-based algorithms exploit properties of the Fourier, wavelet, Hadamard, and other image transforms, making use of the fact that certain deformations manifest themselves more clearly in the transform domain. These methods are used mainly in connection with linear deformation fields. Nevertheless, there are examples of methods that estimate locally linear optical flow using Gabor filters [3][4] and B-spline wavelets [5]. Typical characteristics of the image transforms employed are linearity and independence of the actual image contents. Feature-based algorithms work on a set of characteristic features extracted from the images. The dimensionality of the features is usually drastically smaller than the dimensionality of the original image data. The extraction process is highly non-linear, mostly using thresholding.

In chapter three, registration of 2D-ICE images was studied by working directly with (all) the pixel intensity values on the discrete grid. To improve both the accuracy and computing efficiency, specific masks on the image were applied both in the objective cost function and the optimization process, the importance of which was also identified by the statistical analysis of the speckle noise distributions. In chapter 5, registration of 3D ultrasound images is based on the combination of selected landmark features and of voxel intensity values. The landmark points were interactively tagged on 3D medical images by experienced cardiology experts, a meticulous task

which cannot generally be replaced by automatic preprocessing algorithms.

3.2.2 Search space

One of the important factors to categorize registration algorithms is the search space including all warping (or deformation) functions that are candidate solutions of the registration problem. These classes of deformation (warping) functions thus model the sought for correspondence between the two images to register. The deformation models are strongly characterized by their number of parameters and the spatial extent of the image area influenced by a single parameter.

Parametric, global methods describe the registration mapping using a global deformation model with a relatively small number of parameters. The model mostly consists in expressing the warping function as a linear combination of basis functions, which can be for instance a polynomial basis or a harmonic basis. For these methods, the deformation model corresponding to a specific warp space is as important as the matching criterion being optimized.

Semi-local models use a moderate number of parameters with local influence. A grid of control points is placed over the image. Their spacing corresponds loosely to knot or landmark density.

Finally, there are non-parametric, local methods. The deformation function

sought after is basically unconstrained, or belongs to a very large and unrestricted functional space, e.g., the Sobolev space W_2^2 of twice differentiable functions. We seek the values of this deformation on a very fine grid, usually coinciding with pixel locations. These methods are often formulated as variational techniques, which seek to minimize a scalar functional cost criterion, or (more generally) characterize the deformation function as solutions of partial differential equations (PDE) driven by the image data. The essence of these methods thus lies almost entirely in the cost criterion to minimize or the PDE to solve. Such PDEs have been derived from the optical flow approach (gradient methods) [6], from the viscous fluid model [7][8], from elastic deformations with physical analogs, etc. In chapter 3 of this thesis, registration employs semi-local models where the deformations are modeled by cubic B-spline functions.

In chapter 5, variational methods were applied to solve the diffeomorphic deformation problem in 3D ultrasound images.

Registration in chapter three employs semi-local models where the deformations were described by cubic B spline functions. Variational methods was applied in chapter five to solve the diffeomorphic deformation problem in 3D ultrasound images.

3.2.3 Similarity metrics

The quality of registration is often described by a cost function, involving a predominant term quantifying the matching quality, which we shall call the data term. For feature-based methods it is a mean distance between corresponding features in the

source and target images after warping. Note that the case of interpolation with landmarks can be considered as a limit case, when a maximum weight is given to this distance, constraining it to be zero. If the pairing between source and target features is not known, the iterative closest point algorithm can be used to determine it. For pixel-based methods, the data term is a similarity measure between the two images after warping. Correlation, especially local normalized correlation, is an important cost function because of its probabilistic interpretation. It is rather costly to evaluate and sensitive to noise. Instead, the similarity between images is most often expressed using their difference in l^1 or l^2 norms. For images acquired by different imaging modalities, the local mutual information criterion seems to be superior to correlation criteria, at the expense of more computational complexity. For local criteria, such as local normalized correlation, local variance, or local mutual information, the neighborhood size must be properly chosen. The above mentioned criteria will use image interpolation to calculate the warped version of the test image. In template-based methods, the template can be compared with a specific region in the target image using any of the similarity measures suitable for pixel-based methods. For transform-based methods, the least-squares measure in the transformed domain is often used.

In chapter 3 of this thesis, statistical properties of speckle noises were exploited to build the cost function. While in chapter 5, squared difference of voxel intensities was combined with geometric shape similarity metric to generate the total cost function. Because the correspondences between landmark points are not known, the Hausdorff distance between two discretized surfaces was systematically used.

3.3 Registration of 2D-ICE images

2D-ICE is a primary echocardiography modality to acquire visual information on live cardiac anatomy [43][44][45][46][47][48]. However, ICE-images are always perturbed by a strong “speckle” noise, due to diffuse scattering of ultrasound pulses with randomly dispersed small scatterers, at sound wavelength scales. Speckle noise is non-gaussian with a high noise-to-signal ratio. Speckle decorrelation from frame to frame is a standard assumption in this context, as well as for different pixels within a single frame.

Speckle tracking of heart motion focuses on estimation of the non-linear and often large deformations f_t of the heart muscle between a reference image frame J_r and any other image frame J_t in the same heart cycle, in the presence of speckle noise. Several “speckle tracking” algorithms [62][63][64][65][66] have been applied to heart motion recovery in cardiac echocardiography, including optical flow, block matching, and elastic registration.

Classical optical flow approaches [49] [50] [51] have fairly low performance. Indeed they assume pixel intensities to remain constant between consecutive image frames, and hence are highly sensitive to speckle noise. For the block matching techniques [52] [53] [54] [55] [56], capture ability varies with block size, and it is quite difficult to select an optimum block-size that captures both large and small deformations simultaneously. Elastic registration algorithms [57] [58] have performed better for myocardial motion detection, resulting in good spatial capture ability for non-linear deformations.

Most published cardiac motion detection methods still do not attempt to fully use the statistical features of speckle noise. However a few recent studies [53] [54] [59], have explored maximum likelihood techniques to extract cardiac motion, incorporating a well documented theoretical stochastic model for ultrasound speckle noise. In [59] an elastic registration approach was applied to ICE image data, to compute cardiac motion by minimization of a cost function derived from maximum likelihood principles, and validated the technique by comparison with experimental cardiac strain data acquired by recording displacements of microcrystals surgically inserted in the hearts of animal subjects by Dr D. Khoury’s team at Methodist Hospital.

3.3.1 Medical Goal

Strains describe the active motion of the myocardial muscle. In the presence of ischemic heart disease the affected segments have their contraction altered, however they may move because of the tethering of the neighbor segments.

Strain and strain rate have been calculated using ultrasound data derived from Doppler Velocity Imaging. However this measurement has a disadvantage: it only measures the deformation along the ultrasound beam direction and is very dependent on the signal to noise ratio of the Doppler data. In chapter 3, we explore the possibility of obtaining multidimensional strain using speckle tracking and elastic registration.

We compared regional myocardial strains using the Lagrangian strain S , defined at time t_1 as the relative elongation with respect to the initial distance L_{t_0} between

two crystals fixed in mid myocardium of anterior LV wall, i.e.,

$$S(t_1) = \frac{L_{t_1} - L_{t_0}}{L_{t_0}} \quad (3.1)$$

Our approach achieved a good match between experimental cardiac strain data and strain estimates computed from ICE images. In collaboration with Dr D. Khoury we have undertaken to generalize and deepen the method implemented in [59].

3.3.2 B-Splines approximation

The space of spline functions defined on a Euclidean space is a vector space of smooth linear combinations of a fixed family of polynomials (with fixed degree) restricted to contiguous compact supports. These supports are defined as the basic tiles of a "rectangular" net of nodes (called knots). For a spline of degree n , each one of these basic polynomials has degree n , which would suggest that we need $n + 1$ coefficients to describe each piece. However, there is an additional smoothness constraint that imposes the continuity of the spline and its derivatives up to order $n - 1$ at the knots, so that, effectively, there is only one degree of freedom per basic polynomial. Here, we will only consider splines with uniform knots and unit spacing. The remarkable result, due to Schoenberg [32][34], is that these splines are uniquely characterized in terms of a B-spline expansion

$$s(x) = \sum_{k \in \mathbb{Z}} c(k) \beta^n(x - k) \quad (3.2)$$

which involves the integer shifts of the central B-splines of degree n denoted by $\beta^n(x)$; the parameters of the model are the B-spline coefficients $c(k)$. B-splines,

defined below, are symmetrical, bell-shaped functions constructed from the $(n + 1)$ -fold convolution of a rectangular pulse $\beta^0(x)$

$$\beta^0(x) = \begin{cases} 1, & -\frac{1}{2} < x < \frac{1}{2} \\ \frac{1}{2}, & |x| = \frac{1}{2} \\ 0, & \text{otherwise} \end{cases} \quad (3.3)$$

$$\beta^n(x) = \underbrace{\beta^0(x) * \beta^0(x) * \cdots * \beta^0(x)}_{(n+1) \text{ times}} \quad (3.4)$$

The B-splines of degree 3 (cubic B-splines) is shown in Fig. 3.1. Since the B-spline model (3.2) is linear, studying the properties of the basic atoms can tell us a lot about splines in general. Thanks to this representation, each spline is unambiguously characterized by its sequence of B-spline coefficients $c(k)$, which has the convenient structure of a discrete signal, even though the underlying model is continuous.




Figure 3.1: The centered spline of degree 3: cubic B-splines

The B-splines are the basic building blocks for splines. Their usefulness stems from the fact that they are compactly supported; in fact, they are the shortest possible polynomial splines [33]. The simplest way to obtain an explicit formula is to start by writing its Fourier transform

$$\hat{\beta}^n(\omega) = \left(\frac{\sin(\omega/2)}{\omega/2} \right)^{n+1} = \frac{(e^{j\omega/2} - e^{-j\omega/2})^{n+1}}{(j\omega)^{n+1}} \quad (3.5)$$

where we have expressed the $(n+1)$ -fold convolution 3.2 as a product in the frequency domain.

Then, we interpret the complex exponentials as pure phase factors (time shifts), and obtain the corresponding time domain formula

$$\beta^n(x) = \frac{1}{n!} \sum_{k=0}^{n+1} \binom{n+1}{k} (-1)^k \left(x - k + \frac{n+1}{2} \right)_+^n \quad (3.6)$$

This result clearly shows that $\beta^n(x)$ is a piecewise polynomial of degree n . It also implies that the $(n+1)$ th derivative of $\beta^n(x)$ is differentiable up to order n . For n odd, the knots are on the integers, while for n even, they are at the half integers.

B-splines are very easy to manipulate. For instance, we can obtain derivatives through the following formula

$$\frac{d\beta^n(x)}{dx} = \beta^{n-1}\left(x + \frac{1}{2}\right) - \beta^{n-1}\left(x - \frac{1}{2}\right) \quad (3.7)$$

which reduces the degree by one. Similarly, we compute the integral as

$$\int \beta^n(x) = \sum_{k=0}^{\infty} \beta^{n+1}\left(x - \frac{1}{2} - k\right) \quad (3.8)$$

Once we know the effect of linear operators such as (3.7) or (3.8) on the basis functions, it is a trivial matter to apply them to any spline via the B-spline representation (3.2).

Within the family of polynomial splines, cubic splines tend to be the most popular in applications-perhaps due to their minimum curvature property. Using (3.2), we obtain the following closed-form representation of the cubic B-spline

$$\beta^3(x) = \begin{cases} 2/3 - (1 - |x|/2)x^2, & 0 < |x| \leq 1 \\ (2 - |x|)^3/6, & 1 < |x| < 2 \\ 0, & \textit{otherwise} \end{cases} \quad (3.9)$$

which is often used for performing high-quality interpolation. In later chapters, cubic B-splines are used both to interpolate the image intensities and in modelling the deformation functions, and we use $\beta(x)$ to denote $\beta^3(x)$ for simplicity.

In concrete applications, we use the uniformly-spaced cubic B-splines of the following form

$$g(x) = \sum_{k \in K} c(k) \beta(x/h - k) \quad (3.10)$$

Where K is the set of indices of the spline functions, whose supports intersect with the region of interest in the interpolation; h is the knot spacing (the B-splines will be centered at points kh for $k \in K \subseteq \mathbb{Z}$).

Wavelets methods could be an interesting alternative to spline approximation of deformations. However, exploring the change from splines to wavelets means working in a fairly new numerical framework, which requires a quite substantial amount of programming, and of computing time as well, in order to test and adjust properly

the gradient descent parameters. We have not attempted to explore this alternative approach since the registration results obtained by B-splines approximations were already of quite good quality as compared to medical ground truth, and the total computing time was highly reduced with respect to previous implementations of analogous techniques.

B-splines approximation have the following useful properties:

1. Good approximation properties.
2. Computing speed-Cubic splines have a short compact support of length 4. They are symmetric and piecewise cubic. To evaluate $\beta(x)$ at one particular point, only 5 arithmetic operations (additions or multiplications) and 3 comparisons are needed. In multiple dimensions, where we will use tensor products of cubic splines as basis functions, the computational complexity stays low thanks to separability. The number of operations needed to evaluate g also does not depend on the total number of basis functions (and thus the number of parameters $c(k)$).
3. Plausibility-The spline model corresponds to a wide range of physical situations where the restoring force can be approximated as being linearly dependent on the displacement. In such situations, the generated deformation is physically plausible. It is also a good approximation for cases when a better model is not known, such as the deformation between successive ultrasound images, even though the actual organ deformations are not of mechanical origins.
4. Simplicity-The spline model is linear in the parameters $c(k)$ and polynomial

with respect to the spatial position x . It is thus possible to truncate the Taylor expansion such that it becomes exact in neighborhoods of x with a typical diameter of $h/2$.

5. Scalability-Thanks to the q-scale relation $\beta(x/q) = \sum_k a(k)\beta(x - k)$, where $q \in N$, we have the embedding $V_{hq} \subseteq V_h$; i.e., the transition from a coarse space V_{hq} to a finer space V_h is exact [35].

3.3.3 Heart Anatomy and Mechanism

The heart is surrounded by a double-layered membrane called the pericardium. The heart has four chambers: the two upper chambers called the left and right atria, and the lower chambers called the left and right ventricles. The atria act as reservoirs for venous blood, and also function as small pumps to assist ventricular filling. The ventricles are the major pumping chambers that deliver blood to pulmonary (right ventricle) and systemic circulations (left ventricle). The ventricular wall has an external border (the epicardium) and an inner border (the endocardium). The muscular wall itself is called the myocardium and the wall separating the left and right ventricles is called the interventricular septum. The atria also have an interatrial septum.

Four pressure-operated valves ensure that the blood flows only in one direction and prevent blood from leaking backwards from one chamber to the upstream chamber (valvular regurgitation). The aortic and pulmonic valves are referred to as the

semilunar valves and are located at the downstream sides of the left and right ventricle, respectively. The two atrioventricular (AV) valves, the mitral and tricuspid valve, are located between the atria and ventricles. The leaflets of the atrioventricular valves are connected to the papillary muscles which are, in turn, connected to the walls of the ventricles. During the contraction, the papillary muscles also shorten the valves and preventing an outward movement of AV valves toward the atria that would lead to regurgitation.

3.3.3.1 Cardiac Cycle

The cardiac cycle is the sequence of events that occurs when the heart beats. There are two phases of the cardiac cycle. In the diastole phase, the heart ventricles are relaxed and the heart fills with blood. In the systole phase, the ventricles contract and pump blood to the arteries. One cardiac cycle is completed when the heart fills with blood and the blood is pumped out of the heart. The events of the cardiac cycle described below trace the path of the blood as it enters the heart, is pumped to the lungs, travels back to the heart and is pumped out to the rest of the body.

3.3.3.2 Electrocardiogram

Electrocardiogram (ECG) is a transthoracic interpretation of the electrical activity of the heart over time captured and externally recorded by skin electrodes. It is a noninvasive recording produced by an electrocardiographic device.

The ECG works mostly by detecting and amplifying the tiny electrical changes




Figure 3.2: Cardiac cycle of left ventricle

on the skin caused by the depolarization of the heart muscle during each heart beat. At rest, each heart muscle cell has an electric charge across its outer wall, or cell membrane. Reducing this charge to zero is called depolarization, and activates the cell mechanisms which trigger the cell contraction. During each heartbeat a healthy heart will have an orderly progression of a depolarization wave triggered by the cells of the sinoatrial node ; this wave then spreads out through the atrium, passes through "intrinsic conduction pathways" and finally spreads all over the ventricles. This wave is detected as tiny rises and falls in the voltage between two electrodes placed on the

chest skin on both sides of the heart , and is displayed as a wavy line either on a screen or on paper. This display indicates the overall rhythm of the heart, as well as potential weaknesses in different parts of the heart muscle.

3.3.3.3 Intracardiac Echocardiography

Echocardiography is a procedure using ultrasonic compression waves applied to the chest wall to obtain a graphic record of the heart's position, or the motion of heart sub-structures such as ventricular walls and valves. There are several modes of data acquisition from different anatomic locations, such as: transthoracic (TTE), transesophageal (TPE) and intracardiac echocardiography (ICE) [5]. Identified by their spatial imaging capability, echocardiographic techniques can be also recognized by terms, such as 2-D and 3-D echocardiography. Among them, 2-D transthoracic echocardiography is a widely used ultrasound imaging modality in clinical diagnosis, often referred as the echocardiography standard. During the scanning process of TTE, the transducer (or probe) is placed on the chest wall (or thorax) of the patient, and images are taken by manipulating the transducer position via real-time image console. In this mode, there are relatively few acoustic windows through which the heart can be interrogated transthoracically. The apical and parasternal windows permit the heart to be scanned along its long or short axis, respectively. These two windows are used mainly for the analysis of ventricular function, whereas the suprasternal and subcostal windows are primarily applied for other investigations. For instance, the suprasternal window is well suited for imaging the aorta while the subcostal window allows the imaging of the interatrial septum and the inferior vena

cava. These problems of windows for data acquisition are not present for TPE and ICE, since imaging is done within the body either at the esophagus or within the left ventricle via catheter insertion.




Figure 3.3: Example of Intracardiac Echocardiography (ICE)

Intracardiac echocardiography (ICE) has been accepted as a high spatial resolution imaging modality for the diagnosis of cardiac structure and function [40]. It was first developed on the technical basis of intravascular ultrasound (IVUS) [39], which visualizes the structure of vessel walls using a catheter-based image system. The

ICE catheter has a distal transducer that emits and receives ultrasound pulses. To acquire images using ICE system, an ICE catheter has to be guided into the ventricle through the great vessel. Tomographic views of the cavity are acquired by attaching the ICE catheter to a motor drive unit that enables automatic and continuous rotation of the transducer at a fixed speed. The motor unit is fitted with a custom computer-controlled pullback device with an optical sensor that enabled external automatic and accurate withdrawal of the ICE catheter in desired increments. The ICE catheter and motor unit are connected to an imaging console to acquire continuous echocardiographic images. Fig. 3.3 shows an example of ICE image.

3.4 Diffeomorphic matching of deformable shapes in 3D ultrasound imaging

In elastic registration, the deformation is generated by local linearization of the displacement vector field. Achieving image registration by such displacement vector fields has been an important development in computing non-rigid high dimensional transformations of medical imagery allowing for comparison of anatomy in a standard coordinate system. One of the limitations of this approach is that there are no explicit constraints ensuring that the transformations computed are one-to-one or invertible. Indeed, in some cases folding of the grid over itself can occur thereby destroying the neighborhood structure which is essential for the study of anatomy. This registration method is known as the small deformations approach: anatomically

valid registrations using locally linearized models via displacement vector fields correspond to situations where the two images to be matched can be accurately registered by small spatial deformations of the region of interest. It is of considerable interest to compute transformations which are not only invertible but also preserve soft shapes properties such as smoothness of curves, surfaces or other anatomic features. Therefore, diffeomorphic transformations, which are smooth invertible spatial mappings with smooth inverses, are of considerable interest in this regard. Constraining the spatial deformations of the voxels (or pixels) networks to be diffeomorphisms is a natural choice in the study of anatomic deformations of soft organs since connected sets remain connected, disjoint sets remain disjoint, smoothness of anatomical features such as curves and surfaces is preserved, and coordinates are transformed consistently. Developing a rigorous, quantitative methodology for comparing shapes is a contemporary problem actively investigated in image analysis. A typical application in medical imaging-in particular 3D brain imaging-is the comparison of shapes for analogous anatomic structures between two individuals, and the development of a statistical theory which allows soft organs shapes to be studied across populations. This type of investigation is known as computational anatomy [25]. It is motivated by evidence (for example [26][27][28][29]) of shape differences for soft organs, between characteristically different populations-such as males and females-or populations characterized by specific diseases. From a medical diagnosis point of view, serious information can be gathered about brain diseases, for instance, from the study of the shapes of specific brain substructures, and that ultimately this type of investigation will enable early detection of some specific brain diseases.

Typically the anatomic structure of interest is modeled as a 1, 2 or 3 dimensional submanifold of R^3 , such as a curve (1D), image (2D), surface (2D), or volume (3D). Methodologies for studying shape differences are then developed for these models. A main component in the analysis, after obtaining individual model representations for the subjects being studied, is the computation of diffeomorphic deformations between anatomically homologous substructures between the subjects.

Disorders of the mitral valve are the second most frequent heart problem, cumulating 14 percent of the total number of deaths caused by Valvular Heart Disease each year in the United States, and require elaborate clinical management. Visual and quantitative evaluation of the valve is an important step in the clinical workflow according to experts as knowledge about mitral morphology and dynamics is crucial for interventional planning [30][31]. We apply variational techniques based on Hilbert spaces of diffeomorphic transformations to numerically construct diffeomorphic flows modeling the dynamic deformations between multiple 3D-snapshots of the human mitral valve apparatus. In an abstract setting, the optimal matching is formulated as a minimization problem involving actions of diffeomorphisms on regular Borel measures considered as support measures of the reference and the target submanifolds. The objective functional consists of two parts: quantifying the elastic energy of the dynamically deformed surfaces and the matching quality. To make the problem computationally accessible, we use reproducing kernel Hilbert spaces with radial kernels and weighted sums of Dirac measures which gives rise to diffeomorphic point matching and amounts to the solution of a finite dimensional minimization

problem in quite high-dimensions.

Chapter 4

Multiscale Speckle Tracking in Intracardiac Echocardiography

We studied the myocardial deformation through the registration of IntraCardiac Echocardiography (ICE) images, by utilizing the statistical properties of the speckle noise. Within a framework of parametric elastic registration, the deformation was estimated by optimizing the energy function computed by applying the maximum likelihood approach to the speckle noise model. Instead of using the standard rayleigh model to describe the noise, we generalized it to the parametric model which is image specific. The results of speckle noise estimation show that: for the large area outside the heart, even the generalized rayleigh model cannot be verified at all; for other area where the generalized rayleigh model can be fit, the parameter is quite different on different regions of the images. This conclusion leads us to introduce masks to restrict the cost functional to interested cardiac muscle areas, which represent only a small

fraction of the whole images. On the other hand, the use of masks can also reduce the computing time. Also, the optimization procedure is accelerated by applying the multiresolution method, which also improves the robustness of our algorithm. Furthermore, to get the best tradeoff between accuracy and computing time, we proposed the method of mixture computation where we run the algorithm with no mask applied firstly for several multiresolution levels and then apply masks in the remaining levels. Numerical results showed that application of masks greatly reduced the computation time, and the estimation of speckle noise proved the uselessness of the theoretical rayleigh model outside the mask of myocardium, which confirms that adequate masking by semi-local masks is also essential to improve the accuracy of the algorithm.

4.1 Methodology

4.1.1 Stochastic disparity cost functional

Consider a sequence of 2D-images J_t indexed by time t . The cardiac motion f_t between a reference initial time r and any instant t can be defined by an unknown non-linear spatial deformation f_t moving each pixel x to a new position $f_t(x)$. The reference ICE image J_r is compared to the image J_t after spatial warping by f_t , by computing disparities $M_t(x) = J_t(f_t(x)) - J_r(x)$ at all pixels x . Elastic registration begins by defining a cost functional $Kdisp(M_t)$ to globally penalize the field of disparities $M_t = J_t \circ f_t - J_r$ of disparities. In earlier paper [59] the functional $Kdisp$

was derived from the standard Rayleigh distribution model of speckle noise. We have generalized the maximum likelihood approach to compute, for each specific ICE image sequence, a functional $Kdisp$ parameterized by an image based estimation of the speckle noise characteristics.

4.1.2 Regularization cost functional

Deformation extraction is known to be a so-called “ill-posed” problem having far too many solutions. Mathematical solution of ill-posed problems relies on “regularization” of the potential solutions, to transform the question into a well-posed problem. Hence we define a “smoothness” cost functional $Kreg(f_t)$ penalizing large values of the gradient ∇f_t , and providing the needed regularization term. Finally any candidate deformation f_t for the matching of J_r and J_t is evaluated by a penalization cost functional $K(f_t)$ given by $K(f_t) = Kdisp(M_t) + Kreg(f_t)$. Minimization of the cost functional $K(f_t)$ by multiscale gradient descent provides then an estimate of the deformation f_t .

4.1.3 Estimation of the speckle noise

The theoretical multiplicative model for speckle noise, based on the Rayleigh distribution, as used in [59], becomes parameter free as soon as one considers the ratio of two independent speckle noise values, and hence cannot integrate accurately the actual statistical properties of specific ultrasound image sequences. We use this theoretical model as a preliminary rough approximation of speckle noise likelihood

on several small image patches. We minimize locally the associated cost functional $Kdisp$ by fast gradient descent to obtain local first approximations of f_t and then of disparities M_t , in order to generate direct estimates of the actual speckle noise.

4.1.4 Image specific speckle noise modelling

Our local estimations of speckle noise enable us to compute key statistical characteristics of the speckle, taking into account a parametric model for the unknown software transformation implemented to generate the(logarithmic) B-mode presentation of ICE images. We thus found out that within the myocardium,the Rayleigh model of speckle noise remains valid, but with quite large variances, which are not constant over the whole myocardium. Outside of the myocardium, the speckle noise either does not follow the Rayleigh model at all, or when it seems to follow that model, the local std(standard deviation)s are quite different from those estimated within the heart muscle areas. Moreover we have shown that contrary to the standard non-correlated speckle noise models, the actual speckle noise on ICE image data exhibits strong correlations between neighboring pixels situated at the same radial distance from the ultrasound emitter.

4.1.5 Minimizing image specific cost functionals

Our detailed speckle modelling results shows that the use of local masks and image specific statistical speckle noise modelling are necessary in order to improve accuracy of the probabilistic cost functional $Kdisp$ penalizing disparities. Our approach is

hence, after finalizing the speckle noise modelling, to launch a second run for the final minimization of $K(f_t)$ by gradient descent, but this time with a more accurate version of the cost functional $Kdisp$.

4.1.6 Multiresolution and local Masking

Fast computing of myocardial deformations is of high interest for potential clinical use. We have improved computing efficiency by introducing masks to restrict the cost functional to cardiac muscle areas, which represent a small fraction of the whole ICE images, instead of including all voxels as in [59]. Another key ingredient for faster computing is to implement an optimized coarse to fine multiresolution approach for gradient descent. At coarser resolutions, masks have small sizes, and computing time is strongly reduced, since at coarser scales we also use faster stopping criteria for gradient descent.

4.1.7 Validation against experimental heart strain data

To evaluate performances of these approaches on actual ICE images, we have compared our estimated myocardial strains with experimental sonomicrometry strain data recorded on animal subjects by Dr D. Khoury and his team.

4.2 Mathematical definition of the cost functionals

4.2.1 Cost Functionals

In the spirit of elastic registration methods, we introduce a cost functional $Kdisp(f_t)$ to penalize the residual intensity disparities between reference image and target image after spatial warping f_t . The low signal to noise ratio suggests that probabilistic evaluation of disparities should boost the efficiency of the $Kdisp$ functional. Recent work [59] indicates that maximum likelihood methods do improve motion estimation. This leads to the definition of $(-Kdisp)$ as the probabilistic log-likelihood of the observed residual intensity disparities under a speckle noise stochastic model. Non-linear deformations recovery is known to be an ill-posed inverse problem, and hence requires “regularization” of the unknown deformation f_t . Regularization is achieved here in two ways : firstly by selecting a parametric smooth model (cubic B-splines) for the deformation f_t , and secondly by adding to the cost functional $Kdisp$ a smoothness term $Kreg$ penalizing large values of the deformation gradient ∇f_t).

4.2.2 B-mode ultrasound display and true Tissue Echogenicity

In ultrasound imaging, for each pixel x , the acquisition system aims to recover the “true” intensity $T(x)$ representing “Tissue echogenicity” at x . But the true intensity

$T(x)$ is not directly observable, due to multiplicative corruption by a speckle noise $N(x)$. The echo-acquisition system thus only recovers, by “Envelope Detection”, a noise corrupted version $ED(x)$ of the true $T(x)$. The envelope-detected image ED is linked to speckle noise N and tissue echogenicities T by the multiplicative formula:

$$ED(x) = T(x)N(x) \quad (4.1)$$

In medical ultrasound imaging, the envelope-detected image ED is submitted to logarithmic compression, replacing $ED(x)$ by $LED(x) = \ln(ED(x))$. Then for better display, the log-intensities $LED(x)$ undergo a fixed recoding selected by the internal software of the imaging system, to generate the echocardiographic image frame J displayed in B-mode for visual inspection. We modelled this unknown internal recoding by a linear transformation, involving 2 unknown parameters, a positive re-scale coefficient $\frac{1}{\alpha}$ and an offset θ :

$$J(x) = \frac{1}{\alpha} \ln(ED(x)) + \theta \quad (4.2)$$

The uncorrupted B-mode image I one would ideally like to access or display is given by

$$I(x) = \frac{1}{\alpha} \ln(T(x)) + \theta \quad (4.3)$$

and we hence have the obvious additive relation

$$J(x) = I(x) + S(x) \quad (4.4)$$

where $S(x)$ is the additive noise on observed B-mode ICE images, and is related to the speckle noise by the formula

$$S(x) = \frac{1}{\alpha} \ln(N(x)) \quad (4.5)$$

For brevity we call S the log-speckle noise.

The speckle noise values $N(x)$ at various pixels x are usually modelled as identically distributed independent random variables. Their common distribution is assumed to be the celebrated Rayleigh density $p_N(s)$ with variance σ^2 , given by

$$p_N(s) = \frac{1}{\sigma^2} e^{-s^2/2\sigma^2} \quad \forall s \in R^+ \quad (4.6)$$

4.2.3 Maximum Likelihood approach

Fix two ICE-image frames J_r and J_t acquired at times r and t , where J_r is a reference image. As seen in equation (4.4), we have for all pixels x

$$J_t(x) = I_t(x) + S_t(x); \quad J_r(x) = I_r(x) + S_r(x) \quad (4.7)$$

where $I_t(x)$, $I_r(x)$ are the “true” intensity values, and $S_t(x)$, $S_r(x)$ are the log-speckle values.

All images are indexed by a fixed planar rectangular pixel grid G . Call $f : G \rightarrow G$ the spatial deformation of the pixel grid G capturing the exact apparent motion (between times r and t) of the planar heart section displayed by image I_r .

Ideally we assume that the unknown spatial deformation f of cardiac tissue preserves tissue echogenicity so that for all pixels x

$$I_t(f(x)) = I_r(x)$$

Hence after warping by the unknown true deformation f , the residual intensity

disparities $Disp_t = J_t \circ f - J_r$ for B-mode intensities verify

$$Disp_t(x) = S_t(f(x)) - S_r(x) = \frac{1}{\alpha} \ln(N_t(f(x))/N_r(x)) = \eta(x) \quad (4.8)$$

The variables $\eta(x) = \frac{1}{\alpha} \ln(N_t(x)/N_r(x))$ are independent random variables with the same probability density function (pdf), namely the pdf ρ of $\frac{1}{\alpha}W$ where $W = \ln(Y/Z)$, and where Y, Z are independent random variables having Rayleigh distributions with identical variances. An easy theoretical computation shows that the pdf τ of W is given by

$$\tau(w) = \frac{2e^{2w}}{(e^{2w} + 1)^2} = \frac{1}{2 \cosh^2(w)} \quad \forall w \in R^+ \quad (4.9)$$

Hence all the random disparities $\eta(x)$ generated by noise have the same pdf ρ of $\frac{1}{\alpha}W$ and is given by:

$$\rho(s) = \alpha \tau(\alpha s) \quad \forall s \in R^+ \quad (4.10)$$

We denote the common log-density function $\ln(\rho)$ of the disparity noise variables $\eta(x)$ by

$$\frac{1}{2} \ln(\rho(s)) = \frac{1}{2} \ln(2\alpha) - g(\alpha s) \quad \forall s \in R^+ \quad (4.11)$$

where the function g is given by

$$g(w) = \ln(1 + e^{2w}) - w \quad \forall w \in R^+ \quad (4.12)$$

For myocardial deformation estimation, the only region of interest (ROI) in our ICE images mainly correspond to the myocardial muscle. Hence, we will restrict the observation of disparities $Disp_t$ to the pixels of the ROI, by selecting and fixing at

the reference time r a spatial mask MSK , i.e. a subset MSK of the pixel grid G . Given the true spatial deformation f and the true image intensities I_r and I_t , the restriction to MSK of the random field $Disp_t$ of disparities after warping has the following conditional pdf:

$$p(Disp_t) = \prod_{x \in MSK} \rho(Disp_t(x)) \quad (4.13)$$

The maximum likelihood principle indicates that a good estimator $\hat{f}$ of the true deformation f on our ROI MSK , should maximize in f the conditional likelihood $p(Disp_t)$ given the observed random field $Disp_t$ of disparities.

The average over all x in our ROI MSK of the opposite of the negative of the log likelihood of the disparity field $Disp_t(x)$ hence becomes the natural disparity cost functional $Kdisp$, which our estimator $\hat{f}$ of f will attempt to minimize. We thus define the disparity cost $Kdisp$ by

$$Kdisp(f) = \frac{1}{card\{MSK\}} \sum_{x \in MSK} g(\alpha Disp_t(x)) \quad (4.14)$$

Note that even when f is the true deformation, $Kdisp(f)$ is a random variable, obtained as the average of a large number of independent identically distributed random variables. For large values of $card\{MSK\}$, the Law of Large Numbers proves that $Kdisp(f)$ has a limit tightly linked to the entropy of ρ . Thus when f is the true deformation, the limit of $Kdisp(f)$ for even moderately large regions of interest, is deterministic and equal to

$$H(\rho) = \int_{-\infty}^{\infty} g(\alpha w) \rho(w) dw = 1$$

But for large ROI, we expect the true deformation f to minimize $Kdisp(f)$, so we should expect the minimal disparity cost $\min Kdisp$ to verify $\min Kdisp = Kdisp(\text{true}f) = 1$.

4.2.4 Regularization term $Kreg$ and final cost functional K

For each x in the pixel grid G we define its displacement $u(x)$ under the deformation f by:

$$u(x) = f(x) - x \quad (4.15)$$

Clearly the vector field u immediately determines f and conversely. From now on, the unknown deformation we search to determine will be the displacement vector field u , and the disparity cost $Kdisp(f)$ will be considered as a functional $Kdisp(u)$ of u . Introduce notations for planar coordinates

$$\begin{aligned} x &= (x_1, x_2); & y &= f(x) = (y_1, y_2); \\ u(x) &= (u^1, u^2) = ((y_1 - x_1), (y_2 - x_2)) \end{aligned} \quad (4.16)$$

As explained precisely below we will model the displacement vector field u by a the restriction to the pixel grid G of a smooth vector field u defined on R^2 . The “regularization” term $Kreg(u)$ is a complementary cost functional penalizing the lack of smoothness of the displacement field u and we define $Kreg(u)$ as a squared discrete L_2 norm of $gradient(u)$ explicitly given by :

$$Kreg(u) = \frac{1}{card\{MSK\}} \sum_{x \in MSK} \sum_{i,j=1,2} \left(\frac{\partial u^i}{\partial x_j} \right)^2 \quad (4.17)$$

The final cost functional $K(u)$ to be minimized combines the disparity term $Kdisp(u)$ with the regularization term $Kreg(u)$ as follows

$$K = Kdisp + \gamma Kreg \quad (4.18)$$

where γ is a fixed positive weight to be selected later.

The natural approach to minimize the cost functional $K(u)$ in u is to implement a gradient descent on u , and of course to restrict this gradient descent to an adequate vector space V of smooth displacement fields u on R^2 . We first present the finite dimensional vector spaces V of functions we shall use to approximate u .

4.3 Multi-resolution optimization of the cost functionals

4.3.1 B-spline model for the displacement vector field u

Consider a lattice of knots $L = \{k = (k_1, k_2) | 1 \leq k_1 \leq L_1, 1 \leq k_2 \leq L_2\}$, where the knots k are placed on a regular subgrid of the pixel grid G . Denote by N_1, N_2 the dimensions of the grid G and define the knot spacing $h = (h_1, h_2)$ where $h_1 = N_1/L_1$ and $h_2 = N_2/L_2$ are required to be integers.

The unknown displacement field u is a 2D non-linear transformation which we modelled by a linear combination of cubic B-splines associated to the knots of the lattice L . The unknown vector c of coefficients $c_k = (c_k^1, c_k^2)$ of this linear combination completely defines the displacement field u (and hence the spatial deformation f) by

the explicit formula :

$$u(x) = \sum_{k \in L} c_k B(x, h, k) \quad (4.19)$$

where

$$B(x, h, k) = \beta\left(\frac{x_1}{h_1} - k_1 + \frac{1}{2}\right) \beta\left(\frac{x_2}{h_2} - k_2 + \frac{1}{2}\right) \quad (4.20)$$

and $\beta(s)$ is 1D cubic B-spline function.

During the warping process, we need a continuous version of the discrete image J_t . Using also cubic B-splines, we interpolate the discrete image J_t to extend it from the pixel grid G to a continuous function still denoted J_t :

$$J_t(x) = \sum_{y \in G} b_y B(x - y) \quad (4.21)$$

where $B(x - y) = \beta(x_1 - y_1) \beta(x_2 - y_2)$, and b_y is a set of interpolation coefficients, which can be computed prior to the search for u by solving a linear system. Due to the short support (length 4) of cubic B-splines, almost all off-diagonal elements of the matrix of this linear system are zeros. Also, the b_y can be obtained by an effective filtering algorithm [61] which incurs negligible overhead.

4.3.2 Multi-resolution optimization of cost functionals

To compute the unknown spatial displacement field u matching images J_t and J_r , we need to minimize the cost functional $K(u)$. To minimize $K(u)$ we implement a multiresolution gradient descent in u at coarse to fine successive resolutions. Multi-resolution is used for both the image model and the deformation model.

First, we build an image pyramid composed of gradually coarser versions of the original image by regularly spaced down-sampling. Starting from an initial ICE image of size $N \times N$ (note that $N = 512$ for our real ICE images), we create, by successive dichotomic decimations of pixels, a sequence of increasingly coarser images with sizes $\{N^i \times N^i\}_{i=1,\dots,p}$, where $N^1 = N$, $N^{i+1} = N^i/2$, and we fix the coarsest level $N^p = 64$.

Second, we use multi-resolution also for the B-spline deformation model of u . We start with a coarse B-spline model of u defined by a lattice of knots L_1 with a large initial distance h^1 between knots. We fix an increasing sequence of knots lattices L_i with corresponding knots spacing $\{h^i \times h^i\}_{i=1,\dots,q}$, where $h^{i+1} = 2h^i$ and $h^q = 64$ in the finest level.

Both multi-resolutions do improve accuracy as well as speed of the gradient descent algorithm. Instead of refining both the image and the deformation models at the same rate, we dynamically define a succession of pairs of resolution scales $RS(m) = (N^{i(m)}, h^{j(m)})$ for the image and the deformation models, starting with the coarsest versions of these 2 models. More precisely, once the current pair of resolutions $RS(m)$ has been selected, we interpolate at resolution $RS(m)$ the last computed approximation u^{m-1} of u , just obtained at coarser resolution $RS(m-1)$, and use this interpolated u^{m-1} to initiate and then perform a gradient descent on $K(u)$ computed at the new resolution $RS(m)$. This gradient descent at fixed resolution is detailed below, and terminates automatically at a new approximation u^m of u . We then refine $RS(m)$ into a finer pair of resolutions $RS(m+1)$ by alternate refining of either the image model or the B-spline model. Successive modifications

of resolutions stop when the predefined finest level is reached.

In our actual numerical implementations, we deal with images of size 512×512 , and after intensive simulations we have selected the most efficient sequence of resolutions pairs $RS(m) = (N^{i(m)}, h^{j(m)})$, $m = 1, 2, \dots, 7$, for our ICE images as follows $RS = \{(64, 64), (128, 128), (128, 64), (256, 128), (256, 64), (512, 128), (512, 64)\}$.

4.3.3 Detailed gradient descent at fixed resolution

To minimize the cost functional $K(u)$ we implement a multiresolution gradient descent in u at coarse to fine successive pairs of resolutions $RS(m)$, both for the image model and for the B-spline model. For each fixed pair of resolutions $RS(m)$, the current lattice of B-spline knots L is fixed, and hence u is approximated by its corresponding B-spline expansion, which is defined by an unknown vector of coefficients c_k^i , ($i = 1, 2; k \in L$). After adequate restriction to the current image model, the cost functional to minimize becomes an explicit function $K(c)$ given by the formulas (4.14, 4.18, 4.17) obtained above. To minimize $K(c)$ we implement a gradient descent in c , with the following update rule :

$$\Delta c = -\mu \nabla K(c) \tag{4.22}$$

where $\nabla K(c)$ is the gradient of $K(c)$, μ is the step size. This dynamic step size μ is multiplied by a fixed factor κ after a successful step, i.e a step which actually decreases the cost functional, and is divided by κ after an unsuccessful step. In most of our numerical implementations we selected $\kappa = 10$.

The differential of the cost functional can be derived easily for $(i, j, l = 1, 2; l \neq j)$:

$$\begin{aligned} \frac{\partial K_{disp}}{\partial c_k^i} &= \frac{\alpha}{card\{MSK\}} \sum_{x \in MSK} \frac{\exp(2\alpha Disp(x)) - 1}{\exp(2\alpha Disp(x)) + 1} \\ &\quad \cdot \left. \frac{\partial J_t(f)}{\partial f_i} \right|_{f(x)=u(x)+x} \frac{\partial u^i}{\partial c_k^i} \end{aligned} \quad (4.23)$$

where

$$\frac{\partial J_t(f)}{\partial f_j} = \sum_{y \in G} b_y \beta'(u_j + x_j - y_j) \beta(u_l + x_l - y_l) \quad (4.24)$$

$$\frac{\partial u^i}{\partial c_k^i} = B\left(\frac{x}{h} - k\right) \quad (4.25)$$

We also have

$$\frac{\partial K_{reg}}{\partial c_k^i} = \frac{2}{card\{MSK\}h} \sum_{x \in MSK} \sum_{j=1}^2 \frac{\partial u^i}{\partial x_j} \beta'\left(\frac{x_j}{h_j} - k_j + \frac{1}{2}\right) \beta\left(\frac{x_l}{h_l} - k_l + \frac{1}{2}\right) \quad (4.26)$$

Two simultaneous stopping criteria are implemented to terminate the gradient descent : the descent process is stopped at iteration step i either if the number of iterations i exceeds a preset large integer, or if $\Delta K^i < \epsilon$, where $\Delta K^i = |K(c^i) - K(c^{i-1})|/K(c^{i-1})$ is the relative increment of the cost function at the last iteration step. The size ϵ of the terminating threshold is selected to minimize the computing time, so that we select different values $\epsilon(m)$ for ϵ at each resolution level $RS(m)$. Indeed at coarser resolutions, ϵ may safely take higher values than at finer resolutions.

4.3.4 Image sequence analysis and Masking

To compute the displacement fields u_t of an image sequence $J_t, t = 1, \dots, n$ covering a whole cardiac cycle, we apply successively the registration algorithm to compute an

optimal registration REG_t matching the consecutive pair of images J_{t-1} and J_t for $t = 2, \dots, n$. The global displacement field J_t is then obtained by successively applying the registrations $REG_1, \dots, REG_t$. Each small deformation generated by REG_i is verified to be invertible, therefore the global large deformation is an invertible deformation. For each image J_{t-1} , temporarily considered as a reference image, we determine an initial mask MSK corresponding to the zone of interest in the myocardial area. A few key boundary pixels of MSK can be easily selected interactively and used to generate a mask MSK with a piecewise linear boundary linking these initial vertices.

For software tracking of crystals surgically inserted in animal subjects and experimentally tracked by quite accurate sonomicrometry, we can even select, at each frame time t , a rather small mask which only contains the regions around the crystals. The mask MSK then automatically generates a pyramid of masks $MSK(m)$ at the selected resolution levels $RS(m)$.

The displacement field $u(x)$ is computed only for pixels x in the mask MSK , and the cost functionals K are restricted to the pixels in the mask MSK . Hence for given choices of the multi-resolution sequence $RS(m)$ and of the stopping thresholds such as $\epsilon(m)$, the computing time for the optimization process is approximately proportional to the area of the mask MSK .

4.3.5 Estimation of the log-speckle noise and of the re-scaling parameter

The value of the re-scaling coefficient α must be determined before applying the gradient descent method to minimize $K(u)$. To estimate α we need to estimate the variance of the log-speckle noise η . To solve this, we implement a first gradient descent with the hypothesis $\alpha = 1$ and thus compute a first approximation u of the displacement field. We can then obtain a direct estimate $\hat{\eta}(x)$ of the noise variable $\eta(x) = Disp_t(x)$ by computing the matching error $Disp_t(x)$.

$$\hat{\eta}(x) = J_t(x + u(x)) - J_r(x), \quad \forall x \in MSK \quad (4.27)$$

The explicit pdf of all the $\eta(x)$ has been computed above and immediately yields the equation

$$\alpha = \frac{\pi}{\sqrt{12Var(\eta(x))}} \quad (4.28)$$

The empirical variance of the estimated noise $\hat{\eta}$ can of course be computed by standard statistical formulas on arbitrary patches of pixels. But note that large values of the variance for the ideal residual disparities $\eta(x)$ correspond to small values of α . To avoid strong overestimation of this variance and hence strong underestimation of α we must use more sophisticated “robust estimates” of the empirical variance of $\hat{\eta}$. We implement this important guideline by eliminating a small percentage of the highest outliers in the observed values of $|\hat{\eta}|$, and we also impose an absolute upper bound of 20% for the noise to signal ratio.

4.4 Detailed analysis of the disparity term $Kdisp$ derived from log-likelihood

The disparity cost $Kdisp$ defined above can be written with the preceding notations

$$Kdisp = \frac{1}{card\{MSK\}} \sum_{x \in MSK} g(\alpha Disp_t(x)) \quad (4.29)$$

The “influence function” of the penalty function g has derivatives g', g'' :

$$g'(w) = \frac{\exp(2w) - 1}{\exp(2w) + 1} = \tanh(w) \quad (4.30)$$

$$g''(w) = \frac{1}{\cosh^2(w)} \quad (4.31)$$

Clearly, the absolute value of g' is bounded by 1, which confers robustness of the deformation estimates to the presence of very large disparities (outliers) for the cost functional $Kdisp$.

We can approximate $g(\alpha s)$ by:

$$g(\alpha s) \approx \begin{cases} \alpha s, & s > 2/\alpha \\ \frac{1}{2}\alpha^2 s^2 + \ln(2), & s < 2/\alpha \end{cases} \quad (4.32)$$

We are matching observed images J_t and J_r corresponding to noisy versions of true images I_t and I_r ; we write J and I instead of J_t and I_t to simplify notations, and we call I' and I'' the approximate gradient and hessian of the true image I . Let f be the true deformation $f(x) = x + u(x)$ which we estimate by $\hat{f}(x) = x + \hat{u}(x)$. Define the deformation error field δ by

$$\delta(x) = \hat{f}(x) - f(x) \quad (4.33)$$

After warping by $\hat{f}$ the residual disparity $Disp(x)$ observed at pixel x is

$$Disp(x) = J(\hat{f}(x)) - J_r(x) = I(\hat{f}(x)) - I_r(x) + \eta(x) \quad (4.34)$$

where the random variable $\eta(x)$ is as above the difference of 2 independent log-speckle noise values at pixels x and $\hat{f}(x)$. Since f is the true deformation we have $I(f(x)) = I_r(x)$ for all x and hence

$$Disp(x) = I(\hat{f}(x)) - I(f(x)) + I(f(x)) - I_r(x) + \eta(x) = I(\hat{f}(x)) - I(f(x)) + \eta(x) \quad (4.35)$$

Assume that the deformations errors $\delta(x)$ are small. We use the temporary notations $y = f(x)$, $D^t = \text{transpose of } D$, and apply a 2nd order Taylor approximation of I to obtain

$$Disp(x) \approx I'(y)\delta(x) + \frac{1}{2}\delta(x)^t I''(y)\delta(x) + \eta(x) \quad (4.36)$$

This implies, using Taylor formula for the function g ,

$$g(\alpha Disp(x)) \approx R1(x) + R2(x)(I'(y)\delta(x) + \frac{1}{2}\delta(x)^t I''(y)\delta(x)) + \frac{1}{2}R3(x)(I'(y)\delta(x))^2 \quad (4.37)$$

where $R1, R2, R3$ denote the 3 random fields defined by

$$\begin{aligned} R1_x &= g(\alpha\eta(x)) \\ R2_x &= \alpha'(\alpha\eta(x)) \\ R3_x &= \alpha^2 g''(\alpha\eta(x)) \end{aligned} \quad (4.38)$$

By definition of $Kdisp$ we can write

$$Kdisp(\hat{f}) = \frac{1}{card\{MSK\}} \sum_{x \in MSK} g(\alpha Disp(x)) \quad (4.39)$$

and hence the 2nd order approximation of $g(\alpha Disp(x))$ provide an approximation of $Kdisp(\hat{f})$ as the sum of 3 groups of terms

$$Kdisp(\hat{f}) \approx KA1 + KA2 + KA3 \quad (4.40)$$

where we have

$$\begin{aligned} KA1 &= \frac{1}{card\{MSK\}} \sum_{x \in MSK} R1_x \\ KA2 &= \frac{1}{card\{MSK\}} \sum_{x \in MSK} R2_x I'(y) \delta(x) \\ KA3 &= \frac{1}{2card\{MSK\}} \sum_{x \in MSK} (R2_x \delta(x) * I''(y) \delta(x) + R3_x (I'(y) \delta(x))^2) \end{aligned} \quad (4.41)$$

Using the law of large numbers we have already shown above that

$$KA1 = Kdisp(f) = 1 \quad (4.42)$$

Assume that we are approximating the true transformation f at a reasonably lower resolution than the image data (a better assumption for 3D-images than for 2D-images). To compute $Kdisp$, we can then first average all the terms $Kdisp_x = g(\alpha Disp(x))$ of $Kdisp$ corresponding to the high resolution pixels x of each small image patch $Patch_X$ centered around a low resolution pixel X . This will associate a disparity cost $Kdisp_X$ to each low resolution pixel X . And $Kdisp$ will still be the average of all the $Kdisp_X$ over all the low resolution pixels X in our ROI.

But in the local computation of $Kdisp_X$ the main term may then be estimated by replacing, for all x in $Patch_X$, the values of I', I'', δ at the pixel $f(x)$ by their values at the fixed pixel $f(X)$ and can then be factored out in the averaging over all

x in $Patch_X$. In the averaging of $Kdisp_x$ over all x in $Patch_X$, the random terms $R1_x, R2_x, R3_x$ then appear only through their resp. averages av_1, av_2, av_3 over all x in $Patch_X$, which can be written, for $j=1,2,3$,

$$av_j = \frac{1}{card\{Patch_X\}} \sum_{x \in Patch_X} Rj_x \quad (4.43)$$

In the averaging of the $Kdisp_x$ over all x in $Patch_X$, the 3 groups of terms corresponding to $KA1, KA2, KA3$ now become

$$\begin{aligned} KA1_X &\approx av_1 \\ KA2_X &\approx av_2 I'(f(X))\delta(X) \\ KA3_X &\approx \frac{1}{2} [av_2 \delta(X)^t I''(f(X))\delta(X) + av_3 (I'(X)\delta(X))^2] \end{aligned} \quad (4.44)$$

Since the local disparity cost $Kdisp_X$ verifies

$$Kdisp_X \approx KA1_X + KA2_X + KA3_X \quad (4.45)$$

we obtain the following approximation of $Kdisp_X$ at each low resolution pixel X

$$Kdisp_X \approx av_1 + av_2 I'(f(X))\delta(X) + \frac{1}{2} [av_2 \delta(X)^t I''(f(X))\delta(X) + av_3 (I'(X)\delta(X))^2] \quad (4.46)$$

Call p the cardinal of the set of low resolution pixels $lowMSK = X \in MSK$. By definition of the $Kdisp(X)$, the global disparity cost $Kdisp(\hat{f})$ is the average the $Kdisp(X)$ over $lowMSK$:

$$Kdisp(\hat{f}) = \frac{1}{p} \sum_{X \in lowMSK} Kdisp(X)$$

The estimate $\hat{f}$ of f must minimize $Kdisp(\hat{f})$ and hence for each low resolution pixel $X \in MSK$, all the $Kdisp(X)$ must be approximately minimized by $\delta(X)$. Since $Kdisp(X)$ is approximately a quadratic form in $\delta(X)$, given by equation (4.46) above, this implies that for each low resolution pixels $X \in MSK$, the differential of $Kdisp(X)$ with respect to $\delta(X)$ is approximately equal to 0, so that

$$I'(f(X)) + \left[I''(f(X)) + \frac{av_3}{av_2} (I'(f(X)) * I'(f(X))) \right] \delta(X) \approx 0 \quad \forall X \in lowMSK \quad (4.47)$$

We shall see below that the random variable $\frac{av_3}{av_2}$ is of the form

$$\frac{av_3}{av_2} = \alpha Z$$

where the random variable Z has a fixed pdf which does not depend on X nor on α .

For each low resolution pixel $X \in MSK$, define then the 2×2 matrix $Q(X)$ by

$$Q(X) = I''(f(X)) + \alpha Z (I'(f(X)) * I'(f(X))) \quad (4.48)$$

For each $X \in MSK$ such that $Q(X)$ is invertible we then have the following estimate of the local deformation error $\delta(X)$

$$\hat{f}(x) - f(x) = \delta(X) = -Q(X)^{-1} I'(f(X)) \quad (4.49)$$

To complete this estimate, valid for small deformation errors, we need to evaluate the random variables av_j for $j=1,2,3$.

Note that by the central limit theorem and its known pragmatic rates of convergence, as soon as the integer $q = \text{card}\{Patch_X\}$ is larger than 12 we may consider

that the random variables $av_{j,j} = 1,2,3$, are approximately Gaussian variables with means ERj and variances VRj given by the following probabilistic expected values

$$\begin{aligned} ERj &= E(Rj(x)) \\ VRj &= \frac{1}{q} \text{var}(Rj(x)) \end{aligned} \quad (4.50)$$

where x is an arbitrary pixel since for each $j=1,2,3$, all $Rj(x)$ are identically distributed.

In section 2.4 above, we introduced the fixed random variable $W = \ln(Y/Z)$ which has density $\tau(w) = \frac{1}{2 \cosh^2(w)}$. Since all the random variables $\eta(x)$ have the same pdf ρ as $\frac{1}{\alpha}W$ the constants ERj and VRj are given by,

$$\begin{aligned} ER1 &= E(R1x) = E(g(\alpha\eta(x))) = E(g(W)) = \int_{-\infty}^{\infty} g(w)\tau(w)dw \\ ER2 &= E(R2_x) = E(\alpha g'(\alpha\eta(x))) = \alpha E(g'(W)) = \alpha \int_{-\infty}^{\infty} g'(w)\tau(w)dw \\ ER3 &= E(R3_x) = E(\alpha^2 g''(\alpha\eta(x))) = \alpha^2 E(g''(W)) = \alpha^2 \int_{-\infty}^{\infty} g''(w)\tau(w)dw \end{aligned} \quad (4.51)$$

which yields

$$\begin{aligned} ER1 &= 1 \\ ER2 &= 0 \\ ER3 &= \alpha^2 m_0 \end{aligned} \quad (4.52)$$

where we introduce the numerical constant

$$m_0 = \int_{-\infty}^{\infty} \frac{1}{2 \cosh^4(w)} dw \quad (4.53)$$

Similarly we have

$$\begin{aligned}
VR1 &= \frac{1}{q} \text{var}(R1_x) = \frac{1}{q} \text{var}(g(W)) = m_1^2 \\
VR2 &= \frac{1}{q} \text{var}(R2_x) = \alpha^2 \frac{1}{q} \text{var}(g'(W)) = \alpha^2 m_2^2 \\
VR3 &= \frac{1}{q} \text{var}(R3_x) = \alpha^4 \frac{1}{q} \text{var}(g''(W)) = \alpha^4 m_3^2
\end{aligned} \tag{4.54}$$

where the positive numerical constants m_1, m_2, m_3 are given by

$$\begin{aligned}
m_1^2 &= \frac{1}{q} \int_{-\infty}^{\infty} (g(w) - 1)^2 \tau(w) dw \\
m_2^2 &= \frac{1}{q} \int_{-\infty}^{\infty} g'(w)^2 \tau(w) dw \\
m_3^2 &= \frac{1}{q} \int_{-\infty}^{\infty} (g''(w) - m_0)^2 \tau(w) dw
\end{aligned} \tag{4.55}$$

We thus have the approximations

$$\begin{aligned}
av_1 &\approx 1 + m_1 Z_1 \\
av_2 &\approx \alpha m_2 Z_2 \\
av_3 &\approx \alpha^2 (m_0 + m_3 Z_3)
\end{aligned} \tag{4.56}$$

where Z_1, Z_2, Z_3 are standard gaussian variables

The matrices $Q(X)$ introduced above only involve the ratio $\frac{av_3}{av_2}$ which is given by

$$\frac{av_3}{av_2} \approx \alpha \frac{m_0 + m_3 Z_3}{m_2 Z_2} = \alpha Z \tag{4.57}$$

where the random variable Z is given by $Z = \frac{m_0 + m_3 Z_3}{m_2 Z_2}$. The main qualitative conclusions supported by the preceding approximate formulas are the following.

1. for large values of α , i.e. for small speckle noise variance, the projection of the deformation error $\delta(X)$ in the direction of the image gradient $I'(f(x))$ tends

to 0 as $\frac{1}{\alpha}$, but the error component orthogonal to $I'(f(x))$ does not necessarily vanish, and is essentially controlled by the invertibility of the image Hessian $I''(f(x))$.

2. for small values of α , i.e. for large speckle noise variance, the deformation error $\delta(X)$ is very close to

$$Bias(X) = -(I''(f(X)))^{-1}I'(f(X))$$

Note that if the true image $I(x)$ is locally an affine functional of the pixel x in the neighborhood of pixel X , one has $I''(x) = 0$ and hence for small values of α the bias of the deformation error $\delta(X)$ will tend to be high. Thus local linearities in image intensities should recommend to avoid small values of α .

The preceding bias formula suggests a further improvement to the current method presented here, by adding to the deformation estimate $\hat{f}(x)$ of $f(x)$ a corrective term $I''(\hat{f}(x))^{-1}I'(\hat{f}(x))$ in order to reduce the bias of the deformation estimate.

Chapter 5

Numerical performance of speckle tracking

5.1 Numerical performances on simulated noise corrupted and warped images

5.1.1 Three simulated pairs of images

We have first tested the performance of our speckle tracking algorithms on pairs of artificially generated images I_r and I_t corrupted by simulated log-speckle noise at several signal-to-noise ratios. Three true test images I_t were produced by smoothing a fixed 512x512 real ICE image I_0 with linear Gaussian filters of various window size and variance (denoted by $G(\text{Window size, Variance})$). The three reference cases

tested here were:

1. The true reference image is $I_t = I_0 * G(7, 3)$.
2. The true reference image is $I_t = I_0 * G(13, 5)$.
3. The true reference image is $I_t = I_0 * G(25, 9)$.

From case 1 to case 3, the true reference image I_t becomes smoother, as shown in figure 5.1.1. As a smoothness indicator, the mean absolute values of $\text{gradient}(I_t)$ (computed by a Sobel filter) in these three cases are 22, 14 and 7.5 respectively. The true test image I_t was in each case warped by a known fixed non-linear spatial deformation $f(x) = x + u(x)$ of the pixel grid to generate another true image I_r verifying the exact warping identity $I_t(f(x)) = I_r(x)$. The displacement field $u(x)$ of this simulated warping $f(x)$ was generated by an explicit expansion in 2D cubic B-splines. With previous notations, we used the B-splines corresponding to a lattice of 16x16 knots regularly distributed on our 512x512 pixel grid, and we then defined the displacement field u by :

$$u(x) = \sum_{k \in \mathbb{Z}^2} \beta\left(\frac{x}{h_1} - k_1 + \frac{1}{2}\right) \beta\left(\frac{y}{h_2} - k_2 + \frac{1}{2}\right) \quad (5.1)$$

where $(k_1 = k_2 = 16 \text{ and } h_1 = h_2 = 512/16 = 32)$. The spatial deformation $f(x) = x + u(x)$ is displayed in figure 5.1.1.

In each of our 3 simulated test cases, the true reference image I_r and the corresponding true warped image I_t were separately corrupted by addition of a simulated noise having a log-Rayleigh pdf with known variance, to generate the “observed”

images J_r and J_t . More precisely the corrupted image J is deduced from the true image I by the formula

$$J(x) = I(x) + \frac{1}{\alpha} \ln(N(x)) \quad (5.2)$$

where the random variables $N(x)$ are independent and have the same Rayleigh pdf with variance 10. We performed the preceding simulations for different values of the parameter α , which as seen above determines the theoretical common variance $varLSD$ of all the random log-speckle disparities $\eta(x) = J_t(x + u(x)) - J_r(x)$ by the formula

$$varLSD = var(\eta(x)) = \frac{\pi^2}{12\alpha^2}$$

5.1.2 Performances of deformation estimates on simulated cases

For each one of the 3 cases above, we then started from the pair of simulated images J_r, J_t considered “observed” images and we computed an estimation $\hat{u}$ of the true displacement field u by the multiresolution gradient descent described above. We define the angular errors $\theta(x)$ and the relative magnitude errors $d(x)$ by

$$\begin{aligned} \cos(\theta(x)) &= \frac{\hat{u}(x) \cdot u(x)}{||\hat{u}(x)|| ||u(x)||} \\ d(x) &= \frac{||\hat{u}(x)|| - ||u(x)||}{||u(x)||} \end{aligned} \quad (5.3)$$

Tables 1 and 2 respectively display the average of the angular errors and of the relative magnitude errors, as functions of the α of the log-speckle disparities.





Figure 5.1: Differently smoothed test images (Top: case 1, bottom: case 3)




Figure 5.2: Warping field Example

Table 5.1: Angle error (in degrees)

	$\alpha = 0.1$	$\alpha = 0.5$	$\alpha = 1$	$\alpha = 2$
case1	5.69	4.65	4.80	4.66
case2	6.57	4.71	5.01	4.86
case3	7.52	5.23	5.02	4.94

Table 5.2: Relative magnitude error (in percents)

	$\alpha = 0.1$	$\alpha = 0.5$	$\alpha = 1$	$\alpha = 2$
case1	9.40	8.06	8.26	8.05
case2	10.75	8.32	8.52	8.33
case3	13.39	8.80	8.64	8.44

Note that as the reference image becomes smoother (from case 1 to 3), the estimation error becomes larger. This can be explained by our theoretical analysis above, since our estimate of the bias on deformation errors $\delta(X)$ is given by $Bias(X) = -(I''(f(X)))^{-1}I'(f(X))$. Since I_r is the convolution of image I_0 by a gaussian kernel with large variance σ^2 we have the (quite rough) approximations $I'_r \approx \frac{1}{\sigma}I'_0$ and $I''_r \approx \frac{1}{\sigma^2}I''_0$. This shows that $Bias(X)$ is roughly multiplied by σ and hence increases roughly linearly as the smoothing parameter σ increases.

5.1.3 The case of affine image intensities

We have established in the preceding section that the actual log-speckle disparities in the myocardium have fairly large standard deviation $\frac{\pi}{\sqrt{12\alpha}}$, and take values larger than 10 for $\alpha < 0.1$, while the dynamic range of our ICE images is typically of the order of 200. This led us to evaluate the impact of using various values of α in our multi-resolution gradient descent, for a given fixed variance of the log-speckle disparities.

As we pointed out above local linearities for the true intensity values tend to generate larger bias for the errors on estimated deformations. We have thus tested deformation estimates on a simulated pair of true images I_r, I_t where the reference image intensities $I_r(x)$ were given by a fixed affine function of the pixels x in our (512x512) pixel grid. The true non-linear displacement field u was identical to the field used in section 4.1 above. The images $I_r(x)$ and $I_t(x) = I_r(x + u(x))$ were

Table 5.3: Errors with different α

	$\alpha = 0.1$	$\alpha = 0.5$	$\alpha = 1$	$\alpha = 2$
Angle error (in degrees)	15.69	27.06	7.23	13.96
Relative magnitude error (in percents)	79.80	77.21	69.51	86.46

then separately corrupted by independent additive simulated log-speckle noises corresponding to a real α equal to 0.1, to generate simulated “observed” images J_r and J_t , just as above. We then tested several versions of the gradient descent corresponding to various values of α .

Table 3 displays the average angle errors and relative magnitude errors on the true displacement field u . Note that the performances observed for very small $\alpha = 0.1$ in the gradient descent method are worse than for $\alpha = 1$. The theoretical analysis given above in section 3.6 shows that since the hessian $I_r''(x)$ is here identical to 0, the components of the errors in the direction orthogonal to the fixed gradient vector $I_r'(x)$ cannot be controlled deformation estimates.

5.2 Estimation of speckle noise characteristics on real ICE image sequences

We have estimated the local statistical characteristics of the log-speckle disparities $\eta(x)$ and plotted in figure 5.3 the speckle noise std distribution on the real image, using the algorithm described above in section 4.3.5. For each pixel, we compute the speckle noise in the window of size 60 by 60 pixels centered at it, and then the std can be computed from these 3600 samples. And, we didn’t show the speckle

noise std distribution outside the myocardium because the noise distribution in this large area cannot verify the generalized rayleigh model at all, as shown later. We note that the local stds of the log-speckle disparities $\eta(x)$ are approximately constant within the myocardium, but, for the heart chamber areas, which are blood filled, the log-speckle disparities have much smaller std than in the myocardium. The std distribution of speckle noise is determined by the tissue structure of the myocardium and deserves more anatomical comments.

We perform the Kolmogorov-Smirnov Goodness of Fit test for the speckle noise on 11 different regions $\{Reg_i\}, i = 1, 2, \dots, 11$ of real ICE images (refer to figure 5.5), and compute the Kolmogorov-Smirnov statistic ('KS':defined as the maximum absolute difference between the empirical cumulative distribution function(cdf) and the theoretical cdf. Then, we compute the Signal to Noise Ratio (SNR=Mean intensity/standard deviation of speckle noise) for all these 11 regions. In table 5.4, we show the std, MI(mean intensity), SNR, and KS for all the regions in figure 5.5. We find that the KS statistics for region one to ten is very small, which means the speckle noise distribution can be fitted to generalized rayleigh model very well with associated α value, however, the speckle noise in region 11 has a very large KS statistic when fitted to the generalized rayleigh model.

Since the algorithms of section 4.3.5 enable the direct estimation of all the $\eta(x)$, we have also estimated the common pdf $\hat{\rho}_i$ of the $\eta(x)$ for all the pixels x in region Reg_i after insertion of the proper $varLDS_i$ estimate, and hence of the proper estimate $\hat{\alpha}_i$ in the theoretical pdf ρ of the $\eta(x)$ computed above in (4.28), we also obtain separately the theoretical pdf ρ_i of the $\eta(x)$ in region Reg_i . We have thus confirmed

Table 5.4: Statistical properties of different regions

Regions	STD	MI	SNR	KS
1	10.63	83.13	7.82	0.06
2	10.15	95.51	9.41	0.02
3	13.45	85.52	6.35	0.03
4	14.70	54.35	3.70	0.03
5	12.25	50.35	4.11	0.06
6	11.53	47.43	4.11	0.04
7	11.46	64.80	5.66	0.02
8	10.49	80.97	7.72	0.03
9	11.18	64.76	5.79	0.03
10	7.35	17.29	2.35	0.04
11	2.10	2.24	1.07	0.28

that for all regions Reg_i within the myocardium, the empirical pdf $\hat{\rho}_i$ of the log-speckle disparities has a good fit with the theoretical pdf ρ_i . Figure 5.6 (left) shows the quality of fit for region Reg_8 for example. For the heart chamber areas, although the empirical noise has a smaller std than that of the myocardium, their empirical pdf $\hat{\rho}_{10}$ still have a good fit with the theoretical pdf ρ_{10} .

For the large area (such as Reg_{11}) located outside the heart, the empirical pdf of log-speckle disparities no longer fits their theoretical pdf, as seen in figure 5.6 (right).

The main operational conclusions here are that the theoretical log-Rayleigh models, when adequately parameterized by image specific estimates, are quite correct to separately model the log-speckle disparities within the myocardium and the heart chamber, but are useless outside these regions. In particular this confirms that adequate masking by semi-local masks is essential to improve the accuracy of the cost functionals and of deformation estimates.

The usual assumption inferred from the physics of ultrasound speckle is that the

Table 5.5: Correlation Coefficient

Pixel	1	2	3	4	5	6
Radial	0.28	0.04	0.08	0.07	0.02	0.10
Circumferential	0.85	0.53	0.27	0.11	0.02	0.11

speckle noises $N(x)$ and $N(y)$ at distinct pixels x and y are independent, and hence uncorrelated random variables. This implies that the correlation $cor(x, y)$ of log-speckle disparities $\eta(x)$ and $\eta(y)$ at distinct pixels x, y must be equal to zero. To evaluate this standard assumption we have used our speckle noise estimation algorithms (see section 4.3.5 above) to estimate the log-speckle disparities on real ICE images, and then to computed estimates of the correlations $cor(x, y)$. In particular since the ICE images are acquired by a rotating ultrasound beam centered on a temporarily fixed transducer placed at a central point CP easily identifiable on all ICE images, we have focused the empirical estimations of $cor(x, y)$ on pairs of neighboring pixels (x,y) located on the same circle centered at CP (circumferential correlations) or on the same radial straight line through CP (radial correlations).

As seen in table 5.5, the main conclusions are that, within the myocardium, radial correlations for log-speckle disparities are indeed negligible, but that circumferential correlations are quite strong at short distances. Further research is necessary to analyze the impact of this fact on our cost functionals, and on the performance of the gradient descent algorithm.

After estimation of the speck noise characteristics, we verified the generalized rayleigh model and found that the empirical noise is much stronger than the assumption by standard rayleigh model (factor $\alpha < 0.1$ in our true images). This led

us to try using the actual very small α in the optimization algorithm, however, the result of strain estimation is much worse than using $\alpha = 1$. While we have verified the excellent performance of our algorithm when $\alpha = 1$, we found using small values of α (less than 1) always led to worse result after intensive numerical simulations on our true ICE images. Thus, we conclude that $\alpha = 1$ is a good choice for first round optimization for the ICE images, but a better choice of α might be needed to improve the performance for different kinds of images and the effect of α in the algorithm deserves further discussion.

5.3 Time modulations of log-speckle and ICE image intensities

The initial cost functional introduced above implicitly assumes that between the reference and target images, the true heart muscle deformation mapping x on $x+u(x)$ preserves the true image intensities, so that $I(x) = I(x + u(x))$. To evaluate this commonly used hypothesis we computed the histograms of displayed ICE intensities over small masks for all views available, and found that the empirical histograms $I(x)$ had a small but very concrete shift over time, with *shift values of the median intensities reaching 10% of the mean intensities at half heart cycle*. We conclude that temporal adaptation should be introduced in the cost functional, by *performing local histogram equalizations over time* before implementing the gradient descent for minimization of the cost functional.

For small image patches in our reference ICE view, we analyzed the time evolution of the estimated scale parameters α through the heart cycle. Our estimate indicated a clear time modulation of this log-speckle scale parameter which reached its maximum roughly at mid-cycle.

5.4 Estimated strains versus sonomicrometric ground truth

The sonomicrometric data recorded by Dr. Khoury's team on three groups of ICE image sequences IS_1 , IS_2 , IS_3 , provide effective tracking of two microcrystals surgically inserted in the myocardium, and hence an accurate evaluation of circumferential strain at successive crystal positions.

Call $d(t)$ the true distance between the two crystals at time t . The true strain at time t (at the corresponding crystal position) is then defined by $S(t) = |d(t) - d(0)|/d(0)$. The mean values of these true (relative) strains over the whole image sequences for the 3 data groups IS_1 , IS_2 , IS_3 , are respectively 5.4%, 4.1%, 6.5% , and these strains remain in the respective ranges $[0, 16.8\%]$, $[0, 10.5\%]$, $[0, 17.7\%]$. After reconstruction of the dynamic deformation field as above, we can easily compute an estimate $\hat{d}(t)$ of $d(t)$, and hence an estimate $\hat{S}(t)$ of $S(t)$. We have a very good agreement between these true strains $S(t)$ and their algorithmic registration estimates $\hat{S}(t)$. The estimation errors $Err(t) = \hat{S}(t) - S(t)$ for the 3 data groups IS_1 , IS_2 , IS_3 , had respective mean values 0.4%, 0.4%, 0.5% and standard deviation of 0.76%, 0.71%, 0.92% . These are small errors as compared to the mean true strain values just listed

above.

Our strain estimates clearly improve the accuracy previously reached in [59], where for instance for the data group IS_1 , the mean error of strain estimates was 0.7% with a standard deviation of 1.63%.

This accuracy gain is mainly due to our masking technique and to the use of an image specific cost functional, where the parameters of the speckle noise are estimated at each image pixel. This validates the conclusion of section 7.2 stating that only regions within the myocardium have a good fit with our parameterized stochastic speckle model. The scale parameter γ we have introduced in both the speckle model and cost functional exhibits a clear spatial modulations across the heart regions, therefore adequate masking and image adaptive cost functionals are essential to improve the accuracy of heart deformation estimates.

5.5 Computing time optimization

We have applied the preceding algorithms to three groups of data, each contains a sequence of 180 ICE images acquired on one animal subject by Dr D. Khoury's team at Methodist Hospital, after surgical insertion of 2 microcrystals in the animal heart muscle, and sonomicrometric recording of the displacements of these crystals. For each sequence of 180 ICE images, the brutal computation with no mask applied needs averagely 240 hours, which is a massive computing time, therefore, we try to accelerate the optimization by applying small masks.

Table 5.6 shows the comparison of average computing time (in minutes) per image, where single computation means using no masks through all the seven resolution levels, and mixed computation means using no mask for the first three levels and then applying a small mask (as shown in figure 5.7, whose area is only 6 percent of the whole image) for the rest four levels. To reduce the computing time, we proposed to use the small mask instead of the whole image in the optimization algorithm, because the computing time is approximately proportional to the number of pixels used under the same stopping criteria, furthermore, by using a small mask, we can even use a cruder stopping criteria to arrive at the same accuracy. We here proposed a mixed masks computation scheme, where we set $\epsilon = 10^{-4}$ for the first three levels where no mask is applied and then set $\epsilon = 10^{-3}$ for the remaining four levels when the small mask is applied. We use the small mask only for the final four levels instead of all the seven levels, because using the whole image first in the computation can give us a good initialization for the optimization process because the node lattice is distributed on the whole image, and an inappropriate initialization will lead to waste of computing time or even getting trapped at local minima. To find the best tradeoff between accuracy and speed, we actually tried all the possible mixtures of masks sizes by using no mask in the first one to six levels, and the best solution is using no mask in the first three multiresolution levels and then using small mask in the remaining four levels, and the final gain of computing speed is about 64 times!

The main operational conclusions is that our algorithm should be used for semi-local computations on quite small masks and can then provide accurate estimates of local circumferential and radial strains of the heart muscle. The systematic use of

Table 5.6: Computing time per frame(in minutes)

Different groups	Mixed computation	Single computation
1	1.16	80.68
2	1.40	83.93
3	1.33	82.94

a family of very small masks with slight overlaps and regularly positioned to cover the region of interest, reduces considerably the computing time of the local circumferential and radial strains on the whole image. We are currently systematically implementing and optimizing this semi-local computational scheme. For instance to compute the local strain at all points of a medium ROI covering 25% of the 512x512 image we can cover this area with approximately $16 \times 16 = 256$ local masks of size $20 \times 20 = 400$. On such small masks one can safely use a higher threshold $\epsilon = 10^{-2}$. The computing time on each small mask at accuracy $\epsilon = 10^{-2}$ will be of the order of 2.25 seconds. And the total computing time for the medium size ROI considered here will be roughly $256 \times 2.25 = 576$ seconds instead of the 1570 seconds needed to handle the ROI with a single mask at a necessarily sharper threshold $\epsilon = 10^{-3}$.




Figure 5.3: Speckle noise std distribution(top) and mean intensity(bottom) within the myocardium and heart chamber: at the beginning of systole




Figure 5.4: Speckle noise std distribution(top) and mean intensity(bottom) within the myocardium and heart chamber: at the end of systole




Figure 5.5: Different regions to compute speckle noise variance




Figure 5.6: Empirical pdf v.s. theoretical pdf on different regions: region 8 (top) and region 11 (bottom)




Figure 5.7: Examples of a small mask on real ICE image




Figure 5.8: Examples of deformation vector field

Chapter 6

Speckle tracking in 3D ultrasound images

Clinical diagnosis and therapy planning are increasingly often supported by 3D-imaging modalities, such as MRS (Magnetic Resonance Spectroscopy), PET (Positron Emission Tomography), SPECT (Single Photon Emission Computed Tomography) for functional information, and CT (Computed Tomography), MRI (Magnetic Resonance Imaging), Ultrasound Echography, X-Ray, for anatomical visualization.

Thus clinicians and medical researchers become natural users for 3D-image registration providing voxel to voxel matching of two 3D-images of the same anatomical object obtained by different imaging modalities, at different times, or from different perspectives. The search for a good voxel to voxel correspondence between reference and target images J_r and J_t , is guided by one or several matching quality criteria. Image registration methods were initially designed for 2D-images, for instance to

align tomographic slices of different recordings, but in the last decade, 3D-image registration based on volumetric data sets has become the main technical challenge, and involves much heavier computing resources.

We apply speckle tracking techniques to numerically construct the dynamic deformations between multiple 3D-snapshots of the human mitral valve annulus in the mitral valve apparatus. This apparatus is a biological valve integrated within the heart which has quasi-periodic dynamics tightly correlated to heartbeats enabling an essential periodic obturation of the blood flow to and from the heart. Our starting point is a patient’s specific set of static models of the mitral valve apparatus. These models were generated by image analysis of live 3D-echocardiographic movies at specific heart cycle instants in a recent study [67] involving the Methodist Hospital, Houston, Texas (S. Ben Zekry, S. Little, W. Zoghbi, MDs) as well as a research group on mathematical image analysis led by R. Azencott. Each 3D-movie includes roughly twenty 3D-frames per heartbeat cycle, acquired by ultrasound technology, and represents a high volume of image data corrupted by ‘speckle’ noise. Figure 6.1 shows several 2D slices of one 3D-snapshot viewed from different directions.

6.1 Mitral Valve

The mitral valve is a dual-flap valve in the heart that lies between the left atrium (LA) and the left ventricle (LV). The mitral valve and the tricuspid valve are known collectively as the atrioventricular valves because they lie between the atria and the ventricles of the heart and control the flow of blood. A normally-functioning mitral

valve opens secondary to increased pressure from the left atrium as it fills with blood. As the pressure increases above that of the left ventricle, the valve opens allowing blood to flow into the left ventricle during diastole (early rapid filling and atrial contraction), and closes at the end of atrial contraction to prevent blood flowing back.

Mitral valve disease represents the second most common valvular disease in developed countries [41]. Minimally invasive mitral valve repair procedures are mostly under development or in trial yielding the need for precise knowledge and reliable display of the four-dimensional valve characteristics. Computed tomography and 4D ultrasound are modalities well suited for non-invasive imaging of the heart enabling dynamic four-dimensional evaluation of cardiac structures throughout the cardiac cycle.

The mitral valve's central components are the anterior and posterior leaflet, the annulus and the subvalvular apparatus. The latter two connect the valve to the left ventricular (LV) endocardium. The annulus is a ring-like fibrous entity with a 3D shape resembling a saddle (Fig. 6.2), with the middle portions of the anterior annulus being elevated out of the annular plane towards the left atrium (LA) and merging into the aortic mitral curtain. The aortic mitral curtain ends in the left and right fibrous trigone. The posterior leaflet is divided through slits into several scallops. Multiple nomenclatures exist, most commonly the scallops are named P1 to P3 with opposing segments A1 to A3 [42]. The mitral commissures are the points where both leaflet's free edges join. They do not coincide with the mitral annulus. The subvalvular apparatus consists of the chordae tendinae and the papillary muscles.

The function of the mitral valve is to allow blood to flow only one way, from the LA to the LV, separating them hemodynamically. It is opened by the contracting papillary muscles, which are pulling down the leaflets during diastole and remains closed by the ventricular pressure during systole.

6.2 B-spline model for the displacement vector field u

Consider a lattice of knots $L = \{k = (k_1, k_2, k_3) | 1 \leq k_1 \leq L_1, 1 \leq k_2 \leq L_2, 1 \leq k_3 \leq L_3\}$, where the knots k are placed on a regular subgrid of the pixel grid G . Denote by $N_1 = 224$, $N_2 = 208$, $N_3 = 208$ the dimensions of the grid G and define the knot spacing $h = (h_1, h_2, h_3)$ where $h_1 = N_1/L_1$, $h_2 = N_2/L_2$ and $h_3 = N_3/L_3$.

The unknown displacement field u is a 3D non-linear transformation which we modelled by a linear combination of cubic B-splines associated to the knots of the lattice L . The unknown vector c of coefficients $c_k = (c_k^1, c_k^2, c_k^3)$ of this linear combination completely defines the displacement field u (and hence the spatial deformation f) by the explicit formula :

$$u(x) = \sum_{k \in L} c_k B_3(x, h, k) \quad (6.1)$$

where

$$B_3(x, h, k) = \beta\left(\frac{x_1}{h_1} - k_1 + \frac{1}{2}\right) \beta\left(\frac{x_2}{h_2} - k_2 + \frac{1}{2}\right) \beta\left(\frac{x_3}{h_3} - k_3 + \frac{1}{2}\right) \quad (6.2)$$

and $\beta(s)$ is 1D cubic B-spline function.

During the warping process, we need a continuous version of the discrete image J_t . Using also cubic B-splines, we interpolate the discrete image J_t to extend it from the pixel grid G to a continuous function still denoted J_t :

$$J_t(x) = \sum_{y \in G} b_y B_3(x - y) \quad (6.3)$$

where $B_3(x - y) = \beta(x_1 - y_1)\beta(x_2 - y_2)\beta(x_3 - y_3)$, and b_y is a set of interpolation coefficients, which can be computed prior to the search for u by solving a linear system.

6.3 Multi-resolution optimization of cost functionals

To compute the unknown spatial displacement field u matching images J_t and J_r , we need to minimize the cost functional $K(u)$. To minimize $K(u)$ we implement a multiresolution gradient descent in u at coarse to fine successive resolutions. Multi-resolution is used for both the image model and the deformation model.

First, we build an image pyramid composed of gradually coarser versions of the original image by regularly spaced down-sampling. Starting from an initial ICE image of size $224 \times 208 \times 208$, we create, by successive dichotomic decimations of pixels, a sequence of increasingly coarser images.

Second, we use multi-resolution also for the B-spline deformation model of u . We start with a coarse B-spline model of u defined by a lattice of knots L_1 with a large

initial distance h^1 between knots. We fix an increasing sequence of knots lattices L_i with corresponding knots spacing $\{h^i\}_{i=1,\dots,q}$, where $h^{i+1} = 2h^i$.

Both multi-resolutions do improve accuracy as well as speed of the gradient descent algorithm. Instead of refining both the image and the deformation models at the same rate, we dynamically define a succession of pairs of resolution scales $RS(m) = (N^{i(m)}, h^{j(m)})$ for the image and the deformation models, starting with the coarsest versions of these 2 models. More precisely, once the current pair of resolutions $RS(m)$ has been selected, we interpolate at resolution $RS(m)$ the last computed approximation u^{m-1} of u , just obtained at coarser resolution $RS(m-1)$, and use this interpolated u^{m-1} to initiate and then perform a gradient descent on $K(u)$ computed at the new resolution $RS(m)$. This gradient descent at fixed resolution is detailed below, and terminates automatically at a new approximation u^m of u . We then refine $RS(m)$ into a finer pair of resolutions $RS(m+1)$ by alternate refining of either the image model or the B-spline model. Successive modifications of resolutions stop when the predefined finest level is reached.

6.4 Detailed gradient descent at fixed resolution

To minimize the cost functional $K(u)$ we implement a multiresolution gradient descent in u at coarse to fine successive pairs of resolutions $RS(m)$, both for the image model and for the B-spline model. For each fixed pair of resolutions $RS(m)$, the current lattice of B-spline knots L is fixed, and hence u is approximated by its corresponding B-spline expansion, which is defined by an unknown vector of coefficients

c_k^i , ($i = 1, 2, 3; k \in L$). After adequate restriction to the current image model, the cost functional to minimize becomes an explicit function $K(c)$. To minimize $K(c)$ we implement a gradient descent in c , with the following update rule :

$$\Delta c = -\mu \nabla K(c) \quad (6.4)$$

where $\nabla K(c)$ is the gradient of $K(c)$, μ is the step size. This dynamic step size μ is multiplied by a fixed factor κ after a successful step, i.e a step which actually decreases the cost functional, and is divided by κ after an unsuccessful step. In most of our numerical implementations we selected $\kappa = 10$.

The differential of the cost functional can be derived easily for ($i, j, l, m = 1, 2, 3; l \neq j \neq m$):

$$\begin{aligned} \frac{\partial Kdisp}{\partial c_k^i} &= \frac{\alpha}{card\{MSK\}} \sum_{x \in MSK} \frac{\exp(2\alpha Disp(x)) - 1}{\exp(2\alpha Disp(x)) + 1} \\ &\quad \cdot \left. \frac{\partial J_t(f)}{\partial f_i} \right|_{f(x)=u(x)+x} \frac{\partial u^i}{\partial c_k^i} \end{aligned} \quad (6.5)$$

where

$$\frac{\partial J_t(f)}{\partial f_j} = \sum_{y \in G} b_y \beta'(u_j + x_j - y_j) \beta(u_l + x_l - y_l) \beta(u_m + x_m - y_m) \quad (6.6)$$

$$\frac{\partial u^i}{\partial c_k^i} = B_3 \left(\frac{x}{h} - k \right) \quad (6.7)$$

We also have

$$\frac{\partial Kreg}{\partial c_k^i} = \frac{2}{card\{MSK\}h} \sum_{x \in MSK} \sum_{j, l \neq i}^2 \frac{\partial u^i}{\partial x_j} \beta' \left(\frac{x_i}{h_i} - k_i + \frac{1}{2} \right) \beta \left(\frac{x_j}{h_j} - k_j + \frac{1}{2} \right) \beta \left(\frac{x_l}{h_l} - k_l + \frac{1}{2} \right) \quad (6.8)$$

Two simultaneous stopping criteria are implemented to terminate the gradient descent : the descent process is stopped at iteration step i either if the number of iterations i exceeds a preset large integer, or if $\Delta K^i < \epsilon$, where $\Delta K^i = |K(c^i) - K(c^{i-1})|/K(c^{i-1})$ is the relative increment of the cost function at the last iteration step. The size ϵ of the terminating threshold is selected to minimize the computing time, so that we select different values $\epsilon(m)$ for ϵ at each resolution level $RS(m)$. Indeed at coarser resolutions, ϵ may safely take higher values than at finer resolutions.

6.4.1 Image sequence analysis and Masking

To compute the displacement fields u_t of an image sequence $J_t, t = 1, \dots, n$ covering a whole cardiac cycle, we apply successively the registration algorithm to compute an optimal registration REG_t matching the consecutive pair of images J_{t-1} and J_t for $t = 2, \dots, n$. The global displacement field J_t is then obtained by successively applying the registrations $REG_1, \dots, REG_t$. Each small deformation generated by REG_i is verified to be invertible, therefore the global large deformation is an invertible deformation. For each image J_{t-1} , temporarily considered as a reference image, we determine an initial mask MSK corresponding to the zone of interest in the myocardial area . A few key boundary pixels of MSK can be easily selected interactively and used to generate a mask MSK with a piecewise linear boundary linking these initial vertices.

6.5 Numerical performance on 3D ultrasound images

Our experiment is based on a patient's specific finite set of static models S_j of the mitral valve apparatus. These models were generated by image analysis of live 3D-echocardiographic movies at specific heart cycle instants t_j in a recent study involving the Methodist Hospital, Houston, Texas (S. Ben Zekry, S. Little, W. Zoghbi, MDs) as well as a research group on mathematical image analysis led by R. Azencott. Each 3D-movie includes roughly twenty 3D-frames per heartbeat cycle, acquired by ultrasound technology, and represents a high volume of image data corrupted by 'speckle' noise. The mitral valve models S_j are based on NURBS (non uniform rational B-splines), and were obtained by combining optical flow extraction algorithms with sparse tagging by medical experts. The number of intermediary key heartbeat cycle instants t_j ranged from 3 to 8, and the time range $I = [t_0; t_1]$ covers either a half or a whole heartbeat cycle with total duration between 1/2 to 1 second.

The mitral valve apparatus involves the annulus (a closed thin deformable ring) and two deformable surfaces with boundaries, namely the anterior and posterior leaflets. These mitral leaflets are flexibly linked to the annulus by a subsegment of their boundaries. When the valve is closed, the exterior parts of the leaflets have a common boundary called the coaptation line. The mitral valve apparatus can be viewed as a composite deformable object built from several smooth deformable shapes (see Fig. 6.5), namely a closed curve MA (the mitral annulus), a curve segment COA (the coaptation line), two surfaces AL and PL (the mitral leaflets) with boundaries

$\partial\text{AL} = \text{COA} \cup \text{antMA}$ and $\partial\text{PL} = \text{COA} \cup \text{postMA}$, where antMA and postMA are complementary subsegments of the MA.

In the numerical implementation, we set the regularization coefficient γ to be zero since regularization should be tailored as specific goal is desired and we evaluate here the performance of our algorithm in general. We use tubes of different radius around the annulus as the masks in the multiresolution optimization, an examples of the mask we used are shown in figure 6.5.

We use the annulus model at first frame to generate the mask for optimization, and validate our registration method at every time frame against the annulus models generated from the doctors' tagging at other frames in the 3D image sequences. Fig. 6.5 shows the evolution of cost function in the optimization process. At the end of optimization, the value of cost function is very close to the theoretical value 1. Fig. 6.5 shows the histogram of distances between the reference and the target annulus generated from doctors' tagging by NURBS models. The mean and std of the distances are 11.08 and 3.53 respectively. After optimization, the results show good agreement (Fig. 6.5) between deformed annulus and target annulus with mean 2.05 and std 1.16.




Figure 6.1: Example: 2D slices of 3D images




Figure 6.2: Mitral valve illustration




Figure 6.3: Mitral valve: the middle line is the coaptation line separating the two surfaces. The left surface in the figure is the anterior leaflet and the right surface is the posterior leaflet. The black thick curve is the annulus




Figure 6.4: Examples of mask in 3D image




Figure 6.5: Evolution of cost function in optimization




Figure 6.6: Histogram of distances between reference and target annulus




Figure 6.7: Histogram of distances between deformed and target annulus

Chapter 7

Diffeomorphic Matching and Dynamic Deformable Surfaces in 3D Medical Imaging

We consider optimal matching of submanifolds such as curves and surfaces by a variational approach based on Hilbert spaces of diffeomorphic transformations. In an abstract setting, the optimal matching is formulated as a minimization problem involving actions of diffeomorphisms on regular Borel measures considered as supporting measures of the reference and the target submanifolds. The objective functional consists of two parts measuring the elastic energy of the dynamically deformed surfaces and the quality of the matching. To make the problem computationally accessible, we use reproducing kernel Hilbert spaces with radial kernels

and weighted sums of Dirac measures which gives rise to diffeomorphic point matching and amounts to the solution of a finite dimensional minimization problem. We present a matching algorithm based on the first order necessary optimality conditions which include an initial-value problem for a dynamical system in the trajectories describing the deformation of the surfaces and a final-time problem associated with the adjoint equations. The performance of the algorithm is illustrated by numerical results for examples from medical image analysis.

7.1 Introduction

Clinical diagnosis and therapy planning are increasingly often supported by 3D-imaging modalities, such as MRS (Magnetic Resonance Spectroscopy), PET (Positron Emission Tomography), SPECT (Single Photon Emission Computed Tomography) for functional information, and CT (Computed Tomography), MRI (Magnetic Resonance Imaging), Ultrasound Echography, X-Ray, for anatomical visualization.

Thus clinicians and medical researchers become natural users for automated 3D-image registration providing voxel to voxel matching of two 3D-images of the same anatomical object obtained by different imaging modalities, at different times, or from different perspectives. The search for a good voxel to voxel correspondence between reference and target images J_{ref} and J_{tar} , is guided by one or several matching quality criteria. Image matching is generally achieved by an R^3 -diffeomorphism F matching two given bounded subdomains of the 3D-voxel grid, and can thus be

assigned an elastic energy $EE(F)$ measuring the amount of spatial deformation implemented by F . In 2D or 3D-image registration, typical matching quality criteria involve the differences $difint(z, z') = |J_{tar}(z') - J_{ref}(z)|$ in image intensities at all pairs (z, z') of matched voxels. The Intensity Matching Cost $IMC(F)$ is often defined by fixing some exponent $a > 0$ and summing $difint^a(z, z')$ over all voxels z belonging to the domain of interest in J_{ref} . The search for an optimal registration then becomes a variational problem where one seeks a deformation F minimizing a linear combination of $EE(F)$ and $IMC(F)$. Image registration methods were initially designed for 2D-images, for instance to align tomographic slices of different recordings, but in the last decade, 3D-image registration based on volumetric data sets has become the main technical challenge, and involves much heavier computing resources. Surveys of image registration algorithms can be found in [76, 89, 100].

7.2 Diffeomorphic Shape Matching

7.2.1 Diffeomorphic Matching of Two Shapes in R^3

In most medical imaging applications, different 3D-shapes $B \subset R^3$ of the same deformable organ, such as the brain or the heart, can be assumed to belong to the following family SH3 of *smooth 3D-shapes with boundaries*: We define the family SH3 of *3D-shapes* as the set of all connected open subsets S of R^3 with compact closures $\bar{S}$ such that

- the boundary $\Sigma = \partial S = \bar{S} - S$ of S is a piecewise smooth surface of class

$$C_r, r \geq 3 ;$$

- S coincides with a whole connected component of $R^3 \setminus \Sigma$;
- for each $x \in \Sigma$, there is an open neighborhood U of x in R^3 and a local r -smooth diffeomorphism φ of U onto an open ball $\tilde{U} \subset R^3$ mapping $U \cap \bar{\Sigma}$ onto $\tilde{U} \cap H$, where H is the intersection of either one, two or three closed half-spaces of R^3 .

In medical 3D-imaging, one of the main goals of image registration is to accurately compare two observed 3D-shapes $S_{ref} \in \text{SH3}$ and $S_{tar} \in \text{SH3}$ of the same deformable organ. These 3D-shapes are often extracted from two 3D-images J_{ref} and J_{tar} by various 3D-image analysis techniques such as 3D-segmentation combined with reference points interactive tagging. We refer to Σ_{ref} and Σ_{tar} as the boundaries of S_{ref} and S_{tar} . Since these two 3D-shapes are in SH3, any R^3 -diffeomorphism F such that $F(\Sigma_{ref}) = \Sigma_{tar}$ must necessarily also verify $F(S_{ref}) = S_{tar}$. Hence, in the search for a diffeomorphic matching of two deformable 3D-shapes, whenever these shapes have been already pre-extracted as R^3 -subsets, the matching of image intensities is then often discarded, and the matching quality of candidate R^3 -diffeomorphisms F is focused on suitable geometric distances between the boundary surfaces $F(\Sigma_{ref})$ and Σ_{tar} as well as between the 3D-shapes $F(S_{ref})$ and S_{tar} .

As just indicated, diffeomorphic matching of 3D-shapes S_0, S_1 in R^3 quite naturally involves looking at the matching of their boundaries Σ_0, Σ_1 which are 2D-shapes in R^3 as well as to the matching in R^3 of the boundary 1D-shapes $\partial\Sigma_0, \partial\Sigma_1$, where the definitions of smooth 2D and 1D-shapes in R^3 are similar to the definition of SH3.

7.2.2 Variational Approach

Consider two smooth k -dimensional shapes S_0, S_1 included in R^3 with $k \in \{1, 2, 3\}$. The search for an 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. Such an approach has been actively explored by G. Dupuis, J. Glaunès, U. Grenander, M. Miller, D. Mumford, A. Trounev, L. Younes [75, 80, 85, 86] with applications to a quantified comparison of images of human brains. In these papers, the regularization is achieved through the replacement of the rigid constraint $F(S_0) = S_1$ by a soft constraint based on various geometric 'surface matching' distances $\text{dis}[F(S_0), S_1]$. The unknown diffeomorphism F is restricted to be of the form $F = F^v$, where F^v is generated by integration between times 0 and 1 of some time dependent flow $v = (v_t)$ of smooth R^3 -vector fields v_t , $0 \leq t \leq 1$. The vector fields v_t are required to belong to a Hilbert subspace V of the Banach space C_r^3 of smooth functions from R^3 to R^3 tending to zero at infinity. The Hilbert norm in V is assumed to be bounded by a constant multiple of the Banach norm in C_r^3 . Then, for some fixed constant $\lambda > 0$ one considers the variational problem of finding a vector field flow $v = (v_t)$ minimizing the cost functional

$$J(v) = \int_0^1 \|v_t\|_V^2 dt + \lambda \text{dis}[F^v(S_0), S_1],$$

which linearly combines a kinetic energy term and a surface matching term.

This variational point of view is directly linked (as $\lambda \rightarrow \infty$) to the construction of *geodesics in infinite dimensional Lie groups of diffeomorphisms* in the spirit of ideas

pioneered by Arnold, Ebin, and Marsden who showed (see, e.g., [69]) that for an incompressible fluid, obeying Euler equations, the spatial displacements $F_t(x)$ between times 0 and t of fluid particles emanating from $x \in R^3$ minimize the integral in time and space of the fluid's kinetic energy. The time dependent R^3 -diffeomorphisms F_t define a continuous path in the group of R^3 -diffeomorphisms, and this path is a geodesic $t \rightarrow F_t$ of an infinite dimensional Lie group G of R^3 -diffeomorphisms, endowed with the local Hilbert metric defined by the fluid's kinetic energy on the Lie algebra of G . This Lie algebra is naturally identified with the Hilbert space of smooth vector fields on R^3 defined by fluid velocities at time 0. The classical Euler fluid mechanical equations for the fluid velocities become precisely interpreted as the variational equations characterizing geodesics in G . Natural right-invariant deformation distances on the group G can then be associated to this Riemannian structure (see, e.g., [98]).

For the diffeomorphic matching of two smooth k -dimensional shapes ($k \in \{1, 2, 3\}$) by R^3 -diffeomorphisms, the geometric view just outlined above has been intensively explored in [75, 80, 85, 98] and numerically implemented for comparisons of key anatomic parts of human brains such as the hippocampus and the temporal lobes [75, 85].

7.2.3 Multiple Snapshots of Dynamic Deformable Shapes

In medical domains such as cardiology, urology, gynecology, clinicians routinely use volumetric 3D-echographs to visualize live 3D-movies of deformable organs. However, computerized algorithms to model such soft organs dynamics by explicit time dependent nonlinear 3D-deformations constitute a quite complex and active research target. For deformable anatomic shapes $S(t) \subset R^3$ indexed by time t , current biomedical research often succeeds to extract from medical 3D-movie data at key time frames $t_j, 0 \leq j \leq q$, a sequence of static models $S_j \subset R^3$ of the shapes $S(t_j)$. Given these $q + 1$ shape *snapshots* $S_j = S(t_j) \subset R^3, 0 \leq j \leq q$, a natural goal is to model the time deformations of the shape $S(t)$ by a time dependent family F_t of R^3 -diffeomorphisms such that

$$F_{t_j}(S_0) = S_j \quad \text{for } j = 1, \dots, q, \quad (7.1a)$$

$$F_{t_0} = Id \text{ (identity mapping of } R^3), \quad (7.1b)$$

$$\text{for each fixed } t \quad [F_t(x) - x] \rightarrow 0 \quad \text{as } x \rightarrow \infty \text{ in } R^3. \quad (7.1c)$$

Most publications mentioned above deal with the basic case $q = 1$ where one wants to match, by an R^3 -diffeomorphism, a single pair of static smooth k -dimensional shapes called reference and target shapes .

In this paper, we will focus on the situation where we are given $q + 1$ 'snapshots' $S_j \subset R^3$ which is a natural context for 3D-movies analysis. In particular, we will extend the variational approach described above to the search of time dependent R^3 -diffeomorphisms F_t verifying at fixed intermediary times t_j the $q > 1$ geometric matching constraints listed in (7.1a)-(7.1c). The $q + 1$ given 'snapshots' $S_j \subset R^3$

are typically smooth 3D-shapes belonging to SH3 (see subsection 7.2.1) or piecewise smooth R^3 -submanifolds of lower dimension $k \in \{1, 2\}$.

We will also present a medical application to a dynamic sequence of Mitral Valve snapshots where each snapshot S_j actually belongs to a more general class of *composite deformable objects* S_j which are unions of several bounded piecewise smooth surfaces and curves in R^3 linked by flexible articulations.

7.3 Optimal Matching of Intermediary Snapshots

In this section, we will use standard notation from Lebesgue and Sobolev space theory (cf., e.g., [97]). Moreover, C will denote a generic positive constant not necessarily the same at each occurrence.

7.3.1 Time Dependent Vector Fields With Finite Kinetic Energy

We choose a Hilbert space V of vector fields on R^3 and we consider the associated Hilbert space $L^2(I, V)$ of vector field flows $v : t \rightarrow v_t \in V$, indexed by a time parameter t in the interval $I = [t_0, t_1]$, having finite *kinetic energy* $E(v)$ defined by

$$E(v) := \frac{1}{2} \|v\|_{L^2(I, V)}^2 = \frac{1}{2} \int_{t_0}^{t_1} \|v_t\|_V^2 dt. \quad (7.2)$$

We assume that the Hilbert space V of R^3 -vector fields is continuously embedded in a Sobolev space $W^{s,2}(R^3)^3$ for some $s > 5/2$. By the Sobolev embedding theorem,

$W^{s,2}(R^3)^3$, $s > 5/2$, is continuously embedded in the Banach space $C^{0, s-3/2}(R^3)^3$ of R^3 -vector fields. We note that in this situation, for each $x \in R^3$, the evaluation map $w \rightarrow w(x)$ from V to R^3 is continuous with respect to both the strong and the weak topology on V .

7.3.2 Dynamic System of Diffeomorphic Deformations

For $t \in I$ and $v = (v_t)$ as above, we define the flow of R^3 -diffeomorphisms F_t as the solution of the flow dynamics equations

$$\partial_t F_t = v_t(F_t) \quad , \quad t \in I, \quad (7.3a)$$

$$F_0 = \text{Id}, \quad (7.3b)$$

where Id refers to the identity map of R^3 .

7.3.3 Self-Reproducing Hilbert Spaces

We recall that a symmetric real valued kernel $K(x, x')$ defined for $(x, x') \in R^3 \times R^3$ is called *positive definite* iff for arbitrary vectors $x_n \in R^3$, $n = 1, \dots, N$, the $(N \times N)$ symmetric matrix $K(x_m, x_n)$ is positive definite. We refer to [70, 96] for classic definitions and results on positive definite kernels K and their associated self-reproducing Hilbert spaces. In our framework, the relevant Hilbert space V of R^3 -vector fields is often defined as the classical *self-reproducing Hilbert space* $V = V_K$ of R^3 -vector fields defined by a smooth symmetric bounded positive definite kernel $K(x, x')$ on $R^3 \times R^3$ being invariant under translations. For the definition of V_K , we

consider the family W of R^3 -vector fields $w_{z,u}$, indexed by arbitrary pairs (z, u) in $R^3 \times R^3$, and defined by $w_{z,u}(x) = K(z, x) u$ for all $x \in R^3$. The vector space LW of finite linear combinations of elements of W is endowed with the pre-Hilbertian scalar product

$$\langle w_{z,u}, w_{z',u'} \rangle = K(z, z') \langle u, u' \rangle_{R^3}.$$

The space V_K is then the unique Hilbert space generated by LW.

For many shape matching applications, K can be a *radial Gaussian kernel* K_σ

$$K_\sigma(x, x') = \frac{1}{(2\pi)^{3/2}\sigma^3} \exp\left(-\frac{\|x - x'\|^2}{\sigma^2}\right) \quad (7.4)$$

with a suitable *scale parameter* $\sigma > 0$. In particular, the Sobolev embedding hypothesis above is satisfied for any $s > 5/2$ when $V = V_{K_\sigma}$ is the self-reproducing Hilbert space associated to a radial Gaussian kernel $K = K_\sigma$. The choice $V = V_{K_\sigma}$ seems to be a good pragmatic choice for diffeomorphic shape matching applications as seen in previous studies and in our numerical implementations below.

7.3.4 Matching Quality

The matching quality is measured as a linear combination of geometric matching error and intensity matching error, i.e.,

$$\text{Disp} = \text{Disp}^G + \mu \text{Disp}^I$$

where $\mu \geq 0$ is the weight of combination.

7.3.4.1 Geometric Matching Error

We consider a given sequence of instantaneous 'shape snapshots' $S_j = S_{t_j} \subset R^3$ generated by a deformable shape S_t at instants $t_j, 0 \leq j \leq q$. Typically, each S_j is a k -dimensional submanifold ($k \in \{1, 2, 3\}$) with a boundary that is regularly embedded in R^3 . Let (F_t) be a candidate family of R^3 -diffeomorphisms indexed by time. To compare each given snapshot S_j with the deformed initial shape $F_{t_j}(S_0)$, a key choice is to define smooth non-negative geometric distances $D(S, S')$ quantifying the geometric disparity between pairs of shapes S, S' in R^3 .

The classical Hausdorff disparities $h(S, S')$ and $h(S', S)$ between subsets S, S' of R^3 are defined by

$$h(S, S') = \max_{x \in S} \left(\min_{x' \in S'} |x - x'| \right).$$

They determine the Hausdorff distance by

$$D_h(S, S') = \max(h(S, S'), h(S', S)). \quad (7.5)$$

The Hausdorff distance introduces theoretical complications in the variational framework below, since $h(S, S')$ is not always smooth with respect to small perturbations of S or of S' , but Hausdorff disparities are nevertheless useful in numerical schemes as will be explained below.

In applications, for each snapshot $S_j = S_{t_j}$ one can often identify a well defined set of p *reference points*

$$Z_j = \{z_{j,1}, \dots, z_{j,p}\}$$

such that the Z_j are point-to-point matched by the unknown diffeomorphisms F_t . To force the diffeomorphic matching of these reference point sets, we will use the obvious pointwise disparity functions

$$\text{pointdisp}(V) = \sum_{p=1}^r |F_{t_j}(z_{0,p}) - z_{j,p}|^2.$$

Other efficient disparity functions $D(S, S')$ based on Hilbertian distances between differentiable currents carried by S and S' have been introduced in [85].

7.3.4.2 Intensity Matching Error

We consider at instant $t_0 = T_0$ the shape S_0 as a set of discrete points $\{x_1, \dots, x_n\}$, the intensity matching error at instance T_l ($l \geq 1$) is defined by the mean of squared difference

$$\text{Disp}_l^I = \frac{1}{n} \sum_{i=1}^n (INT_{T_l}(F_{T_l}(x_i)) - INT_{T_{l-1}}(F_{T_{l-1}}(x_i)))^2 \quad (7.6)$$

where $INT_{T_l}(\cdot)$ is the image intensity function at time T_l . In the medical imaging data, the set of instances t_j , $0 \leq j \leq q$ at which we have the geometric information based on the doctors' tagging work is a subset of instances $\{T_0, \dots, T_Q\}$, $Q \geq q$. And the image intensities are provided only at discrete voxels, which means the intensity function is discrete. However, the computation of intensity matching error requires a continuous version of intensity function. Therefore, we interpolate the discrete intensity at discrete voxels to generate a continuous intensity function. There are a lot of popular image intensity interpolation methods with different precision, and for our purpose of simplicity, we select the linear interpolation which practically leads

to satisfactory results with much reduced computing time as compared to high order interpolation methods.

We have now provided all prerequisites to elaborate on the variational formulation of matching intermediary snapshots. As above, we consider a given sequence of $q + 1$ instantaneous 'snapshots' $S_j = S_{t_j} \subset R^3$ generated at fixed intermediary times $t_j, 0 \leq j \leq q$, by a deformable shape S_t with unknown dynamics. Each S_j is typically a bounded piecewise smooth submanifold with a boundary that is regularly embedded in R^3 . We seek a vector field flow $v = (v_t)$ on R^3 belonging to the Hilbert space $L^2(I, V), I := [T_0, T_Q]$ such that for $j = 1, \dots, q$, the R^3 -diffeomorphism F_t (solution of (7.3a),(7.3b)) deforms the initial snapshot S_0 into a submanifold $\hat{S}_j = F_{t_j}(S_0)$ 'coinciding' as good as possible with the given snapshot S_j . We select a particular geometric disparity functional $D(S, S')$ and quantify the constraint matching adequacy of v by the q numerical disparities

$$\text{Disp}_j(v) = 1_{j \in \{t_1, \dots, t_q\}} D(\hat{S}_j, S_j) + \mu \text{Disp}_j^I(v) \quad , \quad j = 1, \dots, Q.$$

We fix Q positive numerical weights $\lambda_j > 0$, and we define the *disparity cost functional* by

$$\text{Disp}(v) := \sum_{j=1}^Q \lambda_j \text{Disp}_j(v). \quad (7.7)$$

These disparity functionals actually belong to a much wider class of functionals $\mathcal{D}(V)$ which we will introduce next:

For $v \in L^2(I, V)$ we denote by $F_t^v \in \text{Diff}(R^3)$ the solution of (7.3a),(7.3b) determined by v . We define $\mathcal{D}(V)$ as the space of all disparity functionals $\text{Disp} : L^2(I, V) \rightarrow R_+$

which are of the form

$$\text{Disp}(v) = \phi(F_{T_1}^v, \dots, F_{T_Q}^v) \quad (7.8)$$

for some fixed, but arbitrary choice of the integer Q , of the instants $T_1, \dots, T_Q$ in R_+ , and of the continuous function $\phi : (\text{Diff}(R^3))^Q \rightarrow R_+$.

7.4 Diffeomorphic point matching

Diffeomorphic point matching [77, 87, 90] is a particular case of diffeomorphic matching of measures that can be derived from the general framework of section 3. In such a framework, let $\hat{S}_j = F_{t_j}^v(S_0)$ be a sequence of Q submanifolds generated at instants T_Q , $1 \leq j \leq q$, from the initial snapshot S_0 by a R^3 -diffeomorphism F_t^v satisfying (7.3a),(7.3b) with unknown flow dynamics $v \in L^2(I, V)$. Let $\hat{X}_j = F_{T_j}^v(X_0) = \{F_{T_j}^v(x_1^0), \dots, F_{T_j}^v(x_{N_0}^0)\}$ be the sequence of q point sets generated by F_t^v at instants T_j , $1 \leq j \leq Q$ from the initial point set X_0 . We denote by $x_n(t) = F_t^v(x_n^0)$, $t \in I$, the corresponding N_0 trajectories emanating from x_n^0 , $1 \leq n \leq N_0$, at $t = 0$. Thus we have $\hat{X}_j = \{x_1(T_j), \dots, x_{N_0}(T_j)\}$, $1 \leq j \leq Q$.

Recall that $x_n(t) = F_t^v(x_n^0)$, $t \in I$, $1 \leq n \leq N_0$, are the solutions of the ODEs

$$\frac{dx_n(t)}{dt} = v_t(x_n(t)), \quad t \in (0, 1], \quad (7.9a)$$

$$x_n(0) = x_n^0. \quad (7.9b)$$

Thus the trajectories $x_n(t)$, $1 \leq n \leq N_0$, and the disparity cost functional $\text{Disp}(v)$, are uniquely determined by the values of v_t taken at N_0 points $x_n(t)$. Taking into

account that $V = V_K$ is a RKHS associated to a radial Gaussian kernel $K = K_{\sigma_0}$, this allows us to restrict the search for $v_t \in V$ to the set of linear combination of $K_{x_n(t)}$, $1 \leq n \leq N_0$, and thus places us in a finite dimensional situation. We look for the flow v_t under the form $v_t = \sum_{n=1}^{N_0} \alpha_n(t) K_{x_n(t)}$, $\alpha_n(t) \in R^3$, which may also be written

$$v_t(x) = \sum_{n=1}^{N_0} K_{\sigma_0}(x_n(t), x) \alpha_n(t) \quad , \quad \forall x \in R^3. \quad (7.10)$$

By the self reproducing property of K , we have

$$\|v_t\|_V^2 = \sum_{n=1}^{N_0} \sum_{n'=1}^{N_0} K_{\sigma_0}(x_n(t), x_{n'}(t)) \alpha_n^T(t) \alpha_{n'}(t).$$

We introduce the matrix-vector notations:

$$x^{(0)} = (x_1^{(0)}, \dots, x_{N_0}^{(0)})^T \in R^{N_0 d}, \quad x(t) = (x_1(t), \dots, x_{N_0}(t))^T \in R^{N_0 d}, \quad t \in I,$$

$$\alpha(t) := (\alpha_1(t), \dots, \alpha_{N_0}(t))^T \in R^{N_0 d}, \quad t \in I,$$

$$A(x(t)) = (A_{nn'}(x(t)))_{n,n'=1}^{N_0} \in R^{N_0 d \times N_0 d}, \quad A_{nn'}(x(t)) := K_{\sigma_0}(x_n(t), x_{n'}(t)) I_d \in R^{d \times d}.$$

It follows that the kinetic energy $E(v)$ defined by (7.2) takes the form

$$E(v) = \frac{1}{2} \int_0^1 \alpha(t)^T A(x(t)) \alpha(t) dt. \quad (7.11)$$

Hence, in terms of $\alpha \in L^2(I, R^{N_0 d})$, the objective functional reads

$$J(\alpha) = \frac{1}{2} \int_0^1 \alpha(t)^T A(x(t)) \alpha(t) dt + \sum_{j=1}^Q \lambda_j \text{Disp}_j(x(T_j)) \quad (7.12)$$

The diffeomorphic point matching amounts to the solution of the optimal control problem

$$\inf_{\alpha \in L^2(I, R^{N_0 d})} J(\alpha), \quad (7.13a)$$

subject to

$$\frac{dx(t)}{dt} = A(x(t)) \alpha(t), \quad t \in (0, 1], \quad (7.13b)$$

$$x(0) = x^{(0)}. \quad (7.13c)$$

The existence of a solution α^* of (7.13a)-(7.13c) follows from Theorem 3.2 in [67].

Assume that $\alpha^*(\cdot)$ is the solution of the optimal control problem (7.13), and that $x^*(\cdot)$ is the corresponding trajectory. Then there exists a function $p^*(\cdot)$, called the adjoint state, such that the triple (x^*, p^*, α^*) satisfies

$$\frac{dx^*(t)}{dt} = A(x^*(t)) \alpha^*(t), \quad t \in (0, 1], \quad (7.14a)$$

$$x^*(0) = x^{(0)}, \quad (7.14b)$$

$$-\frac{dp^*(t)}{dt} = B(x^*(t), \alpha^*(t))^T \left(p^*(t) + \frac{1}{2} \alpha^*(t) \right), \quad t \in (T_{j-1}, T_j), \quad (7.15a)$$

$$p^*(T_Q^+) = 0 \quad (7.15b)$$

$$p^*(T_Q^-) = p^*(T_Q^+) + \lambda_Q \nabla_{x^*(T_Q)} \text{Disp}_Q \quad (7.15c)$$

$$p^*(T_j^-) = p^*(T_j^+) + \lambda_j \nabla_{x^*(T_j)} \text{Disp}_j + \lambda_{j+1} \nabla_{x^*(T_j)} \text{Disp}_{j+1} \quad , \quad j = Q-1, \dots, 1, \quad (7.15d)$$

$$A(x^*(t))(\alpha^*(t) + p^*(t)) = 0, \quad t \in (0, 1]. \quad (7.16)$$

In the adjoint state equation (7.15a), the matrix

$$B(x^*(t), \alpha^*(t)) = \nabla_x (A(x^*(t)) \alpha^*(t)),$$

is given by

$$B(x^*(t), \alpha^*(t)) = B_{nm}(x^*(t), \alpha^*(t))_{n,m=1}^{N_0} \in R^{N_0 d \times N_0 d}$$

$$B_{nm}(x^*(t), \alpha^*(t)) := \alpha_m^*(t)(\nabla_2 K_{\sigma_0}(x_n^*(t), x_m^*(t)))^T + \delta_{nm} \sum_{k=1}^N \alpha_k^*(t)(\nabla_1 K_{\sigma_0}(x_n^*(t), x_k^*(t)))^T,$$

where ∇_x denotes the gradient with respect the argument $x(t)$ and $\nabla_i K_{\sigma_0}(\cdot, \cdot), 1 \leq i \leq 2$, stands for the gradient with respect to the i -th argument of $K_{\sigma_0}(\cdot, \cdot)$.

Moreover, (7.15d) represents the jump discontinuities of $p^*(\cdot)$ at times $t_j, 1 \leq j \leq Q$.

7.5 Numerical solution of diffeomorphic point matching

In this section, we will consider the basic matching algorithm based on the gradient method and time discretizations of the optimality conditions (7.14a),(7.14b) and (7.15a),(7.15d) as well as a variant of it which uses iteration-dependent weighting parameters for the matching term in the objective functional. Further, we will discuss an appropriate initialization of the algorithm.

7.5.1 The basic matching algorithm

For the time discretizations of the optimal control problem (7.13) we introduce a partition Δ_I of I according to

$$\Delta_I := \cup_{j=1}^q \Delta_{I_j}, \quad \Delta_{I_j} := \{T_{j-1} =: T^{L_{j-1}} < T^{L_{j-1}+1} < \dots < T^{L_j-1} < T^{L_j} := T_j\} \quad (7.17)$$

where Δ_{I_j} , $1 \leq j \leq Q$, are subpartitions of intervals $I_j = [T_{j-1}, T_j]$ whose endpoints T_j are key time frames at which the intensity information is given. We set $L_0 := 0$ and $L := L_Q$ and define step sizes $\Delta t^\ell := T_{\ell+1} - T_\ell > 0$, $0 \leq \ell \leq L-1$.

We introduce the discrete control space

$$\mathcal{U}^{\Delta_I} = R^{L \times (N_0 d)}, \quad (7.18)$$

equipped with the inner product

$$(\boldsymbol{\alpha}, \boldsymbol{\beta})_{\Delta_I} = \sum_{\ell=0}^{L-1} \Delta t^\ell \boldsymbol{\alpha}^\ell \cdot \boldsymbol{\beta}^\ell = \sum_{\ell=0}^{L-1} \sum_{n=1}^{N_0} \Delta t^\ell \alpha_n^\ell \cdot \beta_n^\ell,$$

and discretize the state equation (7.14a),(7.14b) and the adjoint state equation (7.15a),(7.15d) by the explicit Euler method. Introducing the notations

$$\mathbf{x} = \{\mathbf{x}^\ell\}_{\ell=0}^L, \quad \mathbf{x}^\ell = \{x_n^\ell\}_{n=1}^{N_0}, \quad x_n^\ell \approx x_n(t^\ell), \quad (7.19a)$$

$$\mathbf{p} = \{\mathbf{p}^\ell\}_{\ell=0}^L, \quad \mathbf{p}^\ell = \{p_n^\ell\}_{n=1}^{N_0}, \quad p_n^\ell \approx p_n(t^\ell), \quad (7.19b)$$

$$\boldsymbol{\alpha} = \{\boldsymbol{\alpha}^\ell\}_{\ell=0}^{L-1}, \quad \boldsymbol{\alpha}^\ell = \{\alpha_n^\ell\}_{n=1}^{N_0}, \quad \alpha_n^\ell \approx \alpha_n(t^\ell), \quad (7.19c)$$

the discretized optimality system reads

$$\frac{\mathbf{x}^{\ell+1} - \mathbf{x}^\ell}{\Delta t^\ell} = A(\mathbf{x}^\ell) \boldsymbol{\alpha}^\ell, \quad \ell = 0, \dots, L-1, \quad (7.20a)$$

$$\mathbf{x}^0 = \mathbf{x}^{(0)}, \quad (7.20b)$$

$$\frac{\mathbf{p}^{\ell-1} - \mathbf{p}^\ell}{\Delta t^\ell} = B(\mathbf{x}^\ell, \boldsymbol{\alpha}^\ell)^T \left(\mathbf{p}^\ell + \frac{1}{2} \boldsymbol{\alpha}^\ell \right), \quad \ell = L_j, \dots, L_{j-1} + 1, \quad (7.21a)$$

$$\mathbf{p}^{L_Q+} = 0, \quad (7.21b)$$

$$\mathbf{p}^{L_Q-} = \mathbf{p}^{L_Q+} + \lambda_Q \nabla_{\mathbf{x}^{L_Q}} \text{Disp}_Q \quad (7.21c)$$

$$\mathbf{p}^{L_j-} = \mathbf{p}^{L_j+} + \lambda_j \nabla_{\mathbf{x}^{L_j}} \text{Disp}_j + \lambda_{j+1} \nabla_{x^*(L_j)} \quad (7.21d)$$

$$A(\mathbf{x}^\ell) (\boldsymbol{\alpha}^\ell + \mathbf{p}^\ell) = \mathbf{0} \quad , \quad 0 \leq \ell \leq L-1. \quad (7.22)$$

It turns out that (7.20)-(7.22) represent the optimality conditions for a discrete minimization problem. In fact, introducing $J^{\Delta_I}(\boldsymbol{\alpha})$ as the discrete objective functional

$$J^{\Delta_I}(\boldsymbol{\alpha}) := \sum_{\ell=0}^{L-1} \frac{\Delta t^\ell}{2} (\boldsymbol{\alpha}^\ell)^T A(\mathbf{x}^\ell) \boldsymbol{\alpha}^\ell + \sum_{j=1}^q \lambda_j \text{Disp}_j(\mathbf{x}^{L_j}), \quad (7.23)$$

we have the following result.

The equations (7.20)-(7.22) are the first order necessary optimality conditions for the finite dimensional minimization problem

$$\min_{\boldsymbol{\alpha} \in \mathcal{U}^{\Delta_I}} J^{\Delta_I}(\boldsymbol{\alpha}) \quad (7.24)$$

subject to the discrete state equations (7.20a),(7.20b).

The proof is the discrete analogue of the proof of Theorem 3.2 in [67] and will thus be omitted.

Let $(\mathbf{x}^*, \mathbf{p}^*, \boldsymbol{\alpha}^*)$ with $\mathbf{x}^* = \{\mathbf{x}_*^\ell\}_{\ell=0}^L$, $\mathbf{p}^* = \{\mathbf{p}_*^\ell\}_{\ell=0}^L$, $\boldsymbol{\alpha}^* = \{\boldsymbol{\alpha}_*^\ell\}_{\ell=0}^L$ satisfy the discrete optimality system (7.20)-(7.22). Then, it holds

$$\mathbf{0} = \nabla J^{\Delta_I}(\boldsymbol{\alpha}^*) \quad (7.25)$$

where

$$\nabla J^{\Delta I}(\boldsymbol{\alpha}^*) = \{\mathbf{g}^\ell\}_{\ell=0}^{L-1}, \quad \mathbf{g}^\ell = A(\mathbf{x}^\ell) (\boldsymbol{\alpha}_*^\ell + \mathbf{p}_*^\ell). \quad (7.26)$$

We observe that

$$\delta J^{\Delta I}(\boldsymbol{\alpha}) = (\nabla J^{\Delta I}(\boldsymbol{\alpha}), \delta \boldsymbol{\alpha})_{\Delta I}. \quad (7.27)$$

From (7.23) we deduce

$$\begin{aligned} \delta J^{\Delta I}(\boldsymbol{\alpha}) &= \sum_{\ell=0}^{L-1} \Delta t^\ell \left((\boldsymbol{\alpha}^\ell)^T A(\mathbf{x}^\ell) \delta \boldsymbol{\alpha}^\ell + \frac{1}{2} (\boldsymbol{\alpha}^\ell)^T B(\mathbf{x}^\ell, \boldsymbol{\alpha}^\ell) \delta \mathbf{x}^\ell \right) \\ &\quad + \sum_{j=1}^q \lambda_j \nabla \text{Disp}_j(\mathbf{x}^{L_j}) \delta \mathbf{x}^{L_j} \end{aligned} \quad (7.28a)$$

where

$$\frac{\delta \mathbf{x}^{\ell+1} - \delta \mathbf{x}^\ell}{\Delta t^\ell} = A(\mathbf{x}^\ell) \delta \boldsymbol{\alpha}^\ell + B(\mathbf{x}^\ell, \boldsymbol{\alpha}^\ell) \delta \mathbf{x}^\ell, \quad \ell = 0, \dots, L-1, \quad (7.28b)$$

$$\delta \mathbf{x}^0 = \mathbf{0}. \quad (7.28c)$$

Multiplying both sides of (7.28b) by $\mathbf{p}^\ell$, partial summation yields

$$\begin{aligned} 0 &= \sum_{\ell=0}^{L-1} \Delta t^\ell \mathbf{p}^\ell \cdot \left(\frac{\delta \mathbf{x}^{\ell+1} - \delta \mathbf{x}^\ell}{\Delta t^\ell} - A(\mathbf{x}^\ell) \delta \boldsymbol{\alpha}^\ell - B(\mathbf{x}^\ell, \boldsymbol{\alpha}^\ell) \delta \mathbf{x}^\ell \right) \\ &= \sum_{\ell=1}^{L-1} \Delta t^\ell \frac{\mathbf{p}^{\ell-1} - \delta \mathbf{p}^\ell}{\Delta t^\ell} \cdot \delta \mathbf{x}^\ell + \mathbf{p}^{L-1} \cdot \delta \mathbf{x}^L - \mathbf{p}^0 \cdot \delta \mathbf{x}^0 - \sum_{\ell=0}^{L-1} \Delta t^\ell \mathbf{p}^\ell \cdot A(\mathbf{x}^\ell) \delta \boldsymbol{\alpha}^\ell \\ &\quad - \sum_{\ell=0}^{L-1} \Delta t^\ell \mathbf{p}^\ell \cdot B(\mathbf{x}^\ell, \boldsymbol{\alpha}^\ell) \delta \mathbf{x}^\ell \end{aligned} \quad (7.29)$$

If we take (7.21a),(7.21d) into account, it follows from (7.29), (7.28a) that

$$\delta J^{\Delta I}(\boldsymbol{\alpha}) = \sum_{\ell=0}^{L-1} \Delta t^\ell A(\mathbf{x}^\ell) (\boldsymbol{\alpha}^\ell + \mathbf{p}^\ell) \cdot \delta \boldsymbol{\alpha}^\ell. \quad (7.30)$$

Since $\delta\boldsymbol{\alpha} = \{\delta\boldsymbol{\alpha}^\ell\}_{\ell=0}^{L-1}$ is arbitrary, (7.27) results in

$$\nabla J^{\Delta_I}(\boldsymbol{\alpha}) = \{\mathbf{g}^\ell\}_{\ell=0}^{L-1}, \quad \mathbf{g}^\ell = A(\mathbf{x}^\ell) (\boldsymbol{\alpha}^\ell + \mathbf{p}^\ell). \quad (7.31)$$

In view of (7.25), (7.26), the discrete minimization problem (7.24) can be solved by a gradient based algorithm operating in $\mathcal{U}^{\Delta_I}$.

7.5.2 Smoothing of the Hausdorff matching term

In general, the Hausdorff distance is not a smooth function. For the gradient method with Armijo line search we use a smoothed version under the assumption that for any point on the deformed manifolds $\hat{S} = F_{t_j}^v(S_0)$, $1 \leq j \leq q$, there exists some neighborhood of points which corresponds to a unique set of neighboring points on the target S_q . We define

$$\tilde{D}_h(S_j, \hat{S}_j) := h_1(X_j, \hat{X}_j) + h_2(\hat{X}_j, X_j), \quad (7.32)$$

$$h_1(X_j, \hat{X}_j) := (N_0 N_j)^{-1} \sum_{x \in X_j} \sum_{k=1}^{N_0} |x - \Phi^{(k)}(x)|^2, \quad (7.33)$$

$$h_2(\hat{X}_j, X_j) := (N_0 N_j)^{-1} \sum_{\hat{x} \in \hat{X}_j} \sum_{k=1}^{N_j} |\hat{x} - \Psi^{(k)}(\hat{x})|^2. \quad (7.34)$$

Here, setting $\Phi^{(0)}(x) = x$ and $\Psi^{(0)}(\hat{x}) = \hat{x}$, the mapping $\Phi^{(k)}(x) : R^3 \rightarrow X_j$, $1 \leq k \leq N_0$, and $\Phi^{(k)}(x) : R^3 \rightarrow X_j$, $1 \leq k \leq N_j$, are given recursively according to

$$\Phi^{(k)}(x) := \min\{|\hat{x} - \Phi^{(0)}(x)| \mid \hat{x} \in \hat{X}_j \setminus \{\Phi^{(k-1)}(x), \dots, \Phi^{(0)}(x)\}\}, 1 \leq k \leq N_0 \quad (7.35)$$

$$\Psi^{(k)}(\hat{x}) := \min\{|x - \Psi^{(0)}(\hat{x})| \mid x \in X_j \setminus \{\Psi^{(k-1)}(\hat{x}), \dots, \Psi^{(0)}(\hat{x})\}\}, 1 \leq k \leq N_j. \quad (7.36)$$

This type of Hausdorff matching will be referred to as total Hausdorff matching.

7.5.3 Choice of scale parameters

The scale parameter $\sigma_0 > 0$ of the radial Gaussian kernel defined the RKHS $V = V_k$ has an impact on the smoothness of the deformations in so far as larger values result in smoother deformations, and vice versa. It can be chosen by means of local radii $r_i, 1 \leq i \leq N_0$, associated with each point x_i on the reference configuration S_0 . In particular, we define

$$r_i := \max_{1 \leq l \leq m} r_{i,l}, \quad r_{i,l} := |x_i - x_i^{(l)}|, 1 \leq l \leq m, \quad (7.37)$$

where the points $x_i^{(l)} \in S_0$ are given recursively according to

$$x_i^{(l)} := \min\{|x - x_i^{(0)}| \mid x \in S_0 \setminus \{x_i^{(l-1)}, \dots, x_i^{(0)}\}\}, x_i^{(0)} := x_i. \quad (7.38)$$

Here, $m \in \mathbb{N}$ is chosen as $m = 2$ for points on a curve. Based on these local radii, we set

$$\sigma_0 := \kappa 2^{-1/2} R, \quad R := \max_{1 \leq i \leq N_0} r_i \quad (7.39)$$

where κ is some constant satisfying $\kappa > 2^{-1/2}$ (e.g., $2^{-1/2} = 1.1$).

The scale parameter can be kept fixed during the whole computation or can be updated by taking into account the positions $\hat{x}_i$ at different times instances. This is realized by defining local radii at any instance as above. In our applications, numerical evidence suggests to keep σ_0 fixed during the whole process, since there is not much difference between the resolution of the points on the intermediaries and the target.

7.5.4 Continuation in the regularization parameter

The weights λ_j in the disparity cost functional are chosen according to $\lambda_j = \lambda, 1 \leq j \leq Q$, for some positive λ . This regularization parameter provides a weighting between the kinetic energy and the disparity cost functional. If λ is small, the regularizing effect of the kinetic energy dominates, whereas large values of λ enforce the matching quality. Hence, for the sake of a good matching quality, one is interested in performing the computation by choosing λ sufficiently large. However, for increasing λ , the optimality system becomes more and more ill-conditioned which may result in the divergence of the gradient method. A convenient remedy to overcome this deficiency is to use an appropriate continuation in λ . The continuation method consists of outer iterations in λ and inner iterations which are chosen as the gradient method with Armijo line search. Progress during the outer iterations is measured in terms of the performance of the intermediary curves and the target curve in terms of geometric matching error (GME_j) and intensity matching error (IME_j) at instances $1 \leq j \leq Q$, which are defined as:

$$GME_j := 90 \text{ percentile of } \{\min_{x_l \in S_j} |x - x_l|, x \in \hat{X}_j\} \quad (7.40)$$

$$IME_j := 100 \frac{\sum_{i=1}^n |INT_{T_j}(F_{T_j}(x_i)) - INT_{T_{j-1}}(F_{T_{j-1}}(x_i))|}{\sum_{i=1}^n INT_{T_{j-1}}(F_{T_{j-1}}(x_i))} \quad (7.41)$$

and we defined the mean errors of them by

$$mean(GME) := \frac{1}{q} \sum_{j=1}^q GME_j \quad (7.42)$$

$$mean(IME) := \frac{1}{Q} \sum_{j=1}^Q IME_j \quad (7.43)$$

7.6 Numerical Results for the Human Mitral Valve Annulus

In the cardiology application, we apply variational techniques to numerically construct diffeomorphic flows modeling the dynamic deformations between multiple 3D-snapshots of the human mitral valve annulus. Our starting point is a patient's specific finite set of $q + 1$ static models S_j of the mitral valve annulus and $Q + 1$ instances with image intensities. These models were generated by image analysis of live 3D-echocardiographic movies at specific heart cycle instants t_j in a recent study involving the Methodist Hospital, Houston, Texas (S. BenZekry, S. Little, W. Zoghbi, MDs) as well as a research group on mathematical image analysis led by R. Azencott. Each 3D-movie includes roughly twenty 3D-frames per heartbeat cycle, acquired by ultrasound technology, and represents a high volume of image data corrupted by 'speckle' noise.

The mitral valve models S_j are based on NURBS(non uniform rational B-splines), and were obtained by combining optical flow extraction algorithms with sparse tagging by medical experts. The number $q + 1$ of intermediary key heartbeat cycle instants t_j ranged from 3 to 8, and time range $I = [T_0, T_1]$ covers either a half or a whole heartbeat cycle with total duration between 1/2 to 1 second.

7.6.1 Diffeomorphic Matching for 5 Annulus Snapshots

We consider the results of diffeomorphic point matching of the annulus in case of the Hausdorff distances (Hausdorff matching) being used for the geometric disparity functional (matching term) and without intensity matching term. We have chosen five snapshots $S_1, S_3, S_5, S_7, S_{10}$ corresponding to time instances $t_1, t_3, t_5, t_7, t_{10}$ of the heart beating cycle. The frame rate is 27Hz. The spline model for the modeling of the annulus allows us to discretize the annulus such that the arc length distance between any two neighboring points is the same. We take 62 points on the reference curve S_1 and final target curves S_3, S_5, S_7, S_{10} . We used trivial initialization $\alpha = 0$, time steps $\Delta t^l = 1/18$. The parameter κ has been chosen as $\kappa = 1.1$ resulting in a value of $\sigma_0 = 3.53$ for the Gaussian kernel K_{σ_0} of the RKHS. The performance of the diffeomorphic matching is measured on one hand by the convergence history of the continuation method in terms of the decrease of quantities of interest ($mean(GME)$ and $mean(IME)$) and on the other hand by the tradeoffs between these two performance measures.

We have chosen $\lambda^{(0)} = 1$ as a start iterate for the regularization parameter in the continuation method and increased it by the factor $\rho = 1.1$. Firstly, we run an experiments by setting $\mu = 0$. The computational performance of the continuation algorithm is evaluated first by the convergence history for the q indicators of the geometric matching accuracy GME_j and for the corresponding values $Disp_j^G$ of the q components of the disparity functional.

We also record and display the tradeoffs between these matching quality indicators and the kinetic energy of the corresponding deformation flows. At the end of each inner iteration of gradient descent with fixed regularization parameter λ , we generate a point on each one of the approximate *Pareto frontiers* displaying matching quality indicators as functions of the kinetic energy, we expect and empirically observe convexity of these approximate Pareto frontiers, viewed as usual as the location of the weak Pareto optima for the pair of competing criteria (kinetic energy versus matching disparity).

The convergence history is documented in Figure 7.2 which displays the decrease of the geometric accuracy indicators GME_j and of the disparity components $Disp_j^G$ for our 4 snapshots. It takes 200 iterations to reach the threshold geometric accuracy required for this application, but we have extended iterations beyond this value to obtain a more complete view of the Pareto frontiers.

The Pareto frontiers for the geometric accuracy indicators GME_j , the Hausdorff disparities $Disp_j^G$, and the global Hausdorff disparity $\sum Disp_j^G$ are shown in Figure 7.3. Note that for the first annulus snapshot the geometric accuracy GME_1 and the Hausdorff disparity $Disp_1^G$ do not exhibit a convex decrease pattern at the beginning of the continuation procedure. Indeed, the currently implemented algorithmic optimization strategy first “focuses” on matching the final snapshot and then successively shifts the focus on the matching of the intermediary snapshots in decreasing order. This is due to the fact that in the backwards adjoint ODE, the snapshots matching errors kick in successively in the same reverse order. Hence, as long as the final snapshots errors are large, the corrections implemented for the first snapshots

remain quite fuzzy, and only turn out to be efficient once the matching errors on the final snapshots have become small enough.

When $\mu = 0$, i.e., intensity information is not applied in the optimization, the final $mean(GME)$ and $mean(IME)$ are 0.4 mm and 13.1 percents respectively, after 400 iterations. In figure 7.1, we plotted the five annulus curves and also trajectories of certain points on the reference annulus. When $\mu = 0.01$, after same number of iterations, the final $mean(GME)$ and $mean(IME)$ are 1.0 mm and 3.5 percents.




Figure 7.1: The five closed curves are five instances of the annulus. The small dotted curve is the reference curve. From bottom to top, the next three curves are the intermediaries and the last one is the final target. The vertical ‘-’ lines are deformations of selected points on the reference curve.




Figure 7.2: Diffeomorphic matching for multiple mitral annulus snapshots, using smoothed Hausdorff distances: Convergence history for the geometric accuracy indicators and for the smoothed Hausdorff disparity components $Disp_j^G$

7.6.2 Diffeomorphic Matching for All Annulus Snapshots in One Cardiac Cycle

We consider here the 9 annulus snapshots acquired at times 1,3,5,7,10,14,18,22,26, which cover one whole cardiac cycle, we generate 8 data sets $D_1, \dots, D_8$ as following:

- D_1 – Time1-3
- D_2 – Time3-5
- D_3 – Time5-7
- D_4 – Time7-10
- D_5 – Time10-14

- D_6 – Time14-18
- D_7 – Time18-22
- D_8 – Time22-26

For each data set D_i , we choose 62 points on each annulus that are located with equal arc length. We use the smoothed Hausdorff disparity, and initialize α by $\alpha = 0$, the parameters θ, σ, ρ are the same as above. We choose $\lambda^{(0)} = 0.2$ as a start iterate for the regularization parameter in the continuation method and increased it by the factor $\rho = 1.1$. We select different μ values for each data set so that the best tradeoff between GME and $mean(IME)$ can be achieved.

For each optimization, we stopped the optimization when GME becomes less than 1 for the first time, and we record the final $mean(IME)$ for each μ and each data set D_i . Numbers of iterations for all experiments are less than 200. The numerical result was shown in table 7.1. It clearly shows that the intensity matching error decreases as the intensity weight μ increases.

For each data set, we compare the intensity curves of the reference annulus and deformed annulus, for $\mu = 0$ and best μ in table 7.1. We plot these curves in figures 7.4, 7.5, 7.6, 7.7, 7.8, 7.9, 7.10, 7.11.

7.6.3 Diffeomorphic Matching of the Anterior Leaflet

We present the performances of diffeomorphic matching for 4 snapshots S_0, S_1, S_2, S_3 of the anterior leaflet, acquired at times 0, 1, 5, 10, using smoothed Hausdorff

	D_1	D_2	D_3	D_4	D_5	D_6	D_7	D_8
$\mu = 0$	9.8	7.3	8.7	7.3	15.5	10.8	4.9	2.2
$\mu = 0.001$	8.2	5.1	6.4	5.2	7.9	8.1	4.3	2.1
$\mu = 0.002$	7.2	4.3	5.3	4.3	6.2	6.5	3.8	2.0
$\mu = 0.003$	-	3.8	4.6	3.8	5.1	5.7	3.5	1.9
$\mu = 0.004$		-	4.1	3.5	-	-	3.3	-
$\mu = 0.005$			3.8	3.1			3.1	

Table 7.1: $mean(IME)$ when $GME \approx 1$.

snapshot disparities without intensity matching term. Initialization for the time dependent vector field is implemented by piecewise constant vector fields for boundary deformations and adequately fitted polynomials in time and space variables for the interior of S_0 . The continuation algorithm starts with $\lambda^{(0)} = 0.1$, and is multiplied by 1.5 at each outer iteration. The scale and termination parameters σ and θ are as above. The following figures 7.12 and 7.13 display satisfactory performance results.

7.6.4 Diffeomorphic Matching for Multiple Snapshots of the Posterior Leaflet

We have similarly implemented the diffeomorphic matching for 4 snapshots S_0, S_1, S_2, S_3 of the posterior leaflet acquired at times 0, 1, 5, 10, using smoothed Hausdorff snapshot disparities without intensity matching term. Time is discretized into 30 equal intervals. The 4 snapshots are discretized by point meshes of approximate cardinals 250 for the initial leaflet S_0 , and 1100 for the other snapshots. The initialization and the choice of the algorithmic parameters are similar to the implementation just described for the anterior leaflet. As already noted above, when one uses the

same value λ for all 3 regularization weights λ_j , the matching quality between S_j and $\hat{S}_j$ improves more slowly for $j = 1$ than for $j = 3$ (see figure 7.14). We have compared this approach to a more adaptive one, where one dynamically adjusts the weights λ_j at each outer iteration by appropriately balancing current values of the 3 Hausdorff disparities. (figure 7.15). The geometric matching accuracies reach a desirable pragmatic threshold slightly faster for the dynamically independent weights, at the cost of some loss of convexity for the Pareto frontiers, which suggests that dynamic adjustment of weights may provide less robustness in the continuation procedure.




Figure 7.3: Diffeomorphic matching for multiple annulus snapshots, using smoothed Hausdorff distances: Pareto frontiers for the geometric accuracy indicators, for the Hausdorff disparities, and for the global Hausdorff disparity




Figure 7.4: Reference(solid) and deformed (dash) annulus intensity curve by arc length: top($\mu = 0$), bottom(best μ), data set D_1 .




Figure 7.5: Reference(solid) and deformed (dash) annulus intensity curve by arc length: top($\mu = 0$), bottom(best μ), data set D_2 .




Figure 7.6: Reference(solid) and deformed (dash) annulus intensity curve by arc length: top($\mu = 0$), bottom(best μ), data set D_3 .




Figure 7.7: Reference(solid) and deformed (dash) annulus intensity curve by arc length: top($\mu = 0$), bottom(best μ), data set D_4 .




Figure 7.8: Reference(solid) and deformed (dash) annulus intensity curve by arc length: top($\mu = 0$), bottom(best μ), data set D_5 .




Figure 7.9: Reference(solid) and deformed (dash) annulus intensity curve by arc length: top($\mu = 0$), bottom(best μ), data set D_6 .





Figure 7.10: Reference(solid) and deformed (dash) annulus intensity curve by arc length: top($\mu = 0$), bottom(best μ), data set D_7 .





Figure 7.11: Reference(solid) and deformed (dash) annulus intensity curve by arc length: top($\mu = 0$), bottom(best μ), data set D_8 .




Figure 7.12: Diffeomorphic matching for four anterior leaflet snapshots: Pareto frontiers for the separate Hausdorff disparities to snapshots (left) and for the global Hausdorff disparity (right)




Figure 7.13: Pareto frontiers for the maximum distances to snapshots (left) and for their 90th percentiles (right)




Figure 7.14: Diffeomorphic matching for four posterior leaflet snapshots: Geometric accuracy indicators and Pareto frontiers for strictly equal regularization weights λ_j




Figure 7.15: Diffeomorphic matching for four posterior leaflet snapshots: Geometric accuracy indicators and Pareto frontiers for dynamically adjusted regularization weights λ_j

Chapter 8

Conclusion

2D-IntraCardiac Echography (ICE) is a primary echocardiography modality to acquire visual information on live cardiac anatomy, by catheter insertion of a transducer within the heart chamber. However, ICE-images sequences are always perturbed by a strong “speckle” noise, due to the diffuse scattering of ultrasound pulses by small scatterers randomly dispersed at sound wavelength scales. Most published cardiac motion detection methods still do not attempt to fully use the statistical features of speckle noise. We have developed an “elastic registration” approach to compute cardiac motion by analysis of ICE-image sequences. The main principles applied here were the modeling of the unknown spatial deformations between two successive views by bicubic B-splines, and the minimization of a cost function derived by a maximum likelihood technique applied to the log-speckle noise. Instead of uniformly modeling the speckle noise by a standard Rayleigh distribution, we build, for each specific ICE image sequence, a cost functional parameterized by an image based local estimation

of the observed speckle noise local statistics.

Since fast computing of myocardiac deformations is of high interest for potential clinical use, we have improved computing efficiency by introducing local image masks to restrict the cost functional to cardiac muscle areas, which represents a small fraction of the whole ICE images. Another key ingredient for fast computing has been to implement an optimized coarse to fine multiresolution approach for gradient descent. At coarser resolutions, masks have small sizes therefore computing time is strongly reduced since at coarser scales we also use faster stopping criteria for gradient descent. Our numerical results matched quite well the experimental sonomicrometry strain data recorded on live animal subjects by physicians.

We have also extended our 2D elastic registration method to the 3D case. Specifically, we apply speckle tracking techniques to numerically track the dynamic deformations between multiple 3D-snapshots of the human mitral valve annulus in the mitral valve apparatus. Our computing results matched quite well the patient specific geometric models of the mitral valve annulus generated by NURBS based on 3D-image tagging by cardiologists.

Although we have introduced multiresolution numerical implementations and local masking techniques to greatly accelerate computing speed, our accelerated registration algorithms are still not rapid enough for fast clinical applications to 3D-echocardiography. One basic obstacle is of course the fairly restricted power available on our current computing server at University of Houston. Another obvious constraint on computing speed is that our algorithms were implemented in Matlab language, which is an interpreted programming language, which is much slower than compiled languages such as C, C++, Fortran, etc.

One mathematical limitation of the elastic registration method is that there are no explicit constraints ensuring that the computed spatial deformations are invertible. Therefore, in the spirit of recent papers on diffeomorphic 3D-image registration ([26][27][28][29]), we have studied registration of 3D-images and of soft shapes in R^3 by R^3 -diffeomorphisms. Our algorithmic method to construct such 3D-diffeomorphisms is developed at length in the joint paper [67], and is based on a variational approach where the unknown diffeomorphisms are generated by integration of smooth vector fields belonging to specific self-reproducing Hilbert spaces. In an abstract setting, the optimal diffeomorphic matching of two submanifolds (surfaces or curves) of R^3 is formulated as a minimization problem involving the actions of diffeomorphisms on the two Borel measures M_1 and M_2 induced by R^3 on the reference and target submanifolds. The objective functional to be minimized is the sum of two terms: the kinetic energy of the deformation and the matching quality of the corresponding shape registration. To make the problem computationally

accessible, we use self-reproducing kernel Hilbert spaces generated by Gaussian radial kernels, and we discretize the measures M_1 and M_2 by finite weighted sums of Dirac measures. The minimization algorithm is reduced to the solution of a finite dimensional minimization problem in very high dimension. The performance of our diffeomorphic tracking in medical 3D-image movies is illustrated by numerical results for the dynamic modeling of the human mitral valve annulus by computer analysis of 3D-echocardiographic image sequences.

Bibliography

- [1] E. H. W. Meijering, W. J. Niessen, and M. A. Viergever, "Quantitative evaluation of convolution-based methods for medical image interpolation," *Medical Image Analysis*, vol. 5, no. 2, pp. 111-126, 2001.
- [2] P. Thevenaz, T. Blu, and M. Unser, "Interpolation revisited," *IEEE Transactions on Medical Imaging*, vol. 19, no. 7, pp. 739-758, July 2000.
- [3] A. Spinei, D. Pellerin, and J. Herault, "Spatio temporal energy-based method for velocity estimation," *Signal Processing*, vol. 65, pp. 347-362, 1998.
- [4] C. Bernard, "Fast optical flow computation with discrete wavelets," *Tech. Rep., Ecole Polytechnique*, Paris, 1997.
- [5] Y. Wu, T. Kanade, J. Cohn, and C. Li, "Image registration using wavelet-based motion model," *Proceedings of Image Registration Workshop, Jacqueline Le Moigne, Ed.*, 1997.
- [6] B. Horn and B. Schunck, "Determining optical flow," *Artificial Intelligence*, vol. 17, pp. 185-203, 1981.
- [7] Gary Christensen, "Deformable Shape Models for Anatomy," *Ph.D. thesis, Washington University, Saint Louis, Mississippi* 1994.
- [8] Gary Christensen, S. Joshi, and M. Miller, "Volumetric transformation of brain anatomy," *IEEE Transactions on Medical Imaging* vol. 16, Dec. 1997.
- [9] C. de Boor, "A practical guide to splines," *Springer-Verlag*, New York, 1978.
- [10] J. Kybic, P. Thevenaz, and M. Unser, "Multiresolution spline warping for EPI registration," *Proceedings of SPIE, Denver, Colorado*, July 1999, vol. 3813, pp. 571-579, SPIE.

- [11] R. Bajcsy and S. Kovacic, "Multiresolution elastic matching," *Computer Vision, Graphics, and Image Processing*, vol. 46, pp. 1-21, 1989.
- [12] L. Brown, "A survey of image registration techniques," *ACM Computing Surveys*, vol. 24, pp: 325-376, 1992.
- [13] P. Van den Elsen, E. Pol, and M. Viergever, "Medical image matching - a review with classification," *IEEE Engineering in Medicine and Biology*, vol. 12, pp: 26-39, 1993.
- [14] C. Maurer Jr. and J. Fitzpatrick, "Interactive Image-Guided Neurosurgery," *American Association of Neurological Surgeons*, Park Ridge, IL., 1993, Ch. A Review of Medical Image Registration.
- [15] J. West, J. Fitzpatrick, M. Wang, B. Dawant, C. Maurer, R. Kessler, and R. Maciunas, "Comparison and evaluation of retrospective intermodality image registration techniques," *Proceedings of the SPIE - The International Society for Optical Engineering, Newport Beach, CA*, pp. 332-347, 1996.
- [16] J. Maintz and M. Viergever, "A survey of medical image registration," *Medical Image Analysis*, vol. 2, pp: 1-36, 1998.
- [17] H. Lester and S. Arridge, "A survey of hierarchical non-linear medical image registration," *Pattern Recognition*, vol. 32, pp: 129-149, 1999.
- [18] D. Rueckert and P. Aljabar, "Nonrigid Registration of Medical Images: Theory, Methods, and Applications," *IEEE Signal Processing Magazine*, vol. 27, pp: 113-119, 2010.
- [19] M. Xu and P.K. Varshney, "A Subspace Method for Fourier-Based Image Registration," *IEEE Geoscience and Remote Sensing Letters*, vol.6, pp: 491-494, 2009.
- [20] S. Liao and A. Chung, "Feature Based Nonrigid Brain MR Image Registration With Symmetric Alpha Stable Filters," *IEEE Transactions on Medical Imaging*, vol. 29, pp: 106-119, 2010.
- [21] M. Hub, M. Kessler, and C. Karger, "A Stochastic Approach to Estimate the Uncertainty Involved in B-Spline Image Registration," *IEEE Transactions on Medical Imaging*, vol. 28, pp: 1708-1716, 2009.
- [22] D. Loeckx, P. Slagmolen, F. Maes, D. Vandermeulen, and P. Suetens, "Nonrigid Image Registration Using Conditional Mutual Information," *IEEE Transactions on Medical Imaging*, vol.29, pp: 19-29, 2010.

- [23] 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.
- [24] I. Yetik and A. Nehorai, "Performance bounds on image registration," *IEEE Transactions on Medical Imaging*, vol.54, pp: 1737-1749 , 2006.
- [25] U. Grenander and M.I. Miller, "Computational anatomy: An emerging discipline," *Quarterly of Applied Mathematics*, pp. 617-C694, 1998.
- [26] J.C. Csernansky, S. Joshi, L. Wang, M. Gado, J. Philip Miller, U. Grenander, and M.I. Miller, "Hippocampal morphometry in schizophrenia by high dimensional brain mapping," *Proceedings of the National Academy of Science*, Vol. 95, pp. 11406-C11411, 1998.
- [27] J.C. Csernansky, L. Wang, S. Joshi, J. Philip Miller, M. Gado, D. Kido, D. McKeel, J.C. Morris, and M.I. Miller, "Early dat is distinguished from aging by high dimensional mapping of the hippocampus," *Neurology*, Vol. 55, Dec., pp. 1636-C1643, 2000.
- [28] C. Davatzikos, M. Vaillant, S. Resnick, J.L. Prince, S. Letovsky, and R.N. Bryan, "A computerized approach for morphological analysis of the corpus callosum," *J. of Comp. Assisted Tomography*, Vol. 20, pp. 88-C97, 1996.
- [29] D.W. Thompson, "On Growth and Form," *Cambridge University Press*, 1917.
- [30] R. Bonow, B. Carabello, K. Chatterjee, A. de Leon, D. Faxon, M. Freed, W. Gaasch, B. Lytle, R. Nishimura, P. OGara, R. ORourke, C. Otto, P. Shah, and J. Shanewise, "ACC/AHA 2006 Guidelines for the Management of Patients With Valvular Heart Disease," *Circulation*, vol. 114, pp: 84-231, 2006.
- [31] A. Vahanian, H. Baumgartner, J. Bax, E. Butchart, R. Dion, G. Filippatos, F. Flachskampf, R. Hall, B. Iung, J. Kasprzak, P. Nataf, P. Tornos, L. Torracca, and A. Wenink, "Guidelines on the management of valvular heart disease: The task force on the management of valvular heart disease of the european society of cardiology," *Eur Heart J*, 28(2), 230-268 (2007).
- [32] I.J. Schoenberg, "Contribution to the problem of approximation of equidistant data by analytic functions," *Quart. Appl. Math.*, vol. 4, pp. 45-99, 112-141, 1946.
- [33] I.J. Schoenberg, "Cardinal Spline Interpolation," *PA: Society of Industrial and Applied mathematics*, 1973.

- [34] M. Unser, "Splines: A Perfect Fit for Signal/Image Processing," *IEEE Signal Processing Magazine*, vol. 16, pp. 22-38, 1999.
- [35] M. Unser, A. Aldroubi, and M. Eden, "Fast B-spline transforms for continuous image representation and interpolation," *IEEE Transactions on Pattern Anal. Mach. Intell.*, vol. 13, 1991.
- [36] T. Blu and M. Unser, "Quantitative Fourier analysis of approximation techniques: Part I-interpolators and projectors," *IEEE Transactions on Signal Processing*, 1999.
- [37] D. J. Dowsett, P. A. Kenny, and R. E. Johnston, "The physics of diagnostic imaging," *Chapman and Hall Medical*, 1998.
- [38] V. Dutt and J. Greenleaf, "Ultrasound echo envelope analysis using homodyned k-distribution signal mode," *Ultrasonic Imaging*, vol. 16, pp.265-287, 1994.
- [39] C. Bruce and P. Friedman, "Intracardiac echocardiography," *Eur. J. Echocardiogr.*, vol. 2, pp.234-244, 2001.
- [40] N. Nanda and A. Miller, "Principles of intracardiac echocardiography," *J. Interv. Card. Electrophysiol.*, vol. 13, pp. 7-10.
- [41] V. Nkomo, J. Gardin, T. Skelton, J. Gottdiener, C. Scott, and M. Enriquez-Sarano, "PBurden of valvular heart diseases: a population-based study," *Lancet*, 2006.
- [42] R. Anderson and M. Kanani, "Mitral valve repair: critical analysis of the anatomy discussed," *Multi-media Manual of Cardiothoracic Surgery*, (2006).
- [43] J. Ren, D. Schwartzman, D. Callans, S. Brode, C. Gottlieb, and F. Marchlinski, "Intracardiac echocardiography (9 mhz) in humans: methods, imaging views and clinical utility," *Ultrasonend Med. Biol.*, vol. 25, pp. 1077-1086, 1999.
- [44] F. Roithinger, P. Steiner, Y. Goseki, K. Liese, D. Scholtz, A. Sippensgroenewegen, P. Ursell, and M. Lesh, "Low-power radiofrequency application and intracardiac echocardiography for creation of continuous left atrial linear lesions," *J. Cardiovasc. Electrophysiol.*, vol. 10, pp. 680-691, 1999.
- [45] P. Friedman, D. Luria, A. Fenton, T. Munger, A. Jahangir, W. Shen, R. Rea, M. Stanton, S. Hammill, and D.L.Packer, "Global right atrial mapping of human atrial flutter: the presence of posteromedial (sinus venosa region) functional block and double potentials. a study in biplane fluoroscopy and intracardiac echocardiography," *Circulation*, vol. 101, pp. 1568-1577, 2000.

- [46] N. Marrouche, D. Martin, O. Wazni, A. Gillinov, A. Klein, M. Bhargava, E. Saad, D. Bash, H. Yamada, W. Jaber, R. Schweikert, P. Tchou, A. Abdul-Karim, W. Saliba, and A. Natale, "Phased-array intracardiac echocardiography monitoring during pulmonary vein isolation in patients with atrial fibrillation: impact on outcome and complications," *Circulation*, vol. 107, pp. 2710-2706, 2003.
- [47] L. Rao, Y. Ling, R. He, A. Gilbert, N. Frangogiannis, J. Wang, S. Nagueh, and D. Khoury, "Integrated multimodal-catheter imaging unveils principal relationships among ventricular electrical activity, anatomy, and function," *Am J Physiol: Heart Circ Physiol*, vol. 294, pp. H1002-H1009, 2008.
- [48] C. Ding, L. Rao, S. Nagueh, and D. Khoury, "Dynamic three dimensional visualization of the left ventricle by intracardiac echocardiography," *Ultrasound Med. Biol.*, vol. 31, pp. 15-C21, 2005.
- [49] J. Barron, D. Fleet, and S. Beauchemin, "Performance of optical flow techniques," *Int. J. Comp. Vis.*, vol. 12, pp. 43-47, 1994.
- [50] P. Baraldi, A. Sarti, C. Lamberti, A. Prandini, and F. Sgallari, "Evaluation of differential optical flow techniques on synthesized echo images," *IEEE Transactions on Biomedical Engineering*, vol. 43, pp. 259-272, 1996.
- [51] M. Sühling, M. Arigovindan, C. Jansen, P. Hunziker, and M. Unser, "Myocardial motion analysis from B-mode echocardiograms," *IEEE Transactions on Image Processing*, vol. 14, pp. 525-536, 2005.
- [52] V. Behar, D. Adam, P. Lysyansky, and Z. Friedman, "The combined effect of non-linear filtration and window size on the accuracy of tissue displacement estimation using detected echo signals," *Ultrasonics*, vol. 41, pp. 743-753, 2004.
- [53] M. G. Strintzis and I. Kokkinidis, "Maximum likelihood motion estimation in ultrasound image sequences," *IEEE Signal Processing Letter*, vol. 4, pp. 156-157, 1997.
- [54] B. Cohen and I. Dinstein, "New maximum likelihood motion estimation schemes for noisy ultrasound images," *Pattern Recogn.*, vol. 35, pp. 455-463, 2002.
- [55] D. Boukerroui, A. Noble, and M. Brady, "Velocity estimation in ultrasound images: a block matching approach," *Proc. 18th Int. Conf. on Inf. Proc. Med. Imag., Ambleside, UK*, pp. 586-598, 2003.
- [56] J. Revell, M. Mirmehdi, and D. McNally, "Computer vision elastography: speckle adaptive motion estimation for elastography using ultrasound sequences," *IEEE Transactions on Medical Imaging*, vol. 24, pp. 755-760, 2005.

- [57] J. Kybic and M. Unser, "Fast parametric elastic image registration," *IEEE Transactions on Image Processing*, vol. 12, pp. 1427-1442, 2003.
- [58] M. Ledesma-Carbayo, J. Kybic, M. Desco, A. Santos, M. Sühling, P. Hunziker, and M. Unser, "Spatio-temporal nonrigid registration for ultrasound cardiac motion estimation," *IEEE Transactions on Medical Imaging*, vol. 24, pp. 1113-1126, 2005.
- [59] Y. Yong, D. Khoury, "Speckle tracking in intracardiac echocardiography for the assessment of myocardial deformation," *IEEE*, submitted.
- [60] M. Unser, A. Aldroubi, and M. Eden, "B-spline signal processing: Part I - Theory," *IEEE Transactions on Signal Processing*, vol. 41, no.2, pp. 821-832, 1993.
- [61] M. Unser, A. Aldroubi, and M. Eden, "B-spline signal processing: Part II - Efficient design and applications," *IEEE Transactions on Signal Processing*, vol. 41, no.2, pp. 834-848, 1993.
- [62] M. Ledesma-Carbayo, J. Kybic, M. Desco, A. Santos, M. Sühling, P. Hunziker, and M. Unser, "Spatio-temporal nonrigid registration for ultrasound cardiac motion estimation," *IEEE Transactions on Signal Processing Medical Imaging*, vol. 24, pp. 1113-1126, 2005.
- [63] S. Langeland, J. Dhooge, P. Wouters, H. Leather, P. Claus, B. Bijnens, and G. Sutherland, "Experimental validation of a new ultrasound method for the simultaneous assessment of radial and longitudinal myocardial deformation independent of insonation angle," *Circulation*, vol. 112, pp. 2157-2162, 2005.
- [64] M. Suffoletto, K. Dohi, M. Cannesson, S. Saba, and J. Gorcsan, "Novel speckle-tracking radial strain from routine black-and-white echocardiographic images to quantify dyssynchrony and predict response to cardiac resynchronization therapy," *Circulation*, vol. 113, pp. 960-968, 2006.
- [65] J. Wang, D. Khoury, Y. Yue, G. Torre-Amione, and S. Nagueh, "Left ventricular untwisting rate by speckle tracking echocardiography," *Circulation*, vol. 116, pp. 2580-2586, 2007.
- [66] B. Pirat, D. Khoury, C. Hartley, L. Tiller, L. Rao, D. Schulz, S. Nagueh, and W. Zoghbi, "A novel feature-tracking echocardiographic method for the quantitation of regional myocardial function," *J Am Coll Cardiol*, vol. 51, pp. 651-659, 2008.
- [67] R. Azencott, R. Glowinski, J. He, A. Jajoo, Y. Li, A. Martynenko, R. Hoppe, S. Benzekry, S. Little and W. Zoghbi, "Diffeomorphic matching and dynamic

- deformable surfaces in 3D medical imaging”, *Computational Methods in Applied Mathematics*, vol. 10, pp. 235-274, 2010.
- [68] N. Arad, N. Dyn, D. Reisfeld, and Y. Yeshurun, “Image warping: application to facial expressions”, *CVGIP: Graphical Models and Image Proc.*, vol. 56, pp. 161-172, 1994.
- [69] V.I. Arnold, “Sur la géométrie différentielle des groupes de Lie de dimension infinie et ses applications à l’hydrodynamique des fluides parfaits”, *Ann. Inst. Fourier (Grenoble)*, vol. 1, pp. 319-361, 1996.
- [70] N. Aronszajn, “Theory of reproducing kernels”, *Transaction of American Math. Soc.*, vol. 68, pp: 337-404, 1950.
- [71] R. Azencott, “Elastic deformations of soft 3D-shapes”, *IBIS Colloquium on Medical image analysis. The Methodist Hospital*, Kiawah Island, 2006
- [72] R. Azencott, R. Glowinski, and A. Ramos, “A controllability approach to shape identification”, *Applied Math. Letters*, vol. 21, pp. 861-865, 2008.
- [73] R. Azencott, S. Alexander, A. Aggarwal, A. Jajoo, S. Jain, Y. Li, S. BenZekry (MD), S.H. Little (MD), and W. Zoghbi (MD), “New parameters to compare Mitral Annulus shapes extracted from 3D-echocardiography”, submitted, 2010
- [74] R. Azencott, A. Jajoo, S. Jain, Y. Li, A. Martynenko, S. BenZekry (MD), S.H. Little (MD), and W. Zoghbi (MD), “3D-echocardiographic movies analysis: dynamic deformable models for Mitral Valve”, in preparation, 2010
- [75] M.F. Beg, M.I. Miller, A. Trouvé, and L. Younes, “Computing large deformations metric mappings via geodesic flows of diffeomorphisms”, *International Journal of Computer Vision*, vol. 61, pp 139-157, 2005.
- [76] L. G. Brown, “A survey on image registration techniques”, *ACM Computing Surveys*, vol. 24, pp. 326-376, 1992.
- [77] H. Chui and A. Rangarajan, “A new point matching algorithm for non-rigid registration”, *Computer Vision and Understanding*, vol. 89, pp. 114-141, 2003.
- [78] M. Chupin, D. Hasboun, S. Baillet, and altri, “Competitive segmentation of the hippocampus and volumetry in Alzheimer disease”, In: *10th meeting of the Organization for Human Brain Mapping*, 2004.

- [79] Y. Cao, M. I. Miller, R. Winslow, and L. Younes, “Large deformation metric mapping of vector fields”, *IEEE Transactions on Medical Imaging*, vol. 24, pp. 1216-1230, 2005.
- [80] P. Dupuis, U. Grenander, and M. I. Miller, “Variational problems on flows of diffeomorphisms for image matching”, *Quart. Appl. Math.*, vol. 56, pp. 587-600, 1998.
- [81] H. Drury, D.V. Essen, M. Corbetta, and A. Snyder, “Surface based analysis of human cerebral cortex”, *Brain Warping*, 337-363, Academic Press, 1999.
- [82] L. Greengard, and J Strain, “Fast Gauss transform”, *SIAM Jour. Sci. Stat. Comp.*, vol. 12, pp. 79-94, 1991.
- [83] J. Glaunès, “Transport par difféomorphismes de points, de mesures et de courants pour la comparaison des formes et l’anatomie numérique”, PhD Thesis, Université Paris 13, 2005.
- [84] J. Glaunès, A. Qiu, M.I. Miller, and L. Younes, “Large deformation diffeomorphic metric curve mapping”, *International Journal on Computer Vision*, vol. 80, pp. 317-336, 2008.
- [85] J. Glaunès, A. Trouvé, and L. Younes, “Diffeomorphic matching of distributions: A new approach for unlabelled point-sets and sub-manifolds matching”, in: *2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’04), Vol. 2*, IEEE Computer Society, Washington, DC, 2004, pp. 712-718.
- [86] J. Glaunès, M. Vaillant, and M. I. Miller, “Landmark matching via large deformation diffeomorphisms on the sphere”, *J. Math. Imaging Vis.*, vol. 20, pp. 179-200, 2003.
- [87] H. Guo, A. Rangarajan, S. Joshi, and L. Younes, “Non-rigid registration of shapes via diffeomorphic point matching”, in: *Proceedings of the 2004 Int. Symposium on Biomedical Imaging: From Nano to Macro*, IEEE Computer Society, Washington, DC, 2004, pp. 924-927.
- [88] H. Guo, A. Rangarajan, and S. Joshi, “Diffeomorphic point matching”, *Handbook of Mathematical Models in Computer Vision*, Springer, Berlin-Heidelberg-New York, 2006, pp. 205-219.
- [89] U. Grenander, and M. I. Miller, “Computational anatomy: an emerging discipline”, *Quart. App. Math.*, vol. 56, pp. 617-694, 1998.

- [90] H. Guo, “Diffeomorphic point matching with applications in medical image analysis”, Ph.D. thesis, Department of Computer and Information Science and Engineering, University of Florida, Gainesville, FL, 2005.
- [91] S. Joshi and M. Miller, “Landmark matching via large deformation diffeomorphisms”, *IEEE Transactions on Image Processing*, vol. 9, pp. 1357-1370, 2000.
- [92] E. Klassen, A. Srivastava, W. Mio, and S. Joshi, “Analysis of planar shapes using geodesic paths on shape spaces”, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 26, pp. 372-383, 2004.
- [93] J.L. Lions, “Controlabilité exacte, perturbation et stabilisation des systèmes distribués”, Masson, Paris, 1988.
- [94] I. Miller, A. Trounev, and L. Younes, “On the metrics and Euler-Lagrange equations of computational anatomy”, *Ann. Review of Biomedical Engineering*, vol. 4, pp. 375-405, 2002.
- [95] M. I. Miller and L. Younes, “Group action, diffeomorphism and matching: A general framework”, *Int. J. Comp. Vis.*, vol. 41, pp. 61-84, 2001.
- [96] S. Saitoh, “em Theory of Reproducing Kernels and its Applications”, *Pitman Research Notes in Mathematics*, Vol. 189, Longman, 1988.
- [97] L. Tartar, “Introduction to Sobolev Spaces and Interpolation Theory”, Springer, Berlin-Heidelberg-New York, 2007.
- [98] A. Trounev, “Diffeomorphisms groups and pattern matching in image analysis”, *Int. J. of Comp. Vis.*, vol. 28, pp. 213-221, 1998.
- [99] L. Younes, “Computable elastic distances between shapes”, *SIAM J. Appl. Math.*, vol. 58, pp. 565-586, 1998.
- [100] B. Zitová and J. Flusser, “Image registration methods: a survey”, *Image and Vision Computing*, vol. 21, pp. 977-1000, 2003.
- [101] E. Zuazua, “Controllability and observability of partial differential equations”, Chapter 7 in *Handbook of Differential Equations*, Vol. 3, Evolutionary Equations, Elsevier, Amsterdam, pp: 527-621, 2007.