

**IMAGE GUIDED RESPIRATORY MOTION
ANALYSIS:
TIME SERIES AND IMAGE REGISTRATION**

by
Dan Ruan

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Electrical Engineering: Systems)
in The University of Michigan
2008

Doctoral Committee:

Professor Jeffrey A. Fessler, Chair
Professor Alfred O. Hero III
Professor Charles R. Meyer
Associate Professor James M. Balter
Assistant Professor Selim Esedoğlu

© Dan Ruan 2008
All Rights Reserved

To Mum, Dad, and my best friend Hua He.

ACKNOWLEDGEMENTS

I thank my previous advisor Professor David A. Castañon at Boston University, who trustingly led me to the door of applied mathematics, and let me wander at will. I enjoyed the freedom of exploration and owe my confidence to perform independent research to his encouragement.

This work would not be possible without my advisor Professor Jeffrey A. Fessler. Jeff is a natural; uncork the bottle and the charm flows. Working with him is fun and enjoyable. I thank Professor James M. Balter for welcoming me to Michigan with open arms, and generously showered me with support and friendship. I am grateful to Professor Alfred O. Hero, whose intellectual power and enlightening guidance have always inspired me to push for better. Much thanks goes to Professor Charles R. Meyer for his expertise, valuable feedback, and patience with a humble graduate student like me. I am indebted to Professor Selim Esedoğlu, from whom I appreciated the professionalism, ambition and passion for truth of a young scientist. I thank Professor Marc. L. Kessler for motivating much of the research work, offering computational support and data access. I have also benefited from discussions with Professor William C. Karl (BU), Professor Jeffrey Rauch (UMich), Professor Susan Murphy (Umich), Professor Paul Keall (Stanford), Professor Steve B. Jiang (UCSC), Dr. Gregory C. Sharp (Mass General Hospital) and many other people over the years.

My colleagues and teammates have always been a source of wisdom for me, they have made my time in Michigan a pleasant experience. I wish to acknowledge my lab-

mates: Matt Jacobson, Anastasia Yendiki, Yingying Zhang, Rongping Zeng, Eran Bashan, Somesh Srivastava, Se Young Chun, Kim Khalsa, Amanda Funai, Hugo Shi, Ted Way and Yong Long. Special thanks goes to Drs. Derek Justice, Raviv Raich, Mark Kliger, with whom I enjoyed many vivid discussions. Michigan Argentina Tango club has been the ultimate fun place for my last year of Ph.D work, and its activities constitutes much of my healthy breaks from work.

Finally, I am indebted to my family. I feel sorry for not being there at his bedside for my late grandfather during the process of finishing this thesis. I thank my parents for their love, care, and nurtue with freedom and critical thinking. My appreciation goes to Dr. Hua He, whose good humor and sensitivity not only brought me much happiness, but also supported me through difficult times. Last but not least, I am grateful to my American grandmother, Ms. Natalie Rammel, who has taught me to appreciate life, and to always smile.

This work was financially supported in part by NIH grant P01-CA59827, International student fellowship and Barbour fellowship at the University of Michigan.

TABLE OF CONTENTS

DEDICATION	ii
ACKNOWLEDGEMENTS	iii
LIST OF FIGURES	vii
LIST OF TABLES	x
LIST OF APPENDICES	xi
ABSTRACT	xii
CHAPTER	
I. Introduction	1
II. Adaptive Ellipse Tracking and its Application in Estimating Respiratory Drifting ¹	3
2.1 Ellipse Fitting Model for Static Data	3
2.2 An Iterative Algorithm for Solving the Generalized Eigen Problem	5
2.3 Adaptive Ellipse Fitting	9
2.3.1 Sliding Window Adaptation	10
2.3.2 Discounting Adaptation	10
2.4 Application to Tracking Respiratory Mean Drift	10
2.4.1 Application Background	11
2.4.2 Experiment Setup	11
2.4.3 Results	13
2.4.4 Discussion	18
2.5 Generalized Fitting Cost for Robust Estimation	19
2.6 Appendix	23
III. Regularized Nonrigid Image Registration	26
3.1 General Optimization Formulation for Regularized Registration	27
3.1.1 Data Dissimilarity (Infidelity) Measure	27
3.2 Tissue-type Dependent Rigidity Regularization	28
3.2.1 Regularization Design	28
3.2.2 Experiment and Test Results	31
3.2.3 Discussions	35
3.3 Discontinuity-Preserving Regularization	37
3.3.1 Indiscriminate Discontinuity Preserving Regularization	39
3.3.2 Discriminative Shear Preserving Regularizer	45

¹This chapter is based on materials from [97,99].

3.3.3	Discussions	59
3.4	Equivalence Between Two L_2 Div- L_1 Curl Regularizations*	60
IV.	Fundamental Performance Analysis in Image Registration Problems: Cramér-Rao Bound and its Variations	64
4.1	Model - the Ideal v.s. Commonly Used	64
4.2	Cramér-Rao Bound and its Asymptotic Behavior	65
4.3	Relating to MCRB	69
4.4	An Alternating Minimization Algorithm	69
4.5	Comparison with Conventional Methods: CRB v.s. M-estimate	70
4.6	A Simple Example	73
V.	Summary and Future Work	91
5.1	Summary	91
5.2	Future Work	92
APPENDIX	94
BIBLIOGRAPHY	127

LIST OF FIGURES

Figure

2.1	Effect of drift compensation for gating system.	12
2.2	Illustration of ellipse fitting performance of the proposed method.	14
2.3	Comparison of moving average (MA) and ellipse fitting estimator for mean position tracking.	15
2.4	Effect of window length L on tracking performance.	16
2.5	Effect of discount factor γ on tracking performance.	16
2.6	Mean tracking for RPM data with window size determined by period estimator.	17
2.7	Overall RMS error (across all patients) as a function of sampling rate.	17
2.8	Evolution of robust fitted ellipse with the gradient projection method.	23
3.1	Illustration of stiffness factor $s(\cdot)$	30
3.2	Different views of the original CT data and tissue information inferred from it.	32
3.3	Deformed source image (green) overlaid with target image (dark blue) for comparison of intensity match.	33
3.4	Geometry extracted from registration results: target (blue) vs. deformed source (white).	34
3.5	Illustration of landmark data on thorax CT.	35
3.6	Registration error for different methods: TPS, BSP and Regularized BSP.	36
3.7	Comparison of 3-dimensional Euclidean landmark registration error.	37
3.8	Truncated quadratic regularization with varying scale.	42
3.9	Registration comparison between Tikhonov (Tik) and Truncated quadratic (TQ) regularizations.	43
3.10	Comparison of deformation fields.	44
3.11	Regularization results for a colliding/folding flow.	47
3.12	Regularization results for a vacuum generating flow.	48

3.13	Regularization results for a sliding flow.	49
3.14	Regularization results for a flow with simultaneously significant divergence and curl components.	50
3.15	Simulated sliding blocks and the ground truth deformation.	50
3.16	Registration performance comparison: registration error v.s. intensity discrepancy.	51
3.17	Best estimation results (relative to the ground truth deformation) from various regularization methods.	52
3.18	Registration results of CT data with various regularizations.	54
3.19	Vertical component of the deformation from CT registration.	55
3.20	Quiver plot overlain with image intensity for L_1 regularized image registration.	55
3.21	Quiver plot overlain with image intensity for L_2 regularized image registration.	56
3.22	Quiver plot overlain with image intensity for discriminately regularized image registration.	56
3.23	Coronal view for 3D discriminate registration.	57
3.24	Sagittal view for 3D discriminate registration.	58
3.25	Illustration for the three cases of feasible region: strictly convex and everywhere differentiable, nonstrictly convex and everywhere differentiable, nondifferentiable.	61
3.26	Equivalence between unconstrained and constrained formulation.	62
4.1	Bias and variance approximation obtained from explicit solution for conventional M-estimate.	77
4.2	Bias and variance approximation for M-estimate obtained from expansion about $(\bar{\alpha}, \bar{z})$	79
4.3	Bias and variance approximation of M-estimate obtained from expansion about $(\check{\alpha}, \check{z})$	83
4.4	Bias and variance approximation of ML-estimate obtained from expansion about $(\bar{\alpha}, \bar{z})$	89
A.1	Proper initialization helps to avoid suboptimal (nonphysical) local minimum.	99
A.2	Clinical significant performance metrics v.s. Root Mean Squared Error (RMSE).	99
A.3	RMSE, Dose error (%), PTV margin (cm), 95% dose coverage of modified cosine model v.s. projection model.	100
A.4	Trajectory fitting with projection model and modified cosine model.	101
A.5	Prediction of breathing trajectory with projection model.	102
B.1	State-space distance and local regression weight assignment.	105
B.2	Typical breathing trajectories: (a) rapid yet regular breath; (b) slow yet irregular breath.	109

B.3	Covariate-response relationship with lag-one augmented state.	109
B.4	Effect of dynamically updating the training atlas.	112
B.5	Comparison of prediction performance using dynamic update with moving windowed and expanding training atlas.	113
B.6	Collective performance comparison for different lookahead lengths.	114
B.7	Collective performance comparison for different sampling rates ψ	114
C.1	Example of a breathing trajectory with respiratory hysteresis.	121
C.2	Scatter plot showing the data samples in augmented external state space with the colors indicating internal AP value.	121
C.3	Correspondence relations in augmented state space and their linear fittings.	122
C.4	Histogram of paired differences between the RMS errors of the direct and the augmented methods.	123
C.5	Estimation performance comparison among different methods.	124
C.6	Estimation error as a function of lag length for state augmentation.	125
C.7	Beam-wise 3D RMSE (mm) for patients 4-6.	126

LIST OF TABLES

Table

2.1	RPM Dataset information	12
2.2	Mean Estimation Performance	16
A.1	Dataset information and Experiment Results	97
B.1	RPM Dataset information.	108
B.2	Comparison of Prediction Performance among Static Training, Dynamic Expanding Training, and Updating Training with Moving Window	110
C.1	Description of study participants.	117
C.2	Estimation Error Table	122

LIST OF APPENDICES

Appendix

A.	A Breathing Pattern Irregularity Index with Projection-based Method ²	95
	A.1 An Irregularity Index based on Projection Distance	95
	A.2 Material and Verification Design	97
	A.3 Results and Discussions	98
	A.4 Summary	102
B.	Real-time Prediction of Respiratory Motion based on Nonparametric Local Regression Methods ³	103
	B.1 Methods	104
	B.1.1 Model Setup and General Notations	104
	B.1.2 Basic Local Weighted Regression	104
	B.1.3 Variations that Potentially Improve Prediction Performance	106
	B.1.4 Baseline Methods for Comparison	107
	B.2 Materials	108
	B.3 Results and Discussions	108
	B.3.1 Scatter Plot in Augmented Space	108
	B.3.2 Local Weighted Regression without Temporal Discounting	109
	B.3.3 Robust Local Regression with Iterative Weighting	110
	B.3.4 The Effect of Dynamically Updating the Training Set	110
	B.3.5 The Effect of Measurement Rate and Lookahead Length	111
	B.4 Conclusion and Future Work	113
C.	Inference of Internal Respiratory Tumor Motion from External Surrogates: A State Space Augmentation Approach in Modeling Hysteresis ⁴	116
	C.1 Methods and Materials	117
	C.1.1 Data Description	117
	C.1.2 A General Correspondence Model	117
	C.1.3 Hysteresis and State Augmentation	119
	C.1.4 Choice of Lag Length	120
	C.1.5 Adaptivity of the Correspondence Map	120
	C.2 Results and Discussions	120
	C.3 Conclusion and Future Work	123

²This chapter is based on materials from [101]

³This chapter is based on materials from [98]

⁴This chapter is based on materials from [100]

ABSTRACT

IMAGE GUIDED RESPIRATORY MOTION ANALYSIS: TIME SERIES AND IMAGE REGISTRATION

by
Dan Ruan

Advisor: Jeffrey A. Fessler

The key feature of image guided radiation therapy (IGRT) systems is improved efficiency in conformal dose delivery by extracting, modeling and predicting tumor movement with imaging techniques. To harvest the benefit of an IGRT system, two major problems have to be solved: motion modeling and image processing. This thesis investigates these issues. In particular, we focus on the application of treating tumors in the thoracic and upper abdominal region, where respiratory motion is the dominant factor for tumor movement. The characteristics of respiration makes motion modeling difficult, as breathing trajectories are semi-periodic with drifting in mean position, frequency and phase. Clinical practice shows large variation of breathing patterns among different individuals, making it necessary to quantify the regularity/reproducibility of a respiratory trace to determine the applicability of certain treatment methods. To this end, we have proposed a subspace projection method to quantitatively evaluate the semi-periodicity of a given observation trace. Extracting tumor location from diagnostic imaging, albeit informative and accurate, incurs radiation dose, which may result in normal tissue complication. To minimize diag-

nostic dose, it is desirable to obtain radiographic observations at low frequency and use external surrogates to infer internal tumor motion without radiation. The sparsity of the internal readout from diagnostic imaging together with the consideration of system latency require methods that predict accurately over a long time range. These concerns introduce additional challenges in applying classic time-series analysis techniques, and special structured models are needed to incorporate prior knowledge (*e.g.*, semi-periodicity in respiratory motion) for improved performance. Physiological hysteresis further complicates external-internal inference and proper modeling is desired to estimating such a relationship. In this thesis, we have investigated regression techniques for real-time tracking and prediction, shape modeling for robust tracking with minimum observation and external-internal inference estimation.

Image processing is another crucial component of IGRT. In particular, accurate tracking and monitoring of tumor evolution, and efficient propagation of dose assignment require accurate image registration. The solution to the registration problem needs to reflect physical priors and constraints. Adopting a regularized optimization setup, we investigated a penalty function design that accommodates tissue-type-dependent elasticity information. To properly account for the sliding effects at motion interfaces, we have studied a class of discontinuity-preserving regularizers that yield smooth solutions in most regions, while allowing discontinuities in the estimated motion field. We have further distinguished two types of singularities in the deformation field, *i.e.*, collision/vacuum generating flow *v.s.* shear with the Helmholtz decomposition. Applying different regularizers to each component discourages the deformation from the first type of unphysical singularities while preserves large shear discontinuities.

The medical image registration field needs good validation and performance evaluation tools. A most general analytical evaluation for image registration is challenging. We have,

however, during the course of this work, performed a preliminary analytical study. We proposed a set of statistical generative models and provided bias and variance estimates for certain estimators. In particular, we have investigated the approximate performance of the maximum-likelihood estimator corresponding to the generative model and the commonly adopted M-estimator. A simple example suggests that the approximation is reasonably accurate.

Our studies in both time series analysis and image registration constitute essential building-blocks for clinical applications such as adaptive treatment. Besides their theoretical interests, it is our sincere hope that with further justifications, the proposed techniques would realize its clinical value, and improve the quality of life for patients.

CHAPTER I

Introduction

The two core components of image guided radiotherapy (IGRT) systems are image processing and radiation delivery. The image processing module extracts tumor status information and feeds it into the treatment delivery system. In particular, the motion of the tumor volume must be tracked and predicted with high accuracy for subsequent localized target treatment; the movement of the whole region under radiation should be monitored to ensure proper dose delivery, to avoid radiation to critical tissues such as heart and spine, and to minimize normal tissue complication probability (NTCP). To this end, an accurate and efficient image registration method is critical. Registration between two (or a sequence of) images estimates the deformation among different image acquisitions, captures the evolution of the region of interest (ROI), and dynamically propagates treatment plans. Despite the vast literature on image registration studies, a good quantitative evaluation tool is unfortunately absent. A performance study to the most generality is challenging, given the nonlinear nature of the registration problem. However, questions then arise as to whether the performance of image registration is limited by a model setup (objective function design) or the behavior of the minimization algorithms deployed (local minima issues). Furthermore, to minimize diagnostic radiation dose, only sparse observations of the internal tumor location are available to the treatment delivery system. For an effective real-time gating system, a reliable mean position estimator is crucial to adaptively control the positioning of the gating window. How to efficiently extract such control information from a minimum amount of data is a key issue. This thesis considers these various questions in three parts.

The first part, dealing with the adaptive mean tracking problem, is the subject of Chapter II. Given a sparsely sampled respiratory trajectory that has drifts in mean, frequency and phase, we aim to extract the mean trace in real time. To solve this loosely defined problem, we resort to a data-based approach which incorporates the semi-periodic nature of breathing motion. In particular, in the state space that is augmented via time lagging, we model the observations as samples in a sequence of time varying ellipses and extract the projection of the center of such ellipses as the real-time estimate of mean position. Formulated as a minimization problem with respect to the algebraic distance, the static ellipse fitting problem can be solved by generalized eigen-decomposition. We introduce a recursive-least squares (RLSE) structured algorithm which naturally leads to a dynamic adaptive solution in a slowly temporal-varying environment. Asymptotic convergence of the proposed algorithm is derived. In addition, we generalize the original least squares fitting problem to a robust estimation setting so that the solution is insensitive to reasonable amounts of outliers, what may be caused by abrupt body movement or noisy data. We prove that the feasible region is a union of two convex sets, analyze the geometry of both the feasible region and the functional value, and apply gradient projection method to solve the adaptive problem. Experimental results with both simulated and clinical data demonstrate feasibility of the proposed methods.

The remainder of the thesis studies the image registration problem - another key component in IGRT. We discuss objective design, optimization issues and quantitative evaluation of registration performance; these aspects are of interest for general image processing as well.

Chapter III focuses on designing regularizers that convey prior information in optimization-based registration for thoracic images. We consider two types of regularization design: one accounts for different rigidity levels for various tissue types and the other accommodates sliding effects along motion boundaries. Tissue-type-rigidity regularization is realized by penalizing the deviation of local transformation Jacobian

from orthogonal; sliding regularization is studied by first generalizing edge-preserving regularization from image denoising problems, and then Helmholtz decomposing the flow to differentiate between the collision/vacuum generating component and the shear discontinuities. By regularizing the divergence and curl components separately, we avoid the first type of unphysical discontinuity, but preserve the latter one that corresponds to shear flow.

Given any image registration method, its results should be validated and the performance evaluated. Unfortunately, despite vast literature on image registration algorithms, validation study mostly relies on either simulation/phantom study or manually placed marker locations. The first approach completely ignores intrinsic information content of the input images, and may result in unreasonably stringent requirement (such as “correct” alignment of uniform areas - the null space of any registration operator). The latter one, on the other hand, is biased towards high gradient regions where human observers can identify reliable correspondence and oblivious of the less obvious clues. As a first step in systematically studying the fundamental performance limit of a registration model, Chapter IV presents a statistical generative model and the corresponding maximum-likelihood (ML) estimator. The bias and variance of this estimator is studied via Cramér-Rao bound analysis. For the commonly employed energy minimization based approaches, the local behavior of the corresponding M-estimate is analyzed using implicit function theorem and Taylor expansion. A simple example suggests reasonable accuracy of the adopted approximation and may lead to wider generality upon further investigation.

The main contributions of this thesis can be summarized as follows.

- A novel data-driven mean tracking model for sparsely sampled semi-periodic data. More specifically, a state augmentation setup and a formulation with algebraic distance that results in closed-form solution from generalized eigen decomposition. An efficient ellipse tracking algorithm based on subspace decomposition that dynamically adapts to slowly varying trends. Conditions and proof for asymptotic convergence of the proposed algorithm. Analysis of the structure of the feasible parameter set. A robust extension of the least squared problem to achieve robustness to outliers, a gradient projection algorithm for solving the optimization problem, and its adaptive generalization.
- A tissue-type dependent regularization that encourages locally rigid behavior, where appropriate.
- An original discontinuity preserving regularization for nonrigid image registration that preserves motion boundaries.
- An original statistical generative model for image registration. Bias and variance analysis for the maximum likelihood estimator. An M-estimate analysis of the conventional energy based registration methods. Empirical comparison with a simple example.
- A new subspace projection based method that quantifies the reproducibility of a temporal trajectory (Appendix A).
- A novel nonparametric local regression method in the augmented state space for real-time prediction of respiratory motion (Appendix B).
- A state space augmentation approach to account for hysteresis for inferring internal tumor motion from external surrogates (Appendix C).

The above remarks describe the major material in this thesis. To maintain the self-containedness of each topic, we provide relevant background at the beginning of each chapter. In the course of this study, we came to certain analyses and preliminary results that are marginally related to the main theme of this thesis, but have potential for integration upon further development. We provide them as optional sections and mark the titles with an asterisk. These sections can be skipped without loss of continuity.

CHAPTER II

Adaptive Ellipse Tracking and its Application in Estimating Respiratory Drifting ¹

Good ellipse fitting methods are desirable in pattern recognition and computer vision. Simple low dimensional shape models are often used to fit noisy high dimensional observation data for increased robustness. Ellipses, as the projection of circular shapes, are common among observations from natural and artificial objects (*e.g.*, human faces, tires, etc), and are among the most interesting shape models [26,44,89]. In addition, ellipses also have potential applications in describing dynamical systems that exhibit semi-periodicity and hysteresis. Using algebraic distance as the data fitting metric, ellipse fitting problems can be formulated in a convex optimization setting, with quadratic constraints. Its solution involves looking for the eigen vector corresponding to the largest eigenvalue in a generalized eigen decomposition problem. In this chapter, we develop an adaptive method to dynamically fit the ellipse model, analyze the convergence of the proposed algorithm, and discuss its application to estimating drifting in respiratory motion. Section 2.1 formulates the ellipse fitting problem into the framework of generalized eigen decomposition. Section 2.2 proposes an iterative algorithm for solving the generalized eigen problem. Section 2.3 considers adaptivity with data stream. Section 2.4 applies the proposed algorithm to tracking mean drifting of respiratory motion. Section 2.6 provides relevant proofs used in this chapter.

2.1 Ellipse Fitting Model for Static Data

We model ellipses using a general quadratic curve equation. Let (x, y) denote the coordinates of a point in the 2-dimensional state space, and define $\mathbf{z} = [x^2 \ xy \ y^2 \ x \ y \ 1]^T$, where superscript T denotes transpose. Then point (x, y) falls on the ellipse parameterized by $\mathbf{a} = [a \ b \ c \ d \ e \ f]^T$ if and only if it satisfies the following quadratic curve equation:

$$(2.1) \quad F(\mathbf{a}, \mathbf{z}) = \mathbf{a}^T \mathbf{z} = ax^2 + bxy + cy^2 + dx + ey + f = 0,$$

with negative discriminant, *i.e.*, $b^2 - 4ac < 0$.

The center (x_0, y_0) of the ellipse parameterized with \mathbf{a} is given by:

$$(2.2) \quad \begin{aligned} x_0 &= \frac{2cd - bf}{b^2 - 4ac} \\ y_0 &= \frac{2af - bd}{b^2 - 4ac}. \end{aligned}$$

From (2.1), a sample \mathbf{z}_i lies on a given ellipse parameterized by \mathbf{a} if and only if $F(\mathbf{a}, \mathbf{z}_i) = 0$. This motivates the use of $F^2(\mathbf{a}, \mathbf{z}_i)$ as a measure of deviation of the sample from the ellipse. This is known as “algebraic distance” which coincides with Euclidean distance in the case F is a plane. It is computationally beneficial to adopt this discrepancy measure so that the collective distances for N samples can be conveniently written in standard matrix form and manipulated with classic least-squares approaches as in

¹This chapter is based on materials from [97,99].

(2.3). For observed samples of the form (x_i, y_i) , $i = 1, 2, \dots, N$, we want to find the ellipse parameter \mathbf{a} that minimizes the following cost function:

$$(2.3) \quad \sum_{i=1}^N F^2(\mathbf{a}, \mathbf{z}_i) = \mathbf{a}^T \mathbf{S} \mathbf{a},$$

where we define the 6×6 empirical correlation matrix $\mathbf{S} \triangleq \sum_{i=1}^N \mathbf{z}_i \mathbf{z}_i^T$.

The minimizer of (2.3) is invariant to a constant scaling applied to \mathbf{a} , so we impose the constraint that $b^2 - 4ac = -1$, or equivalently in matrix form $\mathbf{a}^T \mathbf{C} \mathbf{a} = 1$ with

$$(2.4) \quad \mathbf{C} = \begin{bmatrix} \tilde{\mathbf{C}} & \mathbf{0}_{3 \times 3} \\ \mathbf{0}_{3 \times 3} & \mathbf{0}_{3 \times 3} \end{bmatrix},$$

where $\tilde{\mathbf{C}} \triangleq \begin{bmatrix} 0 & 0 & 2 \\ 0 & -1 & 0 \\ 2 & 0 & 0 \end{bmatrix}$, and $\mathbf{0}_{3 \times 3}$ denotes a 3×3 matrix of zeros. In other words, our ellipse fitting requires minimizing $\mathbf{a}^T \mathbf{S} \mathbf{a}$ subject to the constraint that $\mathbf{a}^T \mathbf{C} \mathbf{a} = 1$.

Introducing the Lagrangian multiplier λ and differentiating, we need to solve the system of equations:

$$(2.5) \quad \begin{aligned} \mathbf{S} \mathbf{a} - \lambda \mathbf{C} \mathbf{a} &= \mathbf{0}, \\ \mathbf{a}^T \mathbf{C} \mathbf{a} &= 1. \end{aligned}$$

We solve this using the generalized eigen-decomposition of the pair (\mathbf{S}, \mathbf{C}) . Because \mathbf{S} has the form of a covariance matrix, it is nonnegative semi-definite. We assume hereafter that there are enough data samples and \mathbf{S} is full-rank, *i.e.*, its eigenvalues are strictly positive. In particular, by Theorem 1 in ² [31], the pair (\mathbf{S}, \mathbf{C}) has exactly one positive generalized eigenvalue and it corresponds to the unique local minimum of the Lagrangian. The corresponding eigen vector is the optimal solution to the ellipse parameter in (2.3). Let (λ, \mathbf{u}) be the solution to the generalized eigenvalue problem $\mathbf{S} \mathbf{a} = \lambda \mathbf{C} \mathbf{a}$ with $\lambda > 0$, then $(\lambda, \mathbf{a}^* = \sqrt{\frac{1}{\mathbf{u}^T \mathbf{C} \mathbf{u}}} \mathbf{u})$ is the solution to the constrained minimization problem in (2.5). The rank deficiency of \mathbf{C} can cause instability issues if a conventional generalized eigen-decomposition algorithm were applied, *e.g.* [92] without caution. If analyzed properly, however, its sparsity may reduce computation.

For later convenience, we first introduce some notations, then restate the corrected lemma 1 in [31] and prove it. We denote the generalized spectra as $\sigma(\mathbf{A}, \mathbf{B}) \triangleq \{\lambda : \mathbf{A} \mathbf{v} = \lambda \mathbf{B} \mathbf{v}\}$. Analogous to the case of a single normal matrix, we define the condition number of a generalized eigen decomposition as $\kappa(\mathbf{A}, \mathbf{B}) = \left| \frac{\lambda_{\max}(\mathbf{A}, \mathbf{B})}{\lambda_{\min}(\mathbf{A}, \mathbf{B})} \right|$ where $\lambda_{\max}(\mathbf{A}, \mathbf{B})$, $\lambda_{\min}(\mathbf{A}, \mathbf{B})$ denote the maximal and minimal (by moduli) generalized eigenvalues of (\mathbf{A}, \mathbf{B}) . The signature ³ $i(\mathbf{A})$ of a real symmetric matrix \mathbf{A} is the number of positive, negative, and zero eigenvalues of the corresponding matrix. ⁴

Lemma II.1. *The number of positive, negative, and zero generalized eigenvalues of (\mathbf{S}, \mathbf{C}) , where $\mathbf{S} \in \mathfrak{R}_{n \times n}$ is a symmetric ⁵ positive definite matrix and $\mathbf{C} \in \mathfrak{R}_{n \times n}$ is symmetric, are the same as the signature of the constraint matrix \mathbf{C} .*

Proof. Symmetric positive definite matrix \mathbf{S} can be decomposed as $\mathbf{Q}^T \mathbf{Q}$ with \mathbf{Q} invertible. We can subsequently rewrite the generalized eigen problem $\mathbf{S} \mathbf{u} = \lambda \mathbf{C} \mathbf{u}$ as $\mathbf{Q}^T \mathbf{Q} \mathbf{u} = \lambda \mathbf{C} \mathbf{u}$. Apply a change of basis $\mathbf{v} = \mathbf{Q} \mathbf{u}$ and get $\mathbf{v} = \lambda (\mathbf{Q}^{-1})^T \mathbf{C} \mathbf{Q}^{-1} \mathbf{v}$. Thus, λ is the eigenvalue of $(\mathbf{Q}^{-1})^T \mathbf{C} \mathbf{Q}^{-1}$. Let $i(\mathbf{C})$ denote the signature of \mathbf{C} , then by Sylvester's law of inertia [131], which states that the signature of \mathbf{C} is invariant under congruence transform, we have $i(\mathbf{C}) = i((\mathbf{Q}^{-1})^T \mathbf{C} \mathbf{Q}^{-1})$. Therefore, the number of positive, negative and zero eigenvalues of $\sigma(\mathbf{S}, \mathbf{C})$ are the same as the signature of \mathbf{C} . \square

²Both the statement and the proof of lemma 1 in [31] are flawed, but the result in theorem 1 is correct.

³Signature is often defined with respect to a symmetric bilinear (quadratic) form; they are the same objects, viewed from different perspective.

⁴In general, eigen decomposition (spectral theorem) applies to normal matrices, which may not necessarily require real symmetry. However, many applications deal with real symmetric matrices, including the one under consideration here.

⁵This is the part missing from [31].

2.2 An Iterative Algorithm for Solving the Generalized Eigen Problem

It is desirable to have an adaptive algorithm in the presence of system variations. In particular, a rank-1 update is needed for the data covariance matrix S every time a new observation becomes available. Assuming sufficient initial data, new data should only mildly perturb the system. Therefore, if we have an iterative algorithm for the static system with enough tolerance to the initial conditions, we shall be confident in using the last state estimate of the one-time-step-lag system as the initialization for the iteration with the new system. In this section, we provide such an iterative algorithm, prove its asymptotic convergence, and discuss the region of attraction.

We first recall some observations from previous discussions:

- \mathbf{a}^* is identical to the generalized eigenvector \mathbf{u} up to a normalization factor that is easy to compute. From now on, we focus on deriving an iteration on the generalized eigen vector, no longer distinguishing between \mathbf{u} and \mathbf{a} , assuming no confusion would result from such notational convenience.
- A simple eigen decomposition of C yields that $\sigma(C) = \{2, 0, 0, 0, -1, -2\}$. From lemma II.1, and the minimization setup, we are interested in tracking only the generalized eigen vector that corresponds to the unique positive generalized eigen value of (S, C) . It is equivalent to finding the eigen vector for the largest eigenvalue, in other words, tracking the dominant eigen subspace.

We first discuss a method based on fixed-point analysis and generalized Rayleigh quotient to compute the solution to generalized eigen problems (A, B) with both A, B full rank.

If (λ, \mathbf{u}) satisfy the generalized eigen decomposition relation:

$$(2.6) \quad A\mathbf{u} = \lambda B\mathbf{u},$$

then we can generalize the Rayleigh-Ritz theorem and show (derivation provided in Section 2.6) that the generalized eigen vectors \mathbf{u} correspond to the stationary points of the energy ratio function:

$$(2.7) \quad J(\mathbf{u}) = \frac{\mathbf{u}^T A \mathbf{u}}{\mathbf{u}^T B \mathbf{u}},$$

and evaluating J at the eigenvectors results in the corresponding generalized eigenvalues.

In fact, the largest generalized eigenvalue λ is the global maximum of J . Substituting $\lambda = J(\mathbf{u}) = \frac{\mathbf{u}^T A \mathbf{u}}{\mathbf{u}^T B \mathbf{u}}$ in (2.6) yields

$$A\mathbf{u} = \frac{\mathbf{u}^T A \mathbf{u}}{\mathbf{u}^T B \mathbf{u}} B\mathbf{u}.$$

By assumption, A is full rank, we may multiply the above expression with A^{-1} on both sides and obtain:

$$(2.8) \quad \mathbf{u} = \frac{\mathbf{u}^T A \mathbf{u}}{\mathbf{u}^T B \mathbf{u}} A^{-1} B \mathbf{u}.$$

Equation (2.8) suggests that the generalized eigen vector is a fixed point for the iterative map

$$(2.9) \quad f : \mathbf{u} \rightarrow \frac{\mathbf{u}^T A \mathbf{u}}{\mathbf{u}^T B \mathbf{u}} A^{-1} B \mathbf{u}.$$

Furthermore, the energy ratio function evaluated at the fixed point is exactly the generalized eigenvalue that corresponds to the fixed point \mathbf{u} .

To use (2.9) to iteratively solve (2.5), we first separate the essential subspace from the nuisance ones, by decomposing the empirical correlation matrix S into block form as follows:

$$S = \begin{bmatrix} E & B \\ B^T & D \end{bmatrix}.$$

We also define the Schur complement of the block D in matrix S as $\tilde{S} \triangleq E - B D^{-1} B^T$. The decomposition of the constraint matrix C is given by (2.4).

We provide an iterative scheme to compute the generalized eigen vector \mathbf{a} and prove its asymptotic convergence.

Theorem II.2. *Let the iteration of \mathbf{a} be given by*

$$(2.10) \quad \mathbf{a}_{n+1} = \eta_n \frac{\mathbf{a}_n^T \mathbf{W} \mathbf{a}_n}{\mathbf{a}_n^T \mathbf{C} \mathbf{a}_n} \mathbf{S}^{-1} \mathbf{C} \mathbf{a}_n + (1 - \eta_n) \mathbf{a}_n,$$

where

$$\mathbf{W} = \begin{bmatrix} \mathbf{E} & 0 \\ 0 & -\mathbf{D} \end{bmatrix},$$

and $\eta_n \in (0, 1)$ is asymptotically bounded above by $\frac{2}{\kappa+1}$ with κ being the condition number of $(\tilde{\mathbf{S}}, \tilde{\mathbf{C}})$ ⁶. Then \mathbf{a}_n converges asymptotically to the eigen vector that corresponds to the unique positive eigenvalue of (\mathbf{S}, \mathbf{C}) .

Proof. We decompose the state estimate \mathbf{a} into the concatenation of two vectors $\mathbf{a} = [\mathbf{a}_1; \mathbf{a}_2]$, and rewrite (2.5) as:

$$(2.11) \quad \begin{aligned} \mathbf{E} \mathbf{a}_1 + \mathbf{B} \mathbf{a}_2 &= \lambda \tilde{\mathbf{C}} \mathbf{a}_1 \\ \mathbf{B}^T \mathbf{a}_1 + \mathbf{D} \mathbf{a}_2 &= 0. \end{aligned}$$

Notice that \mathbf{E} is the autocorrelation matrix of the first three dimensions of the observed data, and is invertible by the assumption that \mathbf{S} is full rank. Being full rank, $\tilde{\mathbf{C}}$ is invertible as well. We can solve the above equation as

$$(2.12) \quad \begin{aligned} \mathbf{a}_2 &= -\mathbf{D}^{-1} \mathbf{B}^T \mathbf{a}_1, \\ (\mathbf{E} - \mathbf{B} \mathbf{D}^{-1} \mathbf{B}^T) \mathbf{a}_1 &= \lambda \tilde{\mathbf{C}} \mathbf{a}_1. \end{aligned}$$

In theorem II.3, we will show that iteration:

$$(2.13) \quad \begin{aligned} \mathbf{a}_{1,n+1} &= \eta_n \frac{\mathbf{a}_{1,n}^T \tilde{\mathbf{S}} \mathbf{a}_{1,n}}{\mathbf{a}_{1,n}^T \tilde{\mathbf{C}} \mathbf{a}_{1,n}} \tilde{\mathbf{S}}^{-1} \tilde{\mathbf{C}} \mathbf{a}_{1,n} + (1 - \eta_n) \mathbf{a}_{1,n} \\ &= \eta_n h(\mathbf{a}_{1,n}, \tilde{\mathbf{S}}, \tilde{\mathbf{C}}) + \mathbf{a}_{1,n}, \end{aligned}$$

where $h(\mathbf{x}, \tilde{\mathbf{S}}, \tilde{\mathbf{C}}) \triangleq \frac{\mathbf{x}^T \tilde{\mathbf{S}} \mathbf{x}}{\mathbf{x}^T \tilde{\mathbf{C}} \mathbf{x}} \tilde{\mathbf{S}}^{-1} \tilde{\mathbf{C}} \mathbf{x} - \mathbf{x}$, converges asymptotically to the solution \mathbf{a}_1 ⁷ of (2.12). Therefore, by letting the covariant \mathbf{a}_2 evolve accordingly as:

$$(2.14) \quad \mathbf{a}_{2,n} = -\mathbf{D}^{-1} \mathbf{B}^T \mathbf{a}_{1,n},$$

we have asymptotic convergence to the only stable stationary point $\mathbf{a} = [\mathbf{a}_1; \mathbf{a}_2]$ of (2.12).

In Lemma II.4, we show that the iteration given in (2.13) and (2.14) is identical to the update equation in (2.10). \square

Theorem II.3. *Iteration according to (2.13) converges to the generalized eigen vector that corresponds to the largest eigenvalue of $(\tilde{\mathbf{S}}, \tilde{\mathbf{C}})$, where $\tilde{\mathbf{S}} = [\mathbf{E} - \mathbf{B} \mathbf{D}^{-1} \mathbf{B}^T]$ is the Schur complement of the block \mathbf{D} in \mathbf{S} .*

The second equation in (2.12) simply states that \mathbf{a}_1 is the generalized eigen vector for the pair $([\mathbf{E} - \mathbf{B} \mathbf{D}^{-1} \mathbf{B}^T], \tilde{\mathbf{C}})$. Observe that $[\mathbf{E} - \mathbf{B} \mathbf{D}^{-1} \mathbf{B}^T]$ is exactly the Schur complement $\tilde{\mathbf{S}}$ of the block \mathbf{D} in matrix \mathbf{S} . Since \mathbf{S} is symmetric positive definite, so is $\tilde{\mathbf{S}} = [\mathbf{E} - \mathbf{B} \mathbf{D}^{-1} \mathbf{B}^T]$. As $\tilde{\mathbf{C}}$ is the only block in \mathbf{C} that contributes to the nonzero spectral components and $\sigma(\tilde{\mathbf{C}}) = \{2, -1, -2\}$, the second equation in (2.12) captures all the nontrivial components of the original generalized eigen decomposition problem (2.6). It

⁶We will see that $\sigma(\tilde{\mathbf{S}}, \tilde{\mathbf{C}}) \subset \sigma(\mathbf{S}, \mathbf{C})$ and it is the subset that contains all non-zero generalized eigenvalues. It is often possible to obtain upper bound for κ by utilizing either prior information or proper training. The role of stepsize (gain) η_n determines the trade-off between convergence and convergence rate. The convergence behavior of vanishing gain ($\eta_n \geq 0$, $\sum_n \eta_n = +\infty$, $\sum_n \eta_n^\alpha < \infty$ for some $\alpha > 1$) is commonly studied in the literature [65, 93], but asymptotic constant gain ($\eta_n \geq 0$, $\eta := \lim_{n \rightarrow \infty} \eta_n > 0$) is more desirable in practice. The condition we have imposed includes that of the decreasing gain, but also admits cases with asymptotic constant gain.

⁷The eigen vector paired with the biggest eigenvalue when regarded as a generalized eigen decomposition problem from the second equation. The original constraint $\mathbf{a}^T \mathbf{C} \mathbf{a} = 1$ translates to $\mathbf{a}_1^T \tilde{\mathbf{C}} \mathbf{a}_1 = 1$, and prevents degenerated results.

immediately follows from Sylvester's Law of Inertia [40] that the generalized eigen spectrum of $(\tilde{\mathbf{S}}, \tilde{\mathbf{C}})$ has the form $\lambda_1 > 0 > \lambda_2 \geq \lambda_3$.

Furthermore, since the generalized spectrum has no zero component, the second equation $\tilde{\mathbf{S}}\mathbf{a}_1 = \lambda\tilde{\mathbf{C}}\mathbf{a}_1$ can be rewritten as

$$\lambda^{-1}\tilde{\mathbf{S}}\mathbf{a}_1 = \tilde{\mathbf{C}}\mathbf{a}_1,$$

which indicates that $\sigma(\tilde{\mathbf{C}}, \tilde{\mathbf{S}}) = \{\lambda^{-1}|\lambda \in \sigma(\tilde{\mathbf{S}}, \tilde{\mathbf{C}})\}$ and the generalized eigen vector for $(\tilde{\mathbf{C}}, \tilde{\mathbf{S}})$ coincides with that for $(\tilde{\mathbf{S}}, \tilde{\mathbf{C}})$ (up to possibly some positive scaling factor) with the pairing determined by the element-wise inversion relation of the spectrum. Again, since there is a unique positive eigenvalue (thus the maximum) of $(\tilde{\mathbf{S}}, \tilde{\mathbf{C}})$, the generalized eigen decomposition of $(\tilde{\mathbf{C}}, \tilde{\mathbf{S}})$ has only one positive element as well, whose corresponding generalized eigen vector is of our interest. We will make use of the above observed relationship in the proof of theorem II.3.

Proof. We consider (2.13) in the framework of generic stochastic approximation algorithms [4] $\mathbf{a}_{1,n+1} = \mathbf{a}_{1,n} + \eta_n h(\mathbf{a}_{1,n})$. To apply the corresponding convergence analysis technique, we need to first justify several assumptions. Viewing (2.13) in the classic adaptive form, we know precisely the mathematical conditions relating to the objects, in particular, the gain η_n , the function h and the state $(\tilde{\mathbf{S}}, \tilde{\mathbf{C}})$ (in our case, since we start with the update equation, there is no residual perturbation involved in the evolution, as opposed to the more general form of adaptive updates). In general, the state is represented by a Markov chain controlled by the parameter to be estimated, and it is assumed that for fixed parameter, the state has to be asymptotically stationary, and its limiting behavior regular in the parameter. In our setting, static collective data \mathbf{S} (we disregard the given constant \mathbf{C}) is used, the duplicate of which can be regarded as the simplest form of Markov chain if preferred. Therefore, the stationarity and regularity condition with respect to \mathbf{a}_1 [4] is trivially satisfied. Furthermore, it is easy to check that the update function $h(\mathbf{a}_1, \tilde{\mathbf{S}}, \tilde{\mathbf{C}})$ is continuously differentiable with respect to \mathbf{a}_1 and regular (locally Lipschitz). Therefore, we are allowed to use ODE based approach as a tool to prove asymptotic convergence. Since $\tilde{\mathbf{S}}$ and $\tilde{\mathbf{C}}$ are both constant, we omit them as argument of h for notational brevity.

We link a continuous time ODE to the discrete time algorithm (2.13) to a first order approximation by:

$$\begin{aligned} \tilde{h}(\mathbf{a}_1(s_n)) &= \frac{\partial}{\partial s} \mathbf{a}_1(s)|_{s=s_n}, \quad \mathbf{a}_{1,n} = \mathbf{a}_1(s_n) \\ &\approx \frac{1}{\eta_n} (\mathbf{a}_{1,n+1} - \mathbf{a}_{1,n}) \\ (2.15) \quad &= \frac{\mathbf{a}_1(s)^T \tilde{\mathbf{S}} \mathbf{a}_1(s)}{\mathbf{a}_1(s)^T \tilde{\mathbf{C}} \mathbf{a}_1(s)} \tilde{\mathbf{S}}^{-1} \tilde{\mathbf{C}} \mathbf{a}_1(s) - \mathbf{a}_1(s). \end{aligned}$$

We represent $\mathbf{a}_1(s)$ as a linear combination of the generalized eigen vectors of $(\tilde{\mathbf{S}}, \tilde{\mathbf{C}})$.

$$(2.16) \quad \mathbf{a}_1(s) = \sum_{k=1}^K \theta_k(s) \mathbf{v}_k.$$

Substituting this parameterization in (2.15) yields a coordinate-wise (with respect to the basis $\{\mathbf{v}_k\}_{k=1}^K$) ODE as:

$$\begin{aligned} \frac{\partial}{\partial s} \theta_k(s) &= \frac{\sum_{k=1}^K \theta_k(s)^2}{\sum_{k=1}^K \frac{1}{\lambda_k} \theta_k(s)^2} \frac{1}{\lambda_k} \theta_k(s) - \theta_k(s) \quad \forall \quad k = 1, 2, \dots, K \\ (2.17) \quad &= \left[\frac{\sum_{k=1}^K \theta_k(s)^2}{\sum_{k=1}^K \frac{1}{\lambda_k} \theta_k(s)^2} \frac{1}{\lambda_k} - 1 \right] \theta_k(s), \end{aligned}$$

where $(\lambda_k, \mathbf{v}_k)$ are the k th generalized eigenvalue and eigenvector of $(\tilde{\mathbf{S}}, \tilde{\mathbf{C}})$, and $\theta_k(s)$ is the k th time (iteration) varying projection coefficient indicating the strength of $\mathbf{a}_1(s)$ along direction \mathbf{v}_k .

We define a region $\Omega = \{\theta = (\theta_1, \dots, \theta_K) \mid |\theta_k| \leq \sqrt{\frac{-\lambda_k}{(K-1)\lambda_1}} |\theta_1| \text{ for } k > 1\}$. In our case, $K = 3$ and $0 > \lambda_2 \geq \lambda_3$. It is easy to check that $\frac{\sum_{k=1}^K \theta_k(s)^2}{\sum_{k=1}^K \lambda_k^{-1} \theta_k(s)^2} \geq 0$ for any $\theta \in \Omega$. For $k > 1$, $\lambda_k < 0$, and (2.17) states that

$$\frac{\partial}{\partial s} \theta_k(s) = \alpha_k(s) \theta_k(s)$$

with $\alpha_k(s) < 0$ for all $k > 1$. It follows from Lyapunov stability theorem [46] that $\theta_k(s) = 0$ as $s \rightarrow \infty$ for $k > 1$. On the other hand, since $\lambda_1 > 0 > \lambda_2 \geq \lambda_3$, we have

$$\sum_{k=1}^K \lambda_k^{-1} \theta_k(s)^2 < \lambda_1^{-1} \theta_1(s)^2 < \lambda_1^{-1} \sum_{k=1}^3 \theta_k(s)^2.$$

Subsequently,

$$\frac{\partial}{\partial s} \theta_1(s) = \left[\frac{\sum_{k=1}^K \theta_k(s)^2}{\sum_{k=1}^K \frac{1}{\lambda_k} \theta_k(s)^2} \frac{1}{\lambda_1} - 1 \right] \theta_1(s) > \left[\frac{\sum_{k=1}^K \theta_k(s)^2}{\lambda_1^{-1} \sum_{k=1}^3 \theta_k(s)^2} \frac{1}{\lambda_1} - 1 \right] \theta_1(s).$$

Unlike the other modes ($k \neq 1$) where origin serves as a stable sink, the magnitude of $\theta_1(s)$ increases as its ODE behaves as

$$\frac{\partial}{\partial s} \theta_1(s) = \alpha_1(s) \theta_1(s),$$

with $\alpha_1(s) > 0$. Therefore, $\mathbf{a}_1 \rightarrow \mathbf{v}_1$ asymptotically.

We have thus far proved the asymptotic convergence if the update follows the ODE. In other words, when the step size η is sufficiently small. Big η values correspond to cruder discretizations of the ODE, and may cause discrepancies between the convergence properties of the ODE and the original update equation (2.13). In order to reveal this effect, we need to explore the pole structure of the dynamic system in both continuous and discrete time. We consider the behavior of $\tilde{h}(\mathbf{a}_1(s))$ in the neighborhood of the stationary point $\mathbf{a}_1(s) = \mathbf{v}_k$, where \mathbf{v}_k is the k -th generalized eigen vector of $(\tilde{\mathbf{S}}, \tilde{\mathbf{C}})$. Local linearization results in

$$A_k = \frac{\partial}{\partial \mathbf{a}_1(s)} \tilde{h}(\mathbf{a}_1(s))|_{\mathbf{a}_1 = \mathbf{v}_k}.$$

It is easy to see that the eigenvalues of A_k are given by $\sigma(A_k) = \left\{ \frac{\lambda_k}{\lambda_j} - 1, \quad j \in \{1, 2, \dots, K\} \setminus \{k\} \right\}$. These are the Laplacian domain poles. The transformation in (2.15) defines a map to Z-domain via $z = \eta s + 1$. We list below all possibilities in mapping the pole patterns in Laplacian domain and Z-domain⁸:

1. When λ_k, λ_j are of the same sign, $s = \frac{\lambda_k}{\lambda_j} - 1 \in (-1, \kappa - 1)$.

When $s > 0$ (s-pole in RHP), $z = \eta s + 1 > 1$ falls outside the unit circle, which corresponds to locally unstable pattern. When $s \in (-1, 0)$ (s-pole in LHP), $-1 < z < 0$ lies inside the unit circle, stable. $s = 0$ corresponds to $z = 1$, for critical stability.

2. When λ_k, λ_j are of opposite signs, and $s = \frac{\lambda_k}{\lambda_j} - 1 \in (-2, -1) \subset LHP$.

In this case, $z = \eta s + 1 \in (-1, 0)$ is inside the unit circle, corresponds to a locally stable pattern.

3. When λ_k, λ_j are of opposite signs, and $s = \frac{\lambda_k}{\lambda_j} - 1 < -2 \subset LHP$.

In this case, the s-pole lies inside the LHP, corresponding to local stability. To avoid discrepancy, we want the mapped z-pole to fall inside unit circle. Recall that with $\tilde{\mathbf{S}}$ and $\tilde{\mathbf{C}}$ both normal, $\kappa(\tilde{\mathbf{S}}, \tilde{\mathbf{C}}) = \left| \frac{\lambda_{\max}(\tilde{\mathbf{S}}, \tilde{\mathbf{C}})}{\lambda_{\min}(\tilde{\mathbf{S}}, \tilde{\mathbf{C}})} \right|$ where λ_{\max} and λ_{\min} are maximal and minimal (by moduli) eigenvalues of the generalized eigen decomposition. Therefore, $s \geq -\kappa - 1$, and with η asymptotically bounded above by $\frac{2}{\kappa+1}$, $z \in (-1, 1)$ corresponds to a local stable pattern.

Therefore, the local stability pattern of the stationary points for the ODE and the update equation (2.13) agree. This links the convergence of the ODE to that of the discrete-time equation, and asymptotic convergence of \mathbf{a}_1 is thus proved. \square

Lemma II.4. *Iteration (2.10) is identical to the set of updates given in (2.13) and (2.14).*

⁸This is very different than the commonly seen eigen decomposition of correlation matrices where spectrum is always positive.

Proof. Recall that \tilde{S} is the Schur complement of block D in

$$S = \begin{bmatrix} E & B \\ B^T & D \end{bmatrix}.$$

Performing matrix inversion in block form results in:

$$(2.18) \quad S^{-1} = \begin{bmatrix} \tilde{S}^{-1} & -\tilde{S}^{-1}BD^{-1} \\ -D^{-1}B^T\tilde{S}^{-1} & D^{-1} + D^{-1}B^T\tilde{S}^{-1}BD^{-1} \end{bmatrix}$$

Setting aside the generalized Rayleigh quotient in (2.13), we observe that the major matrix operations involved in the update can be “extracted from (2.18) as:

$$(2.19) \quad \begin{aligned} \tilde{S}^{-1} &= [I \ 0] S^{-1} \begin{bmatrix} I \\ 0 \end{bmatrix} \\ -D^{-1}B^T\tilde{S}^{-1} &= [0 \ I] S^{-1} \begin{bmatrix} I \\ 0 \end{bmatrix}. \end{aligned}$$

Notice that the sparse structure in C induces the following relations:

$$(2.20) \quad \begin{aligned} Ca &= \begin{bmatrix} \tilde{C}a_1 \\ 0 \end{bmatrix} = \begin{bmatrix} I \\ 0 \end{bmatrix} \tilde{C}a_1 \\ a^T Ca &= a_1^T \tilde{C}a_1 \end{aligned}$$

Therefore, we can rewrite the iteration in (2.13) as:

$$(2.21) \quad a_{:,n+1} = \eta_n \lambda_n S^{-1} C a_{:,n} + (1 - \eta_n) a_{:,n},$$

where λ_n denotes the generalized Rayleigh quotient estimated in the n -th iteration.

Furthermore, we rewrite the numerator of the generalized Rayleigh quotient as:

$$(2.22) \quad \begin{aligned} a_1^T \tilde{S} a_1 &= a_1^T [E - BD^{-1}B^T] a_1 \\ &= a_1^T E a_1 - a_1 B D^{-1} D D^{-1} B^T a_1 \\ &= a_1^T E a_1 - a_2 D a_2, \end{aligned}$$

by the dependence of a_2 on a_1 indicated in (2.14).

Putting the above ingredients back into the iteration yields the compact representation (2.10). \square

2.3 Adaptive Ellipse Fitting

For ellipse fitting, the matrix C describes the shape prior (constraint) of the data, and remains constant. When new data becomes available, the system dynamics are reflected via changes in the empirical covariance matrix S . Formulation (2.10) expresses the update of the state estimate in terms of sub-blocks and inverse of S directly (with no hidden or intermediate transformations as in (2.13)). This enables a straightforward derivation for the update equations when S changes upon the arrival of new data samples. Notice that S takes on the form of empirical covariance, so the diagonal sub-blocks E and D are empirical covariance matrices with respect to their own subspaces and are completely decoupled. In practice, the update of the generalized Rayleigh quotient can be performed accordingly. In essence, the only quantity of real concern in updating the state estimate is S^{-1} . To incorporate the time varying property of the system, we can simply extend the previous results with a hyper-level evolving time tag. To express the time varying property of the system, we use $\mathbf{a}_n(i)$, $\mathbf{S}_n(i)$ etc. to denote the various quantities at a given acquisition time $t = i\Delta t$. For a given i , we rewrite (2.10) as follows:

$$(2.23) \quad \begin{aligned} \mathbf{a}_{n+1}(i) &= \eta_n \frac{\mathbf{a}_n(i)^T \mathbf{W}(i) \mathbf{a}_n(i)}{\mathbf{a}_n(i)^T \mathbf{C} \mathbf{a}_n(i)} (\mathbf{S}(i))^{-1} \mathbf{C} \mathbf{a}_n(i) + (1 - \eta_n) \mathbf{a}_n(i), \quad n = 0, 1, \dots, N_i - 1, \\ \mathbf{a}_0(i+1) &= \mathbf{a}_{N_i}(i), \end{aligned}$$

where N_i denotes the number of iterations used to compute the ellipse parameters at a given time i . The challenge is to compute the inverse of $\mathbf{S}(i+1)$ efficiently, and we provide below efficient rank-one updates for \mathbf{S}^{-1} for both the sliding window adaptation and exponential discount adaptation.

2.3.1 Sliding Window Adaptation

In the sliding window adaptation, we use a constant length sliding window to “mask out” the historical data samples except the ones that are close enough to the time instance of interest. We define $\mathbf{S}(i) = \sum_{j=i-L+1}^i \mathbf{z}_j \mathbf{z}_j^T$ with L indicating the constant window size.

When a new sample pair \mathbf{z}_{i+1} becomes available,

$$\begin{aligned} \mathbf{S}(i+1) &= \sum_{j=i+1-L+1}^{i+1} \mathbf{z}_j \mathbf{z}_j^T \\ (2.24) \quad &= \mathbf{S}(i) - \mathbf{z}_{i-L+1} \mathbf{z}_{i-L+1}^T + \mathbf{z}_{i+1} \mathbf{z}_{i+1}^T. \end{aligned}$$

To compute $\mathbf{S}(i+1)^{-1}$ from $\mathbf{S}(i)^{-1}$, we denote $\mathbf{Q}(i) = \mathbf{S}(i) - \mathbf{z}_{i-L+1} \mathbf{z}_{i-L+1}^T$, so that $\mathbf{S}(i+1) = \mathbf{Q}(i) + \mathbf{z}_{i+1} \mathbf{z}_{i+1}^T$. We invoke the Woodbury matrix identity [39] to compute $\mathbf{S}(i+1)^{-1}$ with two step rank one updates:

$$\begin{aligned} \mathbf{Q}(i)^{-1} &= (\mathbf{S}(i) - \mathbf{z}_{i-L+1} \mathbf{z}_{i-L+1}^T)^{-1} \\ &= \mathbf{S}(i)^{-1} - \mathbf{S}^{-1} \mathbf{z}_{i-L+1} (\mathbf{z}_{i-L+1}^T \mathbf{S}(i)^{-1} \mathbf{z}_{i-L+1} - 1)^{-1} \mathbf{z}_{i-L+1}^T \mathbf{S}^{-1} \\ \mathbf{S}(i+1)^{-1} &= (\mathbf{Q}(i) + \mathbf{z}_{i+1} \mathbf{z}_{i+1}^T)^{-1} \\ (2.25) \quad &= \mathbf{Q}(i)^{-1} - \mathbf{Q}^{-1} \mathbf{z}_{i+1} (\mathbf{z}_{i+1}^T \mathbf{Q}(i)^{-1} \mathbf{z}_{i+1} + 1)^{-1} \mathbf{z}_{i+1}^T \mathbf{Q}^{-1}. \end{aligned}$$

For this procedure to be executable, invertibility of $(\mathbf{x}_{i-L+1}^T \mathbf{S}(i)^{-1} \mathbf{x}_{i-L+1} - 1)$ and $(\mathbf{x}_{i+1}^T \mathbf{Q}(i)^{-1} \mathbf{x}_{i+1} + 1)$ are required. The second one is obvious with $\mathbf{Q} > 0$. We prove the first condition in Section 2.6. This pair of properties (2.25) provides a recursion for $\mathbf{S}(i)$. Substituting into (2.10) yields a recursion in the estimation parameters $\mathbf{a}(i)$.

2.3.2 Discounting Adaptation

As an alternative to a fixed-length sliding window, we can use temporal discounting to emphasize the most recent data. In this case, we define $\mathbf{S}(i) = \frac{1-\gamma}{1-\gamma^i} \sum_{j=1}^i \gamma^{i-j} \mathbf{z}_j \mathbf{z}_j^T$, where $\gamma \in (0, 1)$ is a user-selectable discounting parameter. We can easily write $\mathbf{S}(i+1)$ recursively as:

$$(2.26) \quad \frac{1-\gamma^{i+1}}{1-\gamma} \mathbf{S}(i+1) = \gamma \frac{1-\gamma^i}{1-\gamma} \mathbf{S}(i) + \mathbf{z}_{i+1} \mathbf{z}_{i+1}^T$$

Invoking the matrix inversion lemma yields the recursion for $\mathbf{S}(i+1)^{-1}$:

$$(2.27) \quad \mathbf{S}(i+1)^{-1} = \frac{1-\gamma^{i+1}}{\gamma-\gamma^{i+1}} \mathbf{S}(i)^{-1} - \mathbf{S}(i)^{-1} \mathbf{z}_{i+1} \left\{ \frac{1-\gamma}{\gamma-\gamma^{i+1}} + \mathbf{z}_{i+1}^T \mathbf{S}(i)^{-1} \mathbf{z}_{i+1} \right\}^{-1} \mathbf{z}_{i+1}^T \mathbf{S}(i)^{-1}.$$

Substituting this in (2.23) yields an adaptive ellipse fitting algorithm with temporal discounting.

2.4 Application to Tracking Respiratory Mean Drift

Modeling and predicting tumor motion caused by respiration is challenging due to temporal variations in breathing patterns. Treatment approaches such as gating or adaptive bed adjustment/alignment may not require full knowledge of instantaneous position, but might benefit from tracking the general trend of the motion. One simple method for tracking mean tumor position is to apply moving average filters with window sizes corresponding to the breathing periods. Yet respiratory motion is only semi-periodic, so such methods require reliable phase estimation, which is difficult in the presence of noise. As an alternative, we form a

state vector from the respiration signal values at the current instant and at a previous time, and utilize the algorithms discussed in Section 2.2 and Section 2.3 to dynamically fit ellipse models to the training data and extract the mean position according to (2.2). Ellipse eccentricity and orientation potentially capture hysteresis in respiratory motion. We test the proposed method with simulated breathing traces, as well as with real time-displacement (RPM, Varian) signals. Estimation traces are compared with retrospectively generated moving average results to illustrate the performance of the proposed approach.

2.4.1 Application Background

Accurate modeling and prediction of tumor motion caused by breathing is a challenging problem. Previous studies [37, 53, 109, 125] have noted the difficulty of instantaneous position tracking and prediction. Given such limitations in accuracy, and considering the actual dosimetric impact of small motion variations, treatment approaches such as gating or adaptive bed adjustment/alignment may not require instantaneous position, but might benefit from following trends of the motion, in particular mean position drifting and/or abrupt shifts. Current amplitude-based gating systems compare an instantaneous tumor location measurement with a pre-determined gating window threshold and trigger the treatment beam on/off. A potential modification to such systems would incorporate real-time mean drifting information to (1) adjust bed position to compensate for continuous mild drifting; (2) trigger the treatment beam off upon detection of significant drift. Compensating for mean position drifting could increase effective delivered dose given a fixed treatment margin, or alternatively, it could allow the use of smaller margins to achieve the same dose delivery. Previously, other investigators have shown that there could be only limited gains in trying to eliminate breathing movement completely, and laid the groundwork for consideration of the methodology described here. Engelsman demonstrated that the margin needed for cyclic breathing can be represented as a Gaussian with standard deviation of 0.4 times the amplitude of motion [27]. Wolthaus demonstrated a method for efficiently selecting a mean patient representation from a 4D CT data set [132]. Evidence from these and other investigations [52] hint at the possibility that a ‘tracking’ system that estimates variation in position such as the local mean may provide significant benefit by reducing or eliminating non-periodic trends in motions, while reducing demands on temporal response and acceleration of couch or multileaf collimator-based adjustment systems.

Fig. 2.1 illustrates how real-time knowledge of mean drifting helps to reduce gating margin for the same treatment dose delivery (90% in this simulation). In this example, mean compensation reduces the margin by about 70% compared to traditional static gating approach.

Furthermore, mean drift (or home position motion) is more stable, with slower temporal variation than instantaneous position. This makes it more practical to seek an estimator for this lower order quantity. By imposing smoothness, a good estimator should be less susceptible to noise than instantaneous position trackers.

The seemingly intuitive moving average filter is impractical for real-time application due to (1) the absence of “future” observations at the instant of estimation, and (2) the difficulty of estimating instantaneous phase online from noisy observations.

2.4.2 Experiment Setup

We simulated two sets of data so that we could have “ground-truth” for verification purposes. For the first set of simulations, we used noise-free, strictly periodic data with both ideal sinusoid and modified cosine models [69]. In particular, the discrete sinusoidal and modified cosine waveforms were generated respectively with

$$(2.28) \quad x_i^{\sin} = x(i\Delta t) = x_0 + a \sin(\pi i \Delta t / T - \phi)$$

$$(2.29) \quad x_i^{\text{modified cos}} = x(i\Delta t) = x_0 - a \cos^{2n}(\pi i \Delta t / T - \phi),$$




Figure 2.1: Effect of drift compensation for gating system: respiration trajectory (blue solid line); mean position (red dashed line); gating with static window with 90% delivery coverage (magenta dashdot line); mean drift compensated dynamic gate with 90% coverage (black dotted line).

ID V.S. Parameter	1	2	3	4	5	6	7	8	9	10	11	12
Data Characterization ⁹												
STD	2.91	6.47	13.05	2.83	4.86	2.78	4.30	7.61	2.08	7.72	13.04	6.56
P-P	10.93	25.03	48.91	9.02	13.09	11.47	17.77	26.93	13.14	37.44	38.97	32.54
Period (sec)	4.5	4.6	7.2	5.6	4.4	5.4	4.7	9.7	4.7	4.1	3.1	5.2

Table 2.1: RPM Dataset information

where we used the value $n = 2$. In the second test, we generated a semi-periodic sinusoid function with slow frequency drifting by modulating the *local* frequency with random offset components, as follows

$$(2.30) \quad x_i^{\sin} = x_0 + a \sin \left\{ \pi \Delta t \sum_{k=1}^i (1/T + \delta_k) - \phi \right\}$$

$$(2.31) \quad x_i^{\text{modified cos}} = x_0 - a \cos^{2n} \left\{ \pi \Delta t \sum_{k=1}^i (1/T + \delta_k) - \phi \right\},$$

where the δ_k values were randomly distributed via a Gaussian distribution $N(0, \sigma^2)$ with $\sigma \ll 1/T$. In the simulation, we set period $T = 5$ seconds, $\Delta t = 1/30$ corresponding to a sampling frequency of 30Hz, home position $x_0 = 0$, magnitude $a = 5$ cm, and systematic phase offset $\phi = 0$. Fig. 2.2 shows typical simulation traces.

For real clinical data, we used the Real-Time Position Management (RPM, Varian Medical Systems, Palo Alto, CA) system to obtain the trajectories of external fiducials placed on the chests of 12 patients. The displacement-time relationship was recorded at 30Hz and is assumed to be highly correlated with superior-inferior diaphragm motion [125], which is a major source of respiratory motion for tumors in the chest or lung area. We centered and scaled the unit-less RPM data so that their dynamic range corresponds to typical SI motion for chest and lung tumors [107, 109]. We can thereafter consider the units to be on the order of *mm* for typical thorax tumor motion. Characteristic parameters for the RPM data used in our experiment are reported in Table 2.1.

2.4.3 Results

The fitting methods approximate data in the state space (x, y) by ellipses. It is desirable to have the center of such ellipse, which corresponds to the mean estimator, to be robust to missing data, spurious data, and to input data lengths that differ from the ideal period centered at the time instant of estimation. Fig. 2.2 illustrates both data-abundant cases and the cases where only a segment (3 seconds worth) of arc data is available for fitting. The fitted ellipses are overlaid with the observation samples in the augmented state space. The second column in Fig. 2.2 illustrates that ellipses are reasonable approximations for the scattered observations in the state space. The difference between column 3 and 4 in Fig. 2.2 indicates the change of parameters in the presence of scarce and/or non-centered data. Not only does the ellipse fitting method degrade gracefully with partial data, but also the mean position estimated from this approach is reasonably stable. This empirical study illustrates the feasibility of using the proposed method in mean tracking and prediction.

Adaptive Estimation

We first test the case where we use a fixed interval of the most recent data. In the real time estimation and prediction setting, all the input samples into the estimation algorithm precede the time instant of interest. We also want to emphasize that the windowed history is used to help estimate the ellipse parameters; and it need not have integer multiples of the period. We tested the windowed ellipse fitting with 5 second and 7 second local history length, and report the results in Fig. 2.3. Discount adaptation yield very similar results to the windowed fitting, resulting in virtually overlapping real-time mean tracking curves. We omit them from the figures for visualization clarity. We also plot the outputs of two simple moving average filters with fixed window lengths.

We constructed our simulations to have frequency 0.2Hz for deterministic cases or centered around that for the randomly frequency modulated realizations. Therefore, the “ground-truth” mean motion was zero for all the simulations. The clinical RPM data (Patient 1 in both Table 2.1 and Fig. 2.6) also has approximately the same frequency. Since both the simulated and clinical data lack mean drifting, a good estimator for the mean position should yield very stable (flat) output. When we select the training window size to be the “oracle” (ground-truth value unknown to the algorithm) value of $L = 5$ seconds, which coincides with the signal period, outputs are stable from both the moving average operator and the proposed method¹⁰, as illustrated by the left column in Fig. 2.3. On the other hand, it is impossible to guarantee that the history window size will always match the “true” period. We illustrate the effect of a disagreement, where window size $L = 7$ seconds in the right column in Fig. 2.3. The moving average filter exhibits undesirable oscillations, whereas the ellipse fitting method provides comparable results as in the case of perfect period match.

The size of the sliding window and the discount factor must compromise between response speed (tracking efficiency) and robustness (tracking stability). Even though the ellipse fitting method is not too sensitive to the window size, it is helpful to choose window length L and discount factor γ from a short segment of training data. Fig. 2.4 illustrates the effect of various choices of window length parameter L on mean estimation performance with some RPM data and Fig. 2.5 illustrates the effect of the discount factor γ . For RPM data with relatively long period and slow drifting (as in Fig. 2.4(a) and Fig. 2.5(a)), it is desirable to use a larger window size (and correspondingly weaker discounting, large γ) to take advantage of its robustness. On the other hand, for breathing signals that have relatively short periods and rapid shifts in mean position, such as the one illustrated in Fig. 2.4(b) and Fig. 2.5(b), shorter window lengths and small discount factors are preferable for prompt response to mean changes.

To automatically adjust the sliding window length and the discount parameter, we take a short segment of training data at the beginning of each treatment fraction, and apply a subspace projection-based period estimation method [101]. For the signals in Fig. 2.4, the signal in subplot (a) yields a period estimate of 9.7 seconds and the signal in subplot (b) yields a period estimate of 3.1 seconds. Using the estimated period as the sliding window length and choosing the corresponding discount factor appear to be reasonable based on Fig. 2.4 and Fig. 2.5. We apply this scheme to automatically choose the adaptive parameters for all of the 12 RPM datasets and report the results in Fig. 2.6. For baseline comparison, we collect the complete trajectory,

¹⁰A constant offset (as observed in the modified cosine case) has marginal clinical effect, as long as it is consistent.

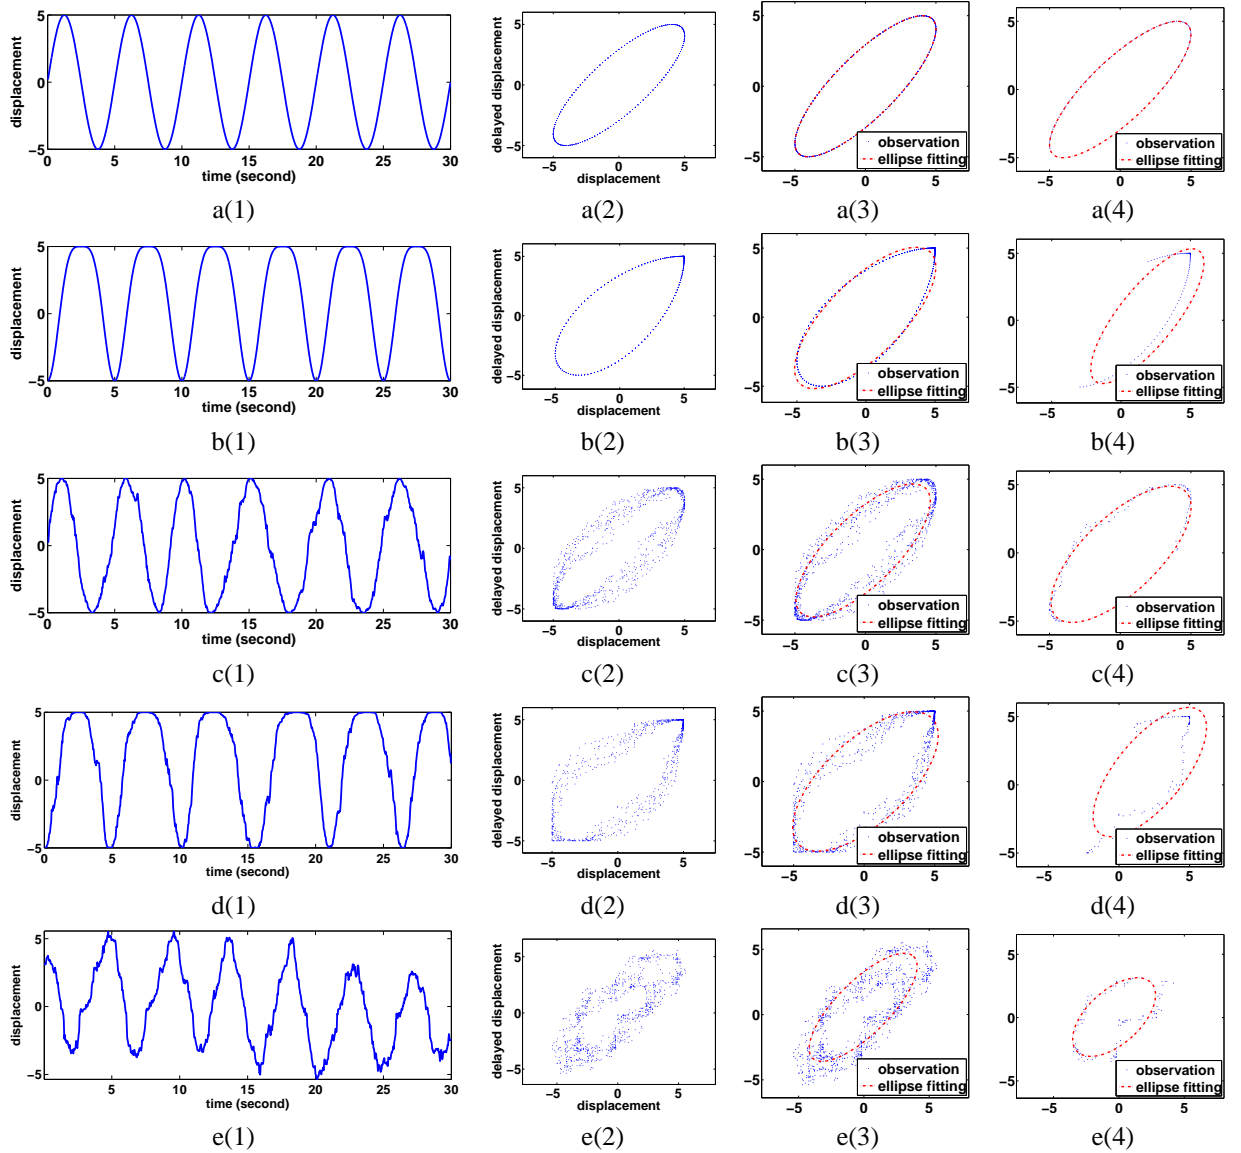


Figure 2.2: Illustration of ellipse fitting performance of the proposed method. Each row corresponds to a different data source: row1 (aX) ideal sinusoid; row2 (bX) ideal modified cosine; (cX) locally modulated (noisy) sinusoid; (dX) locally modulated (noisy) modified cosine; (eX) clinical RPM trace scaled so that P-P \approx 10mm to mimic SI motion. Column-wise: X(1) time-displacement graph; X(2) augmented state space with displacement and its delay ($\tau = 0.5$ seconds); X(3) ellipse fitting (red dashed line) applied to complete dataset; X(4) ellipse fitting (red dashed line) applied to partial dataset.

and apply a moving average filter with the “oracle” window size L to obtain a reasonable “ground-truth”. The deviation of the two adaptive real-time mean position estimator from this “gold standard” (with constant offset compensated) is reported in terms of mean squared error (RMSE) in Table 2.2. Both adaptive methods demonstrate reasonable agreement with the retrospectively obtained “ground-truth”.

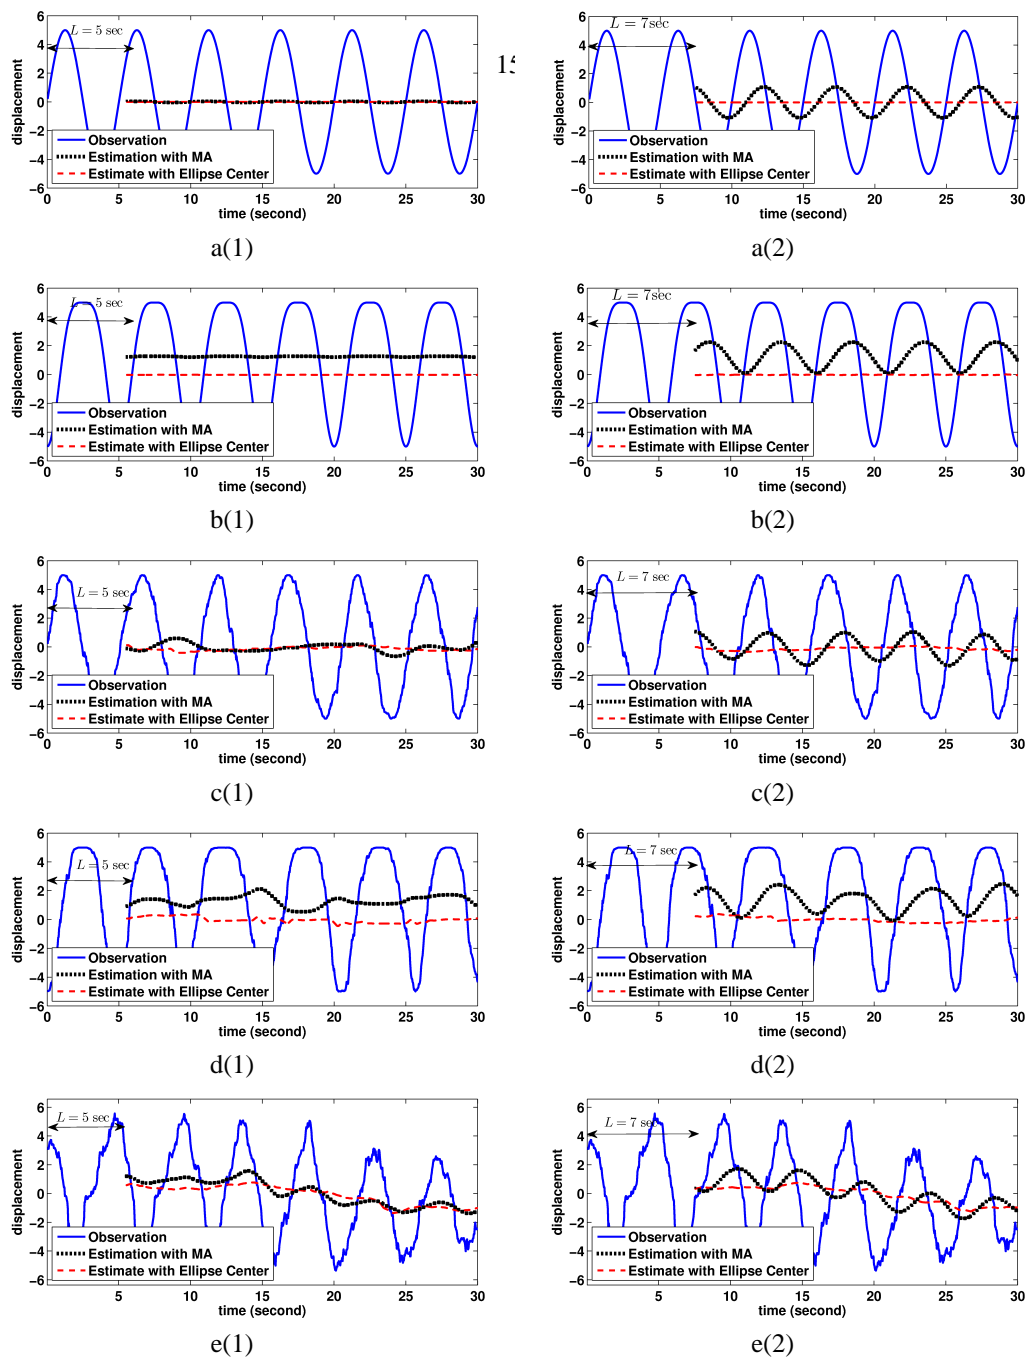


Figure 2.3: Comparison of moving average (MA) and ellipse fitting estimator for mean position tracking: left column X(1): “oracle” history window length: $L = 5$ seconds matches the underlying signal periods exactly; right column X(2): history window length $L = 7$ seconds disagrees with the signal periods. Rows correspond to different data source as in Fig. 2.2. Blue solid line: observation signal; black dotted line: moving average output; red dash dot: output from the ellipse fitting algorithm.

Sensitivity to Sampling Rate

In some cases, it is preferable to obtain observations at a low frequency. This is particularly true when internal tumor motion is extracted from real-time imaging devices that would incur radiation dose. Sparse sampling poses a particular challenge to the conventional mean estimator based on a moving average filter, which is more vulnerable to miss calculation of period length when there are very few samples, resulting in intolerably high variance in mean estimation. We tested the use of sparse real-time observations by sub-




Figure 2.4: Effect of window length L on tracking performance. Solid line: observation; red dashed line: $L = 7$ seconds; green dash-dot line: $L = 5$ seconds; black dotted line: $L = 3$ seconds. (a) RPM with relatively long period; (b) RPM with relatively short period




Figure 2.5: Effect of discount factor γ on tracking performance. Solid line: observation; red dashed line: $\gamma = .99$; green dash-dot line: $\gamma = .97$; black dotted line: $\gamma = .95$. (a) RPM with relatively long period; (b) RPM with relatively short period

ID V.S. Parameter	1	2	3	4	5	6	7	8	9	10	11	12
Period L (sec)	4.5	4.6	7.2	5.6	4.4	5.4	4.7	9.7	4.7	4.1	3.1	5.2
Sliding Win RMSE	0.35	0.77	0.96	0.23	0.68	0.36	0.35	0.90	1.09	1.22	1.21	1.40
Discount factor γ	.978	.979	.986	.982	.978	.982	.979	.990	.979	.976	.968	.981
Discounting RMSE	0.36	0.77	1.08	0.24	0.71	0.35	0.44	1.22	1.54	1.55	2.18	1.39

Table 2.2: Mean Estimation Performance

sampling from the 30Hz signal, applying both windowed and discounted adaptive algorithms to estimate the mean target position, and comparing with the retrospectively generated “true” mean from densely sampled data. Fig. 2.7 illustrates how different observation rates affect overall RMS error across all patients. Both adaptive approaches are quite robust to low sampling rate. In particular, as the windowed adaptation only used historical samples that are within one period, which is normally about 4 – 6 seconds, the observable “break-down” at 1Hz in Fig. 2.7(a) corresponds to estimating the ellipse from 4 – 6 samples only, which is somewhat expected. On the other hand, the discounted adaptation utilizes all previous samples in a weighted fashion, and is naturally less affected by sparse sampling as shown in Fig. 2.7(b).

Setting the Temporal Scale

The size of the window width L and the discount factor γ control the trade-off between response speed and smoothness of the tracking trace in each adaptive algorithm respectively. Even though the ellipse fitting method is robust to missing data (e.g., a partial period), it is still desirable to react more promptly when changes are more frequent (short underlying breathing periods and/or rapid shifts in mean position) and track stably otherwise. For fixed-length sliding window adaptivity, it is preferable to choose a window size

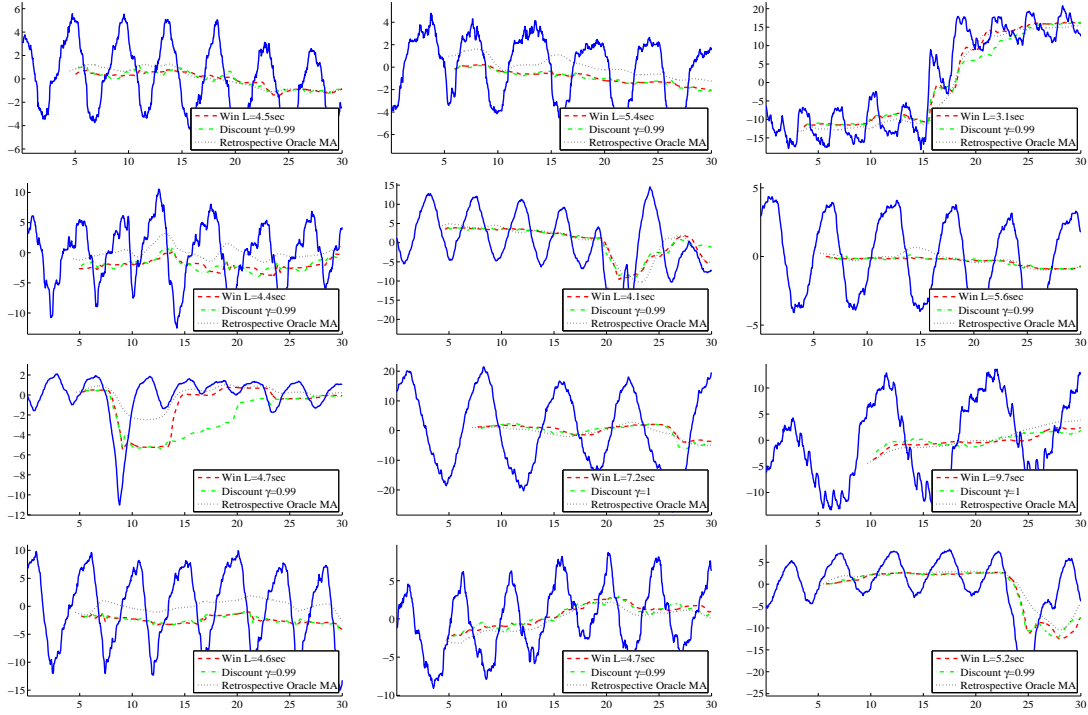


Figure 2.6: Mean tracking for RPM data with window size determined by period estimator. solid line: observed data; black dotted line: retrospective moving average mean estimation with “oracle” period; red dash line: sliding window mean estimator with window size L chosen with period estimation during training phase; green dash-dot line: discounting estimator with discount factor γ chosen such that $\gamma^{L/\Delta t} = 1/20$.




Figure 2.7: Overall RMS error (across all patients) as a function of sampling rate: (a) with windowed ellipse fitting adaptivity; (b) with discounted ellipse fitting adaptivity with discount factor γ chosen such that $\gamma^{L/\Delta t} = 1/20$. Both methods are robust above 2Hz sampling rates.

that roughly matches the “true” period of the signal. Therefore, we use a short segment of training data at the beginning of each treatment fraction, find the closest periodic function to the training segment using a subspace projection method [101] and use the derived period as the fixed window length L .

We could choose the discount factor γ analogously by using *effective memory length*, defined by

$$(2.32) \quad \tilde{L}(i) = \sum_{j=1}^i \gamma^{j-1},$$

because the time unit has a more intuitive physical interpretation. For large i , the efficient memory length is $\tilde{L} = \frac{1}{1-\gamma}$. In other words, we expect the performance of an adaptive mean tracker with discount factor $\gamma = 1 - \frac{1}{L}$ to behave similarly to a sliding window estimator with window size L . In general, however, the discount method should be more stable, but less responsive towards changes than the corresponding sliding window approach with $L = \tilde{L}$ because previous samples are never completely “forgotten”. Thus, we use the period estimated from projection as noted before [101] to find L from 20 seconds of training data, then find γ such that

$$\gamma^{\alpha L} = \beta,$$

where the pair of parameters (α, β) adjusts the decay rate. It has the interpretation that the effect of a given sample decays to β after α periods. We found that in practice $\alpha = 1$ and $\beta = 0.05$ is a reasonable choice and we use these values in later investigations.

2.4.4 Discussion

Although we assumed uniform sampling for simplicity, the proposed approaches easily generalize to nonuniform sampling scenarios, thanks to the robustness of the fitting process. Lower sampling rates should affect the estimation less than the partial datasets tested in Fig. 2.2. Nearly uniform but sparse sampling along the ellipse would increase estimator variance, but should not introduce bias, unlike the partial data case where all the samples are concentrated along an arc segment.

Unlike simple filtering methods, the ellipse fitting method is more objective-oriented: it is specifically designed for estimating time-variant mean of breathing signals. The ellipse model reflects the semi-periodicity of respiratory motion. The fitting process is flexible enough to capture changing trends yet is robust enough to control noisy oscillations. The adaptive algorithms provide efficient updates of the ellipses and allow the users to determine the update rates of the fitting. For adaptive methods using either sliding window or discounting factor, parameter selection involves the trade-off between system response speed and stability. We have suggested one way to adjust the sliding window length L based on the estimated nominal period length, and discussed a connection between the discount factor γ and the “effective memory length” \tilde{L} to provide some guidance about the choice of those parameters. Fast drifting sequences require a more responsive system, and this should be reflected in the corresponding parameter settings. Even though the mean drifting pattern and the respiratory frequency are very often closely correlated, a slow (and regular) breathing pattern may still exhibit abrupt changes, as observed in the upper-left corner of Fig. 2.6. It is possible to resolve this issue by applying the proposed method on a training segment and then investigating the variation pattern of the estimated mean position to further decouple the different causes of the mean position changes. As relation (2.32) only holds asymptotically, and the discounting method is less forgetful than its sliding window counterpart, the discount factor may need to be further reduced to accommodate the more rapidly changing trends.

Our algorithms generalize easily to non-uniformly sampled observations and higher dimensional cases. Commercial solvers for some intermediate steps, such as generalized eigen-decomposition, are available. Clinical experience and physical prior knowledge can help guide choosing either the proper sliding window size or discount factor. In general, both the window size and the discount factor allow real-time adjustment (at the possible cost of more complicated update rules), and could even be tuned intra-fraction, if necessary. The intuitive interpretation of the parameters in terms of window size L , effective memory length \tilde{L} and decaying parameters (α, β) makes the control of those parameters practical.

Practical issues that are worth further investigation include learning of mean position drifting rate, abnormal abrupt change detection, and proper adjustment of the adaptivity pace. This concerns the clinically significant question of “how far we can reliably extrapolate into future based on current observations”. For clinical use, the proposed method needs to be further validated on both external surrogate and internal tumor trajectories, as they may bear different noise properties. Dose effect on various treatment methods and software-hardware cooperation issues should also be studied.

2.5 Generalized Fitting Cost for Robust Estimation

It is often desirable to use a potential function that is robust to the presence of outliers. It is therefore, natural to ask for extension of the squared algebraic distance to robust potential functions such as generalized Gaussian, Huber, Hypergeometric, *etc.* When a more general form of the potential function is to be used, the problem can not be reduced to generalized eigen-decomposition, because the potential is no longer quadratic. General purpose optimization routines need to be studied. Note that this is not a dramatic sacrifice as [90] shows that the generalized eigen decomposition problem can be mapped bijectively into determining whether a matrix $A - \lambda B$ is copositive. Meanwhile, the problem of determining the copositiveness is shown to be NP-complete [82], so is the generalized eigen problem.

Proposition II.5. *The set defined by $\mathbf{a}'\mathbf{C}\mathbf{a} \geq 1$ is a union of two convex sets.*

Proof. Recall the condition for defining the set can be rewritten as $4ac - b^2 \geq 1$ where $\mathbf{a} = [a, b, c, d, e, f]^T$. It is straight forward that $[d, e, f] \in \mathfrak{R}^3$ is a convex subspace. We only need to test the subspace of $[a, b, c]$. Observe that feasible points satisfy $4ac \geq b^2 + 1 > 0$, thus a and c would have the same sign. This naturally split the whole set into two disconnected portions $\mathbf{a}'\mathbf{C}\mathbf{a} \geq 1, a > 0$ and $\mathbf{a}'\mathbf{C}\mathbf{a} \geq 1, a < 0$. Without loss of generality, we concentrate on proving the convexity of the set $U = \{\mathbf{a}'\mathbf{C}\mathbf{a} \geq 1, a > 0\}$ hereafter.

Let (a, b, c) and (x, y, z) are points inside U . Since U is closed, its convexity is implied by ‘‘midpoint convexity’’ [16]. It suffices to test midpoint convexity, which we prove below:

$$(2.33) \quad \begin{aligned} 4\frac{a+x}{2}\frac{c+z}{2} - \left(\frac{b+y}{2}\right)^2 - 1 &= 1/4 \{4ac + 4xz + 4az + 4xc - b^2 - 2by - y^2 - 4\} \\ &\geq 1/2 \{2az + 2xc - by - 1\}. \end{aligned}$$

Notice that

$$(2.34) \quad \begin{aligned} az + xc &\geq 2\sqrt{azxc} = 2\sqrt{(ac)(xz)} \\ &\geq 2\sqrt{\frac{b^2+1}{4}\frac{y^2+1}{4}} \\ &= \frac{1}{2}\sqrt{(b^2+1)(y^2+1)}. \end{aligned}$$

However,

$$(b^2+1)(y^2+1) = b^2y^2 + b^2 + y^2 + 1 \geq b^2y^2 + 2by + 1 = (by+1)^2,$$

thus $\sqrt{(b^2+1)(y^2+1)} \geq by+1$ (This relation holds regardless of the signs of b and y).

Plugging into (2.33) results in

$$4\frac{a+x}{2}\frac{c+z}{2} - \frac{b+y^2}{2} - 1 \geq 0,$$

yielding midpoint convexity of U . □

In principle, we could consider general-purpose optimization techniques to solve a general objective function of the form $\Phi(\mathbf{Z}; \mathbf{a}) = \sum_{i=1}^N \phi(\mathbf{z}_i; \mathbf{a})$, with $\phi = \phi_r(\cdot; \delta) \circ F(\mathbf{z}_i; \mathbf{a}) = \phi_r(\mathbf{a}^T \mathbf{z}_i; \delta)$. ϕ_r may be chosen to be a robust fitting function; it should be positive symmetric about the origin, and equals zero if and only if the argument is zero. Our goal is to solve the general constraint optimization problem:

$$\hat{\mathbf{a}} = \arg \min_{\mathbf{a} \in C} \Phi(\mathbf{a}),$$

where C is a given constraint set.

We will start by considering a simple constrained minimization method called *gradient projection*, which is essentially the gradient descent method with projection on the constraint set at the end of each iteration.

$$(2.35) \quad \mathbf{a}^{(n+1)} = \mathcal{P}_C(\mathbf{a}^{(n)} - \alpha \nabla \Phi(\mathbf{a}^{(n)})),$$

where \mathcal{P}_C denotes the projector on to the convex set C . If stepsize α is chosen appropriately, then for certain families of cost function, the gradient projection method (2.35) converges, as established by Theorem II.6 below. [6, p. 83] analyzes a generalization of the above algorithm.

Theorem II.6. Let C denote a nonempty, convex, closed subset of \mathfrak{R}^{np} . Let $\Phi : \mathfrak{R}^{np} \rightarrow \mathfrak{R}$ be convex and differentiable with gradient $g(\mathbf{x}) \triangleq \nabla\Phi(\mathbf{x})$ satisfying a Lipschitz condition of the form $\|\nabla\Phi(\mathbf{a}) - \nabla\Phi(\tilde{\mathbf{a}})\| \leq L\|\mathbf{a} - \tilde{\mathbf{a}}\|$, $\forall \mathbf{a}, \tilde{\mathbf{a}} \in C$. Suppose the set of minimizers $X^* = \{\mathbf{a}^* \in C : \Phi(\mathbf{a}^*) \leq \Phi(\mathbf{a}), \forall \mathbf{a} \in C\}$ is nonempty. If $0 < \alpha < 2/L$, then the gradient projection algorithm (2.35) converges to some $\mathbf{a}^* \in X^*$.

In our case, the convex half cone is defined by $\mathbf{a}'C\mathbf{a} \geq 1$. Given an initial point $\mathbf{a}_0 = \mathbf{a}^{(n)} - \alpha\nabla\Phi(\mathbf{a}^{(n)})$. If $\mathbf{a}_0 \notin C$, then the projection $\mathbf{a} = \mathcal{P}_C(\mathbf{a}_0)$ has to satisfy:

$$(2.36) \quad \begin{aligned} \mathbf{a}_0 - \mathbf{a} \parallel C\mathbf{a} &\Rightarrow \mathbf{a}_0 - \mathbf{a} = \gamma C\mathbf{a} \\ \mathbf{a}'C\mathbf{a} &= 1, \end{aligned}$$

where $C = \begin{bmatrix} 0 & 0 & 2 \\ 0 & -1 & 0 \\ 2 & 0 & 0 \end{bmatrix}$ and its spectra $\rho(C) = \{-2, -1, 2\}$. There are two scenarios to be considered:

1. When $\gamma \in -\rho(C)^{-1} = \{1/2, 1, -1/2\}$, the linear operator $I + \gamma C$ has a nontrivial null-space of dimension one. The solution \mathbf{a} in that case is obtained as the intersection between a two-dimensional plane (co-dimension one equals the multiplicity of the corresponding eigenvalue) and the cone shape.
2. Let Assuming $\mathbf{a}_\gamma \triangleq [I + \gamma C]^{-1}\mathbf{a}_0$ for $\gamma \notin -\rho(C)^{-1}$. We need to find γ such that

$$\mathbf{a}'_\gamma C \mathbf{a}_\gamma = 1.$$

Let the eigen decomposition of C be $C = V\Lambda V'$ and the above equality can be rewritten as:

$$\mathbf{a}'_0 V [I + \gamma\Lambda]^{-1} \Lambda [I + \gamma\Lambda]^{-1} V' \mathbf{a}_0 = 1.$$

Noting that V is the ‘‘natural’’ coordinate system determined by C and $V'\mathbf{a}_0$ is the representation of \mathbf{a}_0 in that coordinate, we rewrite the problem in the general form of $\mathbf{y}'_0 \tilde{\Lambda}(\gamma) \mathbf{y}_0 = 1$ where we can identify $\mathbf{y}_0 = V\mathbf{a}_0$ and $\tilde{\Lambda}(\gamma)$ is the diagonal scaling $[I + \gamma\Lambda]^{-1} \Lambda [I + \gamma\Lambda]$.

Without further manipulation, this corresponds to finding the roots to a 6th-order polynomial.

We usually desire the resulting projection \mathbf{a} to be close to the starting point \mathbf{a}_0 . It follows from $\mathbf{a}_0 - \mathbf{a} = \gamma C\mathbf{a}$ that

$$\|\mathbf{a}_0 - \mathbf{a}\|_2^2 = \gamma^2 \mathbf{a}' C' C \mathbf{a},$$

and it is straight forward to pick out the \mathbf{a} that is closest to \mathbf{a}_0 in L_2 sense among several (up to 6) candidates.

As stated in Theorem II.6, the stepsize α in (2.35) needs to be upper-bounded by $2/L$ to ensure convergence, where L is the Lipschitz constant for the gradient $g(\mathbf{x})$. In what follows, we will use Huber function as an example to illustrate the procedure of obtaining an upper bound for the Lipschitz constant L . We let $\phi(\mathbf{z}_i; \mathbf{a}) = \phi_h(\cdot; \delta) \circ F(\mathbf{z}_i; \mathbf{a}) = \phi_h(\mathbf{a}^T \mathbf{z}_i; \delta)$ as the fitting measure, where ϕ_h is the huber function given by:

$$(2.37) \quad \phi_h(t; \delta) = \begin{cases} \frac{1}{2}t^2 & |t| < \delta; \\ \delta|t| - \frac{1}{2}\delta^2 & |t| \geq \delta. \end{cases}$$

Notice that this is reasonable fitting measure as $\phi_h \geq 0$ and the equality holds if and only if \mathbf{z}_i falls on the ellipse parameterized by \mathbf{a} .

The column gradient $g(\mathcal{Z}; \mathbf{a}, \delta)$ is given by:

$$(2.38) \quad \begin{aligned} g(\mathcal{Z}; \mathbf{a}, \delta) &= \frac{\partial}{\partial \mathbf{a}} \Phi \\ &= \sum_i \phi'_h(\mathbf{a}^T \mathbf{z}_i; \delta) \mathbf{z}_i, \end{aligned}$$

where $\mathcal{Z} \triangleq [\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_n]$ is the collection of all data points.

Our goal is to find the Lipchitz constant L such that

$$\|g(\mathcal{Z}; \mathbf{a}, \delta) - g(\mathcal{Z}; \tilde{\mathbf{a}}, \delta)\|_2 \leq L \|\mathbf{a} - \tilde{\mathbf{a}}\|_2,$$

for all \mathbf{a} and $\tilde{\mathbf{a}}$ on the feasible set \mathcal{C} .

$$(2.39) \quad \|g(\mathcal{Z}; \mathbf{a}, \delta) - g(\mathcal{Z}; \tilde{\mathbf{a}}, \delta)\|_2 = \left\| \sum_i [\phi'_h(\mathbf{a}^T \mathbf{z}_i; \delta) - \phi'_h(\tilde{\mathbf{a}}^T \mathbf{z}_i; \delta)] \mathbf{z}_i \right\|_2$$

Note that the derivative of huber function is nonnegative with its slope bounded above by unity:

$$\phi'_h(t; \delta) = \begin{cases} t & |t| < \delta \\ \delta & t > \delta \\ -\delta & t < -\delta. \end{cases}$$

Thus $|\phi'_h(t; \delta) - \phi'_h(\tilde{t}; \delta)| \leq |t - \tilde{t}|$. Substituting $t = \mathbf{a}^T \mathbf{z}_i$ and $\tilde{t} = \tilde{\mathbf{a}}^T \mathbf{z}_i$ yields:

$$(2.40) \quad |\phi'_h(\mathbf{a}^T \mathbf{z}_i; \delta) - \phi'_h(\tilde{\mathbf{a}}^T \mathbf{z}_i; \delta)| \leq |\mathbf{a}^T \mathbf{z}_i - \tilde{\mathbf{a}}^T \mathbf{z}_i|$$

Let $\mathbf{c} \triangleq [\phi'_h(\mathbf{a}^T \mathbf{z}_1), \phi'_h(\mathbf{a}^T \mathbf{z}_2), \dots, \phi'_h(\mathbf{a}^T \mathbf{z}_n)]$ and $\tilde{\mathbf{c}} \triangleq [\phi'_h(\tilde{\mathbf{a}}^T \mathbf{z}_1), \phi'_h(\tilde{\mathbf{a}}^T \mathbf{z}_2), \dots, \phi'_h(\tilde{\mathbf{a}}^T \mathbf{z}_n)]$, then (2.39) can be rewritten as:

$$(2.41) \quad \begin{aligned} \|g(\mathcal{Z}; \mathbf{a}, \delta) - g(\mathcal{Z}; \tilde{\mathbf{a}}, \delta)\|_2 &= \sqrt{(\mathbf{c} - \tilde{\mathbf{c}})^T \mathcal{Z}^T \mathcal{Z} (\mathbf{c} - \tilde{\mathbf{c}})} \\ &= \sqrt{\rho(\mathcal{Z}^T \mathcal{Z})} \|\mathbf{c} - \tilde{\mathbf{c}}\|_2. \end{aligned}$$

Substituting the elementwise bound (2.40) into $\|\mathbf{c} - \tilde{\mathbf{c}}\|_2$ yields:

$$(2.42) \quad \begin{aligned} \|\mathbf{c} - \tilde{\mathbf{c}}\|_2 &= \sqrt{\sum_i (\mathbf{c}_i - \tilde{\mathbf{c}}_i)^2} \\ &= \sqrt{(\mathbf{a} - \tilde{\mathbf{a}})^T \mathcal{Z}^T \mathcal{Z} (\mathbf{a} - \tilde{\mathbf{a}})} \\ &= \sqrt{\rho(\mathcal{Z}^T \mathcal{Z})} \|\mathbf{a} - \tilde{\mathbf{a}}\|_2. \end{aligned}$$

Substituting (2.42) into (2.41) yields:

$$\|g(\mathcal{Z}; \mathbf{a}, \delta) - g(\mathcal{Z}; \tilde{\mathbf{a}}, \delta)\|_2 \leq L \|\mathbf{a} - \tilde{\mathbf{a}}\|_2,$$

where $L = \rho(\mathcal{Z}^T \mathcal{Z})$. A loose upperbound for $\rho(\mathcal{Z}^T \mathcal{Z})$ is $\text{trace}\{\mathcal{Z}^T \mathcal{Z}\} = \sum_{i=1}^n \mathbf{z}_i^T \mathbf{z}_i$ as $\mathcal{Z}^T \mathcal{Z}$ is nonnegative definite. This is a reasonable result considering the ‘‘strength’’ of $\nabla \Phi(\mathbf{a})$ incorporates the collective effect of all the data points, and in extreme cases when \mathbf{z}_i are ‘‘aligned’’ would scale as the number of sample points, and $\rho(\mathcal{Z}^T \mathcal{Z}) \approx \text{trace}\{\mathcal{Z}^T \mathcal{Z}\}$. This loose upperbound may be convenient to use when the data are dynamically updated, as it does not require repetitively performing eigen decomposition.

We remark on the structure of the generalized fitting with robust cost here.

- It is reasonable to assume that the general robust fitting objective takes on the form of $\Phi(\mathcal{Z}; \mathbf{a}) = \sum_{i=1}^N \phi_r(\mathbf{z}_i^T \mathbf{a}; \delta)$ where the ϕ_r is some robust cost function and δ controls its shape and scale. Moreover, the symmetry of ϕ_r about the origin in its argument ($\mathbf{a}^T \mathbf{z}_i$) translates naturally to the overall objective Φ . This symmetry has an important geometric implication. Recall that the feasible set of parameters is the union of two convex cones distinguished by the sign of the first element of \mathbf{a} ; together with the above analysis about the geometry of the objective function, we conclude that the graph of the

objective is symmetric about the origin¹¹. Technicality aside, this clears the last bit of reservation one may have towards the applicability of the gradient projection method. Given an initial \mathbf{a}_0 , one can arbitrarily pick a cone (the natural choice would be the one whose first coordinate has the same sign as the element of \mathbf{a}_0), and then perform gradient projection on the chosen cone. Based on whichever minimizer \mathbf{a}^* we obtained, a simple reflection results in $-\mathbf{a}^*$: another minimizer with the same objective function value that resides on the other convex cone.

- Given the iterative structure of the gradient projection method, extension to adaptivity is natural. For static data case, we alternate between two operations: projection onto the feasible set \mathcal{C} and gradient descent in the direction of $\nabla\Phi$. Notice that the inclusion of a new data point only perturbs $\nabla\Phi$ by $\nabla\phi(\mathbf{z}_{\text{new}})$, but does not affect the projection operator $\mathcal{P}_{\mathcal{C}}$. The result from previous iterations should be regarded as an initialization to the updated cost function. More specifically, the adaptive version for the gradient projection algorithm is given by:

$$(2.43) \quad \begin{aligned} \mathbf{a}_{n+1}(i) &= \mathcal{P}_{\mathcal{C}}(\mathbf{a}_n(i) - \alpha \nabla\Phi(\mathbf{a}_n(i))), \quad n = 0, 1, \dots, N_i - 1, \\ \mathbf{a}_0(i+1) &= \mathbf{a}_{N_i}(i), \end{aligned}$$

where i indexes the data samples and n indexes the inner iterations.

- We used algebraic distance to implicitly represent the ellipse to obtain a convex formulation and a simple solution. It is possible to modify the algebraic fit of the ellipses to drive it closer to the geometric solution, which is the minimizer of geometric distance. The idea is to weigh the samples based on a given estimation, leading to a simple iterative mechanism. [11] provides the following interpretation. The algebraic solution \mathbf{a} is the least squares minimizer of Φ . Let $h(\mathbf{z})$ be the geometric distance from the center of the fitted ellipse O_e to \mathbf{z}

$$h(\mathbf{z}) = \|\mathbf{z} - O_e\|_2,$$

and determine \mathbf{p}_i by intersecting the ray from the ellipse's center to \mathbf{z}_i and the ellipse. Then

$$(2.44) \quad \begin{aligned} \phi(\mathbf{z}_i) &= \kappa \left(\frac{h(\mathbf{z}_i)^2}{h(\mathbf{p}_i)^2} - 1 \right) \\ &\approx 2\kappa \frac{h(\mathbf{z}_i) - h(\mathbf{p}_i)}{h(\mathbf{p}_i)}, \quad \text{if } \mathbf{z}_i \approx \mathbf{p}_i, \end{aligned}$$

for some constant κ . Thus one may interpret the algebraic solution as a fit to the ellipse with respect to the relative distances, where distant points are weighted less than near points. This explains why the algebraic solution tends to neglect points far from the center. This is in fact, a desirable trait in many applications where non-eccentric ellipses are favored.

If one prefers to minimize the absolute distance, then data \mathbf{z}_i can be weighted with $h(\mathbf{p}_i)$ for a given estimated ellipse. The resulting estimated ellipse may then be used to update the weight, thus iteratively solving the weighted least squares problem. Naturally, if one is interested in solving the fitting in least squared sense for the geometric distance, then the weight for data \mathbf{z}_i may be set to $d(\mathbf{z}_i)/\phi(\mathbf{z}_i)$ where $d(\mathbf{z}_i)$ is the geometric distance of \mathbf{z}_i from the currently estimated ellipse. The advantage of such iterative weighted least squares scheme is that there is no need to compute Jacobian or Hessian as in the case of a direct nonlinear optimization with respect to geometric distance. The drawback is that its solution generally differs from the minimizer of the geometric distance.

¹¹This is a bit sloppy, since the graph lies in $\mathfrak{R}^{\dim(\mathbf{a})+1}$ dimension, so it should be ideally stated as symmetric with respect to $(\mathbf{0}, \Phi(\mathcal{Z}; \mathbf{0}))$. However, $\mathbf{0}$ is not a feasible point in the domain, so Φ is not defined on that point, which makes this statement illegitimate. A quick remedy would be to redefine Φ as:

$$\tilde{\Phi}(\mathbf{a}) = \begin{cases} \Phi(\mathbf{a}) & \mathbf{a} \in \mathcal{C}; \\ +\infty & \mathbf{a} \notin \mathcal{C}. \end{cases}$$

and the graph of $\tilde{\Phi}$ is symmetric with respect to $(\mathbf{0}, \infty)$.

To harvest the benefit of using robust objectives, we need to choose the parameters for those functions properly. For instance, the threshold parameter for Huber function determines the transition from L_2 cost to L_1 penalty. Without assuming prior knowledge about the mixing probability of normal samples against noise outliers, we determine the parameter by considering the classification sub-problem. In particular, we use the Ostu's method, aiming to best distinguish between the normal and noisy samples. More specifically, after the n th iteration, we examine the distribution of the fitting error and find the value $\delta^{(n+1)}$ that minimizes the within-class variance of the fitting error from the previous iteration $\{e_i = \phi_h(\mathbf{a}^T \mathbf{z}_i; \delta^{(n)})\}$. Mathematically, the threshold parameter at the n th iteration δ^n is selected as the minimizer to

$$\sigma_w^2(\delta) \triangleq w_1(\delta)\sigma_1^2(\delta) + w_2(\delta)\sigma_2^2(\delta),$$

where $w_1 = P(e_i \leq \delta)$ is the probability of normal samples (errors smaller than threshold) under the assumption of threshold δ , $w_2 = 1 - w_1$; σ_i^2 is the empirical variance of each class.

We illustrate the robustness of the proposed method with a simulated example. Noisy samples were uniformly distributed inside the computation region with a roughly elliptical object (the bone contour from a head CT slice). Figure 2.5 illustrates the initial fitting with quadratic minimization and the evolution of the estimated ellipse with the gradient projection method.




Figure 2.8: Evolution of robust fitted ellipse with the gradient projection method: blue dots: observed sample locations; green line: fitted ellipse.

2.6 Appendix

- Proof for the statement about stationary points of the energy ratio function in (2.7).

Claim II.7. *The generalized eigen vectors of (\mathbf{A}, \mathbf{B}) correspond to the stationary points of the energy ratio function*

$$J(\mathbf{u}) = \frac{\mathbf{u}^T \mathbf{A} \mathbf{u}}{\mathbf{u}^T \mathbf{B} \mathbf{u}}.$$

Proof. We compute the stationary point of the energy ratio function $J(\mathbf{u})$, i.e., we set:

$$(2.45) \quad \frac{\partial}{\partial \mathbf{u}} J(\mathbf{u}) = \mathbf{0}^T.$$

The derivative on the LHS of (2.45) turns out to be:

$$\frac{\partial}{\partial \mathbf{u}} J(\mathbf{u}) = \frac{1}{\mathbf{u}^T \mathbf{B} \mathbf{u}} [\mathbf{A} \mathbf{u} \mathbf{u}^T \mathbf{B} \mathbf{u} - \mathbf{B} \mathbf{u} \mathbf{u}^T \mathbf{A} \mathbf{u}].$$

With $\mathbf{A} > 0$, $\mathbf{B} > 0$ as previously assumed, so that $\mathbf{u}^T \mathbf{A} \mathbf{u}$ and $\mathbf{u}^T \mathbf{B} \mathbf{u}$ are simply positive scalars, setting the above expression to zero is equivalent to requiring

$$\mathbf{u}^T \mathbf{B} \mathbf{u} \mathbf{A} \mathbf{u} = \mathbf{u}^T \mathbf{A} \mathbf{u} \mathbf{B} \mathbf{u}.$$

This is exactly the condition for generalized eigen decomposition:

$$\mathbf{A}\mathbf{u} = \frac{\mathbf{u}^T \mathbf{A}\mathbf{u}}{\mathbf{u}^T \mathbf{B}\mathbf{u}} \mathbf{B}\mathbf{u}.$$

Therefore, the generalized eigen vectors are the stationary points for the energy ratio function $J(\mathbf{u})$. Moreover, the evaluated functional values provide the corresponding generalized eigenvalues. This result can be considered as a generalization of the Rayleigh-Ritz theorem. \square

- Derivation for (2.17)

We make use of the relationship between the generalized eigen decomposition $(\tilde{\mathbf{S}}, \tilde{\mathbf{C}})$ and $(\tilde{\mathbf{C}}, \tilde{\mathbf{S}})$. Up to a constant gain, the set of generalized eigen vectors of the two problems coincide, pairing with element-wise inverted spectrum. Since we aim to prove the convergence of the coefficients of the eigen vectors either to zero or really large, the constant scaling can be neglected for the sake of argument clarity. We use the alternative setup of $(\tilde{\mathbf{C}}, \tilde{\mathbf{S}})$ in deriving (2.17) to take advantage of the assumed positive definiteness of $\tilde{\mathbf{S}}$. The generalized eigen decomposition of (\mathbf{A}, \mathbf{B}) with \mathbf{B} being symmetric positive definite indicates the existence of a generalized eigen matrix \mathbf{V} (with columns being the generalized eigen vectors) that can simultaneously diagonalize \mathbf{A} and \mathbf{B} :

$$\begin{aligned} \mathbf{A}\mathbf{V} &= \Lambda \mathbf{B}\mathbf{V}; \\ \mathbf{V}^T \mathbf{B}\mathbf{V} &= \mathbf{I}; \\ \mathbf{V}^T \mathbf{A}\mathbf{V} &= \Lambda. \end{aligned} \tag{2.46}$$

Where Λ is a diagonal matrix whose diagonal elements are the corresponding generalized eigenvalues of (\mathbf{A}, \mathbf{B}) . Indeed, the use of two-stage conventional eigen decomposition to compute generalized eigen decomposition reflects exactly this property. We apply this to $(\tilde{\mathbf{C}}, \tilde{\mathbf{S}})$ and call their eigen matrix \mathbf{V} . Again, \mathbf{V} is also an eigen matrix for $(\tilde{\mathbf{S}}, \tilde{\mathbf{C}})$.

The linear representation in (2.16) can be rewritten as:

$$\mathbf{a}_1(s) = \mathbf{V}\theta(s),$$

where $\theta(s) = [\theta_1(s), \theta_2(s), \dots]^T$.

Substituting in the relevant terms in (2.15), we have

$$\begin{aligned} \mathbf{a}_1(s)^T \tilde{\mathbf{S}} \mathbf{a}_1(s) &= \theta(s)^T \mathbf{V}^T \mathbf{S} \mathbf{V} \theta(s) = \sum_{k=1}^K \theta_k(s)^2 \\ \mathbf{a}_1(s)^T \tilde{\mathbf{C}} \mathbf{a}_1(s) &= \theta(s)^T \mathbf{V}^T \mathbf{C} \mathbf{V} \theta(s) = \theta(s)^T \Lambda^{-1} \theta(s) = \sum_{k=1}^K \lambda_k^{-1} \theta_k(s)^2 \\ \tilde{\mathbf{S}}^{-1} \tilde{\mathbf{C}} \mathbf{a}_1(s) &= \sum_{k=1}^K \lambda_k^{-1} \theta_k(s) \end{aligned} \tag{2.47}$$

The second and third lines in the above derivation also make use of the element-wise inversion relation between the spectra of $(\tilde{\mathbf{C}}, \tilde{\mathbf{S}})$ and $(\tilde{\mathbf{S}}, \tilde{\mathbf{C}})$.

Representing both the LHS and RHS of (2.15) with respect to the basis $\{v_k\}_{k=1,2,\dots,K}$, and we have coordinate-wise equation (2.17).

- Proof of invertibility in (2.25).

Proposition II.8. Let $\mathbf{Q}(t) \triangleq \mathbf{S}(t) - \mathbf{x}_{t-L+1} \mathbf{x}_{t-L+1}^T$, then $(\mathbf{x}_{t-L+1}^T \mathbf{S}(t)^{-1} \mathbf{x}_{t-L+1} - 1)$ is invertible.

Proof. Proving the invertibility of a scalar quantity is the same as showing that it is none-zero. We rewrite the relation between \mathbf{S} and \mathbf{Q} as: $\mathbf{S}(t) = \mathbf{Q}(t) + \mathbf{x}_{t-L+1} \mathbf{x}_{t-L+1}^T$. Invoking the Woodbury matrix inversion lemma, we get:

$$\mathbf{S}(t)^{-1} = \mathbf{Q}(t)^{-1} - \mathbf{Q}(t)^{-1} \mathbf{x}_{t-L+1} (\mathbf{x}_{t-L+1}^T \mathbf{Q}(t)^{-1} \mathbf{x}_{t-L+1} + 1)^{-1} \mathbf{x}_{t-L+1}^T \mathbf{Q}(t)^{-1}. \tag{2.48}$$

Plugging (2.48) into $\mathbf{x}_{t-L+1}^T \mathbf{S}^{-1} \mathbf{x}_{t-L+1} - 1$ yields:

$$\begin{aligned} & \mathbf{x}_{t-L+1}^T \mathbf{S}^{-1} \mathbf{x}_{t-L+1} - 1 \\ &= \mathbf{x}_{t-L+1}^T [\mathbf{Q}(t)^{-1} - \mathbf{Q}(t)^{-1} \mathbf{x}_{t-L+1} (\mathbf{x}_{t-L+1}^T \mathbf{Q}(t)^{-1} \mathbf{x}_{t-L+1} + 1)^{-1} \mathbf{x}_{t-L+1}^T \mathbf{Q}(t)^{-1}] \mathbf{x}_{t-L+1} - 1 \\ &= \mathbf{x}_{t-L+1}^T \mathbf{Q}(t)^{-1} \mathbf{x}_{t-L+1} - \mathbf{x}_{t-L+1}^T \mathbf{Q}(t)^{-1} \mathbf{x}_{t-L+1} \frac{1}{\mathbf{x}_{t-L+1}^T \mathbf{Q}(t)^{-1} \mathbf{x}_{t-L+1} + 1} \mathbf{x}_{t-L+1}^T \mathbf{Q}(t)^{-1} \mathbf{x}_{t-L+1} - 1. \end{aligned}$$

Let $p \triangleq \mathbf{x}_{t-L+1}^T \mathbf{Q}(t)^{-1} \mathbf{x}_{t-L+1}$, then $p > 0$ as $\mathbf{Q}(t) > 0$. We rewrite the expression in (2.49) in terms of p and get

$$p - p \times \frac{1}{p+1} p - 1 = \frac{1}{p+1} [p(p+1) - p^2 - (p+1)] < 0.$$

This result states that $\mathbf{x}_{t-L+1}^T \mathbf{S}(t)^{-1} \mathbf{x}_{t-L+1} - 1 < 0$, thus invertible. In fact, the negativity of this term is not accidental, but a natural consequence of the consistent relation stated below. When $\mathbf{x}_{t-L+1}^T \mathbf{S}(t)^{-1} \mathbf{x}_{t-L+1} - 1 \neq 0$, we could apply the matrix inversion lemma in two different ways (expressing \mathbf{S}^{-1} with \mathbf{Q}^{-1} , and the other way around), and obtain:

$$\begin{aligned} \mathbf{Q}(t)^{-1} - \mathbf{S}(t)^{-1} &= \mathbf{Q}(t)^{-1} \mathbf{x}_{t-L+1} (\mathbf{x}_{t-L+1}^T \mathbf{Q}(t)^{-1} \mathbf{x}_{t-L+1} + 1)^{-1} \mathbf{x}_{t-L+1}^T \mathbf{Q}(t)^{-1} \\ (2.49) \quad &= -\mathbf{S}(t)^{-1} \mathbf{x}_{t-L+1} (\mathbf{x}_{t-L+1}^T \mathbf{S}(t)^{-1} \mathbf{x}_{t-L+1} - 1)^{-1} \mathbf{x}_{t-L+1}^T \mathbf{S}(t)^{-1}. \end{aligned}$$

Because $\mathbf{Q} > 0$, the RHS of line 1 in (2.49) is positive definite. With the minus sign in the front and its quadratic form, line 2 in (2.49) indicates that $\mathbf{x}_{t-L+1}^T \mathbf{S}(t)^{-1} \mathbf{x}_{t-L+1} - 1 < 0$ if it is ever nonzero (otherwise (2.49) cannot be established in the first place).

□

CHAPTER III

Regularized Nonrigid Image Registration: Local Elasticity Penalty, Discontinuity Preserving Flow and Performance Analysis ¹

In medical applications, spatial alignment is often required to properly integrate useful information from separate images [74, 138]. *Registration* is the procedure of retrieving the transformation that maps from the target image's coordinate space to the source image's coordinates.

Registration algorithms can be classified according to the family of transformations. Rigid/affine (global) registration algorithms have only a few degrees of freedom, while nonrigid registration algorithms often have a very high dimensional space of feasible transformations. Usually, rigid registration methods provide satisfactory matching results for individual bone structures, but are in general not descriptive enough for elastic tissues that undergo more free-form deformations.

Nonrigid registration problems can be highly under-determined when transformations of high dimensionality are used, resulting in ill-conditionedness, instability of solutions as well as multiple local optima. Regularizations are usually introduced to alleviate these issues and to effectively incorporate prior physical knowledge into the problem formulation. Regularized nonrigid image registration algorithms usually involve minimizing a cost function, consisting of a dissimilarity measure and a penalty term that discourages undesirable transformations. Conventional regularization methods usually treat the region of interest (ROI) as one single deformable body and homogeneously penalize deviations from smoothness or incompressibility properties of the deformation field [55, 56, 96].

However, homogeneous smoothness regularization has its limitations. In particular, ignoring the elasticity differences between tissue types can cause non-physical results, such as bone warping. Furthermore, isotropic smoothing throughout the ROI blurs motion edges, resulting in artifacts across motion interfaces where sliding effects occur, which are commonly observed between diaphragm and rib cage during respiration.

To address the tissue-dependent elasticity issue, segmentation-based methods were proposed to treat each segmented region of an image independently [51, 64, 128]. These methods rely heavily on precise segmentation and may incur boundary issues with overlapping/vacuum region in the deformed image. Empirical spatial filtering was also used to "correct" the deformation field as a post-processing step [112]. Unfortunately, its deviation from an optimization setup complicates convergence assessment. To study discontinuities in deformation field, some recent research addresses motion field discontinuity problem using variation-based techniques for joint segmentation and estimation [24, 134]. In these methods, smooth regions and singularity set (edges) are devised according to image intensity, and registration aims to align each part respectively. The smoothness and discontinuity in the deformation itself is not addressed directly. We adopt the regularized optimization framework, and propose regularization designs to address the tissue-dependent elasticity and discontinuity preservation issues respectively.

Section 3.1 introduces the regularized registration setup. Section 3.2 provides an approach to incorporating tissue-type-dependent rigidity information into nonrigid registration² and Section 3.3 proposes a class of discontinuity preserving regularizers to address the effects of sliding along motion interfaces. Given the general regularized optimization setup, various optimization techniques can be used. We adopted B-spline

¹This chapter is based on material from [102].

²We proposed this method in 2006 [102] while [113] and [114] independently studied a similar penalty in 2006 and afterwards.

parametrization in Section 3.2 for its natural smoothness, and variational flow in Section 3.3 to better reveal the anisotropic filtering structure. These are specific choices for representing the deformation that are independent of the regularizers themselves, and should not be considered as limitations: in particular, the variational flow solved on rectangular grids can be regarded as a special case of zero-th order B-spline with its support equal to the pixel size. Preliminary results are demonstrated with each approach.

3.1 General Optimization Formulation for Regularized Registration

The goal of nonrigid registration is to find the optimal transformation T^* such that the transformed source image best matches the target. We use $f, g : \Omega \rightarrow \mathfrak{R}$ to denote the intensity map for the source and target images respectively, where d is the image dimensionality, and the open set $\Omega \subset \mathfrak{R}^d$ denotes the physical region of interest (ROI) for registration. Let $T : \Omega \rightarrow \mathfrak{R}^d$ be the transformation. Our goal is to find:

$$\begin{aligned} \hat{T} &= \arg \min_{T \in \Gamma} E(T, f, g) \\ (3.1) \quad &= \arg \min_{T \in \Gamma} \{E_d(g, f \circ T) + E_r(f, g; T)\}, \end{aligned}$$

where the set Γ is the class of admissible transformations. E is the overall objective function that we want to minimize, consisting of two parts: $E_d(g, f \circ T)$ denotes the data dissimilarity measure, also called data infidelity term, and $E_r(f, g; T)$ denotes the regularization term that is applied to penalize undesirable transformations. In the general regularization setting, E_r can also depend on images f and g .

3.1.1 Data Dissimilarity (Infidelity) Measure

Let $\underline{x} \in \Omega$ denote the coordinate (in vector form) of a specific spatial location. We use $T_{\underline{x}}$ to denote the local transformation at location \underline{x} and $\sum_{\underline{x}}(\cdot)$ to denote the summation over a discrete lattice that is a subset of Ω .

Sum of Squared Differences (SSD)

The sum of squared differences is a sensible data dissimilarity metric when the reference and the homologous image are acquired with the same modality with consistent parameters:

$$(3.2) \quad E_{d,SSD} = \sum_{\underline{x}} (g(\underline{x}) - f(T(\underline{x})))^2.$$

This metric has been considered by [58, 61, 62, 85, 115].

Mutual Information (MI)

When different modality images are to be registered, mutual information (MI) is a popular choice, since it does not require explicit knowledge about the intensity mapping between different modalities [20, 72, 75, 88, 116, 126, 130]:

$$\begin{aligned} E_{d,MI} &= -I(g, f \circ T) \\ (3.3) \quad &= -H(g) - H(f \circ T) + H(g, f \circ T), \end{aligned}$$

where $H(\cdot)$ denotes the entropy of a random variable and $H(\cdot, \cdot)$ denotes the joint entropy of two random variables.

In medical image data, we only have access to discrete samples of the intensity. To both improve the smoothness of the dissimilarity measure and approximate its derivative, we use Parzen window to estimate a differentiable entropy from the sample values [25]. Following the setup in [116], the joint discrete Parzen histogram is:

$$(3.4) \quad h^{\text{Paz}}(F, G; T) = \frac{1}{\epsilon_f \epsilon_g} \sum_{\underline{x}} w\left(\frac{F - f(T(\underline{x}))}{\epsilon_f}\right) w\left(\frac{F - g(\underline{x})}{\epsilon_g}\right),$$

for $F \in \mathcal{B}_f$ and $G \in \mathcal{B}_g$, where \mathcal{B}_f and \mathcal{B}_g are discrete sets of intensities associated with the source and target images respectively. $w(\cdot)$ is the Parzen window that integrates to unity, and ϵ_f, ϵ_g control the width of Parzen window in each dimension of the joint histogram.

The data infidelity term (negative mutual information) is computed using the normalized joint discrete Parzen probability $p(F, G; T) \propto h^{\text{PaZ}}(F, G; T)$ as:

$$(3.5) \quad E_{d,MI} = - \sum_{F \in \mathcal{B}_f} \sum_{G \in \mathcal{B}_g} p(F, G; T) \log_2 \frac{p(F, G; T)}{p_f(F; T)p_g(G)},$$

where $p_f(F; T)$ and $p_g(G)$ are obtained by marginalizing the joint probability $p(F, G; T)$ over bins \mathcal{B}_g and \mathcal{B}_f respectively.

Other dissimilarity criteria used in image registration include correlation coefficient and its variation; and landmark matching based comparison. It is also common to combine two or more of the above metric (*e.g.*, SSD and landmark) depending on the applications.

3.2 Tissue-type Dependent Rigidity Regularization

For modeling efficiency, we parametrize the deformation field $\Phi(\underline{x}) \triangleq T(\underline{x}) - \underline{x}$ instead of the transformation T itself. To improve the conditioning of the problem, a roughness penalty is incorporated in terms of the gradients of the deformation Φ , using the squared Frobenius norm $\|\nabla\Phi\|_{Frob}^2$. We define the local tissue rigidity based regularization to be a weighted superposition of local non-rigidity penalty, $\sum_{\underline{x}} \gamma(\underline{x})r(T_{\underline{x}})$. The overall regularizer reads:

$$(3.6) \quad \begin{aligned} E_r(f, g; T) &= E_{nonrigid}(f, g, T) + E_{roughness}(T) \\ &= \sum_{\underline{x}} \{ \gamma(\underline{x})r(T_{\underline{x}}) + \alpha(\underline{x}) \|\nabla\Phi_{\underline{x}}\|_{Frob}^2 \}. \end{aligned}$$

Here, we focus on designing $E_{nonrigid}$, where we will choose $r(T_{\underline{x}})$ to penalize the deviation of the local transformations from being rigid, and $\gamma(\underline{x})$ is the spatially varying weight that reflects local tissue rigidity properties. In particular, $\gamma(\underline{x})$ controls the *local* “trade-off” between intensity match and deformation rigidity. It should be large within bone structures and small within more elastic regions, *e.g.* muscle and fat. We call it “*local stiffness factor*” to reflect this physical interpretation. Correspondingly, the spatially varying “*local smoothness factor*” $\alpha(\underline{x})$ controls the local trade-off between intensity match and deformation smoothness. Since we are mainly interested in spatially varying stiffness property in this work, we set $\alpha(\underline{x})$ to be a constant throughout the ROI for simplicity.

3.2.1 Regularization Design

Local Rigidity Functional

The local rigidity functional $r: (\mathfrak{R}^d \rightarrow \mathfrak{R}^d) \rightarrow \mathfrak{R}^{\geq 0}$ quantifies how much the local transformation deviates from being rigid. We desire the functional r to have the following properties:

- $r(T_{\underline{x}}) = 0$ if and only if $T_{\underline{x}}$ is a rigid transform.³
- The functional r should be invariant to orthogonal coordinate transformation.

To satisfy the first property, we utilize the following arguments:

Lemma III.1. *A necessary and sufficient condition for a transformation T to be rigid at \underline{x} is that its Jacobian matrix $DT_{\underline{x}} \triangleq \nabla T(\underline{x})$ is orthogonal.*

³Here, we equate rigid transformation with the isometry in \mathfrak{R}^d , which by formal definition also includes reflections. However, reflection rarely occurs in practice. Moreover, the roughness penalty described in (3.6) and our choice of a smooth basis for parametrization the deformation field further decreases the chance of a local reflection in the transformation estimate.

The proof follows from the group structure of the isometry on \mathfrak{R}^d , and the fact that the Jacobian operation provides a group homomorphism between the isometry group on \mathfrak{R}^d and the orthogonal group in d -dimension.

Lemma III.1 involves a matrix property, so it suffices to design a penalty that measures how “non-orthogonal” the Jacobian matrix of the local transformation $DT_{\underline{x}}$ is.

We use the following fact:

Lemma III.2. *A necessary and sufficient condition for a matrix $M \in \mathfrak{R}^{d \times d}$ to be orthogonal is that $\|MM^T - I_d\| = 0$, where $\|\cdot\|$ denotes any matrix norm.*

If M is orthogonal, $MM^T = I_d$, and $\|MM^T - I_d\| = 0$ for any norm. On the other hand, for any matrix norm, $\|MM^T - I_d\| = 0$ implies $MM^T = I_d$, which is exactly the definition for a square matrix M to be orthogonal. \square

Therefore, once we define $r(T_{\underline{x}})$ based on $\|DT_{\underline{x}}(DT_{\underline{x}})^T - I_d\|$, the first required property is automatically satisfied.

Lemma III.3. *$\|DT_{\underline{x}}(DT_{\underline{x}})^T - I_d\|$ is invariant under isometric (rigid) transformations.*

Isometric transforms on the coordinate system can be incorporated into the local transformation $T_{\underline{x}}$ by applying the inverse transform. By the chain rule of differentiation, it immediately follows that $D(T_{\underline{x}} \circ g) = DT_{\underline{x}}Dg$. If g is an isometry by assumption, then Dg is an orthogonal matrix, and the invariance result follows from a simple manipulation:

$$(3.7) \quad \begin{aligned} D(T_{\underline{x}} \circ g)D(T_{\underline{x}} \circ g)^T &= DT_{\underline{x}}Dg(Dg)^T(DT_{\underline{x}})^T \\ &= DT_{\underline{x}}(DT_{\underline{x}})^T. \end{aligned}$$

Thus $\|DT_{\underline{x}}(DT_{\underline{x}})^T - I_d\|$ also satisfies the second property above. \square

For simplicity and computation efficiency, we choose to use the squared Frobenius norm, and define the following local rigidity regularization function:

$$(3.8) \quad r(T_{\underline{x}}) \triangleq \frac{1}{2} \|DT_{\underline{x}}(DT_{\underline{x}})^T - I_d\|_{\text{Frob}}^2.$$

Some previous work enforces tissue incompressibility by constraining the Jacobian determinant to be close to unity [55], but a unity valued transformation determinant is only a necessary but not sufficient condition for local rigidity. The combination of Jacobian determinant with its condition number may be a possible alternative, but would require spectral analysis which is computationally demanding. We choose the squared Frobenius norm because it satisfies the two properties above and yet is easy to compute.

Local Stiffness Factor

To design the spatially varying local stiffness factor $\gamma(\underline{x})$, which determines the relative weighting between data infidelity and deviation from rigidity, it would be desirable to have accurate knowledge about mass, elasticity, as well as other mechanical properties. Unfortunately, detailed information is rarely available. Instead, we infer the rigidity level of local tissue from observed CT values. The empirical design could be improved given more precise/specific prior knowledge. We observe that in calibrated X-ray CT images, pixel intensity (CT number) is highly correlated with tissue type information, hence is a good inference source for local rigidity. Therefore, instead of designing a direct map $\gamma: \Omega \rightarrow \mathfrak{R}^+$, we define the local stiffness factor by applying a transfer function $s(\cdot)$ to the image intensity map:

$$\gamma(\underline{x}) = s(f(\underline{x})),$$

where $s: \mathfrak{R} \rightarrow \mathfrak{R}^+$ is a monotone increasing map from the domain of CT number to rigidity level. We choose to use a scaled and shifted hyperbolic tangent function in our application due to its simplicity (two parameters with clear shape meaning) and desirable mapping form: the properly placed sharp rising edge distinguishes bone structures from more elastic tissues, while the saturation behavior is robust to small intensity variations of the same tissue type.

Fig. 3.1 shows the empirical histogram taken from a $192 \times 160 \times 60$ breath-held thorax CT volume with voxel size $0.2 \times 0.2 \times 0.5 \text{cm}^3$. Observations for the tissue type v.s. CT number (in Hounsfield unit) relationship agree with theoretical values [49] in that:

$$\begin{array}{rcl} \text{Air} & : & -1000\text{HU} \\ \left. \begin{array}{l} \text{Fat} \\ \text{Muscle} \end{array} \right\} & : & -100 \sim 60\text{HU} \\ \text{Bones} & : & 250 \sim 1000\text{HU}. \end{array}$$

We choose the location and shape parameters for the hyperbolic function such that the non-rigidity penalty dominates in the bony structures, and is relaxed within elastic tissues.




Figure 3.1: Illustration of stiffness factor $s(\cdot)$. (a) design of functional h based on theoretical tissue-type-to-CT-number map; (b) scaled stiffness factor v.s. tissue type information inferred from empirical histogram.

Parametrization and Optimization

We adopt the widely used tensor product B-spline basis to parametrize both the deformation field Φ [62] and the image intensity. In practice, we often use B-spline $\beta^n(x)$ of order $n = 3$ for both purposes in volumetric registration. B-splines are smooth functions with explicit derivatives [118] and finite support. They are piecewise polynomials and can be recursively constructed by convolution [120, 121].

The deformation for each direction r is represented independently with the corresponding set of B-spline coefficients $\Theta^l = \{\theta_i^l\}$ as follows:

$$(3.9) \quad \Phi^l(\underline{x}; \Theta) = \sum_{i \in N(\underline{x})} \theta_i^l \beta_i(\underline{x}).$$

For volumetric case ($d = 3$), $l \in \{1, 2, 3\}$ represents deformation direction along x, y and z coordinates respectively, and separable B-spline basis is used:

$$\beta_{\underline{i}}(\underline{x}) = \beta\left(\frac{x}{\Delta_x} - i\right) \beta\left(\frac{y}{\Delta_y} - j\right) \beta\left(\frac{z}{\Delta_z} - k\right),$$

where $\underline{i} = (i, j, k)$ denotes the B-spline knot location, $\Delta_x, \Delta_y, \Delta_z$ determines the scale of B-spline in each direction, $\underline{x} = (x, y, z)$ denotes the spatial location, and its neighborhood $N(\cdot)$ is determined by the support of the B-spline basis.

The image model provides a continuous representation of an image given by a set of samples. In fact, only the source image requires interpolation in the formulation considered here:

$$(3.10) \quad f(\underline{x}) = \sum_{i \in \mathcal{N}(\underline{x})} c_i \beta_i(\underline{x}),$$

where the expansion B-spline coefficients c_i are computed from the sample values of G by recursive digital filtering [121].

We utilize a multi-resolution scheme in the registration process, and use gradient descent method at each resolution level to evolve the overall cost function until convergence.

For optimization, we used the derivative of the SSD energy (3.2), given by:

$$(3.11) \quad \frac{\partial}{\partial \theta_i^l} E_{d,SSD} = \sum_{\underline{x}} (g(\underline{x}) - f(T(\underline{x}))) \nabla f|_{T(\underline{x})} \beta^n(\underline{x} - \underline{l}).$$

The derivative for negative mutual information from (3.5) is given by [116]:

$$(3.12) \quad \frac{\partial}{\partial \theta_i^l} E_{d,MI} = - \sum_{F \in \mathcal{B}_f} \sum_{G \in \mathcal{B}_g} \frac{\partial}{\partial \theta_i^l} p(F, G; T) \log_2 \frac{p(F, G; T)}{p_f(F; T)}.$$

The terms involved in evaluating the regularization are:

$$(3.13) \quad \nabla \Phi_{\underline{x}} = \left[\sum_i \theta_i^l \beta_i^j(\underline{x} - \underline{l}) \right]_{(l,j) \in \{1,2,3\} \times \{1,2,3\}},$$

where β_i^j denotes the derivative of the basis function β in the j th direction. Using the derivative property of the B-spline, the derivative of β can be computed analytically [118]:

$$(3.14) \quad \frac{\partial}{\partial x} \beta^n(x) = \beta^{n-1}(x+1/2) - \beta^{n-1}(x-1/2).$$

The local tissue rigidity based penalty term is similarly derived based on the fact that

$$DT_{\underline{x}} = \nabla T_{\underline{x}} = \nabla \Phi_{\underline{x}} + I_d.$$

The derivative of the penalty with respect to deformation parameter θ_i^l can be written as:

$$(3.15) \quad \frac{\partial}{\partial \theta_i^l} \sum_{\underline{x}} \gamma(\underline{x}) r(T_{\underline{x}}) = \sum_{\underline{x}} \gamma(\underline{x}) \text{trace} \left\{ [DT(DT)^T - I_d] \left[\frac{\partial}{\partial \theta_i^l} DT DT^T + DT \frac{\partial}{\partial \theta_i^l} (DT)^T \right] \right\},$$

where we precompute and store $\frac{\partial}{\partial \theta_i^l} DT = \beta^l$ for computation efficiency.

3.2.2 Experiment and Test Results

Experiment One: Geometry Validation by Thresholding

In the first experiment, we tested the proposed approach with two thorax CT scans of the same patient: one at 80% of the vital capacity inhale breath hold (deep inhale breath hold, tidal breathing generally peaks at about 40%) and one at exhale. The scans were $512 \times 512 \times 148$ with voxel size $0.2 \times 0.2 \times 0.5 \text{cm}^3$. We used the deep inhale breath-hold thorax CT image as the target and further cropped it to size $259 \times 175 \times 107$ to reflect the region of interest. Sum of Squared Differences (SSD) was used as dissimilarity metric. Fig. 3.2 shows typical data slices (different views) of the target image, source image and the inferred stiffness map ($h \circ f$). The inferred stiffness map captures rigid structures reasonably well.

We first show the registration results in slice views for pure global rigid, affine transformation, and nonrigid registration with and without nonrigid regularization. The deformed source image is displayed on top of the target image for comparison purposes in Fig. 3.3.




Figure 3.2: Different views of the original data and tissue information inferred from it. Top row [X(1)]: coronal slices; middle row [X(2)]: sagittal slices; bottom row [X(3)]: axial slices. Left column [a(#)]: slices from target image; middle column [b(#)]: slices from source image; right column [c(#)]: slices from inferred stiffness map.

Fig. 3.3 illustrates that nonrigid registration outperforms global rigid/affine model based registration on matching intensity. The advantage is most obvious in regions where organs have undergone extremely elastic deformations, such as the diaphragm. The different performance in the lung area is less noticeable due to the overall low intensity level in lung region, so mismatch in that region is not emphasized in SSD setting. Finally, the introduction of proposed tissue type dependent regularization does not seriously deteriorate intensity matching performance compared to conventional B-spline in general.

To better reveal the geometry of the deformation, we extracted bone structures by thresholding the CT numbers at 250 HU, because they are good indicators of tissue type. Geometry extracted from both the target and the deformed source volumes are overlaid to compare the bone structure alignment in Fig. 3.4.

We can clearly observe nonphysical warping of bones in the deformed source geometry using conventional B-spline based nonrigid registration method without the proposed regularization. This is a typical local optimum situation. Upon localizing the occurrence of this particular “bone warping” phenomena, we can observe that the “pseudo-periodic” structure of the ribs makes the resulted deformation and the desired physical one having comparable intensity dissimilarity (data infidelity) value. On the other hand, since

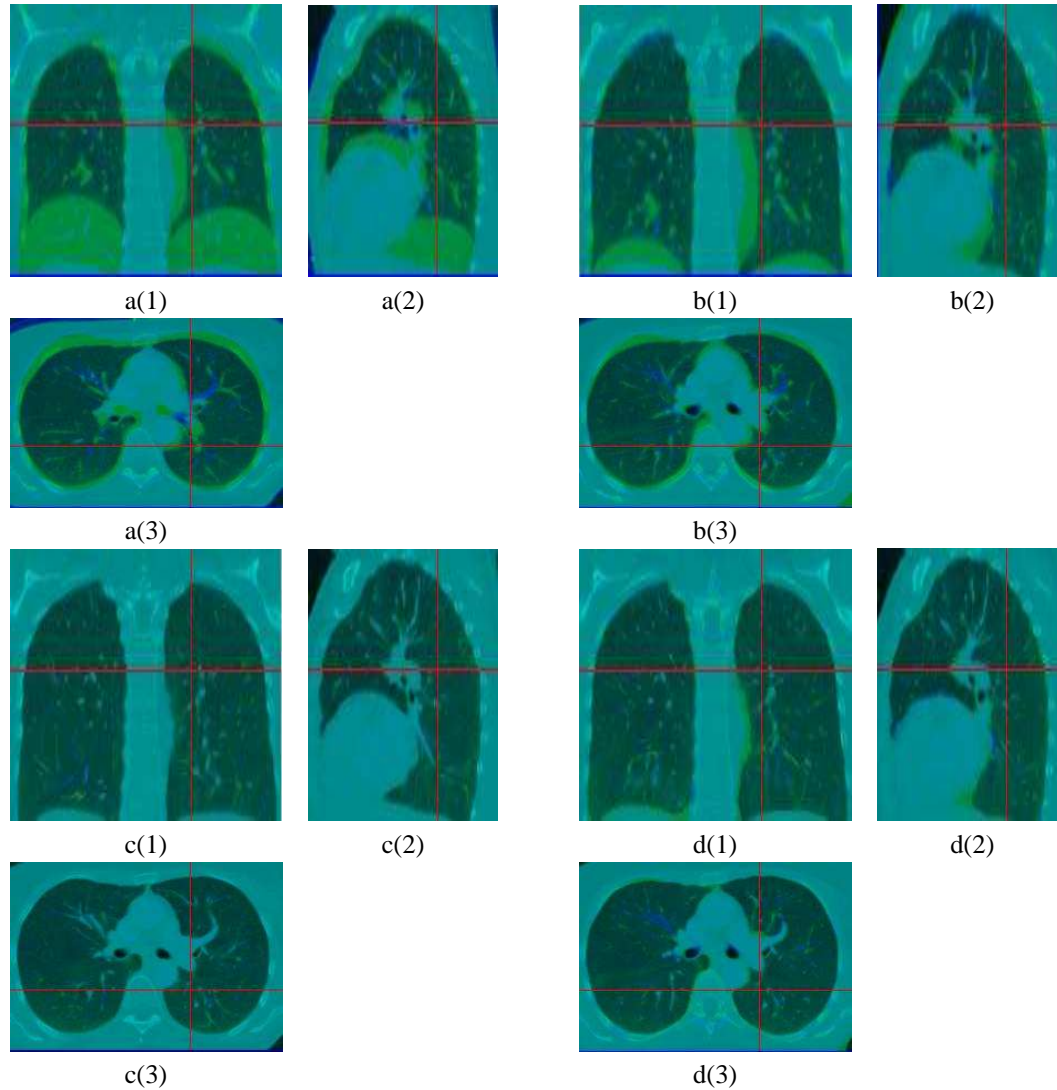


Figure 3.3: Deformed source image (green) overlaid with target image (dark blue) for comparison of intensity match. Different views are indicated with numbers: [X(1)] coronal view; [X(2)] sagittal view; [X(3)] axial view [X(3)]. Different registration method are distinguished with letters: [a(#)] rigid transformation model; [b(#)] affine transformation model; [c(#)] B-Spline registration with smoothness penalty only; [d(#)] B-Spline registration with both proposed regularization.

B-spline is a smooth local basis, together with smoothness regularization to enforce continuity of the deformation field, in regions close to diaphragm/lung region where deformation of more elastic nature occurs, the deformation of bone structures are compromised to resemble those of elastic tissues.

When the proposed regularization is applied, however, the deformation on the bone structures are given an additional “force” to conform to rigid transformation. Fig. 3.4 shows obvious improvements regarding the bone-warping issue.

Experiment Two: Quantitative Validation with Bifurcation Landmarks

In the second experiment, we evaluate the registration accuracy in soft tissue regions as it might be adversely affected by the introduction of the proposed regularization. Sequential thorax CT scans were




Figure 3.4: Geometry extracted from registration results: target (blue) vs. deformed source (white). Left column [a(#)]: B-spline based nonrigid registration with no local rigidity regularization; right column [b(#)] B-spline based nonrigid registration result with proposed local tissue type dependent regularization. Top row [X(1)]: whole ribcage view; bottom row [X(2)]: local zoom-in view around diaphragm neighborhood.

obtained on a helical CT scanner (CT/I, General Electric, Milwaukee, WI) for 11 patients. Two scans were obtained from each patient, one at normal exhale followed immediately by a scan at normal inhale during coached voluntary breath-hold periods of 18-35 seconds. Scans were obtained with a pitch of 2, using a 5mm aperture. The total time spent from the start of the first scan through the completion of the second scan was less than 5 minutes. Images were reviewed by experts to ensure that they were free of breathing-related artifacts in reconstruction. To quantitatively analyze the registration accuracy, we compare the position of known features in the target and source images. A human observer chose six landmarks within the right lung per patient [21]. Landmarks included vascular and bronchial bifurcations, and were nearly uniformly distributed in the ROI. Computed transform from registration algorithms was applied to the landmark coordinate in the target image and compared to the landmark position in the source image coordinate. Fig. 3.5 illustrates some of the manually picked landmarks.

We applied negative mutual information (MI) as the data dissimilarity metric to reflect the general applicability of the proposed methods, even though X-ray CT images are used both as the source and the target image to maximize the consistency of manually picked corresponding landmark pairs. Moreover, landmarks picked at lung bifurcations should fairly characterize the effect of the additional regularization on the soft tissue regions. We compared thin-plate splines (TPS), conventional B-splines and the proposed regularized B-splines in this test. In TPS setup, control points were placed manually on the source and the target dataset. We used the TPS results from [21], where 30 control points were used to align the inhale and exhale CT model of the right lung, with 5 each on 6 specified Superior-Inferior planes in the target dataset. Nelder-Mead simplex algorithm was used to maximize MI for TPS. For the conventional and modified B-spline registration, multi-resolution scheme was used to achieve computation efficiency. In each resolution level, control knots were placed uniformly in the low-pass filtered source image, and B-spline coefficients are updated using gradient descent algorithm until convergence.

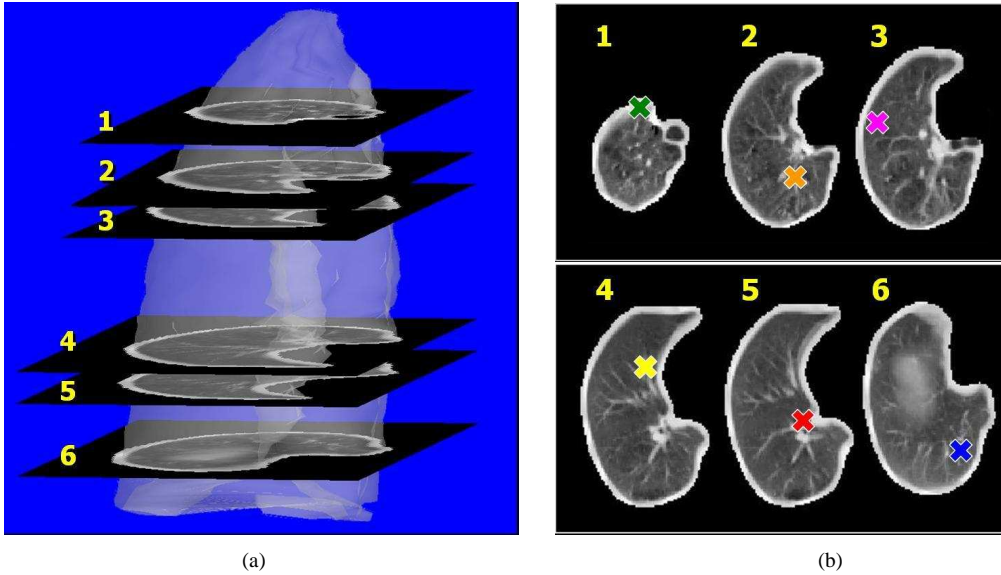


Figure 3.5: Illustration of landmark data on thorax CT: (a) illustration of volumetric data; (b) manual landmark positioning based on bifurcations

We computed the difference between the deformed landmark positions on the source coordinate and the corresponding manually picked target landmark position. Fig. 3.6 shows box plots illustrating median, lower/higher quartile, data extent and outliers to characterize the registration accuracy along each axis: right-left (RL), anterior-posterior (AP), and inferior-superior (IS). The regularized B-spline registration is competitive against thin-plate splines or conventional splines inside the lung. Limitation of human observer due to image resolution (voxel size $0.2 \times 0.2 \times 0.5\text{cm}^3$) and the dominant motion in inferior-superior direction are also reflected in the registration performance.

We also calculated the Euclidean registration error between deformed landmark locations and the manually selected points. In Fig. 3.7, we ordered the patients according to the mean Euclidean error for TPS method, and used box-plot to illustrate the Euclidean error distribution for different methods. Fig. 3.7(d) shows the mean Euclidean error of landmark position estimate for each patient, and a box-plot of the collective Euclidean error for each method is provided in Fig. 3.7(e). Both conventional B-splines and regularized B-splines uniformly outperform the manually assisted thin-plate splines method, whereas performance of the two B-splines based registration methods are comparable. This agrees with the qualitative results in Fig. 3.3 where the proposed regularization appears to preserve the flexibility of the conventional B-splines method in soft tissues. The mean and standard deviation of Euclidean error for regularized B-spline is $M_{R-BSP} = 0.5$ cm and $\sigma_{R-BSP} = 0.48$ cm respectively, on the same order as the slice thickness, and superior to $M_{TPS} = 0.85$ cm and $\sigma_{TPS} = 0.55$ cm from TPS or $M_{BSP} = 0.56$ cm and $\sigma_{BSP} = 0.55$ cm from conventional B-spline.

We used three B-spline resolution levels which took about 100 iterations in the last (finest) resolution level to converge. The computation time for both the conventional B-splines and the regularized B-spline are both in the order of minutes on a standard PC (2.4 GHz CPU and 1G internal memory) running Linux. All programming and visualization in this paper were carried out on the Advanced Visual Systems (AVS) software platform with central modules implemented in C/C++. Including the regularization increased the registration time by less than 20% in most cases.

3.2.3 Discussions

We quantify local non-rigidity by the deviation of the local Jacobian from being orthogonal, measured by a computationally efficient Frobenius norm. We considered both mono- and multi- modality registrations involving a CT image as either the source or target observation. Local tissue rigidity level is inferred by applying a smooth monotone function to the CT values, avoiding explicit segmentation. The smoothness

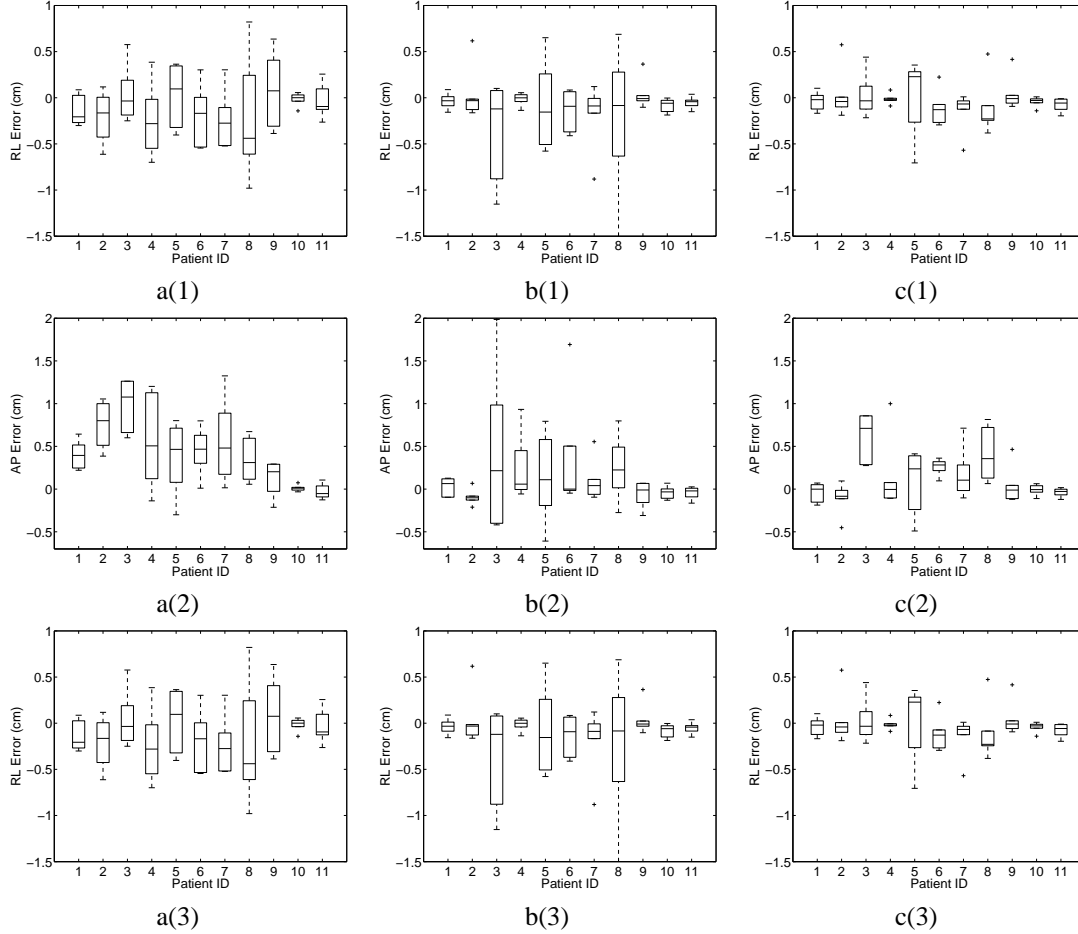


Figure 3.6: Registration error for different methods: TPS, BSP and Regularized BSP. Left column [a(#)]: Thin plate spline registration with manually picked control points; middle column [b(#)] conventional B-spline registration; right column [c(#)] B-spline registration with proposed local tissue type dependent regularization. Top row [X(1)]: right-left (RL) registration error in right-left (RL) direction; middle row [X(2)]: registration error in anterior-posterior (AP) direction; bottom row [X(3)]: registration error in inferior-superior (IS) direction.

of the inference function provides robustness to partial volume effects caused by limited resolution and by multi-resolution schemes deployed to speed up computation.

The proposed regularization design is independent of the user-specified dissimilarity metric and the parametrization of the transformation field. We evaluated registration accuracy using the popular B-spline deformation parametrization, with two different dissimilarity metrics: sum of squared differences (SSD) and negative mutual information (MI). In the first case, we visualized bone geometry in the target and the deformed source image for qualitative assessment. In the second case, we compared deformed landmark locations with manually specified “*ground-truth*” values for quantitative validation. Comparison among thin-plate splines (TPS), conventional B-splines and the proposed method indicates minimal compromise of registration accuracy in soft tissue regions, but significantly improved ribcage registration.

We have performed a preliminary study on incorporating tissue type information into nonrigid registration framework essentially via the introduction of a spatially-varying stiffness coefficient map and use that to adjust the local trade-off between intensity match and rigidity property. This allows inhomogeneous regularization throughout the deformation field estimation. We would like to extend this work in the following aspects:

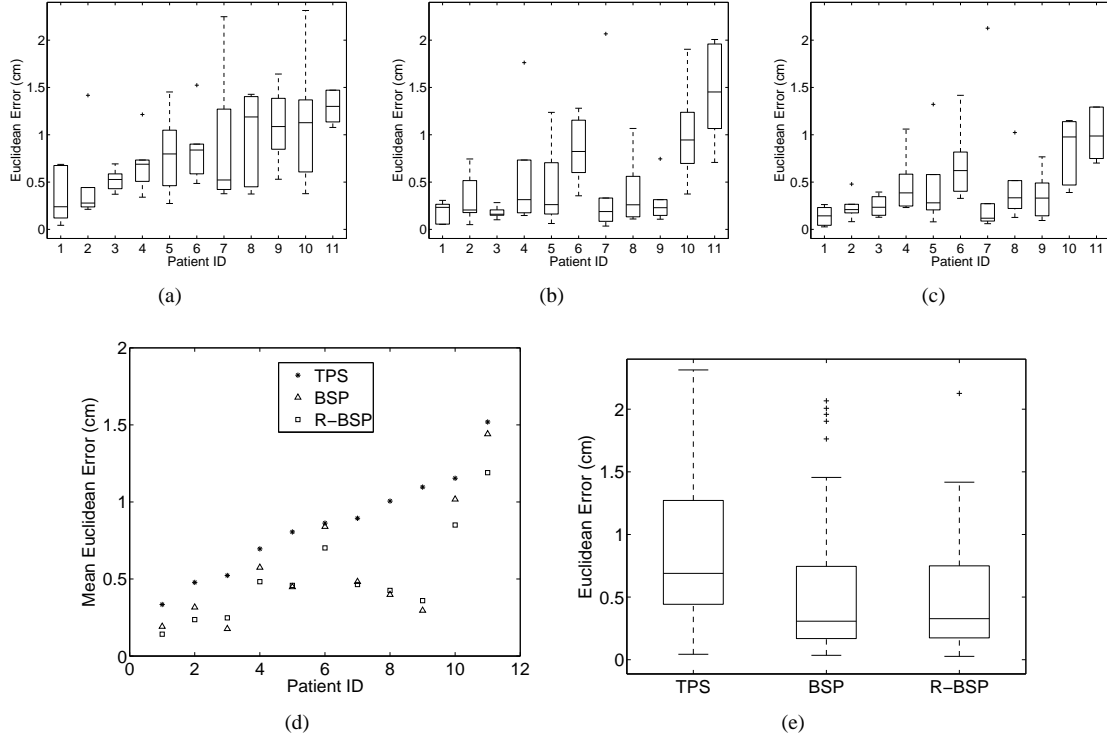


Figure 3.7: Comparison of 3-dimensional Euclidean Error. Top row (left to right): TPS, BSP and Regularized BSP. Bottom left: mean Euclidean error over all landmarks on the same patient; bottom right: box plot of the Euclidean error distribution over all landmarks through all patients.

- We would like to extend the non-homogeneity that we introduced in this work further to non-isotropic setting. In many situations, anatomical structures not only demonstrate tissue-type dependent inhomogeneous deformation, but also directional variations. Examples are bending in head-neck region and the dominant elongation/deformation in up-down direction (vertical direction in sagittal plan) related to breathing motion. These information could be handled in heuristic fashions by non-uniformly placing the B-spline knots and having different knot spacing in different directions. However, these ad-hoc techniques could be tricky in practice and lacks certain theoretical justification. Furthermore, adjusting B-spline knots can only strictly control the deformation level in the 3 vertical plans, which may not be sufficient for some clinical applications. We conjecture that by introducing anisotropic regularization into the optimization framework, we would be able to have a more flexible and straight-forward way to accommodate direction related priors.
- In X-ray Computed Tomography, we designed the stiffness map as the composition of a monotone increasing function with the intensity map, taking advantage of the fact that in this particular modality, intensity is a very reliable reference source for tissue type information. This is not true in general. We would like to explore approaches to address this issue for other modalities in future work.

3.3 Discontinuity-Preserving Regularization

Motivated by the common presence of sliding effects in medical imaging, *e.g.*, the discontinuous motion between diaphragm and ribcage during breathing, we study regularization schemes that preserve discontinuities in the deformation field.

Recent research on image registration that accounts for discontinuities can be classified into two cat-

egories. The first class [23, 24, 134] is based on joint segmentation and registration. In these methods, smooth regions and singularity sets (edges) are devised according to image intensity, and registration aims to align each part respectively. The smoothness and discontinuities in the deformation is not addressed directly. The second category is motivated by edge-preserving image restoration [34, 35, 77]. Several authors [10, 15, 33, 129] have tried to generalize total variation type regularization for vector valued functions. These methods use regularization that combines the total variation from each deformation coordinate. Meanwhile, decomposition and representation of a vector field by velocity potentials and stream functions [38] have motivated flow regularizations with divergence and curl components [42, 135, 137]. Analogous to image denoising, [136] has proposed a convex Hodge decomposition based total variation regularization method to denoise vector fields, resulting in piecewise harmonic flows. This paper is closely related to the latter category, and intends to adapt such principles to design regularizations for medical image registration applications.

There are many ways to extend regularizers, *e.g.*, Tikhonov or total variation (TV), originally developed for scalar fields, to vector flow applications. However, naive extensions may violate the intrinsic structure of the problem, and result in loss of desirable properties. Taking total variation as an example, summing over the total variation in each component direction [10, 33, 129] compromises the rotational invariance with respect to the coordinate system. Section 3.3.1 handles this issue with care and proposes a class of regularizers for vector fields that preserve discontinuities in the deformation field. We provide general analysis of their functional forms, and define some desired properties as a consequence. We derive the descending flow for optimization based on variational calculus and discuss briefly some implementation issues.

Section 3.3.2 further notes that only sliding or shear discontinuity is physical in medical image registration, hence it is necessary to distinguish this class of admissible discontinuity from collision or vacuum creating singularities. To design a regularizer that differentiates between these two types of discontinuities and preserves only large shears, we take advantage of the Helmholtz decomposition, and regularize the divergence and curl components of the vector field differently.

Preliminary result for this work in progress shows promising results.

For clarity, we discuss the derivations for 2D case, yet all analysis generalizes naturally to higher dimensions unless specified otherwise. We represent the deformation vector field $\Phi : \Omega \rightarrow \mathfrak{R}^2$ as $\Phi(\underline{x}) = [u(\underline{x}), v(\underline{x})]^T$, where u and v are directional deformation and assumed to be orthogonal (but do not have to align with the image coordinate (x, y)) in general. As we are mainly interested in geometric regularization for smoothness/discontinuity, the regularization term is taken to be independent of the image. It corresponds to a special case of the regularized registration problem introduced in Section 3.1 with

$$E_r(f, g; T) = E_r(\Phi).$$

A constant weight λ is adopted throughout the whole image to balance the data fidelity and regularization energy. We focus on designing E_r , and assume mono-modality images with L_2 metric as data fidelity measure hereafter. Thus the goal of registration can be formulated as:

$$(3.16) \quad \Phi^* = \arg \min_{\Phi \in \Gamma} E(f, g, \Phi)$$

$$(3.17) \quad = \arg \min_{\Phi \in \Gamma} \{E_d(g, f \circ (I + \Phi)) + \lambda E_r(\Phi)\}.$$

3.3.1 Indiscriminate Discontinuity Preserving Regularization

To encourage smooth deformations in most of the region of interest (ROI), yet admitting some discontinuities requires a “magnitude” measure of the local change of the deformation field, analogous to the norm of image gradient in image restoration. The Jacobian of the deformation Φ at \underline{x} is given by:

$$D\Phi(\underline{x}) = \begin{bmatrix} u_x & u_y \\ v_x & v_y \end{bmatrix}.$$

We propose to use the Frobenius norm of the matrix $D\Phi(\underline{x})$ as the local measure of variation for the deformation field:

$$(3.18) \quad \begin{aligned} |D\Phi|_{\text{Frob}} &= \sqrt{u_x^2 + u_y^2 + v_x^2 + v_y^2} \\ &= \sqrt{|\nabla u|_2^2 + |\nabla v|_2^2}. \end{aligned}$$

This matrix norm is independent of both the image coordinate system (x, y) and the deformation vector field direction (u, v) . For simplicity, we assume that the u and v components of the vector field correspond to the deformation field in x and y directions respectively hereafter. In addition, this measure of “deformation change” introduces coupling among the various directions in the vector fields and reflects the intuition that we observe a “jump” in the deformation field regardless of the specific direction such change occur, unlike the simple coordinate-wise sum used in traditional optical flow regularization [3, 8]. For simplicity, we make matrix Frobenius norm the default notation for $|D\Phi|$ hereafter and drop the subscript.

We consider a class of regularizers with the form:

$$(3.19) \quad E_r(\Phi) = \int \phi(|D\Phi|) d\underline{x}.$$

Applying variational analysis, and assuming Neuman boundary conditions, *i.e.*, $\partial_n u = 0$ and $\partial_n v = 0$ on $\partial\Omega$, we derive the descent flow [3] $w_r = (u_r, v_r)$ of E_r to be as follows:

$$(3.20) \quad \begin{aligned} u_r &= \nabla \cdot \left(\frac{\partial}{\partial \nabla u} \phi \right) \\ &= \nabla \cdot \left(\frac{\phi'(|D\Phi|)}{|D\Phi|} \nabla u \right). \end{aligned}$$

The expression for the update flow v_r for v is similar. For simplicity, we define the “influence function” as $\psi(s) \triangleq \phi'(s)/s$.

To design a proper regularization ϕ that results in edge preserving flow, we interpret the process as anisotropic filtering and decompose the effect of the flow into the normal and tangent directions for *each component of the deformation field*. We derive the regularization flow in u - direction as:

$$(3.21) \quad u_r = \psi(|D\Phi|)(u_{xx} + u_{yy}) + \frac{\phi''(|D\Phi|) - \psi(|D\Phi|)}{|D\Phi|^2} (u_x^2 u_{xx} + 2u_x u_y u_{xy} + u_y^2 u_{yy}).$$

By convention, we denote the second derivatives of u in the tangent (T-) direction and normal (N-) direction as u_{TT} and u_{NN} respectively, with

$$u_{TT} = T^T \nabla^2 u T = \frac{1}{|\nabla u|} (u_x^2 u_{yy} + u_y^2 u_{xx} - 2u_x u_y u_{xy});$$

$$u_{NN} = N^T \nabla^2 u N = \frac{1}{|\nabla u|} (u_x^2 u_{xx} + u_y^2 u_{yy} + 2u_x u_y u_{xy}).$$

Rearranging the terms in (3.21) yields:

$$(3.22) \quad u_r = \psi(|D\Phi|) u_{TT} + |\nabla u|^2 \left(\frac{\phi''(|D\Phi|)}{|D\Phi|^2} - \frac{\psi(|D\Phi|)}{|D\Phi|^2} + \frac{\psi(|D\Phi|)}{|\nabla u|^2} \right) u_{NN}.$$

For 2D case (higher dimension situations have similar structure):

$$\frac{\psi(|D\Phi|)}{|\nabla u|^2} - \frac{\psi(|D\Phi|)}{|D\Phi|^2} = \psi(|D\Phi|) \frac{|\nabla v|^2}{|D\Phi|^2 |\nabla u|^2}.$$

The coupling between u and v in the flow motivates us to consider the contribution of variation in each deformation direction in $|D\Phi|$. We define $\beta_u \triangleq \frac{|\nabla u|^2}{|D\Phi|^2}$ and $\beta_v \triangleq \frac{|\nabla v|^2}{|D\Phi|^2}$. By construction, $\beta \in [0, 1]$ and $\beta_u + \beta_v = 1$. Then (3.22) can be rewritten as:

$$(3.23) \quad u_r = (\phi''(s)\beta_u + \psi(s)\beta_v)u_{NN} + \psi(s)u_{TT},$$

Now we are ready to discuss some desired properties for the function ϕ . This is more complicated than image restoration problems as ϕ is intrinsically a function of both u and v .

- In the presence of small variations in the deformation, ($|D\Phi|$ small implies $|\nabla u|, |\nabla v|$ both small), isotropic smoothing is desirable in each individual deformation direction. It is reasonable to require non-trivial smoothing along the tangent direction:

$$(3.24) \quad \phi'(0) = 0, \quad \text{with } \lim_{s \rightarrow 0^+} \psi(s) > 0.$$

To have isotropic diffusion as $s \rightarrow 0^+$ is equivalent to:

$$\lim_{s \rightarrow 0^+} \beta_v + \beta_u \frac{\phi''(s)}{\psi(s)} = 1.$$

Together with the fact that $\beta_u + \beta_v = 1$, isotropic diffusion for small deformation implies

$$(3.25) \quad \lim_{s \rightarrow 0^+} \psi(s) = \lim_{s \rightarrow 0^+} \phi''(s) > 0.$$

Once the conditions (3.24) and (3.25) are satisfied, the flow (3.23) for small variation reduces to:

$$u_r \approx \phi''(0)\Delta u.$$

The same analysis holds for v_r . We immediately recognize that this diffusion coincides with the isotropic flow from the heat equation.

- In the presence of large variations in deformation (large $|D\Phi|$), it is desirable to diffuse the deformation along the discontinuity, but not across it. We need to keep in mind that the level of discontinuity $|D\Phi|$ takes into account deformation in all directions, and the diffusion process in a certain direction (u or v) is decomposed with respect to its own gradient field. In other words, the diffusion process in u direction is the projection of the joint deformation flow onto that direction. To preserve discontinuity, it suffices to annihilate the coefficients of u_{NN} and v_{NN} for large $|D\Phi|$, and assume non-vanishing coefficients for the tangent flow components.

$$\begin{cases} \lim_{s \rightarrow +\infty} \phi''(s)\beta_u + \psi(s)\beta_v = 0; \\ \lim_{s \rightarrow +\infty} \psi(s) > 0. \end{cases}$$

If one were to insist on the annihilation of the normal flow for all possible combinations of (β_u, β_v) , it would be necessary to require:

$$\lim_{s \rightarrow +\infty} \phi''(s) = 0 \quad \text{and} \quad \lim_{s \rightarrow +\infty} \psi(s) = 0.$$

On the other hand, if $\beta_u \approx 0$, indicating that the variation in x -direction ($|\nabla u|$) is relatively small, isotropic diffusion in that direction would not result in over-smoothing discontinuity and should be

acceptable. With v being the major contributor to the overall discontinuity in $|D\Phi|$, only v_{NN} has to be annihilated. Unfortunately, this again results in a set of incompatible conditions on ϕ :

$$\lim_{s \rightarrow +\infty} \phi''(s) \leq 0 \quad \text{and} \quad \lim_{s \rightarrow +\infty} \psi(s) \geq 0.$$

One possible compromise is to let both terms approach zero as $s \rightarrow +\infty$, but at different rates:

$$(3.26) \quad \begin{cases} \lim_{s \rightarrow +\infty} \phi''(s) = \lim_{s \rightarrow +\infty} \psi(s) = 0; \\ \lim_{s \rightarrow +\infty} \frac{\phi''(s)}{\psi(s)} = 0. \end{cases}$$

Many functions satisfy the above conditions (3.24),(3.25) and (3.26), e.g., the hypersurface minimal function $\phi(s) = \sqrt{1+s^2}$ [3]. Due to the nonconvex nature of registration problems, we are interested in finding only reasonable local minima in general. In the usual case where E_d is nonconvex in Φ , it may be unnecessary to insist on ϕ being convex.

We make a quick comment here:

- $\phi(s) = s^2$ corresponds to the regularization energy:

$$(3.27) \quad E_{r,l_2}(\Phi) = \int \|\nabla u\|^2 + \|\nabla v\|^2 d\underline{x}.$$

This is a natural generalization of Tikhonov regularization in image restoration. It is the same energy that Horn and Schunk [48] introduced in the optical flow setting.

- $\phi(s) = s$ corresponds to the regularization energy:

$$(3.28) \quad E_{r,l_1}(\Phi) = \int \sqrt{\|\nabla u\|^2 + \|\nabla v\|^2} d\underline{x},$$

which can be regarded as a rotationally invariant generalization of the total variation (TV) regularization for flow fields [129].

A Test Setup with Truncated Quadratic Regularizer

For simplicity, we consider mono-modality registration with L_2 norm as the data fidelity measure, *i.e.*,

$$E_d = \frac{1}{2} \int_{\Omega} (g(\underline{x}) - f(\underline{x} + \Phi(\underline{x})))^2,$$

and the corresponding variational descent flow is given by:

$$w_d(\underline{x}) = (g(\underline{x}) - f(\underline{x} + \Phi(\underline{x}))) \nabla f(\underline{x} + \Phi(\underline{x})).$$

For the preliminary test, we use a truncated quadratic [9] as the regularization function:

$$(3.29) \quad \phi(s, \alpha) = \begin{cases} \left(\frac{\alpha_0}{\alpha}\right)^2 s^2 & |s| \leq \alpha \\ \alpha_0^2 & \text{otherwise.} \end{cases}$$

The disadvantage and benefit of this choice are both clear. With strict ‘‘saturation’’ behavior above the scale parameter α , it poses a challenge for optimization. Graduated nonconvexification approaches can be utilized. On the other hand, this formulation provides nice theoretical interpretations. It is natural to introduce a line process [35] which is equivalent to ‘‘labeling’’ the outlier in the robust estimation setting [7].

Notice that (3.29) also provides a simple recipe to extract singularity set S of $|D\Phi|$ from the estimated Φ by thresholding at level α :

$$S = \{\underline{x} : |D\Phi(\underline{x})| > \alpha\}.$$

This may be useful for extracting motion interfaces.

To alleviate the local minima issue due to nonconvexity, we start with a large initial α . This is equivalent to use the conventional Tikhonov regularization (the vector version is more commonly known as Horn and Schunk in optical flow) of the form $E_r = |\nabla u|^2 + |\nabla v|^2$ as $S = \emptyset$ for α large enough. Then the scale parameter α is gradually decreased till the desired tolerance for discontinuity. To speed up the implementation, a multi-resolution scheme is applied.




Figure 3.8: Truncated quadratic regularization with varying scale.

Preliminary Results

We apply the setup described in Section Section 3.3.1 to two coronal CT slices obtained from deep inhale and exhale phases. Proposed regularization results in smooth deformation in homogeneous organ (lung, heart and exterior of rib-cage) and correctly preserves motion interfaces on the boundaries between the diaphragm, heart atria, rib cage and the lungs.






Figure 3.9: Registration comparison between Tikhonov (Tik) and Truncated quadratic (TQ) regularizations.



(a) quiver plot for Tik reg.



(a) quiver plot for TQ reg.

Figure 3.10: Comparison of deformation fields.

3.3.2 Discriminative Shear Preserving Regularizer

The problem of designing regularizer to accommodate sliding effects in medical registration has several distinct traits. First, we prefer to pose the problem in an optimization framework with a single energy functional and obtain the deformation as its optimal, rather than to use regularization to post-process some initial estimate or segmentation. Secondly, the deformation should be fairly smooth except at the sliding sites. Thirdly, dramatic local volume change seldom occurs in physical deformations; in particular, the deformation should neither create collision flow that maps different pixels to the same location (folding) nor generate vacuums. Similarly, within the complement of the sliding surface, shear should be fairly small. On the other hand, we should preserve the large shear at the sliding boundaries. This requires our method to differentiate among different types of discontinuities and regularize them accordingly. Finally, medical image registration involves tissues that are elastic with sliding motion, and we expect nonvanishing divergence and curl components from a physical deformation, so we are not interested in studying the extreme cases of pure solenoidal or irrotational flows. The above prior knowledge can be incorporated by devising a smooth regularization on the divergence component and a regularization on the curl component that preserves large-magnitude. Motivated by [136], we consider the following regularization energy:

$$(3.30) \quad E_{r,\text{divcurl}}(\Phi) = \int_{\Omega} \alpha \|\text{div } \Phi\|^2 + \beta \|\text{curl } \Phi\| d\mathbf{x}.$$

The regularization on curl Φ is reminiscent of total variation. It penalizes small curl values, yet is much more forgiving to large values than the quadratic form. In fact, similar to total variation, one could argue that the proposed functional is unbiased towards shears.

As a simple sanity check, we examine how the proposed regularization energy would drive an initial flow field. We derive the variational direction to descend $E_r(\Phi)$ and use it to evolve the flow. In the absence of data fidelity term, one tradeoff parameter is sufficient, and it is equivalent to descend the following function:

$$(3.31) \quad E_{r,\text{divcurl},2d} = \int_{\Omega} (u_x + v_y)^2 + \gamma |u_y - v_x| d\mathbf{x}.$$

Both divergence and curl operators are linear and invariant to rigid coordinate transformation, so it suffices to check typical cases by aligning the flow to one of the coordinate.

- If $\Phi = (u(x), 0)$, then a large value in u_x would indicate a jump along the direction of the flow, which would potentially causes folding or a vacuum. Locally, we would have $\text{div } \Phi = u_x \neq 0$, yet $\text{curl } \Phi = 0$. Penalizing the L_2 norm of $\text{div } \Phi$ as in (3.30) discourages large u_x values, thereby this helps prevent folding or vacuums.
- If $\Phi = (u(y), 0)$, then a large value in u_y would indicate the presence of shear along the flow which we want the regularization to preserve. In this case, $\text{div } \Phi = 0$ and $\text{curl } \Phi \neq 0$. Regularization with the L_1 norm of $\text{curl } f$ achieves the desired effect of allowing this type of shear.

In 3D, the deformation $\Phi = (u, v, w)$ is decomposed into its divergence and curl components as follows:

$$(3.32) \quad \text{div } \Phi = u_x + v_y + w_z;$$

$$(3.33) \quad \text{curl } \Phi = \det \begin{bmatrix} \mathbf{i} & \mathbf{j} & \mathbf{k} \\ \frac{\partial}{\partial x} & \frac{\partial}{\partial y} & \frac{\partial}{\partial z} \\ u & v & w \end{bmatrix},$$

where $\mathbf{i}, \mathbf{j}, \mathbf{k}$ are the unit vectors for the x -, y -, and z -axes, respectively.

Divergence is still a coordinate independent scalar field. The curl component, however, is a vector field with three coordinates $\Omega \rightarrow \mathfrak{R}^3$, and its direction is determined by the right hand rule. For the regularization function to be rotationally invariant, we use the L_2 length of the curl field pointwise and then integrate over the spatial coordinates, *i.e.*,

$$(3.34) \quad \begin{aligned} E_{r,\text{divcurl},3d}(\Phi) &= \int_{\Omega} (\text{div } \Phi)^2 + \gamma \|\text{curl } \Phi\|_2 d\mathbf{x} \\ &= \int_{\Omega} (u_x + v_y + w_z)^2 + \gamma \sqrt{(w_y - v_z)^2 + (u_z - w_x)^2 + (v_x - u_y)^2} d\mathbf{x}. \end{aligned}$$

Preliminary Experiment

To study the effect of proposed regularization, we first test it on an initial flow field. The flow was evolved along the energy descending direction, which is derived using variational calculus. We approximated the absolute value term with $|u_y - v_x| \approx \sqrt{(u_y - v_x)^2 + \epsilon}$ where ϵ is a small positive constant.

For image registration, the deformed source image needs to be interpolated. We use the fast B-spline interpolation scheme proposed by Unser *et al.* [119–121] with a 4-level multi-resolution structure [122]. The source image is represented as:

$$(3.35) \quad f(x, y) = \sum_{i, j} c_{i, j} \beta^n\left(\frac{x}{\Delta x} - i\right) \beta^n\left(\frac{y}{\Delta y} - j\right),$$

where β^n is the n th-order B-spline basis.

For computational efficiency, we also use linear combinations of tensored m th-order B-spline basis to express the deformation field:

$$(3.36) \quad \begin{aligned} u(\underline{x}) &= \sum_{i, j} d_{i, j}^u \beta^m\left(\frac{x}{m_x} - i\right) \beta^m\left(\frac{y}{m_y} - j\right); \\ v(\underline{x}) &= \sum_{i, j} d_{i, j}^v \beta^m\left(\frac{x}{m_x} - i\right) \beta^m\left(\frac{y}{m_y} - j\right). \end{aligned}$$

The finest level of B-spline deformation basis was chosen to be very narrowly supported (2 pixels) so that its interpolating behavior does not compromise discontinuity preservation within a given tolerance. It is straight forward to generalize the notations in (3.35) and (3.36) to higher dimensions thanks to the tensor structure of the B-spline basis adopted here.

• Regularizing Flow

In general, it is difficult to characterize the solution to a nonlinear registration problem. To study how the regularization energy would bias the registration results, we first examine its effect on a given flow. In particular, we are interested in checking whether it could prevent collision/folding and vacuum creation, yet preserve sliding (shear) discontinuities. Since the proposed penalty can be regarded as a combination of L_2 and L_1 regularization on the divergence and curl component respectively, we also compare with the results of Horn and Schunk (3.27) and total variation (3.28) regularization.

Fig. 3.11 and Fig. 3.12 test the regularization effects on colliding flow and vacuum generating flow respectively. Notice that L_1 regularization preserves large divergence, and is vulnerable to such flow; L_2 regularization successfully smooths the flow to prevent both collision and vacuum. The proposed regularization behaves like L_2 on the divergence, and enjoys similar robustness. Fig. 3.13 presents the results on a pure shear flow. Homogeneous L_2 regularization blurs the motion interface and is inferior to the L_1 regularization. The proposed method is effectively L_1 for curl and preserves big shears as illustrated. Fig. 3.14 provides an example where collision and sliding coexist. It is clear that the proposed method successfully alleviates the collision and preserves the sliding component.

• 2D Sliding Block Registration

In the first registration test, we simulated two blocks sliding against each other over lightly textured still background. As before, we compare the performance using L_1 , L_2 and the proposed regularization. The absence of a quantitative measure of performance has always been an issue in image registration, and a “fair” choice of parameter needs to be made when several models are to be compared. Since the weighting parameters in the energy functional control the tradeoff between data fidelity and regularization, varying their values could lead to very different estimates. Indeed, all models would coincide in the extreme case when the weight of the regularization is set to be zero, as the result is driven solely by data matching. In the absence of a rigorous way to choose the optimal parameters, it is only fair that we compare the models over a range of tradeoff parameter values. Fortunately, we have access to the ground truth deformation with the simulated data ⁴ One may argue that there are

⁴Due to background occlusion, there is no deformation that could match the target perfectly.




Figure 3.11: Regularization results for a colliding/folding flow. First row: (left to right) original image, transformed image, original (unregularized) flow. Second row: resulting flow under different regularizations. Third row: image transformed according to regularized flows.

infinitely many deformations that would generate the same source and target image pair in Fig. 3.15, yet the uniform (within each block) sliding is the most common and natural interpretation for physical motions. We expect this simple simulation to reasonably represent the major features of physical sliding in medical applications.

Fig. 3.16 quantifies the tradeoff between image similarity and regularization, and their effect on estimating the deformation field. For each method under comparison, we vary the tradeoff parameter, and plot the error of estimated deformation map v.s. intensity mismatch. The horizontal axis is the sum of squared difference of intensity values over the computation domain and the vertical axis reports the discrepancy between the estimated deformation and the ground truth flow, measured by sum of squared distance of the error vectors. We observe that for a wide range of data fitting error, the proposed method outperforms the other alternative choices in terms of real registration error. This makes the proposed method appealing. Unlike the oracle ground-truth, the data fitting metric is accessible in practice, and people often choose registration results (or regularization parameters) based on intensity error. Since the proposed regularizer corresponds to a better deformation estimate for the same intensity mismatch, it is a better choice given an fixed error budget in data matching. Furthermore, if one has access to a good parameter choice for all regularizers (or has the luxury of running a few trials and then somehow rate the results), the proposed regularizer has the best performance among all.

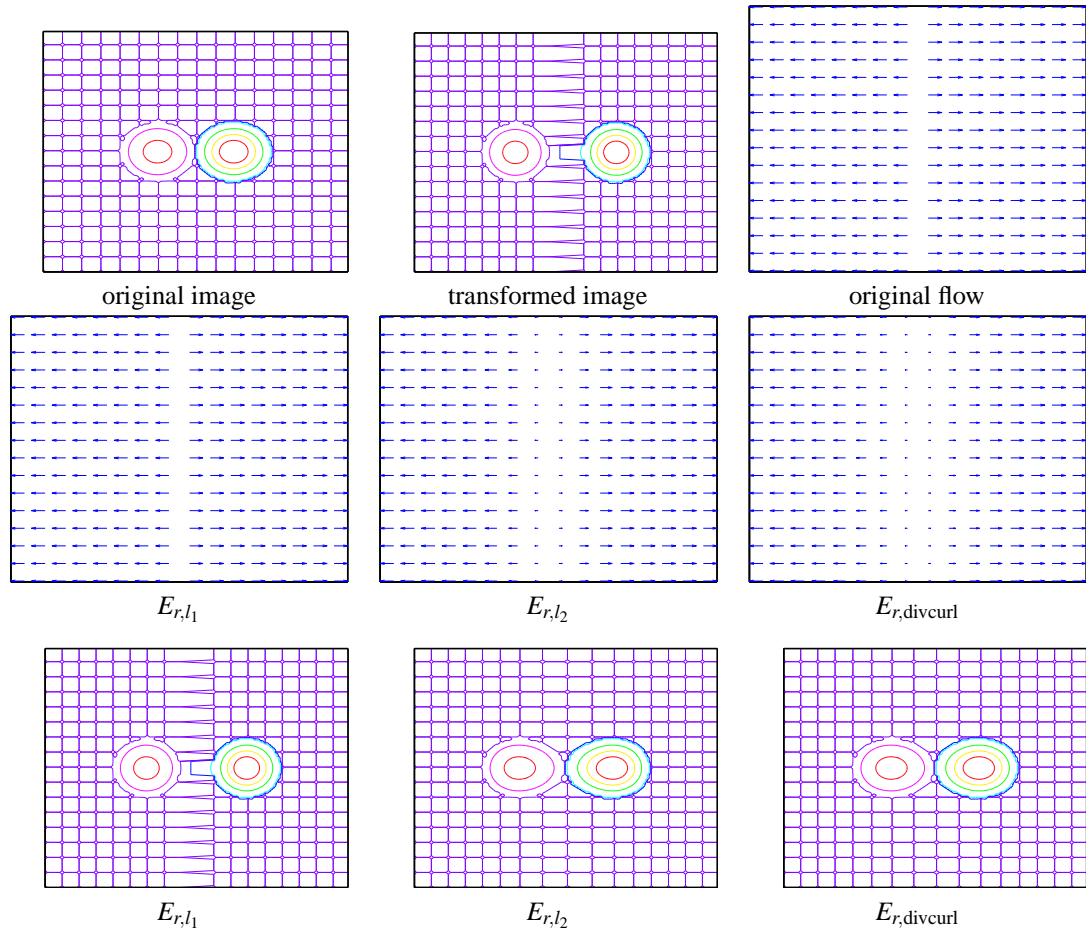


Figure 3.12: Regularization results for a vacuum generating flow. First row: (left to right) original image, transformed image, original (unregularized) flow. Second row: resulting flow under different regularizations. Third row: image transformed according to regularized flows.

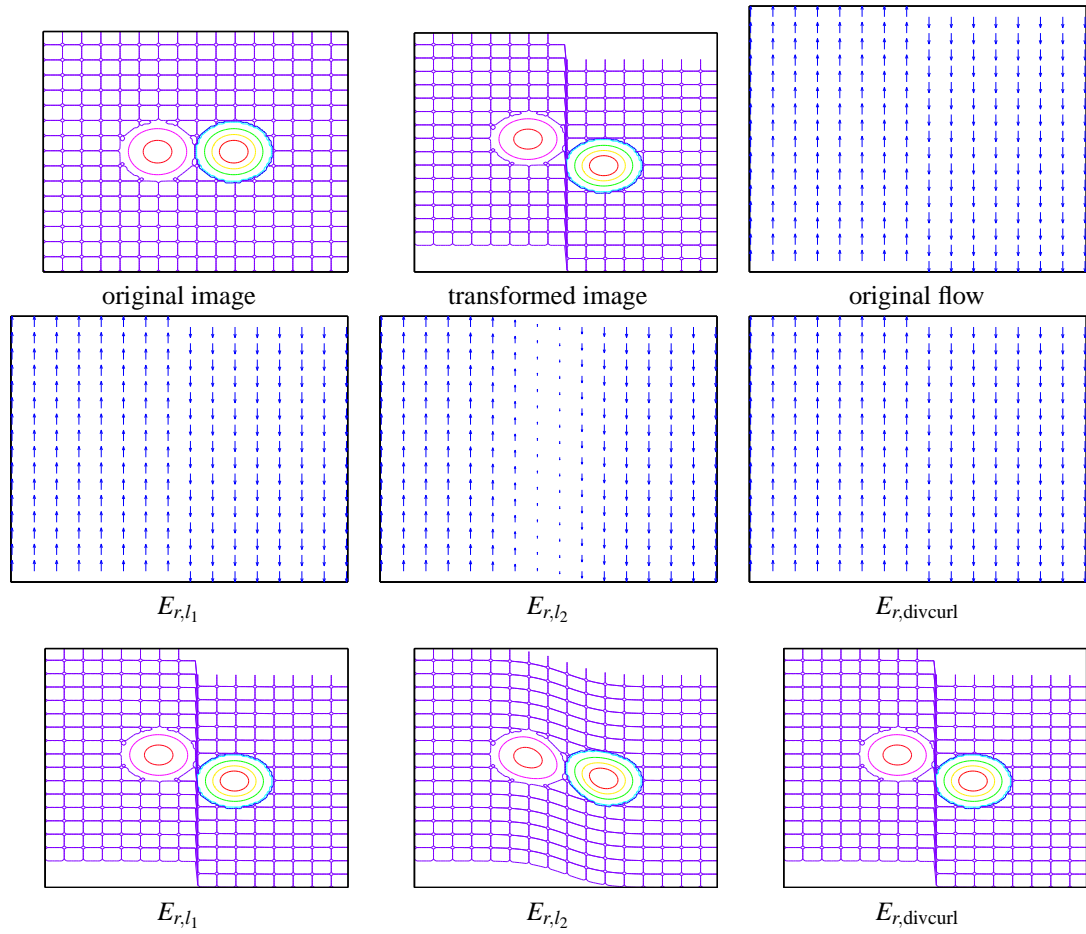


Figure 3.13: Regularization results for a sliding flow. First row: (left to right) original image, transformed image, original (unregularized) flow. Second row: resulting flow under different regularizations. Third row: image transformed according to regularized flows.

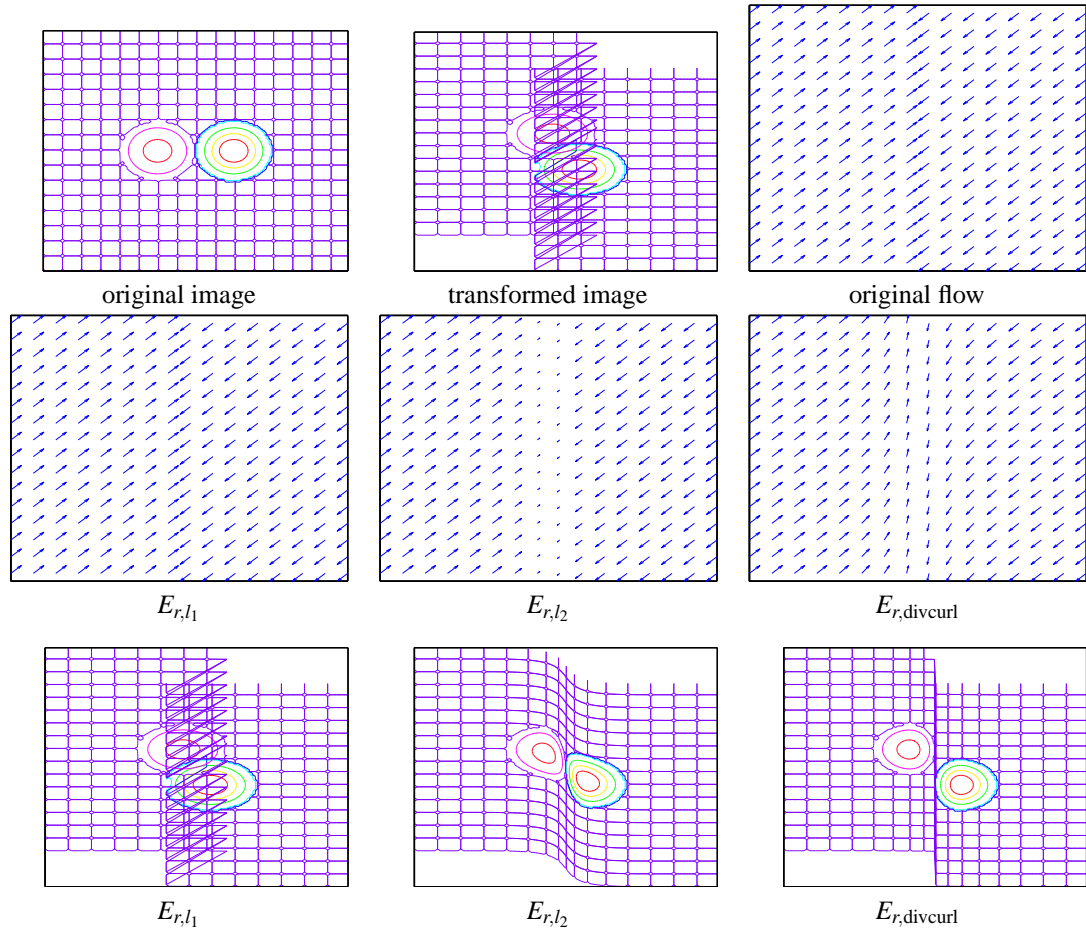


Figure 3.14: Regularization results for a flow with simultaneously significant divergence and curl components. First row: (left to right) original image, transformed image, original (unregularized) flow. Second row: resulting flow under different regularizations. Third row: image transformed according to regularized flows.




Figure 3.15: Simulated sliding blocks and the ground truth deformation.




Figure 3.16: Registration performance comparison: registration error v.s. intensity discrepancy.





Figure 3.17: Best estimation results (relative to the ground truth deformation) from various regularization methods. Column-wise (left to right): (X1) L_1 regularization; (X2) L_2 regularization; (X3) Proposed regularization. Row-wise (top to bottom): (a#) deformed template; (b#) quiver flow; (c #) vertical component of the deformation.

• 2D CT Image Registration

Sliding is widely observed along the boundary of the rib cage and internal organs. In this test, we obtained two breath-hold X-ray CT images of a real oncology patient, scanned at deep inhale (80% vital capacity) and exhale - a common procedure in radiation treatment planning. As a preliminary study, we applied the proposed regularization to register the pair of 2D coronal slices shown in Fig. 3.3.2. As in the sliding block experiment, we compare the proposed method with the L_1 and L_2 regularized results. With real clinical data, there is no ground truth, and the tradeoff parameters were chosen experimentally. For each regularization method, we ran the registration using several different parameter values, and picked the one that achieves a reasonable balance between data fidelity and physical feasibility via visual examination. Fig. 3.18 shows the “best” registration results of all regularization methods. With the chosen parameters, all three regularized registrations provide comparable intensity agreement between the deformed template and the target images. This suggests the fairness of later comparison, as the results can be interpreted as minimization of each regularization energy subject to the intensity match constraint. The warped grid maps and quiver plots [Fig. 3.18(c#) & (d#)] illustrate the advantage of the proposed regularization. In particular, the deformation on both interior and exterior of the rib cage are fairly smooth, and the motion boundary in between is preserved. Interestingly, the proposed regularization also naturally extracts the motion boundary between the lung and the mediastinum, which are affected differently by respiration. As breathing mainly induces motion along superior-inferior direction (vertical in our presentation), we examine that component closely in Fig. 3.19. Inside the thorax, registration result with the proposed regularization demonstrates good continuity, which agrees with the physical interaction between the lungs and diaphragms during respiration. Motion discontinuities are effectively represented along the rib cage and the surface of the mediastinum. Fig. 3.20, Fig. 3.21 and Fig. 3.22 illustrate the deformation field overlain on image intensity to reveal the agreement (or disagreement) of the estimated deformations with motion boundaries.



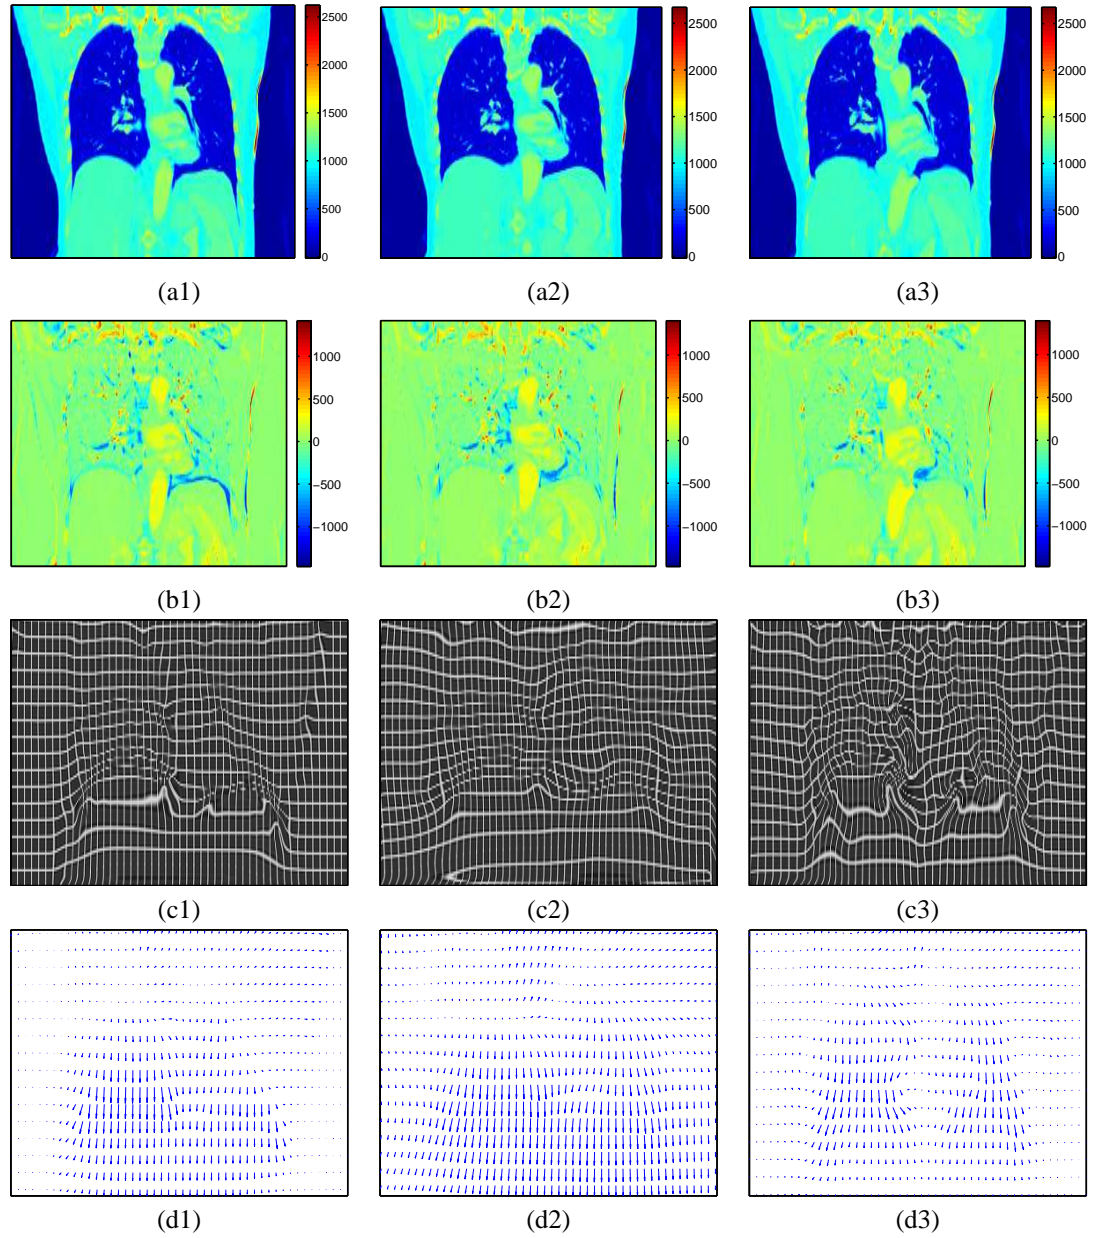


Figure 3.18: Registration results of CT data with various regularizations. Column-wise (left to right): (X1) L_1 regularization; (X2) L_2 regularization; (X3) Proposed regularization. Row-wise (top to bottom): (a#) deformed source; (b#) intensity difference between deformed source and the target images; (c #) warp grid representation of deformation; (d#) quiver plot of deformation.




Figure 3.19: Vertical component of the deformation from CT registration. (a) L_1 regularization; (b) L_2 regularization; (c) Proposed regularization.



Figure 3.20: Quiver plot overlain with image intensity for L_1 regularized image registration.



Figure 3.21: Quiver plot overlay with image intensity for L_2 regularized image registration.



Figure 3.22: Quiver plot overlay with image intensity for discriminately regularized image registration.

- **3D CT Image Registration**

Unfortunately, the proposed functional seems to be insufficient by itself in regularizing 3D registration. In the neighborhood of sliding, the unphysical rolling effects are widely observed. It is possible that this phenomenon is related to the Kelvin-Helmholtz instability for shear flows. Since our algorithm iteratively updates the deformation flow, it behaves as a physical process that evolves the shear flow. In this case, the shear surface is subject to the Kelvin-Helmholtz instability and any small perturbation in the normal direction of the shear surface incurs rolling; in fact, singularities occur in finite time. Also, this stability is more obvious in higher dimensions, which possibly explains the relative benign behavior in the 2D cases. As indicated by generalized fluid flow models, introducing material viscosity and surface tension may help prevent such turbulence. It is also possible that decomposing the overall deformation into a concatenation of small physical steps may alleviate this rolling artifacts, as smaller deformation requires relatively fewer iterations to resolve, and to reach a reasonable result before the singularities form.

Figure 3.23 and Figure 3.24 illustrate the rolling artifacts.

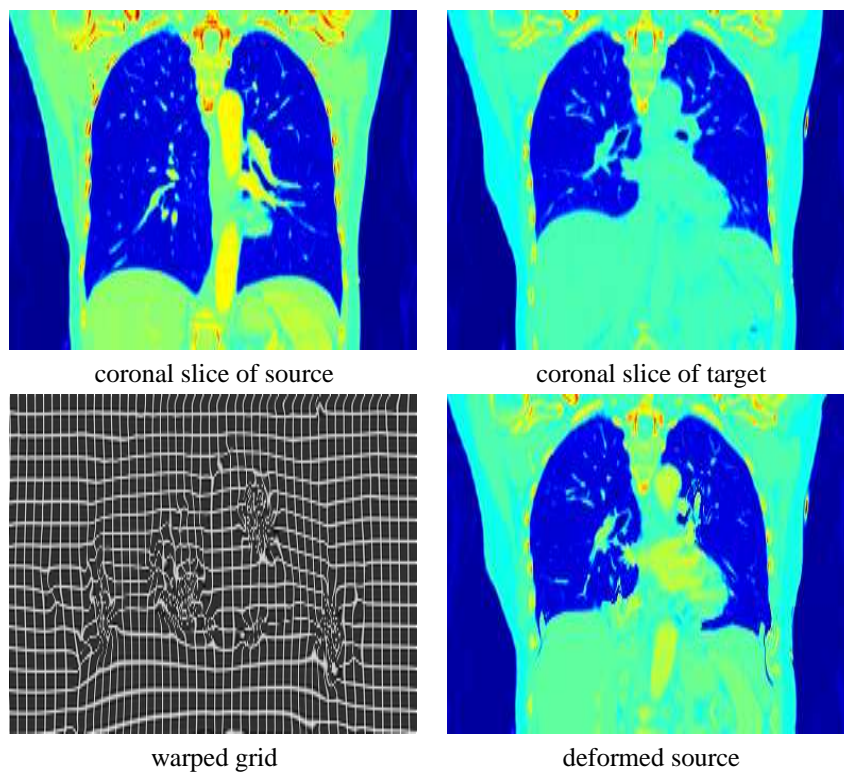


Figure 3.23: Coronal view for 3D discriminate registration. The same slice from source, target and deformed source volume. Since it is a full 3D registration, the deformed source pulls information in all (x, y, z) directions and the source slice is shown for reference purpose only.




Figure 3.24: Sagittal view for 3D discriminate registration. The same slice from source, target and deformed source volume. Since it is a full 3D registration, the deformed source pulls information in all (x, y, z) directions and the source slice is shown for reference purpose only.

3.3.3 Discussions

In Section 3.3.1, we first analyzed the conditions for a general class of regularizers of the form(3.19) in an axiomatic fashion, in the sense that beginning from the assumed behavior of a regularizer, we derived the consequence of these assumptions and the corresponding functional form. Further noticing the necessity of distinguishing among different types of singularities, namely, folding/vacuum v.s. shear, with the latter being the only physically admissible, we designed different regularizers for each component based on the Helmholtz decomposition. In particular, we have used an isotropic L_2 diffusion on the divergence component to enforce volume compatibility (no folding or gap), and adopted an L_1 regularizer on the curl component to preserve large shear.

In fact, we can substitute a more general regularization functional ϕ in place of the L_1 norm, and discuss the conditions on ϕ so that the regularized energy can preserve sliding, as follows

$$(3.37) \quad E_{r,gen}(\Phi) = \int_{\Omega} \alpha \|\operatorname{div} \Phi\|^2 + \beta \phi(\|\operatorname{curl} \Phi\|) d\mathbf{x}.$$

The influence of the regularization ϕ should be such that it penalizes weak curls, corresponding to a smoother deformation field, but preserves the curl if it is strong. We denote $\psi(s) \triangleq \phi'(s)/s$ for $s \in (0, \infty)$, and call it the ‘‘influence function’’ as before.

Mimicking the works in image restoration [3], it is easy to show that ϕ needs to satisfy the following conditions.

1. To suppress small curl values,

$$\phi'(0) = 0, \quad \lim_{s \rightarrow 0^+} \frac{\phi'(s)}{s} = \lim_{s \rightarrow 0^+} \phi''(s) = \phi''(0) > 0.$$

2. To preserve large shear,

$$\lim_{s \rightarrow +\infty} \phi''(s) = \lim_{s \rightarrow +\infty} \frac{\phi'(s)}{s} = 0 \quad \text{and} \quad \lim_{s \rightarrow +\infty} \frac{\phi''(s)}{\phi'(s)/s} = 0.$$

There are many functions that satisfy these properties. In this study, we investigated the truncated quadratic function, which was shown to be the discrete analogue of the Mumford-Shah functional [13]; the absolute value function, which corresponds to regularizing the L_1 norm of the curl component; and the Huber function, which can be regarded as the inf-convolution of L_1 and L_2 functionals [14].

Shear preserving regularization for three dimensional registration needs further investigation.

3.4 Equivalence Between Two L_2 Div- L_1 Curl Regularizations*

Section 3.3 considered discontinuity preserving image registration with energy of the form:

$$(3.38) \quad E(\Phi) = \int (f - g \circ \Phi)^2 dx + \lambda_{div} \int |\operatorname{div} \Phi|^2 dx + \lambda_{curl} \int |\operatorname{curl} \Phi| dx,$$

which behaves as L_2 regularizer on the divergence component to encourage homogeneous smoothness in volume change and L_1 regularizer on the curl component to preserve large shears.

It is immediate that another form has similar properties:

$$(3.39) \quad \tilde{E}(\Phi) = \int (f - g \circ \Phi)^2 dx + \lambda_{div} \sqrt{\int |\operatorname{div} \Phi|^2 dx} + \lambda_{curl} \int |\operatorname{curl} \Phi| dx.$$

One may argue that regularization in (3.39) is one-homogeneous in Φ so that the coefficients λ_{div} and λ_{curl} have the same units. On the other hand, the setup in (3.38) has nice point-wise structure, and direct optimization is easier. If E has a unique minimizer, then the equivalence between these two setups are trivially true; however, the data fidelity term in intensity matching image registration problems is nonconvex, and uniqueness of the minimizer cannot be established in general.

In what follows, we show an approximate equivalence between these two formulations. The main idea of the proof is to first transform the regularized optimization problems into the corresponding constrained version, and then establish equivalence in the constrained setup.

For simplicity, we use the following formulation for energy

$$(3.40) \quad \begin{aligned} E_\lambda(\Phi) &= \left(\int (f - g \circ \Phi)^2 dx + \alpha \int |\operatorname{curl} \Phi| dx \right) + \lambda \int |\operatorname{div} \Phi|^2 dx \\ &= (|f - g \circ \Phi|_2^2 + \alpha |\operatorname{curl} \Phi|_1) + \lambda |\operatorname{div} \Phi|_2^2. \end{aligned}$$

The equivalence between (3.40) and (3.38) is easily shown by identifying $\lambda = \lambda_{div}$ and $\alpha = 1/\lambda_{curl}$.

Given f, g and for a fixed α , we denote $A(\Phi) \triangleq |f - g \circ \Phi|_2^2 + \alpha |\operatorname{curl} \Phi|_1$ and $B(\Phi) \triangleq |\operatorname{div} \Phi|_2^2$, so that $E_\lambda(\Phi) = A(\Phi) + \lambda B(\Phi)$. Since E is not convex in general, there may exist multiple minimizers, and we denote the set of minimizers of E_λ as:

$$C_\lambda = \{\Phi : E(\Phi) \leq E(\Psi), \quad \forall \Psi \in \Gamma\},$$

and the corresponding energy value as $E_\lambda^* = \min E_\lambda(\Psi)$.

Claim III.4. For $\lambda_1 < \lambda_2$. If $C_1 \cap C_2 = \emptyset$, then $B(\Phi_1) > B(\Phi_2)$ for $\Phi_1 \in C_{\lambda_1}$ and $\Phi_2 \in C_{\lambda_2}$.

Proof. Assume not. Take $\Phi_1 \in C_{\lambda_1}$ and $\Phi_2 \in C_{\lambda_2}$ such that $B(\Phi_1) \leq B(\Phi_2)$, then

$$(3.41) \quad \begin{aligned} E_{\lambda_2}(\Phi_1) &= A(\Phi_1) + \lambda_2 B(\Phi_1) \\ &= A(\Phi_1) + \lambda_1 B(\Phi_1) + (\lambda_2 - \lambda_1) B(\Phi_1) \\ &= E_{\lambda_1}^* + (\lambda_2 - \lambda_1) B(\Phi_1) \\ &< E_{\lambda_1}(\Phi_2) + (\lambda_2 - \lambda_1) B(\Phi_2) \\ &= E_{\lambda_2}(\Phi_2). \end{aligned}$$

The third line follows from the fact that $\Phi_1 \in C_{\lambda_1}$; the fourth line from the assumption that $\Phi_2 \notin C_{\lambda_1}$ and $B(\Phi_1) \leq B(\Phi_2)$. This contradicts the assumption that Φ_2 minimizes E_{λ_2} .

In fact, as long as we assume Φ_1, Φ_2 do not simultaneously belong to $C_{\lambda_1} \cap C_{\lambda_2}$, the contradiction would hold.

We now argue that for $\lambda_1 \neq \lambda_2$, it is highly probable that $C_{\lambda_1} \cap C_{\lambda_2} = \emptyset$. The minimizer(s) for the energy $E_\lambda(\Phi) = A(\Phi) + \lambda B(\Phi)$ need to satisfy the Euler-Lagrange equation

$$dA(\Phi) + \lambda dB(\Phi) = 0.$$

If Φ^* simultaneously minimizes E_{λ_1} and E_{λ_2} , *i.e.*, $\Phi \in C_{\lambda_1} \cap C_{\lambda_2}$, then it must be true that $dB(\Phi^*) = 0$. It immediately follows that $dA(\Phi^*) = 0$ as well, so that Φ^* satisfies the Euler-Lagrange equation for all λ and that it is the minimizer for all E_λ . In particular, it is a minimizer for

$$E_{\lambda=0}(\Phi) = A(\Phi).$$

Meanwhile, $dB(\Phi^*) = 0$ implies $\operatorname{div} \Phi = 0$ and $B(\Phi) = 0$. This means that Φ^* is a divergence free minimizer of the energy $|f - g \circ \Phi|_2^2 + \alpha |\operatorname{curl} \Phi|_1$ (which does not penalize divergence at all!). This result also holds the other way around: if there exist a divergence-free element Φ in C_0 , then $\Phi \in C_\lambda \forall \lambda$. The contrapositive of the original statement says that if C_0 has no divergence-free elements, $C_{\lambda_1} \cap C_{\lambda_2} = \emptyset \forall \lambda_1 \neq \lambda_2$.

Now we are ready to state it as a theorem.

Theorem III.5. *If Φ^0 is a divergence free minimizer of $E_0 = |f - g \circ \Phi|_2^2 + \alpha |\operatorname{curl} \Phi|_1$, then it also minimizes $E_\lambda \forall \lambda$, and $E_\lambda^* = E_0^* = A(\Phi^0)$. If there is no such element, then $C_{\lambda_1} \cap C_{\lambda_2} = \emptyset$ for any $\lambda_1 \neq \lambda_2$. Moreover, for $\lambda_1 < \lambda_2$ and any $\Phi_1 \in C_{\lambda_1}$, $\Phi_2 \in C_{\lambda_2}$, $B(\Phi_1) > B(\Phi_2)$.*

Recall the definitions of outer normals in [28], which is closely related to sub-differential in functional minimization.

Definition III.6. Given $p, v \in \mathfrak{R}^N$, let $H(p, v)$ denote the closed half space

$$H(p, v) = \{x \in \mathfrak{R}^N : (x - p) \cdot v \leq 0\}.$$

Given a convex domain $\Omega \subset \mathfrak{R}^N$ and a point $p \in \partial\Omega$, the collection of outer normals to Ω at p is defined as:

$$N_\Omega(p) = \{v \in \mathfrak{R}^N : \Omega \subset H(p, v)\}.$$

Consider the 2-dimensional plane $(A(\Phi), B(\Phi))$, then for given λ , $E_\lambda = A(\Phi) + \lambda B(\Phi)$ is constant along lines of slope $-\frac{1}{\lambda}$. Assume the set $\Omega \triangleq \{(A(\Phi), B(\Phi))\}, \Phi \in \Gamma$ to be convex in \mathfrak{R}^2 (if not, we shall consider its convex hull for now), then we have the following lemma.

Lemma III.7. *$p \in \partial\Omega$ is a minimizer of E_λ if and only if $-(1, \lambda) \in N_\Omega(p)$.*

This is a direct application of the Karush-Kuhn-Tucker (KKT) condition: the gradient of the E_λ is $(1, \lambda)$. To require that ∇E_λ to “point away” from the feasible set is equivalent to the condition we stated with the definition of outer normals.

Furthermore, note that if $\partial\Omega$ is differentiable at p , then $N_\Omega(p)$ contains a single direction, so this is equivalent to the traditional requirement of $-\nabla f = -\nabla g$ where $\Omega = \{x : g(x) \leq 0\}$. If $\partial\Omega$ is differentiable everywhere, then any $p = (a, b) \in \partial\Omega$ minimizes $a + \lambda b$ for at most one λ .




Figure 3.25: Illustration for the three cases of feasible region: strictly convex and everywhere differentiable, nonstrictly convex and everywhere differentiable, nondifferentiable.

- **Case 1.** If Ω is strictly convex⁵, then the minimizer of E_λ for each λ corresponds to exactly one point $p(\lambda) = (a(\lambda), b(\lambda))$ on $\partial\Omega$, and $b(\lambda)$ is strictly monotonically decreasing as a function of λ .

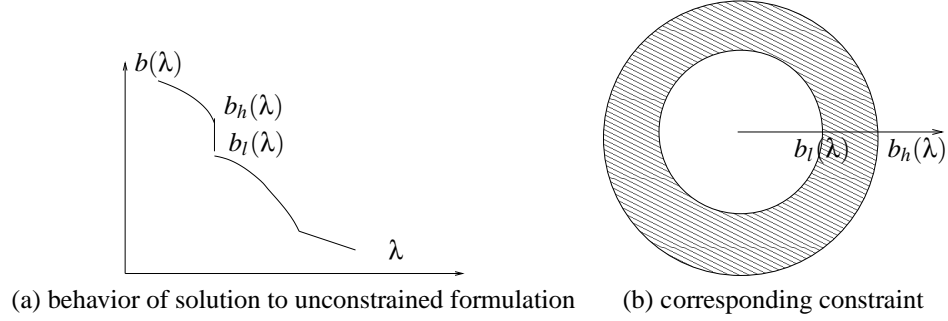


Figure 3.26: Equivalence between unconstrained and constrained formulation: Each λ in the unconstrained formulation maps to a constrained optimization problem $\min A(\Phi)$ over a circle (or ring).

- **Case 2.** If $\partial\Omega$ is everywhere differentiable, but Ω is NOT strictly convex. In particular, if $\partial\Omega$ contains line segments of $a + \lambda b = c_\lambda$ for $\lambda \in \Sigma$, then the minimizer of E_λ for each $\lambda \notin \Sigma$ corresponds to exactly one point $p(\lambda) = (a(\lambda), b(\lambda))$ on $\partial\Omega$; the minimizer of E_λ for each $\lambda \in \Sigma$ corresponds to a line segment on $\partial\Omega$, and the corresponding second coordinates $b(\lambda)$ form a continuum on \mathfrak{R} . In other words, $b(\lambda)$ (now a multiple valued function) has the property that if $\lambda_1 < \lambda_2$, then $b(\lambda_1) > b(\lambda_2)$; and if $b_1, b_2 \in b(\lambda)$, then $b \in b(\lambda)$ for any $b_1 \leq b \leq b_2$.
- **Case 3.** If there exists p where $\partial\Omega$ is nondifferentiable, and suppose that there exists $\lambda_1, \lambda_2 \in \mathfrak{R}^+$ such that $-(1, \lambda_1), -(1, \lambda_2) \in N_\Omega(p)$, then $p = (a, b)$ simultaneously minimizes E_{λ_1} and E_{λ_2} . Combine with our previous argument with the specific definition of E_λ , it must be true that p minimizes all E_λ . In other words, $\Omega \in H(p, v)$ for $v = -(\cos \theta, \sin \theta)$ for $\theta \in [0, \pi/2]$ (since we are only interested in $\lambda \in [0, \infty]$). This corresponds to the situation where a divergence-free registration arises as the minimizer to $|f - g \circ \Phi|_2^2 + \alpha |\text{curl } \Phi|_1$, which is highly unlikely in practice, we omit discussions about this case to avoid too much technicality.

In summary, if Ω is convex, then ⁶, $B(\Phi)$ is either a constant or a continuum for $\Phi \in C_\lambda$ for each fixed λ .

When $\partial\Omega$ is differentiable, $b(\lambda)$ is a strict monotone (potentially multi-valued) function of λ . For better presentation, we define $b_l(\lambda)$ and $b_h(\lambda)$ as the lower and upper bounds of $\{b : a(\lambda) + \lambda b = \min a + \lambda b, (a, b) \in \Omega\}$. Therefore, if $\Phi \in C_\lambda$, then $B(\Phi) \in [b_l(\lambda), b_h(\lambda)]$ and we have the following equivalence:

Theorem III.9. Φ minimizes $E_\lambda = A(\Phi) + \lambda B(\Phi)$ if and only if it solves

$$\begin{aligned} & \min A(\Phi) \\ & \text{s.t. } B(\Phi) = b, \text{ for some } b \in [b_l(\lambda), b_h(\lambda)] \end{aligned} .$$

Proof. Quite straight-forward from previous argument.

If Φ is a minimizer for E_λ , then take $b = B(\Phi)$, and it must be a minimizer for

$$\begin{aligned} & \min A(\Phi) \\ & \text{s.t. } B(\Phi) = b \end{aligned} .$$

On the other hand, if Φ minimizes

$$\begin{aligned} & \min A(\Phi) \\ & \text{s.t. } B(\Phi) = b \end{aligned} .$$

find the λ such that $b \in [b_l(\lambda), b_h(\lambda)]$, and Φ minimizes E_λ . □

Definition III.8. A set of points S in N -dimensional space is strictly convex if for every two points x_1 and x_2 belonging to S such that $x_1 \neq x_2$, the straight line segment joining the two points belongs to S but does not belong to the boundary of S (i.e. all the points on the interior of the straight line must be strictly in the interior of S): that is, a set S is strictly convex if and only if for every $x_1, x_2 \in S$, and scalar v such that $0 < v < 1$ we have $vx_1 + (1-v)x_2 \in$ the interior of S .

⁶With the exception of the rare case where $A(\Phi)$ has a divergence-free minimizer

Since $b_l(\lambda), b_h(\lambda)$ are positive quantities, the constrained setup can be rewritten as:

$$\begin{aligned} & \min A(\Phi) \\ & \text{s.t. } \sqrt{B(\Phi)}=b, \text{ for some } b \in [\sqrt{b_l(\lambda)}, \sqrt{b_h(\lambda)}] \end{aligned} .$$

Replicating all previous argument, we can show that this can be mapped to the unconstrained formulation

$$\min \tilde{E}_\gamma = A(\Phi) + \gamma \sqrt{B(\Phi)} = (\|f - g \circ \Phi\|_2^2 + \alpha |\text{curl } \Phi|_1) + \gamma \|\text{div } \Phi\|_2$$

for some γ .

When Ω is not convex, its convex hull will contain one or more line segments, corresponding to case 2. The optimum will be achieved at the end points of such line segments, which is the intersection between the original Ω and its convex hull. This will affect the results in that for some λ , $b(\lambda)$ may be the union of continuum (or possibly continua) and single value(s). For the constrained formulation, this corresponds to not solving the optimization if the constraint b takes on values in the complement of $\cup_\lambda b(\lambda)$. In that scenario, the minimizer (a, b) lies in the interior of the convex hull of Ω , and does not correspond to the minimizer of unconstrained problem $a + \lambda b$. This is O.K., as we are interested in showing the equivalence of the unconstrained formulations finally, and the missing portions of the constrained space does not contribute to the optimal solutions.

CHAPTER IV

Fundamental Performance Analysis in Image Registration Problems: Cramér-Rao Bound and its Variations

Image registration, as a special form of signal warping, is an important task in image processing. In contrast to the rapid development of algorithmic study in image registration, a standard performance evaluation tool is in general absent, except [95] where the transformation is assumed to be a global translation. It is important to investigate fundamental performance criteria in a principled manner to compare the overall optimality of different estimators for nonrigid registration problems. This chapter presents an observation model for image registration that accounts for image noise more realistically than most formulations, and describes performance analysis based on Cramér-Rao Bound(CRB) and its related variant Modified Cramér-Rao Bound(MCRB). We interpret the result of the commonly used optimization based registration as the M-estimate of the objective function and derive its bias-variance behavior.

4.1 Model - the Ideal v.s. Commonly Used

Generally, image registration methods aim to find the motion in an image sequence $\{z_i\}$, where z_i denotes the i th observation (frame) of an underlying image. In reality, only sampled observations are available, with spatial sample spacing Δ . Therefore, it is natural to use a discrete spatial index to refer to the sampled location. Without loss of generality, we take $z_i[n] = z_i^c(n\Delta)$ where z^c notates the underlying continuous intensity map. Accounting for additive observation noise, we formulate the generative model as:

$$(4.1) \quad z_i[n] = f(n + \tau_i(n)) + \varepsilon_i[n],$$

where it is standard to assume ε_i to be independent identically distributed (i.i.d) Gaussian noise. In principle, the task of registering the observation sequence is to find the deformation sequence of continuous maps $\{\tau_i\}$ for all i . We adopt the parametric setting, and represent the underlying continuous image intensity as a linear combination of a finite number of basis functions b with coefficients $c = \{c_k\}$, i.e., $f(x) = \sum_{k=1}^K c_k b(x, k)$. For simplicity, we focus on pairwise registration which requires estimating one deformation field τ , and drop the subindex in τ_i . Furthermore, we assume the deformation field is properly (sufficiently) parametrized with α , so the estimation performance for deformation and image intensity may be characterized by that of the parameter set (c, α) . For simplicity, we formulate our problem in one dimension, but the analysis generalizes to higher dimensions. The two observed images are modeled as:

$$(4.2) \quad \begin{aligned} z_1[n] &= \sum_{k=1}^K c_k b(n, k) + \varepsilon_1[n], \\ z_2[n] &= \sum_{k=1}^K c_k b(n + \tau_\alpha(n), k) + \varepsilon_2[n] \quad n = 1, 2, \dots, N, \end{aligned}$$

where $\{b(\cdot, k)\}$ are common intensity bases, and τ parametrized by α captures the pointwise deformation. The components of additive noise ε_i are zero mean I.I.D Gaussian with variance σ^2 .

The formulation in (4.2) captures the spatial sampling of the observation, the finite representation of the underlying “true” intensity $\{c_k\}_{k=1}^K$ and the dense pointwise deformation τ .

For comparison purposes, we formulate the estimator in traditional registration setup as the optimal transformation $\hat{\Gamma}$ such that ¹

$$(4.3) \quad \hat{\Gamma} = \arg \min_{\Gamma} D(z_2, z_1 \circ \Gamma),$$

where D is some difference measure, *e.g.*, sum-of-squared-difference (SSD) or mutual information (MI), and Γ indicates the transformation. In this setting, it is implicitly assumed that z_1 (also known as the “source”) is a noise-free version of the true intensity image f , and z_2 (also called the “target”) is a deformed image whose noise properties determines the proper choice of the difference metric. Clearly there is a lack of symmetry regarding the presence of noise in this formulation.

For simplicity, we use sum-of-squared-difference (SSD) as our default choice of the error metric D for (4.3) hereafter, corresponding to the Gaussian noise assumption, as adopted in many practical cases.

4.2 Cramér-Rao Bound and its Asymptotic Behavior

We first reformulate (4.2) in a compact vector form as follows.

$$(4.4) \quad \mathbf{z} = \begin{bmatrix} z_1 \\ z_2 \end{bmatrix} = \begin{bmatrix} A_0 \\ A_\tau \end{bmatrix} \mathbf{c} + \begin{bmatrix} \boldsymbol{\varepsilon}_1 \\ \boldsymbol{\varepsilon}_2 \end{bmatrix} = \mathcal{A}_\tau \mathbf{c} + \boldsymbol{\varepsilon},$$

where $\mathbf{z} \triangleq [z_1(1), \dots, z_1(N), z_2(1), \dots, z_2(N)]^T \in \mathfrak{R}^{2N}$ and $\mathbf{c} \triangleq [c_1, \dots, c_K]^T \in \mathfrak{R}_{\geq 0}^K$ are column vectors by stacking the corresponding elements. The concatenated random noise vector $\boldsymbol{\varepsilon} \sim \mathcal{N}(0, \boldsymbol{\Sigma} = \sigma^2 I_{2N})$. $A_0, A_\tau \in \mathbf{M}^{N \times K}$ have elements $A_0(i, j) = b(i, j)$ and $A_\tau(i, j) = b(i + \tau(i), j)$ for $i = 1, 2, \dots, N, j = 1, 2, \dots, K$. The overall system matrix $\mathcal{A} = [A_0^T, A_\tau^T]^T$. The Cramér-Rao Bound (CRB) is a fundamental lower bound on the variance of any unbiased estimator [123] and serves as a benchmark for estimator performance. When maximum-likelihood (ML) estimators are applied, which are known to be asymptotically unbiased, it is often useful to bound their variance with CRB. In [45], it is suggested that when inverting the Fisher information matrix (FIM) corresponding to the parameter of interest only is not straight-forward, it is feasible to use “complete-parameter” Fisher information matrices. Following a similar logic, we can write (4.4) in a more general form,

$$(4.5) \quad \begin{aligned} \mathbf{z} &= h(\boldsymbol{\tau}_\alpha, \mathbf{c}) + \boldsymbol{\varepsilon} \\ &= h(\boldsymbol{\theta}) + \boldsymbol{\varepsilon}, \end{aligned}$$

where $h(\boldsymbol{\tau}_\alpha, \mathbf{c}) \triangleq \mathcal{A} \mathbf{c}$ and $\boldsymbol{\theta} = [\boldsymbol{\alpha}, \mathbf{c}]$ denotes the “complete-parameter” vector. It follows immediately from the i.i.d Gaussian assumption of noise $\boldsymbol{\varepsilon}$ that the ML estimator $\hat{\boldsymbol{\theta}}_{ML}$ minimizes the L_2 distance between observation \mathbf{z} and system response $h(\boldsymbol{\theta})$ as follows:

$$\hat{\boldsymbol{\theta}}_{ML} = \arg \min_{\boldsymbol{\theta}} \|\mathbf{z} - h(\boldsymbol{\theta})\|_2.$$

Before we delve into the detailed computation, we clarify our goal and the structure of FIM here. We are ultimately interested in the performance of estimators for the deformation parameter $\boldsymbol{\alpha}$, and the image intensity parameter \mathbf{c} is chosen to augment the data to simplify expression. With $\boldsymbol{\theta} \in \mathfrak{R}^{N+K}$, the FIM corresponding to $\boldsymbol{\theta}$ takes on the form:

$$F(\boldsymbol{\theta}^*) = E_{\mathbf{z}|\boldsymbol{\theta}=\boldsymbol{\theta}^*} \left\{ -\frac{\partial^2}{\partial \boldsymbol{\theta}^2} \Lambda(\mathbf{z}|\boldsymbol{\theta})|_{\boldsymbol{\theta}=\boldsymbol{\theta}^*} \right\},$$

¹There is a slight abuse of notation here. The more precise formulation would be: $\hat{\Gamma} = \arg \min_{\Gamma} D(z_2, P(z_1^c \circ \Gamma))$, where z_1^c is the underlying intensity map that agrees with z_1 on sampling grids, and P is the sampling function such that $P(z^c)(n) = z^c(n\Delta)$. Even so, the cost function is still incomplete, as only z_1 is observed and the interpolator $I : z_1 \rightarrow z_1^c$ needs to be specified. The de facto objective function is thus $D(z_2, P(I(z_1) \circ \Gamma))$.

where Λ is the log-likelihood function $\Lambda(z|\theta) \triangleq \log f(z|\theta)$.

Moreover, if we define $J_{x,y} = E \left\{ \left[\frac{\partial}{\partial x} \Lambda(z) \right]^T \left[\frac{\partial}{\partial y} \Lambda(z) \right] \right\}$, then the complete-data FIM can be decomposed into block form as:

$$(4.6) \quad F_{\theta} = \begin{bmatrix} J_{\alpha,\alpha} & J_{\alpha,c} \\ J_{c,\alpha} & J_{c,c} \end{bmatrix}.$$

The sub-block $J_{\tau,\tau}$ is the FIM with respect to the quantity of interest - the deformation parameters. As CRB is the inverse of the FIM, we can invoke the formula for partitioned-matrix inverse [39] to obtain:

$$(4.7) \quad \begin{aligned} \text{CRB}(\alpha) &= [J_{\alpha,\alpha} - J_{\alpha,c} J_{c,c}^{-1} J_{c,\alpha}]^{-1} \\ \text{CRB}(c) &= [J_{c,c} - J_{c,\alpha} J_{\alpha,\alpha}^{-1} J_{\alpha,c}]^{-1}. \end{aligned}$$

This form can be further simplified using its symmetry - a fact that we will utilize later in our computation.

The likelihood function with respect to θ is :

$$f(z; \theta) = \frac{1}{(2\pi)^{2N/2} |\Sigma|^{1/2}} \exp\left(-\frac{1}{2} e^T \Sigma^{-1} e\right),$$

where $e = z - h(\theta) = z - \mathcal{A}(\tau_{\alpha})c$.

The log-likelihood turns out to be:

$$(4.8) \quad \begin{aligned} \Lambda &= \log f(z; \theta) \\ &= -N \log(2\pi) - 2N \log \sigma - \frac{1}{2\sigma^2} \|z - \mathcal{A}(\tau_{\alpha})c\|^2. \end{aligned}$$

Now we compute each term of the FIM.

$$(4.9) \quad \begin{aligned} \nabla_{\tau} \Lambda &= -\frac{1}{2\sigma^2} \nabla_{\tau} \|z - \mathcal{A}(\tau_{\alpha})c\|^2 \\ &= \frac{1}{\sigma^2} (z - \mathcal{A}(\tau_{\alpha})c)^T \nabla_{\tau} (\mathcal{A}(\tau_{\alpha})c) \end{aligned}$$

Notice that

$$(4.10) \quad \begin{aligned} \frac{\partial}{\partial \tau(l)} \{A_{\tau}[n, :]c\} &= \frac{\partial}{\partial \tau(l)} \sum_{k=1}^K c(k) A_{\tau}(n, k) \\ &= \frac{\partial}{\partial \tau(l)} \sum_{k=1}^K c(k) b(n + \tau(n), k) \\ &= \begin{cases} \sum_{k=1}^K c(k) \dot{b}(n + \tau(n), k), & l = n; \\ 0 & \text{else,} \end{cases} \end{aligned}$$

where $\dot{b}(\cdot, \cdot)$ denotes the derivative of $b(\cdot, \cdot)$ with respect to the first variable.

Plugging (4.10) into (4.9), we obtain

$$(4.11) \quad \frac{\partial}{\partial \tau(l)} \Lambda = \frac{1}{\sigma^2} (z_2(l) - A_{\tau}[l, :]c) \sum_{k=1}^K \dot{b}(l + \tau(l), k) c(k).$$

Therefore, the gradient of Λ with respect to τ is:

$$(4.12) \quad \frac{\partial}{\partial \tau} \Lambda = \frac{1}{\sigma^2} [(z_2 - A_{\tau}c) \odot (Dc)]^T = \frac{1}{\sigma^2} [\text{diag}\{Dc\} (z_2 - A_{\tau}(c))]^T,$$

where D is the matrix whose elements are $D(i, j) = \dot{b}(i + \tau(i), j)$, $1 \leq i \leq N$, $1 \leq j \leq K$, and “ \odot ” denotes the Schur/Hadamard product.

By chain rule, the gradient of Λ with respect to α is given by:

$$\begin{aligned} \nabla_{\alpha}\Lambda &= \frac{\partial}{\partial\tau}\Lambda\frac{\partial}{\partial\alpha}\tau \\ (4.13) \quad &= \frac{1}{\sigma^2}[(z_2 - A_{\tau}\mathbf{c}) \odot (D\mathbf{c})]^T \left[\frac{\partial}{\partial\alpha}\tau \right], \end{aligned}$$

where $\frac{\partial}{\partial\alpha}\tau \in \mathbf{M}^{N \times L}$ is the derivative matrix with element $[\frac{\partial}{\partial\alpha}\tau](i, j) = \frac{\partial}{\partial\alpha(j)}\tau(i)$, and L corresponds to the length of the deformation parameter α .

Now we compute the FIM $J_{\alpha, \alpha}$ with

$$\begin{aligned} E \left\{ \frac{\partial^2}{\partial\alpha^2}\Lambda \right\} &= -E \left\{ \left[\frac{\partial}{\partial\alpha}\Lambda \right]^T \left[\frac{\partial}{\partial\alpha}\Lambda \right] \right\} \\ (4.14) \quad &= -\frac{d\tau^T}{d\alpha} E \left\{ \left[\frac{\partial}{\partial\tau}\Lambda \right]^T \left[\frac{\partial}{\partial\tau}\Lambda \right] \right\} \frac{d\tau}{d\alpha} \\ &= -\frac{1}{\sigma^2} \frac{d\tau^T}{d\alpha} E \left\{ \text{diag}\{D\mathbf{c}\} (z_2 - A_{\tau}\mathbf{c}) \left[\frac{1}{\sigma^2} (z_2 - A_{\tau}\mathbf{c})^T \text{diag}\{D\mathbf{c}\} \right]^T \right\} \frac{d\tau}{d\alpha} \\ (4.15) \quad &= -\frac{1}{\sigma^2} \frac{d\tau^T}{d\alpha} \text{diag}^2\{D\mathbf{c}\} \frac{d\tau}{d\alpha} \end{aligned}$$

To calculate $J_{\mathbf{c}, \tau}$ and $J_{\mathbf{c}, \mathbf{c}}$, we take the derivative of Λ with respect to \mathbf{c} :

$$\begin{aligned} \frac{\partial}{\partial\mathbf{c}}\Lambda &= -\frac{1}{2\sigma^2} \frac{\partial}{\partial\mathbf{c}} \|z - \mathcal{A}\mathbf{c}\|^2 \\ (4.16) \quad &= \frac{1}{\sigma^2} (z - \mathcal{A}\mathbf{c})^T \mathcal{A}. \end{aligned}$$

It is now straight forward to compute the entries for the complete FIM:

$$(4.17) \quad E \left\{ \frac{\partial^2}{\partial\mathbf{c}\partial\mathbf{c}}\Lambda \right\} = -\frac{1}{\sigma^2} \mathcal{A}^T \mathcal{A}$$

$$\begin{aligned} E \left\{ \frac{\partial^2}{\partial\tau(l)\partial\mathbf{c}(m)}\Lambda \right\} &= \frac{1}{\sigma^2} E \left\{ -A_{\tau}[l, m] D[l, :] \mathbf{c} + \varepsilon_2(l) D[l, m] \right\} \\ (4.18) \quad &= -\frac{1}{\sigma^2} A_{\tau}[l, m] D[l, :] \mathbf{c}. \end{aligned}$$

The matrix $J_{\alpha, \mathbf{c}}$ can be represented in compact form as:

$$(4.19) \quad E \left\{ \frac{\partial^2}{\partial\alpha\partial\mathbf{c}}\Lambda \right\} = -\frac{1}{\sigma^2} \frac{d\tau^T}{d\alpha} \text{diag}\{D\mathbf{c}\} A_{\tau}.$$

With symmetry, the complete FIM is obtained:

$$(4.20) \quad F_{\theta} = \frac{1}{\sigma^2} \begin{bmatrix} \frac{d\tau^T}{d\alpha} \text{diag}^2\{D\mathbf{c}\} \frac{d\tau}{d\alpha} & \frac{d\tau^T}{d\alpha} \text{diag}\{D\mathbf{c}\} A_{\tau} \\ A_{\tau}^T \text{diag}\{D\mathbf{c}\} \frac{d\tau}{d\alpha} & \mathcal{A}^T \mathcal{A} \end{bmatrix}.$$

As a special case, when τ is parametrized with rect functions, *i.e.*, $\tau(n) = \alpha[n]$, we have $\frac{d\tau}{d\alpha} = I$. The FIM for (τ, \mathbf{c}) is then given by:

$$(4.21) \quad F_{(\tau, \mathbf{c})} = \frac{1}{\sigma^2} \begin{bmatrix} \text{diag}^2\{D\mathbf{c}\} & \text{diag}\{D\mathbf{c}\} A_{\tau} \\ A_{\tau}^T \text{diag}\{D\mathbf{c}\} & \mathcal{A}^T \mathcal{A} \end{bmatrix}.$$

At this point, we make the following observations:

1. With the commonly used model (4.3), it is assumed that the observed source image z_1 corresponds to the ground truth \mathbf{c} . In other words, most existing methods solve for the ML estimator τ with the generative model:

$$(4.22) \quad z_2 = \sum_{k=1}^K c_k b(n + \tau(n), k) + \varepsilon_2(n),$$

by plugging in the c_k 's that best fits z_1 . It is easy to derive the CRB for the log-likelihood function $\Lambda^{\text{com}}(z_2; \tau) = -N/2 \log(2\pi) - N \log \sigma - \frac{1}{2\sigma^2} \|z_2 - A_\tau \mathbf{c}\|^2$. The FIM matrix $F_\tau^{\text{com}} = J_{\tau, \tau}$ as we derived in (4.15). Therefore, $\text{CRB}^{\text{com}}(\tau) = J_{\tau, \tau}^{-1}$. Notice that as $J_{\tau, \mathbf{c}} J_{\mathbf{c}, \mathbf{c}}^{-1} J_{\mathbf{c}, \tau} \geq 0$ ², $\text{CRB}^{\text{com}}(\tau) \leq \text{CRB}(\tau)$ as extra information (known $\{c_k\}$) is assumed in the case of (4.22). In other words, the plug-in operation provides a ‘‘looser’’ bound for the variance than the ‘‘true’’ CRB corresponding to model (4.2).

2. For asymptotically large SNR, *i.e.*, $\sigma^2 \rightarrow 0$, we do expect a decent estimate of \mathbf{c} directly from the source image, assuming no model mismatch in the generative basis. In this case, the plug-in estimator as used in the traditional model, even though not a true ML estimator, is expected to perform similarly to the real ML estimator. Indeed, [76] shows that the ‘‘fake’’ bound approximates the true CRB³.
3. The above points may be interpreted better with a slight modification of the model in (4.2). Instead of i.i.d noise, we may assume that noise level in the two images are not symmetric, more specifically, we assume $\varepsilon_1 \sim \mathcal{N}(0, \sigma_1^2 I_N)$ and $\varepsilon_2 \sim \mathcal{N}(0, \sigma_2^2 I_N)$.

The log-likelihood is given by:

$$(4.23) \quad \Lambda = -\frac{1}{2\sigma_1^2} \|z_1 - A_0 \mathbf{c}\|_2^2 - \frac{1}{2\sigma_2^2} \|z_2 - A_\tau \mathbf{c}\|_2^2 + \text{some constant}.$$

The partial derivatives of the log-likelihood with respect to τ (thus α) is not affected by target image model, and the second-order derivative the log-likelihood with respect to \mathbf{c} is given by:

$$E \left\{ \frac{\partial^2}{\partial \mathbf{c} \partial \mathbf{c}} \Lambda \right\} = -\frac{1}{\sigma_1^2} A_0^T A_0 - \frac{1}{\sigma_2^2} A_\tau^T A_\tau.$$

We thus obtain the complete FIM with respect to (τ, \mathbf{c}) as:

$$(4.24) \quad F_{(\tau, \mathbf{c})} = \begin{bmatrix} \frac{1}{\sigma_2^2} \text{diag}^2\{D\mathbf{c}\} & \frac{1}{\sigma_2^2} \text{diag}\{D\mathbf{c}\} A_\tau \\ \frac{1}{\sigma_2^2} A_\tau^T \text{diag}\{D\mathbf{c}\} & \frac{1}{\sigma_1^2} A_0^T A_0 + \frac{1}{\sigma_2^2} A_\tau^T A_\tau \end{bmatrix}.$$

When $\sigma_1 \rightarrow 0$, corresponding to high SNR in the template image, then $J_{\mathbf{c}, \mathbf{c}} \rightarrow \infty$ and

$$\text{CRB}(\tau) = [J_{\tau, \tau} - J_{\tau, \mathbf{c}} J_{\mathbf{c}, \mathbf{c}}^{-1} J_{\mathbf{c}, \tau}]^{-1} \rightarrow J_{\tau, \tau}^{-1},$$

which reduces to the CRB^{com} .

4. To compute $\text{CRB}(\tau)$ exactly could be challenging, as $\mathcal{A}^T \mathcal{A}$ may not be easy to invert for arbitrary τ . Notice that the sub-matrix A_0 of \mathcal{A} has nice shift-invariant structure, yet A_τ depends on the deformation. In special cases, such as when the whole image (signal) experience uniform transformation $\tau(i) = \text{const}$ for $i = 1, 2, \dots, N$, then $J_{\mathbf{c}, \mathbf{c}}$ is block-shift-invariant, and efficient inversion is possible.
5. As a special case, we consider when the whole image experiences uniform transformation, where a natural parametrization is to use α to describe the global transformation, *i.e.*, $\tau_\alpha(i) = \alpha$ for $\forall i$.

²In most cases, we assume $J_{\tau, \mathbf{c}} J_{\mathbf{c}, \mathbf{c}}^{-1} J_{\mathbf{c}, \tau}$ to be nonsingular, so it is in fact positive definite.

³In particular, the parameter of interest τ is decoupled from the nuisance parameter \mathbf{c} in this case, and the asymptotic behavior of the bound can be shown with ease.

Under the uniform transformation assumption, we have

$$\frac{d\boldsymbol{\tau}}{d\boldsymbol{\alpha}} = \mathbf{1},$$

where $\mathbf{1}$ indicates a column vector (of length N in our case) with all unity elements.

Substituting this relation into (4.15), (4.19) respectively and we obtain:

$$\begin{aligned} F_{\theta} &= \frac{1}{\sigma^2} \begin{bmatrix} \mathbf{1}^T \text{diag}^2\{D\mathbf{c}\} \mathbf{1} & \mathbf{1}^T \text{diag}\{D\mathbf{c}\} A_{\tau} \\ A_{\tau}^T \text{diag}\{D\mathbf{c}\} \mathbf{1} & \mathcal{A}^T \mathcal{A} \end{bmatrix} \\ (4.25) \quad &= \frac{1}{\sigma^2} \begin{bmatrix} \mathbf{c}^T D D^T \mathbf{c} & [D\mathbf{c}]^T A_{\tau} \\ A_{\tau}^T [D\mathbf{c}] & \mathcal{A}^T \mathcal{A} \end{bmatrix}. \end{aligned}$$

4.3 Relating to MCRB

The modified Cramér-Rao Bound(MCRB) was first introduced [22] to resolve the synchronization issues in decoding systems. Rather than seeking the variance around the estimator for the “true” augmented data (“complete data”) which includes both the quantity of interest and the nuisance parameters \mathbf{c} , MCRB choose to look on the other parameters as “unwanted”. Instead of using the true CRB, the MCRB may be regarded as an approximation via “marginalizing” over the nuisance parameters. In fact, MCRB is always lower than CRB, thus a looser bound. In some cases, MCRB approaches the true CRB [76].

The central idea is the following. Instead of computing the true FIM

$$F = E_z \left\{ \left[\frac{\partial}{\partial \boldsymbol{\tau}} \log f(z; \boldsymbol{\tau}) \right]^2 \right\},$$

it uses

$$(4.26) \quad E_{z,\mathbf{c}} \left\{ \left[\frac{\partial}{\partial \boldsymbol{\tau}} \log f(z; \boldsymbol{\tau}, \mathbf{c}) \right]^2 \right\}.$$

The rationale for MCRB is the following:

$$\begin{aligned} E_{z,\mathbf{c}} \{ [\hat{\boldsymbol{\tau}}(z) - \boldsymbol{\tau}]^2 \} &= E_{\mathbf{c}} \{ E_{z|\mathbf{c}} [(\hat{\boldsymbol{\tau}}(z) - \boldsymbol{\tau})^2] \} \\ &\geq E_{\mathbf{c}} \left\{ \frac{1}{E_{z|\mathbf{c}} \left[\left(\frac{\partial}{\partial \boldsymbol{\tau}} \log f(z; \boldsymbol{\tau}, \mathbf{c}) \right)^2 \right]} \right\} \\ &\geq \frac{1}{E_{\mathbf{c}} \left\{ E_{z|\mathbf{c}} \left[\left(\frac{\partial}{\partial \boldsymbol{\tau}} \log f(z; \boldsymbol{\tau}, \mathbf{c}) \right)^2 \right] \right\}} \\ (4.27) \quad &= \frac{1}{E_{z,\mathbf{c}} \left\{ \left[\frac{\partial}{\partial \boldsymbol{\tau}} \log f(z; \boldsymbol{\tau}, \mathbf{c}) \right]^2 \right\}}. \end{aligned}$$

The first inequality comes from the application of CRB to the estimator $\hat{\boldsymbol{\tau}}(z)$ for a fixed \mathbf{c} and second is Jensen’s inequality.

4.4 An Alternating Minimization Algorithm

For registration purposes, we want to minimize the negative log-likelihood in (4.23). We adopt the frequentist perspective and consider the the underlying image intensity f (and thus \mathbf{c}) as fixed unknown. It is natural to ask for the solution of the augmented problem:

$$(\hat{\boldsymbol{\tau}}, \hat{\mathbf{c}}) = \arg \min_{\boldsymbol{\tau}, \mathbf{c}} -\Lambda.$$

We describe an alternating minimization algorithm to solve this problem as follows.

We make the following remarks:

Algorithm 1 Alternating minimization of the negative log-likelihood in (4.23).

- 1: Initialize \hat{c}
- 2: **repeat**
- 3: For given $c = \hat{c}$, minimize $\|z_2 - A_\tau c\|_2$ over τ . This step coincides with conventional registration methods by assuming c known. Obtain $\hat{\tau}$.
- 4: For given $\tau = \hat{\tau}$, minimize $\frac{1}{2\sigma_1^2} \|z_1 - A_0 c\|_2^2 + \frac{1}{2\sigma_2^2} \|z_2 - A_\tau c\|_2^2$. This is a typical quadratic minimization problem, and the solution is given by:

$$(4.28) \quad \hat{c} = \left[\frac{1}{\sigma_1^2} A_0^T A_0 + \frac{1}{\sigma_2^2} A_\tau^T A_\tau \right]^\dagger \left(\frac{1}{\sigma_1^2} A_0^T z_1 + \frac{1}{\sigma_2^2} A_\tau^T z_2 \right),$$

where $(\cdot)^\dagger$ indicates the pseudo-inverse operator for the Gram matrix.

- 5: **until** Some convergence condition is satisfied.
-

- As $\sigma_1 \rightarrow 0$, the contribution of A_0 and z_1 dominates (4.28), and the solution reduces to

$$(4.29) \quad \hat{c} = [A_0^T A_0]^\dagger A_0^T z_1,$$

which corresponds to the conventional method where z_1 is considered to be a highly reliable “template” and the image intensity is solely obtained by fitting z_1 .

- More generally, alternating descent may be used instead of requiring the achieving minimizer at each iteration. This could be particularly beneficial for the step in updating τ conditioned on \hat{c} , as the quadratic form in the other step makes the minimization over c trivial. Relaxing conditional maximization to increment in log-likelihood may have potential computational advantage as well as better behavior to local maxima.
- As $\sigma_1 \rightarrow 0$, the alternating descent algorithm reduces to exactly any conventional descent algorithm in solving (4.3) with l_2 difference metric. In the asymptotic case, the conditional minimization of c given by (4.29) is independent of τ and the whole alternating descent algorithm reduces to using the plug-in estimator (4.29) and descend $-\Lambda$ with respect to τ .

4.5 Comparison with Conventional Methods: CRB v.s. M-estimate

As we have commented briefly in the previous sections, the conventional method estimate the intensity f from the source image z_1 only. With l_2 difference metric, we can write the solution to the conventional method as:

$$(4.30) \quad \begin{aligned} \hat{c} &= \arg \min_c \|z_1 - A_0 c\|_2^2; \\ \hat{\tau} &= \arg \min_\tau \|z_2 - A_\tau \hat{c}\|_2^2, \end{aligned}$$

where z_1, z_2 are discrete observations for the source and target image in vector form, A_0 and A_τ are defined as in (4.4).

The first equation in (4.30) can be solved in closed form given its quadratic form:

$$\hat{c} = A_0^\dagger z_1,$$

and we can rewrite (4.30) as:

$$(4.31) \quad \hat{\tau} = \arg \min_\tau \|z_2 - A_\tau A_0^\dagger z_1\|_2^2.$$

We can also stack the expression as before, and define $\mathcal{A} \triangleq [-A_\tau A_0^\dagger \quad I]$ and write the objective as:

$$(4.32) \quad \hat{\tau} = \arg \min_\tau \Phi(\tau, z) = \|\mathcal{A}(\tau)z\|_2^2.$$

In the following derivations, we will choose the most convenient form and use the above equivalent expressions interchangeably.

Our goal is to derive the covariance of the minimizer defined above and we use similar philosophy as in [30]. By implicit function theorem, the partial derivative of Φ with respect to τ are uniformly zero:

$$(4.33) \quad \frac{\partial}{\partial \tau(i)} \Phi(\tau, z)|_{\tau=\hat{\tau}} = 0, \quad \forall \text{ spatial location } i,$$

for any given data z .

Differentiating (4.33) again with respect to z and applying the chain rule yields:

$$(4.34) \quad \nabla^{20} \Phi(\hat{\tau}(z), z) \nabla_z \hat{\tau}(z) + \nabla^{11} \Phi(\hat{\tau}(z), z) = 0.$$

Where, the components of $\nabla^{20} \Phi(\hat{\tau}(z), z)$ are $\frac{\partial^2}{\partial \tau(i) \partial \tau(j)} \Phi(\hat{\tau}(z), z)$, and the elements of ∇^{11} are $\frac{\partial^2}{\partial \tau(i) \partial z(j)} \Phi(\hat{\tau}(z), z)$. We consider the case when $\nabla^{20} \Phi(\hat{\tau}(z), z)$ is invertible, or more precisely positive definite. This is equivalent to require $\Phi(\hat{\tau}(z), z)$ to be *locally* strictly convex. This assumption is true if the following regularity condition is satisfied: *there \exists a compact neighborhood $\mathcal{N}(\hat{\tau})$ such that $\Phi(\tau, z) > \Phi(\hat{\tau}(z), z)$ for all $\tau \neq \hat{\tau}$.* Then we have:

$$\nabla_Y \hat{\tau}(z) = [-\nabla^{20} \Phi(\hat{\tau}, z)]^{-1} \nabla^{11} \Phi(\tau, z),$$

and the covariance matrix for $\hat{\tau}$ would be $\text{Cov}\{z\}$ transformed by local linearization [91], *i.e.*,

$$\text{Cov}\{\hat{\tau}\} \approx \nabla_z \hat{\tau}(z) \text{Cov}\{z\} [\nabla_z \hat{\tau}(z)]'.$$

By substitution, we obtain

$$(4.35) \quad \text{Cov}\{\hat{\tau}\} \approx [\nabla^{20} \Phi(\hat{\tau}, z)]^{-1} \nabla^{11} \Phi(\hat{\tau}, z) \text{Cov}\{z\} [\nabla^{11} \Phi(\hat{\tau}, z)]' [\nabla^{20} \Phi(\hat{\tau}, z)]^{-1}.$$

We assume the covariance of z to be:

$$(4.36) \quad \text{Cov}\{z\} = \begin{bmatrix} \sigma_1^2 I_N & 0 \\ 0 & \sigma_2^2 I_N \end{bmatrix},$$

so it remains to derive the expressions for $\nabla^{20} \Phi(\hat{\tau}, z)$ and $\nabla^{11} \Phi(\hat{\tau}, z)$.

We first adopt the objective function form in (4.31) to take derivative with respect to $\tau(l)$.

$$(4.37) \quad \frac{\partial}{\partial \tau(l)} \Phi(\tau, z) = \sum_{n=1}^N (A_\tau[n, :] A_0^\dagger z_1 - z_2(n)) \frac{\partial^2}{\partial \tau(l)^2} \{A_\tau(n) A_0^\dagger z_1\}.$$

Similar to (4.10),

$$(4.38) \quad \begin{aligned} \frac{\partial}{\partial \tau(l)} \{A_\tau[n, :] A_0^\dagger z_1\} &= \frac{\partial}{\partial \tau(l)} \sum_{k=1}^K (A_0^\dagger z_1)(k) A_\tau(n, k) \\ &= \frac{\partial}{\partial \tau(l)} \sum_{k=1}^K (A_0^\dagger z_1)(k) b(n + \tau(n), k) \\ &= \begin{cases} \sum_{k=1}^K (A_0^\dagger z_1)(k) \dot{b}(n + \tau(n), k), & l = n; \\ 0 & \text{else,} \end{cases} \end{aligned}$$

where $\dot{b}(\cdot, \cdot)$ denote the derivative of $b(\cdot, \cdot)$ with respect to the first variable.

Plugging (4.38) into the expression in (4.37) yields:

$$(4.39) \quad \frac{\partial}{\partial \tau(l)} \Phi(\tau, z) = (A_\tau[l, :] A_0^\dagger z_1 - z_2(l)) \sum_{k=1}^K (A_0^\dagger z_1)(k) \dot{b}(l + \tau(l), k).$$

To obtain $\nabla^{20}\Phi$, we take derivative with respect to $\tau(n)$. Noticing that $\frac{\partial}{\partial\tau(l)}\Phi$ depends on τ only via $\tau(l)$, we obtain:

$$\frac{\partial^2}{\partial\tau(l)\partial\tau(n)}\Phi(\tau, z) = \begin{cases} \left\{ \sum_{k=1}^K (A_0^\dagger z_1)(k) \dot{b}(l + \tau(l), k) \right\}^2 + \dots \\ + (A_\tau[l, :] A_0^\dagger z_1 - z_2(l)) \sum_{k=1}^K (A_0^\dagger z_1)(k) \ddot{b}(l + \tau(l), k), & l = n; \\ 0 & \text{else.} \end{cases}$$

where $\ddot{b}(\cdot, \cdot)$ denotes the second-order partial derivative with respect to the first argument in $b(\cdot, \cdot)$.

To compute $\nabla^{11}\Phi(\hat{\tau}, z)$, we need to take derivative of (4.39) with respect to each element of z . We perform this by differentiating with respect to the elements in z_1 and z_2 respectively.

Noting that $\frac{\partial}{\partial z_1(n)} [A_0^\dagger z_1](k) = A_0^\dagger[k, n]$, we obtain:

$$(4.40) \quad \begin{aligned} \frac{\partial^2}{\partial\tau(l)\partial z_1(n)}\Phi(\tau, z) &= A_\tau[l, :] A_0^\dagger[:, n] \sum_{k=1}^K (A_0^\dagger z_1)(k) \dot{b}(l + \tau(l), k) + \dots \\ &+ (A_\tau[l, :] A_0^\dagger z_1 - z_2(l)) \sum_{k=1}^K A_0^\dagger[k, n] \dot{b}(l + \tau(l), k). \end{aligned}$$

$$\frac{\partial^2}{\partial\tau(l)\partial z_2(n)}\Phi(\tau, z) = \begin{cases} -\sum_{k=1}^K (A_0^\dagger z_1)(k) \dot{b}(l + \tau(l), k), & l = n; \\ 0 & \text{else.} \end{cases}$$

We assume that at the point of evaluation $(\check{\tau}, \bar{z})$, the samples of the warped z_1^c approximates the observation z_2 , more specifically:

$$A_\tau A_0^\dagger \bar{z}_1 \approx \bar{z}_2.$$

This is a reasonable assumption for most registration results. For simplicity, we denote $\bar{c} \triangleq A_0^\dagger \bar{z}_1$, $\check{D}(i, j) \triangleq \dot{b}(i + \check{\tau}(i), j)$, and the warping map $\mathcal{W} \triangleq A_\tau A_0^\dagger$, then we can rewrite in matrix form:

$$(4.41) \quad \begin{aligned} \nabla^{20}\Phi(\check{\tau}, \bar{z}) &= \text{diag}^2\{\check{D}\bar{c}\} \\ \nabla^{11}\Phi(\check{\tau}, \bar{z}) &= \begin{bmatrix} \text{diag}\{\check{D}\bar{c}\} \mathcal{W} & -\text{diag}\{\check{D}\bar{c}\} \end{bmatrix}. \end{aligned}$$

Plugging (4.41) and (4.36) into the expression for $\text{Cov}\{\hat{\tau}\}$ in (4.35), we obtain:

$$(4.42) \quad \text{Cov}\{\hat{\tau}\}|_{\hat{\tau}=\check{\tau}} \approx \text{diag}\{\check{D}\bar{c}\}^{-1} [\sigma_1^2 \mathcal{W} \mathcal{W}^T + \sigma_2^2 I] \text{diag}\{\check{D}\bar{c}\}^{-1}.$$

Remark: as $\sigma_1^2 \rightarrow 0$, z_1 approaches the noise-free observation of the source image f , and the conventional method should yield the same estimate as the more realistic model. In fact,

$$\text{Cov}_{\sigma_1 \rightarrow 0}\{\hat{\tau}\} = \sigma_2^2 \text{diag}^2\{\check{D}\bar{c}\},$$

which agrees with our previous analysis in (4.24) that $\text{CRB}(\tau) \rightarrow J_{\tau, \check{\tau}}^{-1}$ as z_1 becomes asymptotically noise-free.

It makes sense to compare the covariance prediction for the M-estimate of the conventional method and the Cramér-Rao Bound obtained from the more realistic model from (4.2). For simplicity, we assume that A_0 to be invertible so that $A_0^{-1} = A_0^\dagger$ and consequently the warping map $\mathcal{W} = A_\tau A_0^{-1}$ to be invertible.

To study $\text{CRB}(\tau)$, we plug in $J_{\tau, c}, J_{c, c}$ from (4.24) and obtain:

$$(4.43) \quad \begin{aligned} \text{CRB}(\tau) &= [J_{\tau, \tau} - J_{\tau, c} J_{c, c}^{-1} J_{c, \tau}]^{-1} \\ &= \left\{ \frac{1}{\sigma_2^2} \text{diag}^2\{Dc\} - \frac{1}{\sigma_2^2} \text{diag}\{Dc\} A_\tau \left[\frac{1}{\sigma_1^2} A_0^T A_0 + \frac{1}{\sigma_2^2} A_\tau^T A_\tau \right]^{-1} \frac{1}{\sigma_2^2} A_\tau^T \text{diag}\{Dc\} \right\}^{-1} \\ &= \sigma_2^2 \text{diag}\{Dc\}^{-1} \left\{ I - \frac{1}{\sigma_2^2} A_\tau \left[\frac{1}{\sigma_1^2} A_0^T A_0 + \frac{1}{\sigma_2^2} A_\tau^T A_\tau \right]^{-1} A_\tau^T \right\}^{-1} \text{diag}\{Dc\}^{-1}. \end{aligned}$$

With $A_\tau = \mathcal{W} A_0$, we can write:

$$\frac{1}{\sigma_1^2} A_0^T A_0 + \frac{1}{\sigma_2^2} A_\tau^T A_\tau = \frac{1}{\sigma_1^2} A_0^T A_0 + \frac{1}{\sigma_2^2} A_0^T \mathcal{W}^T \mathcal{W} A_0.$$

The middle part of (4.43) can be rewritten as:

$$\begin{aligned}
& \left\{ I - \frac{1}{\sigma_2^2} A_\tau \left[\frac{1}{\sigma_1^2} A_0^T A_0 + \frac{1}{\sigma_2^2} A_\tau^T A_\tau \right]^{-1} A_\tau^T \right\}^{-1} \\
&= \left\{ I - \sigma_1^2 A_\tau \left[\sigma_2^2 A_0^T A_0 + \sigma_1^2 A_0^T \mathcal{W}^T \mathcal{W} A_0 \right]^{-1} A_\tau^T \right\}^{-1} \\
&= \left\{ I - \sigma_1^2 A_\tau A_0^{-1} \left[\sigma_2^2 I + \sigma_1^2 \mathcal{W}^T \mathcal{W} \right]^{-1} A_0^{-T} A_\tau^T \right\}^{-1} \\
(4.44) \quad &= \left\{ I - \sigma_1^2 \mathcal{W} \left[\sigma_2^2 I + \sigma_1^2 \mathcal{W}^T \mathcal{W} \right]^{-1} \mathcal{W}^T \right\}^{-1}.
\end{aligned}$$

By Woodbury-Sherman-Morrissey identity:

$$\left[\sigma_2^2 I + (\sigma_1 \mathcal{W}) \sigma_1 \mathcal{W}^T \right]^{-1} = \frac{1}{\sigma_2^2} I - \frac{1}{\sigma_2^2} \sigma_1^2 \mathcal{W} \left[I + \frac{\sigma_1^2}{\sigma_2^2} \mathcal{W} \mathcal{W}^T \right]^{-1} \mathcal{W}^T,$$

thus $\sigma_2 \left\{ I - \frac{1}{\sigma_2^2} A_\tau \left[\frac{1}{\sigma_1^2} A_0^T A_0 + \frac{1}{\sigma_2^2} A_\tau^T A_\tau \right]^{-1} A_\tau^T \right\}^{-1} = \sigma_2^2 I + \sigma_1^2 \mathcal{W} \mathcal{W}^T$.

Substituting into (4.43) yields:

$$(4.45) \quad \text{CRB}(\tau) = \text{diag}\{Dc\} (\sigma_2^2 I + \sigma_1^2 \mathcal{W} \mathcal{W}^T) \text{diag}\{Dc\}.$$

This result coincides with the covariance estimate for the M-estimate evaluated at (\check{D}, \bar{c}) in (4.42).

4.6 A Simple Example

This section uses a simple example to illustrate the results from previous sections and also to motivate discussions about performance comparison. In particular, it is expected that the proposed model in (4.4) has advantage over the traditional model in (4.3) as the estimation for c which parametrizes the underlying image intensity should be more reliable, because it combines the information from both the source and the target observations. Consider a simple model

$$(4.46) \quad z = \begin{bmatrix} z_1 \\ z_2 \end{bmatrix} = \begin{bmatrix} I \\ \alpha I \end{bmatrix} c + \begin{bmatrix} \epsilon_1 \\ \epsilon_2 \end{bmatrix},$$

where we assume both z_1 and z_2 are vectors of the same size as the underlying (unknown) c . The scaling parameter α which relates z_1 and z_2 in the noise-free case is the quantity of interest. $\epsilon \sim \mathcal{N}(0, \sigma_1^2 I)$ and $\epsilon_2 \sim \mathcal{N}(0, \sigma_2^2 I)$ are independent Gaussian additive noise.

M-estimator for the Conventional Method

In the conventionally method, the parameter c is estimated solely from observation z_1 :

$$(4.47) \quad \hat{c}(z) = \arg \min_c \|z_1 - c\|_2^2 = z_1.$$

Since $z_1 \sim \mathcal{N}(c, \sigma_1^2 I)$, \hat{c} is an unbiased estimator for c with covariance $\sigma_1^2 I$.

The objective function that $\hat{\alpha}$ minimizes is

$$(4.48) \quad \Phi(\alpha, z) \triangleq \|[\alpha I \quad -I]z\|_2^2 = \|z_2 - \alpha z_1\|_2^2.$$

$$\begin{aligned}
(4.49) \quad \hat{\alpha}(z) &= \arg \min_{\alpha} \Phi(\alpha, z) \\
&= \arg \min_{\alpha} \|z_2 - \alpha \hat{c}\|_2^2 \\
&= \arg \min_{\alpha} \|z_2 - \alpha z_1\|_2^2 \\
&= \frac{z_1^T z_2}{\|z_1\|_2^2}.
\end{aligned}$$

Hereafter, we discuss two approaches in approximating the mean and variance of $\hat{\alpha}$: a direct method based on the explicit solution in (4.49); and an indirect approach that relies on implicit function theorem and M-estimate. The explicit method is straightforward, requires less manipulation, and should be reasonably accurate. On the other hand, explicit solutions are not available in general (as we will see for the ML estimator), so the implicit method is more universally applicable. In this study, the direct method serves as a good baseline reference for approximation performance, and the derivation based on indirect approach is of didactic value.

Direct Approximation of Mean and Variance for the M-estimate

First, we directly approximate the mean and covariance of $\hat{\alpha}$ based on the explicit solution in (4.49). The expected value of $\hat{\alpha}$ from (4.49) is given by:

$$E[\hat{\alpha}] = E \left\{ \frac{(\bar{\mathbf{c}} + \boldsymbol{\varepsilon}_1)^T (\bar{\boldsymbol{\alpha}} \bar{\mathbf{c}} + \boldsymbol{\varepsilon}_2)}{(\bar{\mathbf{c}} + \boldsymbol{\varepsilon}_1)^T (\bar{\mathbf{c}} + \boldsymbol{\varepsilon}_1)} \right\},$$

where $\boldsymbol{\varepsilon}_1 \sim \mathcal{N}(0, \sigma_1^2 I)$ and $\boldsymbol{\varepsilon}_2 \sim \mathcal{N}(0, \sigma_2^2 I)$. We compute the above expression using conditional expectation:

$$\begin{aligned} E[\hat{\alpha}] &= E_{\boldsymbol{\varepsilon}_1} \{ E_{\boldsymbol{\varepsilon}_2} [\hat{\alpha} | \boldsymbol{\varepsilon}_1] \} \\ (4.50) \quad &= \bar{\boldsymbol{\alpha}} E_{\boldsymbol{\varepsilon}_1} \left\{ \frac{(\bar{\mathbf{c}} + \boldsymbol{\varepsilon}_1)^T \bar{\mathbf{c}}}{(\bar{\mathbf{c}} + \boldsymbol{\varepsilon}_1)^T (\bar{\mathbf{c}} + \boldsymbol{\varepsilon}_1)} \right\}. \end{aligned}$$

where the second line follows from the independence between $\boldsymbol{\varepsilon}_1$ and $\boldsymbol{\varepsilon}_2$.

Let c_i denote the i th element of $\bar{\mathbf{c}}$ and e_i denote the i th element of $\boldsymbol{\varepsilon}_1$. Then c_i are constants and e_i are scalar i.i.d Gaussian variables $e_i \sim \mathcal{N}(0, \sigma_1^2)$.

We can rewrite (4.50) as:

$$(4.51) \quad E[\hat{\alpha}] / \bar{\boldsymbol{\alpha}} = E \left\{ \frac{\sum_{i=1}^n (c_i + e_i) c_i}{\sum_{i=1}^n (c_i + e_i)^2} \right\}.$$

Define function $f : \mathfrak{R}^n \rightarrow \mathfrak{R}$ via $f(\mathbf{x}) = \frac{\mathbf{x}^T \mathbf{c}}{\mathbf{x}^T \mathbf{x}}$. We perform second-order Taylor expansion of f around the point $\mathbf{x} = \bar{\mathbf{c}}$ and then take expectation with respect to $\mathbf{x} = \bar{\mathbf{c}} + \boldsymbol{\varepsilon}_1$:

$$\begin{aligned} E[\hat{\alpha}] / \bar{\boldsymbol{\alpha}} &= E[f(\bar{\mathbf{c}}) + \frac{1}{2} (\boldsymbol{\varepsilon}_1 - \bar{\mathbf{c}})^T \nabla_{\mathbf{x}}^2 f(\bar{\mathbf{c}}) (\boldsymbol{\varepsilon}_1 - \bar{\mathbf{c}})] \\ &= 1 + \frac{1}{2} E[(\boldsymbol{\varepsilon}_1 - \bar{\mathbf{c}})^T \nabla_{\mathbf{x}}^2 f(\bar{\mathbf{c}}) (\boldsymbol{\varepsilon}_1 - \bar{\mathbf{c}})] \\ (4.52) \quad &= 1 + \frac{1}{2} E[\boldsymbol{\varepsilon}_1^T \nabla_{\mathbf{x}}^2 f(\bar{\mathbf{c}}) \boldsymbol{\varepsilon}_1]. \end{aligned}$$

Now we focus on the term $E[\boldsymbol{\varepsilon}_1^T \nabla_{\mathbf{x}}^2 f(\bar{\mathbf{c}}) \boldsymbol{\varepsilon}_1]$ whose sign determines the bias. The gradient $\nabla_{\mathbf{x}} f$ and the Hessian $\nabla_{\mathbf{x}}^2 f$ of f are derived as follows:

$$\nabla_{\mathbf{x}} f = \|\mathbf{x}\|_2^{-2} \bar{\mathbf{c}}^T - 2 \|\mathbf{x}\|_2^{-4} (\mathbf{x}^T \bar{\mathbf{c}}) \mathbf{x}^T.$$

The i th element of $\nabla_{\mathbf{x}} f$ is

$$[\nabla_{\mathbf{x}} f]_i = \|\mathbf{x}\|_2^{-2} c_i - 2 \|\mathbf{x}\|_2^{-4} (\mathbf{x}^T \bar{\mathbf{c}}) x_i.$$

Taking derivative with respect to x_j yields:

$$\frac{\partial}{\partial x_j} [\nabla_{\mathbf{x}} f]_i = -2 \|\mathbf{x}\|_2^{-4} c_i x_j - 2 \left\{ -4 \|\mathbf{x}\|_2^{-6} \mathbf{x}^T \bar{\mathbf{c}} x_i x_j + \|\mathbf{x}\|_2^{-4} (x_i c_j + \mathbf{x}^T \bar{\mathbf{c}} \delta[i-j]) \right\},$$

where δ is the Kronecker impulse function defined as

$$\delta[x] = \begin{cases} 1 & x = 0; \\ 0 & \text{otherwise.} \end{cases}$$

The equivalent matrix representation of the Hessian is given by:

$$(4.53) \quad \nabla_{\mathbf{x}}^2 f = 8 \|\mathbf{x}\|_2^{-6} \mathbf{x}^T \bar{\mathbf{c}} \mathbf{x} \mathbf{x}^T - 2 \|\mathbf{x}\|_2^{-4} (\mathbf{x} \bar{\mathbf{c}}^T + \bar{\mathbf{c}} \mathbf{x}^T) - 2 (\mathbf{x}^T \bar{\mathbf{c}}) \|\mathbf{x}\|_2^{-4} \mathbf{I}.$$

We evaluate the Hessian at $\bar{\mathbf{c}}$ and note that $E[\boldsymbol{\varepsilon}_1^T \nabla_{\mathbf{x}}^2 f(\bar{\mathbf{c}}) \boldsymbol{\varepsilon}_1] = \sigma_1^2 \text{trace}\{\nabla^2 f(\bar{\mathbf{c}})\}$ depends only on the diagonal elements of the Hessian, because the noise $\boldsymbol{\varepsilon}_1$ is i.i.d. We obtain:

$$[\nabla_{\mathbf{x}}^2 f(\bar{\mathbf{c}})]_{ii} = 2 \|\bar{\mathbf{c}}\|_2^{-4} (2c_i^2 - \sum_{j=1}^n c_j^2),$$

so that

$$(4.54) \quad \begin{aligned} E[\boldsymbol{\varepsilon}_1^T \nabla_{\mathbf{x}}^2 f(\bar{\mathbf{c}}) \boldsymbol{\varepsilon}_1] &= \sigma_1^2 \sum_{i=1}^n [\nabla_{\mathbf{x}}^2 f(\bar{\mathbf{c}})]_{ii} \\ &= 2\sigma_1^2 \|\bar{\mathbf{c}}\|_2^{-2} (2-n), \end{aligned}$$

which is negative for all $n > 2$.

Subsequently,

$$(4.55) \quad E[\hat{\alpha}] / \bar{\alpha} \approx 1 - (n-2)\sigma_1^2 \|\bar{\mathbf{c}}\|_2^{-2}.$$

As (4.54) describes the difference between $E[\hat{\alpha}/\bar{\alpha}]$ and unity, this indicates that for $n > 2$, $\hat{\alpha}$ is an estimate of $\bar{\alpha}$ that biases towards smaller magnitude.

Similarly, we compute $\text{Var}\{\hat{\alpha}\}$ via $E[\hat{\alpha}^2] - E[\hat{\alpha}]^2$. The correlation reads:

$$E \left\{ \frac{(\bar{\mathbf{c}} + \boldsymbol{\varepsilon}_1)^T (\bar{\alpha} \bar{\mathbf{c}} + \boldsymbol{\varepsilon}_2) (\bar{\alpha} \bar{\mathbf{c}} + \boldsymbol{\varepsilon}_2)^T (\bar{\mathbf{c}} + \boldsymbol{\varepsilon}_1)}{\|\bar{\mathbf{c}} + \boldsymbol{\varepsilon}_1\|_2^4} \right\}.$$

As before, we first use conditional expectation to separate out the uncertainty in $\boldsymbol{\varepsilon}_2$ via:

$$E[\hat{\alpha}^2] = E_{\boldsymbol{\varepsilon}_1} E_{\boldsymbol{\varepsilon}_2} [\hat{\alpha}^2 | \boldsymbol{\varepsilon}_1] = E \left\{ \frac{(\bar{\mathbf{c}} + \boldsymbol{\varepsilon}_1)^T (\bar{\alpha}^2 \bar{\mathbf{c}} \bar{\mathbf{c}}^T + \sigma_2^2 \mathbf{I}) (\bar{\mathbf{c}} + \boldsymbol{\varepsilon}_1)}{\|\bar{\mathbf{c}} + \boldsymbol{\varepsilon}_1\|_2^4} \right\}.$$

Define a deterministic symmetric matrix $H \triangleq (\bar{\alpha}^2 \bar{\mathbf{c}} \bar{\mathbf{c}}^T + \sigma_2^2 \mathbf{I})$ and a function $f(\mathbf{x}) = \frac{\mathbf{x}^T H \mathbf{x}}{\|\mathbf{x}\|_2^4}$, and we aim to find $E[f(\mathbf{x})]$ for $\mathbf{x} = \bar{\mathbf{c}} + \boldsymbol{\varepsilon}_1$. We expand the function $f(\mathbf{x})$ around $\mathbf{x} = \bar{\mathbf{c}}$ and approximate $E[\hat{\alpha}^2]$ via:

$$(4.56) \quad \begin{aligned} E[\hat{\alpha}^2] &\approx \frac{\bar{\mathbf{c}}^T H \bar{\mathbf{c}}}{\|\bar{\mathbf{c}}\|_2^4} + \frac{1}{2} E[(\mathbf{x} - \bar{\mathbf{c}})^T \nabla_{\mathbf{x}}^2 f(\bar{\mathbf{c}}) (\mathbf{x} - \bar{\mathbf{c}})] \\ &= \frac{\bar{\mathbf{c}}^T H \bar{\mathbf{c}}}{\|\bar{\mathbf{c}}\|_2^4} + \frac{1}{2} E[\boldsymbol{\varepsilon}_1^T \nabla_{\mathbf{x}}^2 f(\bar{\mathbf{c}}) \boldsymbol{\varepsilon}_1]. \end{aligned}$$

The deterministic term $f(\bar{\mathbf{c}})$ simplifies to:

$$f(\bar{\mathbf{c}}) = \frac{\bar{\mathbf{c}}^T H \bar{\mathbf{c}}}{\|\bar{\mathbf{c}}\|_2^4} = \bar{\alpha}^2 + \frac{\sigma_2^2}{\|\bar{\mathbf{c}}\|_2^2}.$$

Since $\boldsymbol{\varepsilon}_1$ is componentwise independent, $E[\hat{\alpha}^2]$ only depends on the diagonal element of $\nabla_{\mathbf{x}}^2 f(\bar{\mathbf{c}})$, which we derive as follows.

$$\nabla_{\mathbf{x}} f(\mathbf{x}) = -4 \|\mathbf{x}\|_2^{-6} \mathbf{x}^T (\mathbf{x}^T H \mathbf{x}) + 2 \|\mathbf{x}\|_2^{-4} \mathbf{x}^T H.$$

The i th element of $\nabla_{\mathbf{x}} f(\mathbf{x})$ reads $-4 \|\mathbf{x}\|_2^{-6} x_i (\mathbf{x}^T H \mathbf{x}) + 2 \|\mathbf{x}\|_2^{-4} \mathbf{x}^T H(:, i)$, where $H(:, i)$ indicates the i th column of H . We may explicitly write $\mathbf{x}^T H(:, i) = \sum_j x_j [\bar{\alpha}^2 c_j c_i + \sigma_2^2 \delta[i-j]]$. The second-order derivative is given by:

$$(4.57) \quad \begin{aligned} \frac{\partial^2}{\partial x_i^2} f(\mathbf{x}) &= -4 \|\mathbf{x}\|_2^{-6} [\mathbf{x}^T H \mathbf{x} + 2x_i \mathbf{x}^T H(:, i)] + 24 \|\mathbf{x}\|_2^{-8} x_i^2 \mathbf{x}^T H \mathbf{x} \\ &\quad + 2 \|\mathbf{x}\|_2^{-4} (\bar{\alpha}^2 c_i^2 + \sigma_2^2) - 8 \|\mathbf{x}\|_2^{-6} x_i \mathbf{x}^T H(:, i). \end{aligned}$$

To evaluate $\frac{\partial^2}{\partial x_i^2} f(\mathbf{x})$ at $\mathbf{x} = \bar{\mathbf{c}}$, we use the following relations:

$$\begin{aligned}\bar{\mathbf{c}}^T \mathbf{H}(:, i) &= c_i (\bar{\alpha}^2 \|\bar{\mathbf{c}}\|_2^2 + \sigma_2^2); \\ \bar{\mathbf{c}}^T \mathbf{H} \bar{\mathbf{c}} &= \|\bar{\mathbf{c}}\|_2^2 (\bar{\alpha}^2 \|\bar{\mathbf{c}}\|_2^2 + \sigma_2^2).\end{aligned}$$

Substituting these relations into the expression (4.57) for $\frac{\partial^2}{\partial x_i^2} f(\mathbf{x})$, we obtain:

$$\frac{\partial^2}{\partial x_i^2} f(\mathbf{x})|_{\mathbf{x}=\bar{\mathbf{c}}} = 8 \|\bar{\mathbf{c}}\|^{-6} c_i^2 (\bar{\alpha}^2 \|\bar{\mathbf{c}}\|_2^2 + \sigma_2^2) - 4 \|\bar{\mathbf{c}}\|^{-4} (\bar{\alpha}^2 \|\bar{\mathbf{c}}\|_2^2 + \sigma_2^2) + 2 \|\bar{\mathbf{c}}\|^{-4} (\bar{\alpha}^2 c_i^2 + \sigma_2^2).$$

By the independence of the elements in ε_1 , we obtain:

$$\begin{aligned}E[\varepsilon_1^T \nabla_{\mathbf{x}}^2 f(\bar{\mathbf{c}}) \varepsilon_1] &= \sigma_1^2 \sum_i \frac{\partial^2}{\partial x_i^2} f(\bar{\mathbf{c}}) \\ (4.58) \qquad \qquad \qquad &= \|\bar{\mathbf{c}}\|^{-2} (10 - 4n) \bar{\alpha}^2 \sigma_1^2 + \|\bar{\mathbf{c}}\|^{-4} (8 - 2n) \sigma_1^2 \sigma_2^2.\end{aligned}$$

Substituting this quantity into (4.56) provides:

$$E[\hat{\alpha}^2] \approx \bar{\alpha}^2 + \|\bar{\mathbf{c}}\|^{-2} \sigma_2^2 + \|\bar{\mathbf{c}}\|^{-2} (5 - 2n) \bar{\alpha}^2 \sigma_1^2 + \|\bar{\mathbf{c}}\|^{-4} (4 - n) \sigma_1^2 \sigma_2^2.$$

Together with the estimation for $E[\hat{\alpha}]$ obtained in (4.55), this equation yields an approximation for $\text{Var}\{\hat{\alpha}\}$ as:

$$\begin{aligned}\text{Var}\{\hat{\alpha}\} &= E[\hat{\alpha}^2] - E[\hat{\alpha}]^2 \\ (4.59) \qquad \qquad \qquad &= \|\bar{\mathbf{c}}\|^{-2} (\bar{\alpha}^2 \sigma_1^2 + \sigma_2^2) - \|\bar{\mathbf{c}}\|^{-4} \sigma_1^2 [(n-4) \sigma_2^2 - (n-2)^2 \bar{\alpha}^2 \sigma_1^2].\end{aligned}$$

Expressions (4.55) and (4.59) reveal some interesting structure. For large enough n (in fact for $n > 6$), the variance estimate (4.59) becomes upper-bounded by $\|\bar{\mathbf{c}}\|^{-2} (\bar{\alpha}^2 \sigma_1^2 + \sigma_2^2)$, which we will show later is the Cramér-Rao Bound for the statistical model. This implies that it cannot be unbiased. In fact, the bias quantity measured by $(2 - n) \|\bar{\mathbf{c}}\|^{-2} \sigma_1^2 \bar{\alpha}$ also increases accordingly.

Alternatively, we can follow [30], and use implicit function theorem and Taylor expansion to approximate the bias and variance of $\hat{\alpha}$ as the minimizer of (4.48). The data point $\bar{\mathbf{z}}$ at which to perform Taylor expansion is mainly a choice of convenience rather than considerations of asymptotic behavior. One natural choice of the expansion point would be the noiseless data. Let $\bar{\mathbf{z}}$ denote the noiseless observation $\bar{\mathbf{c}}$ and $\bar{\alpha}$ denote the true parameter values, with $\check{\mathbf{c}}$ and $\check{\alpha}$ denoting the resulting estimates in (4.47) and (4.49) when $\bar{\mathbf{z}}$ is observed. Then $\bar{\mathbf{z}} = [\bar{\mathbf{c}}; \bar{\alpha} \bar{\mathbf{c}}]$, and

$$\begin{aligned}\check{\mathbf{c}} &= \hat{\mathbf{c}}(\bar{\mathbf{z}}) = \bar{\mathbf{c}}; \\ (4.60) \qquad \qquad \qquad \check{\alpha} &= \hat{\alpha}(\bar{\mathbf{z}}) = \frac{\bar{\alpha} \bar{\mathbf{c}}^T \bar{\mathbf{c}}}{\|\bar{\mathbf{c}}\|_2^2} = \bar{\alpha}.\end{aligned}$$

As the minimizer for (4.48), $\hat{\alpha}$ satisfies:

$$\frac{\partial}{\partial \alpha} \Phi(\alpha, \mathbf{z})|_{\alpha=\hat{\alpha}} = 2\mathbf{z}^T \begin{bmatrix} I \\ 0 \end{bmatrix} \begin{bmatrix} \alpha I & -I \end{bmatrix} \mathbf{z} = 0 \quad \forall \mathbf{z}.$$

Taking derivative with respect to \mathbf{z} and invoking the chain rule, we obtain:

$$\frac{\partial^2}{\partial \alpha^2} \Phi \frac{\partial}{\partial \mathbf{z}} \alpha + \frac{\partial^2}{\partial \alpha \partial \mathbf{z}} \Phi = 0,$$

where

$$(4.61) \qquad \qquad \qquad \frac{\partial^2}{\partial \alpha^2} \Phi = 2 \|z_1\|_2^2 = 2\mathbf{z}^T \begin{bmatrix} I & 0 \\ 0 & 0 \end{bmatrix} \mathbf{z},$$

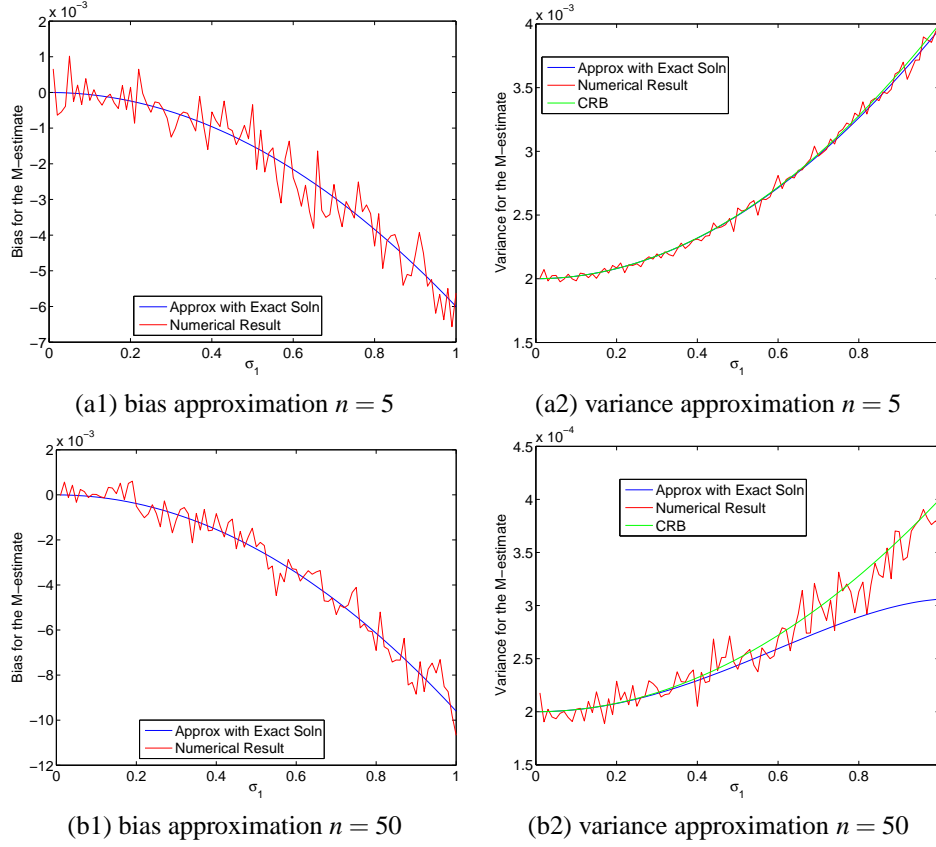


Figure 4.1: Bias and variance approximation obtained from explicit solution for conventional M-estimate.

and

$$(4.62) \quad \frac{\partial^2}{\partial \alpha \partial z} \Phi = 2z^T \left\{ \begin{bmatrix} \alpha I & -I \\ 0 & 0 \end{bmatrix} + \begin{bmatrix} \alpha I & 0 \\ -I & 0 \end{bmatrix} \right\} = 2z^T \begin{bmatrix} 2\alpha I & -I \\ -I & 0 \end{bmatrix}.$$

Therefore,

$$(4.63) \quad \frac{\partial}{\partial z} \hat{\alpha}(z) = -\frac{\partial^2}{\partial \alpha^2} \Phi^{-1} \frac{\partial^2}{\partial \alpha \partial z} \Phi = -\|z_1\|_2^{-2} z^T \begin{bmatrix} 2\alpha I & -I \\ -I & 0 \end{bmatrix}.$$

Evaluating (4.63) at $z = \bar{z}$, we obtain an estimate of covariance $\text{Cov}\{\alpha\}$ at $\hat{\alpha} = \hat{\alpha}z$ as

$$(4.64) \quad \begin{aligned} \text{Cov}\{\hat{\alpha}(z)\} &\approx \frac{\partial}{\partial z} \alpha(\bar{z}) \text{Cov}\{z\} \frac{\partial}{\partial z} \alpha^T(\bar{z}) \\ &= \frac{-1}{\|\bar{c}\|_2^2} \bar{c}^T \begin{bmatrix} \bar{\alpha} I & -I \end{bmatrix} \begin{bmatrix} \sigma_1^2 I & \\ & \sigma_2^2 I \end{bmatrix} \frac{-1}{\|\bar{c}\|_2^2} \begin{bmatrix} \bar{\alpha} I \\ -I \end{bmatrix} \bar{c} \\ &= \frac{\bar{\alpha}^2 \sigma_1^2 + \sigma_2^2}{\|\bar{c}\|_2^2}. \end{aligned}$$

This quantity (4.64) coincides with the Cramér-Rao Bound obtained from the statistical model as we will show later.

To estimate the bias for $\hat{\alpha}$, we present the first and second-order Taylor expansion for $E[\hat{\alpha}]$ as:

$$(4.65) \quad \begin{aligned} E^{(1)}[\hat{\alpha}] &= E[h(z)] \\ &\approx E\{h(\check{z}) + \nabla_z h(\check{z})(z - \check{z})\} \\ &= h(\check{z}) + E\{\nabla_z h(\check{z})(z - \check{z})\}. \end{aligned}$$

$$\begin{aligned}
E^{(2)}[\hat{\alpha}] &\approx E \left\{ h(\check{z}) + \nabla_z h(\check{z})(z - \check{z}) + \frac{1}{2}(z - \check{z})^T \nabla_z^2 h(\check{z})(z - \check{z}) \right\} \\
(4.66) \quad &= h(\check{z}) + E \{ \nabla_z h(\check{z})(z - \check{z}) \} + \frac{1}{2} E \{ (z - \check{z})^T \nabla_z^2 h(\check{z})(z - \check{z}) \}.
\end{aligned}$$

Notice that when \check{z} is chosen to be \bar{z} , the quantity $(z - \check{z})$ is zero mean Gaussian. It follows that the first order term $E \{ \nabla_z h(\check{z})(z - \check{z}) \} = 0$ in (4.65) and (4.66). Therefore, the first order Taylor approximation yields:

$$(4.67) \quad E^{(1)}[\hat{\alpha}] = h(\check{z}) = h(\bar{z}) = \bar{\alpha},$$

corresponding to zero bias.

The second-order approximation (4.66) requires computing $\nabla_z^2 h(\check{z})$, which can be obtained up to second order [30] via:

$$(4.68) \quad \nabla_z^2 h = \left[-\frac{\partial^2}{\partial \alpha^2} \Phi \right]^{-1} \left\{ \frac{\partial^3}{\partial \alpha^3} \Phi \nabla_z h^T \nabla_z h + \frac{\partial^3}{\partial \alpha^2 \partial z} \Phi^T \nabla_z h + \nabla_z h^T \frac{\partial^3}{\partial \alpha^2 \partial z} \Phi + \frac{\partial}{\partial \alpha} \nabla_z^2 \Phi \right\}.$$

Terms involved in the above expression are computed as follows:

$$\frac{\partial^3}{\partial \alpha^3} \Phi = 0.$$

Taking derivative of (4.61) with respect to z yields

$$\frac{\partial^3}{\partial \alpha^2 \partial z} \Phi = 2z^T \begin{bmatrix} I & 0 \\ 0 & 0 \end{bmatrix} = 2 \begin{bmatrix} z_1^T & 0 \end{bmatrix}.$$

Taking derivative of (4.62) with respect to z yields

$$\frac{\partial^3}{\partial \alpha \partial z^2} \Phi = 2 \begin{bmatrix} 2\alpha I & -I \\ -I & 0 \end{bmatrix}.$$

Evaluating at $z = \check{z} = \bar{z}$ and substituting into (4.68) yields:

$$(4.69) \quad \nabla_z^2 h(\bar{z}) = -\frac{1}{2\|\bar{c}\|_2^2} \left\{ \frac{-2}{\|\bar{c}\|_2^2} \begin{bmatrix} \bar{c} \\ 0 \end{bmatrix} \begin{bmatrix} \bar{\alpha} \bar{c}^T & -\bar{c}^T \end{bmatrix} + \frac{-2}{\|\bar{c}\|_2^2} \begin{bmatrix} \bar{\alpha} \bar{c} \\ -\bar{c} \end{bmatrix} \begin{bmatrix} \bar{c}^T & 0 \end{bmatrix} + 2 \begin{bmatrix} 2\bar{\alpha} I & -I \\ -I & 0 \end{bmatrix} \right\}.$$

Since $z - \check{z} \sim \mathcal{N} \left\{ 0, \begin{bmatrix} \sigma_1^2 I & 0 \\ 0 & \sigma_2^2 I \end{bmatrix} \right\}$, the second-order term in (4.66) only involves the diagonal elements of $\nabla_z^2 h(\bar{z})$. We extract the corresponding blocks from (4.69) as:

$$\begin{aligned}
\frac{\partial^2}{\partial z_1(i)^2} h(\bar{z}) &= -\frac{1}{2\|\bar{c}\|_2^2} \left\{ \frac{-4}{\|\bar{c}\|_2^2} \bar{\alpha} c_i^2 + 4\bar{\alpha} \right\}; \\
(4.70) \quad \frac{\partial^2}{\partial z_1(i)^2} h(\bar{z}) &= 0.
\end{aligned}$$

Thus

$$\begin{aligned}
E \{ (z - \bar{z})^T \nabla_z^2 h(\bar{z})(z - \bar{z}) \} &= \sum_i \sigma_1^2 \frac{\partial^2}{\partial z_1(i)^2} h(\bar{z}) \\
&= \frac{\sigma_1^2}{\|\bar{c}\|_2^2} (2\bar{\alpha} - 2\bar{\alpha}n) \\
(4.71) \quad &= 2(1-n)\bar{\alpha} \frac{\sigma_1^2}{\|\bar{c}\|_2^2}.
\end{aligned}$$

It follows that the second-order estimation for $E[\hat{\alpha}]$ is

$$(4.72) \quad E^{(2)}[\hat{\alpha}] = E^{(1)}[\hat{\alpha}] + \frac{1}{2}E\{(z - \bar{z})^T \nabla_{\bar{z}}^2 h(\bar{z})(z - \bar{z})\} = \bar{\alpha} + (1-n) \frac{\sigma_1^2}{\|\bar{c}\|_2^2} \bar{\alpha} = \left\{ 1 + (1-n) \frac{\sigma_1^2}{\|\bar{c}\|_2^2} \right\} \bar{\alpha}.$$

For $n > 1$ and reasonable signal-to-noise ratio, $E^{(2)}[\hat{\alpha}]$ implies shrinkage in magnitude, which WLOG, we refer to as “negative bias” hereafter.

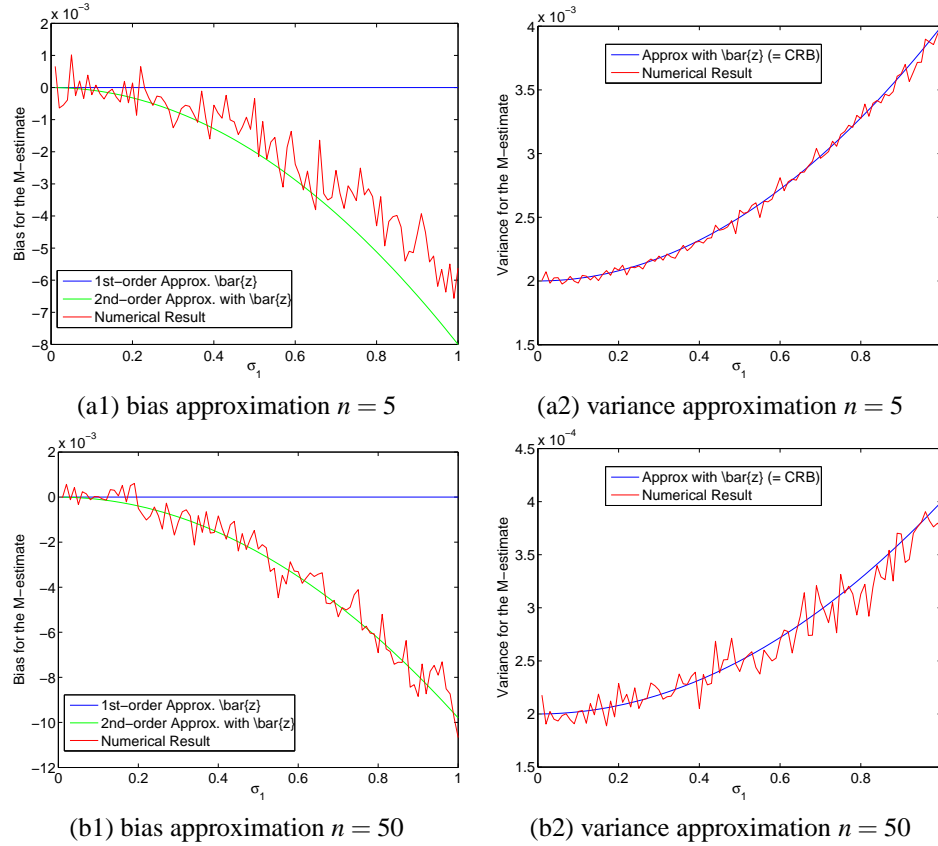


Figure 4.2: Bias and variance approximation for M-estimate obtained from expansion about $(\bar{\alpha}, \bar{z})$.

Notice that the choice of $\bar{z} = \bar{z}$ is mainly due to computation convenience (so that $z - \bar{z}$ is zero mean Gaussian). It is feasible to perform the same routine for different data point \bar{z} . [50, 108] proved that under certain regular conditions, the M-estimate is asymptotically normal with mean $\tilde{\alpha}$ where

$$E\left[\frac{\partial}{\partial \alpha} \Phi(\tilde{\alpha}, z)\right] = 0.$$

Under reasonable regularity conditions, we can exchange the order of expectation and differentiation, and take

$$\frac{\partial}{\partial \alpha} E[\Phi(\tilde{\alpha}, z)] = 0.$$

Note that $\tilde{\alpha}$ can be interpreted as a local minima for an “average” cost function $E[\Phi(\alpha, z)]$, *i.e.*,

$$(4.73) \quad \tilde{\alpha} = \arg \min_{\alpha} E[\Phi(\alpha, z)].$$

The expectation of the objective function with respect to the distribution of the observation noise

$$\begin{aligned}
E[\Phi(\alpha, z)] &= E\left[\left\| \begin{bmatrix} -\alpha I & I \end{bmatrix} \begin{bmatrix} \bar{c} + \varepsilon_1 \\ \bar{\alpha}\bar{c} + \varepsilon_2 \end{bmatrix} \right\|_2^2\right] \\
&= E\left[\begin{bmatrix} \bar{c}^T + \varepsilon_1^T & \bar{\alpha}\bar{c}^T + \varepsilon_2^T \end{bmatrix} \begin{bmatrix} -\alpha I \\ I \end{bmatrix} \begin{bmatrix} -\alpha I & I \end{bmatrix} \begin{bmatrix} \bar{c} + \varepsilon_1 \\ \bar{\alpha}\bar{c} + \varepsilon_2 \end{bmatrix}\right] \\
&= (\alpha - \bar{\alpha})^2 \|\bar{c}\|_2^2 + n(\alpha^2 \sigma_1^2 + \sigma_2^2) \\
(4.74) \quad &= (\|\bar{c}\|_2^2 + n\sigma_1^2)\alpha^2 - 2\bar{\alpha}\|\bar{c}\|_2^2\alpha + \bar{\alpha}^2\|\bar{c}\|_2^2
\end{aligned}$$

is convex quadratic in α and the minimizer reads

$$\begin{aligned}
\check{\alpha} &= \arg \min_{\alpha} E[\Phi(\alpha, z)] \\
(4.75) \quad &= \frac{\|\bar{c}\|_2^2}{\|\bar{c}\|_2^2 + n\sigma_1^2} \bar{\alpha}.
\end{aligned}$$

For simplicity, let $\beta \triangleq \frac{\|\bar{c}\|_2^2 + n\sigma_1^2}{\|\bar{c}\|_2^2}$, then $\check{\alpha} = \frac{1}{\beta} \bar{\alpha}$. Since $\beta > 1$, the expansion point $\check{\alpha}$ is a shrinkage with respect to the true scale $\bar{\alpha}$.

We can construct an expansion point $\check{z} = [\beta\bar{c}; \check{\alpha}\bar{c}]$. Then the minimizer of $\Phi(\check{z}) = \frac{1}{\beta} \bar{\alpha} = \check{\alpha}$, which satisfies the requirement (4.73).

Evaluating (4.63) at (\check{c}, \check{z}) results in:

$$\begin{aligned}
\frac{\partial}{\partial z} \alpha(\check{z}) &= -\|z_1\|_2^{-2} z^T \begin{bmatrix} 2\alpha I & -I \\ -I & 0 \end{bmatrix} \\
&= -\frac{1}{\beta^2 \|\bar{c}\|_2^2} \begin{bmatrix} \beta\bar{c}^T & \check{\alpha}\bar{c} \end{bmatrix} \begin{bmatrix} 2\check{\alpha} I & -I \\ -I & 0 \end{bmatrix} \\
(4.76) \quad &= -\frac{\bar{c}^T}{\beta^2 \|\bar{c}\|_2^2} \begin{bmatrix} \frac{2\beta-1}{\beta} \bar{\alpha} I & -\beta I \end{bmatrix}.
\end{aligned}$$

The approximated covariance of $\hat{\alpha}$ evaluated at the point $(\check{\alpha}, \check{z})$ is given by:

$$\begin{aligned}
\text{Cov}\{\hat{\alpha}\} |_{z=\check{z}, \hat{\alpha}=\check{\alpha}} &= \frac{\partial}{\partial z} \alpha(\check{z}) \text{Cov}\{z\} \frac{\partial}{\partial z} \alpha^T(\check{z}) \\
&= \beta^{-4} \|\bar{c}\|_2^{-4} \bar{c}^T \begin{bmatrix} \frac{2\beta-1}{\beta} \bar{\alpha} I & -\beta I \end{bmatrix} \begin{bmatrix} \sigma_1^2 I & 0 \\ 0 & \sigma_2^2 I \end{bmatrix} \begin{bmatrix} \frac{2\beta-1}{\beta} \bar{\alpha} I \\ -\beta I \end{bmatrix} \bar{c} \\
(4.77) \quad &= \|\bar{c}\|_2^{-2} \beta^{-4} \left(\left(2 - \frac{1}{\beta}\right)^2 \bar{\alpha}^2 \sigma_1^2 + \beta^2 \sigma_2^2 \right).
\end{aligned}$$

We know from previous analysis that the M-estimate is asymptotically unbiased, so its variance is to be bounded below by Cramér-Rao Bound asymptotically. Therefore, it is curious to find whether there exists a consistent relationship between the pre-asymptotic variance in (4.77) and the Cramér-Rao Bound, *i.e.*,

$$(4.78) \quad \|\bar{c}\|_2^{-2} \beta^{-4} \left(\left(2 - \frac{1}{\beta}\right)^2 \bar{\alpha}^2 \sigma_1^2 + \beta^2 \sigma_2^2 \right) \geq \|\bar{c}\|_2^{-2} (\bar{\alpha}^2 \sigma_1^2 + \sigma_2^2)?$$

The quantity on the right-hand-side is the Cramér-Rao Bound obtained from the statistical generative model (to be shown later).

Claim IV.1. *The covariance of the M-estimator is bounded above by the Cramér-Rao Bound. Moreover, it asymptotically approaches the Cramér-Rao Bound as $\sigma_1 \rightarrow 0$.*

Proof. To compare the left and right hand sides in (4.78), it suffices determine the sign of their difference:

$$\text{RHS} - \text{LHS} = \|\bar{c}\|_2^{-2} \beta^{-2} (\beta^6 - 4\beta^2 + 4\beta - 1) \bar{\alpha}^2 \sigma_1^2 + (\beta^4 - 1) \sigma_2^2.$$

For simplicity, we drop the positive quantity $\|\bar{c}\|_2^{-2}$ in later analysis as it does not affect the sign. Let $A \triangleq \bar{\alpha}^2 \sigma_1^2$, $B \triangleq \sigma_2^2$, and we want to determine the sign for:

$$\pi(A, B; \beta) = \beta^{-2}(\beta^6 - 4\beta^2 + 4\beta - 1)A + (\beta^4 - 1)B.$$

The polynomial $(\beta^6 - 4\beta^2 + 4\beta - 1)$ factors into

$$\beta^6 - 4\beta^2 + 4\beta - 1 = (\beta - 1)(\beta^2 + \beta - 1)(\beta^3 + 2\beta - 1).$$

By construction, $\beta > 1$, thus $(\beta^6 - 4\beta^2 + 4\beta - 1) > 0$, so π is linear in A, B with positive coefficients. Meanwhile, A, B are both positive, so $\pi(A, B; \beta) > 0$. This result translates into the claim that in the nondegenerative case ($\sigma_1 \neq 0$), the variance of the M-estimate is bounded above by the Cramér-Rao Bound. It is easy to check that when $\sigma_1 = 0$, the variance equals the Cramér-Rao Bound. \square

Now we approximate $E[\hat{\alpha}]$ with (4.65) and (4.66) by expanding corresponding terms about $(\check{\alpha}, \check{z})$.

The first order coefficient $\nabla_z h$ is obtained in (4.76), and the corresponding first-order approximation for the mean is:

$$\begin{aligned} E^{(1)}[\hat{\alpha}] &= h(\check{z}) + E\nabla_z h(\check{z})(z - \check{z}) \\ &= \frac{\bar{\alpha}}{\beta} + E \left\{ -\frac{\bar{c}^T}{\beta^2 \|\bar{c}\|_2^2} \begin{bmatrix} \frac{2\beta-1}{\beta} \bar{\alpha} I & -\beta I \end{bmatrix} \begin{bmatrix} \bar{c} + \varepsilon_1 - \beta \bar{c} \\ \bar{\alpha} \bar{c} + \varepsilon_2 - \bar{\alpha} \bar{c} \end{bmatrix} \right\} \\ &= \frac{\bar{\alpha}}{\beta} + \frac{\bar{\alpha}}{\beta^3 \|\bar{c}\|_2^2} \bar{c}^T (2\beta - 1)(\beta - 1) \bar{c} \\ &= \frac{\bar{\alpha}}{\beta} \left[1 + \frac{(2\beta - 1)(\beta - 1)}{\beta^2} \right] \\ (4.79) \quad &= \frac{3\beta^2 - 3\beta + 1}{\beta^3} \bar{\alpha}. \end{aligned}$$

Since $\beta > 1$, $(\beta - 1)^3 = \beta^3 - 3\beta^2 + 3\beta - 1 = \beta^3 - (3\beta^2 - 3\beta + 1) > 0$, and $\frac{3\beta^2 - 3\beta + 1}{\beta^3} < 1$. Equivalently, $\frac{E[\hat{\alpha}]}{\bar{\alpha}} < 1$, indicating a shrinkage in magnitude, which agrees qualitatively with the result from exact solution.

Expression in (4.79) can be rewritten as:

$$\begin{aligned} E[\hat{\alpha}] &= \frac{3\beta^2 - 3\beta + 1}{\beta^3} \bar{\alpha} \\ (4.80) \quad &= \left[1 - \frac{(\beta - 1)^3}{\beta^3} \right] \bar{\alpha}. \end{aligned}$$

Denote the signal-to-noise ratio in z_1 as $s \triangleq \frac{\|\bar{c}\|_2^2}{n\sigma_1^2}$ and

$$\frac{E[\hat{\alpha}]}{\bar{\alpha}} = 1 - \frac{1}{(s+1)^3}.$$

To approximate the bias with second-order Taylor expansion, we use (4.68) and evaluate at $(\check{\alpha} = \bar{\alpha}/\beta, \check{z})$.

$$\begin{aligned} \nabla_z^2 h(\check{z}) &= -\frac{1}{\|\beta \bar{c}\|_2^2} \left\{ -\frac{1}{\beta^2 \|\bar{c}\|_2^2} \begin{bmatrix} \beta \bar{c} \\ 0 \end{bmatrix} \begin{bmatrix} \frac{2\beta-1}{\beta} \bar{\alpha} \bar{c}^T & -\beta \bar{c}^T \end{bmatrix} \dots \right. \\ &\quad \left. -\frac{1}{\beta^2 \|\bar{c}\|_2^2} \begin{bmatrix} \frac{2\beta-1}{\beta} \bar{\alpha} \bar{c} \\ -\beta \bar{c} \end{bmatrix} \begin{bmatrix} \beta \bar{c}^T & 0 \end{bmatrix} + \begin{bmatrix} \frac{2\bar{\alpha}}{\beta} I & -I \\ -I & 0 \end{bmatrix} \right\}. \end{aligned}$$

To compute $(z - \check{z})^T \nabla_z^2 h(\check{z})(z - \check{z})$ in (4.66), it suffices to use only the diagonal blocks of $\nabla_z^2 h(\check{z})$, because the components of $z - \check{z} = \begin{bmatrix} \bar{c} + \varepsilon_1 - \beta \bar{c} \\ \bar{\alpha} \bar{c} + \varepsilon_2 - \bar{\alpha} \bar{c} \end{bmatrix} = \begin{bmatrix} (1 - \beta) \bar{c} + \varepsilon_1 \\ \varepsilon_2 \end{bmatrix}$ are independent. Partition $z - \check{z}$

into the deterministic ψ and random part η so that $\psi = \begin{bmatrix} (1-\beta)\bar{c} \\ 0 \end{bmatrix}$ and $\eta = \begin{bmatrix} \varepsilon_1 \\ \varepsilon_2 \end{bmatrix}$. Then the quadratic term in the second-order Taylor expansion in (4.66) can be written as:

$$E[(\psi + \eta)^T \nabla_z^2 h(\check{z})(\psi + \eta)] = \psi^T \nabla_z^2 h(\check{z})\psi + E[\eta^T \nabla_z^2 h(\check{z})\eta],$$

where expectation of cross terms between ψ and η are dropped since η is zero-mean.

The diagonal portion of $\nabla_z^2 h(\check{z})$ reads:

$$(4.81) \quad \nabla_z^2 h(\check{z}) = \frac{2}{\beta^2 \|\bar{c}\|_2^2} \left\{ \frac{1}{\beta^2 \|\bar{c}\|_2^2} \begin{bmatrix} \beta(2\beta-1)\bar{\alpha}\bar{c}\bar{c}^T & \\ & 0 \end{bmatrix} - \begin{bmatrix} \frac{\bar{\alpha}}{\beta} I & \\ & 0 \end{bmatrix} \right\}.$$

It follows that

$$(4.82) \quad \begin{aligned} \psi^T \nabla_z^2 h(\check{z})\psi &= \frac{2(\beta-1)^2}{\|\bar{c}\|_2^2} \left\{ \frac{(\beta-1)\|\bar{c}\|_2^4}{\beta^2 \|\bar{c}\|_2^2} \bar{\alpha} - \frac{\|\bar{c}\|_2^2}{\beta} \bar{\alpha} \right\} \\ &= \frac{2(\beta-1)^2}{\beta^3} \left[\frac{(2\beta-1)}{\beta} - 1 \right] \bar{\alpha} \\ &= \frac{2(\beta-1)^3}{\beta^4} \bar{\alpha}. \end{aligned}$$

$$(4.83) \quad \begin{aligned} \eta^T \nabla_z^2 h(\check{z})\eta &= \frac{2}{\beta^2 \|\bar{c}\|_2^2} \left\{ \frac{\beta(2\beta-1)\sigma_1^2 \|\bar{c}\|_2^2}{\beta^3 \|\bar{c}\|_2^2} \bar{\alpha} - \frac{n\sigma_1^2}{\beta} \bar{\alpha} \right\} \\ &= \frac{2\sigma_1^2}{\beta^6 \|\bar{c}\|_2^2} [(2\beta-1) - n\beta] \bar{\alpha}. \end{aligned}$$

Summing (4.82) and (4.83) yields:

$$(4.84) \quad E[(z - \check{z})^T \nabla_z^2 h(\check{z})(z - \check{z})] = \frac{2(\beta-1)^3}{\beta^4} \bar{\alpha} + \frac{2\sigma_1^2}{\beta^6 \|\bar{c}\|_2^2} [(2\beta-1) - n\beta] \bar{\alpha}.$$

Combining (4.84) with the first order estimation of $E[\hat{\alpha}]$, we obtain the second order approximation for $E[\hat{\alpha}]$ as:

$$(4.85) \quad \begin{aligned} E^{(2)}[\hat{\alpha}] &= h(\check{z}) + E \left\{ \nabla_z h(\check{z})(z - \check{z}) + \frac{1}{2} (z - \check{z})^T \nabla_z^2 h(\check{z})(z - \check{z}) \right\} \\ &= E^{(1)}[\hat{\alpha}] + \frac{1}{2} E[(z - \check{z})^T \nabla_z^2 h(\check{z})(z - \check{z})] \\ &= \left\{ \frac{\beta^3 - (\beta-1)^3}{\beta^3} + \frac{(\beta-1)^3}{\beta^4} \right\} \bar{\alpha} + \frac{(2-n)\beta-1}{\beta^5} \frac{\sigma_1^2}{\|\bar{c}\|_2^2} \bar{\alpha} \\ &= \frac{\beta^4 - (\beta-1)^4}{\beta^4} \bar{\alpha} + \frac{(2-n)\beta-1}{\beta^5} \frac{\sigma_1^2}{\|\bar{c}\|_2^2} \bar{\alpha}. \end{aligned}$$

Recall that $\beta = \frac{\|\bar{c}\|_2^2 + n\sigma_1^2}{\|\bar{c}\|_2^2}$, so for reasonable SNR, $\frac{(2-n)\beta-1}{\beta} \approx 1-n$. Using the $s = \frac{\|\bar{c}\|_2^2}{n\sigma_1^2}$, we can rewrite $E^{(2)}[\hat{\alpha}]$ approximately as:

$$(4.86) \quad E^{(2)}[\hat{\alpha}] = \left[1 - \frac{1}{(s+1)^4} + \frac{(1-n)s^3}{n(1+s)^4} \right] \bar{\alpha}.$$

Notice that when SNR is high (large s), then

$$\begin{aligned}
 E^{(2)}[\hat{\alpha}] &= \left[1 - \frac{1}{(s+1)^4} + \frac{(1-n)s^3}{n(1+s)^4}\right] \bar{\alpha} \\
 &\approx \left[1 + \frac{1-n}{n(1+s)}\right] \bar{\alpha} \\
 &= \left[1 - \frac{1-n}{n} \frac{n\sigma_1^2}{\|\bar{c}\|_2^2 + n\sigma_1^2}\right] \bar{\alpha} \\
 (4.87) \quad &\approx \left[1 + (1-n) \frac{\sigma_1^2}{\|\bar{c}\|_2^2 + n\sigma_1^2}\right] \bar{\alpha},
 \end{aligned}$$

which closely resembles the result (4.72) obtained from expanding about noiseless data \bar{z} . In fact, for high enough SNR, $\frac{\|\bar{c}\|_2^2 + n\sigma_1^2}{\sigma_1^2} \approx \frac{\|\bar{c}\|_2^2}{\sigma_1^2}$ so that (4.87) and (4.72) are approximately equal. This relation is expected, as for small SNR, $\tilde{z} \approx \bar{z}$ and $\check{\alpha} \approx \bar{\alpha}$, the small error analysis is essentially performed on the same neighborhood!

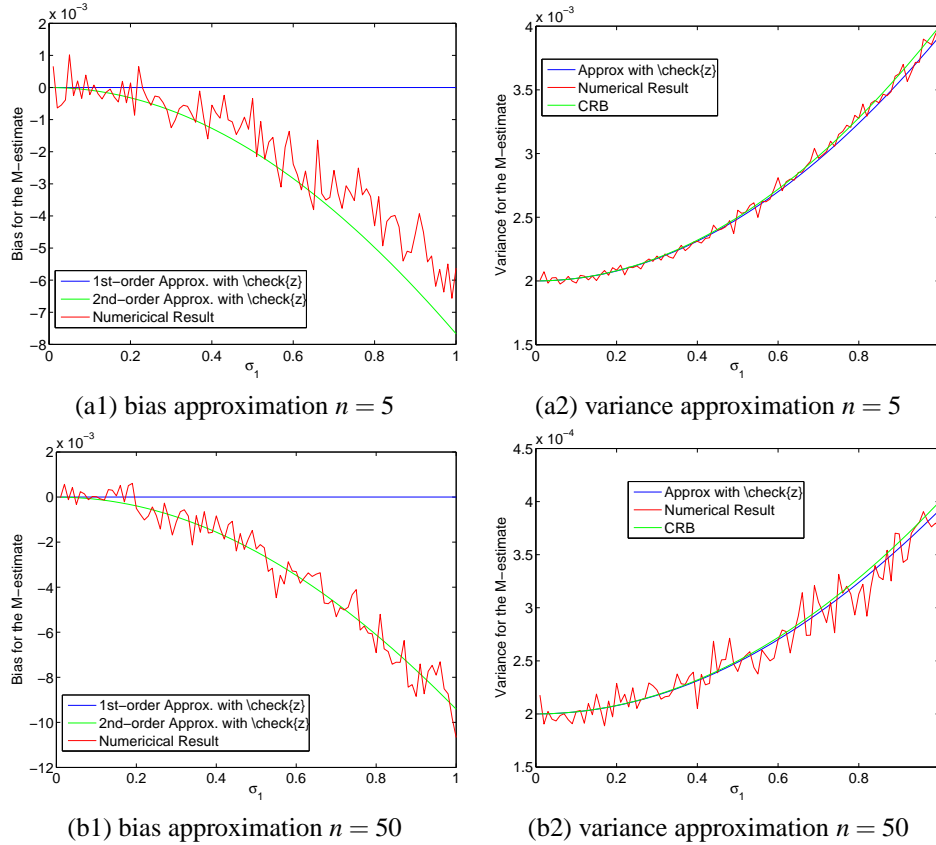


Figure 4.3: Bias and variance approximation of M-estimate obtained from expansion about $(\check{\alpha}, \check{z})$.

ML Estimator for the Statistical Model

The maximum likelihood estimator from (4.23) aims to jointly estimate c and α via:

$$(4.88) \quad [\hat{\alpha}, \hat{c}] = \arg \min_{\alpha, c} \frac{1}{\sigma_1^2} \|z_1 - c\|_2^2 + \frac{1}{\sigma_2^2} \|z_2 - \alpha c\|_2^2.$$

Note that conditioned on α , (4.88) is quadratic in c with the solution $\hat{c}(\alpha, z)$ given by:

$$\begin{aligned}
 \hat{c} &= \left\{ \begin{bmatrix} I \\ \alpha I \end{bmatrix}^T \begin{bmatrix} \frac{1}{\sigma_1^2} I & 0 \\ 0 & \frac{1}{\sigma_2^2} I \end{bmatrix} \begin{bmatrix} I \\ \alpha I \end{bmatrix} \right\}^{-1} \begin{bmatrix} I & \alpha I \end{bmatrix} \begin{bmatrix} \frac{1}{\sigma_1^2} I & 0 \\ 0 & \frac{1}{\sigma_2^2} I \end{bmatrix} z \\
 &= \left(\frac{1}{\sigma_1^2} + \frac{\alpha^2}{\sigma_2^2} \right)^{-1} \left(\frac{1}{\sigma_1^2} z_1 + \frac{\alpha^2}{\sigma_2^2} z_2 \right) \\
 (4.89) \quad &= \frac{1}{\alpha^2 \sigma_1^2 + \sigma_2^2} (\sigma_2^2 z_1 + \alpha \sigma_1^2 z_2).
 \end{aligned}$$

Remark:

- In the limiting case when $\sigma_1 \rightarrow 0$ (with non-vanishing σ_2), z_1 is a noise-free observation of c , it is natural to estimate c solely on z_1 as (4.89) reduces to

$$\lim_{\sigma_1 \rightarrow 0} \hat{c} = z_1,$$

which coincides with (4.47) in the conventional method. On the other hand, as the noise level in z_2 becomes small relative to that in z_1 ($\sigma_2 \rightarrow 0$ with non-vanishing σ_1), the estimate reduces to:

$$\lim_{\sigma_2 \rightarrow 0} \hat{c} = z_2 / \alpha,$$

which corresponds to the case of estimating c solely from z_2 .

More precisely,

$$\begin{aligned}
 \lim_{\sigma_1 \rightarrow 0} \hat{c} &= z_1 \quad \text{as } \sigma_1 / \sigma_2 \rightarrow 0; \\
 \lim_{\sigma_2 \rightarrow 0} \hat{c} &= z_2 / \alpha \quad \text{as } \sigma_1 / \sigma_2 \rightarrow \infty.
 \end{aligned}
 \tag{4.90}$$

- It is easy to check that the estimator in (4.89) is unbiased with variance

$$\text{Var}\{\hat{c}\} = \frac{\sigma_1^2 \sigma_2^2}{\alpha^2 \sigma_1^2 + \sigma_2^2} I = \frac{\sigma_1^2}{1 + \alpha^2 \frac{\sigma_1^2}{\sigma_2^2}} I.$$

It immediately follows that this quantity is upper-bounded by the covariance $\sigma_1^2 I$ of the estimator for c (4.47) resulting from conventional methods.

Now we can plug in the expression of \hat{c} in (4.89) and (4.88) reduces to a minimization problem over α only:

$$\begin{aligned}
 \hat{\alpha} &= \arg \min_{\alpha} \Psi(\alpha, z) \\
 (4.91) \quad &= \arg \min_{\alpha} \frac{1}{\alpha^2 \sigma_1^2 + \sigma_2^2} \|\alpha z_1 - z_2\|_2^2.
 \end{aligned}$$

This function Ψ is nonlinear in α . Note that $\Psi \geq 0$. In the case of noise-free observation $z = \bar{z}$, $\bar{\alpha}$ achieves the zero value and is the global minimizer (we will justify this more precisely later). Therefore, we can utilize the techniques for M-estimate as before, and analyze the behavior of $\hat{\alpha}$ in the neighborhood $\hat{\alpha}(\bar{z}) = \bar{\alpha}$.

Let $\hat{\alpha}$ be the minimizer of the function $\Psi(\alpha, z)$, then it is true that

$$\frac{\partial}{\partial \alpha} \Psi(\alpha, z) = \frac{\partial}{\partial \alpha} \frac{1}{\alpha^2 \sigma_1^2 + \sigma_2^2} \|\alpha z_1 - z_2\|_2^2 = 0 \quad \text{for } \forall z.$$

$$\begin{aligned}
 \frac{\partial}{\partial \alpha} \Psi(\alpha, z) &= \frac{1}{(\alpha^2 \sigma_1^2 + \sigma_2^2)^2} (\alpha z_1 - z_2)^T [2z_1 (\alpha^2 \sigma_1^2 + \sigma_2^2) - 2\alpha \sigma_1^2 (\alpha z_1 - z_2)] \\
 (4.92) \quad &= \frac{2}{(\alpha^2 \sigma_1^2 + \sigma_2^2)^2} z^T \begin{bmatrix} \alpha I \\ -I \end{bmatrix} \begin{bmatrix} \sigma_2^2 I & \alpha \sigma_1^2 \end{bmatrix} z.
 \end{aligned}$$

Let $Q \triangleq \begin{bmatrix} \alpha I \\ -I \end{bmatrix} \begin{bmatrix} \sigma_2^2 I & \alpha \sigma_1^2 \end{bmatrix} = \begin{bmatrix} \alpha \sigma_2^2 I & \alpha^2 \sigma_1^2 I \\ -\sigma_2^2 I & -\alpha \sigma_1^2 I \end{bmatrix}$, then the derivative of $\frac{\partial}{\partial \alpha} \Psi$ with respect to z is given by:

$$(4.93) \quad \begin{aligned} \frac{\partial^2}{\partial \alpha \partial z} \Psi &= \frac{2}{(\alpha^2 \sigma_1^2 + \sigma_2^2)^2} z^T (Q + Q^T) \\ &= \frac{2}{(\alpha^2 \sigma_1^2 + \sigma_2^2)^2} z^T \begin{bmatrix} 2\alpha \sigma_2^2 I & (\alpha^2 \sigma_1^2 - \sigma_2^2) I \\ (\alpha^2 \sigma_1^2 - \sigma_2^2) I & -2\alpha \sigma_1^2 I \end{bmatrix}. \end{aligned}$$

Evaluating (4.93) at $z = \bar{z}$ and $\alpha = \bar{\alpha}$ yields:

$$(4.94) \quad \frac{\partial^2}{\partial \alpha \partial z} \Psi(\bar{\alpha}, \bar{z}) = \frac{2}{\bar{\alpha}^2 \sigma_1^2 + \sigma_2^2} \bar{c}^T \begin{bmatrix} \bar{\alpha} I & -I \end{bmatrix}.$$

Now we compute the derivative of $\frac{\partial}{\partial \alpha} \Psi$ with respect to α and evaluate at the minimizer $\hat{\alpha} = \bar{\alpha}$ with $z = \bar{z}$:

$$\begin{aligned} \frac{\partial^2}{\partial \alpha^2} \Psi &= 2 \frac{\partial}{\partial \alpha} \left\{ \frac{(\alpha z_1 - z_2)^T (\sigma_2^2 z_1 + \alpha \sigma_1^2 z_2)}{(\alpha^2 \sigma_1^2 + \sigma_2^2)^2} \right\} \\ &= 2 \left\{ -2 \frac{2\alpha \sigma_1^2}{(\alpha^2 \sigma_1^2 + \sigma_2^2)^3} (\alpha z_1 - z_2)^T (\sigma_2^2 z_1 + \alpha \sigma_1^2 z_2) \dots \right. \\ &\quad \left. + \frac{1}{(\alpha^2 \sigma_1^2 + \sigma_2^2)^2} [z_1^T (\sigma_2^2 z_1 + \alpha \sigma_1^2 z_2) + (\alpha z_1 - z_2)^T \sigma_1^2 z_2] \right\}. \end{aligned}$$

This is a convenient form to be evaluated at $z = \bar{z}$, and we obtain:

$$(4.95) \quad \frac{\partial^2}{\partial \alpha^2} \Psi(\bar{z}) = \frac{2}{\bar{\alpha}^2 \sigma_1^2 + \sigma_2^2} \|\bar{c}\|_2^2.$$

To prepare for future use, we simplify the general form of (4.95) into:

$$(4.96) \quad \frac{\partial^2}{\partial \alpha^2} \Psi = \frac{2}{(\alpha^2 \sigma_1^2 + \sigma_2^2)^3} z^T \begin{bmatrix} (-3\alpha^2 \sigma_1^2 + \sigma_2^2) \sigma_2^2 I & (3\sigma_2^2 - \alpha^2 \sigma^2) \alpha \sigma_1^2 I \\ (3\sigma_2^2 - \alpha^2 \sigma^2) \alpha \sigma_1^2 I & (3\alpha^2 \sigma_1^2 - \sigma_2^2) \sigma_1^2 I \end{bmatrix} z.$$

Estimating $\frac{\partial}{\partial z} \alpha$ yields:

$$(4.97) \quad \begin{aligned} \frac{\partial}{\partial z} \hat{\alpha}|_{\bar{z}, \bar{\alpha}} &= -\frac{\partial^2}{\partial \alpha^2} \Psi^{-1} \frac{\partial^2}{\partial \alpha \partial z} \Psi \\ &= -\frac{1}{\|\bar{c}\|_2^2} \bar{c}^T \begin{bmatrix} \bar{\alpha} I & -I \end{bmatrix}. \end{aligned}$$

The covariance evaluated at $(\bar{\alpha}, \bar{z})$ is

$$(4.98) \quad \begin{aligned} \text{Cov}\{\hat{\alpha}\}|_{(\bar{z}, \bar{\alpha})} &= \frac{\partial}{\partial z} \alpha(\bar{z}) \text{Cov}\{z\} \frac{\partial}{\partial z} \alpha^T(\bar{z}) \\ &= \|\bar{c}\|_2^{-4} \bar{c}^T \begin{bmatrix} \bar{\alpha} I & -I \end{bmatrix} \begin{bmatrix} \sigma_1^2 I & 0 \\ 0 & \sigma_2^2 I \end{bmatrix} \begin{bmatrix} \bar{\alpha} I \\ -I \end{bmatrix} \bar{c} \\ &= \|\bar{c}\|_2^{-2} (\bar{\alpha}^2 \sigma_1^2 + \sigma_2^2). \end{aligned}$$

Lower Bound for Covariance From Cramér-Rao Bound

The negative log-likelihood is given as the objective function in (4.88). It is straight-forward to compute the sub-matrices for the Fisher-Information Matrix.

$$\frac{\partial}{\partial \alpha} \Lambda = -\frac{1}{\sigma_2^2} (\alpha c - z_2)^T c;$$

$$\begin{aligned}\frac{\partial^2}{\partial \alpha^2} \Lambda &= -\frac{1}{\sigma_2^2} \mathbf{c}^T \mathbf{c}; \\ \frac{\partial^2}{\partial \alpha \partial \mathbf{c}} \Lambda &= \frac{-1}{\sigma_2^2} (2\alpha \mathbf{c}^T - \mathbf{z}_2^T),\end{aligned}$$

resulting in

$$E\left[\frac{\partial^2}{\partial \alpha^2} \mathbf{c}\right] = \frac{-1}{\sigma_2^2} \alpha \mathbf{c}^T.$$

The Fisher-information matrix (FIM) is thus given by:

$$FIM = \frac{1}{\sigma_2^2} \begin{bmatrix} \mathbf{c}^T \mathbf{c} & \alpha \mathbf{c}^T \\ \alpha \mathbf{c} & (\alpha^2 + \frac{\sigma_2^2}{\sigma_1^2}) I \end{bmatrix}.$$

Invoking block-matrix inversion, we obtain:

$$\begin{aligned}(4.99) \quad \text{Cov}\{\hat{\alpha}\} &\geq \sigma^2 [\mathbf{c}^T \mathbf{c} - \alpha \mathbf{c}^T (\alpha^2 + \frac{\sigma_2^2}{\sigma_1^2})^{-1} \alpha \mathbf{c}] \\ &= \|\mathbf{c}\|_2^{-2} (\alpha^2 \sigma_1^2 + \sigma_2^2).\end{aligned}$$

Since the ML estimator is known to be asymptotically unbiased, the coincidence between (4.98) and (4.99) justifies the well-known fact that the ML estimator is asymptotically efficient (thus is asymptotically a uniformly minimal variance and unbiased estimator (UMVUE)).

Approximate Bias of the ML Estimator

Notwithstanding the value of asymptotic analysis for the ML estimator, it is often of great interest to analyze the bias and variance before the estimator enters the asymptotic zone. Hereafter, we focus on deriving analytical approximation for the bias of the ML estimator. As in the covariance analysis previously, we assume the estimate is over continuous parameter's α and is computed by “completely” maximizing the objective function (likelihood in this case) without “stopping rules” that terminates the iterations before the maximum is reached. We derive the approximation using implicit function theorem, the Taylor expansion (with different orders of approximation accuracy), and the chain rule.

The objective function Ψ in (4.91) implicitly defines the M-estimate $\hat{\alpha}$ as a function of \mathbf{z} . Yet the absence of an explicit analytical expression of the form $\hat{\alpha} = h(\mathbf{z})$ (as the one in (4.49)) makes it difficult to study the mean of $\hat{\alpha}$ directly. As in the previous section, we apply Taylor expansion, chain rules and implicit function theorem to estimate the bias with the first and second order approximation given by:

$$(4.100) \quad E[\hat{\alpha}] \approx h(\check{\mathbf{z}}) + E\{\nabla_{\mathbf{z}} h(\check{\mathbf{z}})(\mathbf{z} - \check{\mathbf{z}})\}.$$

$$(4.101) \quad E[\hat{\alpha}] \approx h(\check{\mathbf{z}}) + E\left\{\nabla_{\mathbf{z}} h(\check{\mathbf{z}})(\mathbf{z} - \check{\mathbf{z}}) + \frac{1}{2}(\mathbf{z} - \check{\mathbf{z}})^T \nabla_{\mathbf{z}}^2 h(\check{\mathbf{z}})(\mathbf{z} - \check{\mathbf{z}})\right\}.$$

We now determine the point of expansion $\check{\mathbf{z}}$ and the approximation for first (linear) and second order (Hessian) coefficients $\nabla_{\mathbf{z}} h$, $\nabla_{\mathbf{z}}^2 h$. To obtain the best choice for $\check{\alpha}$

$$(4.102) \quad \check{\alpha} = \arg \min_{\alpha} E[\Psi(\alpha, \mathbf{z})],$$

where $\check{\alpha}$ and $\check{\mathbf{z}}$ in the Taylor expansions are related by $\check{\alpha} = h(\check{\mathbf{z}})$. We compute $E[\Psi(\alpha, \mathbf{z})]$ as follows:

$$E[\Psi(\alpha, \mathbf{z})] = \frac{1}{\alpha^2 \sigma_1^2 + \sigma_2^2} \sum_{i=1}^n (\alpha z_1(i) - z_2(i))^2.$$

For each index i ,

$$\begin{aligned}
E[(\alpha z_1(i) - z_2(i))^2] &= E[\alpha^2 z_1(i)^2 - 2\alpha z_1(i)z_2(i) + z_2(i)^2] \\
&= \alpha^2(\bar{c}_i^2 + \sigma_1^2) - 2\alpha\bar{\alpha}\bar{c}_i^2 + \bar{\alpha}^2\bar{c}_i^2 + \sigma_2^2 \\
(4.103) \qquad \qquad \qquad &= (\alpha^2 - 2\alpha\bar{\alpha} + \bar{\alpha}^2)\bar{c}_i^2 + (\alpha^2\sigma_1^2 + \sigma_2^2),
\end{aligned}$$

where \bar{c}_i and $\bar{\alpha}$ are the underlying “true” parameter values.

Substituting (4.103) yields:

$$(4.104) \qquad E[\Psi(\alpha, z)] = \frac{1}{\alpha^2\sigma_1^2 + \sigma_2^2} (\alpha - \bar{\alpha})^2 \|\bar{c}\|_2^2 + n.$$

Even though $E[\Psi(\alpha, z)]$ is nonlinear in α , its global minimizer is immediately observed as $\alpha = \bar{\alpha}$, because $E[\Psi(\bar{\alpha}, z)] = n$ achieves the lower bound for $E[\Psi(\alpha, z)]$ as a function of α . Thus we have found the proper point to expand around $\check{\alpha} = \bar{\alpha}$.

Note that when noise free data is observed, *i.e.*, $z = \bar{z}$, the minimizer $\hat{\alpha}$ in (4.91) is obtained as:

$$\begin{aligned}
\hat{\alpha}(\bar{z}) &= \arg \min_{\alpha} \frac{1}{\alpha^2\sigma_1^2 + \sigma_2^2} \|\alpha\bar{z}_1 - \bar{z}_2\|_2^2 \\
&= \arg \min_{\alpha} \frac{1}{\alpha^2\sigma_1^2 + \sigma_2^2} \|\alpha\bar{c} - \bar{\alpha}\bar{c}\|_2^2 \\
(4.105) \qquad \qquad \qquad &= \arg \min_{\alpha} \frac{(\alpha - \bar{\alpha})^2 \|\bar{c}\|_2^2}{\alpha^2\sigma_1^2 + \sigma_2^2}.
\end{aligned}$$

Note this function is nonnegative, its global minimizer is obtained at $\alpha = \bar{\alpha}$, *i.e.*, $h(\bar{z}) = \bar{\alpha} = \check{\alpha}$. This indicates that $\check{z} = \bar{z}$ is the proper choice to expand h around, without requiring to know the precise value of $\bar{\alpha}$.

In this case, the bias analysis with first-order Taylor expansion as in (4.100) is simple by noting that $(z - \bar{z}) \sim \mathcal{N}\left(0, \begin{bmatrix} \sigma_1^2 I & \\ & \sigma_2^2 I \end{bmatrix}\right)$, so that

$$\begin{aligned}
E[\hat{\alpha}] &= h(\bar{z}) + E\{\nabla_z h(\bar{z})(z - \bar{z})\} \\
(4.106) \qquad \qquad \qquad &= \bar{\alpha}.
\end{aligned}$$

This states that the estimator is unbiased if we approximate its first moment up to first order dependence on the data.

The first order expansion is usually sufficient in practice and has been extensively used. However, there are situations where (4.100) may be inadequate. We next derive a mean approximation based on the second-order Taylor expansion (4.101) which is expected to be more accurate, but also computationally more intensive.

The first two (0th and 1st order) terms in (4.101) are (4.100), so it suffices to study the Hessian ∇_z^2 .

For scalar α , we follow the simplified expression in [30] to obtain the Hessian of $h(z)$ as:

$$(4.107) \qquad \nabla_z^2 h = \left[-\frac{\partial^2}{\partial \alpha^2} \Psi\right]^{-1} \left\{ \frac{\partial^3}{\partial \alpha^3} \Psi \nabla_z h^T \nabla_z h + \frac{\partial^3}{\partial \alpha^2 \partial z} \Psi^T \nabla_z h + \nabla_z h^T \frac{\partial^3}{\partial \alpha^2 \partial z} \Psi + \frac{\partial}{\partial \alpha} \nabla_z^2 \Psi \right\}.$$

Some of the key ingredients are already available: $\nabla_z h$ is given in (4.97) as well as $\frac{\partial^2}{\partial \alpha^2} \Psi$ in (4.95) (before evaluation) and $\frac{\partial^2}{\partial \alpha \partial z} \Psi$ in (4.93). We still need to compute $\frac{\partial^3}{\partial \alpha^3} \Psi(\bar{\alpha}, \bar{z})$, $\frac{\partial^3}{\partial \alpha^2} \partial z \Psi(\bar{\alpha}, \bar{z})$ and $\frac{\partial}{\partial \alpha} \nabla_z^2 \Psi$. Evaluating (4.95) at $(\bar{\alpha}, \bar{z})$ yields:

$$\frac{\partial^2}{\partial \alpha^2} \Psi(\bar{\alpha}, \bar{z}) = \frac{2\|\bar{c}\|_2^2}{\bar{\alpha}^2\sigma_1^2 + \sigma_2^2}.$$

Taking derivative of (4.96) with respect to z yields:

$$(4.108) \qquad \frac{\partial^3}{\partial \alpha^2 \partial z} \Psi = \frac{4}{(\alpha^2\sigma_1^2 + \sigma_2^2)^3} z^T \begin{bmatrix} (-3\alpha^2\sigma_1^2 + \sigma_2^2)\sigma_2^2 I & (3\sigma_2^2 - \alpha^2\sigma^2)\alpha\sigma_1^2 I \\ (3\sigma_2^2 - \alpha^2\sigma^2)\alpha\sigma_1^2 I & (3\alpha^2\sigma_1^2 - \sigma_2^2)\sigma_1^2 I \end{bmatrix}.$$

Evaluating (4.108) at $(\bar{\alpha}, \bar{z})$ yields:

$$(4.109) \quad \frac{\partial^3}{\partial \alpha^2 \partial z} \Psi(\bar{\alpha}, \bar{z}) = \frac{4}{(\bar{\alpha}^2 \sigma_1^2 + \sigma_2^2)^3} \bar{c}^T \begin{bmatrix} (\sigma_2^4 - \bar{\alpha}^4 \sigma_1^4)I & 2\bar{\alpha} \sigma_1^2 (\bar{\alpha}^2 \sigma_1^2 + \sigma_2^2)I \end{bmatrix}$$

Taking derivative of (4.95) with respect to α yields:

$$(4.110) \quad \begin{aligned} \frac{\partial^3}{\partial \alpha^3} \Psi &= \frac{-12\alpha \sigma_1^2}{(\alpha^2 \sigma_1^2 + \sigma_2^2)^3} z_1^T (\sigma_2^2 z_1 + \alpha \sigma_1^2 z_2) + \frac{2}{(\alpha^2 \sigma_1^2 + \sigma_2^2)^3} \left[-4\alpha \sigma_1^2 z_1^T (\sigma_2^2 z_1 + \alpha \sigma_1^2 z_2) + \dots \right. \\ &\quad \left. + 2\alpha \sigma_1^2 z_1^T (\sigma_2^2 z_1 + \alpha \sigma_1^2 z_2) + 2(\alpha^2 \sigma_1^2 + \sigma_2^2) \sigma_1^2 z_1^T z_2 \right]. \end{aligned}$$

Evaluating (4.110) at $(\bar{\alpha}, \bar{z})$ yields:

$$(4.111) \quad \frac{\partial^3}{\partial \alpha^3} \Psi(\bar{\alpha}, \bar{z}) = \frac{-12\bar{\alpha} \sigma_1^2 \|\bar{c}\|_2^2}{(\bar{\alpha}^2 \sigma_1^2 + \sigma_2^2)^3}.$$

The term $\frac{\partial}{\partial \alpha} \nabla_z^2 \Psi$ is obtained by taking derivative of $\frac{\partial^2}{\partial \alpha \partial z} \Psi$ in (4.93) with respect to z as:

$$(4.112) \quad \frac{\partial}{\partial \alpha} \nabla_z^2 \Psi = 2 \left(\frac{1}{\alpha^2 \sigma_1^2 + \sigma_2^2} \right)^2 \begin{bmatrix} 2\alpha \sigma_2^2 I & (\alpha^2 \sigma_1^2 - \sigma_2^2)I \\ (\alpha \sigma_1^2 - \sigma_2^2)I & -2\alpha \sigma_1^2 I \end{bmatrix}.$$

Evaluating at $\bar{\alpha}$ yields:

$$(4.113) \quad \frac{\partial}{\partial \alpha} \nabla_z^2 \Psi(\bar{\alpha}) = 2 \left(\frac{1}{\bar{\alpha}^2 \sigma_1^2 + \sigma_2^2} \right)^2 \begin{bmatrix} 2\bar{\alpha} \sigma_2^2 I & (\bar{\alpha}^2 \sigma_1^2 - \sigma_2^2)I \\ (\bar{\alpha} \sigma_1^2 - \sigma_2^2)I & -2\bar{\alpha} \sigma_1^2 I \end{bmatrix}.$$

Substituting the expressions of all components into the right-hand-side of (4.107) yields:

$$(4.114) \quad \begin{aligned} \nabla_z^2 h(\bar{z}) &= -\frac{\bar{\alpha}^2 \sigma_1^2 + \sigma_2^2}{2 \|\bar{c}\|_2^2} \left\{ \frac{-12\bar{\alpha} \sigma_1^2}{(\bar{\alpha}^2 \sigma_1^2 + \sigma_2^2)^2 \|\bar{c}\|_2^2} \begin{bmatrix} \bar{\alpha} I \\ -I \end{bmatrix} \bar{c} \bar{c}^T \begin{bmatrix} \bar{\alpha} I & -I \end{bmatrix} + \dots \right. \\ &\quad - \frac{4}{(\bar{\alpha}^2 \sigma_1^2 + \sigma_2^2)^3 \|\bar{c}\|_2^2} \begin{bmatrix} (\sigma_2^4 - \bar{\alpha}^4 \sigma_1^4)I \\ 2\bar{\alpha} \sigma_1^2 (\bar{\alpha}^2 \sigma_1^2 + \sigma_2^2)I \end{bmatrix} \bar{c} \bar{c}^T \begin{bmatrix} \bar{\alpha} I & -I \end{bmatrix} + \dots \\ &\quad \left. - \frac{4}{(\bar{\alpha}^2 \sigma_1^2 + \sigma_2^2)^3 \|\bar{c}\|_2^2} \begin{bmatrix} \bar{\alpha} I \\ -I \end{bmatrix} \bar{c} \bar{c}^T \begin{bmatrix} (\sigma_2^4 - \bar{\alpha}^4 \sigma_1^4)I & 2\bar{\alpha} \sigma_1^2 (\bar{\alpha}^2 \sigma_1^2 + \sigma_2^2)I \end{bmatrix} + \dots \right. \\ &\quad \left. + 2 \left(\frac{1}{\bar{\alpha}^2 \sigma_1^2 + \sigma_2^2} \right)^2 \begin{bmatrix} 2\bar{\alpha} \sigma_2^2 I & (\bar{\alpha}^2 \sigma_1^2 - \sigma_2^2)I \\ (\bar{\alpha} \sigma_1^2 - \sigma_2^2)I & -2\bar{\alpha} \sigma_1^2 I \end{bmatrix} \right\}. \end{aligned}$$

The second order term in (4.101) depends on the Hessian $\nabla_z^2 h(\bar{z})$ via $(z - \bar{z})^T \nabla_z^2 h(\bar{z})(z - \bar{z})$ since $\bar{z} = \bar{z}$, where $z - \bar{z}$ are exactly the noise component $\varepsilon \sim \mathcal{N}\left(0, \begin{bmatrix} \sigma_1^2 I & \\ & \sigma_2^2 I \end{bmatrix}\right)$. Because the elements of ε are mutually independent, $E \{(z - \bar{z})^T \nabla_z^2 h(\bar{z})(z - \bar{z})\}$ only depends on the diagonal elements of the Hessian $\nabla_z^2 h(\bar{z})$.

When a component is located in the z_1 portion of z , the noise component $\varepsilon(i) \sim \mathcal{N}(0, \sigma_1^2)$, and the corresponding element in the Hessian is:

$$(4.115) \quad \frac{\partial^2}{\partial z_1(i)^2} h(\bar{z}) = -\frac{1}{2 \|\bar{c}\|_2^2} \left\{ \frac{-12\bar{\alpha}^3 \sigma_1^2 c_i^2}{(\bar{\alpha}^2 \sigma_1^2 + \sigma_2^2) \|\bar{c}\|_2^2} - \frac{8\bar{\alpha}(\sigma_2^4 - \bar{\alpha}^4 \sigma_1^4) c_i^2}{(\bar{\alpha}^2 \sigma_1^2 + \sigma_2^2)^2 \|\bar{c}\|_2^2} + 4\bar{\alpha} \sigma_2^2 \right\}.$$

Similarly,

$$(4.116) \quad \frac{\partial^2}{\partial z_2(i)^2} h(\bar{z}) = -\frac{1}{2 \|\bar{c}\|_2^2} \left\{ \frac{-12\bar{\alpha} \sigma_1^2 c_i^2}{(\bar{\alpha}^2 \sigma_1^2 + \sigma_2^2) \|\bar{c}\|_2^2} + \frac{16\bar{\alpha} \sigma_1^2 c_i^2}{(\bar{\alpha}^2 \sigma_1^2 + \sigma_2^2) \|\bar{c}\|_2^2} - 4\bar{\alpha} \sigma_1^2 \right\}.$$

Combining the above yields:

$$\begin{aligned}
 E[\varepsilon^T \nabla_{\bar{z}}^2 h(\bar{z}) \varepsilon] &= \sigma_1^2 \sum_{i=1}^n \frac{\partial^2}{\partial z_1(i)^2} h(\bar{z}) + \sigma_2^2 \sum_{i=1}^n \frac{\partial^2}{\partial z_2(i)^2} h(\bar{z}) \\
 &= -\frac{1}{2 \|\bar{c}\|_2^2} \left\{ \frac{-12\bar{\alpha}^3 \sigma_1^4}{(\bar{\alpha}^2 \sigma_1^2 + \sigma_2^2)} - \frac{8\bar{\alpha}(\sigma_2^4 - \bar{\alpha}^4 \sigma_1^4) \sigma_1^2}{(\bar{\alpha}^2 \sigma_1^2 + \sigma_2^2)^2} \right\} \dots \\
 &\quad -\frac{1}{2 \|\bar{c}\|_2^2} \left\{ \frac{-12\bar{\alpha} \sigma_1^2 \sigma_2^2}{(\bar{\alpha}^2 \sigma_1^2 + \sigma_2^2)} + \frac{16\bar{\alpha} \sigma_1^2 \sigma_2^2}{(\bar{\alpha}^2 \sigma_1^2 + \sigma_2^2)} \right\} \\
 (4.117) \qquad \qquad \qquad &= \frac{\bar{\alpha} \sigma_1^2}{\|\bar{c}\|_2^2}.
 \end{aligned}$$

The second order approximation of the estimator yields:

$$E[\hat{\alpha}] / \bar{\alpha} = 1 + \frac{\sigma_1^2}{\|\bar{c}\|_2^2},$$

which indicates a bias toward positive magnitude. Compared with the bias analysis for the conventional M-estimate, the bias of the ML estimate is independent of the data length n , which indicates that even though both estimators are asymptotically unbiased, they approach the asymptotic region with different rate (roughly 1 : n).

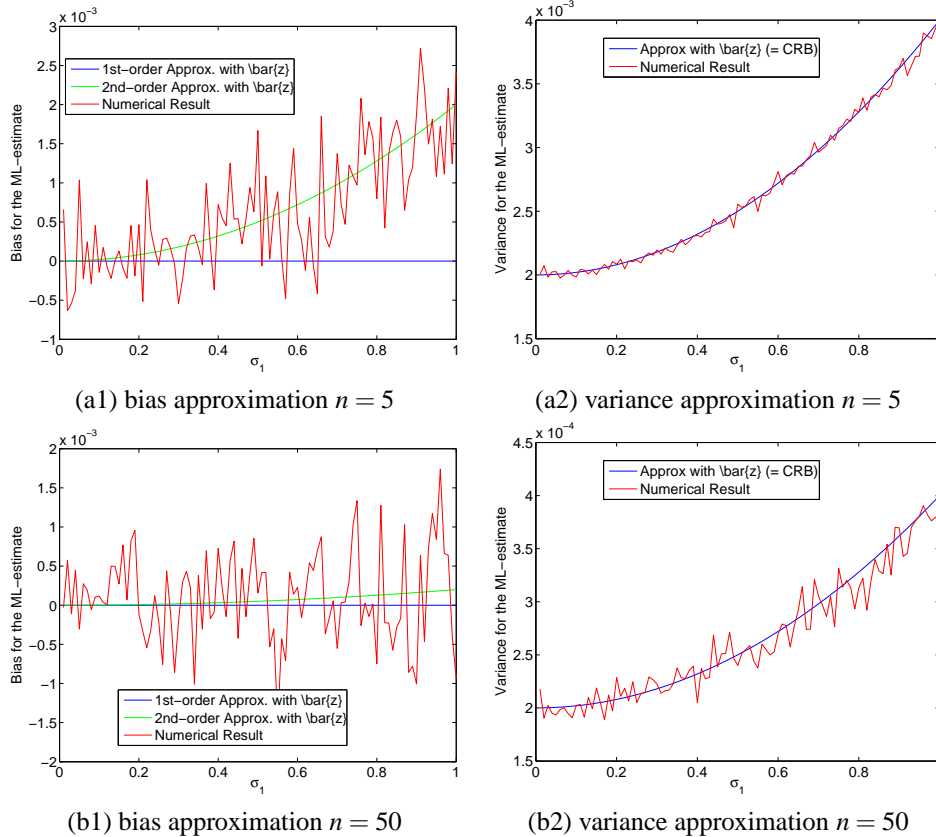


Figure 4.4: Bias and variance approximation of ML-estimate obtained from expansion about $(\bar{\alpha}, \bar{z})$.

In summary, we have tested with a simple example the estimated bias and variance of the conventionally used M-estimate and the ML-estimator from the statistical generative model. With the particular form of the example, the M-estimate can be obtained in closed form, and we have estimated bias and variance

from the explicit solution. To reflect the more general scenario, where such explicit solution is unavailable, we have used implicit function theorem and Taylor expansion to estimate bias and variance up to first and second order. Numerical results demonstrate reasonable agreement of the theoretically predicted values and empirical statistics. Qualitatively, all methods were able to capture the negative bias of the M-estimate, *i.e.*, the estimated parameter is a shrinkage relative to the true value. Furthermore, it could be shown that as an asymptotically unbiased estimator, the variance of the M-estimate is in fact upper bounded, and asymptotically approaches the Cramér-Rao Bound as the SNR increases. The ML-estimator according to the statistical model, being asymptotically UMVUE, has positive bias yet approaches unbiasedness faster (proportional to data length) than the M-estimate. The estimated covariance agrees with the CRB to second order. As a work in progress, this investigation is far from conclusive. More specifically, the ML-estimate demonstrates advantage in that it approaches the asymptotic unbiasedness with a faster rate; yet it has higher variance than the M-estimate in general. This leads to the familiar issue in estimator selection: the (pre-asymptote) bias and variance tradeoff needs to be studied carefully. Numerically, we observe that the M-estimate demonstrates variance close to the Cramér-Rao Bound, so it is possible that by including higher order expansion in estimating the variance, we could obtain an approximate rate at which the variance of the M-estimate approaches the Cramér-Rao Bound. Such information would allow us to reach either a consistent conclusion of the superiority between the M-estimate and the ML-estimate, or a partition of the parameter space so that each estimate would be the method of choice over certain regions.

CHAPTER V

Summary and Future Work

5.1 Summary

We have conducted research addressing two key aspects of image guided respiratory motion analysis: time series analysis to track semi-periodic signal structure from noisy observations and image registration to model motion between inhale-exhale image pairs. To track and predict the slowly varying mean position of a breathing signal, we have proposed a dynamic ellipse tracking method in an augmented state space. Formulated as a minimization problem in terms of algebraic distance, we provided a recursive algorithm for solving the static data case, utilizing stochastic approximation techniques. Assuming slow variations, we presented a natural extension of the recursion to an adaptive framework, to account for newly available samples. To accommodate noisy samples and missing observations, we modified the objective using robust fitting functions instead of the quadratic cost. Having shown that the feasible parameter region is the union of two convex sets and noting about the symmetric structure of the solution, we applied the projection gradient algorithm to solve the minimization problem. Analogous to the quadratic case, we took advantage of the recursive structure of the algorithm and extended it to incorporate adaptivity. To our knowledge, our method was the first to realize complete unsupervised tracking of respiratory motion in the presence of uncertainties in basic pattern, magnitude and phase. It has the potential to improve significantly the performance of both real-time adaptive treatment delivery and real-time gating systems.

For image registration, we focused on designing regularization to incorporate physical priors. In particular, we have proposed to use tissue-type rigidity regularization so that bone and soft tissue structures are regularized differently according to their own elasticity. To account for the commonly observed sliding effects along motion boundaries, we have first proposed a regularizer based on integrating some general functional of the Jacobian magnitude. We derived axiomatically the conditions on such functionals so that discontinuities are preserved. Then we further noticed the necessity to distinguish among different types of singularities, namely, folding and vacuum should be prevented yet shear should be preserved. With this in mind, we utilized the Helmholtz decomposition and regularized the divergence and curl component differently for the deformation field. The experimental results showed that the proposed decomposed regularization effectively combines the advantage of isotropic smoothing as in conventional Horn and Schunk, and discontinuity preserving regularizers such as total variation. Such efficient incorporation of prior knowledge shapes the registration process towards more physical solutions, which leads to better planning and treatment accordingly.

Furthermore, we have initiated a preliminary principled study on the fundamental performance limit of image registration problems. We proposed a statistical generative model to account for the noise effect in both the source and target images. The Cramér-Rao Bound for the corresponding maximum-likelihood estimator was computed. Meanwhile, we interpreted the conventional optimization based image registration results as an M-estimate. Using the implicit function theorem and Taylor expansion to estimate the local curvature of the objective function, we approximated its covariance accordingly. Noting that both the ML and M-estimates are asymptotically unbiased, we studied the pre-asymptotic performance by estimating the mean and covariance of each estimator with finite SNR. With a simple example, we have demonstrated that the bias of the proposed ML estimator decreases faster than the M-estimate as the SNR increases.

This result, unfortunately, is still insufficient to determine the relative superiority of the two estimators under consideration; because both ML and M- estimators are biased in the pre-asymptote region, and their variance is not lower-bounded by the Cramér-Rao Bound. Further investigation is necessary to study the deviation of the covariance from the Cramér-Rao Bound, which can be possibly conducted with higher-order Taylor expansion, similar to the bias analysis.

5.2 Future Work

- We have proposed a general framework for adaptive ellipse tracking. The adaptivity pace controls the balance between response efficiency and output smoothness, and should be determined properly. To this end, we have used a small segment of training data and retrospectively estimated the period with subspace projection method. After that, a static adaptivity parameter value (the window length for sliding adaptivity and the forgetting factor for exponential discounting) is used throughout the course. This is based on the assumption that frequency drifts are slow and that the robustness from ellipse fitting could tolerate the frequency variation. This presumption may be violated for long fractions, since the training segment becomes less correlated with the state as time progresses. An adaptive frequency drifting model is desirable to cope with such situation.
- In robust ellipse fitting, we need to determine the scale parameter for robust objective, *e.g.*, δ in the Huber function. Without assuming prior knowledge about the proportion of outliers relative to the normal samples or their distribution, we have used Otsu’s method to find a threshold value for the residual error and selected the scale parameter accordingly. The scale selection problem falls into the unsupervised classification category, where normal and abnormal samples are to be automatically distinguished. This is worth further investigation.
- We have proved asymptotic convergence with stochastic approximation techniques. Recognizing the similarity between the proposed iteration with the recursive least squares (RLS) algorithm, we believe it is feasible to estimate the error statistics in our model, analogous to the performance analysis work for linear filters.
- In the augmented state space, the distribution of the samples can be viewed as noisy observations of some latent random process. The distribution depends on the respiratory phase and other parameters. It is possible to consider robust statistical quantities such as rank order statistics to implicitly estimate the “center” of the observed cluster in the augmented state space. A potential advantage with such a statistical interpretation is that quantities such as confidence intervals and error distribution may be derived to facilitate the detection of changes in system dynamics.
- We have developed a tissue-type-dependent regularization method, which accounts for inhomogeneity of elasticity among different tissue types. Physically, anatomies not only exhibit inhomogeneous, but also anisotropic deformation properties, such as directional elongation of muscles. It would be desirable to properly incorporate such anisotropic physical prior as well. Furthermore, we have assumed access to an X-ray CT image, and obtained local tissue elastic property subsequently. Alternative methods to classify tissue types will be necessary for other image modalities.
- We have conducted preliminary discontinuity preserving registration for 2D images. We will further investigate the quantitative aspect of the problem and the 3D implementation. In particular, we will study possible solutions, such as introducing viscosity or surface tension regularization, to alleviate the rolling artifacts observed in 3D.
- We would like to further study unsupervised or semi-supervised schemes to choose the regularization parameters in penalized image registration problems.
- In this thesis, we have focused on the pair-wise image registration problems. When multiple frames are available over time, it is natural to extend the current work into a joint estimation setting where the temporal sequence of deformation fields is to be estimated. In this case, temporal correlation should be incorporated to encourage structured solution, such as smooth evolution. Moreover, for image

sequences obtained mainly under respiration-induced motion, this corresponds to an integration of our work in time series analysis that accounts for semi-periodicity and the regularized nonrigid image registration methodology.

- To analyze the fundamental performance limits of image registration, it is necessary to study the pre-asymptote variance for both the ML and M- estimates. The complexity of using high-order Taylor expansion and the limitations of small error analysis gives rise to the question as to whether there are more effective approximation tools for such tasks. This is a challenging topic, but one well worth pursuit.
- It is desirable to utilize the performance analysis of image registration problems to predict the statistical properties of the solutions for a given objective function. Knowledge about the fundamental limitations in image registration may help choose system parameters properly. For example, it is only necessary to obtain images with resolution corresponding to acceptable uncertainty in registration to avoid excessive imaging dose. The threshold for detecting abnormality should be set above the predicted local variance from the performance analysis with normal noise distribution.
- Given the theoretical development in this study, it is our sincere hope that practical benefit could be harvested. This will require thorough study of clinical implications, including effect on various dose metrics, and predictive outcome statistics such as tumor control probability (TCP) and normal tissue complication probability (NTCP).

APPENDIX

In speech processing, audio signal processing, and music synthesis, a *quasiharmonic* signal refers to a waveform that is virtually periodic microscopically, but not necessarily periodic macroscopically. In many biological phenomena, in contrast, there widely exist signals that are virtually periodic, yet demonstrate both microscopic and macroscopic variations. With a little abuse of notation, we use the term “semi-periodic” to describe such class of signals. A typical example of a “semi-periodic” signal is respiratory motion. Respiration is an involuntary action, the cycle of which is regulated through chemoreceptors by the level of CO_2 , O_2 , and PH in the arterial blood. Anatomically, the lungs reside in the thoracic cavity, encased by the liquid-filled intrapleural space. Inhalation requires active participation of respiration muscles, with the diaphragm being the most important. As the diaphragm contracts, and descends, it forces the abdomen inferiorly and anteriorly, increasing the superior-inferior (SI) dimension of the chest cavity. The intercostal muscles pull the ribs superiorly and anteriorly, increasing both the lateral (LR) and anterior-posterior (AP) diameters of the thorax. Exhalation is passive for quiet breathing. Due to the complex respiratory pressure-volume relationship of the lung and chest wall, deflating lung volume is larger than the inflating volume at the same transpulmonary pressure, and breath-in time is typically longer than breath-out time. This commonly observed phenomenon is called *hysteresis*.

With the advent in target conformal radiotherapy, such as Intensity Modulated Radiotherapy (IMRT), it is important to monitor tumor movement with high precision in real-time. Internal tumor position can be extracted from images obtained from electronic portal imaging detector (EPID) or orthogonally placed x-ray flat panels, as in Integrated Radiotherapy Imaging System (IRIS) [54] or CyberKnife Robotic Radiosurgery System (AccuracyTM Inc, Sunnyvale, CA), with or without implanted markers around the tumor region. Moreover, it is desirable to minimize diagnostic imaging dose for safety concerns, and external surrogates such as thermistors, thermocouples, strain gauges, pneumotachographs [60], and infrared skin markers are utilized to infer internal tumor position. In either cases, it is important to characterize the complexity of internal motion, track both instantaneous and long term variation, and predict future tumor position to account for system latency. When external surrogate is applied, it is critical to accurately infer internal tumor position from external observations. Chapter A proposes a scalar complexity index to characterize the irregularity level of a breathing trajectory [101]. Chapter B describes a nonparametric prediction approach based on local regression. Chapter C addresses the hysteresis issue in external-internal inference via state augmentation.

APPENDIX A

A Breathing Pattern Irregularity Index with Projection-based Method

1

Characterization of organ motion is important in radiation therapy, including dose planning and treatment delivery [12, 18, 36, 37, 53]. Tumor motion, especially in lung/liver regions, is highly correlated with breathing patterns. Therefore, an index that characterizes breathing regularity can facilitate treatment planning for tumors in those regions, particularly for individualized treatment planning.

Periodicity has been a major assumption in breathing trajectory analysis, as good reproducibility indicates the potential for a simple structured treatment plan tailored towards the fundamental breathing pattern. Harmonic analysis has been employed widely to characterize respiratory patterns [43, 87, 94]. Peaks of the Fourier spectrum are often used to determine the dominating periodic behavior of the temporal trajectory. Such approaches lack a “goodness” measure, *i.e.*, it is not clear how a periodic signal having the dominant frequency differs from the true trajectory. Consequently, no fundamental periodic pattern is available to judge the soundness of such a result.

We propose a rigorous general framework for periodicity analysis based on subspace projections. For each period within a physiologically reasonable range, a measured breathing signal is projected onto the subspace of all signals having that period to obtain the “best fit” periodic signal in the Least Squared Error (LSE) sense. Residual errors for each such period are then compared to yield the overall best periodic approximation. The estimated trajectory obtained by this “projection” method is therefore the closest periodic signal with respect to observed data. We derived the method in continuous signal space to account for the sampling effect explicitly. We also allow temporal samples to be non-uniformly spaced to offer more freedom for the data acquisition procedure.

A.1 An Irregularity Index based on Projection Distance

Given a set of discrete samples of a breathing trajectory, we want to find the periodic signal that best matches the observation data. This is equivalent to reconstructing a periodic signal of unknown period from its noisy discrete samples. For this problem to be feasible, we assume that there is some maximal frequency component in the signal. This assumption is physiologically reasonable. We thus focus on the subspace of band-limited periodic signals. We formulate the problem in a multilayer optimization setup where we search over all possible periods for the “best-fit” signal. For each period within a reasonable range, the observed breathing trajectory is projected onto the subspace of all band-limited signals having that period to obtain the closest matching periodic function. Projections from each such subspace are then compared to yield the overall best periodic approximation. This method accounts for the discrete temporal sampling explicitly, and allows for the possibility of nonuniform sampling.

We model the observation data y_i as a temporal trajectory sampled at $\{t_i\}_{i=1}^N$ with additive noise:

$$(A.1) \quad y_i = f(t_i) + n_i, \quad i = 1, 2, \dots, N,$$

where f is the unknown ground-truth continuous periodic function whose spectrum has finite support between $[-\gamma, \gamma]$ and n_i denotes the additive noise.

¹This chapter is based on materials from [101]

If $f(t)$ is a band-limited function with period T , then we follow [29] to rewrite it as linear combination of Fourier harmonics:

$$(A.2) \quad f(t) = \sum_{k=-K}^K c_k e^{j\frac{2\pi}{T}kt}, \quad K = \lfloor \frac{T}{2} \rfloor,$$

where c_k 's are the coefficients for Fourier harmonics, and $\lfloor \cdot \rfloor$ denotes the floor function.

Evaluation of the above representation at $\{t_i\}_{i=1}^N$ can be compactly rewritten in vector form as:

$$(A.3) \quad \mathbf{f} = \mathbf{G}_T \mathbf{c},$$

where $\mathbf{f} = [f(t_1), f(t_2), \dots, f(t_N)]$ denotes the discrete samples of the underlying function f ; $\mathbf{c} = [c_{-K}, c_{-K+1}, \dots, c_K]$ is the concatenation of Fourier coefficients; and the matrix \mathbf{G} is defined as

$$(A.4) \quad \mathbf{G}_T(i, k) = e^{j\frac{2\pi}{T}kt_i}.$$

Therefore, given the observed sample trajectory $\mathbf{y} = [y_1, y_2, \dots, y_N]^T$, the optimal period T^* is the solution to the following optimization problem:

$$(A.5) \quad T^* = \arg \min_T \min_{\mathbf{c} \in \mathbb{C}^{2K+1}} \|\mathbf{y} - \mathbf{G}_T \mathbf{c}\|^2,$$

where \mathbb{C}^{2k+1} is the set of vectors of length $(2k+1)$, and $\|\mathbf{y}\|^2 = \sum_{i=1}^N |y_i|^2$. The closest periodic signal to the sampled trajectory in LSE sense is then given by:

$$(A.6) \quad f^*(t) = \sum_{k=-K}^K \hat{c}_k e^{j\frac{2\pi}{T^*}kt}.$$

where $K = \lfloor \frac{T^*}{2} \rfloor$ and \hat{c}_k are obtained as the components of solution to (A.7) below when $T = T^*$.

For a given candidate period T , the bandwidth parameter $K = \lfloor \frac{T}{2} \rfloor$ is a constant, and the inner optimization problem becomes an ordinary least-squares minimization:

$$(A.7) \quad \mathbf{c}_T^* = \arg \min_{\mathbf{c} \in \mathbb{C}^{2K+1}} \|\mathbf{y} - \mathbf{G}_T \mathbf{c}\|^2.$$

From classical optimization theory [68], the optimal \mathbf{c}_T^* of (A.7) satisfies the normal equation:

$$(A.8) \quad (\mathbf{G}_T^* \mathbf{G}_T) \mathbf{y} = \mathbf{G}_T^* \mathbf{c}_T^*,$$

where \mathbf{G}_T^* is the conjugate transpose of \mathbf{G}_T and $\mathbf{G}_T^* \mathbf{G}_T$ is known as the Gram matrix.

Moreover, when the sample size is large enough, specifically $N \geq 2K+1$, which we assume hereafter, \mathbf{G}_T has full column rank, and the $(2k+1) \times (2k+1)$ Gram matrix $\mathbf{G}_T^* \mathbf{G}_T$ is invertible [41]. The optimal solution for equation (A.8) can be written explicitly as:

$$(A.9) \quad \mathbf{c}_T^* = (\mathbf{G}_T^* \mathbf{G}_T)^{-1} \mathbf{G}_T^* \mathbf{y}.$$

At this point, we have solved the inner optimization problem in (A.5) in closed form. The feasible range of periods T in the outer minimization can be designed by incorporating physical knowledge. For instance, normal breathing is expected to have a period between 1 to 10 seconds. Moreover, even though the peak of the Fourier spectrum is not informative enough by itself, it turns out to be a reasonably good initialization for our method. Notice that if exhaustive search over T is to be applied in A.5, we need to evaluate (A.9) and (A.6) for each T of interest. Thus the computation cost depends both on how finely we sample the period parameter T and the range of search. Using a good initial guess for T^* can reduce the search range and thus reduce computation substantially. Also, reasonable initialization helps to prevent the algorithm from falling into nonphysical local minima. Since it is now a simple 1-dimensional optimization problem to find T^* , we use an exhaustive line search over a relatively small interval thanks to a good Fourier-based initialization.

Alternative optimization approaches like multi-resolution or incremental refinement could be used to speed up the process. Due to the use of superposition of harmonics to describe periodic functions, projection to the subspace corresponding to periodic functions with period $2T$ would naturally yield a better data fit than the projection onto the subspace for period T . In other words, a function of period T is certainly a function of period $2T$, but not vice versa. However, the additional descriptive power may not always be desirable, since this could cause over-fitting introduced by noise. Initialization by detecting the peak of the Fourier Spectrum picks out the dominant harmonic component and the algorithm only needs to search over a relatively small neighborhood around that initialization point, with the confidence that the local minimal obtained would be physiologically optimal.

Finally, our proposed irregularity index is the Root Mean Squared Error (RMSE) between the overall optimal periodic signal and the measured trajectory:

$$(A.10) \quad \text{RMSE} = \sqrt{\frac{1}{N} \sum_{i=1}^N \|f^*(t_i) - y_i\|^2}.$$

A.2 Material and Verification Design

We used the Real-Time Position Management (RPM, Varian Medical Systems, Palo Alto, CA) system to obtain the trajectory of an external fiducial placed on each patient’s chest wall. This fiducial tracking system records data in time-displacement pairs that are generally assumed to be highly correlated with superior-inferior diaphragm motion [125]. This system is most useful for treating patients with tumors in the chest or lung area without compromising their breathing.

Twelve such clinical breathing signals were used in this study. The characteristic parameters of this population of data are listed in Table A.1 .

ID V.S. Parameter	1	2	3	4	5	6	7	8	9	10	11	12
Data Characterization ²												
STD (cm)	0.158	0.210	0.266	0.242	0.206	0.259	0.242	0.267	0.283	0.313	0.335	0.202
Breathing Trajectory Fitting with Modified Cosine Model												
period (sec)	4.7	4.6	4.9	5.3	5.3	4.3	4.9	6.4	9.5	5.6	3.0	5.3
RMSE (cm)	0.138	0.171	0.216	0.139	0.193	0.224	0.145	0.208	0.153	0.096	0.337	0.169
dose error (%)	1.667	2.793	3.527	2.092	3.217	3.580	2.402	3.293	2.496	1.454	6.144	2.161
PTV margin (cm)	5.940	5.900	5.523	5.723	5.727	5.859	5.646	5.338	5.724	5.522	5.951	5.835
95% dose coverage	0.909	0.887	0.850	0.904	0.878	0.851	0.906	0.858	0.890	0.938	0.811	0.888
Breathing Trajectory Fitting with Projection Method												
period (sec)	4.7	4.4	4.5	5.4	4.1	4.6	4.7	7.2	9.7	5.6	3.1	5.2
RMSE (cm)	0.135	0.155	0.102	0.132	0.162	0.127	0.115	0.075	0.148	0.090	0.328	0.166
dose error (%)	1.595	2.440	1.638	1.983	2.352	1.721	1.832	1.210	2.471	1.431	6.137	2.066
95% dose coverage	0.915	0.903	0.934	0.903	0.876	0.910	0.924	0.949	0.905	0.942	0.836	0.895
Result for 20sec Training, 10sec Testing												
period (sec) ³	4.2	4.2	4.5	5.2	4.3	4.8	4.8	7.3	9.0	5.7	3.0	5.0
RMSE _{train} (cm) ⁴	0.153	0.151	0.089	0.126	0.082	0.075	0.121	0.042	0.116	0.078	0.228	0.049
RMSE _{rest} (cm) ⁵	0.177	0.256	0.150	0.231	0.318	0.283	0.141	0.147	0.290	0.150	0.580	0.3062

Table A.1: Dataset information and Experiment Results

Under Institutional Review Board (IRB) approved protocol, we have used the RPM system (Varian, Palo Alto, CA) to obtain breathing trace data recorded at 10Hz with duration 30sec from 12 different patients . The recorded RPM data have relative units. To better illustrate the major idea in this paper, we normalize the all the breathing trace data to have uniform zero mean and 1cm peak-to-peak variation. Shifting the mean does not introduce any bias into any treatment simulation since it is a global quantity; while normalizing the amplitudes makes the data more representative of typical tumor motion induced by respiratory motion. The standard deviations of these normalized data are listed in Table A.1. To justify the soundness of the proposed irregularity index, we have virtually simulated a 1-dimensional phantom object of size 5cm that

move according to the observed trajectories, to mimic the behavior of a 5cm size tumor with peak-to-peak motion about 1cm, which is realistic in clinical situations. A single ideal 1-dimensional treatment beam, or in fact, delivery pattern of the same size (5cm) is designed for dose delivery simulations. It has no penumbra, and completely covers the simulated target with uniform radiation intensity. This idealized energy deposition model will be used hereafter to illustrate the potential impact of motion patterns and how they influence energy deposition.

To verify that the proposed “irregularity index” and the fundamental pattern obtained from the the projection model are clinically significant, we have designed three sets of experiments.

First, we show that the Root Mean Squared Error (RMSE), which is a mathematical criterion, is well correlated with clinically critical metrics. In this paper, we use dose error, Planning Target Volume (PTV) margin and 95% dose coverage to characterize performances. In particular, dose error is computed in per cent as the normalized difference between received dose and the ideal dose that corresponds to a perfect overlap between the target and treatment beam throughout the whole treatment procedure; PTV margin is the expansion needed to ensure that the entire clinical target volume receives the prescribed dose; and 95% dose coverage is computed as the portion of the target that receives no less than 95% of the designated dose with no margin. To account for the interplay between target motion and treatment beam adjustment, the phantom object is moved conforming to the observed breathing trace and the treatment beam is scheduled according to a designated pattern. We evaluate both the periodic fundamental pattern extracted with the projection model in A.6 and the one obtained with the optimal commonly used modified cosine model [70, 71] to control the movement of the treatment plan. The modified cosine model assumes that the breathing trajectory conforms to the following formula:

$$(A.11) \quad z(t) = z_0 - a \cos^{2n}(\pi t / \tau - \phi),$$

where z_0, a, n, τ, ϕ are assumed to correspond to exhalation position, motion amplitude, asymmetry degree, period and phase offset respectively, and are parameters to be optimized; $z(t)$ represents the breathing trace index by time.

Dose error, PTV margin and 95% dose coverage are compared against RMSE in both setup to demonstrate the correlation.

Second, we compare the projection-based model with the modified cosine model (A.11) to test the feasibility of the obtained fundamental pattern. RMSE as well as dose error, PTV margin and 95% dose coverage are used for this comparison.

Third, we illustrate the potential clinical use of the proposed method to predict motion induced by respiration. We partitioned the breathing trace into two parts: a training part of duration 20sec and a testing part of duration 10sec. For each breathing trace, the projection model is learned with the training trajectory only, and it is used to “predict” the breathing behavior for the testing portion. This is essentially a test of temporal variance.

A.3 Results and Discussions

Fig. A.1 shows one patient dataset to illustrate the role Fourier-initialization plays in avoiding suboptimal local minima. An exhaustive evaluation for RMSE was carried out over a large range of candidate periods in Fig. A.1 (a). Fig. A.1(b) illustrate the non-physiological optimal obtained without proper prior information, due to reason we discussed previously: harmonic analysis has an inherent bias toward large period. Fig. A.1(c) shows that initializing with peak location of Fourier Spectral (in this example corresponding to $T = 4.3s$) helps to correctly capture the physiologically sound optimal period and enables us to restrict the period search to an even smaller candidate set for further computation efficiency.

To validate the correlation between the root mean squared error (RMSE) and the clinically critical metrics, we plot the performance characteristic parameters (dose error, PTV margin, 95% dose coverage) vs. RMSE in Fig. A.2 for both projection model based motion compensated treatment and modified cosine model based motion compensated treatment. Quantitative results are listed in Table A.1. In both treatment plan simulations, dose error and PTV margin demonstrate an increasing trend as RMSE becomes larger while the 95% dose coverage decreases. This validates the soundness of using RMSE as the index for “performance indicator”.




Figure A.1: Proper initialization helps to avoid suboptimal (nonphysical) local minimum: (a) Exhaustive evaluation of RMSE for difference candidate periods; (b) estimated pattern at $T = 8.2s$, this is nonphysical even though it corresponds to slightly better fitting in RMSE sense; (c) estimated pattern at $T = 4.1s$, the physiologically sound optimal period.

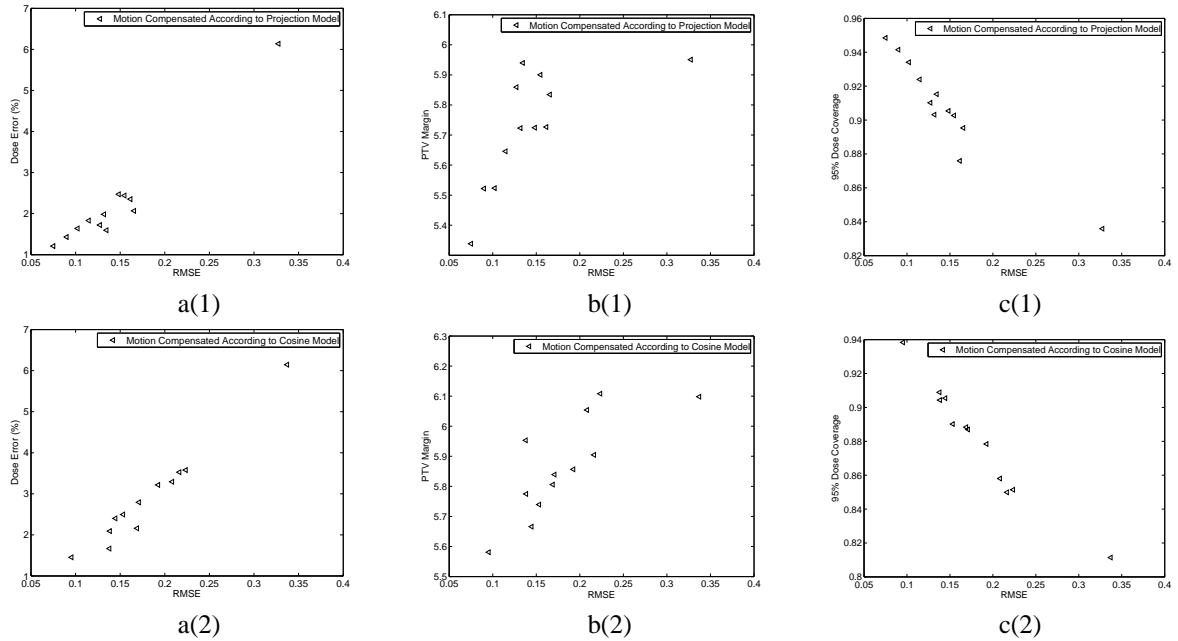


Figure A.2: Clinical significant performance metrics v.s. Root Mean Squared Error (RMSE). Different metrics are indicated with letters [(a#)] dose error (%); [(b#)] PTV margin (cm); [(c#)] 95% dose coverage. Different motion models for commoving the treatment beam are indicated with numbers: [(X1)] projection based model (treatment beam trajectory described as linear combinations of harmonics); [(X2)] modified cosine model.

Moreover, we carry out a comparison between the projection-based model with the commonly used modified cosine model described in Equation (A.11). Fig. A.3 shows the RMSE of the best fit modified cosine model versus the proposed index (RMSE derived from projection model), and it demonstrates that not only does our index capture how well the signal can be approximated by a well-recognized physical model, but the fundamental pattern obtained via the projection model uniformly outperforms the modified cosine model in the LSE sense. For further clinically meaningful justification, we calculate the performance characteristic parameters corresponding to a modified cosine model in Table A.1, and we can observe that our projection model yields lower RMSE, dose error, PTV margin and higher dose coverage than the modified cosine model overall (Fig. A.3). Furthermore, the problem of fitting the data to the model described by Equation (A.11) is highly non-convex with respect to its parameters which incurs two issues: it is extremely sensitive

to initialization due to the numerous local minima; and it is computationally expensive as a nontrivial high dimensional search problem. In contrast, the proposed projection approach offers a closed form solution for the inner optimization problem in (A.5) and is thus simplified to a 1-D line search, it has an obvious advantage in computation efficiency over the modified cosine model.

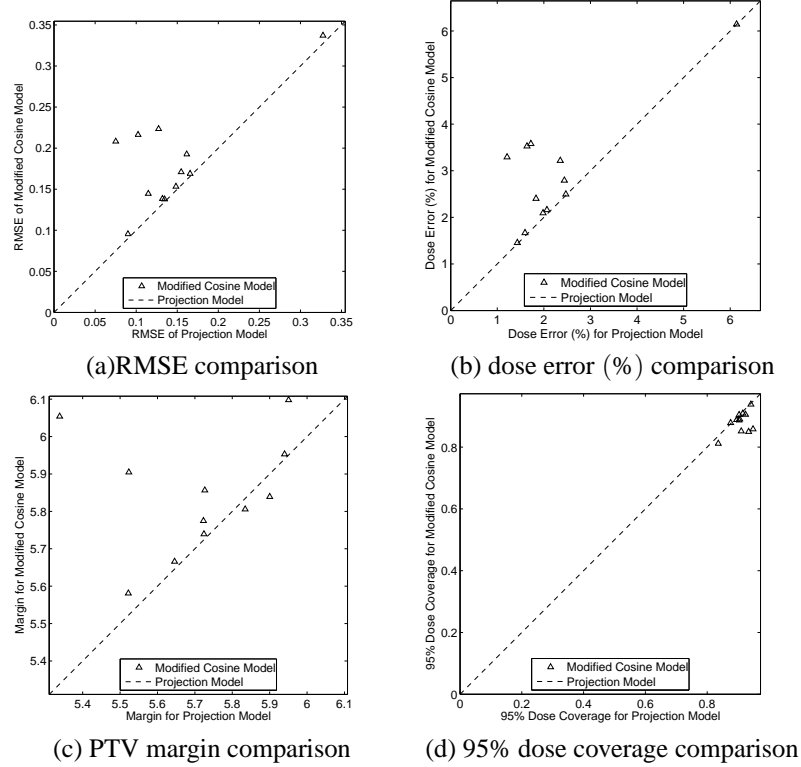


Figure A.3: RMSE, Dose error (%), PTV margin (cm), 95% dose coverage of modified cosine model v.s. projection model.

To further justify the above claims, Fig. A.4 shows some of the fitted trajectories with “optimal” cosine model parameters with their counterparts from the projection-based approach. The fundamental patterns obtained by the projection method do indeed offer a better match than the cosine model. This is a result of the intrinsic “nonparametric” nature of the projection based approach. Described as a linear combination of harmonics, the fundamental pattern has essentially $(2K + 1)$ degrees of freedom where K is determined by the imposed band limit of the physical signal. The modified cosine model, on the other hand, has explicitly assumed no more than 5 degrees of freedom, which has restricted its descriptiveness. For the same reason, our method imposes no symmetry on the fundamental pattern; in particular, the trajectory of inhalation does not have to be the inverse of exhalation, unlike the modified cosine model.

A “good” fit of the breathing trace with a periodic pattern is obtained (low RMSE by the proposed irregularity index) indicates that the breathing trace under examination is highly regular, and vice versa. Similar argument holds for the relationship between “bad” fit (high RMSE) and high irregularity. Instead of examining the combination of a whole bunch of quantities, such as standard deviation of amplitude, mean positions, periods of breathing cycles, etc, this single number (the RMSE) serves as the irregularity index, since it is designed specifically for this purpose. Therefore, observing a low RMSE increases the confidence and feasibility for potential dynamic treatment for the mobile target. In particular, Synchronized Moving Aperture Radiation Therapy (SMART) [84] and similar motion compensation based treatment schemes are potentially applicable. Moreover, the fundamental pattern, which is obtained as a free side-product during the process of estimating period and computing the irregularity index, is a good indicator of what the radiation beam pattern should be, serving the same purpose as average tumor trajectory (ATT) introduced in [84]. In other words, it can be regarded as an alternative derivation of ATT without having to examine individual cycles

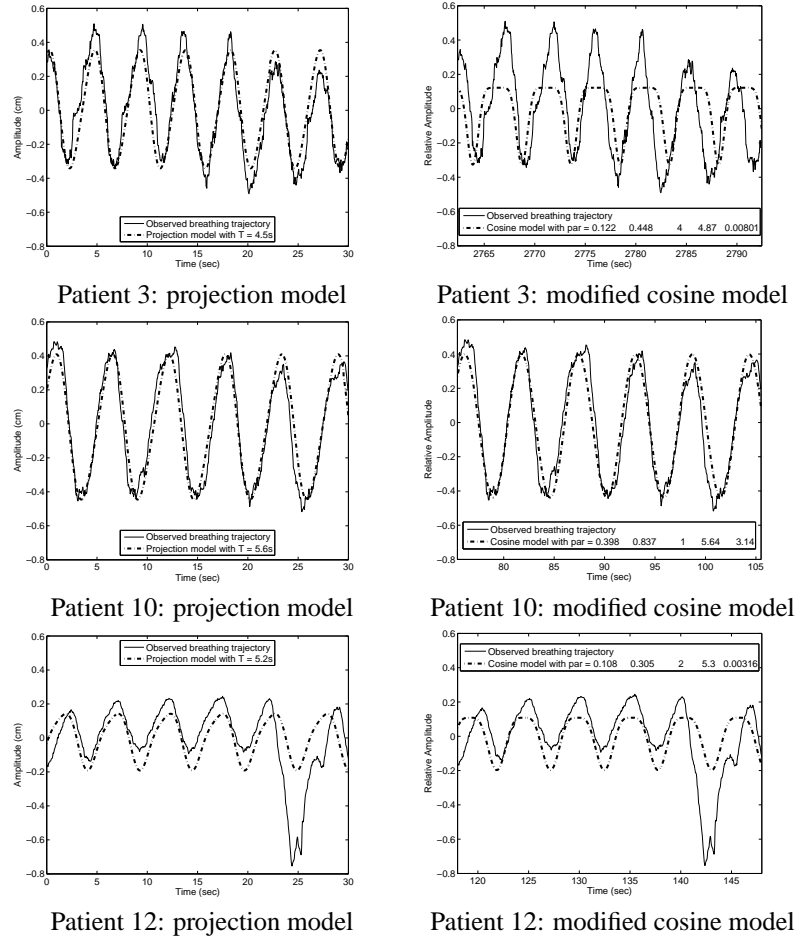


Figure A.4: Trajectory fitting with projection model and modified cosine model. Left column: projection model vs. true trajectory; right column: modified cosine model vs. true trajectory.

too closely. A potential merit of the proposed method for extracting ATT is that it is much less sensitive to additive noise due to its global nature - every sample on the observed breathing trace contributes to the estimation of the fundamental pattern.

To show the potential application of the proposed projection based scheme to predict target motion, we derive the fundamental pattern with the first 20sec of breathing trace (the training portion) and apply it to the remainder of the data - the next 10sec of breathing trajectory is called “testing portion” since it is not seen by the projection model. We illustrate some examples in Fig. A.5. The irregularity indexes derived from the learning portion, the corresponding optimal period and the evaluation of its fit to the the ground-truth trajectory for the testing portion using RMSE are provided in Table A.1.

The quality of the prediction depends on how repetitive the true breathing trajectory is, which again can be measured by the proposed regularity index. When we examine closely the RMSE computed from training portion and test portion, we will see that the latter is uniformly larger, which is expected (since optimization is applied only to training data). Moreover, when we examine across cases, there is a positive correlation between RMSE computed during training and RMSE computed from testing. This indicates RMSE during recent historical trajectory is a good predictor for RMSE, and thus irregularity level for near future. Generally, being a global regularity measure, the proposed index may not capture time varying properties of the breathing signal. This limitation can be overcome by applying the proposed method to smaller sliding time intervals instead of the whole trace. Despite this limitation, the projection model based prediction appears to provide reasonable predictions within approximately a 2 second response window given a sufficiently regular breathing trace. Even though this number is significantly larger than the 0.4




Figure A.5: Prediction of breathing trajectory with projection model.

second discussed in [125], we are *not* claiming that the proposed algorithm is preferable to adaptive filtering, since regularity in breathing trace is a pretty stringent assumption. Modeling of free form breathing is a hard and unsolved problem in general. It is often desirable to have a simple and descriptive model even if some conditions need to be checked in the first place. Moreover, the proposed irregularity index is a convenient tool for such a sanity check. By examining this single index, we can determine whether the breathing trace is regular enough for the periodicity assumption to hold, hence the corresponding prediction or synchronized motion compensation with ATT may be applied.

A.4 Summary

We have derived a general framework to find the closest periodic signal that best matches the temporally sampled observation of breathing trajectory. Experimental results have shown good consistency with physical knowledge and clinically critical parameters as dose percentage error, PTV margin and 95% dose volume. Comparison between the popular modified cosine breathing model and the projection-based approach shows that being consistent with the residual error from fitting the modified cosine model, our approach offers additional computation efficiency and robustness in the optimization process. Furthermore, we get the fundamental breathing pattern which helps to justify the soundness of the results and can serve as a valuable reference in further treatment planning. Potential applications of the fundamental pattern to dynamic motion compensation and prediction are illustrated with preliminary experiments. It is also likely that knowledge of the periodic signal can aid in reconstruction of 4-dimensional computed tomographic models.

In this study, we have focused on finding the optimal periodic signal in the LSE sense. As future work, we would like to investigate alternative metrics that are potentially more tolerant to transient pathological breathing patterns. Also, for a particular treatment planning scheme, some choice of matching metrics could be more suitable than others, and the design of plan-dependent irregularity indexes would be interesting. Finally, we have used in this study the RMSE resulting from the projection method as an irregularity index. Potential variants, for instance, a normalized version, may be more desirable in some applications.

APPENDIX B

Real-time Prediction of Respiratory Motion based on Nonparametric Local Regression Methods ¹

Current developments in radiotherapy such as Cyberknife and Intensity Modulated Radiotherapy (IMRT) offer the potential of precise radiation dose delivery for moving objects. Accurate target volume tracking is necessary for conformal treatment plans to fully utilize their capacity. Image-guided radiotherapy needs to consider system latencies resulting from image acquisition, communication delay, data processing, and mechanical processing. For treatment over multiple fractions, or long procedures, the diagnostic radiation dose can be significant, so it is desirable to reduce the image acquisition rate. To address this issue, hybrid tumour tracking approaches that combine episodic radiographic imaging and continuous monitoring of external surrogates have been investigated [80, 81, 86, 104, 105]. There are two active areas of research related to hybrid tracking: (1) study of feasibility and effective use of external surrogates (including the placement mechanism) such as thermistors, thermocouples, strain gauges, pneumotachographs and infrared skin markers [1, 47, 57, 59, 73, 117, 124]; (2) prediction algorithms [109, 125, 133]. In particular, even if perfect information about the current state is assumed, the lag between observing tumour location and treatment delivery still necessitates having predictors that can “look ahead” enough, yet behave reasonably well even for relatively low input sampling frequencies.

This study belongs to the second category where we are interested in predicting target motion located in the lung area or its vicinity. Such motion is mainly caused by breathing, and exhibits semi-periodicity as observed in normal breathing signals. This is a very active research area [37, 53, 109, 125]. The semi-periodic structure of the breathing signals make explicit modeling challenging, since parametric models often fail to capture local variations. On the other hand, overly flexible models that depend only on temporally local information fail to use correlated historical information. Among the most investigated methods are linear predictors with various covariate lengths, neural networks, and Kalman filters.

We propose a prediction method based on local weighted regression. Adopting a classic approach in modeling dynamical systems, we first generate an augmented state with the most current observation and one or more preceding samples. This augmented state is designed to capture the local dynamics about the time point of interest, and it is used as the covariate for the predictor system. For a pre-specified “look-ahead” length, the target response pattern of the predictor is obtained from the training data. Those state-response pairs form a scatter-plot in a high-dimensional space where we apply locally weighted regression. Intuitively, the predictor infers its response map from the behavior of its neighbors in this state space, since it is probable that they are. The regression weights are designed to reflect the “distance” between the state of interest and the training samples.

For the purpose of real-time tracking and prediction, we adaptively adjust the inference weights to incorporate the decaying temporal correlation among response patterns with longer time lags².

We discuss the proposed methods in detail in Section B.1. The method is applied to clinical RPM data (RPM Varian Medical System, Palo Alto, CA) that is described in Section B.2. We report the test results and the comparison to alternative methods in Section B.3. Finally, we discuss future directions in Section B.4.

¹This chapter is based on materials from [98]

²In fact, this corresponds to augmenting the state with the time index as an extra dimension.

B.1 Methods

In this section, we propose a prediction method based on locally weighted regression. For simplicity, we describe the model in terms of scalar locations, *i.e.*, 1-D observation. The extension to vector observation is straightforward. Section B.1.2 first introduces a primitive version that ignores the change of temporal correlation with time lag, and Section B.1.3 and Section B.1.3 extend it to include time indexing.

B.1.1 Model Setup and General Notations

Let the continuous scalar function $f : \mathfrak{R} \rightarrow \mathfrak{R}$ denote a motion index signal. At time instant τ , we are given a set of discrete samples $\{s_i = f(t(i)), i = 1, 2, \dots, k\}$ of the breathing trajectory prior to τ , with $t(k) < \tau$. For simplicity, we assume that the observed signal is sampled uniformly with frequency ψ Hz, *i.e.*, $t(i+1) - t(i) = 1/\psi$. We assume that the look ahead length is an integer multiple of the sampling interval $1/\psi$ seconds, and for later convenience, we represent it in the discrete unit, *i.e.*, a look-ahead length L indicates a L/ψ seconds prediction. We use p to denote the state dimension used to capture system dynamics. To draw an analogy to ordinary differential equation (ODE) based system, $p = 2$ corresponds to first order difference system with location and approximate velocity; and $p = 3$ corresponds to a second order difference system with the addition of acceleration.

B.1.2 Basic Local Weighted Regression

At current time instant τ , the available observations are s_1, \dots, s_K , where $K \leq \tau \times \psi$. Then for any $i \leq K - L$ we construct length p state vector $\mathbf{x}_i = [s_{i-(p-1)\Delta}, \dots, s_i]$, and response variable $\mathbf{y}_i = s_{i+L}$. The parameter Δ is an integer that indicates the “lag length” used to generate the augmented state. It should be chosen to properly reflect system dynamics: small lags are more sensitive to dynamical change as well as noise; big lags are more robust to the presence of noise yet average out the system dynamics at the same scale. The set of hyper-pairs $(\mathbf{x}_i, \mathbf{y}_i)$ form a scatter plot in the $p + 1$ dimensional space. Assuming that the state thus constructed conveys all the information about system dynamics, then the scatter-plot summarizes the noisy realizations of the prediction map: $g : \mathfrak{R}^p \rightarrow \mathfrak{R}$:

$$(B.1) \quad \hat{\mathbf{y}}_i = g(\mathbf{x}_i),$$

where the predictor g is a smooth function. This is a reasonable model as we do expect the prediction to vary smoothly with the historical trajectory. Our goal is to predict the target location at time $(\tau + L/\psi)$ seconds, which is equivalent to estimating $\hat{\mathbf{y}}_K$.

Respiratory motion is not stationary, in fact, both the system dynamics and its local statistics vary in a semi-periodic fashion. Unfortunately, most existing methods in estimating the prediction map g fail to take this “phase-dependent” phenomena into account, with the exception of [103] and [133] where a discrete set of stage-wise models are constructed and updated adaptively. The idea is to train (or infer) a predictor at a given state with (only) those historical data samples that behave similarly, or vaguely speaking, belong to a similar respiratory stage. Yet the existing stage-wise models require predetermining the number of discrete stages and often involve segmentation-based training. To circumvent these difficulties, we hereafter provide a means to locally estimate g in the state-space neighborhood of \mathbf{x}_K , based on local regression (LOESS) from nonparametric methods in the statistical literatures [19].

Let r be a pre-determined constant³ that specifies the size of the neighborhood whose members affect the estimate in the scattered $(p + 1)$ -dimensional space. Let h_K be the distance from \mathbf{x}_K to the r th nearest neighbor in terms of Euclidean distance in the p -dimensional subspace, *i.e.*, h_K is the r th smallest number among the distance between \mathbf{x}_i , $i = 1, 2, \dots, K - L$ to \mathbf{x}_K . Let $\kappa(\cdot)$ be a symmetric kernel function that satisfies:

1. $\kappa(x) > 0$ for $|x| < 1$ and $\kappa(x) = 0$ for $|x| \geq 1$;
2. $\kappa(-x) = \kappa(x)$;
3. $\kappa(x)$ is a non-increasing function for $x \geq 0$.

³Equivalently, it could also be specified as a ratio with respect to the total number of data points.

We select local inference weight according to:

$$(B.2) \quad w_i = \kappa(h_K^{-1} \|\mathbf{x}_i - \mathbf{x}_K\|).$$

Figure B.1 illustrates the idea of weighting based on distance in state space. For simplicity, one delay tap is used ($p = 2$), so the state $\mathbf{x}_i = [s_{i-\Delta}, s_i]$. The goal is to estimate the response y_K for current state vector \mathbf{x}_K , from available covariate-response pairs $(\mathbf{x}_i, \mathbf{y}_i)$ for $i \leq K - L$. Notice that $K - L < i \leq K$ are not used in the regression, since their response values \mathbf{y}_i are not yet available at time instant K . Distances between current state \mathbf{x}_K and \mathbf{x}_i are computed and the kernel function is used to determine the regression weights w_i as in (B.2). The assigned weights emphasize those training samples that share similar dynamics as the current state of interest as shown in the lower part of Figure B.1.




Figure B.1: State-space distance and local regression weight assignment. Upper subfigure: Illustration of prediction quantities with 1st order dynamic $\mathbf{x}_i = [s_{i-\Delta}, s_i]$, current data point s_K , prediction location $\mathbf{y}_K = s_{K+L}$, available covariant-response pair $(\mathbf{x}_i, \mathbf{y}_i)_{i \leq K-L}$; lower subfigure: distance map (blue dash-dot line) in the state space $d_i = \|\mathbf{x}_i - \mathbf{x}_K\|$ and inferred regression weights w_i (green solid line).

We subsequently estimate g locally using a polynomial of degree d , *i.e.*, we use a predictor of the form $g(\mathbf{x}) = \sum_{q=1}^Q \beta_q z_q(\mathbf{x})$, where $Q = (d+1)^p$ and $z_q(\mathbf{x}) = \prod_{j=1}^p x_j^{d_j}$ and $(d_1, \dots, d_p) \in \{0, 1, \dots, d\}^p$ that corresponds to the base- d representation of q .

We estimate the coefficients of the local polynomial by minimizing the weighted local squared error:

$$(B.3) \quad \begin{aligned} \hat{\beta} &= \arg \min_{\beta} \sum_{i=1}^{K-L} w_i (\mathbf{y}_i - \sum_{q=1}^Q \beta_q z_q(\mathbf{x}_i))^2 \\ &= (Z^T W Z)^{-1} Z^T W Y, \end{aligned}$$

where Z is the $(K-L)$ by Q matrix with elements $z_{iq} = z_q(\mathbf{x}_i)$. The weighting matrix W is a diagonal matrix with $W(i, i) = w_i$. Since the local weight w has a limited bandwidth h_K as designed in (B.2), there

are only r nonzero diagonal elements in the weighting matrix W . Correspondingly, the outer summation $\sum_{i=1}^{K-L}$ can be equivalently written as supported only on a local neighborhood of radius h , *i.e.*, $\sum_{i: \|\mathbf{x}_i - \mathbf{x}_K\| < h}$. Therefore, the data vectors involved have length $r \ll K - L$ rather than $K - L$. It is desirable to choose a small neighborhood size r to decrease computation cost, yet not overly small to sacrifice the regularity of (B.3), *i.e.*, the invertibility of $Z^T W Z$.

For subsequent prediction from a given observation \mathbf{x}_K , we use the estimated polynomial coefficient $\hat{\beta}$:

$$(B.4) \quad \hat{\mathbf{y}}_K = \sum_{q=1}^Q \hat{\beta}_q z_q(\mathbf{x}_K).$$

The algorithmic flow chart is as follows:

Algorithm 2 Predict $\hat{\mathbf{y}}_K$ from $(\mathbf{x}_i, \mathbf{y}_i)_{i \leq K-L}, \mathbf{x}_K$ with local regression.

Select r (size of regression neighborhood), obtain h_K from order statistics of $\|\mathbf{x}_i - \mathbf{x}_K\|$.

Select kernel κ and compute regression weights w_i according to (B.2).

Compute prediction model coefficients $\hat{\beta}$ according to (B.3). For lag-one state augmentation with second order polynomial prediction model, $p = 2, d = 2$, and $Q = 9$, so computing $\hat{\beta}$ requires the inversion of a 9×9 matrix $Z^T W Z$ and then multiplying it by a 9×1 vector.

predict the response $\hat{\mathbf{y}}_K$ using (B.4).

B.1.3 Variations that Potentially Improve Prediction Performance

We now describe two design variations that have the potential to improve prediction performance: using an iterative weighting scheme to increase robustness to outliers in regression (Section B.1.3), and dynamically updating the training atlas to account for temporal variations and/or trends (Section B.1.3 and Section B.1.3).

Robust Local Weighted Regression with Iterative Weight Assignment

It is possible that the training set based on state space distance includes abnormal covariate-response pairs due to noisy observation, or abrupt (and non-repetitive) changes such as patient coughing, and thus they may not be “representative” of the predictor pattern for the given state. To help the local regression method to be robust to such outliers in the $(\mathbf{x}_i, \mathbf{y}_i)$ pairs, we can diminish the weight of a sample covariate-response pair whenever it is inconsistent with the smooth regression from its neighbors. To quantify such inconsistency, we can compare each response value \mathbf{y}_i with its predicted value $\hat{\mathbf{y}}_i = g(\mathbf{x}_i)$. Intuitively, the distance between the observed response \mathbf{y}_i and its estimate $\hat{\mathbf{y}}_i$ indicates how different the particular covariate-response pair behaves than its neighbors. Cleveland [19] has suggested a robust weighting scheme based on a bi-square function B defined as follows:

$$B(x) = \begin{cases} (1 - x^2)^2, & \text{for } |x| < 1 \\ 0, & \text{for } |x| \geq 1. \end{cases}$$

Let $e_i = \mathbf{y}_i - \hat{\mathbf{y}}_i$ be the residual of the observed response from the current fitted value. Let s be the median of the $|e_i|$ for $i = 1, 2, \dots, K - L$. Define the robustness weights by

$$(B.5) \quad \delta_i = B(e_i/6s).$$

The original weight $w_i(\mathbf{x}_j)$ that determines the “contribution” of the i th sample covariate-response pair in estimating the j th response $\hat{\mathbf{y}}_j$ is then modified to be $w_i(\mathbf{x}_j) := \delta_i w_i(\mathbf{x}_j)$, reducing the effect of outliers in fitting the other data points. We apply this re-weighting procedure several times, and use the robust adjusted $w_i(\mathbf{x}_K)$ in place of w_i in (B.3) for estimating the local polynomial coefficient $\hat{\beta}$. This is practical since the δ_i values involved in adjusting the local weight depend only on the i th sample fitting quality, and are independent of the predictor. Plugging $\hat{\beta}$ in (B.4) results in a predicted response value $\hat{\mathbf{y}}_K$. Since the estimation of local polynomial coefficients discounts the effect of outlier samples, the result predictor is expected to be robust to outlier behavior in the “training set” as well. Note that robust local regression could be combined with other methods if needed.

Modified Weight Assignment with Exponentially Discounted Temporal Correlation

Fading memory is present in many natural processes. In breathing trajectories, temporally adjacent sample points tend to be more similar than the sample points further away from one another. To incorporate this property in prediction, we adjust the weights by applying an exponential discount as a function of the temporal distance. Specifically, we modify the weights as follows:

$$(B.6) \quad w_i(\mathbf{x}_j) := \exp(-\alpha|i-j|)w_i(\mathbf{x}_j).$$

The positive constant α determines the decaying rate of influence of one sample on another one as their temporal distance increases. As a special case, $\alpha = 0$ corresponds to no temporal discounting for the sample contributions, but dynamically adds the new samples into the training atlas as they become available.

Temporally Windowed Training Samples

Alternatively, we can modify the weights using a temporal moving window as follows:

$$(B.7) \quad w_i(\mathbf{x}_j) := \begin{cases} w_i(\mathbf{x}_j) & |i-j| < \Gamma \\ 0 & \text{otherwise,} \end{cases}$$

where Γ is the window size. Here only samples that are close enough in time contribute to the local regression with weights determined by (B.2). The length of the window needs to be chosen long enough to guarantee enough samples for the local regression.

B.1.4 Baseline Methods for Comparison

It is desirable to decrease radiation dose due to imaging in image-guided radiotherapy (IGRT). This means we would prefer to predict with low-frequency observation samples (small ψ). On the other hand, it takes time to acquire each observed sample, process it and move the hardware (linac, MLC or cyberknife) accordingly. Thus a system capable of large lookahead lengths is preferable. These two requirements are challenges in prediction, and trade-offs between them need to be considered. More specifically, with lookahead length determined by the limitation of system response, we want to determine the smallest measurement rate that still guarantees certain prediction accuracy. We will study the performance of the proposed method when lookahead lengths and sampling rates are varied, and compare that with some baseline approaches described as follows.

Following [109], we use some commonly used predictors for baseline comparison. We briefly describe their setups and optimization for free parameters in this section.

- Most Recent Sample

This method simply uses the last sample value:

$$\hat{\mathbf{y}}_K = s_K.$$

There are no parameters to be estimated.

- Linear Predictor ⁴

The response is predicted as a linear combination of the previously known positions. This corresponds to a simple model:

$$\hat{\mathbf{y}}_K = \boldsymbol{\beta}^T \mathbf{x}_K + \beta_0.$$

Given a training set, and for a fixed history length the optimal coefficients $\boldsymbol{\beta}, \beta_0$ in terms of mean squared error can be obtained by solving a linear system.

⁴The “linear extrapolation” method described in [109] is a special case of linear prediction.

- Artificial Neural Networks (ANN)

We investigate a multilayer perceptron (MLP) with two feed-forward layers as the ANN predictor [63]. The first layer takes in a fixed history of samples and a constant value 1, linearly transforms the inputs and then uses a sigmoid function to generate the hidden values. The equation for the first layer is

$$h_j(\mathbf{x}) = \frac{1}{1 + \exp(-\gamma_j^T \mathbf{x} + \gamma_{j,0})}.$$

The second layer is chosen to be a simple linear system, and the output is given by

$$\hat{\mathbf{y}}_i = \boldsymbol{\eta}^T \mathbf{h}(\mathbf{x}_i).$$

Parameters $\boldsymbol{\gamma}$ and $\boldsymbol{\eta}$ are estimated from the training set. We use Netlab toolbox [83] to implement ANN in Matlab.

We have also implemented a Kalman Filter for comparison, using Expectation-maximization (EM) method for parameter selection [78], and applied those values for prediction. Our results agree with [109] that the Kalman Filter provides inferior performance compared to ANN. For conciseness, we skip reporting them in this paper. A related research worth noting is the adaptive linear filter model introduced in [125], which can be interpreted as Kalman Filter not in the state, but in linear regression coefficient vector. Unsurprisingly, it shares the limitation of Kalman Filter due to the nonstationarity of respiratory signal.

B.2 Materials

We used the Real-Time Position Management (RPM, Varian Medical Systems, Palo Alto, CA) system to obtain the trajectory of an external fiducial placed on the patient’s chest wall. The recorded displacement-time relationship is believed to be highly correlated with superior-inferior diaphragm motion [125], which is the major source of respiratory motion for tumours in the chest or lung area (the displacements in left-right and anterior-posterior direction are normally on the order of one magnitude lower). To better reflect the behavior of physical superior-inferior motion, the unit-less RPM data were centered and scaled so that their dynamic range matches that with typical SI motion for chest and lung tumours. Table B.1 summarizes the RPM data used in our experiment⁵. Figure B.2 illustrates two typical breathing trajectories.

Subject ID	1	2	3	4	5	6	7	8	9	10
STD	4.96	4.99	3.01	1.99	3.16	1.73	6.27	5.65	2.74	5.29
P-P	25.36	23.65	12.67	11.24	18.72	9.70	28.79	21.89	12.19	21.55

Table B.1: RPM Dataset information.

B.3 Results and Discussions

B.3.1 Scatter Plot in Augmented Space

We first consider a simple $p = 2$ dimensional state vector $\mathbf{x}_i = [s_i, s_{i-\Delta}]$. The response variable is of the form $\mathbf{y}_i = s_{i+L}$. Figure B.3 shows a 3-dimensional scatter plot of $(\mathbf{x}_i, \mathbf{y}_i)$ with the base-line X-Y coordinate reflecting the covariate \mathbf{x}_i and the Z coordinate indicating the corresponding response variable value \mathbf{y}_i . The covariate-response structure is rather smooth, motivating our use of local regression to predict a response from the samples in the neighborhood of the projection onto the X-Y plane. Roughly speaking, the pattern suggests the existence of a conceived functional g that maps the covariate to the response.

We started with a fine sampling rate of $\psi = 30\text{Hz}$ and used only *lag - one* delay with $\Delta = 12$ that corresponds to 0.4 second to augment the state space. We investigate a lookahead length of $L = 30$, which

⁵The data are adjusted to have globally zero mean; average periods are estimated with subspace projection method [101].




Figure B.2: Typical breathing trajectories: (a) rapid yet regular breath; (b) slow yet irregular breath.

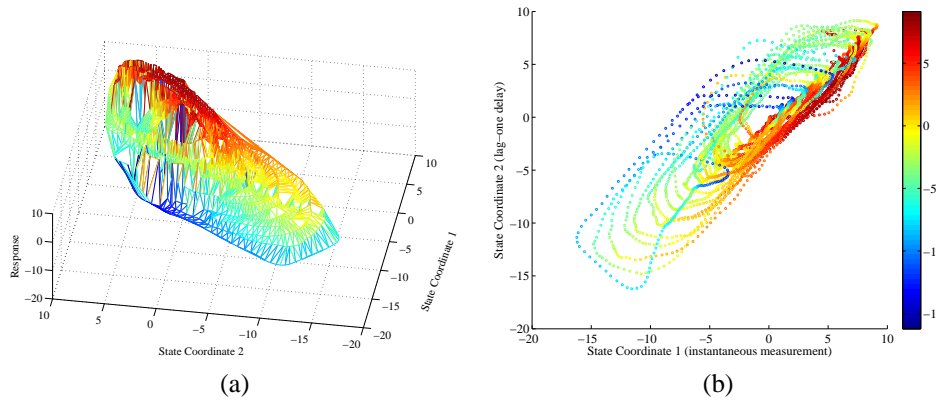


Figure B.3: Covariate-response relationship with lag-one augmented state: (a) 3-dimensional Delaunay tessellation plot; (b) 2-dimensional scatter plot with color indicating the response value.

is equivalent to a 1 second prediction. We used these parameters as defaults in later experiments (*e.g.*, in Section B.3.4). This lookahead length is reported to be difficult by [125] and [109] with a wide spectrum of common prediction techniques. In particular, in the comparative study in [109], the best performance among linear predictors, Kalman filter and artificial neural networks yields a RMSE of about 5mm, with similar data statistics to our rescaled RPM data⁶. Lag-one augmentation corresponds to regression based on the most current sample s_i and one preceding observation $s_{i-\Delta}$, which is the most compact model possible. The temporal lag Δ for augmentation should be chosen to reflect the system dynamics properly and robust enough in the presence of observation noise, and does not have to be unity.

B.3.2 Local Weighted Regression without Temporal Discounting

To illustrate the performance of the simple local weighted regression method described in Section B.1.2, we conduct two simple experiments with the following configurations: we used the “tricube” function [19] as the weighting kernel κ and chose the effective bandwidth so that the local regression is supported on half of the samples. Specifically, we used

$$\kappa(x) = \begin{cases} (1 - |x|^3)^3, & \text{for } |x| < 1 \\ 0, & \text{for } |x| \geq 1. \end{cases}$$

⁶The research conducted in [109] uses 3-dimensional position, which is presumably more complex than this study. However, since motion in SI direction dominates the overall respiratory variation in general, we expect the prediction error to be the major contributor to the overall tracking/prediction performance. Rescaling the RPM data to have similar statistics as typical SI motion, we feel it fair and illuminating to compare quantitatively the performance of our predictor to that of the general 3D predictors.

The neighborhood size r was chosen to be 200, which is equivalent to about 7 seconds worth of samples. Accordingly, h_K is the 200th smallest number among $\|\mathbf{x}_i - \mathbf{x}_K\|$.

B.3.3 Robust Local Regression with Iterative Weighting

We investigated the robust iterative weighting of Section B.1.3, but found that iterative weighting did not significantly change the prediction errors in this experiment. This suggests the absence of dramatic outliers in our experimental data.

B.3.4 The Effect of Dynamically Updating the Training Set

If the training set is determined before the treatment process, and is kept the same thereafter, the corresponding local regression structures are also fixed. This is the “static” inference scenario. It is also possible to “add” (or “substitute” the oldest sample with) new samples into the training set during the treatment process, as new responses become available. We refer to the latter approach as “dynamically updating of the training set”.

The computation for simple local regression is the same regardless of whether we update the training atlas or not, as it uses only the training samples that fall into the neighborhood of the target. On the other hand, when robust local regression with iterative weighting is applied, choosing between static training and dynamic training makes a difference. In the static scenario, the robust weights can be computed offline upon the availability of all the training samples, and are kept the same thereafter. However, if we use dynamic updates, not only does the size of the “atlas” grow with time, but there would also be changes in the robust weights, since the newly available covariate-response samples can potentially change the regression weights for those existing samples whose supports overlap with those of the new ones. This effect can propagate through the whole atlas.

At the cost of possible additional computation, dynamically updating the training atlas admits new information as the time proceeds. This is particularly valuable when the underlying system dynamics demonstrate strong temporal variation, such as frequency change or home position drifting, which are both commonly observed. New samples can either be added to the training set as either simple addition, which corresponds to a collective history case, or substituted for the oldest training sample, as in the windowed training history case. In both cases, experiment results indicate that dynamically updating the training set yields overall superior prediction performance in terms of root mean squared error (RMSE) and mean absolute error (MAE), as we report in Section B.3.4 and Section B.3.4.

Dynamically Expanding the Training set

Using a discount factor $\alpha = 0$ in (B.6) to adjust weight for the training samples up to the most currently available one is equivalent to building a collective atlas that includes all previous covariate-response pairs. Of course, new training pairs are entered into the the atlas as time proceeds. Table B.2 reports the prediction performance for one second lookahead with 5Hz sampling using this dynamic training structure as opposed to a static 20 second training at the beginning of the fraction.

Table B.2: Comparison of Prediction Performance among Static Training, Dynamic Expanding Training, and Updating Training with Moving Window

Subject ID	1	2	3	4	5	6	7	8	9	10	Average
Root Mean Squared Error (RMSE)											
static	9.7	3.6	2.2	1.9	10.8	5.6	4.9	4.2	2.8	4.4	5.0
expand	3.4	2.8	1.6	1.4	2.5	1.3	4.8	2.6	2.1	3.7	2.6
update	2.7	2.5	1.4	1.4	2.6	1.3	4.8	2.5	2.1	3.5	2.5
Mean Absolute Error (MAE)											
static	7.5	2.6	1.7	1.4	3.9	2.6	3.7	2.5	2.1	3.1	3.1
expand	2.6	2.1	1.2	1.1	1.7	1.0	3.5	1.7	1.6	2.6	1.9
update	2.0	2.0	1.1	1.0	1.7	1.0	3.4	1.7	1.4	2.5	1.8

Figure B.4 illustrates improved prediction performance by dynamically expanding the training set for breathing traces that either exhibit mean drifting or pattern changes. There is minimal benefit when the breathing pattern is already fairly regular or irregular with no “trend”, and new observations simply add to the already sufficiently dense training atlas. Change detection may be used to locate some local variations, but this imposes extra complexity ⁷.

Dynamically Updating Training Set with Windowed History

Alternatively, a moving window can be used to update the sample set. This corresponds to substituting the oldest samples with the newly available covariate-response pairs, as discussed in Section B.3.4. We illustrate the effect of this dynamic updating method in Table B.2 and Fig. B.5. A dynamic window of length 20 second is used in all of our experiments. We used the performance of dynamic expansion as a baseline for the windowed study.

For the 20sec training window, the overall prediction performance improves upon the previously discussed dynamic expansion. The level of improvement, though, is much smaller than the one we obtained by going from static training to dynamic expansion. Some trade-offs are expected: for long fractions, it is more likely that the later samples are decoupled from the samples acquired at the very beginning of the procedure, thus moving window method should be favorable; on the other hand, dynamic expansion does not require choosing a window length, and it is almost free of the risk of running into insufficient samples for the local inference, thus has the advantage of being simple and stable. There is little difference in the prediction performance between the two methods from Fig. B.5 except that in the mean drifting case, the windowed update may be slightly better, which is also reflected quantitatively in Table B.2.

B.3.5 The Effect of Measurement Rate and Lookahead Length

We compared the local regression (LOESS) method using expanding training atlas with the baseline approaches described in Section B.1.4. In particular, we compared with most recent sample (MRS), linear prediction (Linear), Kalman Filter (KF) and Artificial Neural Networks (ANN) when lookahead length and sampling rates are varied. Figure B.6 and Fig. B.7 report the results in terms of the collective root mean squared error (RMSE) and mean average error (MAE) across all patients. In general, the prediction errors increase as sampling frequency ψ decreases and/or lookahead length L increases, as expected. Interestingly, the proposed local regression method is insensitive to the sampling rate, and performs almost consistently across different lookahead lengths. Unlike the most recent sample, linear model or Kalman filter, which not only make assumptions about the underlying model structure (linearity), but also try to explicitly solve for the model parameters, LOESS makes none of the above assumptions or effort. The nonparametric nature of the regression avoids assuming a fixed model structure, let alone solving for it. The only requirement is consistent behavior (or *existence* of an underlying functional form). A polynomial of sufficient order approximates this underlying function via fitting samples in the neighborhood of the point of interest. This also explains, to some extent, why ANN outperforms the other approaches [79], as it is a combination of local linear perceptrons, with extra nonlinearity provided by the sigmoid activation function. When lookahead length is short and sampling rate is high, linearity holds approximately, and all methods provide reasonably good prediction. However, when we need to look further ahead, linear models are not sufficient to capture the dynamics, even though the response pattern may still be consistent, and that is where LOESS (and ANN) demonstrates its advantage. Figure B.6 reports the relationship between collective prediction error (across all testing subjects) and lookahead length for sampling rate $\psi = 5Hz$ and Fig. B.7 illustrates how collective prediction error change with different sampling rates. LOESS approach performs competitively with ANN for lookahead length 0.6 seconds, in particular for low sampling frequency, and it demonstrates an obvious advantage for lookahead length 1 second.

⁷Segmentation based tracking/prediction model [133] follows similar logic, yet requires further research to improve robustness and automation.




Figure B.4: Effect of dynamically updating the training atlas: actual signal time history (blue solid line), prediction from static training (black dash-dot line) and prediction from dynamic expanding atlas (red dashed line). (a) breathing with mean drifting (Patient 1); (b) In the presence of changing breathing pattern (Patient 6); (c) with complicated transient interrupting regular breathing (Patient 8); (d) quasi-regular breathing pattern (Patient 10).

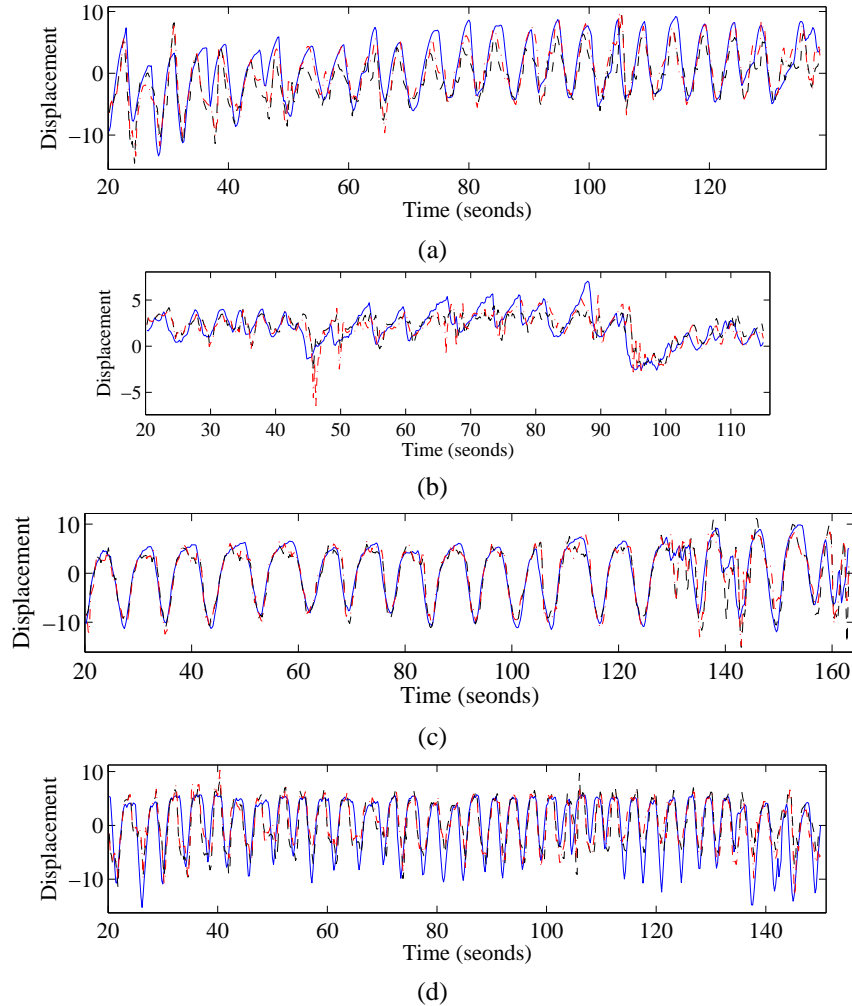


Figure B.5: Comparison of prediction performance using dynamic update with moving windowed and expanding training atlas: actual signal time history (blue solid line), dynamic expanding training (black dash-dot line), moving window adaptive training (red dashed line).

B.4 Conclusion and Future Work

In this paper, we have proposed a local regression based method to predict respiratory motion. We compared the proposed method and conventional approaches such as most recent sample, linear model, Kalman filtering, and artificial neural networks. The proposed method had lower prediction error than the others for tasks requiring long lookahead prediction. We have also discussed extensions and variations of the basic method to provide robustness to outliers that may be caused by low SNR or miss-tracking. We studied the prediction performance with different error metrics (RMSE and MAE) for various combinations of lookahead length and sampling frequency. The proposed method showed the most advantage for long lookahead lengths and low sampling rates.

We have discussed the challenge of choosing a good discount factor for weight adjustment in local regression in Section B.1.3 and discussed the two simple cases corresponding to either no forgetting or inference from windowed historical sample. The proper choice of the temporal discount factor depends on the variation of the underlying breathing pattern, and automatic schemes should be investigated.

As observed in our experiment, various phases of respiratory motion are predicted with different accuracies. Respiratory motion demonstrates obvious non-stationarity: the system variation at the transition phase could be very different than that during extreme tidal stages (end-inhale or end-exhale). From an-




Figure B.6: Collective performance comparison for different lookahead lengths. With sampling rate $\psi = 5\text{Hz}$, (a) root mean squared error (RMSE), (b) mean absolute error (MAE).

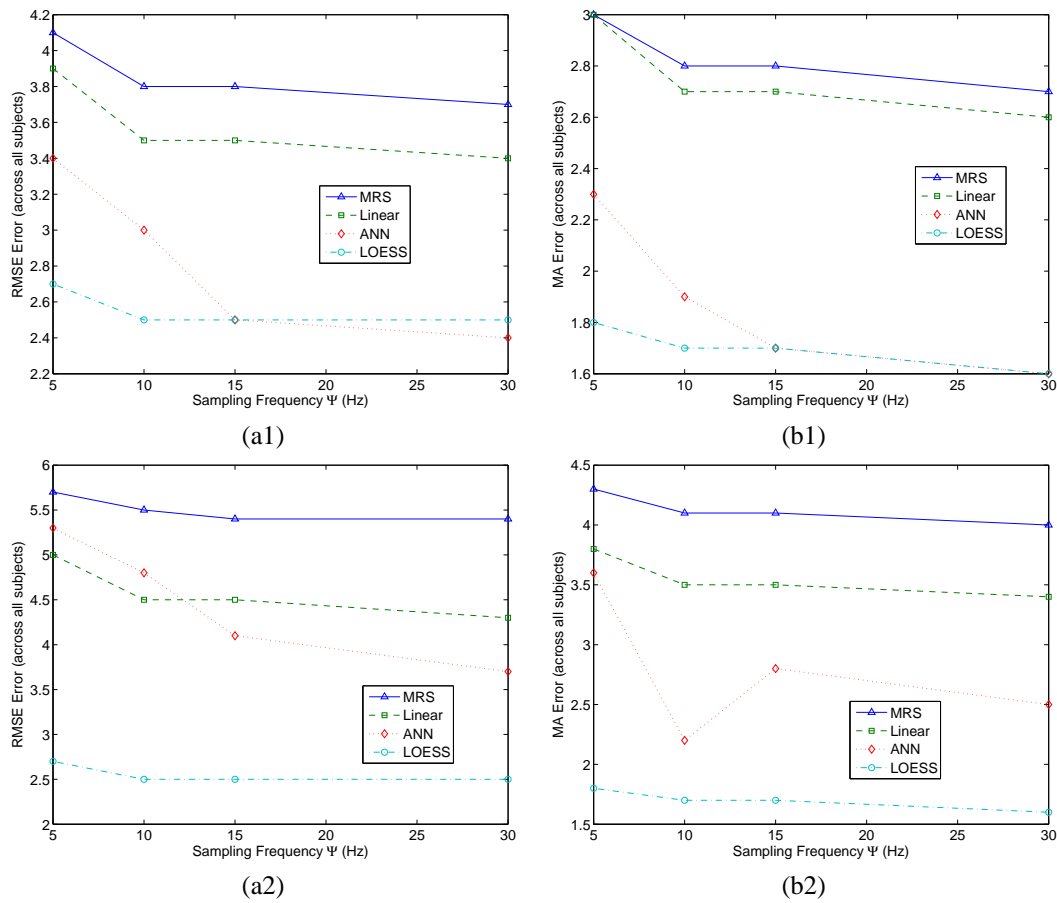


Figure B.7: Collective performance comparison for different sampling rates ψ . Left column [a#] RMSE, right column [b#] MAE. Top row [X1] illustrates the results from a lookahead length of 0.6 second and bottom row [X2] shows the results when lookahead length is 1 second.

other perspective, if we examine the signal-to-noise ratio (SNR) over a windowed portion of the signal, SNR would change as the window covers different stages of the breathing: SNR would be relatively high

during transition stage, as the signal variation is big relative to noise, while SNR is low at the plateau stages, which correspond to end of inhale or exhale. These observations motivate a potential research topic: if we aim at homogeneous prediction performance throughout the breathing trajectory, it may be necessary to use adaptive sampling. More precisely, denser sampling may help where prediction uncertainty is big whereas sparser sampling should suffice where prediction is more reliable. This is a topic for future study.

The dynamics of respiratory motions change over various stages of breathing, and makes general prediction difficult. Models using state dependent transition probabilities have been investigated for stochastic tracking [103], and explicit segmentation was also studied [133]. Our proposed method uses local kernel regression to capture this variation implicitly by essentially limiting inference to a neighborhood of training samples that are expected to behave similarly. Intuitively, this is almost equivalent to training a local model at each state of interest. Since the state distance (and thus the inference weight) is assigned with respect to Euclidean distance in state space, it is important that clustering with this distance reflect dynamic similarity. This is expected in most cases, except when home (mean) drifting is high both in frequency and displacement value. In the exceptional case of dramatic mean drift, samples belonging to different breathing stages may be clustered together. One straight-forward remedy would be to incorporate mean drifting compensation in the inference weight. A robust mean tracking algorithm for respiratory motion is provided in [99] that outputs mean position estimates for both the training samples and the state of interest. We expect improved accuracy by accounting for mean position drifting. We plan to conduct further experiments and analyze this effect in more detail in the future.

APPENDIX C

Inference of Internal Respiratory Tumor Motion from External Surrogates: A State Space Augmentation Approach in Modeling Hysteresis ¹

Respiratory motion affects tumours in the thorax and abdomen. In particular, breathing is the major reason for intrafractional tumour motion for lung cancer patients. It is important to monitor such motion during radiotherapy treatment to ensure the accurate delivery of radiation dose in motion-compensated Intensity Modulated Radiotherapy (IMRT). Fluoroscopic imaging or portal imaging can monitor tumour motion during the treatment process. To reduce x-ray exposure, hybrid tumour tracking approaches that combine episodic radiographic imaging and continuous external surrogates have been investigated widely [80, 81, 86, 104, 105]. Using external surrogates to infer internal tumour motion assumes that there is consistent relationship between internal and external motion.

Hysteresis is typical in lung tumour movements, with the tumour taking a different path during inhale and exhale. Inhalation normally takes longer than exhalation, and the deflating lung volume exceeds the inflating volume at the same trans-pulmonary pressure [57]. Respiratory hysteresis makes inferring internal tumour locations from external surrogate signals challenging. Most of the external surrogate systems, such as thermistors, thermocouples, strain gauges, pneumotachographs [60], and infrared skin markers as applied in the Varian Real-time Position ManagementTM (RPM) system (Varian Medical Systems, Palo Alto, CA), provide one-dimensional signals, whose instantaneous amplitude (or displacement) alone does not provide sufficient information about the specific breathing stages.

Previous studies about correspondence between internal tumour motion and external surrogates can be classified into two categories. One class of studies investigates the correlation between the two signals to justify the feasibility of using certain types of surrogates, or compare different surrogate options (including the placement mechanism) [1, 47, 59, 73, 117, 124, 127]. Alternatively, some other studies assume *a priori* the existence of a strong correlation between internal and external signals, and aim to estimate the correspondence map [106]. We adopt the latter perspective and study with a general setup the correspondence maps that take the external surrogate trace as input and output estimates of the internal tumour location, including, but not restricted to linear relations as reflected by the correlation coefficient and its variants. The presence of respiratory hysteresis makes this a challenging problem, as the same external surrogate position can reflect different internal tumour locations during different phases. Existing methods address hysteresis by first separating empirically the breathing trajectories into two distinct “directions” (inhale v.s. exhale), and then constructing a piecewise phase-dependent map [66, 67, 106, 107]. However, subdividing the breathing into inhale and exhale phases often requires manual intervention, and is infeasible for real-time application, because a breathing “peak” or “trough” can be only be identified retrospectively.

In this study, we propose to use a simple state augmentation of the external surrogate signal. Augmenting the state space with self-delayed observation bestows the model with “memory”, which is an alternative way to characterize the “path-dependence” property of hysteretic systems. This procedure captures system dynamics, and embeds the breathing phase information implicitly into the framework. We then provide the solution to a general class of parametric inference models with the augmented observations. As special cases,

¹This chapter is based on materials from [100]

we derive optimal solutions for the parameters of linear and quadratic correspondence models. Furthermore, given a training internal/external dataset, we demonstrate a computationally efficient approach to choose a patient-specific (or fraction-dependent) augmentation scheme. Generalization to adaptive correspondence models follows naturally. We test the proposed approach on synchronized recordings of internal gold marker trajectories and external fiducial marker locations [5].

Section C.1 describes the clinical data used for this test, discusses the challenges caused by hysteresis in converting the external surrogate position directly to internal tumour location and presents the proposed method. A general correspondence model is formulated with polynomial models as an example. Optimal model parameters are derived and generalization is given to accommodate adaptivity. Section C.2 reports testing results followed by discussions. Section C.3 concludes this study with a brief summary.

C.1 Methods and Materials

C.1.1 Data Description

To study the internal/external motion correspondence, we obtained synchronized recordings of internal tumour motion trajectories and external fiducial marker locations. The paired trajectories from eight lung cancer patients were collected with a Mitsubishi real-time radiation therapy (RTRT) system at the Radiation Oncology Clinic at the Nippon Telegraph and Telephone Corporation (NTT) hospital in Sapporo, Japan. Two to four 1.5mm diameter gold ball bearings (bb's) were implanted in or near the tumour [111] and these internal markers were tracked in real time with diagnostic x-ray fluoroscopy [110]. External surrogate signals were obtained with the AZ-733V external respiratory gating system (Anzai Medical, Tokyo, Japan) integrated with the RTRT system. It uses a laser source and a detector, both attached to the treatment couch with the beam placed orthogonal to the patient's abdominal skin surface. The device calculates the change in the surface amplitude by measuring the relative position of the reflected light [5] and outputs a one dimensional relative position measurement of the abdominal surface. The data acquisition rate for the entire system is 30 frames per second. Table C.1 describes the study participants. All patients included in this analysis had peak-to-peak marker motion greater than 1cm. The KV fluoroscopy + Anzai system took multiple readings for each fraction from several treatment field configurations to account for obscured x-ray views as the gantry rotated. The recording lengths varied between 20 and 250 seconds with an average of 82 seconds. There are in total 128 readings, 46 of which were longer than 100 seconds.

Patient	Gender	Age	Tumour Pathology	# of bb's	Tumour Site	Prescribed Dose (Gy)	# of Fractions
1	F	47	Adenocarcinoma	4	R S7	N/A	1
2	F	70	Adenocarcinoma	3	L S6	N/A	1
3	F	71	Adenocarcinoma	2	R S5	N/A	1
4	F	47	Adenocarcinoma	3	R S4	48	8
5	M	81	Squamous cell carcinoma	3	R S6b	48	4
5	M	81	Squamous cell carcinoma	3	R S6b	40	8
6	M	61	small cell lung cancer	3	R S10	40	8
7	M	68	Squamous cell carcinoma	3	R S6	48	4
8	M	85	Adenocarcinoma	3	R S8	48	4

Table C.1: Description of study participants. Patients 1-3 were brought in for data acquisition purposes only, so there is no prescription dose. Patient 5 was treated twice at the same site, with two months between treatments. The tumour site is indicated using the common anatomical notation for lung segmentation: S1-3 is upper lobe, S4-5 is middle lobe and S6-10 is lower lobe.

C.1.2 A General Correspondence Model

To minimize diagnostic imaging dose in IGRT systems, it is important to infer internal tumour location from external surrogates. In principle, we could use a correspondence model that observes a trajectory \vec{r} of the scalar external surrogate r up to time instant n to infer the 3-dimensional internal tumour position

$\mathbf{p} = (x, y, z)$. We denote the collective surrogate information available at time n as $\bar{\mathbf{r}}(n) \triangleq \{\mathbf{r}(m) : 0 \leq m \leq n\}$. However, it is challenging to estimate such a map that estimates the internal tumour position from the complete collection of historical surrogate data, since the length of the input variable grows to infinity as time progresses. A more practical choice is to use some much more compact quantity \mathbf{r} that captures sufficient information from $\bar{\mathbf{r}}$ for inference. With internal and external motion both being smooth, it is reasonable to approximate $\mathbf{p}(\mathbf{r})$ using polynomials. Therefore, we focus on estimating a class of correspondence models that are linear in their coefficients as follows:

$$(C.1) \quad \hat{\mathbf{p}}(\mathbf{r}) = \mathbf{A}\mathbf{f}(\mathbf{r}),$$

where \mathbf{f} is a vector function of external surrogate \mathbf{r} ; all model parameters to be optimized are contained in the coefficient matrix \mathbf{A} . In particular, two simple correspondence models, *i.e.*, a linear model and a quadratic model introduced in [106] are special cases of the form given in (C.1).

Linear models assume each coordinate of internal motion is affine in $\mathbf{r} = r(t)$. This corresponds to the case where

$$(C.2) \quad \mathbf{f}(\mathbf{r}) = \begin{bmatrix} r \\ 1 \end{bmatrix} \quad \text{and} \quad \mathbf{A} = \begin{bmatrix} b_x & c_x \\ b_y & c_y \\ b_z & c_z \end{bmatrix}.$$

Quadratic models map the external surrogate to each coordinate of internal motion via a quadratic relation. It can be expressed in (C.1) with

$$(C.3) \quad \mathbf{f}(\mathbf{r}) = \begin{bmatrix} r^2 \\ r \\ 1 \end{bmatrix} \quad \text{and} \quad \mathbf{A} = \begin{bmatrix} b_x & c_x & d_x \\ b_y & c_y & d_y \\ b_z & c_z & d_z \end{bmatrix}.$$

The expression in (C.1) is linear in the model coefficients \mathbf{A} and yields a closed form optimal solution in the least squared error (LSE) sense. Given N sample points $(\mathbf{r}_n, \mathbf{p}_n)$, $n = 1, 2, \dots, N$, the solution to the LSE problem:

$$(C.4) \quad \hat{\mathbf{A}} = \arg \min_{\mathbf{A}} E(\mathbf{A}),$$

where $E(\mathbf{A}) = \sum_{n=1}^N \|\mathbf{p}_n - \mathbf{A}\mathbf{f}(\mathbf{r}_n)\|^2$, is given by solving the normal equation [68], and

$$(C.5) \quad \hat{\mathbf{A}} = \mathbf{P}^T \mathbf{F} (\mathbf{F}^T \mathbf{F})^{-1},$$

where $\mathbf{F} = \begin{bmatrix} \mathbf{f}(\mathbf{r}_1)^T \\ \vdots \\ \mathbf{f}(\mathbf{r}_N)^T \end{bmatrix}$ and $\mathbf{P} = \begin{bmatrix} \mathbf{p}_1^T \\ \vdots \\ \mathbf{p}_N^T \end{bmatrix}$. The corresponding residual is given by:

$$(C.6) \quad \begin{aligned} \Delta \mathbf{P} &\triangleq \mathbf{P} - \mathbf{F} \hat{\mathbf{A}}^T \\ &= (\mathbf{I} - \mathbf{F} (\mathbf{F}^T \mathbf{F})^{-1} \mathbf{F}^T) \mathbf{P}, \end{aligned}$$

with overall residual error (summed over all 3-dimensions) as:

$$(C.7) \quad \begin{aligned} E(\hat{\mathbf{A}}) &= \text{trace}\{\Delta \mathbf{P}^T \Delta \mathbf{P}\} \\ &= \text{trace}\{\mathbf{P}^T (\mathbf{I} - \mathbf{F} (\mathbf{F}^T \mathbf{F})^{-1} \mathbf{F}^T) \mathbf{P}\}. \end{aligned}$$

It may be preferable to have simpler models (with fewer free parameters) over more complicated models at the cost of small sacrifice in data fitting performance. This model selection preference can be incorporated into the optimization setting by modifying the objective function as:

$$(C.8) \quad \tilde{E}(\mathbf{A}) = E(\mathbf{A}) + \lambda R(\#\mathbf{A}),$$

where $\#\mathbf{A}$ denotes the number of free parameters in the coefficient matrix \mathbf{A} , and R is a monotonically increasing function that assigns higher costs to more complicated models. The regularization weight λ controls the tradeoff between the data fitting $E(\mathbf{A})$ and the preference for lower-order models. A simple example of R would be the linear function $R(\#\mathbf{A}) = \#\mathbf{A}$, which directly penalizes the number of components in \mathbf{A} ; this is equivalent to the Akaike Information Criterion [2]. Using the closed form optimal solution (C.5) and the expression for optimal residual error (C.7) for a given fixed model structure, the modified objective function can be minimized in two layers. We say two inference models have the same “*model structure*” if they only differ in parameter values. It follows immediately that all models with the same structure has equal number of degrees of freedom, thus the same complexity regularization $R(\#\mathbf{A})$ in (C.8). Therefore, to minimize over models of different complexity, it is natural to choose the “best” parameter setting within each model structure (with fixed degrees of freedom thus a constant complexity penalty), and then compare across structures. Within each class, minimizer of the complexity penalized objective $\tilde{E}(\hat{\mathbf{A}})$ is the same as that of $E(\hat{\mathbf{A}})$, and can be solved and evaluated efficiently using the closed form optimal solution (C.5) and expression for optimal residual error (C.7). This motivates the two-layer hierarchical algorithm shown below for finding the optimal solution within K candidate model structures $\mathcal{C} = \cup_{i=1}^K \{C_i\}$.

Algorithm 3 Two-layer Optimization Routine for Solving $\hat{\mathbf{A}} = \arg \min \tilde{E}(\mathbf{A})$ (C.8).

```

1:  $\tilde{E} \leftarrow +\infty$ ;  $i_{\text{opt}} \leftarrow 0$ ;  $\hat{\mathbf{A}} \leftarrow []$ .
2: for  $i = 0$  to  $K$  do
3:   Choose model structure  $C_i$  from the collection of models  $\mathcal{C}$ ,
4:   Compute  $R_i = R(\#\mathbf{A})$  for structure  $C_i$ ;
5:   Compute  $\hat{\mathbf{A}}_i$  within class  $C_i$  according to (C.5) and its residual error  $E(\hat{\mathbf{A}}_i)$  from (C.7).
6:   if  $E(\hat{\mathbf{A}}_i) + R_i < \tilde{E}$  then
7:      $\tilde{E} \leftarrow E(\hat{\mathbf{A}}_i) + R_i$ ;
8:      $i_{\text{opt}} \leftarrow i$ ;
9:      $\hat{\mathbf{A}} \leftarrow \hat{\mathbf{A}}_i$ .
10:  end if
11: end for

```

C.1.3 Hysteresis and State Augmentation

Conventional methods that explicitly segment the breathing process into inhale and exhale phases have their limitations, as physical phase transitions (and delays) occur continuously rather than as discrete jumps. To circumvent the intrinsic difficulty of estimating breathing phases, we study the system dynamics directly, expecting them to sufficiently convey phase information. In a discretely observed system, one usually captures the system dynamics with time-lagged samples. For the sake of simplicity and to avoid over-parameterization, we restrict this study to a single lag. The proposed method generalizes to multiple-lag models naturally.

Given a discrete-time external surrogate $r(n)$, $n = 1, 2, \dots, N$, we augment each external surrogate state with a time τ (in discrete unit) delayed sample, *i.e.*, $\mathbf{r}(n) \triangleq (r(n), r(n - \tau))$. This augmentation captures first-order system dynamics, as the difference between $r(n)$ and $r(n - \tau)$ can be regarded as a measure of average local velocity. As \mathbf{r} is uniquely determined by \bar{r} , it fits into the general formulation (C.1). We apply the methods provided in Section C.1.2 to estimate the coefficients for the augmented model. To demonstrate the idea, we establish a linear model that is comparable to (C.2) and a quadratic model analogous to (C.3).

The augmented linear model (in \mathbf{r}) represents each internal coordinate as a linear combination of $r(n)$, $r(n - \tau)$ and a constant offset, corresponding to:

$$(C.9) \quad \hat{\mathbf{p}} = \mathbf{A}\mathbf{f}(\mathbf{r}), \quad \text{where} \quad \mathbf{f}(\mathbf{r}) = \begin{bmatrix} r(n) \\ r(n - \tau) \\ 1 \end{bmatrix}$$

with a 3×3 coefficient matrix \mathbf{A} .

The augmented quadratic model (in \mathbf{r}) estimates each internal coordinate as a linear combination of $r^2(n), r(n)r(n-\tau), r^2(n-\tau), r(n), r(n-\tau), 1$, corresponding to:

$$(C.10) \quad \hat{\mathbf{p}} = \mathbf{A}\mathbf{f}(\mathbf{r}), \quad \text{where} \quad \mathbf{f}(\mathbf{r}) = \begin{bmatrix} r^2(n) \\ r^2(n-\tau) \\ r(n)r(n-\tau) \\ r(n) \\ r(n-\tau) \\ 1 \end{bmatrix}$$

with a 3×6 coefficient matrix \mathbf{A} .

In both cases, linearity in \mathbf{A} results in the closed form solution given by (C.5) with the corresponding \mathbf{F} respectively.

C.1.4 Choice of Lag Length

The delay τ should be chosen properly, since too long a lag provides minimal local dynamic information and too short a lag makes the estimation sensitive to observation noise. For inference purposes, we desire a lag that maximally resolves the ambiguity in the estimated correspondence map. We choose the lag that minimizes the fitting error for training data:

$$(C.11) \quad \hat{\tau} = \arg \min_{\tau} E(\hat{\mathbf{A}}(\tau)),$$

with the objective function E defined in (C.4). The coefficients $\hat{\mathbf{A}}$ and the error E depend on τ because \mathbf{f} contains both the current external surrogate displacement $r(n)$ and its lagged state $r(n-\tau)$.

Equations in (C.6) and (C.7) provides a closed-form expression for $E(\hat{\mathbf{A}}(\tau))$ for each given τ . The optimization problem (C.11) simplifies to a simple one-dimensional line search that we solve by searching over an interval with the corresponding delay time between 0 (no lag) and about half of an average breathing period.

C.1.5 Adaptivity of the Correspondence Map

Adaptivity may be useful to accommodate gradual changes in the correspondence models, due to drifting or variations in patients' breathing. In the case of linear and quadratic models, the operation in (C.5) involves inverting fairly small matrices (3×3 and 6×6 respectively), so direct inversion is numerically feasible. However, when more complicated models with higher degrees of freedom are used, it is desirable to reduce computation by applying recursive algorithms that modify current estimates based on newly available data. The key to recursively updating (C.5) is to avoid recomputing $(\mathbf{F}^T \mathbf{F})^{-1}$ from scratch every time. This is effectively the inversion of empirical correlation matrix with observation \mathbf{f}_i . [99] provides rank-one update equations for sliding window and exponential discount adaptivities.

C.2 Results and Discussions

To illustrate the challenges caused by hysteresis, Fig. C.1 shows an example of the relationship between internal tumour location obtained by fluoroscopic imaging and an external surrogate from an abdominal surface measurement as described in Section C.1.1. We depict only the anterior-posterior (AP) coordinate against the surrogate signal, as this axis demonstrates the strongest hysteresis for this test subject. The optimal linear and quadratic correspondence maps [106] provide reasonable inference of internal tumour motion from external surrogates, yet they fail to describe the breathing-phase dependency of an ideal correspondence map. In fact, any function that tries to map the scalar $r(n)$ to \mathbf{p} would experience the same problem, since this is a one-to-multiple relation with hysteresis.

Figure C.2 illustrates the internal tumour location in the anterior-posterior (AP) direction v.s. the state augmented external surrogates for $\tau = 45$, which corresponds to a 1.5 second delay for 30Hz sampling rate. The scatter-plot in Fig. C.2(a) represents each data sample in the $(r(n), r(n-\tau))$ space with a circle, and




Figure C.1: Example of a breathing trajectory with respiratory hysteresis.




Figure C.2: Scatter plot showing the data samples in augmented external state space with the colors indicating internal AP value. Locally consistent colored samples suggests the potential of resolving hysteretic ambiguity by distinguishing among different respiratory phases implicitly with state augmentation.

uses color (or intensity if viewed in gray-scale) to depict the internal AP coordinate values (in mm) from fluoroscopic readout. The one-to-multiple discrepancy appears largely resolved as different colored circles are not overlaid on each other, suggesting the existence of a single-valued inference map.

To illustrate the idea of model fitting in augmented state space, we first apply the simple linear model in (C.9) to the dataset shown in Fig. C.1 with a lag length of 1.5 seconds (which may not be optimal), and illustrate the results in Fig. C.3. Even though there are still noticeable differences between the observed internal coordinates in the upper row of Fig. C.3 and their linear fit in the bottom row, the aggregated estimation error (across all patients and fractions) reduced to 1.74 mm from 2.01 mm with direct linear fitting as in (C.2) and 1.93 mm with direct quadratic fitting as in (C.3). In particular, we observe noticeable decreases in estimation error in the AP direction, where hysteric ambiguity is the most significant. Table C.2 reports the Root Mean Squared Error (RMSE) in each direction respectively for the linear and quadratic model, with and without state augmentation². Figure C.4 reports the paired (across patient/fraction) differences between the RMS error of the direct methods and the augmented methods. The RMSE difference between direct linear and augmented linear methods has mean 0.14mm and a median of 0.11mm; the RMSE difference between direct quadratic and augmented quadratic method has mean 0.17mm and a median of 0.15mm. To assess statistical significance, we performed a paired student-t test with the null hypothesis that the performance of the

²For comparison purposes, we have also computed estimate from the 5th-order polynomial model with direct method, which has the same degrees of freedom (18 parameters) as the augmented quadratic model. Its estimation error is 0.75, 1.25 and 1.11 (mm) in LR, SI and AP direction respectively, with a 3D RMSE equals 1.83mm. A paired student t-test between the RMSE for the 5th-order polynomial model and the augmented quadratic model yields a p -value of 1.06×10^{-10} , which indicates statistically significant error reduction by the augmented quadratic model. This shows that the improved performance of the proposed method is not a direct consequence of increased degrees of freedom, but should rather be attributed to its capability of resolving hysteretic ambiguity via state augmentation.

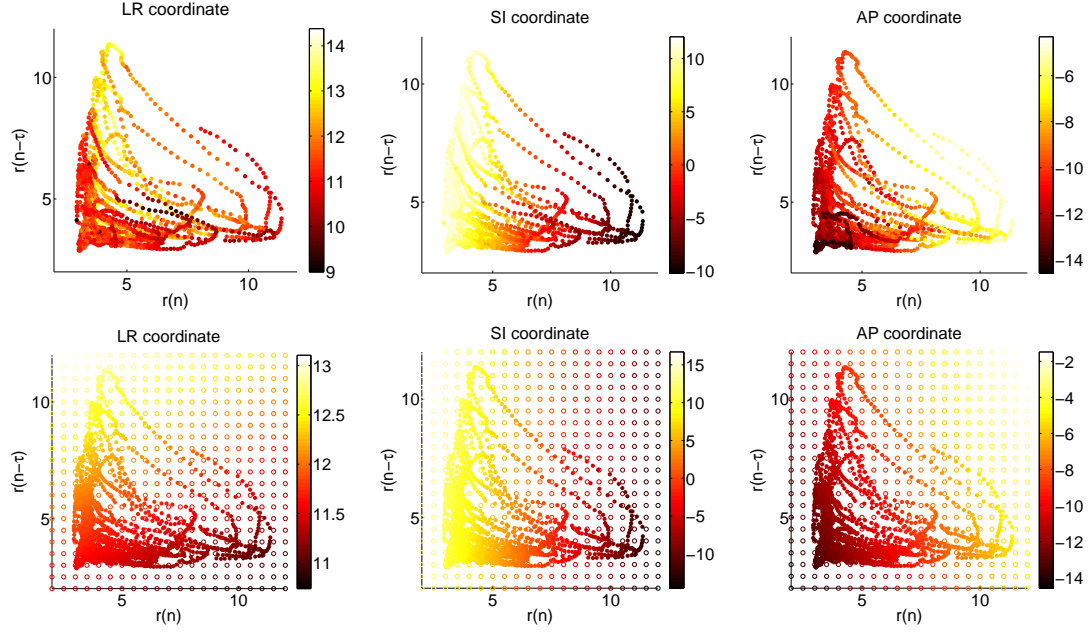


Figure C.3: Correspondence relations in augmented state space and their linear fittings. Upper row: internal tumour coordinate v.s. augmented state for observed samples with colors indicating internal AP value; bottom row: estimates of tumour coordinate via linear fit with hollow circles depicting modeled hypersurface evaluated at regular grid points and solid circles for the evaluation at the sample locations, with colors indicating estimated AP value.

direct and augmented methods do not differ. The p values for the linear method and the quadratic method are 4.96×10^{-13} and 4.08×10^{-18} respectively, demonstrating that the error reductions were statistically significant.

	LR (mm)	SI (mm)	AP (mm)	3D (mm)
Direct Linear	0.80	1.45	1.13	2.01
Direct Quadratic	0.79	1.35	1.13	1.93
Aug. Linear	0.75	1.30	0.87	1.74
Aug. Quadratic	0.74	1.18	0.84	1.63

Table C.2: Estimation Error Table

Figure C.5 shows the estimated time series of these four approaches for converting external surrogates to internal tumour locations. The higher-order models were more descriptive with the extra degrees of freedom, as demonstrated by the relative performance of quadratic models and linear models within each class respectively. State augmentation enables varying response patterns during different stages of breathing as indicated implicitly by the system dynamics.

As discussed in Section C.1.4, to properly choose the lag length τ , we use a short training set with internal-external pairs to compute offline the estimation performance $E(\hat{A}(\tau))$ defined in (C.4) as a function of the lag length τ . In practice, the lag length does not have to be the exact optimum in (C.11); values near that optimum should sufficiently convey system dynamics. Reasonable insensitivity in the choice of lag length τ is desirable as this value is determined prior to the treatment and remains fixed subsequently. Figure C.6 illustrates that the estimation error is a smooth function of the lag length, which suggests the desired robustness. For both the linear correspondence model (C.9) and the second-order polynomial model (C.10) with state augmentation, the optimal τ corresponds to about 1.7 – 1.8 seconds delay. Without this knowledge, our previous experiments used 1.5 seconds delay to augment the state space (Fig. C.2-Fig. C.5), and still yielded plausible results. The asymmetric slopes in Fig. C.11 around the optimal $\hat{\tau}$ suggests that it




Figure C.4: Histogram of paired differences between the RMS errors of the direct and the augmented methods: (a) difference between the RMSE of the direct linear approach and augmented linear approach; (b) difference between the RMSE of the direct quadratic approach and augmented quadratic approach.

may be preferable to use a relatively small time delay in the absence of precise information.

Assuming that the choice of lag length is robust to inter-patient and inter-fraction variations, we used a fixed lag length equivalent to 1.5 seconds delay for simplicity, and illustrate in Fig. C.7 the beam-wise 3D RMSE for patients 4,5 and 6, whose treatment extended over multiple days. The minimum RMS error for non-compensated treatment, which corresponds to a constant estimate at the retrospective mean value, is also shown for reference purposes. These results confirm that the augmented methods consistently exhibit lower error.

Adaptivity is most beneficial for irregular respiration traces. Our test data had relatively regular breathing patterns, so inclusion of adaptivity improved the estimation accuracy only slightly.

C.3 Conclusion and Future Work

We have proposed a method to map external surrogate signals to internal tumour positions. Breathing-phase dependent response patterns due to hysteresis are incorporated implicitly by using a simple state augmentation technique to capture system dynamics. We introduced a general class of correspondence models that are linear in model parameters, with linear and quadratic (in external surrogate) models as special cases. We described closed-form expressions for both the optimal model parameters and the corresponding error value. Based on the latter, we further investigated the proper choice of lag length in state augmentation, and argued its relative robustness. Test results on clinical data demonstrated reduced inference error over the direct linear and polynomial models.

The number of degrees of freedom in a correspondence model determines the trade-off between flexibility and robustness. We seek a model that is descriptive enough to fit the data without undesired sensitivity to observation noise, also known as “overfitting”. The proposed method may have more degrees of freedom than previous methods due to state augmentation. On the other hand, because it incorporates breathing-stage information implicitly, it can use all available internal-external correspondence pairs, without subdividing the training data as required for piecewise models [66, 67, 106]. In principle, using all the data may compensate for the possible increased sensitivity caused by the extra flexibility. The choice among different complexity levels in augmented models is still open. Both the number of augmentations and the model degree contribute to the overall complexity. Further studies should investigate methods for properly penalizing model complexity based on information criteria as explained in Section C.1.2.

Many research groups have observed phase shifts between external surrogate signal and internal tumour motions [17, 32]. Typically, this phase shift was to be avoided to obtain higher internal-external correlation. However, it is possible to compensate for consistent phase shift, to simplify and improve the correspondence map estimation. In particular, we can artificially synchronize the internal-external phase by shifting one

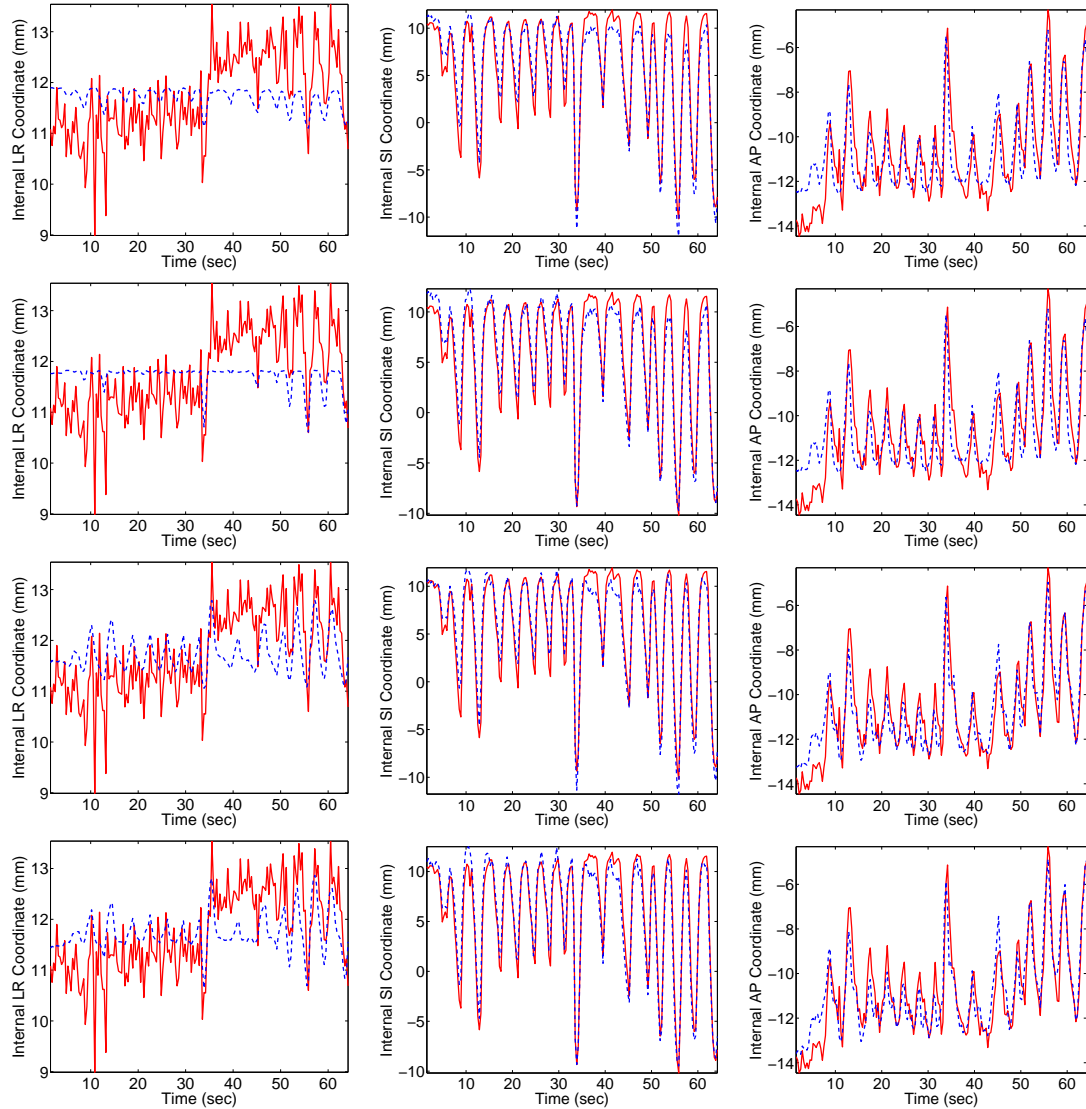


Figure C.5: Estimation performance comparison among different methods. Red-solid line depicts the internal tumour position obtained from fluoroscopic imaging, and dashed blue line provides estimated quantities from external surrogates. Each column represents one internal motion coordinate. Each row indicates the time series generated with one estimation method: (1st row) direct linear; (2nd row) direct polynomial; (3rd row) augmented linear; (bottom row) augmented polynomial.

of them according to a constant offset estimated from training data. We will further study phase-offset estimation and its use in external-internal inference in the future.

This work is a preliminary study to validate the existence of a reasonably simple correspondence map and the possibility to estimate it with high accuracy. In practice, internal-external pairs are obtained at a much slower rate. Correspondence maps must be extracted from sparse imaging data and applied to continuously obtained external surrogate signals to estimate the internal tumour locations. Our method can serve as a critical module in this overall framework, yet intensive simulations and validations are further required.

Even though our test data did not exhibit dramatic improvements when using adaptive model estimation, model updates in response to changes are necessary in general. Pursuing this direction requires more thorough analysis of breathing motion variations, change detection and model adaptive rate.




Figure C.6: Estimation error as a function of lag length for state augmentation: linear fit (solid line); quadratic fit (dashed line).




Figure C.7: Beam-wise 3D RMSE (mm) for patients 4-6: minimum non-surrogate (blue circle-dashed); linear inference (green square-dashed); polynomial inference (red star-dashed); augmented linear inference (cyan triangle-dashed); augmented polynomial (magenta diamond-dashed). Non-uniform tick locations along the x-axis indicate the number of beams applied each individual on treatment day.

BIBLIOGRAPHY

BIBLIOGRAPHY

- [1] S. Ahn, B. Yi, Y. Suh, J. Kim, S. Lee, S. Shin, and E. Choi. A feasibility study on the prediction of tumour location in the lung from skin motion. *Br. J. Radiol.*, 77:588–96, 2004.
- [2] H. Akaike. A new look at the statistical model identification. *IEEE Trans. Auto. Control*, 19(6):716–23, December 1974.
- [3] G. Aubert and P. Kornprobst. *Mathematical Problems in Image Processing*. Applied Mathematical Sciences. Springer, 2004.
- [4] A. Benveniste, M. Metivier, and P. Priouret. *Adaptive algorithms and stochastic approximation*. Springer-Verlag, 1990.
- [5] R. I. Berbeco, S. Nishioka, H. Shirato, G. T. Y. Chen, and S. B. Jiang. Residual motion of lung tumours in gated radiotherapy with external respiratory surrogates. *Phys. Med. Biol.*, 50(16):3655–68, August 2005.
- [6] D. P. Bertsekas. *Constrained optimization and Lagrange multiplier methods*. Academic-Press, New York, 1982.
- [7] M. J. Black and A. Rangarajan. On the unification of line processes, outlier rejection, and robust statistics with applications in early vision. *Intl. J. Comp. Vision*, 19(1):57–91, July 1996.
- [8] M. J. Black, G. Sapiro, D. H. Marimont, and D. Heeger. Robust anisotropic diffusion. *IEEE Trans. Im. Proc.*, 7(3):421–32, March 1998.
- [9] A. Blake and A. Zisserman. *Visual Reconstruction*. MIT Press, 1987.
- [10] P. Blomgren and T. F. Chan. Color TV: total variation methods for restoration of vector-valued images. *IEEE Trans. Img. Proc.*, 7(3):304–9, Mar 1998.
- [11] F. Bookstein. Fitting conic sections to scattered data. *Comp. Vision, Graphics and Image Processing*, 9:56–71, 1979.
- [12] T. Bortfeld, K. Jokivarsi, M. Goitein, J. Kung, and S_i Jiang. Effects of intra-fraction motion on IMRT dose delivery: statistical analysis and simulation. *Phys. Med. Biol.*, 47:2203–2220, 2002.
- [13] A. Chambolle. Image segmentation by variational methods: Mumford and Shah functional and the discrete approximations. *SIAM J. Appl. Math.*, 55(3):827–63, June 1995.
- [14] A. Chambolle and P. L. Lions. Image recovery via total variation minimization and related problems. *Numer. Math.*, 76(2):167–88, 1997.
- [15] T. F. Chan and J. Shen. *Image processing and analysis*. Cambridge Univ. Press, 2005.
- [16] E. W. Cheney. *Introduction to approximation theory*. American Mathematical Society, 1999.
- [17] P. M. Chi, P. Balter, D. Luo, R. Mohan, and T. Pan. Relation of external surface to internal tumor motion studies iwth cine CT. *Med. Phys.*, 33(9):3116–23, Sep. 2006.

- [18] C. S. Chui, E. Yorke, and L. Hong. The effects of intra-fraction organ motion on the delivery of intensity-modulated field with a multileaf collimator. *Med. Phys.*, 30:1736–46, 2003.
- [19] W. S. Cleveland. Robust locally weighted regression and smoothing scatterplots. *J. Amer. Stat. Asso.*, 74(368):829–36, 1979.
- [20] A. Collignon, F. Maes, D. Delaere, D. Vandermeulen, P. Suetens, and G. Marchal. Automated multimodality image registration using information theory. In Y. Bizais, C. Barillot, and R. D. Paola, editors, *Information Processing in Medical Im.*, pages 263–74. Kluwer, Dordrecht, 1995.
- [21] M. M. Coselmon, J. M. Balter, D. L. McShan, and M. L. Kessler. Mutual information based CT registration of the lung at exhale and inhale breathing states using thin-plate splines. *Med. Phys.*, 31(11):2942–8, November 2004.
- [22] A. D’Andrea, U. Mengalia, and R. Reggiannini. The modified Cramer-Rao bound and its application to synchronization problems. *IEEE Trans. Commun.*, 42(234):1391–99, Feb-Apr. 1994.
- [23] M. Droske and W. Ring. A Mumford-Shah level-set approach for geometric image registration. *SIAM J. Appl. Math.*, 66(6):2127–48, 2006.
- [24] M. Droske and M. Rumpf. Multiscale joint segmentation and registration of image morphology. *IEEE Trans. Patt. Anal. Mech. Intell.*, 29(12):2181–2194, Dec. 2007.
- [25] R. O. Duda, P. E. Hart, and D. G. Stork. *Pattern classification*. Wiley, New York, 2001.
- [26] T. Ellis, A. Abbood, and B. Brillault. Ellipse detection and matching with uncertainty. *Image and vision computing*, 10(2):271–6, 1992.
- [27] M. Engelsman, G. C. Sharp, T. Bortfeld, R. Onimaru, and H. Shirato. How much margin reduction is possible through gating or breath hold? *Phys. Med. Biol.*, 50(3):477–90, February 2005.
- [28] S. Esedoglu and S. J. Osher. Decomposition of images by the anisotropic Rudin-Osher-Fatemi model. *Comm. Pure. & Appl. Math.*, 57:1609–26, 2004.
- [29] H. G. Feichtinger, K. Gröchenig, and T. Strohmer. Efficient numerical methods in non-uniform sampling theory. *Numerische Mathematik*, 69(4):423–40, February 1995.
- [30] J. A. Fessler. Mean and variance of implicitly defined biased estimators (such as penalized maximum likelihood): Applications to tomography. *IEEE Trans. Im. Proc.*, 5(3):493–506, March 1996.
- [31] A. Fitzgibbon, M. Pilu, and R. B. Fisher. Direct least square fitting of ellipses. *IEEE Trans. Patt. Anal. Mach. Int.*, 21(5):476–80, May 1999.
- [32] E. C. Ford, G. S. Mageras, E. Yorke, and C. C. Ling. Respiration-correlated spiral ct: a method of measuring respiratory-induced anatomic motion for radiation treatment planning. *Med. Phys.*, 30:88–97, 2003.
- [33] C. Frohn-Schauf, S. Henn, and K. Witsch. Multigrid based total variation image registration. *Computing and Visualization in Science*, 11(2):101–13, March 2008.
- [34] D. Geman and G. Reynolds. Constrained restoration and the recovery of discontinuities. *IEEE Trans. Patt. Anal. Mach. Int.*, 14(3):367–83, March 1992.
- [35] S. Geman and D. Geman. Stochastic relaxation, Gibbs distributions, and Bayesian restoration of images. *IEEE Trans. Patt. Anal. Mach. Int.*, 6(6):721–41, November 1984.
- [36] R. George, P. Keall, V. Kini, S. Vedam, J. Siebers and Q. Wu, M. Lauterbach, D. Arthur, and R. Mohan. Quantifying the effect of intrafraction motion during breast IMRT planning and dose delivery. *Med. Phys.*, 30:552–62, 2003.

- [37] D. P. Gierga, G. T. Y. Chen, J. H. Kung, M. Betke, J. Lombardi, and C. G. Willett. Quantification of respiration-induced abdominal tumor motion and its impact on IMRT dose distributions. *Int. J. Radiat. Oncol. Biol. Phys.*, 58(5):1584–95, April 2004.
- [38] V. Girault and P. A. Raviart. *Finite element methods for navier-stokes equations*. springer, 1986.
- [39] G. H. Golub and C. F. Van Loan. *Matrix computations*. Johns Hopkins Univ. Press, 2 edition, 1989.
- [40] I. S. Gradshteyn and I. M. Ryzhik. *Tables of integrals, series and products*. Academic Press, 6th edition, 2000.
- [41] K. Gröchenig. A discrete theory of irregular sampling. *Linear Algebra and its Applications*, 193:129–50, November 1993.
- [42] S. Gupta and J. Prince. Stochastic models for div-curl optical flow methods. *Sig. Proc. Letters.*, 3(2):32–4, 1996.
- [43] R. P. Hamalainen and A. Kettunen. Stability of Fourier coefficients in relation to changes in respiratory air flow patterns. *Med. Eng. and Phys.*, 22(10):733–739, Dec 2000.
- [44] R. Haralick and L. Shapiro. *Computer and Robot Vision*. Addison-Wesley, Reading, MA, 1992.
- [45] A. O. Hero and J. A. Fessler. A recursive algorithm for computing Cramer-Rao-type bounds on estimator covariance. *IEEE Trans. Info. Theory*, 40(4):1205–10, July 1994.
- [46] M. Hirsch, S. Smale, and R. Devaney. *Differential equations, dynamical systems, and an introduction to Chaos*. Academic Press, 2003.
- [47] J. D. Hoisak, K. E. Sixel, R. Tirona, P. C. Cheung, and P. P. Pignol. Correlation of lung tumor motion with external surrogate indicator of respiration. *Int. J. Radiat. Oncol., Biol., Phys.*, 60(4):1298–306, Nov 2004.
- [48] B. Horn and B. G. Schunck. Determining optical flow. *Artif. Intell.*, 18(1-3):185–203, August 1981.
- [49] J. Hsieh. *Computed tomography: Principles, design, artifacts, and recent advances*. SPIE, Bellingham, 2003.
- [50] P. J. Huber. *Robust statistics*. Wiley, New York, 1981.
- [51] R. H. Huesman, G. J. Klein, J. A. Kimdon, C. Kuo, and S. Majumdar. Deformable registration of multimodal data including rigid structures. *IEEE Trans. Nuc. Sci.*, 50(3):389–92, June 2003.
- [52] G. D. Hugo, D. Yan, and J. Liang. Population and patient-specific target margins for 4D adaptive radiotherapy to account for intra- and inter-fraction variation in lung tumour position. *Phys. Med. Biol.*, 52(1):257–74, January 2007.
- [53] S. Jiang, C. Pope, K. Al Jarrah, J. Kung, T. Bortfeld, and G. Chen. An experimental investigation on intra-fractional organ motion effects in lung IMRT treatments. *Phys. Med. Biol.*, 48:1773–1784, 2003.
- [54] S. Jiang, G. Sharp, R. Berbeco, G. Chen, H. Mostafavi, and A. Jeung. Development of an integrated radiotherapy imaging system (IRIS). *Int. J. Radiat. Oncol., Biol., Phys.*, 60(1, Supp1):S611, 2004.
- [55] B. Karaçali and C. Davatzikos. Topology preservation and regularity in estimated deformation fields. In *Information Processing in Medical Im.*, pages 426–37, 2003.
- [56] B. Karacali and C. Davatzikos. Estimating topology preserving and smooth displacement fields. *IEEE Trans. Med. Imag.*, 23(7):868–80, July 2004.

- [57] P. J. Keall, G. S. Mageras, J. M. Balter, R. S. Emery, K. M. Forster, S. B. Jiang, J. M. Kapatoes, H. D. Kubo, D. A. Low, M. J. Murphy, B. R. Murray, C. R. Ramsey, M. B. van Herk, S. Vedam, J. W. Wong, and E. Yorke. The management of respiratory motion in radiation oncology. Technical report, AAPM Task Group 76, 2006.
- [58] G. J. Klein, B. W. Reutter, and R. H. Huesman. Four-dimensional affine registration models for respiratory-gated PET. *IEEE Trans. Nuc. Sci.*, 48(3):756–60, June 2001.
- [59] N. Koch, H. H. Liu, G. Starkschall, M. Jacobson, K. Forster, Z. Liao, R. Komaki, and C. W. Stevens. Evaluation of internal lung motion for respiratory-gated radiotherapy using mri: Part i-correlating internal lung motion with skin fiducial motion. *Int. J. Radiat. Oncol., Biol., Phys.*, 60(5):1459–1472, 2004.
- [60] H. D. Kubo and B. C. Hill. Respiration gated radiotherapy treatment: a technical study. *Phys. Med. Biol.*, 41:83–91, 1996.
- [61] J. Kybic, P. Thevenaz, A. Nirkko, and M. Unser. Unwarping of unidirectionally distorted EPI images. *IEEE Trans. Med. Imag.*, 19(2):80–93, February 2000.
- [62] J. Kybic and M. Unser. Fast parametric elastic image registration. *IEEE Trans. Im. Proc.*, 12(11):1427–42, November 2003.
- [63] R. P. Lippmann. An introduction to computing with neural nets. *IEEE ASSP Mag.*, 4(2):4–22, April 1987.
- [64] J. A. Little, D. L. G. Hill, and D. J. Hawkes. Deformations incorporating rigid structures. *Comp. Vision & Im. Understanding*, 66(2):223–32, May 1997.
- [65] L. Ljung and T. Soderström. *Theory and practice of recursive identification*. MIT Press, Cambridge, 1983.
- [66] D. A. Low, P. J. Parikh, W. Lu, J. F. Dempsey, S. H. Wahab, J. P. Hubenschmidt, M. M. Nystrom, M. Handoko, and J. D. Bradley. Novel breathing motion model for radiotherapy. *Int. J. Radiat. Oncol., Biol., Phys.*, 63(3):921–9, Nov. 2005.
- [67] W. Lu, P. J. Parikh, I. M. E. Naqa, M. M. Nystrom, J. P. Hubenschmidt, S. H. Wahab, S. Mutic, A. K. Singh, G. E. Christensen, J. D. Bradley, and D. A. Low. Quantitation of the reconstruction quality of a four-dimensional computed tomography process for lung cancer patients. *Med. Phys.*, 32(4):890–901, April 2005.
- [68] D. G. Luenberger. *Optimization by vector space methods*. Wiley, New York, 1969.
- [69] A. E. Lujan, J. M. Balter, and R. K. T. Haken. A method for incorporating organ motion due to breathing into 3D dose calculations in the liver: Sensitivity to variations in motion. *Med. Phys.*, 30(10):2643–9, October 2003.
- [70] A. E. Lujan, J. M. Balter, and R. K. Ten Haken. A method for incorporating organ motion due to breathing into 3D dose calculations in the liver: sensitivity to variations in motion. *Med. Phys.*, 30(10):2643–2649, Oct 2003.
- [71] A. E. Lujan, E. W. Larsen, J. M. Balter, and R. K. Ten Haken. A method for incorporating organ motion due to breathing into 3D dose calculations. *Med. Phys.*, 26(5):715–20, May 1999.
- [72] F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens. Multimodality image registration by maximization of mutual information. *IEEE Trans. Med. Imag.*, 16(2):187–98, April 1997.
- [73] G. S. Mageras, A. Pevsner, E. D. Yorke, K. E. Rosenzweig, E. C. Ford, A. Hertanto, S. M. Larson, D. M. Lovelock, Y. E. erdi, S. A. Nehmeh, J. L. Humm, and C. C. Ling. Measurement of lung tumor motion using respiration-correlated CT. *Int. J. Radiat. Oncol., Biol., Phys.*, 60:933–41, 2004.

- [74] J. B. A. Maintz and M. A. Viergever. A survey of medical image registration. *Med. Im. Anal.*, 2(1):1–36, March 1998.
- [75] D. Mattes, D. R. Haynor, H. Vesselle, T. K. Lewellen, and W. Eubank. PET-CT image registration in the chest using free-form deformations. *IEEE Trans. Med. Imag.*, 22(1):120–8, January 2003.
- [76] M. Moeneclaey. On the true and the modified Cramer-Rao bounds for the estimation of a scalar parameter in the presence of nuisance parameters. *IEEE Trans. Comm.*, 46(11):1536–44, November 1998.
- [77] D. Mumford and J. Shah. Optimal approximations by piecewise smooth functions and associated variational problems. *Comm. Pure Appl. Math.*, 42(5):577–685, May 1989.
- [78] K. Murphy. Kalman filter toolbox for matlab. <http://www.ai.mit.edu/~murphyk/Software/Kalman/kalman.html>.
- [79] M. Murphy and S. Dieterich. Comparative performance of linear and nonlinear neural networks to predict irregular breathing. *Phys. Med. Biol.*, 51:5903–14, Oct. 2006.
- [80] M. J. Murphy. Tracking moving organs in real time. *Semin. Radiat. Oncol.*, 14(1):91–100, 2004.
- [81] M. J. Murphy, J. Jalden, and M. Isaksson. Adaptive filtering to predict lung tumor breathing motion during image-guided radiation therapy. *16th International Congress on Computer-assisted Radiology and Surgery (CARS)*, 2002.
- [82] K. G. Murty and S. N. Kabadi. Some NP-complete problems in quadratic and nonlinear programming. *Math. Programming*, 39:117–29, 1987.
- [83] I. Nabney and C. Bishop. Netlab neural network software. <http://www.ncrg.aston.ac.uk/netlab/>, 2003.
- [84] T. Neicu, H. Shirato, Y. Seppenwoolde, and S. B. Jiang. Synchronized moving aperture radiation therapy (smart): average tumour trajectory for lung patients. *Phys. Med. Biol.*, 48(5):587–598, 2003.
- [85] V. Noblet, C. Heinrich, F. Heitz, and J-P. Armspach. 3-D deformable image registration: a topology preservation scheme based on hierarchical deformation models and interval analysis optimization. *IEEE Trans. Im. Proc.*, 14(5):553–66, May 2005.
- [86] C. Ozhasoglu and M. J. Murphy. Issues in respiratory motion compensation during external beam radiotherapy. *Int. J. Radiation Oncology Biol. Phys.*, 52(5):1389–1399, 2002.
- [87] R. Peslin, C. Gallina, C. Saunier, and C. Duvivier. Fourier analysis versus multiple linear regression to analyse pressure-flow data during artificial ventilation. *Eur. Respir J.*, 7:2241–45, 1994.
- [88] J. P. W. Pluim, J. B. A. Maintz, and M. A. Viergever. Mutual-information-based registration of medical images: a survey. *IEEE Trans. Med. Imag.*, 22(8):986–1004, August 2003.
- [89] J. Porrill. Fitting ellipses and predicting confidence envelopes using a bias corrected kalman filter. *Image and Vision Computing*, 8(1):37–41, 1990.
- [90] J. C. Preisig. Copositivity and the minimization of quadratic functions with nonnegativity and quadratic equality constraints. *SIAM J. Cont. Opt.*, 34(4):1135–50, July 1996.
- [91] C. R. Rao. *Linear statistical inference and its applications*. Wiley, New York, 1973.
- [92] Y. N. Rao, J. C. Principe, and T. F. Wong. Fast RLS-like algorithm for generalized eigendecomposition and its applications. *J. VLSI Sig. Proc.*, 37:333–344, 2004.
- [93] P. A. Regalia. *Adaptive IIR Filtering in signal processing and control*. Marcel Dekker, 1995.
- [94] C. Riviere, A. Thakral, I. I. Iordachita, G. Mitroi, and D. Stoianovici. Predicting respiratory motion for active canceling during percutaneous needle insertion. In *Proc. 23rd Annual Intl. Conf. IEEE Enginnering in Medicine and Biology Society*, pages 3477–80, Oct. 2001.

- [95] D. Robinson and P. Milanfar. Fundamental performance limits in image registration. *IEEE Trans. Im. Proc.*, 13(9):1185–99, September 2004.
- [96] T. Rohlfing, C. R. Maurer, D. A. Bluemke, and M. A. Jacobs. Volume-preserving nonrigid registration of MR breast images using free-form deformation with an incompressibility constraint. *IEEE Trans. Med. Imag.*, 22(6):730–741, June 2003.
- [97] D. Ruan and J. A. Fessler. Adaptive ellipse tracking and a convergence proof. Technical Report 382, Comm. and Sign. Proc. Lab., Dept. of EECS, Univ. of Michigan, Ann Arbor, MI, 48109-2122, May 2007.
- [98] D. Ruan, J. A. Fessler, and J. Balter. Real-time prediction of respiratory motion based on local regression methods. *Phys. Med. Biol.*, 52:7135–52, 2007.
- [99] D. Ruan, J. A. Fessler, and J. M. Balter. Mean position tracking of respiratory motion. *Med. Phys.*, 35(2), 2008.
- [100] D. Ruan, J. A. Fessler, J. M. Balter, R. I. Berbeco, S. Nishioka, and H. Shirato. Inference of hysteretic respiratory tumour motion from external surrogates: A state augmentation approach. *Phys. Med. Biol.*, 53(11):2923–36, June 2008.
- [101] D. Ruan, J. A. Fessler, J. M. Balter, and J. Sonke. Exploring breathing pattern irregularity with projection-based method. *Med. Phys.*, 33(7):2491–2499, Jul. 2006.
- [102] D. Ruan, J. A. Fessler, M. Roberson, J. Balter, and M. Kessler. Nonrigid registration using regularization that accommodates local tissue rigidity. In *Proc. SPIE 6144, Medical Imaging 2006: Image Proc.*, page 614412, 2006.
- [103] D. Ruan, D. A. Casta non, G. Sharp, and S. Jiang. Real-time tumor tracking with interactive multiple model filter. CenSSIS '03.
- [104] A. Schweikard, G. Glosser, M. Bodduluri, and M. J. Murphy. Robotic motion compensation for respiratory movement during radiosurgery. *Computer Aided Surgery*, 5(4):263–277, 2000.
- [105] A. Schweikard, H. Shiomi, and J. Adler. Respiration tracking in radiosurgery. *Med. Phys.*, 31(10):2738–41, October 2004.
- [106] Y. Seppenwoolde, R. I. Berbeco, S. Nishioka, H. Shirato, and B. Heijmen. Accuracy of tumor motion compensation algorithm from a robotic respiratory tracking system: a simulation study. *Med. Phys.*, 34(7):2774–84, Jul. 2007.
- [107] Y. Seppenwoolde, H. Shirato, K. Kitamura, S. Shimizu, M. . Herk, J. V. Lebesque, and K. Miyasaka. Precise and real-time measurement of 3D tumor motion in lung due to breathing and heartbeat, measured during radiotherapy. *Int. J. Radiat. Oncol. Biol. Phys.*, 53(4):822–34, July 2002.
- [108] R. J. Serfling. *Approximation theorems of mathematical statistics*. Wiley, New York, 1980.
- [109] G. C. Sharp, S. B. Jiang, S. Shimizu, and H. Shirato. Prediction of respiratory tumour motion for real-time image-guided radiotherapy. *Phys. Med. Biol.*, 49(3):425–40, February 2004.
- [110] H. Shirato et al. Physical aspects of a real-time tumor-tracking system for gated radiotherapy. *Int. J. Radiat. Oncol. Biol. Phys.*, 48(4):1187–95, November 2000.
- [111] H. Shirato, T. Harada, T. Harabayashi, K. Hida, H. Endo, K. Kitamura, R. Onimaru, K. Yamazaki, N. Kurauchi, T. Shimizu, N. Shinohara, M. Matsushita, H. Dosaka-Akita, and K. Miyasaka. Feasibility of insertion/implantation of 2.0 mm diameter gold internal fiducial markers for precise setup and real-time tumor tracking in radiotherapy. *Int. J. Radiat. Oncol., Biol., Phys.*, 56(1):240–247, 2003.
- [112] M. Staring, S. Klein, and J. P. Pluim. Nonrigid registration with adaptive, content-based filtering of the deformation field. In *Proc. SPIE 5747, Medical Imaging 2005: Image Proc.*, pages 212–21, 2005.

- [113] M. Staring, S. Klein, and J. P. Pluim. Nonrigid registration using a rigidity constraint. In *Proc. SPIE 6144, Medical Imaging 2006: Image Proc.*, 2006.
- [114] M. Staring, S. Klein, and J. P. W. Pluim. A rigidity penalty term for nonrigid registration. *Med. Phys.*, 34(11):4098–108, November 2007.
- [115] P. Thevenaz, U. E. Ruttimann, and M. Unser. A pyramid approach to subpixel registration based on intensity. *IEEE Trans. Im. Proc.*, 7(1):27–41, January 1998.
- [116] P. Thevenaz and M. Unser. Optimization of mutual information for multiresolution image registration. *IEEE Trans. Im. Proc.*, 9(12):2083–99, December 2000.
- [117] Y. Tsunashima, T. Sakae, Y. Shioyama, K. Kagei, T. Terunuma, A. Nohtomi, and Y. Akine. Correlation between the respiratory waveform measured using a respiratory sensor and 3D tumor position in gated radiotherapy. *Int. J. Radiat. Oncol., Biol. Phys.*, 60:951–8, 2004.
- [118] M. Unser. Splines: A perfect fit for signal and image processing. *spmag*, 16(6):22–38, November 1999.
- [119] M. Unser, A. Aldroubi, and M. Eden. Fast B-spline transforms for continuous image representation and interpolation. *IEEE Trans. Patt. Anal. Mach. Int.*, 13(3):277–85, March 1991.
- [120] M. Unser, A. Aldroubi, and M. Eden. B-spline signal processing: Part I—theory. *IEEE Trans. Sig. Proc.*, 41(2):821–33, February 1993.
- [121] M. Unser, A. Aldroubi, and M. Eden. B-spline signal processing: Part II—efficient design and applications. *IEEE Trans. Sig. Proc.*, 41(2):834–48, February 1993.
- [122] M. Unser, A. Aldroubi, and M. Eden. The L_2 polynomial spline pyramid. *IEEE Trans. Patt. Anal. Mach. Int.*, 15(4):364–379, April 1993.
- [123] H. L. Van Trees. *Detection, estimation, and modulation theory*. Wiley, New York, 1968.
- [124] S. Vedam, V. Kini, P. Keall, V. Ramakrishnan, H. Mostafavi, and R. Mohan. Quantifying the predictability of diaphragm motion during respiration with a noninvasive external marker. *Med. Phys.*, 30(4):505–13, April 2003.
- [125] S. S. Vedam, P. J. Keall, A. Docef, D. A. Todor, V. R. Kini, and R. Mohan. Predicting respiratory motion for four-dimensional radiotherapy. *Med. Phys.*, 31(8):2274–2283, Aug. 2004.
- [126] P. Viola. *Alignment by maximization of mutual information*. PhD thesis, Dept. of EECS, MIT, Cambridge, MA, 1995.
- [127] O. L. Wade. Movement of the thoracic cage and diaphragm in respiration. *J. Physiol.*, pages 193–212, 1954.
- [128] K. Wang, Y. He, and H. Qin. Incorporating rigid structures in non-rigid registration using triangular b-splines. *Variational, Geometric and Level Set Methods in Computer Vision (VLSM)*, 3752:235–246, Oct 2005.
- [129] J. Weickert. On discontinuity-preserving optic flow. In S. Orphanoudakis, P. Trahanias, J. Crowley, and N. Katevas, editors, *Proc. Computer Vision and Mobile Robotics Workshop*, pages 115–22, Santorini, Sep. 1998.
- [130] W. M. Wells, P. Viola, H. Atsumi, S. Nakajima, and R. Kikinis. Multi-modal volume registration by maximization of mutual information. *Med. Im. Anal.*, 1(1):35–51, March 1996.
- [131] J. H. Wilkinson. *The algebraic eigenvalue problem*. Clarendon Press, Oxford, England, 1965.

- [132] J. W. H. Wolthaus, C. Schneider, J.-J. Sonke, M. van Herk, J. S. A. Belderbos, M. M. G. Rossi, J. V. Lebesque, and E. M. F. Damen. Mid-ventilation CT scan construction from four-dimensional respiration-correlated CT scans for radiotherapy planning of lung cancer patients. *Int. J. Radiat. Oncol. Biol. Phys.*, 65(5):1560–71, August 2006.
- [133] H. Wu, G. C. Sharp, B. Salzberg, D. Kaeli, H. Shirato, and S. B. Jiang. A finite state model for respiratory motion analysis in image guided radiation therapy. *Phys. Med. Biol.*, pages 5357–5372, Nov. 2004.
- [134] A. Yezzi, L. Zilei, and T. Kapur. A variational framework for integrating segmentation and registration through active contours. *Med. Im. Anal.*, 7(2):171–85, June 2003.
- [135] J. Yuan, C. Schnörr, and E. Mémin. Discrete orthogonal decomposition and variational fluid flow estimation. *J. Math. Imaging Vis.*, 28(1):67–80, 2007.
- [136] J. Yuan, C. Schnörr, and G. Steidl. Convex hodge decomposition and regularization of image flows. Preprint, University of Heidelberg, 2008.
- [137] J. Yuan, C. Schnörr, and G. Steidl. Simultaneous optical flow estimation and decomposition. *SIAM J. Sci. Comput.*, 29(6):2283–304, 2007.
- [138] B. Zitová and J. Flusser. Image registration methods: a survey. *Im. and Vision Computing*, 21(11):977–1000, October 2003.