

© 2011

Jonathan Chappelow

ALL RIGHTS RESERVED

MULTIMODAL IMAGE REGISTRATION USING
MULTIVARIATE INFORMATION THEORETIC
SIMILARITY MEASURES: APPLICATIONS TO
PROSTATE CANCER DIAGNOSIS AND TARGETED
TREATMENT

BY JONATHAN CHAPPELOW

A dissertation submitted to the
Graduate School—New Brunswick
Rutgers, The State University of New Jersey
and
The Graduate School of Biomedical Sciences
University of Medicine and Dentistry of New Jersey
in partial fulfillment of the requirements for the
Degree of Doctor of Philosophy
Graduate Program in Biomedical Engineering

Written under the direction of

Anant Madabhushi

and approved by

New Brunswick, New Jersey

May, 2011

ABSTRACT OF THE DISSERTATION

Multimodal Image Registration using Multivariate Information Theoretic Similarity Measures: Applications to Prostate Cancer Diagnosis and Targeted Treatment

by Jonathan Chappelow

Dissertation Director: Anant Madabhushi

Multimodal and multiprotocol image registration refers to the process of alignment of two or more images obtained from different imaging modalities (e.g. digitized histology and MRI) and protocols (e.g. T2-w and PD-w MRI). Registration is a critical component in medical applications including image guided surgery, image fusion for cancer diagnosis and treatment planning, and automated tissue annotation. However, registration is often complicated on account of differences in both the image intensities and the shape of the underlying anatomy. For example, non-linear differences in the overall shape of the prostate between *in vivo* MRI and *ex vivo* whole mount histology (WMH) often exist as a result of the presence of an endorectal coil during pre-operative MR imaging and deformations to the specimen during slide preparation.

To overcome these challenges, we present new registration techniques termed Combined Feature Ensemble Mutual Information (COFEMI) and Collection of Image-derived Non-linear Attributes for Registration Using Splines (COLLINARUS). The goal COFEMI is to provide a similarity measure that is driven by unique low level textural features, for registration that is more robust to intensity artifacts and modality differences than measures restricted to intensities alone. COLLINARUS offers the robustness of COFEMI

to artifacts and modality differences, while allowing fully automated non-linear image warping at multiple scales via a hierarchical B-spline mesh grid. In addition, since routine clinical imaging procedures often involve the acquisition of multiple imaging protocols, we present a technique termed Multi-attribute Combined Mutual Information (MACAMI) to leverage the availability of multiple image sets to improve registration.

We apply our registration techniques to a unique clinical dataset comprising 150 sets of *in vivo* MRI and post-operative WMH images from 25 patient studies in order to retrospectively establish the spatial extent of prostate cancer (CaP) on structural (T2-w) and functional (DCE) *in vivo* MRI. Accurate mapping of CaP on MRI is used to facilitate the development and evaluation of a system for computer-assisted detection (CAD) of CaP on multiprotocol MRI. We also demonstrate our registration and CAD algorithms in developing radiation therapy treatment plans that provide dose escalation to CaP by elastically registering diagnostic MRI with planning CT.

Preface

This dissertation represents the collective published and unpublished works of the author. Chapters 2-7 are primarily composed from the content of the technical reports, conference papers [1–5], and peer-reviewed journal articles [6, 7], written by the author of this dissertation over the course of his thesis work.

Acknowledgements

Thanks to all the members of the Laboratory for Computational Imaging and Bioinformatics (LCIB) in the Biomedical Engineering Department at Rutgers University. The spirit of collaboration among the members of LCIB was essential in the execution of this work. Special thanks go to my advisor, Professor Anant Madabhushi, whose time and patience for his students are seemingly illimitable.

Without the considerable contributions of our clinical collaborators, in particular Dr. B. Nicolas Bloch of Boston University, Dr. John Tomaszewski and Dr. Mark Rosen of the University of Pennsylvania, none of these studies would have been possible.

Finally, sincere thanks to each of my committee members for their valuable input and guidance in this research.

Dedication

To Natalia, for seeing the best in me.

I choose to live and to
Grow, take and give and to
Move, learn and love and to
Cry, kill and die and to
Be paranoid and to
Lie, hate and fear and to
Do what it takes to move through.

I choose to live and to
Lie, kill and give and to
Die, learn and love and to
Do what it takes to step through.

See my shadow changing,
Stretching up and over me.
Soften this old armor.
Hoping I can clear the way
By stepping through my shadow,
Coming out the other side.
Step into the shadow.
Forty six and two are just ahead of me.

from “Forty-six & 2” by Tool

Table of Contents

Abstract	ii
Preface	iv
Acknowledgements	v
Dedication	vi
List of Tables	xiv
List of Figures	xviii
1. Introduction	1
1.1. Overview	1
1.2. Computer-Assisted Detection of Prostate Cancer	4
1.3. Previous Work	8
1.4. Overview of Contributions	10
1.5. Notation	12
2. Registration of <i>ex vivo</i> Whole-Mount Histology and <i>ex vivo</i> MRI of Prostate Specimens	13
2.1. Introduction to Image Registration	13
2.1.1. Image Similarity-based Multimodal Image Registration Background	13
2.1.2. Motivation for Multiple Image Features in Image Registration . .	15
2.1.3. Combined Feature Ensemble Mutual Information	16
2.1.4. Quantitative Evaluation of Registration Tasks	17
2.2. Theory of Combined Mutual Information	19
2.2.1. Mutual Information	19
2.2.2. Generalized Mutual Information	20

2.2.3.	Combined Mutual Information (CMI)	21
2.3.	Combined Feature Ensemble Mutual Information (COFEMI)	23
2.3.1.	Brief Overview of COFEMI	23
2.3.2.	Feature Extraction	25
	Gradient Features	25
	Non-steerable Gradient Features	25
	Steerable Gradient Features	26
	First Order Statistical Features	27
	Second Order Statistical Features	27
2.3.3.	Feature Ensemble Selection	28
2.3.4.	Registration by Maximization of CMI	29
2.4.	Quantitative Evaluation Methods	29
2.4.1.	Evaluation using Anatomical Region Labels	29
2.4.2.	Evaluation via Target Registration Error	30
2.4.3.	Quantitative Evaluation via Skerl Evaluation Protocol	31
2.5.	Results of Automated Affine Registration	33
2.5.1.	Real and Synthetic Image Data	33
2.5.2.	Experiment 1: Synthetic Multiprotocol BrainWeb Data	35
	Introducing Different Levels of Bias Field and Noise	35
	Experiment 1a: TRE and Region Overlap	35
	Qualitative Results	36
	Quantitative Results	37
	Experiment 1b: Skerl Evaluation Protocol	39
	Properties Independent of Trace Radius	42
	Number of Minima versus Trace Radius	44
	Risk of Non-convergence versus Trace Radius	44
2.5.3.	Experiment 2: Multimodal Prostate Registration	46
	Experiment 2 Overview	46
	Qualitative Evaluation	47

Quantitative Evaluation	49
2.6. Concluding Remarks	50
3. Improving Supervised CAD on <i>ex vivo</i> MRI via Registration	53
3.1. Importance of CAD Training Data Accuracy	53
3.2. Overview of Experimental Design	56
3.3. Overview of Registration Methods	58
3.3.1. Notation	58
3.3.2. Manual Mapping of Cancer from Histology to MRI	59
3.3.3. Overview of Registration Methods	59
Maximization of Mutual Information (MI)	60
Combined Feature Ensemble Mutual Information	61
Affine Registration via MI and COFEMI	62
Elastic Registration via MI, COFEMI, and TPS	62
3.4. Supervised Computer-Aided Classification	63
3.4.1. Feature Extraction	63
Gradient Features	64
Statistical Features	64
3.4.2. Bayesian Classifier and Independent Component Analysis	65
3.5. Application to Prostate CAD	67
3.5.1. Supervised Classifier Training	67
3.5.2. Evaluation of Supervised CAD Classifier	67
3.6. Results of Semi-automated Elastic Registration and CAD Evaluation . .	68
3.6.1. Qualitative Evaluation of Registration Schemes	68
3.6.2. Quantitative Results	70
3.7. Concluding Remarks	70
4. Registration of <i>ex vivo</i> Whole-Mount Histology and <i>in vivo</i> MRI using Multiple Texture Features	72
4.1. Challenges of Automated Elastic Multimodal Image Registration	72

4.2. Collection of Image-derived Non-linear Attributes for Registration Using Splines (COLLINARUS)	77
4.2.1. Overview	77
4.2.2. Notation	78
4.2.3. General Registration Framework	78
4.2.4. COLLINARUS Non-Rigid Registration	79
4.2.5. Registration Evaluation	81
4.3. Results of Automated Elastic Image Registration	81
4.3.1. Data Sets	81
4.3.2. Synthetic Brain Data	82
4.3.3. Clinical Multi-Modal Prostate Data	84
Prostate Registration Task	84
Three Time Point (3TP) DCE Cancer Maps	85
Prostate Registration Results	86
4.4. Concluding Remarks	88
5. Registration of <i>ex vivo</i> Whole-Mount Histology and <i>in vivo</i> MRI using Multiparametric Images	90
5.1. Background and Motivation	90
5.2. Previous Work	93
5.3. Novel Contributions and Significance	97
5.4. Methods	100
5.4.1. Theory on Mutual Information and MACMI	100
Mutual Information Between Scalar-Valued Images	100
Mutual Information Between High-dimensional (Multi-Attribute) Images	100
5.4.2. Framework for Registration of Multiple Images using MACMI . .	101
Algorithm	101
Instance of MACMI for Registration of 4 Images	104

5.4.3.	Requirements for Implementation of MACMI	105
5.5.	Experimental Design	105
5.5.1.	Synthetic Multiprotocol Brain MRI	106
	Data Description	106
	Registration Experiment	106
	Registration Evaluation	107
5.5.2.	Clinical Multimodal Prostate MRI and Histology	107
	Data Description	107
	Registration Experiments	109
	Registration Evaluation	110
5.6.	Results and Discussion	110
5.6.1.	Synthetic Brain Registration	110
5.6.2.	Clinical Prostate Registration	111
	Mapping CaP extent from WMH onto <i>in vivo</i> T2-w and T1-w MRI	111
	Mapping CaP Extent from WMH onto ADC, T2-w and T1-w MRI	114
5.7.	Concluding Remarks	115
6. Application of Elastic Image Registration and CAD to Improving		
IMRT		118
6.1.	Introduction to Radiation Therapy Planning	118
6.2.	Computer-Assisted Targeted Therapy (CATT)	122
6.2.1.	Data Description and Preprocessing	122
6.2.2.	Localization of Cancer on Diagnostic MRI	123
	Texture feature extraction	123
	First order statistical features	123
	Gradient features	124
	Second order statistical features	124
	Non-linear dimensionality reduction	125
	Replicated k-means clustering	126

6.2.3.	Elastic Registration of Diagnostic MRI, Planning MRI, and CT .	127
6.2.4.	Mapping of Tumor Location from Diagnostic MRI onto CT . . .	129
6.2.5.	Dose Plan Generation	129
6.3.	Results of CATT	130
6.3.1.	Unsupervised CAD Classifier to Identify Tumor Labels on Diagnostic MRI	130
6.3.2.	Multiprotocol and Multimodal Image Registration	131
6.3.3.	Tumor Mapping onto CT	132
6.3.4.	Dose Planning	134
6.4.	Concluding Remarks	134

7. HistoStitcher: An Interactive Program for Accurate and Rapid Reconstruction of Digitized Whole Histological Sections from Tissue Fragments 135

7.1.	Need for Efficient Digital Reconstruction of Whole Histological Images .	135
7.2.	Challenges in Digital Reconstruction of High Resolution Whole Histology	138
7.3.	Features and Operation of HistoStitcher	140
7.3.1.	Software Overview and Workflow	140
7.3.2.	Notation	142
7.3.3.	Determining the Optimal Transformation from Control Point Pairs	143
	Unconstrained Affine: Anisotropic Scale Allowing Reflection . . .	144
	Isotropic Scaling with No Reflection	146
	Isotropic Scaling allowing Reflection	147
	Rotation and Translation: No Scaling	149
	Enforcing Reflection	149
7.3.4.	Transformation of Moving Image	150
7.3.5.	Image Reconstruction	150
7.3.6.	Considerations for High Resolution Images	151
7.3.7.	Reference Stitching: Manual Alignment of Fragments in Photoshop	151

7.4. Criteria for Stitching Quality and Evaluation by Multiple Expert Consensus	152
7.5. Results of HistoStitcher	153
7.5.1. Experiment 1: Reconstructing High Resolution Whole Prostate Sections from Quadrants	153
7.5.2. Experiment 2: Comparison with Manual Stitching of High Reso- lution Sections	156
7.5.3. Experiment 3: Evaluation of Stitching Quality via Multiple Experts	159
7.6. Concluding Remarks	161
8. Conclusions	162
Vita	164
References	165

List of Tables

1.1.	Common notation and symbols used throughout this document.	12
2.1.	List of notation and symbols specific to Chapter 2.	24
2.2.	Description of 311 image derived features used by COFEMI and corresponding parameters for each texture feature class.	26
2.3.	Description of properties $\phi^{S1}-\phi^{S5}$ of Ψ calculated by the Skerl protocol [8] for evaluation in the context of rigid registration. The reader is referred to [8] for additional details on the Skerl evaluation protocol. . .	33
2.4.	Description of image data sets used in this study for evaluation of COFEMI.	34
2.5.	p -values for paired t -tests comparing ϕ^{TRE} values for MI and COFEMI for 30 different combinations of simulated noise (σ) and field inhomogeneity (b). These values correspond to the ϕ^{TRE} differences shown in Figure 2.7(c) averaged across 30 pairs of $\mathcal{T}1$ and $\mathcal{P}_{b,\sigma}^r$ for each b, σ	39
2.6.	p -values for paired t -tests comparing ϕ^{TRE} values for CMI^{grad} and COFEMI for 30 different combinations of simulated noise (σ) and field inhomogeneity (b). These values correspond to the ϕ^{TRE} differences shown in Figure 2.7(e) averaged across 30 pairs of $\mathcal{T}1$ and $\mathcal{P}_{b,\sigma}^r$ for each b, σ	40
2.7.	p -values for paired t -tests comparing ϕ^{ROR} values for MI and COFEMI for 30 different combinations of simulated noise (σ) and field inhomogeneity (b). These values correspond to the ϕ^{ROR} differences shown in Figure 2.7(d) averaged across 30 pairs of $\mathcal{T}1$ and $\mathcal{P}_{b,\sigma}^r$ for each b, σ	40
2.8.	p -values for paired t -tests comparing ϕ^{ROR} values for CMI^{grad} and COFEMI for 30 different combinations of simulated noise (σ) and field inhomogeneity (b). These values correspond to the ϕ^{ROR} differences shown in Figure 2.7(f) averaged across 30 pairs of $\mathcal{T}1$ and $\mathcal{P}_{b,\sigma}^r$ for each b, σ	40

2.9. Similarity measures Ψ compared to COFEMI by the Skerl evaluation protocol, in terms of two generic images \mathcal{A} and \mathcal{B} with respective intensities $f^{\mathcal{A}}(C)$ and $f^{\mathcal{B}}(c)$ at coordinate c , and average intensities $\bar{f}^{\mathcal{A}} = \frac{1}{ C } \sum_{c \in C} f^{\mathcal{A}}(c)$ and $\bar{f}^{\mathcal{B}} = \frac{1}{ C } \sum_{c \in C} f^{\mathcal{B}}(c)$. Note that p_{MI} represents a prior gray level distribution determined from two images that are known to be in correct alignment.	41
2.10. Φ^{ROR} , Φ^{HD} , and Φ^{MAD} values, averaged over 26 pairs of \mathcal{M} and \mathcal{H} , and corresponding to registration via MI, CMI^{grad} , and COFEMI, are given for prostate regions $G(\mathcal{M}^r)$ and $G(\mathcal{H})$ and boundaries $G^b(\mathcal{H})$ and $G^b(\mathcal{M}^r)$	49
2.11. p -values for t -tests comparing ϕ^{ROR} , ϕ^{DH} , and ϕ^{MAD} values corresponding to the average values presented in Table 2.10, indicating consistently higher overlap and lower boundary distances obtained from COFEMI versus MI, and COFEMI versus CMI^{grad}	49
2.12. Approximate computation times (in seconds) for the different modules involved in the operation of the COFEMI registration scheme. Note that calculation of the full, comprehensive feature set and selection of the optimal ensemble are performed only once initially (offline) for a given modality pair. To register successive test image pairs (online), it is only required to calculate the features comprising the selected ensemble. Computation times were obtained for images of size 256×256 running on an Intel Core 2 Duo 2.4GHz PC running 64-bit MATLAB [®]	51
3.1. List of notation and symbols specific to Chapter 3.	59

4.1.	Comparison of non-rigid registration accuracy for COLLINARUS and MI-MLS alignment of $n = 11$ pairs of synthetic PD MRI and T1-w MRI brain images. Error of recovered deformation field in terms of mean absolute difference (F_{mad}) and root mean squares (F_{rms}) measures the deviation of the registration-derived deformation field from the known field. Units of F_{mad} and F_{rms} are mm. Euclidean distance between the original undeformed PD MRI and recovered PD MRI obtained by non-rigid registration (D_{L2}), measures the dissimilarity between the registration result and the ideal result. The mutual information and entropy correlation coefficient between the recovered and original PD MRI sections (S_{MI} and S_{ECC}) indicate how well the recovered image resembles the original, ideal result. Each of F_{mad} , F_{rms} , D_{L2} , S_{MI} and S_{ECC} indicate that COLLINARUS more accurately recovered the original undeformed PD MRI compared with MI-MLS.	84
5.1.	Summary of the synthetic and clinical datasets registered by MACMI. Pixels are square.	106
5.2.	Comparison of elastic registration accuracy for MACMI and pairwise MI alignment of $n = 20$ pairs of synthetic PD MRI with co-registered T1-w and T2-w MRI brain images. The measures illustrated below correspond to (i) error of recovered deformation field (in mm) in terms of F_{mad} and F_{rms} , and (ii) distance (D_{L2}) between the undeformed and recovered PD MRI. MACMI results are significantly more accurate compared to either PW approach (p -values for both tests shown).	111
5.3.	Comparison of elastic registration accuracy for MACMI, PW-T2, and PW-DCE for $n = 25$ patient studies. The measures illustrated below correspond to mean similarity in terms of the total MI of all registered WMH slices $\tilde{\mathcal{H}}$, obtained by MACMI, PW-T2, or PW-DCE (columns of table), with either T2-w MRI or DCE MRI (rows).	112
6.1.	Summary of the prostate image data sets acquired for each of 10 patients.	122
7.1.	List of notation and symbols specific to Appendix 7.	143

7.2.	Comparison of HistoStitcher [©] to manipulation via Photoshop for the task of reassembling a whole prostate histological section from four separate slide images of histology quadrants. HistoStitcher [©] is demonstrated to be fast, memory efficient, and capable of producing better quality stitches of very high resolution images.	158
7.3.	Average and standard deviation of scores for 6 sections reconstructed using both Photoshop and HistoStitcher [©]	158

List of Figures

1.1.	(a) <i>ex vivo</i> prostate MRI, (b) tumor ground truth from histology, (c) CAD result, (d) superposition of 3 human expert segmentations of CAP on (a). Note that (i) CAD performance is comparable to human experts and (ii) the low levels of inter-observer agreement in (d). The bright areas in (d) are the only areas where all expert segmentations agree. . . .	6
1.2.	(a) Manually determined ground truth from histology. (b) Representation of tumor texture features in reduced dimensional space. Cluster outliers shown in red on (b) represent incorrectly classified tissue also shown in red on (c).	7
2.1.	Overview of the operation of COFEMI. Automatic selection of an optimal ensemble of textural features is central to the COFEMI method for multimodal registration.	16
2.2.	Venn diagrams with the shaded regions representing the information theoretic quantities (a) MI (I_2) of 2 random variables or images \mathcal{A}_1 and \mathcal{A}_2 , (b) higher order MI (I_3) of 3 random variables \mathcal{A}_1 , \mathcal{A}_2 , and \mathcal{A}_3 , and (c) CMI (I_2^*) between \mathcal{A}_1 and an ensemble of two simultaneously observed random variables, $\mathcal{A}_2\mathcal{A}_3$. (d) With the addition of a fourth random variable \mathcal{A}_4 , the fourth order MI (I_4) decreases from the quantity in (b). (e) CMI with the addition of a fourth random variable increases as the knowledge of \mathcal{A}_4 further reduces the uncertainty about \mathcal{A}_1	20

2.3.	(a) A plot of CMI of intensity image \mathcal{A}_1 with 16 $(n - 1)$ dimensional ensembles comprising intensity image \mathcal{A}_2 and $(n - 2)$ corresponding textural feature images, demonstrates that CMI is bounded by $S(\mathcal{A}_1)$ (horizontal line), the information content of \mathcal{A}_1 . (b) Plot of the sum of pairwise MI (SMI) of \mathcal{A}_1 with each of $\mathcal{A}_2, \dots, \mathcal{A}_n$ aggregates redundant information about \mathcal{A}_1 and is unbounded. (c) Plot of GMI of $\mathcal{A}_1, \dots, \mathcal{A}_n$ illustrates information shared by all images becomes less as additional features are considered. The upper bound on the value of CMI in (a) suggests that it is a more intuitive formulation for integrating multiple semi-independent feature images derived from \mathcal{A}_2 , compared to either (b) linear combination of multiple MI measurements or (c) GMI. For the synthetic example considered in this figure, intensity images \mathcal{A}_1 and \mathcal{A}_2 are in correct alignment.	22
2.4.	Flowchart describing the operation of COFEMI for the registration of a target intensity image \mathcal{B} to a stationary template intensity image \mathcal{A} . . .	24
2.5.	The quantitative evaluation protocol proposed by Skerl involves systematically transforming an image along “lines” in transformation parameter hyperspace. (a) A PD-w MR image shown at the two extrema of deformations ($m_s = -M/2$ and $m_s = M/2$) on a line n_s . (b) The corresponding T2-w MR image shown in its stationary position superposed on the deformed PD-w MR images in (a). The T2-w MRI remains stationary, while (c) values of image similarity (Ψ) are calculated at positions $m_s \in \{-M/2, \dots, M/2\}$ along each line n_s . Properties of Ψ are derived from the derivative (light blue) and local maxima (dashed vertical) along the profile of a measure’s trace. (d) Location and number of local minima (solid vertical) and the global maximum (dash-dotted vertical) are also used in the calculation of the properties described in Table 2.3. . .	32

- 2.6. Synthetic multiprotocol MRI of the brain is used to compare MI and COFEMI for two different combinations of simulated inhomogeneity and noise. (a) $\mathcal{T}1$ and corresponding (b) $\mathcal{P}_{b,\sigma}$ with simulated misalignment, high inhomogeneity strength ($b = 2.5$) and low noise ($\sigma = 1$). The boundary of the skull on $\mathcal{T}1$ ((a),(e)) and $\mathcal{P}_{b,\sigma}$ ((b)-(d), (e)-(h)) is shown in green and red, respectively. (c) MI registration result $\mathcal{P}_{b,\sigma}^r$ is susceptible to inhomogeneity in the lower region of $\mathcal{P}_{b,\sigma}$, as evidenced by misalignment of red and green contours. (d) COFEMI is able to compensate for the artifacts and noise, as evidenced by the nearly perfect alignment of the red and green contours. (e)-(h) For pair of $\mathcal{T}1$ and $\mathcal{P}_{b,\sigma}$ with low inhomogeneity strength $b = 1.0$ and higher noise $\sigma = 2$ 36
- 2.7. (a) Registration failure rates of MI over 5 levels of noise and 6 levels of field inhomogeneity ($b \in \{0, 0.5, 1.0, 1.5, 2.0, 2.5\}, \sigma \in \{0, 1, \dots, 4\}$) in $\mathcal{P}_{b,\sigma}$, where failure is $\phi^{TRE} > 1.5\text{mm}$. (b) Failure rates of COFEMI. (c) Plot of $\Delta\phi^{TRE} = \phi_{\text{MI}}^{TRE} - \phi_{\text{COFEMI}}^{TRE}$. (d) Plot of $\Delta\phi^{ROR} = \phi_{\text{COFEMI}}^{ROR} - \phi_{\text{MI}}^{ROR}$. CMI^{grad} is similarly compared to COFEMI by (e) $\Delta\phi^{TRE} = \phi_{\text{CMI}^{grad}}^{TRE} - \phi_{\text{COFEMI}}^{TRE}$, and (f) $\Delta\phi^{ROR} = \phi_{\text{COFEMI}}^{ROR} - \phi_{\text{CMI}^{grad}}^{ROR}$. COFEMI is consistently more robust to the presence of bias field and noise artifacts, compared with MI or CMI^{grad}. 38
- 2.8. Comparison of 9 similarity measures Ψ^u , $u \in \{\text{COFEMI}, \text{MI}, \text{NMI}, \text{ECC}, \text{EOH}, \text{JE}, \text{SSD}, \text{CC}, \text{PSMI}\}$ in terms of the 5 properties proposed by Skerl [8]. (a) ACC indicates that global maximum of all similarity measures correspond very closely to the ground truth registration. (b) CR, (c) NOM, and (e) DO for all Ψ^u . (d) NOM for only the MI-based measures in (c) is shown magnified for clearer comparison. RON was very high for measures $u \in \{\text{EOH}, \text{JE}, \text{SSD}, \text{CC}\}$, thus RON is shown for only MI-based measures in (f), indicating superior optimization characteristics of Ψ^{COFEMI} . The supervised PSMI method demonstrates marginally better performance in terms of NOM and RON versus COFEMI. 43

- 2.9. NOM versus R , the distance of deformation from ground truth registration. (a) NOM for $\Psi^u, u \in \{EOH, JE, SSD, CC\}$, demonstrate a large number of minima compared with (b) NOM for $\Psi^u, u \in \{COFEMI, MI, NMI, ECC, PSMI\}$. Note the scale difference between (a) and (b). Only PSMI demonstrates any improvement over COFEMI, which can only be seen as the marginal difference at the highest value of R in (b). 44
- 2.10. (a) RON box-plot for all similarity measures $\Psi^u, u \in \{COFEMI, MI, NMI, ECC, EOH, JE, SSD, CC, PSMI\}$. (b) Compact box-plot for $\Psi^u, u \in \{COFEMI, MI, NMI, ECC, PSMI\}$. Note the scale difference between (a) and (b). (c) RON averaged over N lines versus R for $\Psi^u, u \in \{EOH, JE, SSD, CC\}$ and (d) for $u \in \{COFEMI, MI, NMI, ECC, PSMI\}$. As with NOM, MI-based measures are superior, while Ψ^{COFEMI} and Ψ^{PSMI} have the best convergence characteristics. 45
- 2.11. (a) 4 Tesla *ex vivo* MR image (\mathcal{M}), (b) corresponding histological section \mathcal{H} , and (c)-(h) MR feature scenes \mathcal{M}_β corresponding to Haralick correlation ($\kappa = 8$) and inverse difference moment ($\kappa = 5$), mean intensity ($\kappa = 5$), range ($\kappa = 5$), range ($\kappa = 8$), and Haralick correlation ($\kappa = 5$). 46
- 2.12. (a)-(d) Four consecutive histology sections \mathcal{H} with the prostate boundary shown in green and (e)-(h) corresponding unregistered MRI slices \mathcal{M} with the prostate boundary $G^b(\mathcal{M})$ accentuated in red and the green boundary $G^b(\mathcal{H})$ from (a)-(d) overlaid. Misalignment of the red and green boundaries in (e)-(h) indicates the initial misalignment between $G^b(\mathcal{H})$ and $G^b(\mathcal{M})$. (i)-(l) Registered MRI slices \mathcal{M}^r generated by intensity-based MI demonstrate misalignment near hypointense regions at the top of \mathcal{M} , as evidenced by the space between the red and green contours. (m)-(p) Improved accuracy of COFEMI is evidenced by closer coincidence of $G^b(\mathcal{H})$ and $G^b(\mathcal{M}^r)$, particularly at the locations in (i)-(l) where MI exhibited misalignment. 48

3.1.	(a) MRI slice of a prostatectomy specimen with cancer present. (b) Likelihood image generated by a CAD system applied to (a), and (c) the prediction for cancer presence obtained by thresholding (b) at 50% confidence level and superposing on (a). (d) The corresponding histology slice shown labeled for cancer extent by H&E stain (black arrow). (e) Likelihood image calculated by CAD after removal of erroneous training instances, and (f) the resulting binary prediction.	54
3.2.	Outline of registration-based supervised classifier training system applied to CAD for prostate MRI with histological ground truth for cancer (top left). Labeling of MRI training data is achieved by registration of histology with MRI, followed by directly mapping histological ground truth to MRI, which are used to train the supervised classifier which drives the CAD system (bottom row). Training labels are also established by expert manual delineation, upon which another classifier is trained (top row). The classifiers are then evaluated with respect to each other (right).	57
3.3.	MRI intensity image (a) and corresponding independent component feature images obtained via ICA (b)-(h) of feature images (not shown). Maximal statistical independence between independent components is maximized to facilitate naïve combination in Bayes' rule.	65
3.4.	(a) Histology with ground truth indicated by the black arrow is registered to (e) MRI to transfer histological cancer map to MRI. Resulting mappings obtained via MI (c), COFEMI+TPS (d) and COFEMI (f). Overlays of MRI with histology registered via MI and COFEMI+TPS shown in (g) and (h), respectively, indicate the more accurate alignment with elastic vs. affine deformations. (b) Manually determined labels on MRI are shown for comparison.	69

3.5.	Accuracy of classifiers Ψ_{1-5} represented by the average ROC curves in	
	(a). Quantitative comparison of Ψ_{1-5} in terms of AUC (not shown)	
	indicates a statistically significant improvement from training labels es-	
	tablished by registration vs. expert delineation. Comparison in terms of	
	(b) PPV suggests an improvement in labeling using elastic registration.	71
4.1.	Comparison of MI and feature-driven COFEMI rigid registration of im-	
	ages with strong bias field inhomogeneity artifacts. (a) A T2 MR brain	
	image is registered to (b) the corresponding T1 MRI using (c) intensity-	
	based MI and (f) COFEMI using second order (d) correlation and (e)	
	inverse difference moment features. Green contours in (c) and (f) repre-	
	sent the boundary of the T2 brain MRI of (a) overlaid onto the registered	
	target. Red outlines accentuate the boundaries in the registration result.	
	Use of textural feature images by COFEMI was shown to improve regis-	
	tration of multiprotocol images with heavy intensity artifacts.	73
4.2.	A two step COFEMI-driven rigid and non-rigid registration methodology	
	applied in this study to perform automated alignment of two intensity	
	images. Initial global alignment is performed using COFEMI to optimize	
	an affine transformation of the target intensity image. Subsequently,	
	non-rigid registration via COLLINARUS is performed to determine the	
	remaining local deformations.	75
4.3.	Application of the COLLINARUS automated feature driven non-rigid	
	registration technique to alignment of (a) <i>ex vivo</i> whole mount histology	
	(WMH), (b) <i>in vivo</i> T2-w MRI and (c) <i>in vivo</i> DCE-MRI images of the	
	prostate and annotation of cancer on multiprotocol MRI. (a) Histopatho-	
	logic staining of whole-mount sections of a prostate with cancer provides	
	the “gold standard” for cancer extent. Non-rigid registration via COL-	
	LINARUS of (a) WMH to (b) corresponding <i>in vivo</i> MRI obtained prior	
	to resection allows the histological cancer map to be transferred onto	
	(b). (c) Corresponding DCE-MRI is registered to (b) by COFEMI rigid	
	registration to establish a map of cancer on (c).	76

- 4.4. (a) Synthetic T1-w MRI section, and (e) corresponding PD MRI section with simulated noise and bias field inhomogeneity artifacts. A deformation field, demonstrated on a grid in (f), is applied to (e) PD MRI to generate (b) a warped PD MRI section. Both MI and COFEMI are used to drive a non-linear B-spline based registration of (b) to (a). A correctly transformed PD MRI section would closely resemble (f). The results of (e) MI-driven and (g) COLLINARUS registration appear similar, while representations of the deformation field error magnitudes in (d) and (h) illustrate the greater error of MI compared with COFEMI. 83
- 4.5. Registration of WMH to 3T in vivo T2-w MRI and DCE-MRI of the prostate with cancer. (a) WMH with cancer extent delineated (dotted lines) is registered to (b) corresponding T2-w MRI using COLLINARUS to generate (c) transformed WMH in spatial alignment with the prostate in (b). (d) T2-w MRI with the cancer extent mapped from (c) superimposed in green. (e) DCE-MRI registered with (d) by an affine transformation, shown with cancer extent (green) mapped from (d) T2-w MRI. (g) The commonly used 3TP representation of the DCE data in (e), which provides an independent estimate of CaP extent, demonstrates that the cancer extent mapped by registration is accurate. (g)-(l) Similar results are demonstrated for a different set of multimodal prostate images from another study. 87

- 5.1. Ultrasound imagery of the prostate provides poor soft tissue resolution, while (b) high resolution MRI (*ex vivo* image shown) shows internal anatomical details of the prostate with greater clarity. Ground truth for CaP extent is obtained only through histopathologic analysis of (c) the corresponding Haematoxylin and Eosin stained tissue section. The histopathologic CaP extent (dark purple) on (c) can be mapped onto the MRI in (b) by either (d) manual labeling the MRI using histology as a visual reference, or (e) automatically mapping CaP extent from (c) via image registration. Note that the morphology of the CaP extent is better preserved in the mapping from (c) onto (e) as compared to (d). 92
- 5.2. Registration of an *ex vivo* prostate histology (\mathcal{A}) image to corresponding *in vivo* T2-w (\mathcal{B}) and T1-w (\mathcal{C}) MR images can be achieved in different ways. Two possible approaches using PW registration involve (a) PW alignment of histology to each individual MRI protocol ($\mathcal{A} \rightarrow \mathcal{B}$ and $\mathcal{A} \rightarrow \mathcal{C}$), or (b) alignment of multiprotocol MRI ($\mathcal{C} \rightarrow \mathcal{B}$) and alignment of histology to just T2-w MRI ($\mathcal{A} \rightarrow \mathcal{B}$). In the latter case, T1-w MRI would be in implicit alignment with histology at the end of the two registration steps. Alternatively, (c) a multi-attribute image registration scheme involves initial PW alignment of images from the same modality (T1-w and T2-w MRI) as in (b), followed by alignment of histology to a multi-attribute image comprising the registered multiprotocol MRI via a similarity measure defined for high dimensional data. 96
- 5.3. Establishing disease signatures on *in vivo* multiprotocol MRI using MACMI for registration of (a) WMH sections to corresponding (d) T2-w and (e) T1-w MRI. Alignment of the T2-w and T1-w MRI allows generation of (c) the multi-attribute MRI comprised of co-registered multiprotocol MRI. MACMI is used to aligned (a) WMH to (c) the multi-attribute MRI. CaP extent on the (b) elastically registered WMH is mapped directly onto both MR images in (c). 98

- 5.4. (a) Graphical representation of the organization of 4 images ($\mathcal{Z}_1, \dots, \mathcal{Z}_4$) within a family of image sets ($\mathbf{Z} = \{\{\mathcal{Z}_1\}, \{\mathcal{Z}_2\}, \{\{\mathcal{Z}_3\}, \{\mathcal{Z}_4\}\}\}$), and the application of the *MACMIreg* algorithm for alignment of all 4 images. Since only \mathbf{Z}_3 contains subsets (i.e. $|\mathbf{Z}_3| = 2$), line 3 of *MACMIreg* in (a) begins a new instance of the algorithm in (b) with $\mathbf{Z} = \{\{\mathcal{Z}_3\}, \{\mathcal{Z}_4\}\}$ as the input. The instance in (b) brings \mathcal{Z}_3 and \mathcal{Z}_4 into alignment, and returns the ensemble $[\mathcal{Z}_3, \hat{\mathcal{Z}}_4]$ to the instance in (a). The instance in (a) first bring \mathcal{Z}_1 and \mathcal{Z}_2 into alignment, and then align the ensemble of registered images from \mathbf{Z}_3 ($\epsilon_0 = [\mathcal{Z}_3, \hat{\mathcal{Z}}_4]$) with the registered images of \mathbf{Z}_1 and \mathbf{Z}_2 ($\epsilon = [\mathcal{Z}_1, \tilde{\mathcal{Z}}_2]$). At each registration step (line 13), a transformation \mathbf{T} is determined by $\arg\max_{\mathbf{T}}[I_2^*(\epsilon, \mathbf{T}(\epsilon_0))]$, and ϵ is then expanded by $\epsilon \leftarrow \tilde{\epsilon}_0 = \mathbf{T}(\epsilon_0)$ (lines 14 and 15 of *MACMIreg*). The output of (a), containing all of the co-registered images in \mathbf{Z} , is $\epsilon = [\mathcal{Z}_1, \tilde{\mathcal{Z}}_2, \tilde{\mathcal{Z}}_3, \tilde{\mathcal{Z}}_4]$. . 103
- 5.5. (a) 3T *in vivo* T2-w MRI of a prostate with a clearly visible DIL (shown in blue box) and magnified in (b) with a manual estimate of CaP extent (green). (c) Closest corresponding WMH slice with CaP ground truth (dotted line) and urethra (orange). (d)-(e) T2-w MRI with estimate of CaP extent (red) as mapped from (f) WMH via elastic registration using only T2-w MRI. (h)-(i) T2-w MRI with CaP estimate from (j) WMH registered to DCE (T1-w) MRI (co-registered to T2-w MRI). (k)-(l) Registration using both T2-w and DCE MRI via MACMI results in closer agreement of the registration-derived CaP extent (red) and the manual estimate (green). The verumontanum of the urethra (orange contour) is also shown on the registered WMH images in (f), (j), and (m). 113

5.6.	Using MACMI to include ADC MRI in the elastic registration of (a) histology to each of (d) T2-w, (e) DCE (T1-w) and (f) ADC MRI. Prior to elastic registration of histology, (d)-(e) T2-w and T1-w MRI were first successfully aligned via MI, as seen by (c) the checkerboard overlay of T2-w MRI and registered T1-MRI. (e) ADC was then registered to both T2-w and T1-w MRI via MACMI. (b) Elastically registered histology was obtained using the co-registered multiprotocol MRI via MACMI, and CaP extent was mapped onto (d)-(f) MRI (red lines).	114
6.1.	(a) Diagnostic MRI which may allow for identification of CaP extent. The diagnostic MRI must be registered to (b) planning CT so that a targeted dose plan may be generated. (c) Planning MRI can be used as an intermediary to facilitate alignment between (a) diagnostic MRI and (b) planning CT. (d) Registered diagnostic MRI (estimated cancer label in red) in alignment with planning CT.	120
6.2.	Overview of the CATT system. An unsupervised CAD system is used to identify suspect regions on the diagnostic MRI. A two step registration procedure is performed, which involves alignment of diagnostic MRI to planning MRI, followed by alignment of multiprotocol MRI to the planning CT. This allows for mapping the tumor regions onto CT, allowing for a targeted radiotherapy dose plan to be generated.	121
6.3.	(a) Original diagnostic MRI image with prostate boundary manually delineated (green). (b),(c) The median and Haralick correlation features of the 60 total textural features extracted from (a). Note the enhanced visibility of the hypointense lesion in (b) that is not visible in (a). (d) RGB visualization of 3 dimensional embedded feature space obtained by application of the non-linear dimensionality reduction scheme to the full texture space of the prostate voxels. (f) Map of the cancer region obtained by replicated k -means clustering in the embedded feature space in (e).	124

6.4.	(a) 3D visualization of the orientations and shapes of the 7 dynamically shaped radiation beams used to design the IMRT dose plan. (b) Axial view of a CT slice with the 7 beams and the resulting dose map thresholded at 7920 cGy (minimum dose for the PTV) overlaid. (c) Adjacent CT slice with the dose map thresholded at 8640 cGy (minimum dose for the cPTV) and boundaries for the PTV and cancer overlaid.	129
6.5.	(a) Diagnostic MRI is affinely registered to (b) planning MRI. (c) The registered diagnostic MRI is shown in (d) overlaid onto (b). (e)-(h) Checkerboard patterns of four additional pairs of aligned diagnostic and planning MRI demonstrate accuracy of registration, as evidenced the continuity of prostate capsule and internal structures across checks. . . .	131
6.6.	Elastic registration of both planning and diagnostic MRI to CT. (a)-(d) Overlay of all three imaging modalities of four slices from a single prostate study. (e)-(h) Deformed grid patterns for slices corresponding to (a)-(d) illustrate the non-linear nature of the deformation field required to align the prostate on diagnostic MRI to the visible gland on CT.	132
6.7.	(a), (b) Two different CT studies shown with cancer labels Ω^{CT} (white) mapped from MRI by registration and the resulting cPTV (red line). Resulting dose intensity maps for the slices in (a), (b) are shown thresholded at the minimum dose for PTV (7920 cGy, outer orange line) in (c), (d) and thresholded at the minimum dose for cancer (8640 cGy, inner red line) in (e), (f). These demonstrate dose escalation to the cancer PTV while maintaining dose to the prostate PTV. Note that the original dose plan has not been shown.	133

7.1.	(a) Reconstructed pseudo-whole-mount histology from a prostate specimen with cancer (delineated in dark purple) may be used to identify signatures for disease on pre-operative <i>in vivo</i> imagery, similar to the MRI image shown in (b), by registering the two images. (c) The registered MRI slice that is in alignment with whole histology in (a) allows mapping of cancer onto MRI (green). (d) Quadrants comprising the whole histology section in (a) must be first reconstructed into a pseudo whole-mount section to facilitate improved annotation across cuts and registration with MRI in (b).	136
7.2.	The HistoStitcher [©] graphical user interface where two histology pieces from the same section are digitally combined by selecting control point pairs along common boundaries. An optimal transformation of one fragment onto the the other is automatically determined based on user specified constraints.	140
7.3.	(a)-(d) Histology slide images of fragments (quadrants) of a whole section of a prostate specimen are reassembled using HistoStitcher [©] . (e) Half prostate histology section reconstructed by stitching quadrants in (a) and (b). (f) Three-quarters prostate histology section from stitching of quadrant in (c) with half in (e). (g) Final reconstructed pseudo-whole prostate histology section following stitching of quadrant in (d) with (e).	141

7.4.	(a) Photograph of specimen face prior to cutting into quadrants and mounting on slides. (b) Manually reassembled whole histology section obtained using Photoshop with low resolution images (final stitched image is 2,796 x 2,358 pixels). (c) HistoStitcher [©] -reassembled whole histology section from high resolution images (final stitched image is 48k x 41k pixels). Both stitches in (b) and (c) are validated by the photograph of the section taken prior to cutting. (d)-(g) Close up views of regions I-IV over the seams of the HistoStitcher [©] -reassembled image in (c). (h)-(k) Close up views of the regions nearest to I-IV on the manually-reassembled image in (b). The seams in (d)-(g) contain smaller gaps and better continuity of internal tissue structures compared to the seams in (h)-(k).	154
7.5.	Comparison of whole histology reassembly via Photoshop vs. HistoStitcher [©] . (a) Inadequate stitching is obtained using Photoshop on 4x resolution images (final image is 13975 x 13675 pixels), as computer system memory of 2GB was insufficient to introduce the transparency layers required to bring the images closer. (b) Stitching results using Photoshop on low resolution (about 1x magnification) images (final image was 2706 x 2244 pixels). (c) HistoStitcher [©] result on high resolution (8x magnification) images (final image is about 22k x 18k). (d)-(f) Zoomed regions I-III of manually stitched image (b), compared to (g)-(i) same regions of HistoStitcher [©] result (c), demonstrating both more contiguous stitching with minimal gaps and improved continuity of tissue structures across the stitch using HistoStitcher [©]	157
7.6.	Three prostate histology sections reconstructed using Photoshop (left column) and HistoStitcher [©] (right column). Stitching quality scores for these sections are given in Table 7.3 under section numbers 6 for (a) and (b), 1 for (c) and (d), and 4 for (e) and (f).	160

Chapter 1

Introduction

1.1 Overview

Multimodal and multiprotocol image registration refers to the process of alignment of two images obtained from different imaging modalities (e.g. digitized histology and MRI) and protocols (e.g. T2-weighted and PD MRI), utilizing either rigid or non-rigid coordinate system transformations. Both processes are critical components in a range of applications, including image guided surgery [9–11], multimodal image fusion for cancer diagnosis and treatment planning [12], and automated tissue annotation [13]. However, registration of multimodal imagery has posed a more challenging task compared with alignment of images from the same modality or protocol on account of differences in both image intensities and shape of the underlying anatomy. The first of these hindrances, dissimilar intensities between modalities, arises as a consequence of the measurement of orthogonal sources of information such as functional (SPECT) and structural (CT/MRI) imagery [12], as well as on account of other factors such as intensity artifacts, resolution differences, and weak correspondence of observed structural details. For example, artifacts such as inhomogeneity and non-standardness in MRI [14,15], speckle and shadowing in ultrasound [16], and beam hardening and photon scatter in x-ray represent common but particularly formidable hindrances. Interprotocol image registration tasks, such as alignment of high resolution MRI and interventional MRI (iMRI) [9] for guided treatment, are also hindered, although to a lesser extent, by reduced image similarity owing primarily to resolution differences.

Further, non-linear shape differences are common between real multimodal biomedical image data sets, and must be corrected using require elastic transformations. For example, registration of images of highly deformable tissues such as in the breast have

been shown to require flexible non-rigid techniques [17]. Similarly, non-linear differences in the overall shape of the prostate between *in vivo* MRI and *ex vivo* whole mount histology (WMH) have been shown to exist as a result of (1) the presence of an endorectal coil during MR imaging and (2) deformations to the histological specimen as a result of fixation and sectioning [18, 19]. Consequently, achieving correct alignment of such imagery requires elastic transformations to overcome the non-linear shape differences. Most automated registration techniques rely upon intensity-based similarity measures, which have been shown to be wanting for robustness across highly dissimilar modalities and in the presence of artifacts [3].

To overcome the challenges of both non-linear deformations and intensity artifacts, we present new registration techniques termed Combined Feature Ensemble Mutual Information (COFEMI) and Collection of Image-derived Non-linear Attributes for Registration Using Splines (COLLINARUS). The goal of the COFEMI technique is to provide a similarity measure that is driven by unique low level textural features, for registration that is more robust to intensity artifacts and modality differences than similarity measures restricted to intensities alone. The operation of COFEMI is described in detail in Chapter 2. COLLINARUS offers the robustness of COFEMI to artifacts and modality differences, while allowing fully automated non-linear image warping at multiple scales via a hierarchical B-spline mesh grid optimization scheme. The operation of COLLINARUS is described in detail in Chapter 4. In addition, since routine clinical imaging procedures also involve the acquisition of multiple imaging protocols, we present a technique termed Multi-attribute Combined Mutual Information (MACAMI) to leverage the availability of multiple image sets to improve registration. The operation of MACMI is described in detail in Chapter 5.

The registration schemes were developed to perform an automated tissue annotation task that is designed to facilitate the development and evaluation of a novel system for computer-assisted detection (CAD) of prostate cancer on multiprotocol MRI. The development of a multimodal CAD system requires ground truth labels for cancer on each modality to characterize malignant tissue. Since these MRI pixel labels are usually obtained by manual delineation, they can be extremely time consuming to generate and

subject to errors and bias of the expert performing the annotation, which adversely affects CAD performance. Therefore, to improve labeling and hence CAD classifier accuracy, we perform alignment of *in vivo* images with corresponding *ex vivo* whole-mount histology (WMH), the source of the cancer "gold standard", via automated multimodal image registration using our novel registration techniques. Following registration, cancer ground truth on registered WMH is directly mapped onto corresponding MRI studies.

We apply our registration techniques to a unique dataset combining multimodal images comprising *in vivo/ex vivo* MRI and whole mount histology (WMH) in order to establish the spatial extent of prostate cancer (CaP) on multi-functional structural (T2-weighted) and functional (DCE) *in vivo* MRI. Elastic registration using our feature-driven registration technique COFEMI is demonstrated in conjunction with a spline-based warping for registration of WMH and *ex vivo* MRI. A CAD system for *ex vivo* MRI was trained using the registration-established CaP labels and demonstrated to provide accuracy greater than CAD trained using manually established labels, as described in Chapter 3. Automated elastic registration and mapping of CaP labels from WMH to *in vivo* T2-w and DCE MRI via COLLINARUS was then demonstrated for 150 sets of clinical MRI and WMH images from 25 patient studies. Finally, for the same 25 patient studies, we also demonstrate elastic registration of the WMH to both the T2-w and DCE MRI protocols simultaneously using MACMI to combine and utilize the information channels associated with the two MRI protocols.

To demonstrate the utility of our elastic registration algorithms for improving treatment possibilities, the CAD system was then applied to detection of lesions on an independent data cohort comprising corresponding T2-w MRI and computerized tomography (CT) acquired for radiation therapy planning. Having determined suspect regions on MRI, we leveraged our registration algorithms to align the MR data to the CT studies, hence establishing the location of CaP on CT. These suspected regions were then used to generate intensity modulated radiation therapy (IMRT) plans providing dose escalation to the detected lesion, while maintaining the usual dose to the remainder of the prostate. We refer to this combination of image registration and CAD

in radiation therapy planning as Computer-Assisted Targeted Therapy (CATT), and describe the operation of CATT in detail in Chapter 6.

In the following section, we describe in detail the requirements and challenges in developing a CAD system for accurately identifying the spatial location of prostate cancer on *in vivo clinical*. We also introduce the novel image registration techniques that were developed in this thesis to address these challenges.

1.2 Computer-Assisted Detection of Prostate Cancer

Cancer of the prostate (CaP) is the most common malignancy among men with 217,730 new cases estimated to have occurred in 2010 alone, including 32,050 fatalities (American Cancer Society). Early diagnosis of CaP expands the possibilities for treatment and provides the best hope for curing it [20], however no image-based protocol exists for its detection. The current protocol for CaP detection and diagnosis is a screening test based on elevated levels of the prostate-specific antigen (PSA) in the blood [21], followed by a needle biopsy guided using transrectal ultrasound (TRUS) to locate the prostate *in vivo* [22], for patients with high PSA levels. However, TRUS biopsies have been associated with low CaP detection accuracy due to (a) the low specificity of the PSA test, and (b) lack of contrast between benign and malignant tissues in ultrasound (US) [23]. Further, since no *a priori* knowledge of tumor location is available, biopsy locations are chosen at random.

High resolution MRI has been demonstrated to show better separation of normal benign tissue from cancer compared to ultrasound [22, 24, 25]. However, human visual recognition has proven problematical as many benign features, such as atrophic glands and stromal over-growth, exhibit similar texture and intensity characteristics to malignant tissue. In response to the limitations associated with visual detection of prostate cancer, methods using image processing and machine learning have recently been applied to the task of improving detection. We have demonstrated [26–29] the efficacy of a computer-aided detection (CAD) system for detection of prostate cancer from high resolution *ex vivo* MR of prostate specimens, and showed that its detection sensitivity

and specificity were comparable to trained radiologists. One of the challenges in extending CAD to *in vivo* is determining the spatial extent and location of tumor on *in vivo* MR in order to train the system to detect lesions. In light of low inter-observer agreement in the task of determination of ground truth on *ex vivo* MR by visually mapping histology ground truth, the quality of manually determined ground truth on *in vivo* MR would be even more unsuitable.

The overall goal is to allow the construction of a system for detecting prostatic adenocarcinoma in high resolution *in vivo* MRI, thus providing a powerful tool for prostate cancer screening and staging. In this research, we wish to develop the image registration methods necessary to facilitate the use of CAD on *in vivo* MR. It is our general hypothesis that such a system would help (i) reduce the number of unnecessary biopsies by avoiding biopsying everyone with a high PSA, (ii) decrease the number of false positives associated with TRUS by improving the accuracy of needle insertions during biopsy procedures, and (iii) improve therapy through targeted dose conformality and reduced dosage to healthy tissue. In addition, the ability to characterize the 3D morphology of CAP will improve the accuracy of targeted radiation therapy.

The CAD system developed in [26] produces a likelihood map for tumor presence (Figure 1.1(c)) at every point on a high resolution *ex vivo* MR image by stochastically combining 3D image texture features. Sophisticated classification [28] techniques have been developed for the system, and novel implementations of machine learning methods such as graph embedding [29] have been utilized in the fine tuning of its detection capabilities. On *ex vivo* trials, the CAD system performed comparably to trained experts in terms of accuracy where the expert segmentations were determined by visual examination of the MRI images. The system has also demonstrated [28] significantly lower intra-system variability (of CAD accuracy with change in system parameter) compared to experts' intra- and inter-observer variability (Figure 1.1(d)). Three human experts' determinations of CAP presence on *ex vivo* MR are shown superimposed on Figure 1.1(d). Note that there is little inter-observer agreement (white) and that expert determinations are highly inaccurate unless the determination is made using histology as a reference as in Figure 1.1(b). The CAD result in Figure 1.1(c) is comparable to




Figure 1.1: (a) *ex vivo* prostate MRI, (b) tumor ground truth from histology, (c) CAD result, (d) superposition of 3 human expert segmentations of CAP on (a). Note that (i) CAD performance is comparable to human experts and (ii) the low levels of inter-observer agreement in (d). The bright areas in (d) are the only areas where all expert segmentations agree.

tumor ground truth determined from histology. We wish to extend the CAD system to cancer detection *in vivo*, and we plan to address several key challenges in making this important transition.

There are important limitations of the current CAD system that must be addressed to implement a CAD for *in vivo* MR. The CAD system depends on the availability of accurately delineated tumor boundaries on MRI datasets for training of the system’s tissue classification parameters and for evaluating the CAD-produced tumor segmentations for accuracy against a “ground truth” for tumor presence. Previously, however, this *ex vivo* MR ground truth was determined by manually mapping the tumor boundary from the histological sections onto the MR data, a procedure that is prone to errors and subject to observer bias (Figure 1.1(b) and (d)). Further, expert mappings of ground truth on MR are not based on precisely corresponding histology slices, but approximations of corresponding data. Delineation of accurate ground truth will improve accuracy of *ex vivo* tumor detection by improving prior feature distributions used for tissue classification. Previous studies have demonstrated the consequences of inaccuracies in ground truth by detecting classification errors using the graph embedding techniques [29]. This method was demonstrated to be able to correct for errors in labeling by experts and thus yield a refined set of expert labeled cancer instances. Errors in tumor segmentation in Figure 1.2(a) are detected using graph embedding to identify outlier pixels (Figure 1.2(b)). The outliers in the graph embedding result correspond




Figure 1.2: (a) Manually determined ground truth from histology. (b) Representation of tumor texture features in reduced dimensional space. Cluster outliers shown in red on (b) represent incorrectly classified tissue also shown in red on (c).

to mislabeled regions shown in Figure 1.2(c). These discrepancies need to be addressed to determine CAP ground truth on *ex vivo* MR from histology.

Since the tumor ground truth is similarly absent on *in vivo* MR data sets, an analogous procedure for mapping the ground truth from histological sections and *ex vivo* MRI to corresponding *in vivo* MRI is required for training and evaluating the CAD model for detection of cancer *in vivo*. The process of physically slicing and quartering of histology sections results in significant deformation and tissue loss. In order to correctly match material points on histology with MR, it is necessary to account for these deformations using elastic transformations. Given these issues, it is necessary to reliably establish *ex vivo* ground truth in a robust, unbiased manner.

The lower spatial resolution of *in vivo* MRI compared to *ex vivo* MRI makes determination of distinct image features difficult, hindering the ability of expert observers to accurately align the *in vivo* MR slices to histology. Manual registration of *in vivo* MR directly to histology is complicated by the absence of co-oriented *in vivo* MR and histology slices and by the presence of out-of-plane deformations in *in vivo* prostate MR data. In addition, the decreased signal-to-noise ratio of *in vivo* MR poses a problem for implementing automated registration techniques as the structural information contained in the data is much less than that of the *ex vivo* data sets. To address these problems and build a CAD system for detection of prostate cancer from high resolution *in vivo* MR data, we propose the following,

1. Since tumor ground truth is determined from histology, we propose elastic registration and mapping ground truth from whole-mount histology to *ex vivo* MRI and *in vivo* MRI of the prostate.
2. The use of textural features derived from MRI intensities to improve the performance of image similarity measures in the registration procedure.
3. The use of multiprotocol (T2-w and DCE) MRI as an adjunct to scalar image intensities to enhance the registration procedure.

1.3 Previous Work

Manual labeling of structures on radiological images is often used as a surrogate of ground truth for evaluating CAD algorithms, however we [29,30] and other researchers [31–33] have demonstrated that expert segmentations are often corrupted by a significant number of mislabeled instances. These errors in the ground truth surrogates could severely compromise the efficacy CAD models as well as the validity of CAD evaluation results. Several researchers have demonstrated [34,35] that utilization of one or even multiple expert segmentations is not suitable as a basis for ground truth. We propose the use of image registration to map the histology ground truth onto the *ex vivo* MR images. Tumor regions determined on *ex vivo* MR will then be mapped onto *in vivo* MR, via an elastic image registration technique. Although some researchers have attempted to use whole mount histological prostatectomy specimens to demonstrate the use of high resolution *in vivo* MR in visually identifying CAP [36,37], little work exists on the task of image registration of prostate histology and MR on account of the numerous difficulties involved. In [28] we attempted to map tumor extent determined on histological sections to MR by maintaining the *ex vivo* MR slices and histological slices in the same plane of section. Areas visually determined to be affected by cancer were manually mapped to *ex vivo* MR. Manual registration is approximate at best and does not account for the processes of tissue fixation and slicing of histology sections which results in deformation and tissue loss. A robust registration algorithm will require localized deformations to be applied. Multimodality problems such as this one are also

faced with data that have different slice thicknesses and spatial resolution, such that the histology slices, while of high spatial resolution within the plane of the image, do not directly correspond to specific slices within the *ex vivo* MR volume. A volumetric deformation algorithm is necessary to overcome this aspect of multimodality image registration.

Mutual information (MI) has recently [38–40] been introduced as a similarity measure for voxel-based registration of multimodality images, due in part to the robustness of the measure to intensity differences between modalities and protocols. In its capacity as a global image similarity measure, MI is often used as an objective function for rigid image registration followed by feature-driven elastic deformations. Recently, some researchers have developed warping algorithms such as finite element models [11], free form deformations [12, 17], thin-plate-splines (TPS) [41], and Gaussian kernels [42] to accomplish elastic registration in sequence with rigid registration by maximization of mutual information. MI has also been utilized in the elastic registration step as a metric to validate TPS deformations in 2D images [43] and more recently in volumetric data [44]. Porter et al. [10] and Moskalik et al. [45] have attempted to register histological and US images of the prostate. In [10] rudimentary corrections such as rescaling the aspect ratio of the images were used to account for histological deformations. Lee et al. [12] have utilized manually defined anatomical landmarks to guide non-rigid registration between histological prostate sections with low resolution *in vivo* MRI. Since an external array coil and not an endorectal coil was used in the MR imaging, no prostate deformation occurred. Consequently a simple rigid registration combined with deformation of landmarks was used to warp the histology onto the MRI via a TPS transformation. In many studies anatomical landmarks have been used to aid the registration procedure [46–50], however automatic selection of landmarks for registering prostate imagery has not received much attention. Generic methods for selecting point correspondences and performing non-rigid deformations have been described in [46]. Our proposed method will integrate automatic anatomic fiducial determination with an elastic registration technique to map tumor regions from histological sections onto the *ex vivo* MRI sections.

The mapping of tumor regions from *ex vivo* onto *in vivo* MRI is complicated by presence of large prostate deformations due largely to the endorectal (ER) coil, presence of which is necessary to maximize signal intensity. In addition, the *in vivo* MR protocol is subject to significantly higher noise and MR-related artifacts compared to *ex vivo*. Alterovitz et al. [51] have studied ER coil induced deformations *in vivo* and have shown that they produce significant dimensional and volumetric changes. Registration of prostate CT images subject to large deformations due to variable presence bowel gas has been addressed by Foskey [52]. In [11] Bharatha et al. presented an alignment strategy for registering segmented pre-operative (1.5 T) and intra-operative (0.5 T) MR images of the prostate using an active surface algorithm. du Bois d’Aische et al. [49] proposed an improvement to [11] by using a MI based non-rigid registration algorithm to match the internal structures of the prostate. Both [11, 49] however, only looked at registering *in vivo* MRI to *in vivo* MRI and not the more difficult task of registration of *in vivo* MRI to *ex vivo* MRI or histology. We present a novel approach using elastic deformation driven by multiparametric image information to register the *in vivo* and *ex vivo* MRI with histology.

More detailed descriptions of previous work on registration of histology and MR images of the prostate, and the problems of elastic and multimodal image registration in general, are provided in Chapters 2-5.

1.4 Overview of Contributions

To address the challenges of automated multimodal elastic image registration, we have developed a set novel registration methodologies. Our methods are aimed at improving the image similarity measure, which drives the optimization of the image transformation, by (1) incorporating multiple independent textural features, and (2) leveraging the availability of multifunctional image sets. First, we have developed an improved similarity measure for multimodal images based on the combination of multiple calculated textural features. We have implemented the texture-driven similarity measure within both affine and elastic registration frameworks. This technique, termed Combined Feature Ensemble Mutual information (COFEMI) [53], is described in detail and

demonstrated for the affine registration of *ex vivo* histology and *ex vivo* MRI in Chapter 2. The application of COFEMI for mapping regions of cancer from histology onto *ex vivo* MRI, and subsequently training and evaluating a supervised classifier for detection of prostate cancer on *ex vivo* MRI, is demonstrated in Chapter 3. In this study, we have also developed a novel pixelwise CAD system using independent component analysis (ICA) of multiple textural features to drive a naïve Bayesian classifier for the detection of cancer [5]. In Chapter 4, we demonstrate a novel automated elastic registration scheme, termed Collection of Non-linear Attributes for Registration Using Splines (COLLINARUS) [1], which implements the COFEMI technique within an elastic registration framework. The application of COLLINARUS for mapping regions of cancer from histology directly onto *in vivo* (pre-operative) MRI is demonstrated for 25 patient studies with prostate cancer comprising 150 pairs of T2-w MRI and whole mount histology images. Our second primary contribution is the development of a technique for incorporating information from multiple co-registered image channels, such as multiple MRI protocol image series that are commonly acquired as part of routine clinical practice, in a multivariate similarity measure to improve registration performance. This technique, termed Multi-Attribute Combined Mutual Information (MACMI) [2, 7], is described in detail and demonstrated for the elastic registration of *ex vivo* histology and multiprotocol (T2-w and DCE) *in vivo* MRI in Chapter 5. For 25 patient studies, MACMI is applied to mapping spatial extent of cancer onto both *in vivo* T2-w MRI and *in vivo* DCE MRI by simultaneous elastic registration with 150 histology images.

In addition, we have leveraged our image registration and CAD techniques for creating intensity modulated radiation therapy (IMRT) plans providing dose escalation to the dominant intraprostatic lesion, while maintaining dose to the entirety of the prostate and minimizing dose to neighboring healthy tissue.

To facilitate both registration and annotation of whole-mount histological sections, we have also developed a software package for digital reconstruction of whole-mount sections from tissue fragments such as the quadrants employed in Chapters 2 and 3. The software, referred to as HistoStitcher[®], provides a graphical user interface for interactive reassembly of adjacent fragments and a set of routines for efficient and flexible image

transformation and stitching. The design and operation of HistoStitcher is described in detail in Chapter 7.

1.5 Notation

Table 1.1 below lists the common symbols used throughout this document. Notation and symbols specific to the studies in each of the following chapter are listed in separate tables.

Table 1.1: Common notation and symbols used throughout this document.

Symbol	Description
\mathcal{A}	2D template intensity image
\mathcal{B}	2D target intensity image
C^A	Coordinate set defined for \mathcal{A}
C^B	Coordinate set defined for \mathcal{B}
$f_A(c)$	Value of image \mathcal{A} at pixel $c \in C^A$
$f_B(c)$	Value of image \mathcal{B} at pixel $c \in C^B$
\mathcal{B}^r	Registered target \mathcal{B}
Φ_β	Feature operator where $\beta \in \{1, \dots, n\}$
$f^{\Phi_\beta}(\mathbf{c})$	Feature value at \mathbf{c} for Φ_β
$\mathbf{f}(\mathbf{c})$	Feature vector $[f^{\Phi_\beta}(\mathbf{c}) \beta \in \{1, \dots, n\}]$ at \mathbf{c}
\mathbf{T}^{ap}	Known, applied transformation
\mathbf{T}^{co}	Corrective transformation

Chapter 2

Registration of *ex vivo* Whole-Mount Histology and *ex vivo* MRI of Prostate Specimens

2.1 Introduction to Image Registration

2.1.1 Image Similarity-based Multimodal Image Registration Background

Multimodal and multiprotocol image registration refers to the process of alignment of two images obtained from different imaging modalities (e.g. PET and MRI) and protocols (e.g. T1-w and T2-w MRI), respectively. Both processes are critical components of medical applications ranging from routine clinical tasks to cutting edge visualization and diagnostic tools, including image guided surgery [9–11], multimodal image fusion for cancer diagnosis and treatment planning [12], and automated image based tissue/region annotation [13]. However, registration of multimodal imagery is a significantly more challenging task compared to alignment of images from the same modality or protocol, since image intensities between modalities may not reflect the same structural attributes. More specifically, considerable differences in appearance between modalities may arise as a consequence of the measurement of orthogonal sources of information such as functional (SPECT) and structural (CT/MRI) imagery [12]. Other important sources of dissimilarity include image intensity artifacts and resolution differences. Artifacts such as bias field inhomogeneity and intensity non-standardness in MRI [15], speckle and shadowing in ultrasound [16], and beam hardening and photon scatter in X-ray represent common but particularly formidable hindrances to multimodal image registration. Interprotocol image registration tasks [9], are similarly hindered by reduced image similarity owing primarily to differences in image resolution and fidelity.

The registration methodology presented in this study is geared toward addressing (1) a lack of precise correspondence between anatomical fiducials observable in the different modality images and (2) the presence of modality-specific artifacts in one image but not the other(s) being registered.

Numerous multimodal image registration methods have focused on the development of voxel-based similarity measures to establish image alignment, yet the optimal mechanism for incorporating the information encoded by the image intensities and their spatial arrangements remains an open issue [54]. The spirit of voxel-wise similarity measures (Ψ) is that when images \mathcal{A} and \mathcal{B} are brought into alignment via an appropriate transformation \mathbf{T}_{max} , obtained over multiple possible transformations (\mathbf{T}), their similarity is maximized.

$$\mathbf{T}_{max} = \underset{\mathbf{T}}{\operatorname{argmax}} [\Psi(\mathcal{A}, \mathbf{T}(\mathcal{B}))] \quad (2.1)$$

However, when images \mathcal{A} and \mathcal{B} are obtained from different imaging modalities or protocols, a non-linear relationship between voxel intensities in \mathcal{A} and \mathcal{B} may exist. For this reason, non-linear, statistical measures of image similarity or interdependence such as mutual information (MI) [38, 55], entropy correlation coefficient (ECC) [38], and normalized MI (NMI) [56] are uniquely suited to multimodal image registration tasks [39]. Despite the general robustness of information theoretic quantities in handling the non-linear relationship of image intensities between dissimilar modalities, these measures are limited in their ability to handle intensity artifacts and highly dissimilar modalities. The shortcomings of MI are primarily a consequence of the limitations of image intensities to overcome these challenges.

Recently, some researchers [57–59] have attempted to address the failure of MI and other intensity-based similarity measures using low-level image-based textural features to incorporate information regarding the spatial relationships of image intensities. Plum, et al. [57] proposed incorporating edge information by weighting MI or NMI values by a measure of gradient vector agreement. Studholme, et al. [58] used a form of MI to consider connected component region labels as an image feature. Rueckert, et al. [59] derived a multivariate counterpart of MI to characterize spatial co-occurrence of

intensities. The form of MI used in [58] and [59], which we refer to as combined mutual information (CMI), uses higher order entropy calculations to account for redundancies between feature values and the image intensities from which they are derived. The use of CMI for inclusion of semi-independent feature images that are in implicit alignment with an intensity image is a more intuitive choice of higher-order MI compared to previously presented formulations which involved linearly combining multiple pairwise image similarity values or other *ad hoc* formulations [60, 61]. These CMI formulations [58, 59] however lack (1) quantitative justification for choice of features to complement image intensity in driving the registration, and (2) a rationally designed method for selecting appropriate features for an arbitrary registration task, suggesting that these schemes may not be readily generalized to a broad range of multimodal registration tasks. The theory behind previously proposed information theoretic quantities [39, 58] is described in detail in Section 2.2.

2.1.2 Motivation for Multiple Image Features in Image Registration

Multiple textural features derived from image intensities can provide valuable structural descriptions not readily discernible by use of a single image feature such as image intensity. In addition, the use of textural attributes allows for alternate image intensity representations that are more robust to imaging artifacts [14, 16]. In [14] and [62], the utility of textural features in discrimination of posterior acoustic shadowing from the lesion area in breast ultrasound was demonstrated. Similarly, spatial image features calculated from statistics on gradients were successfully used in [63] to improved retrospective correction of intensity inhomogeneity artifacts on MRI. The utility of image texture attributes has also been shown to improve accuracy of image point matching tasks [42, 64, 65] when used in conjunction with information theoretic techniques. It is thus intuitive that multiple textural attributes can also provide the same benefits of improved image correspondence and robustness to artifacts in multimodal registration tasks. For example, given two intensity images \mathcal{A} and \mathcal{B} from different modalities, multiple textural representations of \mathcal{A} or \mathcal{B} may (a) not be plagued to the same extent by intensity artifacts as \mathcal{A} or \mathcal{B} , and may (b) improve correspondence of structural details

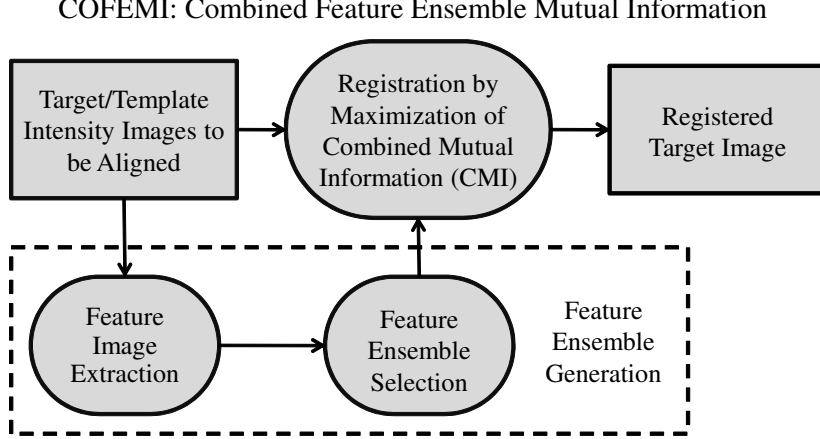


Figure 2.1: Overview of the operation of COFEMI. Automatic selection of an optimal ensemble of textural features is central to the COFEMI method for multimodal registration.

between the two modalities. The value of a similarity measure that is able to incorporate the information encoded by these textural features will hence correlate more closely with degree of alignment of \mathcal{A} and \mathcal{B} .

2.1.3 Combined Feature Ensemble Mutual Information

The registration scheme presented in this study, Combined Feature Ensemble Mutual Information (COFEMI), employs an algorithm for (i) optimal textural feature selection and (ii) integration of multiple image derived attributes to create an improved similarity measure to drive multimodal image registration. The COFEMI method operates by dynamically selecting maximally informative textural representations from a larger, comprehensive feature set, and then incorporates them via an information theoretic formulation tailored to capture non-redundant information between the two images to be registered. The flowchart in Figure 2.1 shows the modules comprising the COFEMI technique. From the image(s) \mathcal{A}, \mathcal{B} to be registered, a large number of corresponding textural images are generated as candidates for inclusion into a feature ensemble. These feature images are obtained via application of (i) gradient, (ii) first order statistical, and (iii) second order statistical operators on the images to be registered (Section 2.3.2). Since only a subset of the most informative features are required to provide the

additional information to formulate a well defined similarity measure for multimodal image registration, COFEMI selects textural feature images for inclusion into an optimal ensemble via a novel joint entropy maximization technique (Section 2.3.3). The feature selection criteria include (1) the descriptiveness of the feature images in terms of information content and (2) the distinctiveness of an image attribute to be included in the ensemble with respect to the other existing feature images in the ensemble. The COFEMI feature selection criteria are similar to previous feature selection approaches that consider feature redundancy and feature relevance [66, 67], as quantified by MI. Note however, that to the best of our knowledge, COFEMI is the only multimodal image registration scheme to employ intelligent feature selection to generate an optimal global similarity measure. Following feature ensemble generation, correct image alignment is achieved by optimizing some transformation \mathbf{T} via maximization of the CMI multivariate similarity measure of the corresponding image ensembles.

2.1.4 Quantitative Evaluation of Registration Tasks

The decision to utilize a particular method for evaluation of registration accuracy or similarity measure performance depends on availability and type of ground truth. When the ground truth is unknown, evaluation is usually qualitative and consists of visual assessment of alignment quality. Similarity measure values themselves are often used in the absence of ground truth to infer alignment quality, although it is generally not appropriate to evaluate the performance of a similarity measure via the same or a related measure (e.g. using MI values to quantify accuracy of MI-based registration) [68]. On the other hand, if ground truth in the form of anatomical object segmentations is available, meaningful estimates of registration accuracy can be defined in terms of boundary alignment and overlap of corresponding structures on both images. Manually placed fiducial markers may also provide an estimate of registration accuracy by calculating target registration error (TRE) [69]. Further, with synthetic image data, such as is available from the multiprotocol MR simulated brain database at the Montreal Neurological Institute [70], definitive quantitative evaluation is possible because ground truth for alignment is known. With such data, registration performance may be measured

directly in terms of TRE, which may be calculated over multiple registration trials. Alternatively, the utility of a similarity measure as a registration objective function may be characterized by inspecting the behavior of the measure over a range of misalignments [38]. Using this approach, systematic protocols for thorough characterization of similarity measures have been proposed by Skerl, *et al.* [8, 71].

The evaluation techniques used in this study make use of ground truth for alignment, and region segmentations to infer alignment accuracy. In 3 sets of experiments designed to evaluate robustness to modality differences and intensity artifacts, a total of 8 unique quantitative evaluation measures are calculated to characterize COFEMI. On 30 synthetic MR images with known alignment ground truth, registration error from MI, COFEMI, and an image gradient-driven extension of MI is compared for 6 simulated inhomogeneity levels and 5 noise levels. We also implement the evaluation protocol proposed by Skerl, *et al.* [8] to compare 8 image similarity measures with COFEMI in terms of 5 properties designed to characterize registration capabilities without solving Equation (2.1). Finally, on an additional 26 real multimodal (MRI and whole mount histology) prostate images, for which alignment ground truth is not available, MI, COFEMI, and the gradient extended MI are compared in terms of region overlap and boundary alignment.

The organization of the rest of the study is as follows. Section 2.2 reviews the information theory used in MI and related multivariate formulations, including the CMI measure considered for COFEMI. Section 2.3 describes the COFEMI registration scheme in detail, including the feature calculation and selection schemes, and the optimization approach. Section 2.4 describes the methods used to evaluate and compare the performance of COFEMI with other methods. In Section 2.5 we present and discuss the results of registration evaluation applied to two data sets, one synthetic and one clinical. Concluding remarks and future directions are presented in Section 2.6.

2.2 Theory of Combined Mutual Information

2.2.1 Mutual Information

MI is often defined in terms of Shannon entropy [72], a measure of information content of a random variable. Equation 2.2 gives the marginal entropy, $S(\mathcal{A}_1)$, of image \mathcal{A}_1 in terms of its gray level probability distribution $p(a_1)$,

$$S(\mathcal{A}_1) = - \sum_{a_1} p(a_1) \log p(a_1), \quad (2.2)$$

where $a_1 \in \{0, \dots, f^{max}\}$ and f^{max} is the highest gray level intensity in image \mathcal{A}_1 . While $S(\mathcal{A}_1)$ describes the information content of \mathcal{A}_1 , the joint entropy $S(\mathcal{A}_1\mathcal{A}_2)$ of an ensemble of two images \mathcal{A}_1 and \mathcal{A}_2 describes the information gained by simultaneous knowledge of both images ($\mathcal{A}_1\mathcal{A}_2$). This may be expressed as in [39, 72],

$$S(\mathcal{A}_1\mathcal{A}_2) = - \sum_{a_1, a_2} p(a_1, a_2) \log p(a_1, a_2), \quad (2.3)$$

where a_2 represents the gray levels in \mathcal{A}_2 . Thus, when image \mathcal{A}_1 best explains image \mathcal{A}_2 , joint entropy is minimized to $\max[S(\mathcal{A}_1), S(\mathcal{A}_2)]$. Equation 2.4 below is a common formulation of MI of a pair of images (or random variables) in terms of the marginal and joint entropies wherein $I_2(\mathcal{A}_1, \mathcal{A}_2)$ is maximized by minimizing joint entropy $S(\mathcal{A}_1\mathcal{A}_2)$ and maintaining the marginal entropies $S(\mathcal{A}_1)$ and $S(\mathcal{A}_2)$.

$$I_2(\mathcal{A}_1, \mathcal{A}_2) = S(\mathcal{A}_1) + S(\mathcal{A}_2) - S(\mathcal{A}_1\mathcal{A}_2), \quad (2.4)$$

where $I_2(\mathcal{A}_1, \mathcal{A}_2)$ describes the interdependence of 2 variables, or gray levels of a pair of images [39]. A Venn diagram representation of $S(\mathcal{A}_1)$ and $S(\mathcal{A}_2)$ is shown in Figure 2.2(a) where the shaded region represents $I_2(\mathcal{A}_1, \mathcal{A}_2)$. As $I_2(\mathcal{A}_1, \mathcal{A}_2)$ increases, the uncertainty about \mathcal{A}_1 given \mathcal{A}_2 decreases. As such, it is assumed that the global MI maximum will occur at the point of precise registration, when maximal uncertainty about intensities of \mathcal{A}_1 is explained by \mathcal{A}_2 . These properties of MI contrast with those of traditional coefficients of correlation, which measure linear relationships rather than

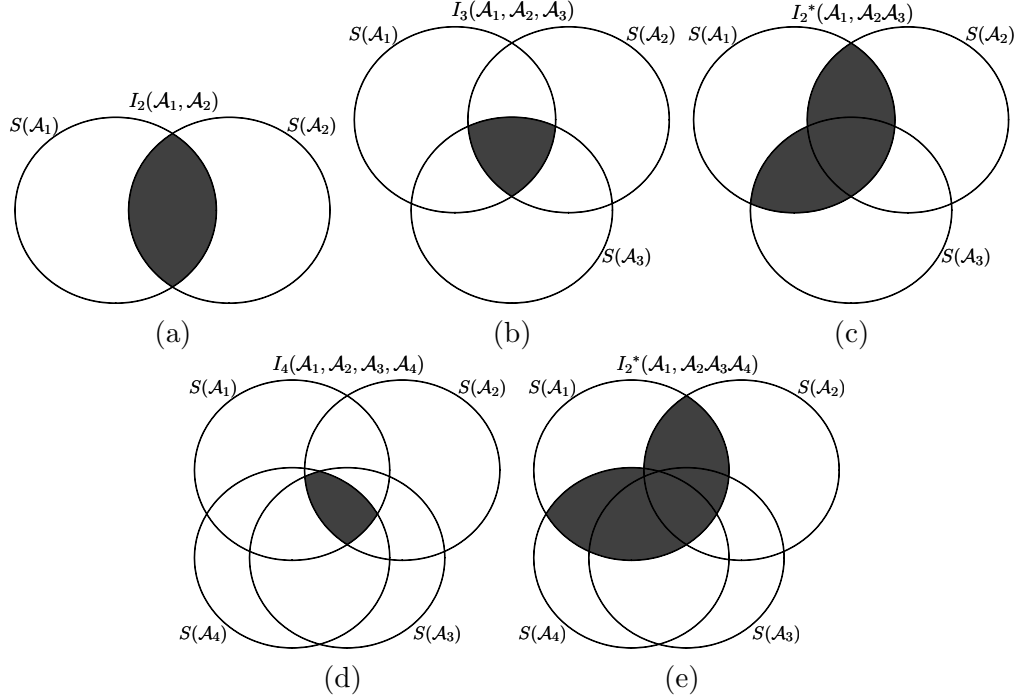


Figure 2.2: Venn diagrams with the shaded regions representing the information theoretic quantities (a) MI (I_2) of 2 random variables or images \mathcal{A}_1 and \mathcal{A}_2 , (b) higher order MI (I_3) of 3 random variables \mathcal{A}_1 , \mathcal{A}_2 , and \mathcal{A}_3 , and (c) CMI (I_2^*) between \mathcal{A}_1 and an ensemble of two simultaneously observed random variables, $\mathcal{A}_2\mathcal{A}_3$. (d) With the addition of a fourth random variable \mathcal{A}_4 , the fourth order MI (I_4) decreases from the quantity in (b). (e) CMI with the addition of a fourth random variable increases as the knowledge of \mathcal{A}_4 further reduces the uncertainty about \mathcal{A}_1 .

a more general functional relationship [39]. For these reasons, MI is considered the most robust intensity-based similarity measure for multimodal images [39]. Note however, that the spatial organization of pixels is not taken into account with MI.

2.2.2 Generalized Mutual Information

Generalized (higher-order) MI (GMI) may also be expressed as a summation of multiple entropy terms, and represents the intersecting or “interaction” [73, 74] information of multiple images or variables. The MI of three random variables \mathcal{A}_1 , \mathcal{A}_2 , and \mathcal{A}_3 can be written in terms of Shannon entropy as [39],

$$I_3(\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3) = S(\mathcal{A}_1) + S(\mathcal{A}_2) + S(\mathcal{A}_3) - S(\mathcal{A}_1\mathcal{A}_2) - S(\mathcal{A}_1\mathcal{A}_3) - S(\mathcal{A}_2\mathcal{A}_3) + S(\mathcal{A}_1\mathcal{A}_2\mathcal{A}_3). \quad (2.5)$$

The formula may be generalized for n random variables $\mathcal{A}_1, \dots, \mathcal{A}_n$ as [75],

$$I_n(\mathcal{A}_1, \dots, \mathcal{A}_n) = \sum_{s=1}^n (-1)^{s+1} \sum_{i_1 < \dots < i_s} S(\mathcal{A}_{i_1} \cdots \mathcal{A}_{i_s}), \quad (2.6)$$

for all combinations $\{i_1, \dots, i_s\} \in \{1, \dots, n\}$. GMI which is used to refer to $I_n, n \geq 2$, can only decrease with the addition of variables, as information shared by all variables decreases. For example, Figures 2.2(a), (b), and (d) illustrate I_n for $n \in \{2, 3, 4\}$ semi-independent variables. The addition of \mathcal{A}_3 and \mathcal{A}_4 has the effect of decreasing the common information from the quantity I_2 as shown in Figure 2.2(a).

Since an intensity image and any textural feature representations are in implicit alignment and semi-independent, they may be considered as a single high dimensional observation rather than multiple independent variables since such images are related. However, where the texture feature \mathcal{A}_3 is derived from and in alignment with \mathcal{A}_2 , GMI in the example shown in Figure 2.2(b) would treat \mathcal{A}_3 as independent from \mathcal{A}_2 . GMI is thus useful only in a registration scenario where each image is intended to move independently. Further, the formulation of GMI shown in Equation 2.6 is not necessarily a nonnegative quantity with a clear interpretation for registration. Despite these inherent limitations of GMI, equivalent formulations of GMI expressed in the form of Kullback-Leibler distance between the joint gray level distribution and the product of the marginal distributions [39] are often used to incorporate image-derived features [61].

2.2.3 Combined Mutual Information (CMI)

The goal of CMI is not to measure the intersecting information of multiple random variables or images $(\mathcal{A}_1, \dots, \mathcal{A}_n)$, as with GMI, but to quantify the combined information content encoded by one multivariate observation (e.g. $\mathcal{A}_1 \cdots \mathcal{A}_n$) with respect to another (e.g. $\mathcal{B}_1 \cdots \mathcal{B}_n$). In the most simple case, the CMI (I_2^*) that a single image \mathcal{A}_1 shares with an ensemble of two semi-independent images, \mathcal{B}_1 and \mathcal{B}_2 , is defined by the equation,

$$I_2^*(\mathcal{A}_1, \mathcal{B}_1 \mathcal{B}_2) = S(\mathcal{A}_1) + S(\mathcal{B}_1 \mathcal{B}_2) - S(\mathcal{A}_1 \mathcal{B}_1 \mathcal{B}_2). \quad (2.7)$$




Figure 2.3: (a) A plot of CMI of intensity image \mathcal{A}_1 with 16 $(n - 1)$ dimensional ensembles comprising intensity image \mathcal{A}_2 and $(n - 2)$ corresponding textural feature images, demonstrates that CMI is bounded by $S(\mathcal{A}_1)$ (horizontal line), the information content of \mathcal{A}_1 . (b) Plot of the sum of pairwise MI (SMI) of \mathcal{A}_1 with each of $\mathcal{A}_2, \dots, \mathcal{A}_n$ aggregates redundant information about \mathcal{A}_1 and is unbounded. (c) Plot of GMI of $\mathcal{A}_1, \dots, \mathcal{A}_n$ illustrates information shared by all images becomes less as additional features are considered. The upper bound on the value of CMI in (a) suggests that it is a more intuitive formulation for integrating multiple semi-independent feature images derived from \mathcal{A}_2 , compared to either (b) linear combination of multiple MI measurements or (c) GMI. For the synthetic example considered in this figure, intensity images \mathcal{A}_1 and \mathcal{A}_2 are in correct alignment.

This formulation utilizes higher order joint entropy calculations to measure only the unique (non-redundant) information provided by an additional image, \mathcal{B}_2 , about \mathcal{A}_1 . Significantly, by considering \mathcal{B}_1 and \mathcal{B}_2 as simultaneously measured semi-independent variables in the single multidimensional ensemble $\mathcal{B}_1\mathcal{B}_2$, any dependence that exists between \mathcal{B}_1 and \mathcal{B}_2 is discounted. The generalized form of CMI of the n dimensional ensemble $\epsilon_n^{\mathcal{A}} = \mathcal{A}_1 \cdots \mathcal{A}_n$ with the m dimensional ensemble $\epsilon_m^{\mathcal{B}} = \mathcal{B}_1 \cdots \mathcal{B}_m$ is then written here as,

$$I_2^*(\epsilon_n^{\mathcal{A}}, \epsilon_m^{\mathcal{B}}) = S(\epsilon_n^{\mathcal{A}}) + S(\epsilon_m^{\mathcal{B}}) - S(\epsilon_n^{\mathcal{A}} \epsilon_m^{\mathcal{B}}). \quad (2.8)$$

To compare GMI ($I_3(\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3)$ in Figure 2.2(b)) with CMI, the quantity in Equation 2.7 is presented in Figure 2.2(c) in terms of only \mathcal{A} as $I_2^*(\mathcal{A}_1, \mathcal{A}_2\mathcal{A}_3)$, which is the CMI of \mathcal{A}_1 with the ensemble $\mathcal{A}_2\mathcal{A}_3$. Unlike GMI, CMI can be seen to increase with the addition of the new variable \mathcal{A}_3 that contains unique information. Figures 2.2(d) and (e) extend this example to a fourth random variable \mathcal{A}_4 on GMI ($I_4(\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3, \mathcal{A}_4)$) of all four variables and on CMI ($I_2^*(\mathcal{A}_1, \mathcal{A}_2\mathcal{A}_3\mathcal{A}_4)$) of \mathcal{A}_1 with the ensemble $\mathcal{A}_2\mathcal{A}_3\mathcal{A}_4$.

Again it can be seen that additional variables only reduce GMI and increase CMI.

Now consider the terms ϵ_n^A and ϵ_m^B in Equation 2.8 to represent two image ensembles, each comprising an intensity image and multiple (n, m) intensity-derived feature images. To compare the behavior of CMI and GMI for such variable length ensembles, we again restrict the comparison to components of ϵ_n^A , and consider the quantity $I_2^*(\mathcal{A}_1, \mathcal{A}_2 \cdots \mathcal{A}_{n+1})$ for CMI, and $I_n(\mathcal{A}_1, \dots, \mathcal{A}_{n+1})$ for GMI in Figure 2.3 for multiple values of n . In addition to CMI and GMI, Figure 2.3 also compares the quantity obtained by summation of MI (SMI) values for pairs of images ($\sum_{i=2}^{n+1} I_2(\mathcal{A}_1, \mathcal{A}_i)$). SMI has also been previously utilized for combining multiple sources of information [60]. As seen in Figure 2.3(a), as the ensemble grows with the introduction of images, CMI approaches an asymptote equal to $S(\mathcal{A}_1)$, which is the total information content or uncertainty about \mathcal{A}_1 . On the other hand, SMI (Figure 2.3(b)) over image pairs $((\mathcal{A}_1, \mathcal{A}_2), \dots, (\mathcal{A}_1, \mathcal{A}_n))$ increases in an unbounded fashion as intersecting information provided by $\mathcal{A}_2, \mathcal{A}_3, \dots, \mathcal{A}_n$ about \mathcal{A}_1 is recounted. Conversely, GMI (Figure 2.3(c)) drops to zero as the information content shared by each $\mathcal{A}_1, \dots, \mathcal{A}_n$ becomes too small to measure. In summary, CMI has the desirable property that $I_2^*(\mathcal{A}_1, \mathcal{A}_2 \cdots \mathcal{A}_n) \leq S(\mathcal{A}_1) = I_2(\mathcal{A}_1, \mathcal{A}_1)$.

2.3 Combined Feature Ensemble Mutual Information (COFEMI)

2.3.1 Brief Overview of COFEMI

A flowchart showing the operation of the COFEMI registration technique for alignment of a moving target image \mathcal{B} with a stationary template image \mathcal{A} is shown in Figure 2.4. Briefly, these steps include,

Step 1. Feature Extraction (Section 2.3.2): COFEMI calculates two extensive sets of n feature images from each intensity image \mathcal{A}, \mathcal{B} for which additional structural description is required to improve multimodal image similarity.

Step 2. Feature Ensemble Selection (Section 2.3.3): Using the information maximization algorithm described in Section 2.3.3, COFEMI constructs from the feature

We define template and target intensity images $\mathcal{A} = (C, f^{\mathcal{A}}(c))$ and $\mathcal{B} = (C, f^{\mathcal{B}}(c))$, both defined on some coordinate system C and with intensity values $f^{\mathcal{A}}(c)$ and $f^{\mathcal{B}}(c)$ at every location $c \in C$. Image operators Φ_β , $\beta \in \{1, \dots, n\}$, are defined to calculate from \mathcal{A} and \mathcal{B} the n feature scenes denoted by $\mathcal{A}_\beta = (C, f_\beta^{\mathcal{A}})$, and $\mathcal{B}_\beta = (C, f_\beta^{\mathcal{B}})$, respectively. Hence $f_\beta^{\mathcal{B}}(c)$ is the resulting feature value at pixel c associated with application of feature operator Φ_β to \mathcal{B} . The optimal feature ensembles, denoted by $\pi_l^{\mathcal{A}}$ and $\pi_k^{\mathcal{B}}$, of l and k images are uniquely composed such that $\pi_l^{\mathcal{A}} = \mathcal{A}_{\gamma_1} \mathcal{A}_{\gamma_2} \dots \mathcal{A}_{\gamma_k}$ where $\{\gamma_1, \gamma_2, \dots, \gamma_l\} \in \{1, \dots, n\}$, and $\pi_k^{\mathcal{B}} = \mathcal{B}_{\alpha_1} \mathcal{B}_{\alpha_2} \dots \mathcal{B}_{\alpha_k}$ where $\{\alpha_1, \alpha_2, \dots, \alpha_k\} \in \{1, \dots, n\}$. Hence, registration of \mathcal{A} and \mathcal{B} is achieved by optimizing some spatial transformation \mathbf{T} via maximization of the CMI between the ensembles $\mathcal{A}\pi_l^{\mathcal{A}}$ and $\mathcal{B}\pi_k^{\mathcal{B}}$ to obtain \mathbf{T}_{COFEMI} , which may be stated formally as,

$$\mathbf{T}_{COFEMI} = \underset{\mathbf{T}}{\operatorname{argmax}} [I_2^*(\mathcal{A}\pi_l^{\mathcal{A}}, \mathbf{T}(\mathcal{B}\pi_k^{\mathcal{B}}))] . \quad (2.9)$$

A list of the main notations employed in this study is given in Table 1.1.

2.3.2 Feature Extraction

The textural operators Φ_β , $\beta \in \{1, \dots, n\}$ include both linear and non-linear image operations, and are drawn from 3 general classes of texture features including (i) gradient, (ii) first order statistical, and (iii) second order statistical texture features. A description of the 311 features considered here as part of COFEMI is given in Table 2.2.

Gradient Features

These features include 144 steerable and 11 non-steerable linear gradient operators, whereby steerable filters refer to operators specified by an orientation parameter.

Non-steerable Gradient Features

Eleven non-steerable gradient features were obtained using the Sobel and Kirsch edge operators [76], and discrete first order spatial derivatives. Three Sobel linear operators

Feature Class	Parameters	Total Features
Gradient (Sobel, Kirsch, Discrete Derivatives, Gabor)	scale ($ \sigma = 3$) orientation ($ \lambda = 6$) frequency ($ \varphi = 8$)	155
First Order Statistics (Mean, Median, Standard Deviation, Range)	window size ($ \kappa = 3$)	12
Second Order Statistics (Energy, Entropy, Inertia, Correlation, Inverse Difference Moment, Information Correlation (two), Sum Average, Sum Variance, Sum Entropy, Difference Average, Difference Variance, Difference Entropy, Shade, Prominence, Variance)	window size ($ \kappa = 3$) number of bins ($ L = 3$)	144

Table 2.2: Description of 311 image derived features used by COFEMI and corresponding parameters for each texture feature class.

were defined to detect the strength of horizontal, vertical, and diagonal edges. Four Kirsch linear operations were utilized as approximations to the strength of edges normal to lines oriented $0, \frac{\pi}{4}, \frac{\pi}{2}$, and $\frac{3\pi}{4}$ radians from the horizontal. The Sobel and Kirsch filter responses were obtained by convolution of \mathcal{A} and \mathcal{B} with each linear operator. Four additional edge estimates are obtained by combinations of three simple first order spatial derivative operators, including the vertical, horizontal, and diagonal spatial derivatives. The derivative filter responses were again obtained by convolution of \mathcal{A} and \mathcal{B} with the linear operators. The gradient magnitude of the image was calculated as the L2 norm of the horizontal and vertical derivative responses at each pixel.

Steerable Gradient Features

Gabor gradient operators [77] comprise the steerable class of gradient calculations. The linear Gabor kernel used here is defined by the product of a 2D Gaussian function with a cosine in 2D space $c = (x, y)$.

$$\Phi_{Gab} = e^{\frac{(x'+y')}{\sigma}} \cos(2\frac{\pi}{\lambda}x'), \quad (2.10)$$

where λ is the wavelength of the sinusoid and hence controls the spatial frequency (scale) of the oscillations. The width of the Gaussian envelope σ is used to define filters of size $\kappa \times \kappa$ where $\kappa = 5\sigma$. Orientation of the filter is affected by φ through the following coordinate transformations $x' = x \cos \varphi + y \sin \varphi$ and $y' = x \sin \varphi - y \cos \varphi$, where $1 \leq x \leq \kappa$, and $1 \leq y \leq \kappa$. A set of 144 Gabor filters were generated for $\sigma \in \{1, 1.6, 3\}$ ($\kappa \in \{5, 7, 15\}$), $\lambda \in \{\frac{\pi}{2\sqrt{2}}, \frac{\pi}{4}, \frac{\pi}{4\sqrt{2}}, \frac{\pi}{8}, \frac{\pi}{8\sqrt{2}}, \frac{\pi}{16}\}$, and $\varphi \in \{0, \frac{\pi}{8}, \frac{\pi}{4}, \frac{3\pi}{8}, \frac{\pi}{2}, \frac{5\pi}{8}, \frac{3\pi}{4}, \frac{7\pi}{8}\}$. A total of 144 Gabor feature images for each \mathcal{A} and \mathcal{B} were obtained by convolution of Φ_{Gab} with \mathcal{A} and \mathcal{B} for each of σ, λ, φ .

The application of each of the above steerable and non-steerable gradient operators Φ_β , $\beta \in \{1, \dots, 155\}$ to \mathcal{A} and \mathcal{B} produces corresponding gradient scenes of \mathcal{A} and \mathcal{B} .

First Order Statistical Features

Four first order statistical feature operations (mean, median, standard deviation, and range of image intensities) are defined from the gray level distributions of pixels within $\kappa \times \kappa$ local neighborhoods of pixels $\mathcal{N}_\kappa(c)$ centered about each $c \in C$ and of size $\kappa \in \{3, 5, 7\}$. The mean feature operation was defined as the average gray level of all pixels $d \in \mathcal{N}_\kappa(c)$. Each of the median, standard deviation, and range feature operations are defined via pixel-wise, non-linear calculations using the same κ -neighborhoods to sample local gray level distributions. For example, the range value at a pixel c is defined as $\max_{d \in \mathcal{N}_\kappa(c)} [f(d)] - \min_{d \in \mathcal{N}_\kappa(c)} [f(d)]$. These four feature operations, calculated for all three κ values, (Φ_β , $\beta \in \{156, \dots, 167\}$) when applied to \mathcal{A} and \mathcal{B} result in 12 corresponding feature scenes each.

Second Order Statistical Features

We define 144 unique second order statistical feature operations from the class of Haralick features [78] for characterization of spatial gray level dependence. An $L \times L$ spatial gray level co-occurrence matrix $P_\kappa(c)$, where $L = \max_{c \in C} [f(c)]$, is first defined. The value at any location $[u, v]$ in the matrix $P_\kappa(c)$ represents the frequency with which two distinct pixels c and $d \in \mathcal{N}_\kappa(c)$ with associated image intensities $f(c) = u$, $f(d) = v$ are adjacent (within the same 8-neighborhood within $\mathcal{N}_\kappa(c)$). The 16 Haralick features

listed in Table 2.2 were calculated from $P_\kappa(c)$ at every $c \in C$, for $\kappa \in \{3, 5, 8\}$, and $L \in \{64, 128, 256\}$ to obtain a total of 144 second order statistical feature operators $(\Phi_\beta, \beta \in \{168, \dots, 311\})$.

2.3.3 Feature Ensemble Selection

The feature ensemble selection method applied to generate ensembles $\pi_l^{\mathcal{A}}$ and $\pi_k^{\mathcal{B}}$ corresponding to intensity images \mathcal{A} and \mathcal{B} is based on the interpretation of joint entropy $S(\mathcal{B}\pi_k^{\mathcal{B}})$ of an ensemble of images $\mathcal{B}\pi_k^{\mathcal{B}}$ as the total information content of the ensemble. Thus, an optimal ensemble $\pi_k^{\mathcal{B}}$ of length k for CMI-based registration would contain maximum joint entropy $S(\mathcal{B}\pi_k^{\mathcal{B}})$. Further, it follows that maximization of $S(\mathcal{B}\pi_k^{\mathcal{B}})$ for an ensemble of length $k \ll n$ also simultaneously enforces the criteria of feature distinctiveness. Determining the ensemble of length k that is optimal in terms of $S(\mathcal{B}\pi_k^{\mathcal{B}})$ is however not a trivial task. A brute force approach to selecting π_k for $k = 5$ from $n = 311$ features would involve computing the joint entropy of $\binom{311}{5}$ ensembles (over 10^{10}). Consequently, we present the following feature selection algorithm (*CMIfeatures*), which iteratively incorporates the most distinct features (in terms of joint entropy) into an optimal feature ensemble.

Algorithm *CMIfeatures*

Input: \mathcal{B} , k , $\mathcal{B}_\beta, \beta \in \{1, \dots, n\}$.

Output: $\pi_k^{\mathcal{B}}$.

begin

0. Initialize $\pi_k^{\mathcal{B}}$ and Ω as empty queues;
1. *for* $i = 1$ to k
2. *for* $\beta = 1$ to n
3. *If* \mathcal{B}_β is present *then*
4. Insert \mathcal{B}_β into $\pi_k^{\mathcal{B}}$;
5. $\Omega[\beta] = S(\mathcal{B}\pi_k^{\mathcal{B}})$;
6. Remove \mathcal{B}_β from $\pi_k^{\mathcal{B}}$;
7. *endif*;
8. *endfor*;
9. $o = \operatorname{argmax}_\beta [\Omega[\beta]]$;
10. Insert \mathcal{B}_o into $\pi_k^{\mathcal{B}}$;
11. *endfor*;

end

The central idea of the algorithm is to compute the joint Shannon entropy between

\mathcal{B} and $\pi_k^{\mathcal{B}}$ when a single feature image \mathcal{B}_β is inserted into $\pi_k^{\mathcal{B}}$ in a pairwise fashion over all $\beta \in \{1, \dots, n\}$. At the end of each set of n iterations, the feature image \mathcal{B}_β contributing the most information over the previously chosen features and \mathcal{B} is inserted into the queue $\pi_k^{\mathcal{B}}$ and removed from further consideration. The process is repeated until $k - 1$ additional unique feature images are inserted into the ensemble $\pi_k^{\mathcal{B}}$.

2.3.4 Registration by Maximization of CMI

Following feature extraction and selection of ensembles $\pi_l^{\mathcal{A}}$ and $\pi_k^{\mathcal{B}}$, the CMI multivariate similarity measure is used to incorporate the non-redundant information encoded by the implicitly registered feature images within the ensembles. As such, the optimal image transformation \mathbf{T}_{COFEMI} required to align \mathcal{B} with \mathcal{A} is found by $\mathbf{T}_{COFEMI} = \operatorname{argmax}_{\mathbf{T}} I_2^*(\mathcal{A}\pi_l^{\mathcal{A}}, \mathbf{T}(\mathcal{B}\pi_k^{\mathcal{B}}))$ (denoted by $\Psi^{COFEMI}(\mathcal{A}, \mathcal{B})$ for simplicity), and the registered target $\mathcal{B}^r = \mathbf{T}(\mathcal{B})$ is obtained. A Nelder-Mead simplex algorithm [79] was used in our tests to optimize parameters for rotation, translation and scaling in \mathbf{T} .

2.4 Quantitative Evaluation Methods

Quantitative evaluation is performed using three methods: (1) overlap and boundary alignment of anatomical structures (Section 2.4.1), (2) TRE measured over multiple registration trials (Section 2.4.2), and (3) an optimizer-independent evaluation protocol presented by Skerl (Section 2.4.3). Sections 2.4.1 and 2.4.2 describe the methods involving registration trials (maximizing Equation 2.1 for different images and misalignments), while Section 2.4.3 describes the methods for characterizing similarity measure Ψ without optimization of \mathbf{T} (Equation 2.1).

2.4.1 Evaluation using Anatomical Region Labels

When ground truth for correct alignment is not available, registration performance can be evaluated via pre-segmented anatomical structures to calculate (1) alignment of the boundaries and (2) overlap of the segmented region. We denote as $G(\mathcal{A})$ and $G(\mathcal{B})$, the sets of pixels corresponding to prominent regions on the stationary image \mathcal{A} and on the

target image \mathcal{B} . $G(\mathcal{B}^r)$ denotes the region $G(\mathcal{B})$ on the registered image \mathcal{B}^r . The region overlap ratio (*ROR*) [80] of corresponding regions $G(\mathcal{A})$ and $G(\mathcal{B}^r)$ is calculated as,

$$\phi^{ROR}(G(\mathcal{A}), G(\mathcal{B}^r)) = \frac{|G(\mathcal{A}) \cap G(\mathcal{B}^r)|}{|G(\mathcal{A}) \cup G(\mathcal{B}^r)|}. \quad (2.11)$$

where $|G|$ denotes the cardinality of set G . We also denote as $G^b(\mathcal{A})$ and $G^b(\mathcal{B}^r)$ as the boundary pixels on $G(\mathcal{A})$ and $G(\mathcal{B}^r)$, respectively. The Hausdorff distance metric [81] (ϕ^{HD}) is calculated between sets $G^b(\mathcal{A})$ and $G^b(\mathcal{B}^r)$ to quantify the accuracy of alignment of $G(\mathcal{A})$ and $G(\mathcal{B})$ as in Equation 2.12.

$$\phi^{HD}(G^b(\mathcal{A}), G^b(\mathcal{B}^r)) = \max \left[\max_{c \in G^b(\mathcal{A})} \left[\min_{d \in G^b(\mathcal{B}^r)} \|c - d\| \right], \max_{d \in G^b(\mathcal{B}^r)} \left[\min_{c \in G^b(\mathcal{A})} \|c - d\| \right] \right] \quad (2.12)$$

where $\|\cdot\|$ denotes the L2 distance in \mathbb{R}^2 space. A low Hausdorff distance measure $\phi^{HD}(G^b(\mathcal{A}), G^b(\mathcal{B}^r))$ for a registration result indicates a low degree of misalignment along $G^b(\mathcal{A})$ and $G^b(\mathcal{B}^r)$ and hence reflects a high degree of alignment between \mathcal{A} and \mathcal{B}^r . The Mean Absolute Distances (MAD) ϕ^{MAD1} and ϕ^{MAD2} , which represent the average of the shortest distances from each point on $G^b(\mathcal{A})$ to $G^b(\mathcal{B}^r)$ and from $G^b(\mathcal{B}^r)$ to $G^b(\mathcal{A})$, respectively, are also calculated. We define $\phi^{MAD}(G^b(\mathcal{A}), G^b(\mathcal{B}^r))$ as the average of ϕ^{MAD1} and ϕ^{MAD2} .

2.4.2 Evaluation via Target Registration Error

TRE values are obtained by applying a known deformation (\mathbf{T}^{ap}) away from known alignment, followed by execution of the registration algorithm to recover the initial correct alignment via a corrective deformation (\mathbf{T}^{co}). The TRE measure is thus calculated for each pixel $c \in C$ as the distance between the common undeformed coordinates (c) in \mathcal{B} and the coordinates that have been moved by the applied and recovered transformations ($c^r = \mathbf{T}^{co}(\mathbf{T}^{ap}(c))$ in \mathcal{B}^r) by the equation,

$$TRE(c) = \|c - c^r\|. \quad (2.13)$$

The average TRE over all $c \in C$ in an image pair $(\mathcal{B}, \mathcal{B}^r)$ is then calculated as,

$$\phi^{TRE}(\mathbf{T}^{co}) = \frac{1}{|C|} \sum_{c \in C} TRE(c). \quad (2.14)$$

2.4.3 Quantitative Evaluation via Skerl Evaluation Protocol

To evaluate the efficacy of a similarity measure Ψ for rigid registration in the absence of confounding factors such as strategy for optimization of \mathbf{T} in the maximization of Equation (2.1), we implement the evaluation protocol presented in [8]. The Skerl evaluation protocol, which also requires images for which a “gold standard” for alignment is known, calculates 5 unique properties to characterize Ψ . The spirit of the method is to systematically sample a normalized linear transformation space at regular intervals, and then to apply the transformations in order to characterize Ψ at various degrees of misalignment of an image pair $(\mathcal{B}, \mathcal{B}^r)$. The space of image transformations is sampled at M evenly spaced points $m_s \in \{-M/2, \dots, M/2\}$, or positions, along each of N total lines $n_s \in \{1, \dots, N\}$ drawn in the hyperspace spanned by the transformation parameters. The hyperspace used here is 3-dimensional, where each position (n_s, m_s) represents a transformation specified by $(\theta_{n_s, m_s}^1, \theta_{n_s, m_s}^2, \theta_{n_s, m_s}^3)$ for rotation, horizontal translation, and vertical translation.

A parameter space line $n_s \in \{1, \dots, N\}$ is illustrated in Figure 2.5(a), where a PD-w MR image is shown transformed from its original position by the two transformations at the extrema of parameter line n_s . At each point m_s along this line, the 3 parameters $(\theta_{n_s, m_s}^1, \theta_{n_s, m_s}^2, \theta_{n_s, m_s}^3)$ are smoothly varied to generate a range of misalignment from ground truth at $m_s = 0$ (the black dot in Figure 2.5(a)). A stationary T2-w MR image of the same brain slice is shown in Figure 2.5(b) relative to the same deformed PD-w MRI slices in Figure 2.5(a). Next, by computing Ψ for the two images at each (n_s, m_s) , a total of N similarity “traces” each comprising M individual similarity values are obtained from which fitness of Ψ is characterized. The trace of Ψ for line n_s is shown in bold in Figure 2.5(c). The primary quantities that are extracted from each of the N traces are the location and number of local minima, and the slopes of the traces, shown via the dashed and the light weight lines, respectively, in Figure 2.5(c). A close-up of

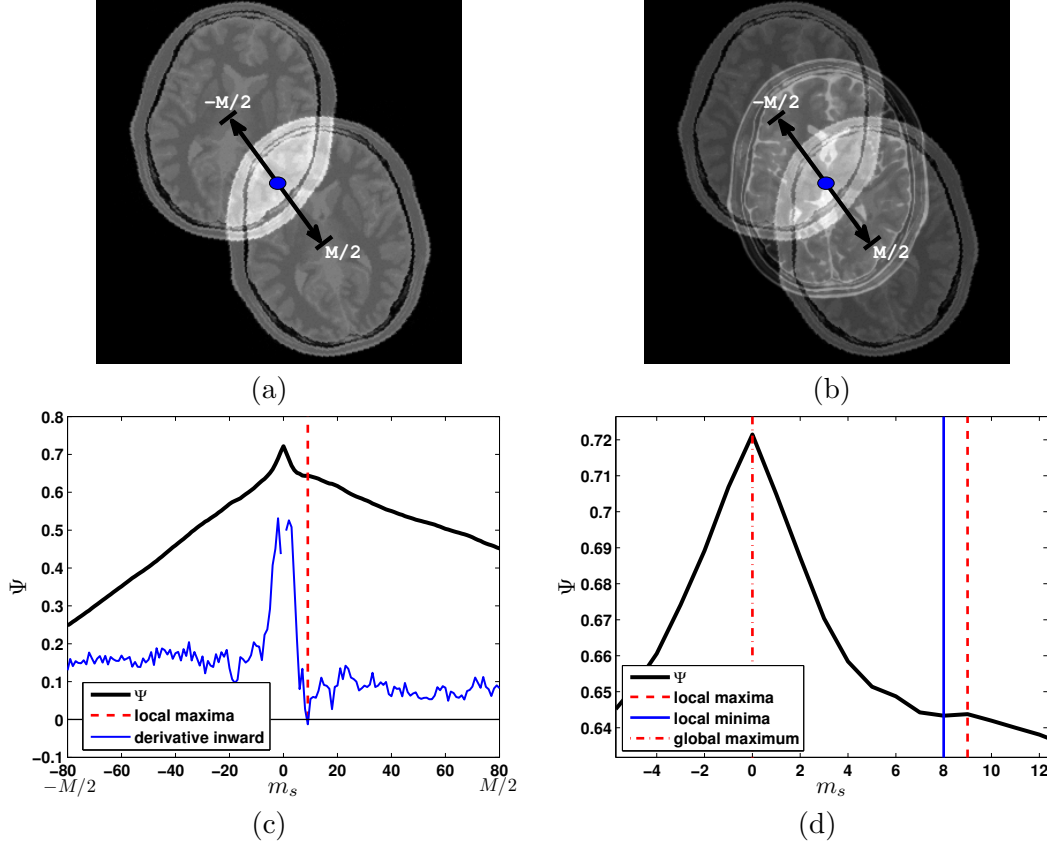


Figure 2.5: The quantitative evaluation protocol proposed by Skerl involves systematically transforming an image along “lines” in transformation parameter hyperspace. (a) A PD-w MR image shown at the two extrema of deformations ($m_s = -M/2$ and $m_s = M/2$) on a line n_s . (b) The corresponding T2-w MR image shown in its stationary position superposed on the deformed PD-w MR images in (a). The T2-w MRI remains stationary, while (c) values of image similarity (Ψ) are calculated at positions $m_s \in \{-M/2, \dots, M/2\}$ along each line n_s . Properties of Ψ are derived from the derivative (light blue) and local maxima (dashed vertical) along the profile of a measure’s trace. (d) Location and number of local minima (solid vertical) and the global maximum (dash-dotted vertical) are also used in the calculation of the properties described in Table 2.3.

Property	Description
Accuracy (ACC), ϕ^{S1}	Distance from the observed global maximum to ground truth. Captures the extent to which the maximum of Ψ corresponds to correct alignment; lower ACC is better.
Capture Range (CR), ϕ^{S2}	Distance from the global maximum to the nearest local minimum. Captures the extent of initial misalignment that can be present and still achieve correct alignment. Higher CR is better.
Distinctiveness of the Optimum (DO), ϕ^{S3}	Number of local minima with respect to distance from the global maximum. “Peakedness” of Ψ about the global maximum. Higher DO is better.
Number of Minima (NOM), ϕ^{S4}	Number of local minima with respect to distance from the global maximum. Possible locations for Ψ to find a local optimum instead of the global optimum. Lower NOM is better.
Risk of Non-convergence (RON), ϕ^{S5}	Integral of the slopes of the similarity measure that are positive away from the global maximum. Reflects the likelihood that an optimization scheme will move away from the global maximum. Lower RON is better.

Table 2.3: Description of properties ϕ^{S1} – ϕ^{S5} of Ψ calculated by the Skerl protocol [8] for evaluation in the context of rigid registration. The reader is referred to [8] for additional details on the Skerl evaluation protocol.

the peak of Ψ is shown in Figure 2.5(d), where a local minimum and its associated local maximum can be seen at 8 mm from the ground truth alignment.

A total of five properties proposed in [8] to quantify the ability of Ψ to drive rigid registration include accuracy, capture range, distinctiveness of the global optimum, number of local minima, and risk of non-convergence, and are calculated for each of the N similarity measure traces (Table 2.3). We refer the reader to [8] for details on the formulation of these properties as well as the protocol for defining and sampling the normalized parameter space.

2.5 Results of Automated Affine Registration

2.5.1 Real and Synthetic Image Data

Table 2.4 summarizes the real and synthetic image data sets used for quantitative evaluation of COFEMI, and the availability and type of ground truth for each data set, and the evaluation techniques (Section 2.4) that are applicable for each data set.

Image Data Set 1: The image data set denoted by \mathcal{D}^1 represents multiprotocol

Data Set	Protocols	Parameters	Ground Truth Availability	Evaluation Measures
\mathcal{D}^1 (synthetic brain MRI)	T1-w ($\mathcal{T}1$), T2-w ($\mathcal{T}2$), PD-w (\mathcal{P})	30 2D sections of a normal study, 256×256 pixels	Brain gray matter regions $G(\mathcal{T}1)$ and $G(\mathcal{P})$. Exact misalignment \mathbf{T}^{ap} between each $\mathcal{T}1$, $\mathcal{T}2$ and \mathcal{P} known.	Gray matter region overlap $\phi^{ROR}(G(\mathcal{T}1), G(\mathcal{P}))$. Target registration error $\phi^{TRE}(\mathbf{T}^{co})$ of whole brain. Skerl's similarity measure properties $\phi^{S1} - \phi^{S5}$.
\mathcal{D}^2 (clinical prostate)	histology (\mathcal{H}), <i>ex vivo</i> MR (\mathcal{M})	2 3D volumes comprising 26 2D sections, 256×256 pixels	Prostate regions $G(\mathcal{M})$ and $G(\mathcal{H})$, and prostate contours $G^b(\mathcal{M})$ and $G^b(\mathcal{H})$. No alignment truth.	Hausdorff distance of prostate boundary $\phi^{HD}(G^b(\mathcal{M}^r), G^b(\mathcal{H}))$. Region overlap of prostate $\phi^{ROR}(G(\mathcal{M}^r), G(\mathcal{H}))$.

Table 2.4: Description of image data sets used in this study for evaluation of COFEMI.

synthetic brain MRI (T1-w, T2-w, and PD-w and denoted via scenes $\mathcal{T}1$, $\mathcal{T}2$ and \mathcal{P} , respectively) from BrainWeb [82]. Each 3D MRI brain volume has dimensions $181 \times 217 \times 181$ with voxel size of 1mm^3 . Ground truth for correct alignment between $\mathcal{T}1$, $\mathcal{T}2$, and \mathcal{P} is implicit in the simulated data, allowing for quantitative evaluation of different registration schemes as described in Sections 2.4.2 and 2.4.3.

Image Data Set 2: The image data set denoted by \mathcal{D}^2 contains a total of 26 2D 4 Tesla *ex vivo* MR and digitized histology images from two prostate patient studies. Details of the data acquisition are described in [28]. Haematoxylin and Eosin stained 2D digitized histological images, denoted by \mathcal{H} , were generated by conversion to gray scale and resampling from their original resolution of 0.07 mm/pixel (approximately 2000×2000 pixels) to match the resolution of the 2D MRI slices (256×256 pixels), denoted by \mathcal{M} . Since there were roughly twice the number of MRI slices as there were histological sections for each patient study, the MRI sections corresponding to the histological sections were identified by an expert by visually inspecting \mathcal{H} and \mathcal{M} . Since only a total of 26 histological sections were available from the 2 glands, only 26 corresponding MRI slices were considered.

2.5.2 Experiment 1: Synthetic Multiprotocol BrainWeb Data

Introducing Different Levels of Bias Field and Noise

Simulated noise and bias field inhomogeneity were artificially introduced into \mathcal{P} by,

$$f^*(c) = (f(c) \times g^b(c)) + \eta_\sigma(c), \quad (2.15)$$

where $g^b(c)$ and $\eta_\sigma(c)$ represent bias field and noise functions, respectively. Bias field $g^b(c) \in [0, 1]$ was obtained via the BrainWeb simulated MRI database, and noise $\eta_\sigma(c)$ sampled from Gaussian distributions with mean 0 and variances σ . From each of the 30 \mathcal{P} sections, an additional 30 scenes $\mathcal{P}_{b,\sigma}$ with $g^b(c)$, $b \in \{0, 0.5, 1.0, 1.5, 2.0, 2.5\}$ and $\eta_\sigma(c)$, $\sigma \in \{0, 1, \dots, 4\}$, were generated, for a total of 900 $\mathcal{P}_{b,\sigma}$ scenes.

Experiment 1a: TRE and Region Overlap

Since the $\mathcal{T}1$ and \mathcal{P} images generated by the BrainWeb MRI simulator [82] are in implicit alignment, evaluation of registration accuracy was performed as described in Section 2.4.2 by imposing a known deformation \mathbf{T}^{ap} to $\mathcal{P}_{b,\sigma}$, followed by execution of a registration scheme to determine the transformation \mathbf{T}^{co} required to recover the original alignment. Using 30 pairs of corresponding $\mathcal{T}1$ and $\mathcal{P}_{b,\sigma}$ images from set \mathcal{D}^1 , a total of 900 registration trials are performed over 5 different σ and 6 different b . In this experiment, three registration methods MI, COFEMI, and CMI^{grad} were individually applied to determine \mathbf{T}^{co} . Since common feature driven registration approaches have used image gradients [57, 60, 83], we investigate gradient-driven MI variant by denoting CMI^{grad} as the scheme using similarity measure $I_2^*(\mathcal{T}1, \mathcal{P} \mathcal{P}_{grad})$, where \mathcal{P}_{grad} is obtained from the gradient magnitude feature operator. MI registration was performed using 128 gray level bins, while CMI^{grad} utilized 48 gray levels. For COFEMI registration, a feature ensemble of length $k = 2$ was generated by the *CMIfeatures* algorithm, and 42 gray levels were used. The numbers of gray levels were determined empirically to be optimal for the respective registration schemes. Following registration by MI, CMI^{grad} , and COFEMI to determine \mathbf{T}^{co} and hence $\mathcal{P}_{b,\sigma}^r$, both $\phi^{TRE}(\mathbf{T}^{co})$ and $\phi^{ROR}(G(\mathcal{P}_{b,\sigma}^r), G(\mathcal{T}1))$




Figure 2.6: Synthetic multiprotocol MRI of the brain is used to compare MI and COFEMI for two different combinations of simulated inhomogeneity and noise. (a) $\mathcal{T}1$ and corresponding (b) $\mathcal{P}_{b,\sigma}$ with simulated misalignment, high inhomogeneity strength ($b = 2.5$) and low noise ($\sigma = 1$). The boundary of the skull on $\mathcal{T}1$ ((a),(e)) and $\mathcal{P}_{b,\sigma}$ ((b)-(d), (e)-(h)) is shown in green and red, respectively. (c) MI registration result $\mathcal{P}_{b,\sigma}^r$ is susceptible to inhomogeneity in the lower region of $\mathcal{P}_{b,\sigma}$, as evidenced by misalignment of red and green contours. (d) COFEMI is able to compensate for the artifacts and noise, as evidenced by the nearly perfect alignment of the red and green contours. (e)-(h) For pair of $\mathcal{T}1$ and $\mathcal{P}_{b,\sigma}$ with low inhomogeneity strength $b = 1.0$ and higher noise $\sigma = 2$.

were calculated, $G(\mathcal{T}1)$ and $G(\mathcal{P}_{b,\sigma}^r)$ denoting the sets of pixels representing the gray matter region on $\mathcal{T}1$ and \mathcal{P} respectively and where $G(\mathcal{P}_{b,\sigma}^r) = \mathbf{T}^{co}(\mathbf{T}^{ap}(G(\mathcal{P}_{b,\sigma})))$. Values for ϕ^{TRE} and ϕ^{ROR} are calculated for each of 900 registration trials.

Qualitative Results

Figure 2.6 presents a comparison of registration via MI and COFEMI for two different slice pairs from \mathcal{P} and $\mathcal{T}1$. In Figures 2.6(a) and (e), two different stationary $\mathcal{T}1$ images are shown with the boundary of the skull outlined in green. The corresponding $\mathcal{P}_{b,\sigma}$ sections subject to \mathbf{T}^{ap} and simulated noise and inhomogeneity are shown in Figures

2.6(b) and (f), where the red contour represents the skull boundary. The image intensities in Figures 2.6(b) and (f), are generated by Equation 2.15 using $(b = 2.5, \sigma = 1)$ and $(b = 1.0, \sigma = 2)$, respectively. Since accurate registration of \mathcal{P} to $\mathcal{T}1$ should result in close coincidence of the red and green boundaries, boundary alignment can be used to infer the accuracy of alignment. The results of MI-based registration of $\mathcal{P}_{b,\sigma}$ to $\mathcal{T}1$ are shown in Figures 2.6(c) and (g). Clearly, inhomogeneity in $\mathcal{P}_{b,\sigma}$ has caused misalignment at the bottom of the scene (Figure 2.6(c)). Yet with only modest inhomogeneity artifact ($b = 1.0$), as in Figure 2.6(f), MI-based registration is still unable to correct for the rotational misalignment, and registration error is apparent in Figure 2.6(g). On the other hand, the results of COFEMI-based registration (Figures 2.6(d) and (h)) demonstrate a near perfect alignment of the red and green boundaries, indicating an improved alignment over that attained by MI.

Quantitative Results

Figure 2.7 presents an analysis of the failure rates over all 900 registration trials for MI and COFEMI. A registration failure is defined as one where $\phi^{TRE} > 1.5$ mm. COFEMI is more robust to increased noise levels or bias field inhomogeneity artifacts than MI (Figure 2.7(a)). Failure of MI rigid registration is most clearly influenced by bias field strength, while higher noise levels in the presence of moderate to high inhomogeneity are also detrimental to MI-based registration. Figure 2.7(c) shows the differences in ϕ^{TRE} for MI – COFEMI, and highlights the robustness of COFEMI to artifacts compared with MI. Figure 2.7(d) shows a similar plot comparing differences in ϕ^{ROR} for COFEMI – MI since higher ϕ^{ROR} values correspond to better accuracy, and complements the ϕ^{TRE} results of Figure 2.7(c). Figures 2.7(e) and (f) present a similar comparison of ϕ^{TRE} and ϕ^{ROR} values from COFEMI and CMI^{grad} , where COFEMI again demonstrates its advantage when noise and inhomogeneity artifacts are present.

For each of the 30 simulated noise-inhomogeneity levels considered, the average values of ϕ^{TRE} from COFEMI were lower than ϕ^{TRE} from MI. The ϕ^{TRE} values for COFEMI and MI were compared using a paired t -test under the null hypothesis that the ϕ^{TRE} values are equal for COFEMI and MI. The p -values for each of the




Figure 2.7: (a) Registration failure rates of MI over 5 levels of noise and 6 levels of field inhomogeneity ($b \in \{0, 0.5, 1.0, 1.5, 2.0, 2.5\}, \sigma \in \{0, 1, \dots, 4\}$) in $\mathcal{P}_{b,\sigma}$, where failure is $\phi^{TRE} > 1.5\text{mm}$. (b) Failure rates of COFEMI. (c) Plot of $\Delta\phi^{TRE} = \phi_{\text{MI}}^{TRE} - \phi_{\text{COFEMI}}^{TRE}$. (d) Plot of $\Delta\phi^{ROR} = \phi_{\text{COFEMI}}^{ROR} - \phi_{\text{MI}}^{ROR}$. CMI^{grad} is similarly compared to COFEMI by (e) $\Delta\phi^{TRE} = \phi_{\text{CMI}^{grad}}^{TRE} - \phi_{\text{COFEMI}}^{TRE}$, and (f) $\Delta\phi^{ROR} = \phi_{\text{COFEMI}}^{ROR} - \phi_{\text{CMI}^{grad}}^{ROR}$. COFEMI is consistently more robust to the presence of bias field and noise artifacts, compared with MI or CMI^{grad}.

Noise/Field (σ, b)	$\sigma = 0$	$\sigma = 1$	$\sigma = 2$	$\sigma = 3$	$\sigma = 4$
$b = 0.0$	0.1595	0.1158	0.3052	0.1789	0.3545
$b = 0.5$	0.6270	0.1214	0.2936	0.2280	0.4607
$b = 1.0$	0.9914	0.3966	0.0867	0.5132	0.2329
$b = 1.5$	0.0196	0.1949	0.0056	0.2602	0.3895
$b = 2.0$	0.1508	0.0289	0.5195	0.1250	0.0180
$b = 2.5$	0.1255	0.2135	0.0446	0.0041	0.0290

Table 2.5: p -values for paired t -tests comparing ϕ^{TRE} values for MI and COFEMI for 30 different combinations of simulated noise (σ) and field inhomogeneity (b). These values correspond to the ϕ^{TRE} differences shown in Figure 2.7(c) averaged across 30 pairs of $\mathcal{T}1$ and $\mathcal{P}_{b,\sigma}^r$ for each b, σ .

30 noise-inhomogeneity levels are listed in Table 2.5, where instances corresponding to $p < 0.05$ (shown in bold) were considered to indicate statistically significant differences in ϕ^{TRE} between COFEMI and MI. Table 2.6 presents p -values for the comparison of ϕ^{TRE} values between COFEMI and CMI^{grad}. Complementary to ϕ^{TRE} , average $\phi^{ROR}(G(\mathcal{T}1), G(\mathcal{P}_{b,\sigma}^r))$ values from COFEMI were higher than those from MI for each of the 30 simulated noise-inhomogeneity combinations. The ϕ^{ROR} values for COFEMI and MI were similarly compared using a paired t -test under the null hypothesis that the ϕ^{ROR} values are equal for COFEMI and MI. Table 2.7 presents the p -values obtained over the 30 slice pairs for each of the 30 noise-inhomogeneity levels, where instances corresponding to $p < 0.05$ (shown in bold) were considered to indicate statistically significant differences in ϕ^{ROR} between COFEMI and MI. Table 2.8 presents p -values for the comparison of ϕ^{ROR} values between COFEMI and CMI^{grad}. These results for ϕ^{TRE} and ϕ^{ROR} indicate that (1) COFEMI provides equal or better registration performance for images with and without serious artifacts and that (2) a statistically significant improvement in registration performance is obtained using COFEMI over MI or CMI^{grad} for images with moderate to high levels of inhomogeneity.

Experiment 1b: Skerl Evaluation Protocol

The evaluation protocol described in Section 2.4.3 was applied to 30 pairs of \mathcal{P} and $\mathcal{T}2$ from \mathcal{D}^1 to calculate 5 additional properties $\phi^{S1}-\phi^{S5}$. Since the data in \mathcal{D}^1 was simulated using a similar ICBM protocol [84] as in [8], the parameter space normalization

Noise/Field (σ, b)	$\sigma = 0$	$\sigma = 1$	$\sigma = 2$	$\sigma = 3$	$\sigma = 4$
$b = 0.0$	0.2134	0.7029	0.2019	0.9859	0.1605
$b = 0.5$	0.8155	0.3623	0.0851	0.1503	0.6249
$b = 1.0$	0.3745	0.9372	0.3605	0.6450	0.5929
$b = 1.5$	0.7566	0.3224	0.1773	0.0546	0.5559
$b = 2.0$	0.4552	0.6838	0.1822	0.4206	0.0497
$b = 2.5$	0.3782	0.0693	0.0034	0.0166	0.0036

Table 2.6: p -values for paired t -tests comparing ϕ^{TRE} values for CMI^{grad} and COFEMI for 30 different combinations of simulated noise (σ) and field inhomogeneity (b). These values correspond to the ϕ^{TRE} differences shown in Figure 2.7(e) averaged across 30 pairs of $\mathcal{T}1$ and $\mathcal{P}_{b,\sigma}^r$ for each b, σ .

Noise/Field (σ, b)	$\sigma = 0$	$\sigma = 1$	$\sigma = 2$	$\sigma = 3$	$\sigma = 4$
$b = 0.0$	0.1183	0.1752	0.2933	0.1982	0.5281
$b = 0.5$	0.3661	0.2496	0.4450	0.1070	0.7237
$b = 1.0$	0.5934	0.8980	0.1637	0.9800	0.6305
$b = 1.5$	0.0568	0.3829	0.2170	0.6346	0.6700
$b = 2.0$	0.2396	0.0695	0.0490	0.2147	0.0127
$b = 2.5$	0.2732	0.2585	0.0436	0.0026	0.0207

Table 2.7: p -values for paired t -tests comparing ϕ^{ROR} values for MI and COFEMI for 30 different combinations of simulated noise (σ) and field inhomogeneity (b). These values correspond to the ϕ^{ROR} differences shown in Figure 2.7(d) averaged across 30 pairs of $\mathcal{T}1$ and $\mathcal{P}_{b,\sigma}^r$ for each b, σ .

Noise/Field (σ, b)	$\sigma = 0$	$\sigma = 1$	$\sigma = 2$	$\sigma = 3$	$\sigma = 4$
$b = 0.0$	0.1877	0.7110	0.1932	0.8598	0.1580
$b = 0.5$	0.9551	0.3000	0.1576	0.1663	0.9636
$b = 1.0$	0.5231	0.8746	0.6066	0.7910	0.4856
$b = 1.5$	0.7489	0.2913	0.2073	0.0800	0.8226
$b = 2.0$	0.3428	0.7609	0.2635	0.6177	0.0395
$b = 2.5$	0.5035	0.0906	0.0049	0.0195	0.0053

Table 2.8: p -values for paired t -tests comparing ϕ^{ROR} values for CMI^{grad} and COFEMI for 30 different combinations of simulated noise (σ) and field inhomogeneity (b). These values correspond to the ϕ^{ROR} differences shown in Figure 2.7(f) averaged across 30 pairs of $\mathcal{T}1$ and $\mathcal{P}_{b,\sigma}^r$ for each b, σ .

Similarity Measure	Relationship Characterized	Formula
Correlation Coefficient (CC)	Linear Intensity Differences	$\Psi^{CC}(\mathcal{A}, \mathcal{B}) = \frac{\sum_{c \in C} (f^{\mathcal{A}}(c) - \bar{f}^{\mathcal{A}})(f^{\mathcal{B}}(c) - \bar{f}^{\mathcal{B}})}{\sqrt{\sum_{c \in C} (f^{\mathcal{A}}(c) - \bar{f}^{\mathcal{A}})^2 \sum_{c \in C} (f^{\mathcal{B}}(c) - \bar{f}^{\mathcal{B}})^2}}$
Sums of Squared Differences (SSD)	Linear Intensity Differences	$\Psi^{SSD}(\mathcal{A}, \mathcal{B}) = \sum_{c \in C} (f^{\mathcal{A}}(c) - f^{\mathcal{B}}(c))^2$
Mutual Information (MI)	Joint Gray Level Distribution	$\Psi^{MI}(\mathcal{A}, \mathcal{B}) = S(\mathcal{A}) + S(\mathcal{B}) - S(\mathcal{AB})$
Normalized Mutual Information (NMI)	Joint Gray Level Distribution	$\Psi^{NMI}(\mathcal{A}, \mathcal{B}) = \frac{S(\mathcal{A}) + S(\mathcal{B})}{S(\mathcal{AB})}$
Entropy Correlation Coefficient (ECC)	Joint Gray Level Distribution	$\Psi^{ECC}(\mathcal{A}, \mathcal{B}) = 2 - 2 \frac{S(\mathcal{AB})}{S(\mathcal{A}) + S(\mathcal{B})}$
Negative Joint Entropy (JE)	Joint Gray Level Distribution	$\Psi^{JE}(\mathcal{A}, \mathcal{B}) = -S(\mathcal{AB})$
Energy of Joint Histogram (EOH)	Joint Gray Level Distribution	$\Psi^{EOH}(\mathcal{A}, \mathcal{B}) = \sum_{f^{\mathcal{A}}} \sum_{f^{\mathcal{B}}} p^2(f^{\mathcal{A}}, f^{\mathcal{B}})$
Point Similarity Measure Based on MI (PSMI)	Joint Gray Level Distribution	$\Psi^{PSMI}(\mathcal{A}, \mathcal{B}) = \sum_a \sum_b p(a, b) PS_{MI}(a, b),$ where $PS_{MI}(a, b) = \log \left(\frac{p_{MI}(a, b)}{p_{MI}(a)p_{MI}(b)} \right)$

Table 2.9: Similarity measures Ψ compared to COFEMI by the Skerl evaluation protocol, in terms of two generic images \mathcal{A} and \mathcal{B} with respective intensities $f^{\mathcal{A}}(c)$ and $f^{\mathcal{B}}(c)$ at coordinate c , and average intensities $\bar{f}^{\mathcal{A}} = \frac{1}{|C|} \sum_{c \in C} f^{\mathcal{A}}(c)$ and $\bar{f}^{\mathcal{B}} = \frac{1}{|C|} \sum_{c \in C} f^{\mathcal{B}}(c)$. Note that p_{MI} represents a prior gray level distribution determined from two images that are known to be in correct alignment.

was performed as in [8]. The 3-dimensional parameter hyperspace is sampled on $N = 40$ random, unique lines with $M = 300$ evenly space points per line. Gaussian noise of variance $\sigma = 4$ was introduced into both $\mathcal{T}2$ and \mathcal{P} . Simulated bias field inhomogeneity with a sharpness of $b = 2$ was used to modulate \mathcal{P} intensities as described in Section 2.5.2. Hence, the scenes $\mathcal{T}2_{\sigma}$ and $\mathcal{P}_{b,\sigma}$ with $\sigma = 4, b = 2$ were considered.

In addition to the feature-driven similarity measure provided by COFEMI (Ψ^{COFEMI}) and described in Section 2.3.4, eight additional similarity measures were compared. These include negative joint entropy (Ψ^{JE}), MI (Ψ^{MI}), NMI (Ψ^{NMI}), ECC (Ψ^{ECC}), energy of the joint histogram (Ψ^{EOH}), and a point similarity measure based on MI [85]

(Ψ^{PSMI}). Two linear similarity measures, correlation coefficient (Ψ^{CC}) and sums of squared distance (Ψ^{SSD}), were also compared. The measures and their formulations are given in Table 2.9.

Properties Independent of Trace Radius

For each Ψ^u , $u \in \{COFEMI, MI, NMI, ECC, EOH, JE, SSD, CC, PSMI\}$, values for each of the 5 properties $\phi^{S1}-\phi^{S5}$ (Table 2.3) were calculated along $N = 40$ transformation parameter lines. A comparison of each Ψ^u in terms of these 5 properties is presented in Figure 2.8. As originally observed in [8], ACC is the least informative measure for synthetic data, where the ground truth is heavily favored as the global maximum. As may be discerned by Figure 2.8(a), the average ACC is less than the size of a pixel. On the other hand, CR indicates a significant difference between MI-based measures, which do not suffer from local minima until at least 40 mm of misalignment. CC, SSD, EOH, and JE, are similarly prone to local minima even when very close to ground truth. Consequently, MI, NMI, ECC, PSMI, and COFEMI are able to attain the global maximum in spite of a larger initial misalignment, where COFEMI and PSMI exhibits the highest capture range. As with several other results, PSMI exhibits a marginal superiority attributable to the use of a prior gray level distribution p_{MI} determined from a training image set in correct alignment. Figure 2.8(c) provides a comparison of NOM values for all measures, while Figure 2.8(d) magnifies the vertical axis to compare just the MI-based measures ($\Psi^u, u \in \{COFEMI, MI, NMI, ECC, PSMI\}$). NOM values reflect the CR results, in that the MI-based measures are vastly superior in terms of absence of local minima. Over all N lines COFEMI only had a single local minimum, while MI had three. Figure 2.8(e) presents the average DO values over all N lines, which clearly show that COFEMI performs favorably compared to all other measures, including the PSMI measure which requires off-line training. The average RON is shown in Figure 2.8(f) for just the MI-based measures, as RON for the other measures are very high. These RON values indicate that, in addition to few local minima, any gradients of Ψ^u that point away from the global maximum are small and uncommon for Ψ^{COFEMI} .

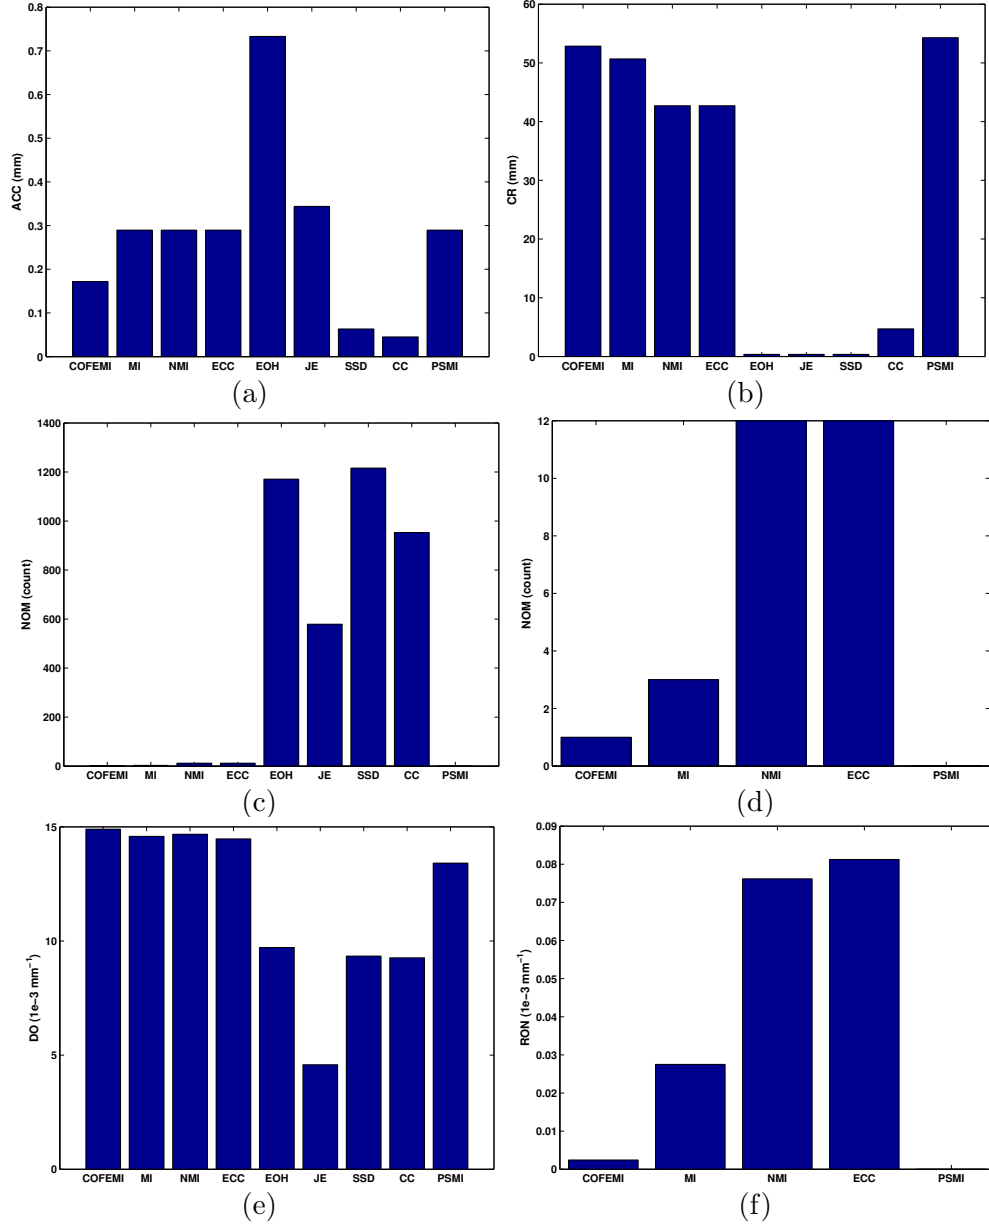


Figure 2.8: Comparison of 9 similarity measures Ψ^u , $u \in \{COFEMI, MI, NMI, ECC, EOH, JE, SSD, CC, PSMI\}$ in terms of the 5 properties proposed by Skerl [8]. (a) ACC indicates that global maximum of all similarity measures correspond very closely to the ground truth registration. (b) CR, (c) NOM, and (e) DO for all Ψ^u . (d) NOM for only the MI-based measures in (c) is shown magnified for clearer comparison. RON was very high for measures $u \in \{EOH, JE, SSD, CC\}$, thus RON is shown for only MI-based measures in (f), indicating superior optimization characteristics of Ψ^{COFEMI} . The supervised PSMI method demonstrates marginally better performance in terms of NOM and RON versus COFEMI.




Figure 2.9: NOM versus R , the distance of deformation from ground truth registration. (a) NOM for $\Psi^u, u \in \{EOH, JE, SSD, CC\}$, demonstrate a large number of minima compared with (b) NOM for $\Psi^u, u \in \{COFEMI, MI, NMI, ECC, PSMI\}$. Note the scale difference between (a) and (b). Only PSMI demonstrates any improvement over COFEMI, which can only be seen as the marginal difference at the highest value of R in (b).

Number of Minima versus Trace Radius

The number of local minima can also be visualized as a function of distance (R) from ground truth at line position $m_s = 0$. In Figure 2.9, NOM is plotted against the magnitude of misalignment, or equivalently the number of steps away from ground truth at $m_s = 0$ averaged over all N parameter lines. Figure 2.9(a) shows NOM versus R for EOH, JE, SSD, and CC, which are two orders of magnitude greater than NOM for the MI-based measures shown in Figure 2.9(b). It is important to note the X-axis in Figure 2.9(b), which begins at 38 mm from ground truth. It can be seen that NOM values for ECC and NMI are equivalent; a logical consequence of their related formulations. Each of MI, COFEMI, and PSMI do not suffer from local minima until R is much larger ($R > 50$ mm). Again, these results suggest that COFEMI is superior over MI, while COFEMI and the supervised PSMI technique demonstrate the fewest local minima at most values of R .

Risk of Non-convergence versus Trace Radius

Figure 2.10(a) illustrates the comparison of all similarity measures in terms of RON values obtained for each N lines. As previously stated, the measures EOH, NJE, SSD,

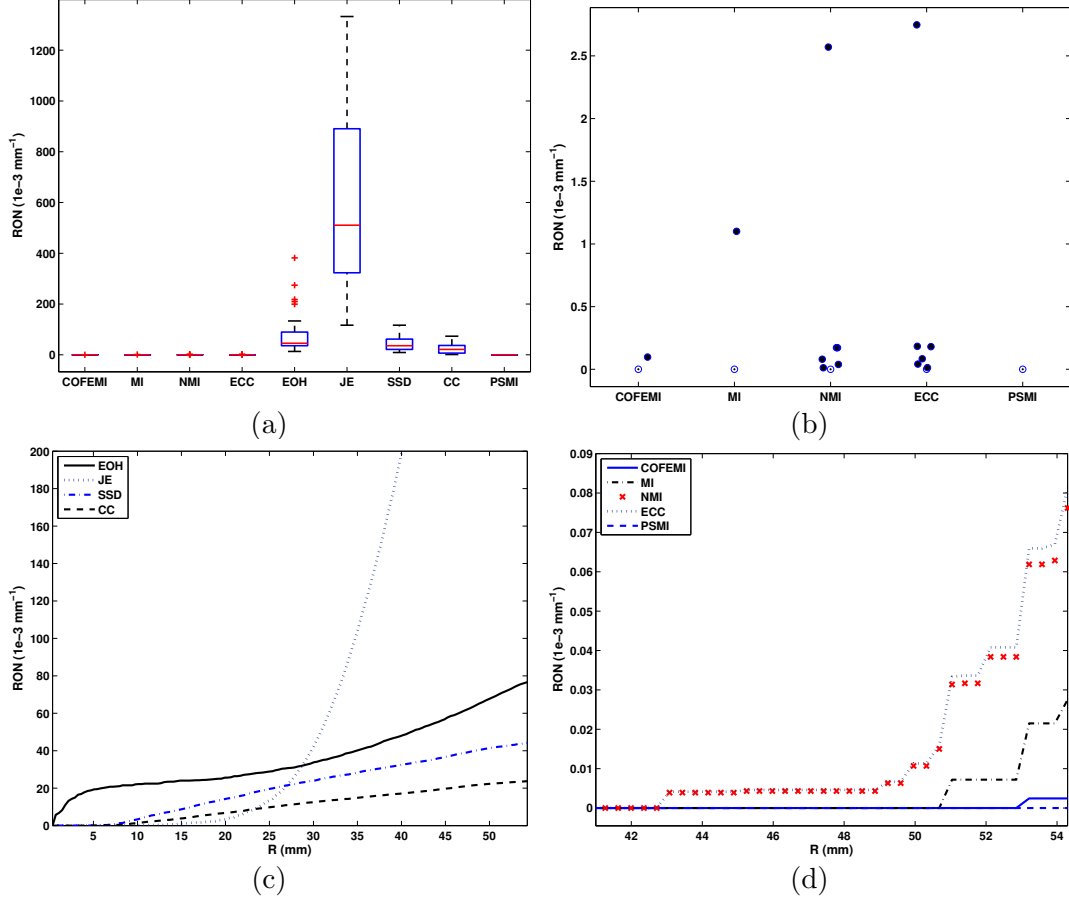


Figure 2.10: (a) RON box-plot for all similarity measures Ψ^u , $u \in \{COFEMI, MI, NMI, ECC, EOH, JE, SSD, CC, PSMI\}$. (b) Compact box-plot for Ψ^u , $u \in \{COFEMI, MI, NMI, ECC, PSMI\}$. Note the scale difference between (a) and (b). (c) RON averaged over N lines versus R for Ψ^u , $u \in \{EOH, JE, SSD, CC\}$ and (d) for $u \in \{COFEMI, MI, NMI, ECC, PSMI\}$. As with NOM, MI-based measures are superior, while Ψ^{COFEMI} and Ψ^{PSMI} have the best convergence characteristics.

and CC have a very high RON as R becomes large, while the RON for MI-based measures remains relatively low, as can be seen in Figure 2.10(b) by the few outlying lines with non-zero RON values. In Figure 2.10(c), RON is plotted against the R for the measures EOH, JE, SSD, and CC. As with NOM, the MI-based measures in Figure 2.10(d) exhibit RON values several orders of magnitude smaller than the other measures. Again, NMI closely tracks ECC, while COFEMI and PSMI exhibit the best performance.




Figure 2.11: (a) 4 Tesla *ex vivo* MR image (\mathcal{M}), (b) corresponding histological section \mathcal{H} , and (c)-(h) MR feature scenes \mathcal{M}_β corresponding to Haralick correlation ($\kappa = 8$) and inverse difference moment ($\kappa = 5$), mean intensity ($\kappa = 5$), range ($\kappa = 5$), range ($\kappa = 8$), and Haralick correlation ($\kappa = 5$).

2.5.3 Experiment 2: Multimodal Prostate Registration

Experiment 2 Overview

The multimodal prostate data set \mathcal{D}^2 was used to compare the performance of COFEMI with MI for the task of aligning the real *ex vivo* T2-w MRI with the whole mount histological sections. The objective of this experiment was to compare COFEMI and MI in terms of their ability to map known CaP extent from H&E stained \mathcal{H} onto the corresponding \mathcal{M} . Each MRI section \mathcal{M} was aligned to corresponding \mathcal{H} sections to obtain a registered MR image \mathcal{M}^r via an affine transformation determined using COFEMI and also by maximization of MI (Equation 2.4). Intensity-based MI registration was performed using 128 gray level bins, empirically chosen for maximal performance. In this multimodal task, \mathcal{H} possesses clear and extensive structural details, while the MRI sections are affected by noise, inhomogeneity artifacts, and lower spatial resolution [3]. Therefore, feature scenes were only calculated (Section 2.3.2) from \mathcal{M} and the feature ensemble selection algorithm (Section 2.3.3) was used to choose 4 maximally unique and informative feature images $\pi_4^{\mathcal{M}} = \mathcal{M}_{\alpha_1} \cdots \mathcal{M}_{\alpha_4}$, where $\{\alpha_1, \dots, \alpha_4\} \in \{1, \dots, n\}$.

Figures 2.11(c)-(h) show some of the first and second order statistical feature scenes $\mathcal{M}_\beta, \beta \in \{1, \dots, n\}$ initially calculated by COFEMI from \mathcal{M} (Figure. 2.11(a)). Each subsequently chosen $\mathcal{M}_{\alpha_1}, \dots, \mathcal{M}_{\alpha_4}$ contributes to describing the corresponding \mathcal{H} (Figure 2.11(b)), by addressing the resolution disparity between \mathcal{H} and \mathcal{M} and by providing representations of \mathcal{M} that can overcome the intensity inhomogeneity, clearly visible at the top of \mathcal{M} . COFEMI was then performed using Haralick Correlation ($\kappa = 5$), mean ($\kappa = 5$), and range ($\kappa = 5$) features selected by *CMIfeatures*. For \mathcal{H} , $\pi_l^{\mathcal{H}} = \emptyset$ was used. MI-based registration, which corresponds to $k = 0$, was also performed to align \mathcal{M} and \mathcal{H} .

Qualitative Evaluation

Figure 2.12 shows 4 consecutive pairs of corresponding histology and MRI slices of the prostate. The boundary of the prostate region $G(\mathcal{H})$ on histology ($G^b(\mathcal{H})$) is shown in green on Figures 2.12(a)-(d), while the red contours in Figures 2.12(e)-(h) represent the prostate boundary in the unregistered MR images ($G^b(\mathcal{M})$). With $G^b(\mathcal{H})$ overlaid onto the MRI with $G^b(\mathcal{M})$, it is easy to visualize the degree of misalignment between \mathcal{H} and \mathcal{M} in terms of the distance between the red and green boundaries. The transformed MRI intensity images \mathcal{M}^r obtained by MI registration in Figures 2.12(i)-(l) show improved alignment of registered boundaries $G^b(\mathcal{M}^r)$ and $G^b(\mathcal{H})$ over the unregistered boundaries. However, noticeable misalignment is observable at the tops of the gland in Figure 2.12(i)-(l), where hypointense regions hinder image similarity and hence correct alignment. On the other hand, COFEMI provided consistently superior alignment using all four feature ensembles. This performance improvement is illustrated in Figures 2.12(m)-(p) for the feature ensemble comprising range, correlation, and mean features (automatically chosen by *CMIfeatures* for $k = 3$). As evidenced by the closer coincidence of $G^b(\mathcal{M}^r)$ and $G^b(\mathcal{H})$, COFEMI is clearly robust to the inhomogeneity of tissue intensities visible in the MR intensity images, which caused the misregistration associated with conventional MI (Figures 2.12(i)-(l)).




Figure 2.12: (a)-(d) Four consecutive histology sections \mathcal{H} with the prostate boundary shown in green and (e)-(h) corresponding unregistered MRI slices \mathcal{M} with the prostate boundary $G^b(\mathcal{M})$ accentuated in red and the green boundary $G^b(\mathcal{H})$ from (a)-(d) overlaid. Misalignment of the red and green boundaries in (e)-(h) indicates the initial misalignment between $G^b(\mathcal{H})$ and $G^b(\mathcal{M})$. (i)-(l) Registered MRI slices \mathcal{M}^r generated by intensity-based MI demonstrate misalignment near hypointense regions at the top of \mathcal{M} , as evidenced by the space between the red and green contours. (m)-(p) Improved accuracy of COFEMI is evidenced by closer coincidence of $G^b(\mathcal{H})$ and $G^b(\mathcal{M}^r)$, particularly at the locations in (i)-(l) where MI exhibited misalignment.

Measure	ϕ^{ROR}	ϕ^{HD}	ϕ^{MAD}
MI	0.924	19.843	3.417
CMI ^{grad}	0.932	19.819	3.1023
COFEMI ($k = 3$)	0.940	18.296	2.786

Table 2.10: Φ^{ROR} , Φ^{HD} , and Φ^{MAD} values, averaged over 26 pairs of \mathcal{M} and \mathcal{H} , and corresponding to registration via MI, CMI^{grad}, and COFEMI, are given for prostate regions $G(\mathcal{M}^r)$ and $G(\mathcal{H})$ and boundaries $G^b(\mathcal{H})$ and $G^b(\mathcal{M}^r)$.

Comparison and Measure	ϕ^{ROR}	ϕ^{HD}	ϕ^{MAD}
MI <i>vs</i> CMI ^{grad}	0.0406	0.9610	0.1157
MI <i>vs</i> COFEMI	7×10^{-4}	0.033	0.0120
CMI ^{grad} <i>vs</i> COFEMI	0.0127	0.003	0.0577

Table 2.11: p -values for t -tests comparing ϕ^{ROR} , ϕ^{HD} , and ϕ^{MAD} values corresponding to the average values presented in Table 2.10, indicating consistently higher overlap and lower boundary distances obtained from COFEMI versus MI, and COFEMI versus CMI^{grad}.

Quantitative Evaluation

Since no ground truth for correct alignment is available for the real multimodal prostate data, quantitative evaluation of registration was done in terms of (1) Hausdorff distance ϕ^{HD} , and (2) ϕ^{MAD} between $G^b(\mathcal{H})$ and $G^b(\mathcal{M}^r)$, and (3) overlap ϕ^{ROR} of $G(\mathcal{H})$ and $G(\mathcal{M}^r)$. Table 2.10 below quantitatively summarizes the multimodal prostate image registration results in terms of average $\phi^{ROR}(G(\mathcal{M}^r), G(\mathcal{H}))$, $\phi^{HD}(G^b(\mathcal{M}^r), G^b(\mathcal{H}))$, and $\phi^{MAD}(G^b(\mathcal{M}^r), G^b(\mathcal{H}))$ values over 26 pairs of \mathcal{M} and \mathcal{H} . Average values of ϕ^{ROR} were higher from COFEMI than from MI or CMI^{grad}, while the average values of ϕ^{ROR} were higher from CMI^{grad} compared to MI (COFEMI > CMI^{grad} > MI). Similarly, average values of ϕ^{HD} and ϕ^{MAD} were lower from COFEMI than from MI or CMI^{grad}, and lower from CMI^{grad} compared to MI (COFEMI < CMI^{grad} < MI). The values of ϕ^{ROR} , ϕ^{HD} and ϕ^{MAD} for COFEMI, CMI^{grad} and MI were compared in a pairwise fashion using paired t -tests under the null hypothesis that the ϕ^{TRE} , ϕ^{HD} and ϕ^{MAD} values are equal for each of COFEMI, CMI^{grad} and MI. These values, shown in Table 2.11, indicate significantly higher ϕ^{ROR} and significantly lower ϕ^{HD} and ϕ^{MAD} with COFEMI compared to both MI and CMI^{grad}.

2.6 Concluding Remarks

In this study we presented a new feature-driven multimodal image registration scheme called COFEMI that improves upon existing MI-based techniques via automatically generating an ensemble of feature images. An intelligent feature selection step means the scheme can be tailored to any registration task. While feature space representations have been used in different forms [57–61, 86] to encode spatial information and enhance multimodal image registration performance, the COFEMI technique is the first to provide a method for automatic selection of an optimal ensemble of multiple features specific to the registration task at hand and which are less susceptible to intensity artifacts and to structural differences across modalities. COFEMI introduces the following significant improvements over previous registration techniques,

1. Robustness to intensity artifacts such as inhomogeneity and noise.
2. Robustness to modality differences that otherwise weaken correspondence of anatomical structures between intensity images.
3. Natural generalizability to diverse multimodal data sets with different requirements for structural description via image-derived low-level textural features.

We demonstrate the ability of COFEMI to qualitatively and quantitatively improve registration accuracy over intensity-based MI on a multimodal prostate data set and a synthetic multiprotocol brain MR data set. The findings of our experiments are summarized as follows,

- In terms of Hausdorff distance and the region overlap ratio of the prostate and its boundaries on 26 pairs of corresponding MR and histology images, COFEMI significantly outperformed both MI and CMI^{grad} in terms of accuracy.
- A comparison of TRE and ROR values for the brain region on 30 pairs of synthetic multiprotocol MR images over 5 levels of noise and 6 levels of inhomogeneity, totaling 900 registration trials, similarly demonstrated the performance improvement of COFEMI over MI.

Offline Operations for a Modality Pair		Online Operations	
Comprehensive Feature Calculation	Feature Ensemble Selection	Ensemble Generation for Test Image Pair	Affine Registration
110	4	10	45

Table 2.12: Approximate computation times (in seconds) for the different modules involved in the operation of the COFEMI registration scheme. Note that calculation of the full, comprehensive feature set and selection of the optimal ensemble are performed only once initially (offline) for a given modality pair. To register successive test image pairs (online), it is only required to calculate the features comprising the selected ensemble. Computation times were obtained for images of size 256×256 running on an Intel Core 2 Duo 2.4GHz PC running 64-bit MATLAB[®].

- A comprehensive comparison of 8 similarity measures with COFEMI using the evaluation protocol proposed by Skerl in conjunction with 30 pairs of synthetic T2-w and PD-w MR images quantitatively validated the superiority of COFEMI as a similarity measure for rigid registration in terms of 5 different similarity measure properties.

While our technique is currently presented in the context of affine registration, COFEMI can also be used to drive elastic registration techniques, such as free form deformations (FFD) [17], and we are presently investigating the use of COFEMI in a hierarchical FFD scheme. COFEMI may also be implemented in conjunction with any deformation model, intensity interpolation technique, parameter optimization scheme, or entropy estimation method. Thus, COFEMI may be utilized to provide a customized feature-based objective function for a FFD method such as [17], or simply to drive the optimization of affine transformation parameters. The later application is considered here to evaluate COFEMI in the absence of confounding factors associated with complicated optimization strategies, while the former non-linear application is presently being investigated as a powerful tool for automated deformable multimodal image registration. Table 2.12 lists typical computation times (in seconds) for the primary modules of COFEMI in an affine transformation framework.

While any technique for joint entropy estimation may be used, COFEMI registration is performed here using high order joint histograms with 32 gray level bins for a single feature image, 16 bins for two feature images, and 10 bins for up to 4 features. These

values were selected both to limit the total number of bins in the histogram, and empirically for performance, as with MI. Joint probability distribution estimates used in the calculation of both MI and CMI were obtained via the normalized histograms of intensity and feature values. Obtaining accurate estimates of higher order joint entropy represents the largest limitation associated with the COFEMI technique. However, with recent advances in information theory [87], high order entropy can be accurately calculated for even very sparse joint histograms.

Chapter 3

Improving Supervised CAD on *ex vivo* MRI via Registration

3.1 Importance of CAD Training Data Accuracy

Supervised classifiers are the central component of computer-aided diagnosis (CAD) systems for medical imagery, in which the goal is to classify observations (e.g. pixels) as belonging to a particular class (e.g. cancer or non-cancer) based on prior observations. A crucial step in the implementation of such classifiers is obtaining object labels that accurately represent the targeted tissue class. The most common way of obtaining such class labels involves manual labeling by an expert (e.g. radiologist). Although, even with trained experts, manual annotations are time consuming, variable, and error prone [28]. The adverse effect of any mislabeled training instances on supervised classifier performance has been repeatedly observed [29, 32, 88] in terms of accuracy and robustness. As such, avoidance of incorrectly labeled instances is of utmost importance in building an effective supervised CAD classifier.

Consider the MRI image of a prostate gland is shown in Figure 3.1(a) from which a CAD system [28] driven by a Bayesian classifier calculates the likelihood that each pixel belongs to the cancer class (Figure 3.1(b)) and obtaining a binary prediction for cancer presence (Figure 3.1(c)). Figure 3.1(e) shows the likelihood image obtained by CAD after removal of erroneous training instances, indicating that removal of even a small portion of mislabeled data resulted in a significant improvement in classification accuracy as can be discerned in the increase in contrast in the posterior likelihood scenes, indicating enhanced separation between classes (compare Figures 3.1(e) and (b)). Figure 3.1(d) shows a reconstructed WMH section of the prostate obtained via radical prostatectomy which corresponds to the MRI slice in Figure 3.1(a).




Figure 3.1: (a) MRI slice of a prostatectomy specimen with cancer present. (b) Likelihood image generated by a CAD system applied to (a), and (c) the prediction for cancer presence obtained by thresholding (b) at 50% confidence level and superposing on (a). (d) The corresponding histology slice shown labeled for cancer extent by H&E stain (black arrow). (e) Likelihood image calculated by CAD after removal of erroneous training instances, and (f) the resulting binary prediction.

Delineating regions of CaP on MRI is difficult even for radiologists; fortunately, images from multiple modalities can be utilized to assist in labeling the CAD targeted modality. In this study, histology provides a modality with readily available ground truth for CaP extent, found with the aid of haematoxylin and eosin (H&E) staining. For instance, manual mapping of histological ground truth onto MRI using anatomical features as visual cues has been previously demonstrated in [28], where CaP extent on MRI was obtained by examination of WMH and radiological sections by a pathologist and a radiologist simultaneously. For example, the ground truth for CaP extent provided by the WMH section in Figure 3.1(d) is manually mapped by the simultaneous efforts of both experts onto the corresponding MRI in Figure 3.1(a). Despite possession of empirical domain knowledge by the experts performing the mapping, this process is also time consuming and subject to observer bias. Note that this task is akin to manual registration or alignment of the two modalities, whereby a mapping between

corresponding pixels is determined. Thus, it is clear that in order to accurately establish training instances using multimodal imagery, it is necessary to achieve accurate registration.

In this view of the manual mapping process, we present an automated labeling scheme whereby the histological ground truth is automatically mapped directly to MR from aligned WMH sections via image registration techniques. The optimal alignment between two images \mathcal{A} and \mathcal{B} is obtained when the pixels representing the same material points on the anatomical object or region of interest (C_a in \mathcal{A} , C_b in \mathcal{B}) are mapped into the same spatial coordinates ($C_b = T(C_a)$). The result of successful image registration is a pair of images $\hat{\mathcal{A}}$ and \mathcal{B} where $\hat{\mathcal{A}}$ represents \mathcal{A} subjected to the determined transformation T , and pixels C_a in $\hat{\mathcal{A}}$ and C_b in \mathcal{B} represent the same underlying anatomical features of the region of interest. If we consider the set of pixels that represents the CaP extent on histology, it is clear that determining the transform T , which specifies the mapping of pixels on WMH to MRI, can provide a precise ground truth estimate for CaP on MRI.

In summary, we present a method for automated training of a supervised classifier by establishing accurate training instances via registration of multimodal images. We demonstrate the method for a specific CAD application where the objective is to automatically identify CaP on high resolution *ex vivo* MRI. We further quantitatively demonstrate that the object class labels established by registration are more reliable compared to expert annotations even when (1) considerable image differences exist between modalities and (2) the manual segmentations are meticulously obtained by multiple experts analyzing both modalities simultaneously. The principal contributions of this work are to demonstrate,

- a novel method for automatically training a supervised classifier when multimodal data is present,
- the ability of multimodal registration methods to accurately annotate object labels with significantly greater accuracy compared to expert delineations in spite of differences between disparate modalities (e.g. histology and MRI) which generally

hinder successful registration, and

- the potential use of classifier performance as a benchmark for evaluating registration performance.

The rest of this paper is organized as follows. In Section 3.2 we outline the design of the supervised classifier training technique and its evaluation. In Section 3.3 we describe the registration methodologies, the classifier and the automated training method, and the cross-validation procedure used to compare the accuracy of multiple classifiers. In Section 3.6 we present the results of each supervised classifier. Concluding remarks and future research directions are presented in Section 3.7.

3.2 Overview of Experimental Design

There has been relatively little work in the area of prostate registration, less so in the context of multimodal registration. Registration of prostate histology with ultrasound images has been addressed recently by Taylor [19] who utilized surface matching, Moskalik [45] who fit the prostate boundary with an ellipsoid to find an affine mapping and Porter [10] who used blood vessels as fiducial markers. Lee [12] performed non-rigid registration of CT/MRI with histological sections for a visualization system. In this study, we consider two rigid and two non-rigid multimodal registration methods for aligning prostate histology with MRI. Finding a transformation T which can bring the pixels representing the prostate on histology and MRI into spatial correspondence is hindered by:

1. Dissimilarity of the modalities and the structural characteristics captured by the modalities.
2. The complicated nature of the deformations to which the anatomical object is subjected.

Namely, the processes of quartering, sectioning and fixation of the gland in formalin for acquisition of digitized histological images subject the tissue to irregular and extensive deformation and alter the underlying tissue.

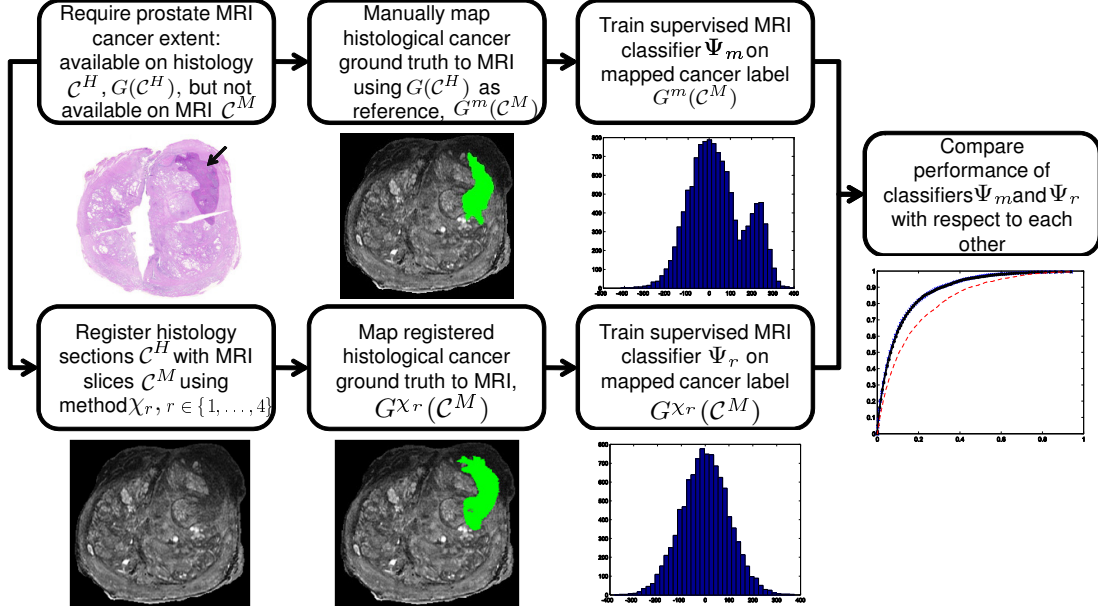


Figure 3.2: Outline of registration-based supervised classifier training system applied to CAD for prostate MRI with histological ground truth for cancer (top left). Labeling of MRI training data is achieved by registration of histology with MRI, followed by directly mapping histological ground truth to MRI, which are used to train the supervised classifier which drives the CAD system (bottom row). Training labels are also established by expert manual delineation, upon which another classifier is trained (top row). The classifiers are then evaluated with respect to each other (right).

Most current robust non-linear multimodal image registration techniques operate by initially determining an affine or rigid alignment as a precursor to an elastic deformation step [17,89,90]. An accurate initial global alignment is often important for the success of the elastic alignment step, primarily because physically realistic deformations are most accurately modeled when required deformations are small. Thus, in [3] we presented our registration technique, termed combined feature ensemble mutual information or COFEMI, with emphasis on obtaining an accurate global affine precursor to elastic deformation. In this study, we demonstrate two similarity measures in the optimization of an affine transformation: MI and COFEMI, both followed by elastic transformation. To compare the influence of the affine alignments on the final alignment, we utilize a control point driven thin plate splines (TPS) deformation model to elastically transform histology onto MRI. Thus, the design of our supervised classifier training method and its evaluation against expert-based training are outlined in Figure 3.2 using the notation defined in Section 3.3.1 for a set of multimodal images:

1. Automated affine registration of MRI and histology is performed via MI and COFEMI.
2. Subsequent TPS warping is applied following affine alignment to achieve the final image alignments (MI+TPS and COFEMI+TPS) using prominent anatomical landmarks as control points.
3. The ground truth for spatial extent of cancer is directly mapped onto the CAD-targeted modality (MRI).
4. Manual delineations of cancer on MRI are obtained from radiologists using visual cues from histology.
5. The pixels automatically labeled as cancer and those delineated by experts are used as training instances in separate supervised classifiers.
6. Accuracies of the classifiers trained using the 4 registration methods and the expert delineation are compared in terms of ROC curves using a 5-fold randomized cross validation procedure.

In this manner, we demonstrate that in spite of the challenges in registration of prostate MRI and histology, the automated CaP annotations thus obtained are significantly more accurate compared to expert delineation of CaP on MRI obtained with aid from corresponding stained histological sections and prior domain knowledge. We further demonstrate that despite the non-linear nature of the deformations even simple rigid registration methods are more accurate compared to expert annotation, yielding better quality classifier training data.

3.3 Overview of Registration Methods

3.3.1 Notation

We define a 2D image by the notation $\mathcal{C} = (C, f)$ where C is a coordinate grid, called the image domain, and f is a function that assigns to every spatial location $\mathbf{c} \in C$ an intensity value $f(\mathbf{c})$ in 2D space $\mathbf{c} = (x, y)$. Under this convention, the 2D slices of the 3D *ex vivo* MRI dataset are defined by $\mathcal{C}^M = (C, f^M)$, while the corresponding

histology slices are defined by $\mathcal{C}^H = (C, f^H)$. The notations and symbols used in the rest of this paper are listed in Table 3.1.

Table 3.1: List of notation and symbols specific to Chapter 3.

Symbol	Description
\mathcal{C}^M	2D MRI image
\mathcal{C}^H	2D Histology image
$G(\mathcal{C}^H)$	CaP pixels from H&E staining on \mathcal{C}^H
$G^m(\mathcal{C}^M)$	CaP pixels found manually on \mathcal{C}^M
χ_r	Registration method where $r \in \{1, 2, 3, 4\}$
$G^{\chi_r}(\mathcal{C}^M)$	$G(\mathcal{C}^H)$ on \mathcal{C}^M by method χ_r
\mathcal{C}^{Φ_β}	MRI feature image for operator Φ_β

3.3.2 Manual Mapping of Cancer from Histology to MRI

Histological sections were microscopically examined to determine cancer extent on \mathcal{C}^H denoted by the set of pixels $G(\mathcal{C}^H)$. Manual expert segmentations of cancer on *ex vivo* MRI, denoted by the set of pixels $G^m(\mathcal{C}^M)$, are obtained by visually mapping $G(\mathcal{C}^H)$ onto \mathcal{C}^M . Note that the manual delineation of cancer $G^m(\mathcal{C}^M)$ on \mathcal{C}^M was guided by $G(\mathcal{C}^H)$ on \mathcal{C}^H along with visual cues from the anatomical features of the prostate on both \mathcal{C}^H and \mathcal{C}^M . Since there were roughly twice the number of MRI slices as there were histological sections for each patient study, the MRI sections corresponding to the histological sections were manually identified by an expert by identifying anatomical features visible on both \mathcal{C}^H and \mathcal{C}^M . A total of 26 MRI-histological sections were considered in this study.

3.3.3 Overview of Registration Methods

Most robust rigid/affine multimodal registration techniques utilize similarity measures to bring images into alignment, the assumption being that when the images are in alignment, their similarity is maximized. The optimal alignment of the two images \mathcal{A} and \mathcal{B} obtained by the transformation T is given by

$$T_{max} = \operatorname{argmax}_T \psi(T\mathcal{A}, \mathcal{B}), \quad (3.1)$$

where ψ is some voxel-based similarity measure and $T\mathcal{A}$ represents the image \mathcal{A} subject to transformation T . The designation of \mathcal{A} and \mathcal{B} as moving (target) and stationary (template) images is user dependent. For the case of \mathcal{A} as the target of deformations, we further denote $\hat{\mathcal{A}}$ as the result of the initial affine registration step and $\tilde{\mathcal{A}}$ as the result of the subsequent elastic registration step. The four effective affine and elastic registration techniques (denoted $\chi_r \in \{\text{MI}, \text{MI+TPS}, \text{COFEMI}, \text{COFEMI+TPS}\}$) evaluated in this study are described in the following sections.

Maximization of Mutual Information (MI)

The most common similarity measures used for voxel-wise image registration are the sums of squared distance (SSD), correlation coefficient (CC), mutual information (MI), related information theoretic quantities such as entropy correlation coefficient (ECC), normalized mutual information (NMI) [39]. Since no linear relationship between voxel intensities in \mathcal{A} and \mathcal{B} may exist with multimodal imagery, mutual information and other statistical measures of image similarity or interdependence have been widely utilized for multimodal image registration tasks. Registration by maximization of MI can be achieved by calculating the MI of two images in terms of Shannon entropy [39], a measure of information content of a random variable. Equation 3.2 is a formulation of MI in terms of the marginal and joint entropies wherein the MI of a pair of images or random variables, $I_2(\mathcal{A}, \mathcal{B})$, is maximized by minimizing joint entropy $S(\mathcal{AB})$ and maintaining the marginal entropies $S(\mathcal{A})$ and $S(\mathcal{B})$. Where ensemble \mathcal{AB} represents simultaneous knowledge of both images, MI is given by,

$$I_2(\mathcal{A}, \mathcal{B}) = S(\mathcal{A}) + S(\mathcal{B}) - S(\mathcal{AB}), \quad (3.2)$$

where $I_2(\mathcal{A}, \mathcal{B})$ describes the interdependence of multiple variables or gray levels of a set of images [39]. Thus, when $I_2(\mathcal{A}, \mathcal{B})$ increases, the uncertainty about \mathcal{A} given \mathcal{B} decreases. As per Equation 3.1, it is assumed that the global MI maximum will occur at the point of precise registration.

Combined Feature Ensemble Mutual Information

Both MI and NMI can be ill-defined and highly non-convex measures for dissimilar modalities [38,39,91]. Hence, in [3] we presented the combined feature ensemble mutual information (COFEMI) registration technique, which incorporates additional information in the form of image features which are complementary to image intensity in order to increase the robustness of the similarity measure to modality differences and image artifacts, and was demonstrated to be superior compared to MI in terms of robustness to images with artifacts and dissimilar modalities. The COFEMI technique utilizes the notion of combined mutual information (CMI) to incorporate additional information in the form of image features. The CMI that two semi-independent images, \mathcal{B} and \mathcal{B}' , contain about a third image, \mathcal{A} , is defined as $I_2(\mathcal{A}, \mathcal{B}\mathcal{B}') = S(\mathcal{A}) + S(\mathcal{B}\mathcal{B}') - S(\mathcal{A}\mathcal{B}\mathcal{B}')$ [75]. This formulation allows the incorporation of only the unique (non-redundant) information provided by an additional image, \mathcal{B}' , about \mathcal{A} . Hence, the generalized form of CMI for n additional images is,

$$I_2(\mathcal{A}, \mathcal{B}\mathcal{B}'_1 \cdots \mathcal{B}'_n) = S(\mathcal{A}) + S(\mathcal{B}\mathcal{B}'_1 \cdots \mathcal{B}'_n) - S(\mathcal{A}\mathcal{B}\mathcal{B}'_1 \cdots \mathcal{B}'_n). \quad (3.3)$$

It can be seen that CMI incorporates only the unique information of additional images toward describing \mathcal{A} , thus enhancing but not overweighting the similarity measure with redundant information. Therefore, it will always be the case that $I_2(\mathcal{A}, \mathcal{B}\mathcal{B}'_1 \cdots \mathcal{B}'_n) \leq S(\mathcal{A}) = I_2(\mathcal{A}, \mathcal{A})$. The intuition behind using CMI is that one or more of the feature images $\mathcal{B}'_1, \mathcal{B}'_2, \dots, \mathcal{B}'_n$ derived from the intensity image \mathcal{B} will be relatively more immune to image intensity artifacts such as bias field inhomogeneity and intensity non-linearity compared to \mathcal{B} . In addition one or more of $\mathcal{B}'_1, \mathcal{B}'_2, \dots, \mathcal{B}'_n$ will potentially provide additional structural description of \mathcal{A} not readily discernible from intensity image \mathcal{B} . By choosing appropriate statistical features calculated over windows of \mathcal{B} , one may introduce information pertaining to the spatial arrangement of intensity values into the similarity measure. The feature extraction and selection algorithms described in [3] were used to generate a feature ensemble $\pi_k \subset \{\mathcal{B}'_1, \mathcal{B}'_2, \dots, \mathcal{B}'_n\}$ of length k used to drive the registration procedure by maximization of $I_2(\mathcal{A}, \mathcal{B}\pi_k)$. Details of feature extraction

are given in Section 3.4.1.

Affine Registration via MI and COFEMI

Images \mathcal{A} and \mathcal{B} are initially aligned using an affine transformation with a total of five parameters for rotation (θ), translation (dx, dy), and scaling (σ_x, σ_y). The affine transformation matrix (\mathbf{M}) is constructed by combination of three coordinate transformation matrices corresponding to rotation in the plane of the image (\mathbf{R}), translation in the X and Y directions (\mathbf{T}_{xy}), and horizontal and vertical scaling (\mathbf{S}_{xy}). The individual affine transformation matrices are composed in terms of the five parameters ($\theta, dx, dy, \sigma_x, \sigma_y$) by common geometric coordinate transformations equations. Homogeneous transformation matrices and coordinate representations were utilized to apply translations in the same matrix operation as scaling and rotation. Thus for a homogeneous coordinate $\mathbf{c}_h = (x, y, 1)^T$, the transformed coordinate is determined by $\mathbf{c}'_h = \mathbf{M} \mathbf{c}_h$. Deformations are applied in the order of rotation, scaling, and then translation by the product $\mathbf{M} = \mathbf{T}_{xy} \mathbf{S}_{xy} \mathbf{R}$. Both COFEMI and MI registration techniques (χ_1 and χ_2) utilize NN interpolation to avoid artifacts associated with linear and partial volume interpolation. MI-based registration is achieved by optimizing (via a Nelder-Mead simplex algorithm tuned for differences in parameter scale) the affine transformation \mathbf{M} by maximization of MI of \mathcal{A} with \mathcal{B} to obtain the optimal transformation \mathbf{M}^{x_1} . The COFEMI procedure operated by maximization of CMI of \mathcal{A} and ensemble $\mathcal{B}\pi_k$, resulting in the optimal rigid transformation \mathbf{M}^{x_2} to best align \mathcal{B} with \mathcal{A} by maximizing $I_2(\hat{\mathcal{A}}, \mathcal{B}\pi_k)$. The common coordinate frame C is maintained.

Elastic Registration via MI, COFEMI, and TPS

Thin plate splines (TPS) [92] were used to model the non-linear portion of deformation required to bring \mathcal{A} and \mathcal{B} into alignment. That is, having obtained $\hat{\mathcal{A}}$ via affine registration, a TPS defined deformation field was used to generate the final alignment result $\tilde{\mathcal{A}}$ from $\hat{\mathcal{A}}$. Using the approach outlined by Bookstein [92], the equation describing the surface of a thin material plate is used as an interpolant to define in-plane deformations based on the translation of arbitrary control points. Thus, to define a

warping of image $\hat{\mathcal{A}}$ onto image \mathcal{B} , two sets of control points \mathbf{c} and \mathbf{d} representing the same anatomical features on images \mathcal{A} and \mathcal{B} , respectively, are selected by manually identifying salient anatomical landmarks. The landmarks are structural features that are clearly visible on both modalities, such as large duct centers, the urethra, features along the outer boundary of the prostate and along the transition between peripheral and central zones. For any slice pair, $\{\mathbf{c}_1, \dots, \mathbf{c}_p\}, \{\mathbf{d}_1, \dots, \mathbf{d}_p\} \subset C$ where the number of control point pairs range from $7 \leq p \leq 12$. Anatomical control points are chosen on the common undeformed images such that precisely the same anatomical features are used to define the elastic deformations that follow the different affine alignments obtained by MI and COFEMI. The principal warp determined from the solution of the thin plate spline subject to the displacement of points \mathbf{c} and \mathbf{d} is applied to obtain the overall affine and elastic deformation fields \mathbf{M}^{χ_3} and \mathbf{M}^{χ_4} . Similarly, $\tilde{\mathcal{B}}$ may be obtained from $\hat{\mathcal{B}}$ if \mathcal{A} is desired to remain stationary. We refer to elastic registration following MI- and COFEMI-based affine registration as MI+TPS and COFEMI+TPS.

3.4 Supervised Computer-Aided Classification

3.4.1 Feature Extraction

We compute a total of $n = 311$ unique feature images from each MRI intensity image \mathcal{C}^M , the CAD-targeted modality, which corresponds to image \mathcal{B} in the generalized notation used throughout Section 3.3.3 above. These feature representations were chosen for their class-discriminative abilities as previously demonstrated in prostate MRI CAD [28]. These features were also chosen to compose the ensemble used in the COFEMI registration scheme described in Section 3.3.3 since several of these feature representations were found to be less susceptible to image artifacts such as bias field inhomogeneity [3]. Thus, for each \mathcal{B} , feature images $\mathcal{B}^{\Phi_\beta} = (C, f^{\Phi_\beta})$ are calculated by applying feature operators Φ_β , $\beta \in \{1, \dots, n\}$ within a local neighborhood associated with every $\mathbf{c} \in C$. Thus $f^{\Phi_\beta}(\mathbf{c})$ represents the feature value associated with feature operator Φ_β at pixel \mathbf{c} in \mathcal{B} , and $\mathbf{f} = [f^{\Phi_\beta}(\mathbf{c}) | \beta \in \{1, \dots, n\}]$ is the feature vector of length n for pixel \mathbf{c} .

Gradient Features

Gradient features are calculated using steerable and non-steerable linear gradient operators. Eleven non-steerable gradient features were obtained using Sobel, Kirsch and derivative operations. Gabor gradient operators [77] comprise the steerable class of gradient calculations. The magnitude of the Gabor filter response was defined for every pixel $\mathbf{c} = (x, y)$ by $f^{Gab}(\mathbf{c}) = e^{(x'+y')/2\pi} \cos(2\pi\lambda x')$ where λ is the spatial frequency (scale) of the sinusoid. The filter was applied on a $\kappa \times \kappa$ square neighborhood $\mathcal{N}_{\mathbf{c},\kappa}$ associated with every pixel $\mathbf{c} \in C$ and centered on \mathbf{c} . The orientation of the filter is affected by φ through the coordinate transformations: $x' = x \cos \varphi + y \sin \varphi$ and $y' = x \sin \varphi - y \cos \varphi$. Gabor gradient features were calculated for $\kappa \in \{5, 8, 15\}$, $\varphi \in \{0, \frac{\pi}{8}, \frac{\pi}{4}, \frac{3\pi}{8}, \frac{\pi}{2}, \frac{5\pi}{8}, \frac{3\pi}{4}, \frac{7\pi}{8}\}$ and $\lambda \in \{\frac{\pi}{2\sqrt{2}}, \frac{\pi}{4}, \frac{\pi}{4\sqrt{2}}, \frac{\pi}{8}, \frac{\pi}{8\sqrt{2}}, \frac{\pi}{16}\}$.

Statistical Features

Four first order statistical features, including mean, median, standard deviation, and range, and were calculated as described in [28]. To calculate the second order statistical feature images, we define any pixel $\mathbf{d} \in \mathcal{N}_{\mathbf{c},\kappa}$ as a κ neighbor of \mathbf{c} . We now compute a $M \times M$ co-occurrence matrix $P_{\mathbf{c},\kappa}$ associated with $\mathcal{N}_{\mathbf{c},\kappa}$, where M is the maximum gray scale intensity over all \mathcal{B} . The value at any location $[u, v]$, where $u, v \in \{1, \dots, M\}$, in the matrix $P_{\mathbf{c},\kappa}[u, v]$ represents the frequency with which two distinct pixels $\mathbf{c}, \mathbf{d} \in \mathcal{N}_{\mathbf{c},\kappa}$ with associated image intensities $f(\mathbf{c}) = u$, $f(\mathbf{d}) = v$ are adjacent (within the same 8-neighborhood within $\mathcal{N}_{\mathbf{c},\kappa}$). A total of 16 Haralick features including energy, entropy, inertia, correlation, inverse difference moment, two information correlation measures, sum average, sum variance, sum entropy, different average, difference variance, difference entropy, shade, prominence and variance were extracted at every pixel $\mathbf{c} \in C$, from $P_{\mathbf{c},\kappa}$, for $\kappa \in \{5, 8, 15\}$, and $M \in \{64, 128, 256\}$.




Figure 3.3: MRI intensity image (a) and corresponding independent component feature images obtained via ICA (b)-(h) of feature images (not shown). Maximal statistical independence between independent components is maximized to facilitate naïve combination in Bayes' rule.

3.4.2 Bayesian Classifier and Independent Component Analysis

The naïve form of Bayes' rule defines the likelihood of observing some class ω_1 given the multivariate feature vector $\mathbf{f}(\mathbf{c})$ of dimensionality n and is determined by,

$$P(\omega_1|\mathbf{f}(\mathbf{c})) = \frac{P(\omega_1)p(\mathbf{f}(\mathbf{c})|\omega_1)}{P(\omega_1)p(\mathbf{f}(\mathbf{c})|\omega_1) + P(\omega_0)p(\mathbf{f}(\mathbf{c})|\omega_0)}. \quad (3.4)$$

where the prior probabilities of occurrence of the two classes are $P(\omega_1)$ and $P(\omega_0)$, and $p(\mathbf{f}(\mathbf{c})|\omega_1)$ and $p(\mathbf{f}(\mathbf{c})|\omega_0)$ are the *a priori* class conditional distributions of $\mathbf{f}(\mathbf{c})$. Estimation of $p(\mathbf{f}(\mathbf{c})|\omega_1)$ and $p(\mathbf{f}(\mathbf{c})|\omega_0)$ is difficult or impossible when the distributions are of a high dimensionality. However, if the conditional probabilities of each dimension of $\mathbf{f}(\mathbf{c})$ can be made statistically independent by some linear transformation W (i.e. $\phi = W\mathbf{f}$) the distribution of the new random variable ϕ of dimensionality h can be written $p(\phi) = p(\phi_1, \dots, \phi_k) = \prod_i^h p(\phi_i)$. Hence, independent component analysis (ICA) is used to reduce the feature vector \mathbf{f} of size $n = 311$ to a vector $\phi = \{\phi_1, \dots, \phi_h\}$ of size $h \ll n$ with statistically independent distributions. That is, for the entire set of pixels C of image \mathcal{B} , ICA determines a $h \times n$ unmixing matrix W such that

the distributions of the rows of ϕ are maximally independent. The h independent components $\phi_i(\mathbf{c})$ of pixels $\mathbf{c} \in G^{X_r}$ are then used to generate *a priori* distributions $p(\phi_i(\mathbf{c})|\omega_1)$ for the target class ω_1 , while pixels $\mathbf{c} \in \{C - G^{X_r}\}$ are used to generate distributions for class ω_0 . Hence, Bayes' rule can then be used to obtain the *a posteriori* probability distribution $P(\omega_1|\phi(\mathbf{c}))$ of observing ω_1 for the linearly independent feature vector $\phi(\mathbf{c})$ at each pixel $\mathbf{c} \in C$ by

$$P(\omega_1|\phi(\mathbf{c})) = \frac{P(\omega_1) \prod_i^h p(\phi_i(\mathbf{c})|\omega_1)}{P(\omega_1) \prod_i^h p(\phi_i(\mathbf{c})|\omega_1) + P(\omega_0) \prod_i^h p(\phi_i(\mathbf{c})|\omega_0)}. \quad (3.5)$$

The probability images $\mathcal{L} = (C, P(\omega_1|\phi(\mathbf{c})))$, represent the joint *a posteriori* likelihood of $\mathbf{c} \in C$ belonging to class ω_1 obtained for each training image \mathcal{B} , which correspond to MRI slices (\mathcal{C}^M) in this study. Importantly, since the class of features represented by the individual feature operators Φ_β has been demonstrated to be discriminative in previous studies [28], the independent component distributions $P(\omega_1|\phi_i(\mathbf{c}))$ will contain discriminative potential.

The Jade ICA implementation [93] was used to employ a measure of independence in approximation to MI, a natural measure of statistical dependence of multiple distributions related to measures of non-Gaussianity that are commonly used in ICA algorithms. In this view of independence, the maximum number of independent component images that can be obtained from the feature set with a minimal degree of dependence between the components is determined via calculation of pairwise MI between the resulting independent component images as described in [94]. Estimation of further “independent” components beyond the maximum results in a sharp increase in the magnitude of MI between multiple components, indicating that additional independent sources are not identifiable in the feature set. In this manner, the value of h is determined. Figure 3.3(b)-(h) shows $h = 7$ independent component images derived from feature images of \mathcal{C}^M in Figure 3.3(a), each of which were tested for independence by paired MI calculations and inspection of the joint distributions.

3.5 Application to Prostate CAD

3.5.1 Supervised Classifier Training

The registration-based supervised classifier training technique described above was used to train a CAD system for detection of cancer in MRI images of the prostate (\mathcal{C}^M). In the registration procedure we treat \mathcal{C}^M as the stationary template image (\mathcal{A}) and \mathcal{C}^H as the target image (\mathcal{B}). The four effective affine and elastic registration techniques ($\chi_r \in \{\text{MI}, \text{MI+TPS}, \text{COFEMI}, \text{COFEMI+TPS}\}$) described in Section 3.3.3 were each used to establish corresponding region labels $G^{\chi_1}(\mathcal{C}^M)$, $G^{\chi_2}(\mathcal{C}^M)$, $G^{\chi_3}(\mathcal{C}^M)$ and $G^{\chi_4}(\mathcal{C}^M)$ on MRI. The region labels were each used to establish the membership of each pixel in \mathcal{C}^M to either the cancer or non-cancer class ($\mathbf{c} \hookrightarrow \omega_c$ or $\mathbf{c} \hookrightarrow \omega_{nc}$), and four supervised Bayesian classifiers Ψ_1 , Ψ_2 , Ψ_3 and Ψ_4 were trained as described in Section 3.4.2. The manual delineation of cancer $G^m(\mathcal{C}^M)$ was used to train a fifth Bayesian classifier Ψ_5 .

3.5.2 Evaluation of Supervised CAD Classifier

Since no definitive ground truth is available on \mathcal{C}^M , testing and evaluation of classifier accuracy was done using seven ground truth surrogates, which correspond to the five individual cancer labels ($G^{\chi_1}(\mathcal{C}^M)$, $G^{\chi_2}(\mathcal{C}^M)$, $G^{\chi_3}(\mathcal{C}^M)$ and $G^{\chi_4}(\mathcal{C}^M)$ and $G^m(\mathcal{C}^M)$), the union ($\bigcup G^\Omega$ where $\Omega \in \{\chi_1, \chi_2, \chi_3, \chi_4, m\}$) and the intersection ($\bigcap G^\Omega$) of the labels. These 7 ground truth estimates are represented as $\mathbf{\Gamma}_{1-7}$. For each classifier Ψ_{1-5} and ground truth estimate $\mathbf{\Gamma}_{1-7}$, totaling 35 combinations, receiver operating characteristic (ROC) curves are generated, and the area under the ROC curves (AUC) and positive predictive value (PPV) are calculated. The vertical axis of the ROC curve is the true positive rate (TPR) or sensitivity and the horizontal axis is the false positive rate (FPR) or 1-specificity. Each point on the curve corresponds to the sensitivity and specificity of detection of classifier Ψ_α , $\alpha \in \{1, \dots, 5\}$, at threshold ρ . If $P(\omega_c | \phi(\mathbf{c})) \geq \rho$ then \mathbf{c} is classified as ω_c else ω_{nc} to obtain a binary prediction $G_{\rho, \Psi}^{CAD}$ for threshold ρ and classifier

Ψ . For each $G_{\rho,\Psi}^{CAD}$ and $\mathbf{\Gamma}_{1-7}$, sensitivity (SN) and specificity (SP) are calculated as

$$SN_{\rho,\Psi,\Gamma} = \frac{|G_{\rho,\Psi}^{CAD} \cap \mathbf{\Gamma}|}{|G_{\rho,\Psi}^{CAD} \cap \mathbf{\Gamma}| + |\mathbf{\Gamma} - G_{\rho,\Psi}^{CAD} \cap \mathbf{\Gamma}|} \quad (3.6)$$

$$SP_{\rho,\Psi,\Gamma} = \frac{|(C - G_{\rho,\Psi}^{CAD}) \cap (C - \mathbf{\Gamma})|}{|(C - G_{\rho,\Psi}^{CAD}) \cap (C - \mathbf{\Gamma})| + |G_{\rho,\Psi}^{CAD} - \mathbf{\Gamma}|} \quad (3.7)$$

where $|\cdot|$ indicates the cardinality of a set of pixels. A 5-fold randomized cross-validation procedure is used to compare the accuracy of each Ψ_{1-5} with respect to each $\mathbf{\Gamma}_{1-7}$, whereby from the 26 image pairs with $G(\mathcal{C}^H)$ mapped to \mathcal{C}^M by each method, 21 MR images are randomly chosen from which a classifier is trained. The remaining images are used as a test set upon which classifier accuracy is evaluated. Randomized cross-validation is repeated 50 times for different training and testing slices \mathcal{C}^M .

Average ROC curves for each classifier were generated for display purposes by fitting a smooth polynomial and averaging the ROCs over each test image and $\mathbf{\Gamma}_{1-7}$. To avoid introducing bias toward classifiers derived from registration-established training data, which would occur if curves from all four $\mathbf{\Gamma}_{1-4}$ are averaged, each average curve was generated using a subset of $\mathbf{\Gamma}_{1-7}$: the expert determined ground truth ($\mathbf{\Gamma}_5$), the intersection and union ground truths ($\mathbf{\Gamma}_{6-7}$) and one automatically mapped ground truth estimate ($\mathbf{\Gamma}_{1-4}$ corresponding to the Ψ being evaluated). Values of AUC and PPV for each classifier pair were compared in a similar manner.

3.6 Results of Semi-automated Elastic Registration and CAD Evaluation

3.6.1 Qualitative Evaluation of Registration Schemes

For each of the 26 histological and MRI images, alignment of corresponding \mathcal{C}^M and \mathcal{C}^H slices and establishment of training instances in \mathcal{C}^M was performed using each of the four registration techniques. Figure 3.4 shows the results of the registration procedures for a single corresponding pair of images \mathcal{C}^H (Fig. 3.4(a)) and \mathcal{C}^M (Fig. 3.4(e)). The manual expert delineation of cancer on MRI ($G^m(\mathcal{C}^M)$) is shown in green in Fig.




Figure 3.4: (a) Histology with ground truth indicated by the black arrow is registered to (e) MRI to transfer histological cancer map to MRI. Resulting mappings obtained via MI (c), COFEMI+TPS (d) and COFEMI (f). Overlays of MRI with histology registered via MI and COFEMI+TPS shown in (g) and (h), respectively, indicate the more accurate alignment with elastic vs. affine deformations. (b) Manually determined labels on MRI are shown for comparison.

3.4(b). Histological ground truth $G(\mathcal{C}^H)$, indicated by the arrow in Fig. 3.4(a), is shown mapped to $G^{x_1}(\mathcal{C}^M)$, $G^{x_4}(\mathcal{C}^M)$ and $G^{x_3}(\mathcal{C}^M)$ in Figs. 3.4(c),(d) and (f). Comparing $G^{x_1}(\mathcal{C}^M)$ and $G^{x_3}(\mathcal{C}^M)$ as shown in Figs. 3.4(c) and (f) indicates similar results with fewer obvious mislabeled instances in $G^{x_3}(\mathcal{C}^M)$. However, the difference in accuracy of the registration methods is best illustrated as an overlay of \mathcal{C}^M with the whole transformed image \mathcal{C}^H as in Figs. 3.4(g) and (h) for the methods MI and COFEMI+TPS. Notice the considerable region of non-overlapping tissue in Figure 3.4(g) compared with the nearly ideal alignment in Figure 3.4(h). It is clear that a non-linear elastic registration procedure is superior in this task. Over all slice pairs, COFEMI was found to produce qualitatively more accurate alignments compared with MI, supporting previous results [3]. Further, the associated elastic methods MI+TPS and COFEMI+TPS produced the most qualitatively accurate and realistic alignments overall.

3.6.2 Quantitative Results

From the registration results, the supervised classifiers Ψ_{1-5} were generated using $h = 7$ independent component images obtained from the full set of feature images. Using the randomized cross-validation procedure described above, Ψ_{1-5} were compared for the ground truth estimates Γ_{1-7} in terms of PPV and AUC. Figure 3.5(a) shows average ROC curves for each classifier Ψ_{1-5} derived from the different training data labels ($G^{\chi_1}(\mathcal{C}^M)$, $G^{\chi_2}(\mathcal{C}^M)$, $G^{\chi_3}(\mathcal{C}^M)$, $G^{\chi_4}(\mathcal{C}^M)$ and $G^m(\mathcal{C}^M)$). The dashed line in Figure 3.5 represents the ROC curve for classifier Ψ_5 averaged over different CaP ground truth estimates as described in Section 3.5.2. Clearly, overall accuracy of the classifiers Ψ_{1-4} trained from registration-derived training labels is greater than the accuracy of Ψ_5 , which is trained using expert labels. Values for AUC and PPV were then calculated from the original ROC curves. Differences between classifiers Ψ_{1-5} are compared in terms of PPV values as illustrated in Figure 3.5(b), which shows the average PPV for each Ψ_{1-5} . The difference in PPV between Ψ_5 (expert) and each Ψ_{2-4} (MI+TPS, COFEMI, COFEMI+TPS) were statistically significant ($p < 0.05$) for each $\Gamma_{1-4,6-7}$, while for Γ_5 the values for Ψ_5 were not significantly greater than any Ψ_{1-4} . The AUC of each Ψ_{1-4} (MI, MI+TPS, COFEMI, COFEMI+TPS) was significantly ($p < 0.05$) greater than the AUC of Ψ_5 (expert) for the same Γ s as for PPV, while Ψ_5 again displayed no significant advantage for any Γ .

3.7 Concluding Remarks

In this paper we demonstrated the need for use of multi-modal registration methods for establishing accurate ground truth for classifier training. We demonstrate that even simple MI-based schemes are more accurate in mapping spatial extent of cancer from one modality to another compared to expert labeling in spite of inter-modal differences and intensity artifacts which hinder automated registration techniques. We quantitatively demonstrated with the use of 4 different registration methods (MI, COFEMI, MI+TPS, COFEMI+TPS) to align 26 pairs of *ex vivo* MRI with histology images of the prostate upon which cancer ground truth is available and provide improved registration accuracy

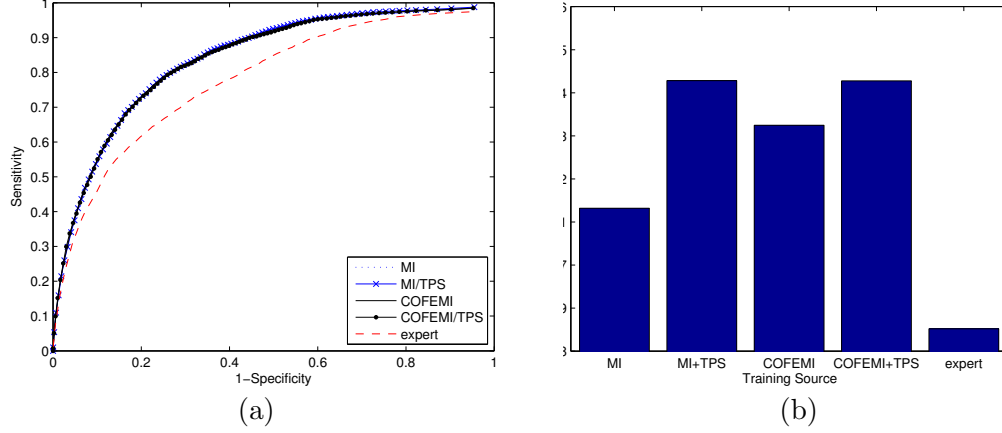


Figure 3.5: Accuracy of classifiers Ψ_{1-5} represented by the average ROC curves in (a). Quantitative comparison of Ψ_{1-5} in terms of AUC (not shown) indicates a statistically significant improvement from training labels established by registration vs. expert delineation. Comparison in terms of (b) PPV suggests an improvement in labeling using elastic registration.

compared to expert delineations. The 5 labels for CaP were used to train 5 classifiers and these were evaluated against 7 different ground truth surrogates. A quantitative performance improvement in the supervised classifier component of a CAD system was observed using the cancer labels obtained by registration versus labels manually determined by an trained expert pathologist with a histological reference.

This method for training of a supervised classifier by multimodal registration is applicable to any scenario in which multimodal data with ground truth is available. By providing accurate and repeatable training labels, registration overcomes several of the problems with manual delineation of training labels including human error, observer bias and variability. Further, the method of comparing classifier performance where the classifiers have been trained on the registration results may be used as a means of comparing registration performance. Some steps in this direction were undertaken in this paper.

Chapter 4

Registration of *ex vivo* Whole-Mount Histology and *in vivo* MRI using Multiple Texture Features

4.1 Challenges of Automated Elastic Multimodal Image Registration

Multimodal and multiprotocol image registration refers to the process of alignment of two images obtained from different imaging modalities (e.g. digitized histology and MRI) and protocols (e.g. T2-weighted and PD MRI), utilizing either rigid or non-rigid coordinate system transformations. Both processes are critical components in a range of applications, including image guided surgery [9–11], multimodal image fusion for cancer diagnosis and treatment planning [12], and automated tissue annotation [13]. However, registration of multimodal imagery has posed a more challenging task compared with alignment of images from the same modality or protocol on account of differences in both image intensities and shape of the underlying anatomy. The first of these hindrances, dissimilar intensities between modalities, arises as a consequence of the measurement of orthogonal sources of information such as functional (SPECT) and structural (CT/MRI) imagery [12], as well as on account of other factors such as intensity artifacts, resolution differences, and weak correspondence of observed structural details. We have previously addressed these challenges in the context of rigid registration using our feature-driven registration scheme termed combined feature ensemble mutual information (COFEMI) [3, 53]. The goal of the COFEMI technique is to provide a similarity measure that is driven by unique low level textural features, for registration that is more robust to intensity artifacts and modality differences than similarity measures restricted to intensities alone. For example, the multiprotocol MRI in Fig. 4.1 which contains strong bias field artifact on T1 MRI are registered using both conventional intensity-based MI and with COFEMI. The features in Figs. 4.1(e)

and (e) clearly demonstrate robustness to artifacts, and hence provide improved registration with COFEMI as in Fig. 4.1(f). We refer the reader to [3] for demonstration and further description of the technique.




Figure 4.1: Comparison of MI and feature-driven COFEMI rigid registration of images with strong bias field inhomogeneity artifacts. (a) A T2 MR brain image is registered to (b) the corresponding T1 MRI using (c) intensity-based MI and (f) COFEMI using second order (d) correlation and (e) inverse difference moment features. Green contours in (c) and (f) represent the boundary of the T2 brain MRI of (a) overlaid onto the registered target. Red outlines accentuate the boundaries in the registration result. Use of textural feature images by COFEMI was shown to improve registration of multiprotocol images with heavy intensity artifacts.

While accurate rigid registration is a valuable precursor to more complex transformations, and rigid image transformations are often sufficient to model many deformations in biomedical imagery, non-linear shape differences are common between real multimodal biomedical image data sets. For example, registration of images of highly deformable tissues such as in the breast have been shown to require flexible non-rigid techniques [17]. Similarly, non-linear differences in the overall shape of the prostate between *in vivo* MRI and *ex vivo* whole mount histology (WMH) have been shown to exist as a result of (1) the presence of an endorectal coil during MR imaging and (2)

deformations to the histological specimen as a result of fixation and sectioning [18, 19]. Consequently, achieving correct alignment of such imagery requires elastic transformations to overcome the non-linear shape differences. The free form deformation (FFD) technique proposed by Rueckert in [17] has been demonstrated to provide a flexible automated framework for non-rigid registration by using any similarity measure to drive registration. However, this technique relies upon intensity-based similarity measures, which have been shown to be wanting for robustness across highly dissimilar modalities and in the presence of artifacts [3]. Thin plate splines (TPS) warping methods are common, but involve identification of anatomical fiducials, a difficult task that is usually performed manually.

To overcome the challenges of both non-linear deformations and intensity artifacts simultaneously, we present a new technique termed Feature Ensemble Multi-level Splines (COLLINARUS). Our new COLLINARUS non-rigid registration scheme offers the robustness of COFEMI to artifacts and modality differences, while allowing fully automated non-linear image warping at multiple scales via a hierarchical B-spline mesh grid optimization scheme. An overview of the registration methodology used in this paper to demonstrate COLLINARUS is presented in Fig. 4.2, whereby feature ensembles drive both rigid and non-rigid registration of an intensity image that is the target for transformation, onto a template intensity image that remains stationary. As previously described [3], COFEMI is used to drive an initial rigid registration step to correct large scale translations, rotations, and differences in image scale. The transformed target intensity image that results from rigid registration is then registered in a non-linear fashion via COLLINARUS to the template image. Registration by COLLINARUS is critical to account for local deformations that cannot be modeled by any linear coordinate transformation. Since COFEMI and COLLINARUS involve maximization of a similarity measure, each step is fully automated.

We developed the COLLINARUS scheme to perform an automated tissue annotation task that is designed to facilitate the development and evaluation of a novel system for computer-assisted detection (CAD) of prostate cancer on multi-protocol MRI [28]. The development of a multimodal CAD system that operates upon *in vivo* imagery requires




Figure 4.2: A two step COFEMI-driven rigid and non-rigid registration methodology applied in this study to perform automated alignment of two intensity images. Initial global alignment is performed using COFEMI to optimize an affine transformation of the target intensity image. Subsequently, non-rigid registration via COLLINARUS is performed to determine the remaining local deformations.

ground truth labels for cancer on each modality to characterize malignant tissue. Since these MRI pixel labels are usually obtained by manual delineation of cancer, they can be extremely time consuming to generate and subject to errors and bias of the expert performing the annotation. The deleterious effect of such errors in training labels on MRI CAD has been demonstrated [29, 88]. Therefore, to improve labeling and hence CAD classifier accuracy, alignment of *in vivo* imagery with corresponding *ex vivo* whole mount histology (WMH), the source of the cancer “gold standard”, may be performed via automated multimodal image registration. The use of COFEMI for automated rigid registration has been previously demonstrated on *ex vivo* MRI [5]. In the current study, we present the non-rigid spatial registration of *in vivo* T2-w MRI, *in vivo* DCE MRI, and *ex vivo* whole mount prostate histology slices, followed by mapping of the “gold standard” from histology onto both MRI protocol images. A diagram of the multimodal prostate registration task performed in this paper is shown in Fig. 4.3. The *ex vivo* WMH containing the “gold standard” label for cancer shown in Fig. 4.3(a) is registered to the corresponding *in vivo* T2-w MRI section via the COLLINARUS non-linear registration technique. The transformed WMH section shown in the bottom of Fig. 4.3(a) contains a cancer map that is then transferred directly onto the T2-w MRI as illustrated at the bottom of Fig. 4.3(b). The new non-rigid COLLINARUS

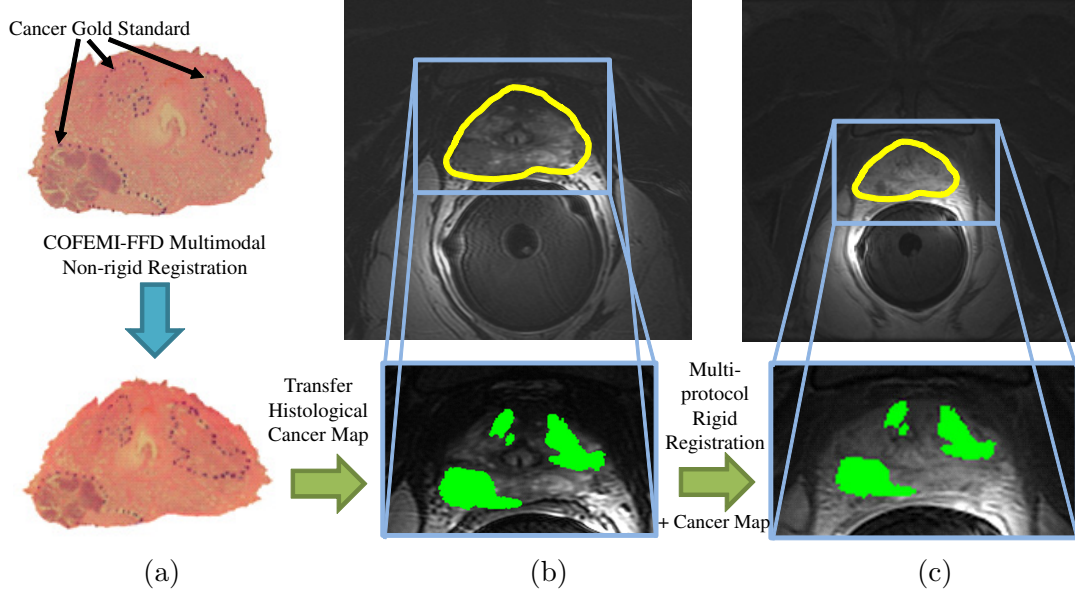


Figure 4.3: Application of the COLLINARUS automated feature driven non-rigid registration technique to alignment of (a) *ex vivo* whole mount histology (WMH), (b) *in vivo* T2-w MRI and (c) *in vivo* DCE-MRI images of the prostate and annotation of cancer on multiprotocol MRI. (a) Histopathologic staining of whole-mount sections of a prostate with cancer provides the “gold standard” for cancer extent. Non-rigid registration via COLLINARUS of (a) WMH to (b) corresponding *in vivo* MRI obtained prior to resection allows the histological cancer map to be transferred onto (b). (c) Corresponding DCE-MRI is registered to (b) by COFEMI rigid registration to establish a map of cancer on (c).

registration technique will overcome the limitations of rigid deformation models, while providing similar improvements in efficiency and accuracy of cancer delineation on *in vivo* multiprotocol MRI. This will allow the structural appearance and functional properties of cancer to be accurately characterized for the development and evaluation of *in vivo* multiprotocol CAD applications.

Qualitative and quantitative evaluation of the COLLINARUS scheme is performed on a set of real multimodal prostate images and on a set of synthetic multiprotocol brain images. Multimodal prostate image registration is performed as described above for 6 clinical data sets comprising a total of 21 groups of *in vivo* T2-w MRI, DCE-MRI, and *ex vivo* WMH images with cancer present. Evaluation of prostate registration is performed by comparison with 3TP DCE mappings, the industry standard for DCE-MRI analysis, and by measures of prostate region overlap. The set of synthetic multiprotocol images, acquired from the BrainWeb Simulated Brain Database [82], comprises 11

pairs of T1-w and proton density (PD) MRI of the brain. The synthetic registration task was generated by applying a known non-linear warping to the PD MRI, hence misaligning T1-w MRI from PD MRI. Non-rigid registration was then performed to recover the original, correct alignment of each image pair. Quantitative evaluation of brain registration was performed by direct comparison of (1) the recovered deformation field to the applied field and (2) the original undeformed and recovered PD MRI. For each of the data sets, COLLINARUS is compared with the MI-driven counterpart of the B-spline technique.

The primary novel contributions of this work are,

- A new method termed COLLINARUS for automated non-rigid image registration.
- Use of textural feature image ensembles in a non-rigid registration technique for robustness to artifacts and modality differences.
- A multimodal rigid and non-rigid registration scheme that provides superior registration accuracy compared to use of MI-driven counterparts.

The rest of the paper is organized as follows. In Section 4.2, the COLLINARUS registration technique is described. In Section 4.3, the results of the real and synthetic registration tasks are described for both COLLINARUS and MI-MLS. Concluding remarks are given in Section [82].

4.2 Collection of Image-derived Non-linear Attributes for Registration Using Splines (COLLINARUS)

4.2.1 Overview

The new registration technique referred to as Collection of Image-derived Non-linear Attributes for Registration Using Splines, or COLLINARUS, consists of three primary components,

1. A robust feature ensemble-driven similarity measure derived the COFEMI [3] scheme,

2. A flexible non-linear image warping model based on B-splines and,
3. A variable spline grid size approach for optimizing a multi-scale local image warping.

4.2.2 Notation

Define a stationary template image as $\mathcal{A} = (C, f^A)$, where C is a set of coordinates $c \in C$ and $f^A(c)$ is the intensity value of \mathcal{A} at location c . A target image $\mathcal{B} = (C, f^B)$ is similarly defined with intensities $f^B(c)$ on the same coordinate grid C . The goal of the registration task is to provide a coordinate transformation $\mathbf{T}(c)$, $\forall c \in C$, that describes the mapping of each point on a registered target image \mathcal{B}^r to the template intensity image \mathcal{A} . We can then define $\mathcal{B}^r = (C, f^*(c))$, where $f^*(c) = g(\mathbf{T}(c), f^B)$ represents an interpolation function used to provide intensity values at location $\mathbf{T}(c)$ using the underlying intensity map f^B . We can further define a generic image transformation Φ to represent the application of \mathbf{T} at each $c \in C$, such that $\mathcal{B}^r = \Phi(\mathcal{B}, \mathbf{T})$.

4.2.3 General Registration Framework

We demonstrate COLLINARUS in a two stage rigid and non-rigid registration framework, whereby COFEMI and COLLINARUS are used in the rigid and non-rigid components, respectively. As described in Fig. 4.2, registration of \mathcal{A} and \mathcal{B} may be performed by first determining a global, rigid transformation \mathbf{T}_{rigid} , followed by an local, elastic transformation $\mathbf{T}_{elastic}$. The global transformation is determined by maximizing,

$$\mathbf{T}_{rigid} = \underset{\mathbf{T}}{\operatorname{argmax}} \psi(\mathcal{A}, \Phi(\mathcal{B}, \mathbf{T})), \quad (4.1)$$

where ψ is an image similarity measure such as conventional intensity-driven MI or the feature ensemble-driven measure from COFEMI. Application of \mathbf{T}_{rigid} to \mathcal{B} gives the rigidly registered target image \mathcal{B}^r by,

$$\mathcal{B}^r = \Phi(\mathcal{B}, \mathbf{T}_{rigid}). \quad (4.2)$$

The elastic transformation $\mathbf{T}_{elastic}$ and the final registered target image \mathcal{B}^{nr} are then determined, again by maximization of the similarity measure ψ , by,

$$\mathbf{T}_{elastic} = \underset{\mathbf{T}}{\operatorname{argmax}} \psi(\mathcal{A}, \Phi(\mathcal{B}^r, \mathbf{T})), \text{ and} \quad (4.3)$$

$$\mathcal{B}^{nr} = \Phi(\mathcal{B}^r, \mathbf{T}_{elastic}). \quad (4.4)$$

A unified coordinate transformation may then be defined as the successive application of the coordinate transformations \mathbf{T}_{rigid} and $\mathbf{T}_{elastic}$,

$$\mathbf{T}^*(c) \equiv \mathbf{T}_{elastic}(\mathbf{T}_{rigid}(c)) \quad (4.5)$$

and the non-rigid registration result generated directly by,

$$\mathcal{B}^{nr} = \Phi(\mathcal{B}, \mathbf{T}^*) \quad (4.6)$$

For the implementation of the above methodology used in this study, we define \mathbf{T}_{rigid} as an affine coordinate transformation. Details of the multi-scale optimization of $\mathbf{T}_{elastic}$ are described in the following section.

4.2.4 COLLINARUS Non-Rigid Registration

The COLLINARUS non-rigid registration technique achieves the optimization of $\mathbf{T}_{elastic}$ in Eqn. (4.3) by synergy of the following concepts,

1. Using a feature ensemble-driven similarity measure for ψ , obtained via the techniques for feature extraction, selection, and combination that are associated with COFEMI.
2. Defining $\mathbf{T}_{elastic}$ as the 2-D tensor product of the cubic B-splines [41,95] to allow local elastic image deformations.
3. Maximization of ψ using a multi-level control point grid approach to achieve B-spline deformations at multiple scales.

The feature ensemble-driven similarity measure used for ψ is obtained by the COFEMI technique. The primary components of COFEMI are (1) extraction of an exhaustive set of low level textural feature images, (2) selection of a highly descriptive ensembles of textural features from the intensity images using the *CMIfeatures* algorithm described in [3], and (3) incorporation of the feature ensembles by combined mutual information (CMI), a form of MI derived to compare two multivariate observations. The measure ψ used by COLLINARUS is thus defined as in [3] by,

$$\psi(\mathcal{A}, \mathcal{B}) = CMI(\mathcal{A}\pi^A, \mathcal{B}\pi^B), \quad (4.7)$$

where π^A and π^B are the selected feature ensembles, and each $\mathcal{A}\pi^A$ and $\mathcal{B}\pi^B$ also represent distinct ensembles. Thus, the COLLINARUS non-rigid registration scheme involves optimization of $\mathbf{T}_{elastic}$ by Eqn. (4.3) via (4.7), whereas the COFEMI rigid registration scheme involves optimization of \mathbf{T}_{rigid} by Eqn. (4.1) via (4.7).

By defining $\mathbf{T}_{elastic}$ for COLLINARUS in terms of the cubic B-splines basis functions, COLLINARUS is capable of defining local elastic deformations without the use of anatomical fiducial markers. The coordinate transformation used by COLLINARUS is instead defined in terms of a regularly spaced control point mesh of size $n_x \times n_y$, the displacements of which are used to define a coordinate transformation $\mathbf{T}_{elastic}$ according to the 2D tensor product of B-spline basis functions [41].

The multiresolution image warping method employed by COLLINARUS is achieved by a multi-level spline grid optimization approach, whereby the number of grid points are modulated via n_x and n_y . Spline deformations defined with successively finer control point meshes are then combined into a single non-linear transformation. The idea behind this approach is to exploit the local neighborhood influence of B-splines grid so as to model successively smaller and more local deformations. This technique is thus capable of modeling deformations of varying magnitude. For a total of L transformations, $\mathbf{T}_{elastic}^l$ is defined at multiple control point spacing levels $l \in \{1, \dots, L\}$ with corresponding mesh sizes $n_{x,l} \times n_{y,l}$. At each level l , the transformation $\mathbf{T}_{elastic}^l$ is determined as in Eqn. 4.3 by maximization of similarity measure ψ , where the displacements

of each of the $n_{x,l} \times n_{y,l}$ control points are the free parameters. Each $\mathbf{T}_{elastic}^l$ is applied successively to form the final elastic transformation $\mathbf{T}_{elastic}$.

4.2.5 Registration Evaluation

Evaluation of registration accuracy can be performed easily if the correct coordinate transformation, \mathbf{T}^{ap} , is known. First, the magnitude of error in the transformation \mathbf{T}^{co} determined by registration can be quantified in terms of mean absolute difference (MAD) ($F_{mad}(\mathbf{T}^{co})$) and root mean squared (RMS) error ($F_{rms}(\mathbf{T}^{co})$) from \mathbf{T}^{ap} ,

$$F_{mad}(\mathbf{T}^*) = \frac{1}{N} \sum_{c \in C} \|\mathbf{T}^{co}(c) - \mathbf{T}^{ap}(c)\|_{L2} \quad (4.8)$$

$$F_{rms}(\mathbf{T}^*) = \sqrt{\frac{1}{N} \sum_{c \in C} \|\mathbf{T}^{co}(c) - \mathbf{T}^{ap}(c)\|_{L2}^2}. \quad (4.9)$$

Further, the desired transformed target image \mathcal{B}' may be obtained by from the known correct transformation \mathbf{T}^{ap} by,

$$\mathcal{B}' = \Phi(\mathcal{B}, \mathbf{T}^{ap}), \quad (4.10)$$

and compared directly with the resulting target image \mathcal{B}^{nr} actually obtained from registration. Three measures are used to compare \mathcal{B}^{nr} with \mathcal{B}' , including L2 distance (D_{L2}), MI (S_{MI}), and entropy correlation coefficient (S_{ECC}) [38].

4.3 Results of Automated Elastic Image Registration

4.3.1 Data Sets

Synthetic Data. Synthetic brain MRI [82] were obtained from BrainWeb, comprising corresponding simulated T1-w and PD MR brain volumes of dimensions $181 \times 217 \times 181$ with voxel size of 1mm^3 . We denote the T1-w and PD MRI images as \mathcal{C}^{T1} and \mathcal{C}^{PD} , respectively. Ground truth for correct alignment between \mathcal{C}^{T1} and \mathcal{C}^{PD} is implicit in the simulated data, allowing use of the evaluation methods described in Sec. 4.2.5.

Clinical Data. Clinical *in vivo* multiprotocol (T2-w and DCE) 3T MRI images with

WMH sections of the prostate were acquired to establish a map for spatial extent of cancer on both T2-w MRI and DCE-MRI. For 6 clinical data sets comprising *in vivo* T2-w MRI, DCE-MRI, and WMH, a total of 21 corresponding images with cancer present were considered. Cancer extent on histology was first obtained via microscopic analysis of haematoxylin and eosin stained tissue. Slices of T2-w MRI that correspond with the available WMH sections were identified by visual inspection by an expert radiologist. The slices of the T2-w MRI and DCE-MRI volumes are in implicit correspondence (but not 2D alignment) since the multiprotocol MRI data was acquired during a single scanning session in which the patient remained stationary.

4.3.2 Synthetic Brain Data

The synthetic T1-w and PD MRI brain data was used to perform quantitative analysis of registration accuracy under simulated noise and intensity inhomogeneity. Since the \mathcal{C}^{T1} and \mathcal{C}^{PD} images generated by the BrainWeb MRI simulator are in implicit alignment, evaluation of registration accuracy was performed as described in Sec. 4.2.5 by imposing a known transformation T' to each coordinate of \mathcal{C}^{PD} , followed by execution of COLLINARUS to determine the transformation T^* required to recover the original alignment. For 11 pairs of corresponding \mathcal{C}^{T1} and \mathcal{C}^{PD} images, registration was performed using COLLINARUS and an analogous MI-driven B-spline registration scheme. Fig. 4.4 demonstrates the registration of one T1-w MRI section with a PD MRI section. The T1-w MRI in Fig. 4.4(a) is initially in alignment with the PD MRI in 4.4(b), which contains noise and simulated field inhomogeneity. The non-linear deformation illustrated in Fig. 4.4(c) by the deformed grid is then applied to generate the deformed PD MRI in Fig. 4.4(d). MI-driven B-spline registration is then performed to obtain the PD MRI image in Fig. 4.4(e). A textural feature calculated by COLLINARUS from Fig. 4.4(b) is shown in Fig. 4.4(g), demonstrating the diminished effect of inhomogeneity on the feature image. The registration result from COLLINARUS is shown in Fig. 4.4(h). While the MI-based and COLLINARUS results in Figs. 4.4(e) and (h) appear similar, deformation field error magnitude images shown in Fig. 4.4(f) and (i) clearly indicate that T^* obtained from COLLINARUS contains far less error than the




Figure 4.4: (a) Synthetic T1-w MRI section, and (e) corresponding PD MRI section with simulated noise and bias field inhomogeneity artifacts. A deformation field, demonstrated on a grid in (f), is applied to (e) PD MRI to generate (b) a warped PD MRI section. Both MI and COFEMI are used to drive a non-linear B-spline based registration of (b) to (a). A correctly transformed PD MRI section would closely resemble (f). The results of (e) MI-driven and (g) COLLINARUS registration appear similar, while representations of the deformation field error magnitudes in (d) and (h) illustrate the greater error of MI compared with COFEMI.

transformation obtained MI spline registration.

The quantities F_{mad} and F_{rms} are calculated from \mathbf{T}^* obtained from COLLINARUS and MI-MLI from Eqns. (4.9), as well as D_{L2} , S_{MI} , and S_{ECC} . The average values of each quantity for $n = 11$ image pairs are given in Table 4.1, along with p -values for student's t-tests. The average values of both F_{mad} and F_{rms} were significantly lower for COLLINARUS, indicating less error in the recovered deformation field determined by \mathbf{T}^{co} . The average values of both S_{MI} , and S_{ECC} were significantly higher for COLLINARUS, indicating greater similarity between the recovered PD MRI and the correct undeformed image. Similarly the distance D_{L2} was lower for COLLINARUS, indicating greater similarity between COLLINARUS recovered and correct PD images.

	F_{mad}	F_{rms}	D_{L2}	S_{MI}	S_{ECC}
MI-MLS	0.9585	2.2201	1.88e+03	2.8803	0.5104
COLLINARUS	0.8330	1.9406	1.56e+03	3.0709	0.5437
p ($n = 11$)	0.0075	0.0097	9.97e-05	8.76e-05	9.97e-05

Table 4.1: Comparison of non-rigid registration accuracy for COLLINARUS and MI-MLS alignment of $n = 11$ pairs of synthetic PD MRI and T1-w MRI brain images. Error of recovered deformation field in terms of mean absolute difference (F_{mad}) and root mean squares (F_{rms}) measures the deviation of the registration-derived deformation field from the known field. Units of F_{mad} and F_{rms} are mm. Euclidean distance between the original undeformed PD MRI and recovered PD MRI obtained by non-rigid registration (D_{L2}), measures the dissimilarity between the registration result and the ideal result. The mutual information and entropy correlation coefficient between the recovered and original PD MRI sections (S_{MI} and S_{ECC}) indicate how well the recovered image resembles the original, ideal result. Each of F_{mad} , F_{rms} , D_{L2} , S_{MI} and S_{ECC} indicate that COLLINARUS more accurately recovered the original undeformed PD MRI compared with MI-MLS.

4.3.3 Clinical Multi-Modal Prostate Data

Prostate Registration Task

The “gold standard” for cancer presence, which is available on the whole mount histological (WMH) images, is mapped onto both *in vivo* T2-w MRI and DCE-MRI by alignment of each of the modalities. In this task, large differences in the overall shape of the prostate exist between WMH and *in vivo* MRI as a result of (1) the presence of an endorectal coil during MR imaging and (2) deformations to the histological specimen as a result of fixation and sectioning. Consequently, achieving correct alignment of WMH and MRI requires elastic transformations to overcome the non-linear shape differences. Thus, a multi-stage rigid and non-rigid registration procedure utilizing COLLINARUS was implemented to align the WMH, T2-w MRI, and DCE-MRI. The main steps are described below:

1. Initial affine registration of the WMH target image to the *in vivo* T2-w MRI image via the COFEMI multiple feature-driven registration technique.
2. Non-rigid registration of rigidly registered WMH image from step 1 onto T2-w MRI using the COLLINARUS technique.
3. Combine the resulting affine and non-rigid transformations, mapping pixels from

WMH onto the T2-w MRI.

4. Affine registration of multiprotocol images (T2-w MRI and DCE) via maximization of mutual information (MI), bringing all modalities and protocols in to spatial alignment.

In step 2, the B-spline derived warping from COLLINARUS allows for modeling of the local deformations that result from the presence of the endorectal coil required for high resolution *in vivo* MRI of the prostate. On the other hand, since T2-w MRI and DCE-MRI were acquired during the same scanning session, only a rigid transformation was required in step 4 to compensate for resolution and bounding box differences, as well as any small patient movements that may have occurred between acquisition of the two protocol volumes. The combined non-linear transformation obtained in step 3 was applied to the histopathologic cancer label, hence bringing the label into the coordinate frame of T2-w MRI. The final determined affine transformation was then applied to the histopathologic cancer label on T2-w MRI, thus generating the label for cancer extent on DCE-MRI.

Three Time Point (3TP) DCE Cancer Maps

The commonly used 3 time point (3TP) representation of the DCE-MRI can provide an independent estimate of cancer extent against which the registration-established cancer masks are compared. Most current efforts in computer-aided diagnosis of CaP from DCE-MRI involve pharmacokinetic curve fitting such as in the 3 Time Point (3TP) scheme [96]. Based on the curve/model fits these schemes attempt to identify wash-in and wash-out points, i.e. time points at which the lesion begins to take up and flush out the contrast agent. Lesions are then identified as benign, malignant or indeterminate based on the rate of the contrast agent uptake and wash out. Red, blue and green colors are used to represent different classes based on the ratio $w = \frac{\text{Rate of wash-in}}{\text{Rate of wash-out}}$ of the contrast agent uptake. When w is close to 1, the corresponding pixel is identified as cancerous area (red), when w is close to zero, the pixel is identified as benign (blue), and green pixels are those are identified as indeterminate.

Prostate Registration Results

Corresponding sections from a WMH slice, *in vivo* T2-w MRI, and a single time frame of DCE-MRI imagery from the same patient are shown in Figures 4.5(a) and (b) for a prostate with cancer, where the boundary of the prostate on T2-w MRI has been outlined in yellow. Using the automated non-rigid registration method described above, the WMH section in Fig. 4.5(a) is warped into alignment with the prostate region in Fig. 4.5(b). The cancer extent is then mapped directly onto T2-w MRI, as shown in Figure 4.5(d) to display the prostate more clearly. Having established cancer extent on T2-w MRI (Fig. 4.5(d)), the T2-w MRI image is registered to the DCE-MRI image. Finally, the CaP extent is mapped from T2-w MRI (Fig. 4.5(d)) onto DCE-MRI, as in Figure 4.5(e). For visual comparison, we calculate the 3TP color representation from the DCE time series, as shown in Fig. 4.5(f) for the same slice, providing an independent means of evaluating the CaP labels mapped by our registration technique. Prostate overlap between modalities and protocols, and comparison of mapped cancer extent (green) with 3TP cancer extent (red) indicates excellent overall alignment between modalities and protocols obtained by COLLINARUS, and an accurate mapping of cancer on images from both MRI protocols.

Since the correct transformation required to bring the images into alignment is not known, registration accuracy is evaluated in terms of how well the region of the images representing the prostate overlaps between the aligned images (the overlap ratio). The intuition of this measure is that if the prostate image regions completely overlap between the modalities, the registration and hence the cancer mapping is highly accurate. The overlap ratio for two images of the prostate is defined as the ratio of the number of pixel coordinates that represent the prostate in both modalities to the total number of pixel coordinates representing the prostate in either modality (i.e. the prostate region intersection-to-union pixel count ratio). The mean overlap ratio for the prostate in the pairs of registered WMH and T2-w MRI was 0.9261, indicating that prostate in the two aligned modalities occupies nearly the same spatial coordinates. Similarly for the pairs of registered T2-w MRI and DCE-MRI images, a high overlap of 0.8964 was achieved.




Figure 4.5: Registration of WMH to 3T in vivo T2-w MRI and DCE-MRI of the prostate with cancer. (a) WMH with cancer extent delineated (dotted lines) is registered to (b) corresponding T2-w MRI using COLLINARUS to generate (c) transformed WMH in spatial alignment with the prostate in (b). (d) T2-w MRI with the cancer extent mapped from (c) superimposed in green. (e) DCE-MRI registered with (d) by an affine transformation, shown with cancer extent (green) mapped from (d) T2-w MRI. (f) The commonly used 3TP representation of the DCE data in (e), which provides an independent estimate of CaP extent, demonstrates that the cancer extent mapped by registration is accurate. (g)-(l) Similar results are demonstrated for a different set of multimodal prostate images from another study.

4.4 Concluding Remarks

We have demonstrated a new method for fully automatic non-rigid multimodal/multiprotocol image registration that combines textural feature ensembles in a similarity measure to drive a multilevel B-spline based image warping scheme. The robustness to modality differences offered by the COFEMI technique for feature ensemble selection and combination, and the flexibility of B-splines to model non-linear deformations are leveraged in this study to provide a powerful tool for automated multimodal image registration. Our method was used to successfully register a unique data set comprising WMH, in vivo T2-w MRI, and in vivo DCE-MRI images of the prostate, and subsequently map histopathologic CaP extent onto the images from both in vivo MRI protocols. A comparison of the CaP labels mapped onto DCE-MRI with the independent 3TP representation suggests that the labels are established accurately by the registration procedure. We have thus presented a robust, accurate means for aligning and thus facilitating fusion of structural and functional data. The primary contributions of our method are, The primary novel contributions of this work are,

- A new method termed COLLINARUS that provides flexible and robust automated non-rigid multimodal image registration.
- Use of textural feature image ensembles to drive a non-rigid registration technique and provide robustness to artifacts and modality differences.
- Superior non-rigid registration accuracy compared with similar MI-driven techniques.
- Application of our technique in automatically determining the spatial extent of CaP by registration of multiprotocol 3T in vivo clinical MRI images of the prostate, with histology containing cancer ground truth. Our work presented here represents the first time a fully automated technique has been presented for registration of histology with *in vivo* MRI of the prostate.

Our experiments demonstrated on clinical and synthetic data indicated that COLLINARUS provides greater registration accuracy compared with similar MI-driven techniques. As such, the COLLINARUS technique will have broad applicability for automated registration of multimodal and multiprotocol images, including structural and functional data.

Chapter 5

Registration of *ex vivo* Whole-Mount Histology and *in vivo* MRI using Multiparametric Images

5.1 Background and Motivation

Recently, magnetic resonance (MR) imaging (MRI) has emerged as a promising modality for detection of prostate cancer (CaP), with several studies showing that 3 Tesla (T) endorectal *in vivo* T2-weighted (T2-w) imaging yields significantly higher contrast and resolution compared to ultrasound (US) [97]. For example, Fig. 5.1(a) shows a typical *in vivo* US image of a prostate in which internal anatomical details, such as the urethra, ducts, and hyperplasia, are barely discernible, while in the segmented T2-w MR image shown in Fig. 5.1(b), internal anatomical details within the prostate are clearly visible. An additional advantage offered by MRI is the ability to use different acquisition protocols to capture orthogonal sources of information, including functional (dynamic-contrast enhanced (DCE)), metabolic (magnetic resonance spectroscopy (MRS)), vascular (diffusion weighted imaging (DWI)), and structural (T2-w) attributes. Since multiple protocols can be acquired in the same scanning session, little additional setup time is required.

The use of multiprotocol MRI for CaP diagnosis has been shown to improve detection sensitivity and specificity compared to the use of a single MR imaging protocol [98–100]. Previous studies have demonstrated improved CaP detection sensitivity and specificity by simultaneous use of multiple MRI protocols, including DCE and T2-w MRI [101], MRS and T2-w MRI [102], and DWI with both T2-w [103] and DCE MRI [104]. Since the current clinical diagnostic protocol involves no image-based detection of CaP, the ability to utilize *in vivo* multiprotocol diagnostic images for detection and localization of CaP *in vivo* would have clear implications for (1) non-invasive

image-based screening, (2) targeted biopsies, and (3) conformal radiation therapy.

If the spatial extent for CaP on multiprotocol *in vivo* radiological imaging can be accurately delineated, it may then be possible to define specific imaging parameters with the greatest diagnostic accuracy in reliably characterizing CaP on *in vivo* clinical, radiologic images. The definition of such image signatures would be invaluable in building (a) a computer-assisted disease detection system [5, 28, 102, 105], or (b) spatial disease atlases which could serve as training and educational tools for medical students, radiology residents, and fellows. However, direct annotation of disease extent on MRI is often challenging even for experienced radiologists. Thus, to reliably ascertain extent of CaP on *in vivo* radiological images, it is necessary to utilize *ex vivo* tissue specimens, upon which “ground truth” estimates of CaP extent may be established by histopathologic inspection. (In the context of patients diagnosed with CaP and scheduled for radical prostatectomy (RP), in several centers in the U.S., pre-operative imaging is performed to identify presence of extracapsular spread. [106]) Figure 5.1(c) shows a whole-mount histology (WMH) section of a RP specimen on which cancerous tissue has been manually annotated (dark purple) following microscopic examination of the excised gland.

Spatial correlation of diseased regions on histology and MRI may be performed by (a) visually identifying and labeling corresponding structures on each modality [28, 107–109] or (b) using a semi- or fully-automated image registration procedure [2, 110–112]. For example, the spatial extent of CaP on MRI, obtained by manually labeling the *in vivo* image while visually referencing the histology, is shown in green in Fig. 5.1(d). On the other hand, the disease extent established by elastic registration of the MR and histology images in Figs. 5.1(b) and 5.1(c) is shown in Fig. 5.1(e). Note that the shape of the disease mask mapped onto the MRI in Fig. 5.1(e) more closely resembles the histopathological ground truth for CaP extent in Fig. 5.1(b) compared to the manually annotated region shown in Fig. 5.1(d). Thus, with an accurate registration technique, CaP extent on MRI can be established with greater accuracy, efficiency, and consistency compared to manual labeling.

Registration of images from different modalities such as histology and radiology




Figure 5.1: Ultrasound imagery of the prostate provides poor soft tissue resolution, while (b) high resolution MRI (*ex vivo* image shown) shows internal anatomical details of the prostate with greater clarity. Ground truth for CaP extent is obtained only through histopathologic analysis of (c) the corresponding Haematoxylin and Eosin stained tissue section. The histopathologic CaP extent (dark purple) on (c) can be mapped onto the MRI in (b) by either (d) manual labeling the MRI using histology as a visual reference, or (e) automatically mapping CaP extent from (c) via image registration. Note that the morphology of the CaP extent is better preserved in the mapping from (c) onto (e) as compared to (d).

is complicated on account of the vastly different image characteristics of the individual modalities. For example, the appearance of tissue and anatomical structures (e.g. hyperplasia, urethra, ducts) on MRI and histology are considerably different, as may be appreciated from Figs. 5.1(b) and 5.1(c). The shape of the WMH is also significantly altered due to uneven tissue fixation, gland slicing and sectioning, resulting in duct dilation, gland deformation, and tissue loss. Traditional intensity-based similarity measures, such as mutual information (MI) are typically inadequate to robustly and automatically register images from two such significantly dissimilar modalities. There have been several efforts to complement intensity information with alternative image information, including image gradients [57], co-occurrence information [59], color [61] and image segmentations [58] in conjunction with MI variants, specifically adapted to incorporate these additional channels of information. Similar to the use of calculated features to complement image intensity, it may also be advantageous to leverage additional imaging protocols that may be available. For instance, multiprotocol MR imaging is part of standard clinical practice at a number of medical centers for disease diagnosis and treatment [113,114]. These additional channels may provide complementary structural, metabolic, and functional data to complement image intensity for the registration process.

In this paper, we present an information theoretic approach, Multi-Attribute Combined Mutual Information (MACMI), to simultaneously utilize all available imaging channels, such as registered multiprotocol imagery (or image features calculated from the original images), in the registration of several images. In this work we demonstrate the application of MACMI for establishing spatial extent of CaP on radiological imaging via registration of annotated *ex vivo* histology sections with corresponding multiprotocol *in vivo* MRI.

5.2 Previous Work

Registration of prostate images is important for (a) the planning, guidance and retrospective evaluation of radiotherapy procedures [11, 52, 115, 116], (b) the fusion of diagnostic images for improved CaP detection accuracy [12, 117–119], and (c) the automated annotation of radiological images via histopathologic correlation [111, 120, 121]. Intra-modality prostate image registration has been utilized by [122] for image-guided prostate surgery via elastic registration of pre- and intraoperative prostate MRI, and by [52] for monitoring therapy related changes over time by registration of serial CT images. [12] investigated the use of multimodal prostate images (MRI, CT, and SPECT) to characterize CaP, performing automated multimodal registration of MRI to SPECT via MI, and manual registration of MRI to CT using a graphical user interface. Automated rigid registration of MRI with CT has been addressed by several groups [117–119], while automated elastic registration has only recently been addressed by [123] using a spatially constrained B-spline. Methods for registration of MRI and 3D transrectal ultrasound (TRUS) for real-time MRI-guided prostate biopsy have been presented by both [115] and [116].

The unique set of challenges associated with registration of *ex vivo* histology and multiprotocol MRI of the prostate has begun to be addressed by several recent studies; these studies have however primarily been in the context of high resolution *ex vivo* MRI [5, 112, 120, 124]. [120] utilized a thin plate spline (TPS) to model the elastic 2D deformations of histology to *ex vivo* MRI of prostate specimens using control points. However, “block face” photographs of thick tissue sections of the prostate, taken prior

to microtome slicing and slide preparation, were used to facilitate correction of the non-linear tissue deformations and in the creation of a histology volume. While these photographs allowed [120] to overcome the non-linear deformations to histology and address the issue of slice correspondences (identifying a one-to-one relationship between histology sections and slices in the MRI volume), such photographs are not generally acquired as part of routine clinical practice. Recently [112] aligned *ex vivo* MRI and histology sections of a prostate specimen with the aid of precise cancer labels on *both* modalities to improve their objective function for registration. As part of an integrated registration and segmentation strategy, the required cancer extent on MRI was established via a pixel-wise supervised classifier. The authors did not however address how classifier errors would affect registration accuracy. While high resolution (4 T) *ex vivo* MRI provided sufficient segmentation accuracy in [112], it is not clear that disease extent can reliably be established on *in vivo* clinical images. [124] also performed registration of *ex vivo* MRI and histology sections using pairs of automatically detected control points on each modality. However, automated identification of a large number of landmark pairs across *ex vivo* WMH and *in vivo* MRI (of lower image resolution and quality compared to the *ex vivo* MRI used in [124]) may not be feasible.

Registration of clinical *in vivo* radiologic images with WMH of the prostate has also been recently addressed [1, 12, 111, 121]. [12] performed histopathologic validation of CaP estimates on CT and MRI by 2D elastic registration of histology sections with CT and MRI. Manually identified control points (anatomical landmarks), placed primarily along the gland boundary, were used to define a TPS interpolant. In a rat brain study, [121] also leveraged available block face photographs of the gland prior to sectioning (similar to [120]) to generate an histology volume for 3D registration. Their approach also utilized an intermediate *ex vivo* MRI series, to which WMH was aligned via a TPS-based approach using manually selected initial control points, followed by MI-driven refinement of the coordinates of the control points. Subsequently, *ex vivo* MRI was registered to *in vivo* MRI, thus indirectly aligning the *in vivo* MRI and histology slices of the rat brains. [111] extended this approach to the human prostate and to include multiprotocol (T2-w and DWI) *in vivo* MRI and PET, again using

block face photographs and *ex vivo* MRI as an intermediate. While the works of Park and Meyer [111, 121] successfully address the need to rely on approximate slice correspondences, neither block face photographs nor *ex vivo* MRI of prostate specimens are usually available in the course of the clinical workflow. This might also explain why only two patient studies were employed in [111], and one rat brain slide in [121].

Alignment of more than two images or volumes representing very different structural or functional attributes of the same object is not well studied. One approach to registration of multiple images is to take a groupwise (GW) approach, whereby all images are simultaneously aligned, usually to some reference image or coordinate frame. The limitations of fully-GW approaches are that they either (1) involve optimization problems with many degrees of freedom arising from multiple simultaneous transformations, or (2) are limited to images with similar intensity and/or deformation characteristics. In the GW registration method presented by Bhatia [125], all images contribute to the same histogram used for entropy calculation. This limits the technique to images of the same modality. On the other hand, the GW method of Studholme [126] utilizes a high dimensional distribution suitable for multimodal data, but the use of a dense deformation field requires a constraint that penalizes deformations that deviate from an average deformation. However, in the context of large deformation fields, such as might be present between *ex vivo* and *in vivo* images, this technique is restrictive. Other methods require repeated refinement of individual transformations prior to convergence [127]. More recently, Balci [128] performed simultaneous inter-patient registration of a large number (50) of brain MRI scans using a sum of univariate (1D) entropy values (“Stack Entropy”) calculated at every pixel location. However, since this cost function requires a large number of images to calculate entropy at each pixel location, it is suited only for registration of very large populations of images from the same modality, as opposed to a smaller number of multimodal images from a single patient. Thus, while a GW approach is generally preferable to several pairwise (PW) registration steps, GW methods can be restrictive or computationally prohibitive.

Given the challenges and constraints of GW registration approaches, simple sequential PW registration steps are most commonly performed to bring multiple images from




Figure 5.2: Registration of an *ex vivo* prostate histology (\mathcal{A}) image to corresponding *in vivo* T2-w (\mathcal{B}) and T1-w (\mathcal{C}) MR images can be achieved in different ways. Two possible approaches using PW registration involve (a) PW alignment of histology to each individual MRI protocol ($\mathcal{A} \rightarrow \mathcal{B}$ and $\mathcal{A} \rightarrow \mathcal{C}$), or (b) alignment of multiprotocol MRI ($\mathcal{C} \rightarrow \mathcal{B}$) and alignment of histology to just T2-w MRI ($\mathcal{A} \rightarrow \mathcal{B}$). In the latter case, T1-w MRI would be in implicit alignment with histology at the end of the two registration steps. Alternatively, (c) a multi-attribute image registration scheme involves initial PW alignment of images from the same modality (T1-w and T2-w MRI) as in (b), followed by alignment of histology to a multi-attribute image comprising the registered multiprotocol MRI via a similarity measure defined for high dimensional data.

different modalities into alignment. This was the method of choice for the prostate work in [12] and [1] where several modalities are registered in steps using two images at a time. Figures 5.2(a) and 5.2(b) illustrate two possible approaches to PW registration of an *ex vivo* prostate histology (\mathcal{A}) image to corresponding *in vivo* T2-w MR (\mathcal{B}) and *in vivo* T1-w (single time point of a DCE series) (\mathcal{C}) MR images. Figure 5.2(a) illustrates the case where image \mathcal{A} is independently registered to both \mathcal{B} and \mathcal{C} . Figure 5.2(b) illustrates a scenario where the multiprotocol MR images are co-registered by alignment of \mathcal{C} to \mathcal{B} , and \mathcal{A} is registered to just \mathcal{B} , thus bringing \mathcal{A} into alignment with both \mathcal{B} and \mathcal{C} . For this particular set of multimodal images, the approach illustrated in Fig. 5.2(b) is preferable to that shown in Fig. 5.2(a) since alignment between multiprotocol MRI is less complicated compared to multimodal alignment of *ex vivo* histology and *in vivo* MRI. The latter approach involves dealing with highly elastic deformations and dissimilar intensities. In both instances (illustrated in Figs. 5.2(a) and (b)), the two registration steps are independent and utilize only two images at a time.

5.3 Novel Contributions and Significance

Both of the PW approaches illustrated in Figs. 5.2(a) and (b) consider only two images at a time. Hence, they exploit only a fraction of the available data in driving each registration step. Further, in subsequent alignment steps, it is necessary to select only a single image from the set of co-registered images for use as a reference. A more effective approach is to exploit all the information acquired from prior alignment steps to drive the subsequent registration operations. As illustrated in Fig. 5.2(c), following registration of \mathcal{C} to \mathcal{B} , both the newly aligned images could be considered in unison (as a “multi-attribute” image) to drive the registration with \mathcal{A} . Such an approach would exploit the fact that \mathcal{C} and \mathcal{B} are (a) in implicit alignment, and (b) represent different and informative image attributes (in this case structural and functional). This approach is akin to previous studies that have used additional textural and gradient feature images [57] to complement image intensity in order to improve image registration.

Multivariate formulations of MI have been shown to be useful in incorporating multiple image attributes (e.g. texture) [1, 59]. Thus, multivariate MI may also be

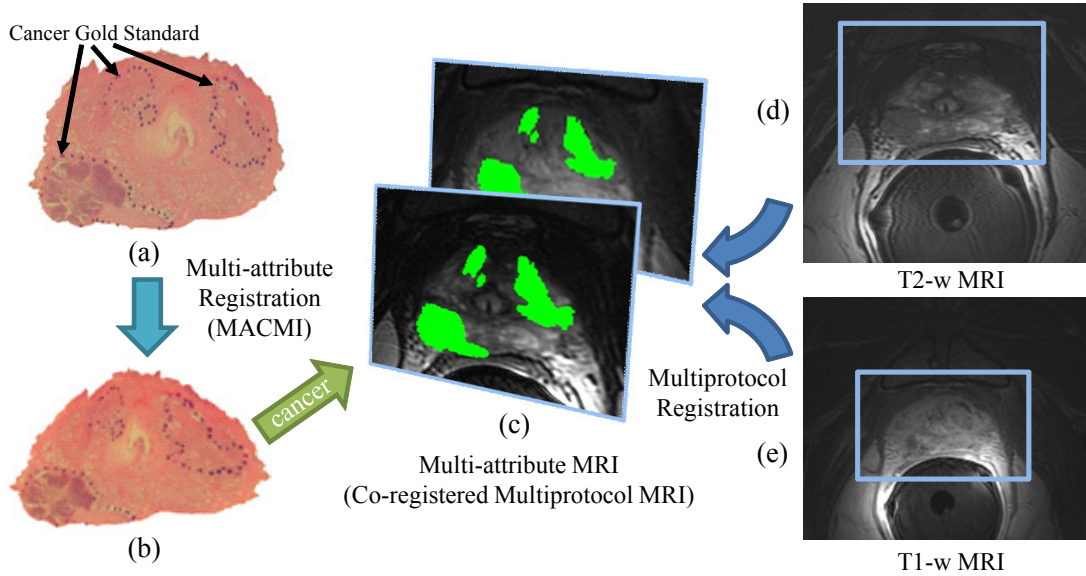


Figure 5.3: Establishing disease signatures on *in vivo* multiprotocol MRI using MACMI for registration of (a) WMH sections to corresponding (d) T2-w and (e) T1-w MRI. Alignment of the T2-w and T1-w MRI allows generation of (c) the multi-attribute MRI comprised of co-registered multiprotocol MRI. MACMI is used to align (a) WMH to (c) the multi-attribute MRI. CaP extent on the (b) elastically registered WMH is mapped directly onto both MR images in (c).

applied in the context of applications where multiprotocol imaging (e.g. T2-w, T1-w MRI) needs to be registered to another modality (e.g. histology).

The novel contribution of this work is a formal quantitative image registration framework, which we refer to as Multi-Attribute Combined MI (MACMI). MACMI allows for incorporation of multiple modalities, protocols or even feature images, in an automated registration scheme, facilitated by the use of multivariate MI. MACMI is distinct from previous GW approaches in that it handles images that can significantly vary in terms of image intensities (e.g. multimodal data) and deformation characteristics (e.g. *in vivo* to *ex vivo*). Additionally, it involves a simple (low degree of freedom) optimization procedure whereby individual image transformations are determined in sequence. The use of an information theoretic similarity measure is central to the ability of MACMI to handle (1) multimodal data, which may contain non-linear relationships between the image intensities of different modalities, and (2) high dimensional (multi-attribute)

observations, which may contain redundancies between the attributes that can be discounted via joint entropy. Finally, by employing a sequential approach to alignment of multiple images within the multi-attribute representation, each successive optimization procedure remains as simple (in terms of degrees of freedom) as with the conventional PW registration.

In this work, we evaluate MACMI in the context of a clinical problem involving multiprotocol (T2-w and DCE) MRI of the prostate and WMH sections; the sections being digitized following gland resection. Registration of WMH to corresponding multiprotocol MRI is then performed to map CaP extent from *ex vivo* WMH (previously delineated by pathologists) onto *in vivo* MRI. This procedure involves (1) initial affine alignment of the T2-w and DCE (T1-w) images (Figs. 5.3(d) and 5.3(e)) using MI to generate a multi-attribute MR image (Fig. 5.3(c)), followed by (2) multimodal elastic registration of WMH with the multi-attribute MRI.

Our scheme for registration of *in vivo* MRI and *ex vivo* WMH of the prostate is distinct from previous related efforts [12,111,121] in that (1) information from all *in vivo* imaging protocols is being utilized simultaneously to drive the process of automated elastic registration with histology, (2) no additional, intermediate *ex vivo* radiology or gross histology images need be obtained in addition to the clinically acquired *in vivo* MRI series, and (3) no point correspondences are required to be identified manually or automatically.

For the registration of 150 corresponding sets of prostates images from 25 patient studies with T2-w and DCE MRI, we quantitatively compare MACMI to a PW registration approach using conventional MI (Fig. 5.2(b)). For 15 patients for which apparent diffusion coefficient (ADC) MRI was also acquired, we further demonstrate MACMI for including the third MR protocol in the elastic registration of histology with all three MRI series for 85 sets of images (*in vivo* T2-w, DCE and ADC MRI slices, and *ex vivo* WMH sections). We also quantitatively evaluate MACMI on a synthetic brain MRI study from BrainWeb¹, whereby T1-w and T2-w MR images are registered

¹<http://www.bic.mni.mcgill.ca/brainweb/>

to PD MRI, where both the T1-w and T2-w protocols are simultaneously considered. MACMI is compared to PW MI-based approaches where T1-w and T2-w MR images are individually registered with PD MRI.

5.4 Methods

5.4.1 Theory on Mutual Information and MACMI

Mutual Information Between Scalar-Valued Images

Equation 5.1 below is a common formulation of MI for a pair of images (or random variables) $\mathcal{A}_1, \mathcal{A}_2$ in terms of Shannon entropy.

$$I_2(\mathcal{A}_1, \mathcal{A}_2) = S(\mathcal{A}_1) + S(\mathcal{A}_2) - S(\mathcal{A}_1, \mathcal{A}_2), \quad (5.1)$$

where $I_2(\mathcal{A}_1, \mathcal{A}_2)$ describes the interdependence of 2 variables, or intensity values of a pair of images [57]. As $I_2(\mathcal{A}_1, \mathcal{A}_2)$ increases, the uncertainty about \mathcal{A}_1 given \mathcal{A}_2 decreases. Thus, it is assumed that the global MI maximum will occur at the point of precise alignment, when maximal uncertainty about intensities of \mathcal{A}_1 can be explained by \mathcal{A}_2 .

Mutual Information Between High-dimensional (Multi-Attribute) Images

The conventional MI formulation can be extended to high dimensional observations by combining the multiple dimensions or attributes via high order joint entropy calculations. We refer to this application of MI as multi-attribute combined MI (MACMI) to distinguish it from conventional applications of MI and higher order MI, and denote it as I_2^* . Unlike the more familiar higher order MI ($I_n, n \geq 2$), the goal of MACMI is not to measure only the intersecting information between multiple sources ($\mathcal{A}_1, \dots, \mathcal{A}_n$), but to quantify the combined predictive value of one multivariate source (e.g. $[\mathcal{A}_1, \dots, \mathcal{A}_n]$) with respect to another (e.g. $[\mathcal{B}_1, \dots, \mathcal{B}_n]$). Here we introduce the notion of an image ensemble as the concatenation of n intensity-valued images ($\mathcal{I}_1, \dots, \mathcal{I}_n$) into an n -dimensional (multi-attribute) image, denoted as $[\mathcal{I}_1, \dots, \mathcal{I}_n]$. In the simplest case,

the MI (I_2^*) that a single image \mathcal{A}_1 shares with an ensemble of two other images, \mathcal{B}_1 and \mathcal{B}_2 , is,

$$I_2^*(\mathcal{A}_1, [\mathcal{B}_1, \mathcal{B}_2]) = S(\mathcal{A}_1) + S(\mathcal{B}_1, \mathcal{B}_2) - S(\mathcal{A}_1, \mathcal{B}_1, \mathcal{B}_2). \quad (5.2)$$

By considering \mathcal{B}_1 and \mathcal{B}_2 as simultaneously measured semi-independent variables in the multidimensional ensemble $[\mathcal{B}_1, \mathcal{B}_2]$, any dependence that exists between \mathcal{B}_1 and \mathcal{B}_2 may be discounted and I_2^* remains bounded by the smaller of $S(\mathcal{A}_1)$ and $S(\mathcal{B}_1, \mathcal{B}_2)$. The generalized form of MI between the n dimensional ensemble $\epsilon_n^{\mathcal{A}} = [\mathcal{A}_1, \dots, \mathcal{A}_n]$ with the m dimensional ensemble $\epsilon_m^{\mathcal{B}} = [\mathcal{B}_1, \dots, \mathcal{B}_m]$ is,

$$I_2^*(\epsilon_n^{\mathcal{A}}, \epsilon_m^{\mathcal{B}}) = S(\epsilon_n^{\mathcal{A}}) + S(\epsilon_m^{\mathcal{B}}) - S(\epsilon_n^{\mathcal{A}}, \epsilon_m^{\mathcal{B}}). \quad (5.3)$$

Thus, MACMI accomplishes fusion of the multiple dimensions of a multi-attribute image, allowing only intersecting information between two such images (e.g. $\epsilon_n^{\mathcal{A}}$ and $\epsilon_m^{\mathcal{B}}$) to be quantified. Calculation of $I_2^*(\epsilon_n^{\mathcal{A}}, \epsilon_m^{\mathcal{B}})$ is discussed in Section 5.4.3.

5.4.2 Framework for Registration of Multiple Images using MACMI

In Section 5.4.2, we present a generalized algorithm (*MACMIreg*) for performing registration of m images $\mathcal{Z}_1, \dots, \mathcal{Z}_m$ in a specific order. The order is specified using a hierarchical organization of the images within a family of sets \mathbf{Z} , and by progressively aligning and accumulating the registered images into a single ensemble ϵ . In Section 5.4.2 we illustrate the operation of the algorithm for 4 images ($\mathcal{Z}_1, \mathcal{Z}_2, \mathcal{Z}_3, \mathcal{Z}_4$), where \mathbf{Z} is structured to register \mathcal{Z}_1 with \mathcal{Z}_2 , and \mathcal{Z}_3 with \mathcal{Z}_4 , prior to alignment of the two resulting ensembles.

Algorithm

Consider a family of sets \mathbf{Z} that contains $m \geq 2$ images $\mathcal{Z}_1, \dots, \mathcal{Z}_m$ distributed throughout $n \leq m$ ordered subsets $\mathbf{Z}_j, j \in I$, where $I = \{1, \dots, n\}$, (i.e. $\bigcup_{j \in I} \mathbf{Z}_j = \{\mathcal{Z}_1, \dots, \mathcal{Z}_m\}$ and $\bigcap_{j \in I} \mathbf{Z}_j = \emptyset$). Each subset $\mathbf{Z}_j, j \in \{1, \dots, n\}$ may also be a family (i.e. have subsets of its own) or simply an ordered set of registered images. For example, if

$\mathbf{Z}_j = \{\{\mathcal{Z}_1^{(j)}, \mathcal{Z}_2^{(j)}\}, \{\mathcal{Z}_3^{(j)}\}, \{\mathcal{Z}_4^{(j)}\}\}$, we define \mathbf{Z}_j as a family of $|\mathbf{Z}_j| = 3$ subsets, containing a total of $k = \|\mathbf{Z}_j\| = 4$ images. We further denote the ensemble of all k images in \mathbf{Z}_j as $\varepsilon = \langle \mathbf{Z}_j \rangle = [\mathcal{Z}_1^{(j)}, \dots, \mathcal{Z}_k^{(j)}]$. By organizing the m images into a hierarchy of subsets within the family \mathbf{Z} , the order in which the images are registered and combined into multi-attribute images is determined. The procedure for alignment of all images (within and between each \mathbf{Z}_j) into a single ensemble ε of registered images is described in the following recursive algorithm *MACMIreg*. Here we define the notation $\varepsilon \leftarrow \varepsilon_d$ as the expansion of an n -dimensional multi-attribute image (ensemble) ε into an $(d + n)$ -dimensional ensemble by concatenation with a d -dimensional ensemble ε_d . We also denote $\mathcal{Z}_j \leftarrow \varepsilon$ as the assignment of each the m dimensions (intensity images) in ε to the existing m total members of \mathcal{Z}_j (independent of the organization of images within the family structure), thus replacing or updating \mathcal{Z}_j with the contents of ε .

Algorithm *MACMIreg*

Input: $\mathbf{Z} = \{\mathbf{Z}_1, \dots, \mathbf{Z}_n\}$, $n \geq 1$.

Output: ε .

Auxiliary Data Structures: Index k, j, α ; Image ensemble ε_0 .

begin

0. *for* $j = 1$ to n *do*
1. $k = |\mathbf{Z}_j|$;
2. *if* $k > 1$ *then*
3. Obtain ensemble $\varepsilon_0 = \text{MACMIreg}(\mathbf{Z}_j)$;
4. Update $\mathbf{Z}_j \leftarrow \varepsilon_0$;
5. *endif*;
6. *endfor*;
7. Initialize ε as an empty ensemble;
8. $\varepsilon \leftarrow [\mathcal{Z}_1, \dots, \mathcal{Z}_k]$, $k = \|\mathbf{Z}_1\|$;
9. $\alpha = k + 1$;
10. *for* $j = 2$ to n *do*
11. $k = \|\mathbf{Z}_j\|$;
12. $\varepsilon_0 = [\mathcal{Z}_\alpha, \dots, \mathcal{Z}_{\alpha+k}]$;
13. Obtain $\mathbf{T} = \text{argmax}_{\mathbf{T}} [I_2^*(\varepsilon, \mathbf{T}(\varepsilon_0))]$;
14. Obtain $\tilde{\varepsilon}_0 = \mathbf{T}(\varepsilon_0) = [\tilde{\mathcal{Z}}_\alpha, \dots, \tilde{\mathcal{Z}}_{\alpha+k}]$;
15. $\varepsilon \leftarrow \tilde{\varepsilon}_0$;
16. $\alpha = \alpha + k + 1$;
17. *endfor*;

end

Lines 1-6 of *MACMIreg* use recursive calls to *MACMIreg* to register the images

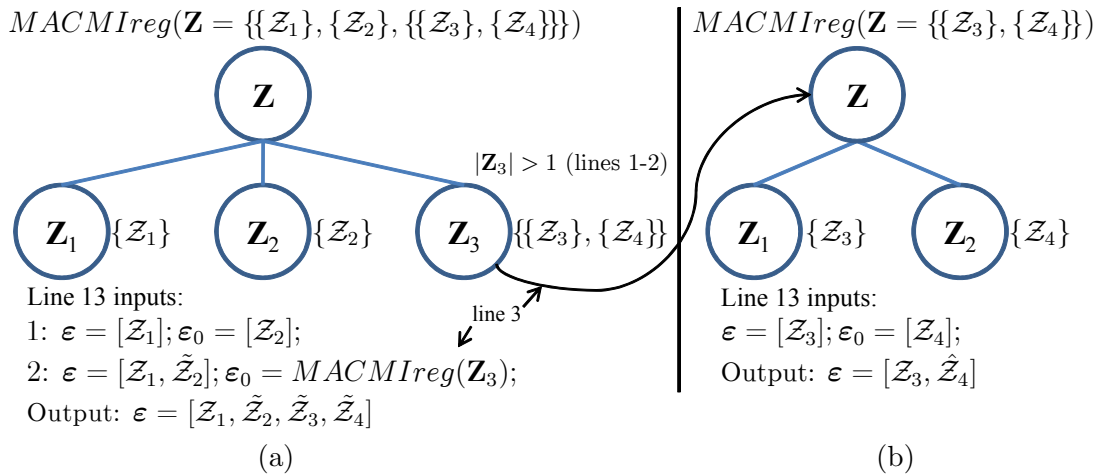


Figure 5.4: (a) Graphical representation of the organization of 4 images ($\mathbf{Z}_1, \dots, \mathbf{Z}_4$) within a family of image sets ($\mathbf{Z} = \{\{\mathbf{Z}_1\}, \{\mathbf{Z}_2\}, \{\{\mathbf{Z}_3\}, \{\mathbf{Z}_4\}\}\}$), and the application of the *MACMIreg* algorithm for alignment of all 4 images. Since only \mathbf{Z}_3 contains subsets (i.e. $|\mathbf{Z}_3| = 2$), line 3 of *MACMIreg* in (a) begins a new instance of the algorithm in (b) with $\mathbf{Z} = \{\{\mathbf{Z}_3\}, \{\mathbf{Z}_4\}\}$ as the input. The instance in (b) brings \mathbf{Z}_3 and \mathbf{Z}_4 into alignment, and returns the ensemble $[\mathbf{Z}_3, \hat{\mathbf{Z}}_4]$ to the instance in (a). The instance in (a) first bring \mathbf{Z}_1 and \mathbf{Z}_2 into alignment, and then align the ensemble of registered images from \mathbf{Z}_3 ($\epsilon_0 = [\mathbf{Z}_3, \hat{\mathbf{Z}}_4]$) with the registered images of \mathbf{Z}_1 and \mathbf{Z}_2 ($\epsilon = [\mathbf{Z}_1, \tilde{\mathbf{Z}}_2]$). At each registration step (line 13), a transformation \mathbf{T} is determined by $\arg\max_{\mathbf{T}}[I_2^*(\epsilon, \mathbf{T}(\epsilon_0))]$, and ϵ is then expanded by $\epsilon \leftarrow \tilde{\epsilon}_0 = \mathbf{T}(\epsilon_0)$ (lines 14 and 15 of *MACMIreg*). The output of (a), containing all of the co-registered images in \mathbf{Z} , is $\epsilon = [\mathbf{Z}_1, \tilde{\mathbf{Z}}_2, \tilde{\mathbf{Z}}_3, \tilde{\mathbf{Z}}_4]$.

within each \mathbf{Z}_j containing more than 1 image. When *MACMIreg*(\mathbf{Z}_j) is executed on line 3, the algorithm is recursively instantiated in order to co-register the images within the subset \mathbf{Z}_j and any of its subsets, returning the registered images within ensemble ϵ . Line 4 then updates each \mathbf{Z}_j by replacing its constituent elements with the co-registered member images contained within ϵ . Lines 7-17 of *MACMIreg* perform the registration between the multi-attribute images generated from each \mathbf{Z}_j , each of which now comprise only co-registered images (or a single image) following lines 1-6 of the algorithm. A spatial transformation \mathbf{T} of the current moving image ensemble ϵ_0 into alignment with the stationary, growing ensemble ϵ is determined on line 13. The registered ensemble $\tilde{\epsilon}_0$, obtained via \mathbf{T} on line 14, is then combined with ϵ on line 15. The algorithm continues to align each subsequent \mathbf{Z}_j with the expanding reference ensemble ϵ .

Instance of MACMI for Registration of 4 Images

The operation of the MACMI algorithm is illustrated in Fig. 5.4 for a scenario involving registration of four images $(\mathcal{Z}_1, \mathcal{Z}_2, \mathcal{Z}_3, \mathcal{Z}_4)$, of which \mathcal{Z}_3 and \mathcal{Z}_4 are designated to be co-registered prior to alignment with \mathcal{Z}_1 and \mathcal{Z}_2 . In this generic example, \mathcal{Z}_1 and \mathcal{Z}_2 could represent images from two different modalities such as CT and PET, and \mathcal{Z}_3 and \mathcal{Z}_4 could represent multiprotocol images from the same modality such as T1-w and proton density (PD) MRI that may be in proximal alignment through hardware configuration or through prior application of a registration routine.

The operation of $MACMIreg(\mathbf{Z})$ for $\mathbf{Z} = \{\{\mathcal{Z}_1\}, \{\mathcal{Z}_2\}, \{\{\mathcal{Z}_3\}, \{\mathcal{Z}_4\}\}\}$ begins by registration of images within each \mathbf{Z}_j , $j \in \{1, 2, 3\}$, where only $\mathbf{Z}_3 = \{\{\mathcal{Z}_3\}, \{\mathcal{Z}_4\}\}$ contains more than one image. Thus, as in Fig. 5.4(b), $MACMIreg(\mathbf{Z}_3)$ is called to register \mathcal{Z}_3 to \mathcal{Z}_4 and update \mathbf{Z}_3 with ensemble $\boldsymbol{\varepsilon} = [\mathcal{Z}_1, \hat{\mathcal{Z}}_2]$ (lines 3-4 of $MACMIreg$). Having registered the images within each \mathbf{Z}_j (lines 1-6 of $MACMIreg$), all images in \mathbf{Z} are registered as in Fig. 5.4(a) in two steps (lines 7-17 of $MACMIreg$). At each registration step, an optimal spatial transformation \mathbf{T} of $\boldsymbol{\varepsilon}_0$ to $\boldsymbol{\varepsilon}$ is determined by $\arg\max_{\mathbf{T}}[I_2^*(\boldsymbol{\varepsilon}, \mathbf{T}(\boldsymbol{\varepsilon}_0))]$ (line 13 of $MACMIreg$), and $\boldsymbol{\varepsilon}$ is then expanded by $\boldsymbol{\varepsilon} \leftarrow \tilde{\boldsymbol{\varepsilon}}_0 = \mathbf{T}(\boldsymbol{\varepsilon}_0)$ (lines 14-15 of $MACMIreg$). Thus, \mathbf{Z}_2 is first registered to \mathbf{Z}_1 where $\boldsymbol{\varepsilon} = [\mathcal{Z}_1]$ and $\boldsymbol{\varepsilon}_0 = [\mathcal{Z}_2]$, and $\boldsymbol{\varepsilon} \leftarrow \tilde{\boldsymbol{\varepsilon}}_0 = \tilde{\mathcal{Z}}_2 := \mathbf{T}(\mathcal{Z}_2)$. Next, \mathbf{Z}_3 is registered to \mathbf{Z}_1 (and implicitly \mathbf{Z}_2) where $\boldsymbol{\varepsilon} = [\mathcal{Z}_1, \tilde{\mathcal{Z}}_2]$ and $\boldsymbol{\varepsilon}_0 = [\mathcal{Z}_3, \hat{\mathcal{Z}}_4]$ (the output of $MACMIreg(\mathbf{Z}_3)$), and $\boldsymbol{\varepsilon} \leftarrow \tilde{\boldsymbol{\varepsilon}}_0 = [\tilde{\mathcal{Z}}_3, \tilde{\mathcal{Z}}_4] := [\mathbf{T}(\mathcal{Z}_3), \mathbf{T}(\mathcal{Z}_4)]$. The final output is $\boldsymbol{\varepsilon} = [\mathcal{Z}_1, \tilde{\mathcal{Z}}_2, \tilde{\mathcal{Z}}_3, \tilde{\mathcal{Z}}_4]$, comprising all of the co-registered images in \mathbf{Z} .

The use of both \mathcal{Z}_3 and \mathcal{Z}_4 (and both \mathcal{Z}_1 and \mathcal{Z}_2) in the final registration step has the following benefits, (1) avoids potential ambiguity in choosing between \mathcal{Z}_3 and \mathcal{Z}_4 (between \mathcal{Z}_1 and \mathcal{Z}_2), and (2) potentially provides improved alignment versus use of just \mathcal{Z}_3 or \mathcal{Z}_4 (\mathcal{Z}_1 or \mathcal{Z}_2) individually. The advantage of MACMI is that it yields cumulative incorporation of all images, while allowing flexibility to choose the order of multi-attribute image construction. Implementation of MACMI within a complete registration framework is described in Section 5.4.3 below.

5.4.3 Requirements for Implementation of MACMI

MACMI can be utilized to leverage multiple image sources in nearly any registration application by selecting the following components based on domain requirements:

1. MI estimation for high dimensional data: The most straightforward approach to estimating $I_2^*(\epsilon_n^A, \epsilon_m^B)$ is to formulate the density estimates from high dimensional histograms. While histogram-based techniques are feasible and effective for up to 4-dimensional observations with appropriate bin size, as demonstrated in [58] and [59], higher dimensionality necessitates an alternate estimate of entropy or MI, such as those based on entropic spanning graphs or related quantities such as α -MI [87].
2. Image transformation model(s): Since MACMI only dictates the construction of the objective function, MACMI is agnostic to the deformation model. Further, different deformation models may be used for each image since the individual image transformations are performed in independent steps.
3. Optimization scheme to find a maximum of $I_2^*(\epsilon_n^A, \epsilon_m^B)$: If the analytical gradient can be derived, as demonstrated for α -MI in [129], an efficient stochastic gradient descent method can be used. In the absence of analytical gradients of $I_2^*(\epsilon_n^A, \epsilon_m^B)$, methods including direct search (e.g. downhill simplex), quasi-Newton (e.g. Broyden-Fletcher-Goldfarb-Shanno (BFGS)), and other finite difference-based schemes can be employed.

Specific implementation details employed in this work are described in the following section.

5.5 Experimental Design

A summary of the clinical prostate and synthetic brain image data sets investigated in this study is presented in Table 5.1.

Table 5.1: Summary of the synthetic and clinical datasets registered by MACMI. Pixels are square.

Data Set	Description	Modalities
S^s	Synthetic multiprotocol brain MRI from BrainWeb	T2-w, T1-w, PD MRI
S_1^c	Clinical multiprotocol MRI and histology of prostate	<i>in vivo</i> T2-w, DCE MRI, <i>ex vivo</i> WMH
S_2^c	Clinical multiprotocol MRI and histology of prostate	ADC, T2-w, DCE MRI, <i>ex vivo</i> WMH
Data Set	Dimensions	Studies (Images)
S^s	181×217 (1 mm pixel)	1 (20)
S_1^c	T2-w MRI: 512×512 (0.230-0.280 mm pixel)	25 (150)
S_2^c	T2-w MRI: 512×512 (0.230-0.280 mm pixel)	15 (85)

5.5.1 Synthetic Multiprotocol Brain MRI

Data Description

To quantitatively evaluate the performance of MACMI, we consider a synthetic registration task using a data set S^s comprising 20 2D multiprotocol (T1-w, T2-w, and PD) MRI slices from the BrainWeb simulated brain database [82]. We denote the T1-w, T2-w, and PD MRI slices as $\mathcal{T}1$, $\mathcal{T}2$ and \mathcal{P} , respectively.

Registration Experiment

Since synthetic brain MRI volumes $\mathcal{T}1$, $\mathcal{T}2$ and \mathcal{P} are initially in alignment, we apply a known non-linear deformation (\mathbf{T}^{ap}) to \mathcal{P} to generate an image \mathcal{P}^d , for which misalignment from $\mathcal{T}1$ and $\mathcal{T}2$ is known. The objective of the registration task is to recover the initial correct alignment via a corrective deformation (\mathbf{T}^{co}). We denote the recovered \mathcal{P} slice as \mathcal{P}^r . Image transformation is implemented using an elastic Free Form Deformation (FFD) model with a hierarchical mesh grid spacing scheme, as described in [59]. Three mesh grid levels were defined for vertex spacings of approximately $\{36, 27, 18\}$ mm at each level ($n_{x,y} \in \{(6, 5), (8, 7), (12, 10)\}$ moving control points on zero-padded images). MACMI is performed in a manner similar to the scenario illustrated in Fig.

5.2(c), whereby \mathcal{P}^d (instead of WMH) is registered to the multi-attribute image comprising the co-registered sections $\mathcal{T}1$ and $\mathcal{T}2$ via the recovered transformation,

$$\mathbf{T}_{MACMI}^{co} = \operatorname{argmax}_{\mathbf{T}} \left[I_2^*(\varepsilon(\mathcal{T}1\mathcal{T}2), \mathbf{T}(\mathcal{P}^d)) \right], \quad (5.4)$$

where $\varepsilon(\mathcal{T}1\mathcal{T}2)$ represents the ensemble ε of $\mathcal{T}1$ and $\mathcal{T}2$. In order to compare with MACMI, registration is also performed using objective functions defined by the MI of (1) \mathcal{P}^d with $\mathcal{T}1$, and (2) \mathcal{P}^d with $\mathcal{T}2$. Thus, two additional \mathcal{P}^r images are obtained by $\mathbf{T}_{PW1}^{co} = \operatorname{argmax}_{\mathbf{T}} [I_2(\mathcal{T}1, \mathbf{T}(\mathcal{P}^d))]$, and by $\mathbf{T}_{PW2}^{co} = \operatorname{argmax}_{\mathbf{T}} [I_2(\mathcal{T}2, \mathbf{T}(\mathcal{P}^d))]$. Estimation of I_2 and I_2^* was achieved using 2D and 3D probability density estimates obtained using histograms with 128 and 62 graylevel bins, respectively. The number of bins in each case were chosen empirically for reliable optimization of Equation 5.4 using a Nelder-Mead simplex algorithm [79].

Registration Evaluation

For the synthetic data, quantitative evaluation of registration accuracy can be performed easily since the correct coordinate transformation, \mathbf{T}^{ap} , is known. The magnitude of error in the transformation \mathbf{T}^{co} determined by registration can be quantified in terms of mean absolute difference (MAD) ($F_{mad}(\mathbf{T}^{co})$) and root mean squared (RMS) error ($F_{rms}(\mathbf{T}^{co})$) from \mathbf{T}^{ap} . Both MAD and RMS error are computed over the N total image pixels c in the common coordinate frame C of $\mathcal{T}1$, $\mathcal{T}2$, and \mathcal{P} , and can be expressed as in Equation (4.9). Further, the original \mathcal{P} is compared directly with the resulting \mathcal{P}^r using L^2 distance (D_{L2}) as the similarity measure.

5.5.2 Clinical Multimodal Prostate MRI and Histology

Data Description

We address the registration of two prostate data sets, S_1^c and S_2^c , comprising multimodal (3 T *in vivo* MRI and histology) and multiprotocol (T2-w, DCE, and ADC MRI) images (see [113] for details on acquisition). Set S_1^c comprises 150 corresponding *ex vivo* WMH sections with CaP delineated and their closest corresponding 3 T *in vivo* T2-w and

DCE MRI slices over all 25 patient studies. Set S_2^c is a subset of 15 patients (85 image sets) in S_1^c for which DWI was also acquired and ADC maps calculated. We denote the T2-w MRI, DCE MRI, ADC MRI, and WMH images as \mathcal{S} , \mathcal{F} , \mathcal{D} and \mathcal{H} respectively. Each DCE MRI series comprises a sequence of T1-w gradient echo MRI volumes acquired following bolus injection of the gadopentetate dimeglumine contrast agent at 0.1 mmol/kg of body weight. Two pre- and five post-contrast T1-w MRI volumes were obtained at a temporal resolution of 95 seconds. Maximal enhancement was generally obtained at the third post-contrast time point (denoted as \mathcal{F}^3), which was designated for use in the registration routines. No *ex vivo* MRI or gross pathology photographs were acquired.

Following RP and prior to sectioning, the excised prostate was embedded in a paraffin block while maintaining the orientation to keep the urethra perpendicular to the plane of slicing. This procedure facilitates the identification of a corresponding *in vivo* 2D axial MRI slice for each 2D histology slice. Preparation of the digitized WMH sections proceeds as follows: (1) the excised prostate is cut into sections that are 3-4 mm thick by slicing axial sections from the paraffin block using a circular blade, (2) a microtome is used to further cut the sections into thin slices that are about 5 μm thick, and (3) a single thin slice from each 3-4 mm thick section is chosen and digitally scanned. Each slide is then examined under a light microscope using up to 40x apparent magnification to identify and delineate the regions of CaP. As a result of this slide preparation process, spacing of the digital WMH slides is both coarse and irregular, ranging from 3-8 mm. Further, the non-linear tissue deformations and artifacts introduced during microtome slicing are independent between slices, and no block face photographs or fiducial markers are available to correct the distortions. Therefore it is not feasible to accurately construct a 3D histology volume suitable for a 3D registration procedure without significant modifications to the established clinical routine [6, 130]. Thus for each WMH slice with disease, the closest corresponding T2-w MRI slice was visually identified by an expert radiologist.

Registration Experiments

As previously described, the goal of this task is to register WMH to each MRI protocol in order to map CaP extent onto MRI. Since each MRI series was acquired in sequence and with minimal movement, 3D affine registration of the DCE to T2-w MRI volume was performed for each of the 25 patients in S_1^c (and S_2^c) via MI between the 3D T2-w MRI volume and the 3D T1-w MRI volume corresponding to the third post-contrast time point of the 4D DCE MRI volume, using 256 graylevel bins for the joint histogram. For each axial slice $\tilde{\mathcal{F}}$ of the registered 4D DCE MRI volume, a multi-attribute image representation $\epsilon(\mathcal{S}\tilde{\mathcal{F}}^3)$ was generated, as illustrated in Figs. 5.3(c)-(e). Prior to 2D elastic registration of each \mathcal{H} to each $\epsilon(\mathcal{S}\tilde{\mathcal{F}}^3)$, the prostate capsule on the corresponding \mathcal{S} images (6 slices on average) for each patient was roughly delineated and the extraneous tissue masked out during registration. This was done to facilitate a rough global localization of the prostate on \mathcal{S} relative to \mathcal{H} .

For data set S_1^c , automatic FFD registration of WMH to the MRI by MACMI is then performed for each of the 150 \mathcal{H} slices using $I_2^*(\epsilon(\mathcal{S}\tilde{\mathcal{F}}^3), \mathbf{T}^e(\mathcal{H}))$, resulting in a warped WMH $\tilde{\mathcal{H}} = \mathbf{T}^e(\mathcal{H})$, as shown in Fig. 5.3(b). The CaP extent is then mapped onto \mathcal{S} and $\tilde{\mathcal{F}}^3$ by \mathbf{T}^e , as shown in green on Fig. 5.3(c). In addition to MACMI, the elastic registration of \mathcal{H} to the coordinate frame of \mathcal{S} is also performed using (1) the conventional MI of \mathcal{H} with \mathcal{S} ($I_2(\mathcal{S}, \mathbf{T}^e(\mathcal{H}))$), and (2) the conventional MI of \mathcal{H} with $\tilde{\mathcal{F}}^3$ ($I_2(\tilde{\mathcal{F}}^3, \mathbf{T}^e(\mathcal{H}))$). We refer to these PW registration approaches as PW-T2 and PW-DCE, respectively, and denote corresponding transformations as \mathbf{T}_{MACMI}^e , \mathbf{T}_{PW-T2}^e , and \mathbf{T}_{PW-DCE}^e .

For data set S_2^c , the 3D ADC volume for each patient is also registered using a 3D affine transformation to the co-registered T2-w and DCE MRI volumes via a multi-attribute volume (i.e. the volume composed of axial images $\epsilon(\mathcal{S}\tilde{\mathcal{F}}^3)$), hence generating a registered ADC volume comprising slices $\tilde{\mathcal{D}}$. Automated FFD registration of WMH to MRI is then performed for each set of corresponding \mathcal{H} and \mathcal{S} slices by MACMI, which also considers $\tilde{\mathcal{F}}^3$ and $\tilde{\mathcal{D}}$ via $I_2^*(\epsilon(\mathcal{S}\tilde{\mathcal{F}}^3\tilde{\mathcal{D}}), \mathbf{T}^e(\mathcal{H}))$. CaP extent is then mapped to each of \mathcal{S} , $\tilde{\mathcal{F}}^3$ and $\tilde{\mathcal{D}}$. PW registration using just \mathcal{H} and \mathcal{D} is not performed as the DWI

protocol provides insufficient spatial resolution and anatomical detail for multimodal correlation based on image intensities alone.

Registration Evaluation

Since ground truth for alignment of the clinical prostate data is not known or easily determinable, evaluation is performed by calculating similarities of both \mathcal{S} and $\tilde{\mathcal{F}}^3$ with $\mathbf{T}_{MACMI}^e(\mathcal{H})$, $\mathbf{T}_{PW-T2}^e(\mathcal{H})$, and $\mathbf{T}_{PW-DCE}^e(\mathcal{H})$. To exclude the extraneous tissue outside the prostate in registration evaluation, similarity is calculated in terms of the MI (Equation 5.1) for only the area in the image containing the prostate. Qualitative evaluation is also performed by visually comparing the CaP extent mapped from histology onto T2-w MRI via \mathbf{T}_{MACMI}^e , \mathbf{T}_{PW-T2}^e , and \mathbf{T}_{PW-DCE}^e . An estimate of CaP extent manually established by a radiologist on select slices of the T2-w MRI is used as the ground truth. This is done for only those specific MR images where CaP extent could be reliably delineated. It is important to note these slices are not representative of all T2-w MR images with CaP, since easily delineated lesions are generally associated with dense tumors (i.e. those with compact cellular arrangements) [131] and those occurring only in the peripheral zone [132]. Further, delineation of the CaP boundary is more challenging than discerning the presence of a lesion.

5.6 Results and Discussion

5.6.1 Synthetic Brain Registration

Table 5.2 presents a comparison of the evaluation measures F_{mad} , F_{rms} , and D_{L2} for transformations obtained in elastic registration of the $n = 20$ multiprotocol MRI slices using MACMI (\mathbf{T}_{MACMI}^{co}) and both PW registration approaches (\mathbf{T}_{PW1}^{co} , \mathbf{T}_{PW2}^{co}). The values of F_{mad} were compared between \mathbf{T}_{MACMI}^{co} and \mathbf{T}_{PW1}^{co} using a paired t -test under the null hypothesis that there was no difference in F_{mad} between \mathbf{T}_{MACMI}^{co} and \mathbf{T}_{PW1}^{co} . The values of F_{rms} and D_{L2} were also compared between \mathbf{T}_{MACMI}^{co} and \mathbf{T}_{PW1}^{co} . Similarly, the values of F_{mad} , F_{rms} , and D_{L2} were also compared between \mathbf{T}_{MACMI}^{co} and \mathbf{T}_{PW2}^{co} (second row of table). MACMI achieves better performance in terms of each

Table 5.2: Comparison of elastic registration accuracy for MACMI and pairwise MI alignment of $n = 20$ pairs of synthetic PD MRI with co-registered T1-w and T2-w MRI brain images. The measures illustrated below correspond to (i) error of recovered deformation field (in mm) in terms of F_{mad} and F_{rms} , and (ii) distance (D_{L2}) between the undeformed and recovered PD MRI. MACMI results are significantly more accurate compared to either PW approach (p -values for both tests shown).

	F_{mad}	F_{rms}	D_{L2}
\mathbf{T}_{PW1}^{co} (T1-PD)	0.9117	2.1407	1.83e+03
\mathbf{T}_{PW2}^{co} (T2-PD)	0.9506	2.0248	2.35e+03
\mathbf{T}_{MACMI}^{co} (ϵ (T1T2)-PD)	0.8348	1.9307	1.71e+03
p (\mathbf{T}_{PW1}^{co} vs. \mathbf{T}_{MACMI}^{co})	0.0817	0.0578	0.0174
p (\mathbf{T}_{PW2}^{co} vs. \mathbf{T}_{MACMI}^{co})	0.0013	0.2020	1.8e-10

measure, with significantly lower error ($p < 0.05$ for $n = 20$) compared to one or both PW methods.

5.6.2 Clinical Prostate Registration

Mapping CaP extent from WMH onto *in vivo* T2-w and T1-w MRI

Figure 5.5 illustrates the registered WMH images and the corresponding T2-w MRI along with the contours of the mapped CaP extent and the urethra (verumontanum), for a single set of corresponding images in S_1^c (and S_2^c). The \mathcal{S} slice is shown in Fig. 5.5(a) with the region containing the dominant intraprostatic lesion (DIL) shown in the blue box and the urethra outlined in orange. An expert delineation of the DIL (green contour) is shown in Fig. 5.5(b), while the original WMH image with CaP ground truth (dotted line) is shown in Fig. 5.5(c). As described in Section 5.5.2, affine registration of T2-w and DCE MRI protocol volumes (containing images \mathcal{S} with \mathcal{F}) is performed prior to elastic registration of \mathcal{H} . Elastic registration of the \mathcal{H} slice (shown in Fig. 5.5(c)) to \mathcal{S} in Fig. 5.5(a) is then performed individually using each method (PW-T2, PW-DCE, MACMI). The elastically warped \mathcal{H} , $\tilde{\mathcal{H}}$, obtained using only \mathcal{S} (PW-T2) is shown in Fig. 5.5(f), and the CaP extent is shown mapped onto \mathcal{S} in Figs. 5.5(d) and 5.5(e). $\tilde{\mathcal{H}}$ obtained using only \mathcal{F} (PW-DCE) is shown in Fig. 5.5(j), and the CaP extent is again shown mapped onto \mathcal{S} (Figs. 5.5(h) and 5.5(i)). Finally, $\tilde{\mathcal{H}}$ obtained using both \mathcal{S} and \mathcal{F} (MACMI) is shown in Fig. 5.5(m), and the CaP extent is again shown on \mathcal{S} (Figs.

Table 5.3: Comparison of elastic registration accuracy for MACMI, PW-T2, and PW-DCE for $n = 25$ patient studies. The measures illustrated below correspond to mean similarity in terms of the total MI of all registered WMH slices $\tilde{\mathcal{H}}$, obtained by MACMI, PW-T2, or PW-DCE (columns of table), with either T2-w MRI or DCE MRI (rows).

mean MI with WMH	registration objective function			p (vs. MACMI)	
MR protocol	MACMI	PW-T2	PW-DCE	PW-T2	PW-DCE
T2-w: $\bar{I}_2(\mathcal{S}, \tilde{\mathcal{H}})$	0.3378	0.3339	0.3297	0.0099	0.0006
DCE: $\bar{I}_2(\tilde{\mathcal{F}}^3, \tilde{\mathcal{H}})$	0.3155	0.3085	0.3102	4.0e-4	0.0014

5.5(k) and 5.5(l)). The contours of the mapped CaP extent on MRI suggest that (1) accurate elastic registration of WMH directly to corresponding *in vivo* MRI is feasible using the described FFD framework and (2) MACMI outperforms PW application of MI using data from only one of the several available MRI protocols. Note again that while the DIL shown in Fig. 5.5 is useful for qualitative evaluation, identifying such a clearly bounded lesion on T2-w MRI is rare. The positions of the urethra (orange contours) on $\tilde{\mathcal{H}}$ in Figs. 5.5(f), 5.5(j), and 5.5(m) also illustrate improved alignment of the image interior via MACMI compared to PW-T2 and PW-DCE. For example, note the improved urethral positioning via MACMI in Fig. 5.5(m) compared to the left misalignment by PW-T2 in Fig. 5.5(f).

The results in Fig. 5.5 qualitatively demonstrate that the general FFD-based registration paradigm described in this paper is capable of generating good alignment between *in vivo* MRI and *ex vivo* WMH prostate sections without the use of any additional *ex vivo* MRI series or gross histology (“block face”) photographs. Table 5.3 quantitatively illustrates the advantage of using MACMI over a single MRI protocol. For each of the $n = 25$ patient studies in S_1^c , registration accuracy for each of MACMI, PW-T2, and PW-DCE was approximated by the total MI of the elastically registered histology slices $\tilde{\mathcal{H}}$ with all of the corresponding (1) \mathcal{S} and (2) $\tilde{\mathcal{F}}^3$ slices, both of which are in the same coordinate frame as \mathcal{S} . Table 5.3 lists the average MI value \bar{I}_2 over $n = 25$ patients of all registered WMH slices $\tilde{\mathcal{H}}$ obtained by MACMI, PW-T2, and PW-DCE (columns) with either \mathcal{S} or $\tilde{\mathcal{F}}^3$ (rows). The values of $\bar{I}_2(\mathcal{S}, \tilde{\mathcal{H}})$ (and $\bar{I}_2(\tilde{\mathcal{F}}^3, \tilde{\mathcal{H}})$) were compared between MACMI and PW-T2 using a paired t -test under the null hypothesis that there was no difference in $\bar{I}_2(\mathcal{S}, \tilde{\mathcal{H}})$ (and $\bar{I}_2(\tilde{\mathcal{F}}^3, \tilde{\mathcal{H}})$) between MACMI and




Figure 5.5: (a) 3T *in vivo* T2-w MRI of a prostate with a clearly visible DIL (shown in blue box) and magnified in (b) with a manual estimate of CaP extent (green). (c) Closest corresponding WMH slice with CaP ground truth (dotted line) and urethra (orange). (d)-(e) T2-w MRI with estimate of CaP extent (red) as mapped from (f) WMH via elastic registration using only T2-w MRI. (h)-(i) T2-w MRI with CaP estimate from (j) WMH registered to DCE (T1-w) MRI (co-registered to T2-w MRI). (k)-(l) Registration using both T2-w and DCE MRI via MACMI results in closer agreement of the registration-derived CaP extent (red) and the manual estimate (green). The verumontanum of the urethra (orange contour) is also shown on the registered WMH images in (f), (j), and (m).




Figure 5.6: Using MACMI to include ADC MRI in the elastic registration of (a) histology to each of (d) T2-w, (e) DCE (T1-w) and (f) ADC MRI. Prior to elastic registration of histology, (d)-(e) T2-w and T1-w MRI were first successfully aligned via MI, as seen by (c) the checkerboard overlay of T2-w MRI and registered T1-MRI. (e) ADC was then registered to both T2-w and T1-w MRI via MACMI. (b) Elastically registered histology was obtained using the co-registered multiprotocol MRI via MACMI, and CaP extent was mapped onto (d)-(f) MRI (red lines).

PW-T2. The comparisons of $\bar{I}_2(\mathcal{S}, \tilde{\mathcal{H}})$ and $\bar{I}_2(\tilde{\mathcal{F}}^3, \tilde{\mathcal{H}})$ were also made between MACMI and PW-DCE. In all comparisons, MACMI demonstrated significant ($p < 0.05$) improvement over both PW-T2 and PW-DCE, despite these PW methods using MI as their objective function.

Mapping CaP Extent from WMH onto ADC, T2-w and T1-w MRI

For the $m = 15$ patient studies in S_2^c for which ADC maps were also obtained, MACMI was applied in both (1) the 3D affine alignment of \mathcal{D} to $\varepsilon(\mathcal{S}\tilde{\mathcal{F}}^3)$ and (2) the 2D elastic alignment of \mathcal{H} to $\varepsilon(\mathcal{S}\tilde{\mathcal{F}}^3\tilde{\mathcal{D}})$. Figures 5.6(a) and 5.6(b) show the original \mathcal{H} and the warped \mathcal{H} ($\mathbf{T}^e(\mathcal{H})$), following elastic alignment with $\varepsilon(\mathcal{S}\tilde{\mathcal{F}}^3\tilde{\mathcal{D}})$. Figure 5.6(c) shows a

checker board visualization of the two co-registered \mathcal{S} and $\tilde{\mathcal{F}}^3$ slices (shown in Figs. 5.6(d)-(e)) after the 3D registration of the T2-w and DCE MRI volumes performed for the studies in S_1^c . For each of the studies in $S_2^c \subset S_1^c$, the second 3D affine multiprotocol registration step was performed to align the ADC volume to the T2-w and DCE MRI volumes. This enables the generation of aligned $\tilde{\mathcal{D}}$ images (Fig. 5.6(f)) for each \mathcal{S} (and $\tilde{\mathcal{F}}^3$). \mathcal{S} and the registered $\tilde{\mathcal{F}}^3$ and $\tilde{\mathcal{D}}$ images are shown in Figs. 5.6(d)-(f) with the contour of the mapped CaP extent from $\mathbf{T}^e(\mathcal{H})$ (Fig. 5.6(b)). It was observed that the inclusion of $\tilde{\mathcal{D}}$ in $\epsilon(\mathcal{S}\tilde{\mathcal{F}}^3\tilde{\mathcal{D}})$ had little effect on the resulting alignment with \mathcal{H} when compared with the results obtained using only \mathcal{S} and $\tilde{\mathcal{F}}^3$ as in Section 5.6.2 above (comparison not shown).

5.7 Concluding Remarks

Signatures for disease on multimodal *in vivo* imaging may be used to develop systems for computer-assisted detection of cancer, or to assist in the training of medical students, radiology residents, and fellows. To establish *in vivo* radiological imaging signatures for prostate cancer, an accurate estimate of ground truth for cancer extent on each of the imaging modalities is necessary. In the context of certain anatomic regions and diseases, the spatial extent of disease may be obtained by spatial correlation or registration of histology sections with corresponding *in vivo* images. In this paper, we presented a new method termed Multi-Attribute Combined Mutual Information (MACMI) within an automated elastic FFD registration framework for alignment of images from multiple *in vivo* acquisition protocols with corresponding *ex vivo* whole-mount histology (WMH) sections. Our approach to registration of *in vivo* multiprotocol radiology images and *ex vivo* WMH of the prostate using MACMI is distinct from previous related efforts [12, 111, 121] in that (1) information from all *in vivo* image sources is being utilized simultaneously to drive the automated elastic registration with WMH, (2) no additional, intermediate *ex vivo* radiology or gross histology images need to be obtained (this approach does not disrupt routine clinical workflow), and (3) no point correspondences are required to be identified manually or automatically. This last advantage is particularly relevant in the context of *in vivo* MR images where visual identification of

anatomical landmarks is a challenge even for experts.

MACMI performs registration of several images by incorporating multiple image sources using an information theoretic approach. These may include different modalities, acquisition protocols, or image features. For the clinical application discussed in this paper, MACMI facilitates the use of all available *in vivo* prostate imaging protocols acquired during the standard clinical routine, in order to perform automated elastic registration with the *ex vivo* WMH. Unlike fully-groupwise registration techniques, the optimization problem remains simple while accommodating both highly dissimilar modalities and large deformations of variable magnitude.

We demonstrated the use of MACMI for registration of 150 multimodal (WMH, T2-w and DCE MRI) prostate image sets from 25 patients, 85 sets of WMH, T2-w, ADC and DCE MRI from 15 patients, and 20 sets of synthetic T1-w, T2-w, and PD MR brain images. Statistically significant improvement in registration accuracy was observed in using MACMI to simultaneously register PD MRI to both T1-w and T2-w MRI, compared to pairwise registration of PD to T1-w or T2-w MRI, for the synthetic dataset. Qualitative examination of alignment between multiprotocol clinical prostate MRI and histology suggested improved performance via MACMI over pairwise MI. The inclusion of ADC MRI in the multi-attribute registration had little effect on the resulting alignment with WMH when compared with the results obtained using only T2-w and T1-w MRI (Section 5.6.2). Nevertheless, it is possible that the use of MACMI in the registration of ADC to the ensemble of registered T2-w and T1-w MRI helped achieve more consistent multiprotocol MRI alignment than if either protocol were used alone. We intend to investigate this application of MACMI further in the future.

While we utilized histograms for density estimation, other techniques, such as entropic graphs [87], can be applied for larger numbers of images. However, independent of the implemented estimation method, MACMI affects information theoretic fusion of multiple image sources by computing multivariate MI between multi-attribute images constructed as ensembles of co-registered images. It is important to note that in the absence of a predetermined order for combining images, MACMI may still be applied by combining images in a completely arbitrary order. Even in this scenario, MACMI

still represents an improvement over fully-pairwise registration by utilizing all registered images. Future work will investigate the influence of the order of multi-attribute image construction on alignment accuracy.

Chapter 6

Application of Elastic Image Registration and CAD to Improving IMRT

6.1 Introduction to Radiation Therapy Planning

Treatment of prostate cancer (CaP) by targeted radiotherapy requires the use of computerized tomography (CT) to formulate a dose plan. Successful conformal planning can reduce rectal and bladder toxicity by more accurately targeting the prostate, in turn allowing dose escalation to the planning target volume (PTV) and more effective treatment. Localization of the dominant intraprostatic lesion (DIL) can be used to create dose plans with even less exposure to benign tissue. More importantly, a more focused dose plan can afford significant dose escalation to the tumor, potentially providing greatly reduced rates of recurrence. However, CT does not provide good tumor localization [133]. Further, CT has been shown to overestimate the prostate volume and provide inaccurate discrimination between base and apex and surrounding structures [134]. As such, there is a need for improved image information in planning of guided therapy.

Magnetic resonance imaging (MRI) of the prostate has been shown to provide improved resolution of intraprostatic structures and the prostate boundary compared to CT [117] and ultrasound [22, 25]. We have previously demonstrated the utility of a computer-aided detection (CAD) system for prostate cancer on high resolution *ex vivo* MRI [28], as well as on *in vivo* multiprotocol MRI [101, 102, 135]. Pickett demonstrated the combined use of MRI and magnetic resonance spectroscopy (MRS) for tumor detection to provide escalated dose to the DIL [136], although their study was limited to looking within the peripheral zone and only involved rudimentary alignment.

In order to utilize MRI for dose planning, it is necessary to align, or register, MRI

with CT so that MRI-derived diagnoses may be mapped onto CT. Therefore, registration of MRI with CT of the prostate has been investigated by a few groups [117–119]. One approach [117] has involved using bones or implanted gold fiducials as landmarks for alignment. Such intraprostatic marker-based registration techniques have been limited by time requirements and uncertainty associated with both the implantation and identification of marker centers on images. While surface-based registration using the iterative closest point (ICP) algorithm have been shown in [118] to be superior to fiducial-based methods, manual segmentation of the prostate on both CT and MRI is required for this type of approach. In addition to manual segmentation of the prostate on both CT and MRI being time consuming, the methods in [117] and [118] were limited to rigid body transformations. Recently, a method for semiautomated registration of post-implant CT and MRI was proposed by Vidakovic [119]. The technique exploited the presence of seed and insertion-needle tracks to achieve superior-inferior alignment prior to an automated registration step via maximization of the normalized mutual information (NMI) similarity measure. Besides being limited to rigid-body transformations, these registration methods are subject to interobserver variability.

Significant non-linear deformations may occur between the acquisition of planning CT and diagnostic MRI, due to the use of an endorectal coil (needed for diagnostic MR imaging) and variability in the size of the bladder, both of which push against the prostate. As such, a robust elastic registration framework is necessary for accurate and efficient delineation of prostate and tumor boundary on planning CT images using high resolution diagnostic MR. Direct registration of diagnostic MRI with planning CT images, such as those shown in Figures 6.1(a) and (b), is complicated on account of both resolution and modality differences. Resolution differences are associated with differences in field of view (FOV), where diagnostic MRI has a small FOV centered on the prostate, and the planning CT has a wider FOV as well as reduced resolution to resolve soft tissue details [117, 133, 134]. However, a second image set (planning MRI) can be acquired in the same scanning session with a similar FOV and patient positioning to that of the CT.

In this paper, we present a complete quantitative framework for Computer-Assisted




Figure 6.1: (a) Diagnostic MRI which may allow for identification of CaP extent. The diagnostic MRI must be registered to (b) planning CT so that a targeted dose plan may be generated. (c) Planning MRI can be used as an intermediary to facilitate alignment between (a) diagnostic MRI and (b) planning CT. (d) Registered diagnostic MRI (estimated cancer label in red) in alignment with planning CT.

Targeted Therapy (CATT). The CATT planning method presented here is comprised of three main components, (1) an unsupervised texture-driven classifier for identifying suspected locations on diagnostic MRI, (2) an elastic registration method for alignment of diagnostic MRI, planning MRI and CT images of the prostate, and (3) mapping of tumor location onto CT, and generation of a dose plan that is targeted at the tumor location. The CAD system utilizes multiple textural features to enhance characterization of tissue and facilitates the discrimination of benign and malignant regions of the prostate. Figure 6.1(c) shows the planning MRI that corresponds to the diagnostic MRI (Figure 6.1(a)) and planning CT (Figure 6.1(b)). Using planning MRI, we present

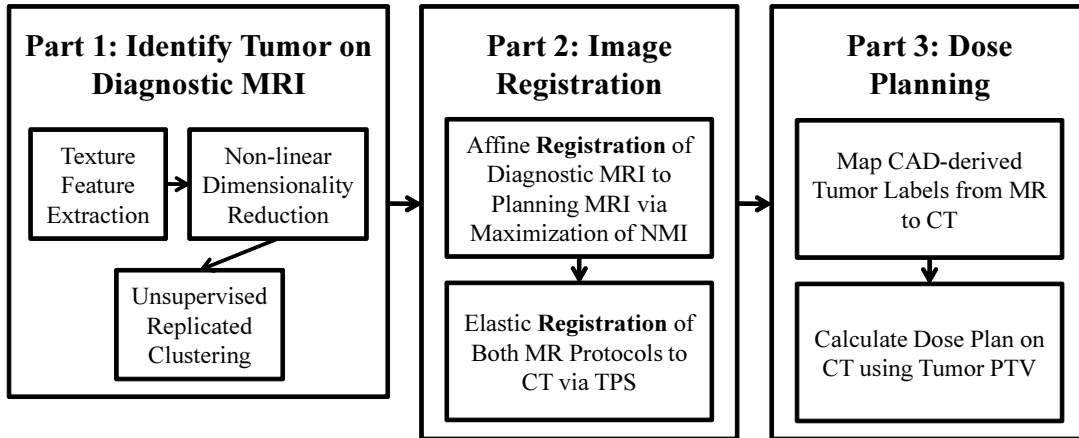


Figure 6.2: Overview of the CATT system. An unsupervised CAD system is used to identify suspect regions on the diagnostic MRI. A two step registration procedure is performed, which involves alignment of diagnostic MRI to planning MRI, followed by alignment of multiprotocol MRI to the planning CT. This allows for mapping the tumor regions onto CT, allowing for a targeted radiotherapy dose plan to be generated.

a novel two-step elastic registration procedure to overcome the modality differences between CT and MRI and bring diagnostic MRI into alignment with planning CT. In the first step, the registration of diagnostic and planning MRI is driven by exploiting the similarity of two MRI datasets to overcome FOV differences. Following coregistration of both MRI protocols, an elastic registration procedure is employed to align the planning MRI to CT, and consequently the diagnostic MRI to CT. This two step process allows for overcoming the modality differences that might otherwise complicate the multimodal registration task (i.e. mapping tumor labels directly from diagnostic MRI to planning CT). For each slice on which the prostate is visible, 2D transformations are determined to align the corresponding slices from each modality. We adopt this approach compared to a volumetric approach on account of the difficulty in establishing point correspondences in 3D, and due to the highly anisotropic voxels of diagnostic and planning MRI (see Table 6.1). The CATT system components and workflow are illustrated in Figure 6.2.

The rest of the paper is organized as follows. In Section 6.2 we describe the components of the CATT framework, including the unsupervised CAD classifier for localization of cancer, the elastic registration procedure for alignment of diagnostic MRI to CT (via the planning MRI), and the mapping of cancer onto CT for improved dose

planning. In Section 6.3 we present the results of the CATT framework for improved dose planning on 79 sets of CT and multiprotocol MR images from 10 prostate studies, and generate radiotherapy dose plans based on target volumes defined using both the prostate boundary and the suspected cancer region. Concluding remarks and future directions are presented in Section 6.4.

6.2 Computer-Assisted Targeted Therapy (CATT)

6.2.1 Data Description and Preprocessing

For 10 patients with prostate cancer, scheduled to undergo radiotherapy treatment, two MRI and one CT image sets were acquired. These image data sets are described in Table 6.1, where the large FOV differences between diagnostic MRI (row 1) and planning MRI and CT (rows 2,3) should be noted. The CT image set is a planning, or simulation study that is used to determine attenuation characteristics necessary to formulate a radiotherapy dose plan capable of delivering sufficient levels of radiation to the targeted volumes. As such, a large FOV is necessary to encompass the entire body.

Table 6.1: Summary of the prostate image data sets acquired for each of 10 patients.

Set Number	Slice Notation	Modality	Description/Purpose
1	\mathcal{I}^d	T2-w MRI	Diagnostic
2	\mathcal{I}^p	T2-w MRI	Treatment Planning
3	\mathcal{I}^{CT}	CT	Planning/Simulation

Set Number	Dimensions (mm ³)	Voxel Size (mm ³)
1	$120 \times 120 \times 107$	$0.5 \times 0.5 \times 4$
2	$340 \times 340 \times 256$	$1.3 \times 1.3 \times 8$
3	$500 \times 500 \times 320$	$1 \times 1 \times 2$

In the remainder of this paper, we denote a 2D slice of a volume as $\mathcal{I} = (C, f)$, where C is a finite 2D rectangular grid of pixels and $f(c)$ is the image intensity at each pixel $c \in C$. Slices of diagnostic MRI, planning MRI, and planning CT are thus defined as \mathcal{I}^d , \mathcal{I}^p and \mathcal{I}^{CT} , respectively. Note that these images are defined on independent coordinate grids C^d , C^p and C^{CT} , and it is the goal of any registration technique to determine spatial transformations that map locations in C^d and C^p to corresponding locations in C^{CT} .

Several preprocessing steps were performed. Correction of MRI bias field inhomogeneity was first performed using the automatic low-pass filter based technique presented by Cohen in [137]. Delineation of the prostate boundary on the diagnostic MRI volume was then manually performed (Figure 6.3(a)) to allow for application of the unsupervised classifier for CaP detection to only the voxels within the prostate. The region identified as the prostate on a slice \mathcal{I}^d is denoted by the set of pixels $G^d \subset C^d$. Finally, the slices of the planning MRI volume containing the prostate were identified. On average, the prostate was found to be visible on about 6 slices of each planning MRI volume.

6.2.2 Localization of Cancer on Diagnostic MRI

The location of cancer on diagnostic MRI is identified using a computer-aided detection (CAD) system comprising the following components, (1) extraction of multiple descriptive textural features at every spatial location (voxel) within the prostate, (2) non-linear dimensionality reduction on the texture feature space via graph embedding [138] to project the textural signatures associated with each voxel into a reduced-dimensional sub-space, and (3) replicated k -means clustering [102] to reliably partition voxels of the prostate into distinct tissue classes, including cancer. These components are briefly described below.

Texture feature extraction

We have previously shown the utility of texture features in distinguishing cancerous from non-cancerous regions in the prostate [5]. We extract 60 unique features from three classes of 3D texture attributes. These include, (1) first order statistical attributes calculated on graylevel distributions, (2) gradient operators, and (3) second order statistical quantities calculated on distributions of spatial co-occurrence of graylevels [78].

First order statistical features

Each of mean, standard deviation, median and range, are calculated over a neighborhood $\mathcal{N}_{c,\kappa}$ of size $\kappa \in \{5, 7\}$ centered on each voxel.




Figure 6.3: (a) Original diagnostic MRI image with prostate boundary manually delineated (green). (b),(c) The median and Haralick correlation features of the 60 total textural features extracted from (a). Note the enhanced visibility of the hypointense lesion in (b) that is not visible in (a). (d) RGB visualization of 3 dimensional embedded feature space obtained by application of the non-linear dimensionality reduction scheme to the full texture space of the prostate voxels. (f) Map of the cancer region obtained by replicated k -means clustering in the embedded feature space in (e).

Gradient features

Gradient features comprising the responses from 9 distinct Sobel operators, oriented along each of the 3 coordinate axes and 6 diagonals, and from the spatial derivatives in each direction plus the gradient magnitude are calculated.

Second order statistical features

Haralick or second order statistical features characterize spatial co-occurrence [78], where any pixel $d \in \mathcal{N}_{c,\kappa}$ is defined as a κ neighbor of c . A $M \times M$ co-occurrence matrix $P_{c,\kappa}$ associated with $\mathcal{N}_{c,\kappa}$ is then computed, where M is the chosen number of gray level bins. The value at any location $[u, v]$, where $u, v \in \{1, \dots, M\}$, in the matrix $P_{c,\kappa}[u, v]$ represents the frequency with which two distinct pixels with associated image

intensities $f(c) = u$, $f(d) = v$ are adjacent. From $P_{c,\kappa}$, the 13 Haralick features comprising energy, entropy, inertia, correlation, inverse difference moment, two information correlation measures, sum average, sum variance, sum entropy, different average, difference variance and difference entropy were extracted at every pixel $c \in C$, for $\kappa \in \{5, 7\}$, and $M \in \{64, 128\}$.

Each pixel $c \in C$ is now associated with a high dimensional texture feature vector $\mathbf{F}(c)$ rather than a scalar intensity $f(c)$. Details of the calculation of these textural features described in [5]. Figures 6.3 (b),(c) show the median and Haralick correlation feature images corresponding to the intensity image in Figure 6.3 (a). Note that the hypointense lesion in the medial section of the gland is more easily discernible on the texture maps (Figures 6.3 (b),(c)) compared to the intensity image in Figure 6.3 (a)

Non-linear dimensionality reduction

Object-class discriminability can be improved by projecting the data into a reduced-dimensional embedding space, thus allowing for greater separation of the data classes. To reduce the high dimensional texture space $\mathbf{F}(c)$ to a lower dimensional representation $\mathbf{S}(c)$ for $c \in G^d$, graph embedding is employed via the normalized cuts algorithm [138]. The aim of graph embedding [29, 138] is to find a vector $\mathbf{S}(c), \forall c \in G^d$ such that the relative adjacency of any two pixels $c_i, c_j \in G^d, i, j \in \{1, \dots, |G^d|\}$ in high dimensional space is preserved in the low dimensional space. To compute $\mathbf{S}(c), \forall c \in G^d$, an affinity matrix $W \in \Re^{N \times N}$, where $N = |G^d|$ is defined as,

$$W(i, j) = e^{-\|\mathbf{F}(c_i) - \mathbf{F}(c_j)\|_2}, \quad (6.1)$$

for each $c_i, c_j \in G^d$. The embedding space $\mathbf{S}(c)$, for each $c \in G^d$, for $m \ll 60$ dimensions, is defined by the eigenvectors \mathbf{X} corresponding to the m smallest eigenvalues of $(D - W)\mathbf{X} = \lambda D\mathbf{X}$, where D is a diagonal matrix with elements $D(i, i) = \sum_j W(i, j)$. Thus, $\mathbf{S}(c_i)$ is obtained by extracting the first m components of the i^{th} row of \mathbf{X} . Since linear projections are not utilized to compute these projections (such as done by PCA) [139], graph embedding is capable of discovering inherent non-linearity in the

data. The reduced dimensional space can be visualized for $m = 3$ by representing every location $c \in G^d$ on \mathcal{I}^d by its embedding coordinates $\mathbf{S}(c)$ and scaling these values to display as an RGB image, as shown overlaid on a diagnostic MRI slice in Figure 6.3(d). Note how dissimilar regions can be more easily discerned in Figure 6.3(d), where similar colors represent the same class, as compared to Figure 6.3(a).

Replicated k-means clustering

Unsupervised replicated clustering is used to partition each pixel $c \in G^d$ into one of k classes based on the embedding representation $\mathbf{S}(c)$ obtained from corresponding textural descriptions. The k -means algorithm provides an efficient and unsupervised method for clustering observations. However, when random initialization of cluster centroids is used, k -means is not deterministic. Replicated clustering has been shown to provide more stable clusters [102] by selecting from T independent runs of k -means, the clustering with the smallest average intra-class variance. Defining $V_t^q \subset G^d$ as the set of pixels belonging to cluster $q \in \{1, \dots, k\}$ from k -means replication $t \in \{1, \dots, T\}$, the cluster centroid is defined by $\mathbf{S}_t^q = \frac{1}{|V_t^q|} \sum_{c \in V_t^q} \mathbf{S}(c)$. The optimal clusters $V_\tau^q \subset G^d$, $\tau \in \{1, \dots, T\}$ are found by

$$\tau = \underset{t}{\operatorname{argmin}} \left[\frac{1}{k} \sum_{c \in V_t^q, q} \frac{1}{|V_t^q|} \|\mathbf{S}(c) - \mathbf{S}_t^q\| \right]. \quad (6.2)$$

Therefore, replicated k -means is utilized here to classify $\mathbf{S}(c)$, $\forall c \in G^d$ as one of k classes. We select $k = 6$ to allow clustering of many dissimilar tissue classes, including benign epithelium, stroma and hyperplasia, cancer, atrophy, and structures such as blood vessels and the urethra. Figure 6.3(c) shows a tumor region identified as one of the clusters generated using the replicated k -means technique. Note that since the technique is unsupervised, the cluster corresponding to cancer must be manually selected based on appearance and spatial configuration of the cluster within the prostate. In our experiments this was done by an expert radiologist. The set of pixels belonging to the cluster representing cancer is denoted $\Omega^d \subset G^d \subset C^d$.

6.2.3 Elastic Registration of Diagnostic MRI, Planning MRI, and CT

The goal of registration is to determine the spatial transformations $\mathbf{T}^{\text{dC}}(c)$ and $\mathbf{T}^{\text{pC}}(e)$ in order to map coordinates $c \in C^{\text{d}}$ and $e \in C^{\text{p}}$ onto the corresponding locations in C^{CT} . This allows for the transformation of both \mathcal{I}^{d} and \mathcal{I}^{p} into alignment with \mathcal{I}^{CT} . A two stage registration approach is used to first determine $\mathbf{T}^{\text{dp}}(c)$, the transformation from each location $c \in C^{\text{d}}$ to a corresponding $e \in C^{\text{p}}$, followed by $\mathbf{T}^{\text{pC}}(e)$, the transformation from each location $e \in C^{\text{p}}$ into C^{CT} . The registration of \mathcal{I}^{d} to \mathcal{I}^{CT} involves the steps below.

1. **Multiprotocol Registration:** Automated affine registration of \mathcal{I}^{d} to \mathcal{I}^{p} via maximization of NMI [56] is performed to determine the mapping $\mathbf{T}^{\text{dp}} : C^{\text{d}} \mapsto C^{\text{p}}$ by

$$\mathbf{T}^{\text{dp}} = \underset{\mathbf{T}}{\operatorname{argmax}} \left[\text{NMI}(\mathcal{I}^{\text{p}}, \mathbf{T}(\mathcal{I}^{\text{d}})) \right], \quad (6.3)$$

where $\mathbf{T}(\mathcal{I}^{\text{d}}) = (C^{\text{p}}, f^{\text{d}})$ so that every $c \in C^{\text{d}}$ is mapped to new spatial location $e \in C^{\text{p}}$ (i.e. $\mathbf{T}(c) \Rightarrow e$ and $c \Leftarrow \mathbf{T}^{-1}(e)$). The NMI between \mathcal{I}^{p} and $\mathbf{T}(\mathcal{I}^{\text{d}})$ is defined as,

$$\text{NMI}(\mathcal{I}^{\text{p}}, \mathbf{T}(\mathcal{I}^{\text{d}})) = \frac{H(\mathcal{I}^{\text{p}}) + H(\mathbf{T}(\mathcal{I}^{\text{d}}))}{H(\mathcal{I}^{\text{p}}, \mathbf{T}(\mathcal{I}^{\text{d}}))}, \quad (6.4)$$

in terms of the marginal and joint entropies,

$$H(\mathcal{I}^{\text{p}}) = - \sum_{e \in C^{\text{p}}} p^{\text{p}}(f^{\text{p}}(e)) \log p^{\text{p}}(f^{\text{p}}(e)), \quad (6.5)$$

$$H(\mathbf{T}(\mathcal{I}^{\text{d}})) = - \sum_{e \in C^{\text{p}}} p^{\text{d}}(f^{\text{d}}(\mathbf{T}^{-1}(e))) \log p^{\text{d}}(f^{\text{d}}(\mathbf{T}^{-1}(e))), \text{ and} \quad (6.6)$$

$$H(\mathcal{I}^{\text{p}}, \mathbf{T}(\mathcal{I}^{\text{d}})) = - \sum_{e \in C^{\text{p}}} p^{\text{pd}}(f^{\text{p}}(e), f^{\text{d}}(\mathbf{T}^{-1}(e))) \log p^{\text{pd}}(f^{\text{p}}(e), f^{\text{d}}(\mathbf{T}^{-1}(e))), \quad (6.7)$$

where $p^{\text{p}}(\cdot)$ and $p^{\text{d}}(\cdot)$ are the graylevel probability density estimates, and $p^{\text{pd}}(\cdot, \cdot)$ is the joint density estimate. Note that the similarity measure is calculated over the pixels in C^{p} , the coordinate grid of planning MRI. Despite the small FOV of \mathcal{I}^{d} and the large FOV of \mathcal{I}^{p} (Figures 6.1(a),(c)), an affine transformation is sufficient since the two MRI protocols are acquired in the same scanning session

with minimal patient movement between acquisitions. Since both T2-weighted MRI protocols are also similar in terms of intensity characteristics, NMI is effective in establishing optimal spatial alignment.

2. **Multimodal Registration:** Elastic registration of \mathcal{I}^P to \mathcal{I}^{CT} using control point-driven thin plate splines (TPS) [92] to define mapping $\mathbf{T}^{PC} : C^P \rightarrow C^{CT}$ is performed. An intuitive point-and-click graphical interface was developed and employed for identifying pairs of corresponding spatial locations between \mathcal{I}^P and \mathcal{I}^{CT} . Having nearly the same FOV and similar spatial resolution, landmarks on \mathcal{I}^P , including the femoral head, pelvic bone and prostate capsule, are identifiable on the corresponding \mathcal{I}^{CT} . For example, in Figures 6.1(b),(c) corresponding \mathcal{I}^{CT} and \mathcal{I}^P images are shown, where the wide FOV of planning MRI encompasses peripheral anatomical features such as the hip bones and spine, which are not visible on \mathcal{I}^d (see Figure 6.1(a)). Note that while $\mathbf{T}^{dp}(c)$ determined in Stage 1 is implemented as an affine transformation, $\mathbf{T}^{PC}(e)$ is a non-parametric deformation field, elastically mapping each coordinate in C^P to C^{CT} .

3. **Combination of transformations:** The direct mapping $\mathbf{T}^{dC}(c) : C^d \mapsto C^{CT}$ of coordinates of \mathcal{I}^d to \mathcal{I}^{CT} is obtained by the successive application of \mathbf{T}^{dp} and \mathbf{T}^{PC} ,

$$\mathbf{T}^{dC}(c) = \mathbf{T}^{PC}(\mathbf{T}^{dp}(c)).$$

Thus, using $\mathbf{T}^{dC}(c)$, each coordinate in C^d is mapped into C^{CT} .

In summary, the procedure described above is used to obtain the following spatial transformations: (1) \mathbf{T}^{dp} mapping from C^d to C^P , (2) \mathbf{T}^{PC} mapping from C^P to C^{CT} , and (3) \mathbf{T}^{dC} mapping from C^d to C^{CT} . Using \mathbf{T}^{dC} and \mathbf{T}^{PC} , the diagnostic and planning MRI that are in alignment with CT are obtained as $\tilde{\mathcal{I}}^d = \mathbf{T}^{dC}(\mathcal{I}^d) = (C^{CT}, f^d)$ and $\tilde{\mathcal{I}}^P = \mathbf{T}^{PC}(\mathcal{I}^P) = (C^{CT}, f^P)$, respectively. For example, Figure 6.1(d) shows $\tilde{\mathcal{I}}^d$ obtained by aligning with the corresponding \mathcal{I}^{CT} . Note the non-linear nature of \mathbf{T}^{dC} is clearly visible in Figure 6.1(d).

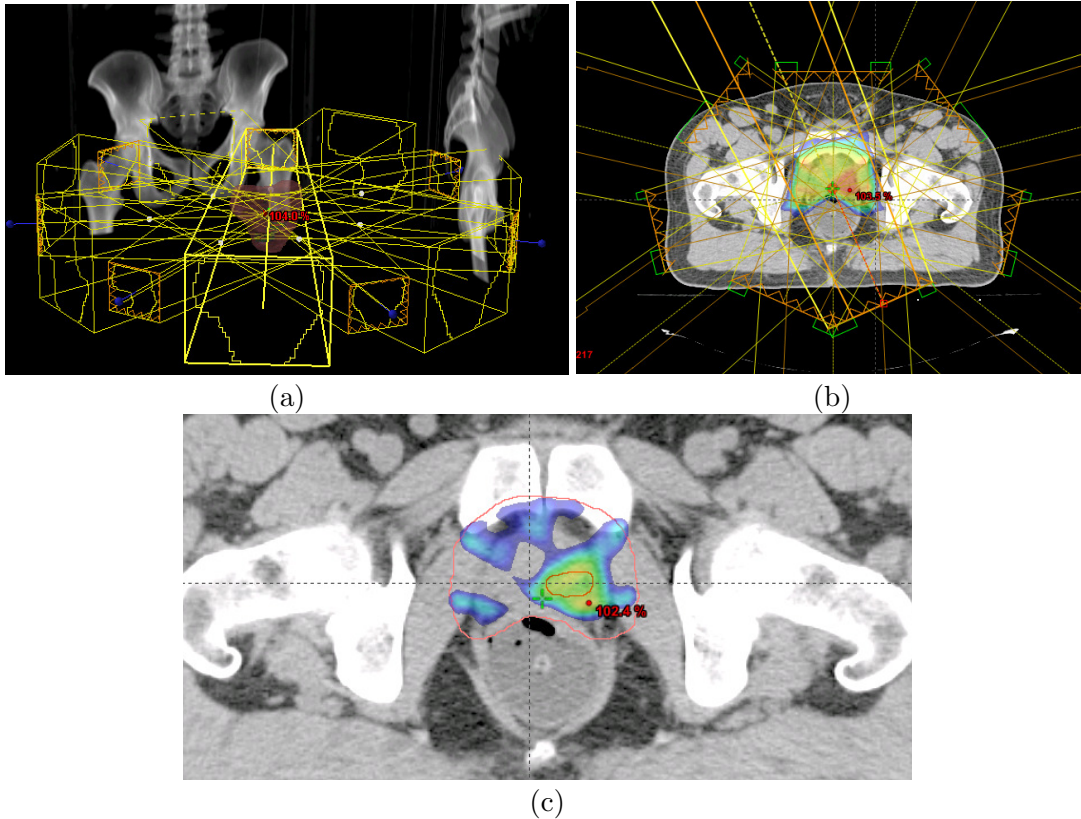


Figure 6.4: (a) 3D visualization of the orientations and shapes of the 7 dynamically shaped radiation beams used to design the IMRT dose plan. (b) Axial view of a CT slice with the 7 beams and the resulting dose map thresholded at 7920 cGy (minimum dose for the PTV) overlaid. (c) Adjacent CT slice with the dose map thresholded at 8640 cGy (minimum dose for the cPTV) and boundaries for the PTV and cancer overlaid.

6.2.4 Mapping of Tumor Location from Diagnostic MRI onto CT

After registration, the CAD result, represented by a set of spatial locations labeled as Ω^d on \mathcal{I}^d , is mapped via the transformation \mathbf{T}^{dC} to set Ω^{CT} on \mathcal{I}^{CT} . For example, the label Ω^d shown in red on \mathcal{I}^d (Figure 6.3(e)) is mapped onto CT slice \mathcal{I}^{CT} (Figure 6.3(f)) via \mathbf{T}^{dC} .

6.2.5 Dose Plan Generation

Two target volumes are defined and used to generate a single IMRT dose plan using the Varian[©] Eclipse software. These are the planning target volume (PTV), which is based on the radiologist's outline of the prostate on CT plus margins, and a cancer PTV (cPTV), based on Ω^{CT} plus margins. The IMRT plan utilizes 7 beam angles and

a dynamic multileaf collimator (MLC) to shape the beam and deliver no less than 7920 cGy to the PTV and no less than 8640 cGy to the cPTV. Figure 6.4(a) shows a 3D view of the 7 beam angles and the shapes of the MLC at various points over the course of the treatment. An axial view of the beams are shown in Figure 6.4(b) with the resulting dose map thresholded at 7920 cGy overlaid onto the CT slice. The increased dose to the cPTV is demonstrated in Figure 6.4(c) where the dose map thresholded at 8640 cGy covers the cPTV, which is indicated by the inner boundary. Note that while dose escalation to the cancer is achieved, the rest of the PTV receives no less radiation, hence ensuring that the entire prostate receives a full dose.

6.3 Results of CATT

6.3.1 Unsupervised CAD Classifier to Identify Tumor Labels on Diagnostic MRI

The unsupervised CAD classifier (Section 6.2.2) was applied to identify pixels belonging to the cancer class (G^d) for each of 10 diagnostic MRI studies. Figures 6.3(a)-(c) show a 2D section of diagnostic MRI of the prostate and two representative textural feature images. The features shown in Figures 6.3(b),(c) are the median calculated for each pixel c with neighborhoods $\mathcal{N}_{c,\kappa}$ of size $\kappa = 5$ and the Haralick correlation for $\mathcal{N}_{c,\kappa}$ of size $\kappa = 7$. Shown in Figure 6.3(d) is a color representation of the 3 dimensional feature space $\mathbf{S}(c)$ obtained by application of the graph embedding scheme to $\mathbf{F}(c)$ for each $c \in G^d$. The corresponding cancer label Ω^d was obtained for each of the 10 prostate volumes by applying the replicated k -means clustering technique to all $\mathbf{S}(c)$, $c \in G^d$. The identified cancer cluster Ω^d , manually identified by a radiologist, is shown in Figure 6.3(e). No quantitative evaluation of CAD accuracy is possible was not available for these patients; in previous work [1] we identified ground truth for prostate cancer on MRI by registering the imaging with corresponding whole mount *ex vivo* histological sections.

6.3.2 Multiprotocol and Multimodal Image Registration

Registration was performed as described in Section 6.2.3 for the 79 data sets of \mathcal{I}^d , \mathcal{I}^p , and \mathcal{I}^{CT} from the 10 patients. Figures 6.5(a) and (b) show a corresponding pair of diagnostic and planning MRI slices (\mathcal{I}^d and \mathcal{I}^p). Figure 6.5(c) shows $\tilde{\mathcal{I}}^d$, \mathcal{I}^d after automated affine registration to \mathcal{I}^p , while the fusion of $\tilde{\mathcal{I}}^d$ and \mathcal{I}^p is shown as a blended overlay in Figure 6.5(d). In Figures 6.5(e)-(h), the registration of four different \mathcal{I}^d and \mathcal{I}^p slices is demonstrated as checkerboard patterns of \mathcal{I}^p and $\tilde{\mathcal{I}}^d$. The continuity of the internal structures in the prostate and surrounding tissues between the registered images can be appreciated in Figure 6.5(e)-(h). Clearly, by sequentially acquiring the diagnostic and planning MRI volumes in the same imaging session, an affine transformation using the NMI similarity measure appears to be sufficient to bring \mathcal{I}^d and \mathcal{I}^p into alignment.




Figure 6.5: (a) Diagnostic MRI is affinely registered to (b) planning MRI. (c) The registered diagnostic MRI is shown in (d) overlaid onto (b). (e)-(h) Checkerboard patterns of four additional pairs of aligned diagnostic and planning MRI demonstrate accuracy of registration, as evidenced the continuity of prostate capsule and internal structures across checks.

Figures 6.6(a)-(d) demonstrate the elastic registration results of 4 sets of corresponding \mathcal{I}^d , \mathcal{I}^p and \mathcal{I}^{CT} images as overlays of \mathcal{I}^{CT} and the registered images $\tilde{\mathcal{I}}^d$ and $\tilde{\mathcal{I}}^p$. The alignment of internal structures is evident in the vicinity of the prostate, between the hip joints (within the FOV of the diagnostic MRI). The non-linearity of the image transformations required to align MRI with CT (Figures 6.6(a)-(d)) is clearly demonstrated by the grids shown in Figures 6.6(e)-(h).




Figure 6.6: Elastic registration of both planning and diagnostic MRI to CT. (a)-(d) Overlay of all three imaging modalities of four slices from a single prostate study. (e)-(h) Deformed grid patterns for slices corresponding to (a)-(d) illustrate the non-linear nature of the deformation field required to align the prostate on diagnostic MRI to the visible gland on CT.

6.3.3 Tumor Mapping onto CT

The transformations \mathbf{T}^{dC} determined between each \mathcal{I}^d and \mathcal{I}^{CT} slice are applied to the corresponding cancer label Ω^d of each \mathcal{I}^d to obtain new spatial locations Ω^{CT} on each \mathcal{I}^{CT} . Figures 6.7(a), (b) show two different CT slices with Ω^{CT} as a white mask and the resulting cPTV boundary as a red line.




Figure 6.7: (a), (b) Two different CT studies shown with cancer labels Ω^{CT} (white) mapped from MRI by registration and the resulting cPTV (red line). Resulting dose intensity maps for the slices in (a), (b) are shown thresholded at the minimum dose for PTV (7920 cGy, outer orange line) in (c), (d) and thresholded at the minimum dose for cancer (8640 cGy, inner red line) in (e), (f). These demonstrate dose escalation to the cancer PTV while maintaining dose to the prostate PTV. Note that the original dose plan has not been shown.

6.3.4 Dose Planning

From the CT volumes with PTV and cPTV delineated (shown as red lines in Figures 6.7(a),(b)), dose plans were generated for each of the 10 patient studies. Figures 6.7(c) and (d) show the radiation dosage maps for the slices in Figures 6.7(a) and (b). To illustrate the dose to the PTV, the radiation intensity is thresholded at 7920 cGY in Figures 6.7(c) and (d), while to illustrate the dose escalation to the cPTV, the radiation intensity is thresholded at 8640 cGY in Figures 6.7(e) and (f).

6.4 Concluding Remarks

With high resolution diagnostic MRI, it is now possible to visually identify presence and extent of prostate cancer. Additionally, computer-aided classifiers for detecting disease extent on MRI have been developed. The tumor location identified on the MRI by CAD classifiers was mapped onto CT via a two step registration methodology, allowing formulation of a more precise dose plan. The implications of a dose plan targeted only at the tumor include (1) reduced exposure to the bladder, rectum, and other benign tissues, and (2) dose escalation to the tumor. Thus, the approach described here could translate to a real reduction in side effects and increased efficacy of radiotherapy treatments.

It is important to note that there is no straightforward way of evaluating the accuracy of tumor localization on CT. Validation of the computer-assisted dose plans would involve a long-term, clinical trial whereby CAD-assisted therapy is compared against conventional therapy targeted at the entire prostate without the benefit of CAD. The work presented here will pave the way for such a clinical trial.

For future work, we are currently implementing a fully automated registration paradigm using free form deformations, such as previously described in our work on MR-histology registration [1]. Another avenue of future work will focus on automated prostate segmentation using a previously developed feature-driven active contour technique [140]. Finally, we plan to investigate the use of a supervised classifier to drive the CAD system, which will remove the interactive component of the current clustering scheme and potentially provide more accurate cancer labels.

Chapter 7

HistoStitcher: An Interactive Program for Accurate and Rapid Reconstruction of Digitized Whole Histological Sections from Tissue Fragments

7.1 Need for Efficient Digital Reconstruction of Whole Histological Images

Histological sections of tissue specimens obtained via biopsy or surgical excision, such as lumpectomy or radical prostatectomy, are used for identifying presence and extent of disease, and if resolution is sufficient, the grade and malignancy of the disease [141–143]. In cases where the patient is scheduled for surgical resection of the diseased tissue or gland, it may be valuable to map disease extent from *ex vivo* histology sections onto *in vivo* radiological imaging, which is commonly acquired prior to excision, to discover signatures for disease on *in vivo* imaging [5, 28, 107]. For example, men with prostate cancer confirmed by biopsy and who are scheduled to undergo radical prostatectomy may receive pre-operative MRI of the prostate for identifying the presence of extra capsular spread. By registering these pre-operative MRI with corresponding *ex vivo* histology sections obtained from the radical prostatectomy specimen (see Figures 7.1(a) and (b)), pathological regions on histology (purple regions in Figures 7.1(a)) are then mapped onto the registered MRI (shown in green on Figure 7.1(c)). To achieve the required histology annotation and multi-modal registration steps described above, whole histological sections of entire slices of the specimen rather than fragmented sections of tissue are required.

Obtaining whole-mount histological sections of large specimens, such as an enlarged prostate gland, while maintaining tissue integrity is not always feasible. For large specimens, a single glass slide is often too small to mount the entire contiguous tissue




Figure 7.1: (a) Reconstructed pseudo-whole-mount histology from a prostate specimen with cancer (delineated in dark purple) may be used to identify signatures for disease on pre-operative *in vivo* imagery, similar to the MRI image shown in (b), by registering the two images. (c) The registered MRI slice that is in alignment with whole histology in (a) allows mapping of cancer onto MRI (green). (d) Quadrants comprising the whole histology section in (a) must be first reconstructed into a pseudo whole-mount section to facilitate improved annotation across cuts and registration with MRI in (b).

section, necessitating the cutting of the section into smaller fragments and mounting them onto multiple slides. Further, it is difficult to slice large specimens thin enough to achieve the very high spatial resolution required for accurate annotation without the use of specialized procedures and microtomes. For example, obtaining a whole slide for the specimen shown in Figure 7.1(b) with a thickness of only $4\mu\text{m}$ is difficult, prone to specimen damage, and inefficient, requiring exorbitant amounts of preparation time. As a result of these challenges, it is often preferable to adopt the much simpler approach of cutting large tissue specimens into smaller fragments and preparing multiple slides for separate analysis, such as quadrant sections shown in Figures 7.1(d).

Tissue fragments spread across multiple slides presents a significant hindrance for both (1) digital annotation of disease extent by a pathologist, which requires integration of visual cues across disjoint section boundaries, and (2) image processing tasks, such as registration of whole histological sections with *in vivo* imaging modalities. Annotation of multiple slides by a pathologist may be performed separately, however when the targeted pathology of interest crosses the boundary between slides, the process is complicated as the pathologist must repeatedly alternate between adjacent slides. As with any strenuous manual task, the annotations obtained in this manner may be inaccurate or unreliable. In addition to tissue annotation, whole histology slides are valuable in facilitating correlation with *in vivo* imagery by image registration, and much

work has been done on registration of whole-mount histology with *in vivo* radiological imaging [4, 19, 112, 144]. On the other hand, registration of smaller histologic image fragments to a sub-image or region of the *in vivo* data has not been widely investigated and is likely to be a more challenging registration task. Thus, in order to streamline and improve pathologist slide annotation, and facilitate image processing tasks requiring whole histological sections, it becomes useful to reconstruct a pseudo-whole-mount histological section from multiple individual fragments [28].

With the spread and growing acceptance of digital pathology [141–143], it is feasible for high resolution whole-mount sections to be digitally reconstructed from the images of the smaller fragments. Digitally reconstructed whole-mount histological sections would not only facilitate a variety of image processing tasks, but if pathologists can perform digital annotation of disease extent on high resolution histology images on computer monitors, annotations would be greatly improved compared to manual labeling of slides with felt pen or drawing on standardized examination sheets. Annotation could be improved in terms of both accuracy and efficiency using digital images by allowing labeling at any level of detail and for editing and revision of the markup. For example, in Figure 7.1(a) are the digitized histology quadrants of a section from a prostatectomy specimen with cancer. The closest corresponding MRI slice is shown in Figure 7.1(b). By digitally combining, or stitching, the quadrants a pseudo-whole-mount section can be generated (see Figure 7.1(c)), upon which efficient pathologist annotation can be performed on a high resolution computer monitor. The resulting cancer label established by analysis of the reconstructed pseudo-whole-mount section is shown in purple, and can be seen to cross the cut between the original quadrants. Finally, having a reconstructed whole histology section, registration with corresponding *in vivo* imagery, such as MRI, can be performed to achieve mapping of spatial extent of disease (in this case prostate cancer) from the annotated whole histology onto corresponding radiological *in vivo* imaging. For example, the aligned MRI slice obtained in a previous study [5] is shown in Figure 7.1(d) with the mapped histopathologic cancer label shown in green. In this paper, we present a software utility called HistoStitcher[®] for computerized reconstruction of a whole histological section from digital images of the multiple slides of tissue fragments.

While we demonstrate the utility of HistoStitcher[®] in the context of prostate sections in this work, the program is just as readily applicable to other domains and applications such as reassembly of whole histology sections of breast lumpectomy or mastectomy.

The rest of the paper is organized as follows. In Section 7.2 we discuss the challenges and requirements associated with digital reconstruction of high resolution histopathologic images. In Section 7.3.1 we present an overview of the HistoStitcher[®] software, including the interface and workflow. In Section 7.3.3 we detail the mathematical methods used to determine the optimal image transformation based on user-defined control points. In Section 7.3.4 we describe the application of the transformation to the full resolution histology and the process of stitching to create a larger image. In Section 7.3.6 we describe the features of HistoStitcher[®] that facilitate operation on very large images. In Section 7.4 we present formalized evaluation criteria and a 6 point scoring scheme for comparing and assessing the quality digitally reconstructed histology sections. In Section 7.5 we demonstrate the use of HistoStitcher[®] for reassembly of quadrants of prostate histology sections, and compare the results to manually reassembled (using Photoshop) sections and to block face photographs of the gland taken prior to cutting into quadrants. The reconstruction evaluation scheme is applied to compare the quality of stitches generated using HistoStitcher[®] to corresponding stitches generated using Photoshop for 6 prostate additional prostate studies. In Section 7.6 we present concluding remarks and enumerate further applications of HistoStitcher[®].

7.2 Challenges in Digital Reconstruction of High Resolution Whole Histology

At institutions where high resolution digitization of excised tissue specimens are performed, an imaging software like Photoshop^{®1} is typically employed to digitally reconstruct a whole histology image when the tissue is fragmented across multiple slides. The use of conventional image editing programs for reassembly of high-resolution histology images is both inefficient and inadequate on several accounts. Firstly, these programs

¹<http://www.adobe.com/photoshop/>

do not allow for simultaneous rotation, translation and scaling, which makes the digital stitching process using these programs tedious and the resulting transformation difficult to parameterize and record. Secondly, operating on high resolution (10x magnification) histology is extremely memory intensive. For example, a single digital histology image, which can be larger than 15,000 x 15,000 RGB pixels or 675 MBytes, necessitates careful memory handling to perform even basic image manipulations. The size of the images underscores the need for efficient parameterization of the transformations, which must be easily stored and applied to the full size images in chunks at a later time.

Despite the shortcomings of commercially available image editing programs for the task, digital reassembly of the fragments of whole histological sections requires expert interaction to ensure accurate reconstruction, where the corresponding edges of adjacent sections are aligned or “stitched” to preserve the continuity of anatomical structures of the tissue. Common automatic techniques for image stitching generally rely on significant overlap between adjacent images, such as with photograph [145–147] or microscopy [148–151] mosaic generation. However no such overlap exists between adjacent histology fragments. Automated edge matching techniques [152–155], which are generally designed for rigid objects such as puzzles and broken tiles, break down on account of uneven distortion and loss of tissue along the edges of histological fragments, leading to highly dissimilar contours of the two edges that were originally adjacent in the contiguous section. For example, the prostate histology quadrants shown in Figure 7.1(d) do not possess distinctive curvatures or other features on the contours of the adjacent edges, either at a macro or micro scale, which can be reliably used to align the edges. Many large macroscopic structures within the tissue are however visible that cross between quadrants. Thus, instead of trying to characterize overall edge similarity, the pathologist must observe the organization of tissue architecture and infer the continuity of structures across the cut. At these points along the cut where structural continuity is clear, the corresponding points along the two adjacent contours can simply be marked and used to perform the stitching.

HistoStitcher[®], our interactive software package presented in this paper, adopts this approach based on identification of pairs of control points since (1) it provides a simple

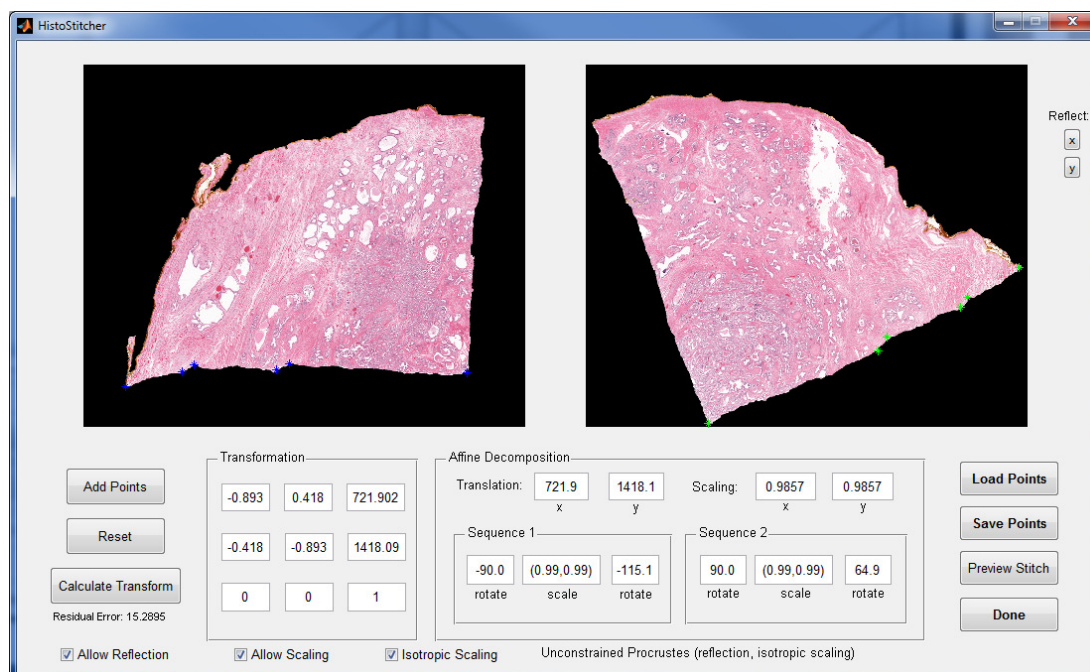


Figure 7.2: The HistoStitcher[®] graphical user interface where two histology pieces from the same section are digitally combined by selecting control point pairs along common boundaries. An optimal transformation of one fragment onto the other is automatically determined based on user specified constraints.

way for a pathologist to drive the reassembly process, (2) as few as three pairs of control points are sufficient to obtain a transformation, and (3) the transformation is easily parameterized, stored, and applied to high resolution images. While the underlying techniques employed by HistoStitcher[®] for image transformation using control points are not novel per se, their application to interactive digital reassembly of histology fragments and subsequent visualization of whole sections is. The design and operation of HistoStitcher[®] software is described in Section 7.3 below.

7.3 Features and Operation of HistoStitcher

7.3.1 Software Overview and Workflow

The HistoStitcher[®] software is comprised of an intuitive graphical user interface (GUI) and a set of computational routines for reassembly of a single stitched histological section from two images containing the adjacent fragments. A screenshot of the main HistoStitcher[®] GUI with two prostate histology quadrants loaded is shown in Figure 7.2.




Figure 7.3: (a)-(d) Histology slide images of fragments (quadrants) of a whole section of a prostate specimen are reassembled using HistoStitcher[®]. (e) Half prostate histology section reconstructed by stitching quadrants in (a) and (b). (f) Three-quarters prostate histology section from stitching of quadrant in (c) with half in (e). (g) Final reconstructed pseudo-whole prostate histology section following stitching of quadrant in (d) with (e).

The operation of the program for stitching two adjacent histology fragments involves the following steps,

1. Load two adjacent images: the stationary image (left image) and the moving image (right image) that is to be transformed so that common edges are in the best possible alignment.
2. Select pairs of corresponding anatomical landmarks (points) that exist along the cut separating two adjacent sections. (Add Points button).
3. Specify constraints on the image transformation, including reflection, scale, and scale isotropy. (Tick boxes at bottom).
4. Automatically calculate the optimal coordinate transformation to minimize the error (in the least-squares sense) between pairs of control points. (Calculate Transform button).

5. Transform the moving image, bringing both images into a common coordinate frame.
6. Combine the images to generate the stitched image, expanding the image canvas as necessary. (Preview Stitch button).
7. Return the parameterized linear coordinate transformation and the full-resolution stitched image. (Done button).
8. Continue to Step 1 with the next adjacent fragment and the newly stitched image.

This general approach to reassembling a whole histology section from several fragments, such as the four quadrants of the prostate histology section shown in Figure 7.1(d), involves the cumulative stitching of two images at a time until all fragments have been combined. The process of stitching these quadrants, shown in Figures 7.3(a)-(d), proceeds as follows. First, two adjacent fragments (Figures 7.3(a) and (b)) are stitched using the steps above to generate a larger histology image (Figure 7.3(e)). The newly assembled semi-circular histology section is then stitched with the next histology fragment (Figure 7.3(c)) to generate a larger histology image (Figure 7.3(f)) comprising three quarters of the original whole section. Finally, the last histology fragment is stitched with the three-quarters section (Figure 7.3(f)) to generate the pseudo-whole prostate histology section, as shown in Figure 7.3(g).

7.3.2 Notation

The pairs of anatomical landmarks identified using the HistoStitcher[®] GUI (shown as blue and green stars in Figure 7.2) correspond to two sets of N control points, denoted by the homogeneous coordinate matrices $\mathbf{X} = [\mathbf{x}; \mathbf{y}; \mathbf{1}]$ (moving) and $\mathbf{U} = [\mathbf{u}; \mathbf{v}; \mathbf{1}]$ (stationary), where $\mathbf{x}, \mathbf{y}, \mathbf{u}, \mathbf{v}, \mathbf{1}$ are row vectors of length N . We define the stationary and moving histology fragment images as $\mathcal{A} = (f_A, C^A)$ and $\mathcal{B} = (f_B, C^B)$, where $f_A(c)$ and $f_B(c)$ are the image intensity or RGB values at each pixel c in rectangular coordinate sets C^A and C^B , respectively. We denote the transformation of \mathcal{B} as $\beta = \mathbf{T} \circ \mathcal{B} = (f_\beta, C^\beta)$, defined by a new coordinate set C^β and intensity function $f_\beta(c)$ for each

Table 7.1: List of notation and symbols specific to Appendix 7.

Symbol	Description
\mathbf{X}	Control point matrix on moving image
\mathbf{U}	Control point matrix on stationary image
\mathbf{x}, \mathbf{y}	X - and Y -axis components of \mathbf{X}
\mathbf{u}, \mathbf{v}	X - and Y -axis components of \mathbf{U}
$\hat{\mathbf{X}}$	Mean-centered \mathbf{X}
$\hat{\mathbf{U}}$	Mean-centered \mathbf{U}
$\tilde{\mathbf{X}}$	Normalized \mathbf{X}
$\tilde{\mathbf{U}}$	Normalized \mathbf{U}
\mathbf{T}	Linear coordinate transformation
\mathbf{R}	Orthonormal rotation matrix
\mathbf{t}	Translation vector
s	Scale factor
$\bar{\mathbf{X}}$	Mean coordinate of \mathbf{X}
$\bar{\mathbf{U}}$	Mean coordinate of \mathbf{U}
\mathcal{A}	Stationary image on HistoStitcher [®] left
\mathcal{B}	Moving image on HistoStitcher [®] right
β	Image \mathcal{B} transformed by \mathbf{T}
C^β	Coordinate set for β
$f_\beta(c)$	Value of image β at pixel $c \in C^\beta$
α	Transparency value for background pixels
\mathcal{S}	Stitched image comprising \mathcal{A} and β
C^S	Coordinate set for \mathcal{S} s.t. $(C^A \cup C^\beta) \subset C^S$
$f_S(c)$	Value of image \mathcal{S} at pixel $c \in C^S$
N	Number of control points

coordinate $c \in C^\beta$. The stitched image comprising \mathcal{A} and β is denoted $\mathcal{S} = (f_S, C^S)$ and defined on the coordinate frame C^S . A comprehensive list of the main notations employed in this paper is given in Table 7.1.

7.3.3 Determining the Optimal Transformation from Control Point Pairs

The goal of this section is to determine a linear coordinate transformation \mathbf{T} that minimizes the error e in the transformation of \mathbf{X} via $\mathbf{U} = \mathbf{TX} + e$, subject to variable constraints on scaling (anisotropic, isotropic, none) and reflection (allowed, forced, none). To evaluate the accuracy of control point placement, the residual error in the mapping of \mathbf{X} onto \mathbf{U} is calculated as $\|\mathbf{TX} - \mathbf{U}\|^2$ and displayed on the GUI (shown

beneath the “Calculate Transform” button in Figure 7.2).

Depending on the handling of the histology sections during creation and imaging of the slides, different constraints on the number of degrees of freedom (DOF) in the linear coordinate transformation, which is ultimately applied to the entire moving image, may be necessary. It is often necessary to reflect (mirror image) a section if the slice was flipped upside-down when it was scanned. Scaling may be necessary to correct for variable degrees of tissue shrinkage between fragments, or if the magnification used to acquire all sections was not exactly the same. Further anisotropic scaling (two independent scale factors on orthogonal axes) may be necessary when the slicing process causes deformations that are not equal in each direction. However, to prevent overfitting of the coordinates, especially when few or collinear landmarks were identified, it may be necessary to reduce the degrees of freedom by constraining or disallowing scaling and/or reflection. For example, when only one pair of edges are being stitched, the control points are mainly collinear and isotropic scaling should be used to prevent overfitting. When two or more pairs of edges are being stitched, the points are not collinear and anisotropic scaling can be safely used. However, if the object is truly rigid, such as if bone were being stitched, anisotropic scaling would never be used. The mathematical techniques used to find the optimal transformation \mathbf{T} subject to the various constraints are described in Sections 7.3.3-7.3.3.

Unconstrained Affine: Anisotropic Scale Allowing Reflection

Maximal flexibility in the transformation of the moving image (shown in the right side of the GUI) is provided by allowing two independent scale parameters (anisotropic scaling) and reflection, in addition to rotation and translation, to align the selected control points. In the HistoStitcher[®] interface, this configuration corresponds to selecting the “Allow Reflection” and “Allow Scaling” check boxes, and unselecting the “Isotropic Scaling” checkbox. To achieve an alignment with these transformation possibilities, an

unconstrained affine transformation \mathbf{T}_a of \mathbf{X} is defined as,

$$\mathbf{T}_a \mathbf{X} = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \\ \mathbf{1} \end{bmatrix} \quad (7.1)$$

where $a_{11}, a_{12}, a_{13}, a_{21}, a_{22}, a_{23}$ are the 6 independent parameters that comprise \mathbf{T}_a . The four parameters $a_{11}, a_{12}, a_{21}, a_{22}$ collectively control the scale, rotations and any reflection. The parameters a_{13} and a_{23} represent the translations of the moving image that occur after any scale, rotation and reflection operations.

Following selection of the control point pairs \mathbf{U} and \mathbf{X} and selecting the appropriate constraints, the user can then click the “Calculate Transform” button, directing HistoStitcher[©] to solve for \mathbf{T}_a . HistoStitcher[©] finds the optimal transformation that minimizes the error between \mathbf{U} and $\mathbf{T}_a \mathbf{X}$,

$$E(\mathbf{T}_a) = \|\mathbf{T}_a \mathbf{X} - \mathbf{U}\|^2, \quad (7.2)$$

where $\|\cdot\|$ is the Frobenius norm. Equation (7.2) can be minimized by expanding in terms of $a_{11}, a_{12}, a_{13}, a_{21}, a_{22}, a_{23}$ as,

$$E(\mathbf{T}_a) = \sum_{i=1}^N (a_{11}x_i + a_{12}y_i + a_{13} - u_i)^2 + (a_{21}x_i + a_{22}y_i + a_{23} - v_i)^2, \quad (7.3)$$

and differentiating $E(\mathbf{T}_a)$ with respect to each of the six parameters and setting the resulting equations to zero. This produces the following linear system,

$$\begin{bmatrix} \sum x_i^2 & \sum x_i y_i & \sum x_i & 0 & 0 & 0 \\ \sum x_i y_i & \sum y_i^2 & \sum y_i & 0 & 0 & 0 \\ \sum x_i & \sum y_i & \sum 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & \sum x_i^2 & \sum x_i y_i & \sum x_i \\ 0 & 0 & 0 & \sum x_i y_i & \sum y_i^2 & \sum y_i \\ 0 & 0 & 0 & \sum x_i & \sum y_i & \sum 1 \end{bmatrix} \begin{bmatrix} a_{11} \\ a_{12} \\ a_{13} \\ a_{21} \\ a_{22} \\ a_{23} \end{bmatrix} = \begin{bmatrix} \sum u_i x_i \\ \sum u_i y_i \\ \sum u_i \\ \sum v_i x_i \\ \sum v_i y_i \\ \sum v_i \end{bmatrix}, \quad (7.4)$$

where \sum represents the sum over each element i (and j). Solving the system for $[a_{11} \ a_{12} \ a_{13} \ a_{21} \ a_{22} \ a_{23}]^T$ provides the parameters of the optimal transformation \mathbf{T}_a , which is then displayed in the HistoStitcher[©] GUI in the “Transformation” box beneath the stationary image on the left. To allow the operator to assess the accuracy of control point placement, the residual error of the minimization (calculated as $\|\mathbf{T}_a \mathbf{X} - \mathbf{U}\|^2$) is displayed immediately below the “Calculate Transform” button.

Isotropic Scaling with No Reflection

It may often be appropriate to constrain the transformation of the moving image by allowing only a single scale parameter (isotropic scaling), and disallowing any reflection (mirroring). To select this configuration, the operator checks the box for “Isotropic Scaling” and unchecks the box for “Allow Reflection”. Constraining the linear transformation to a single isotropic scaling factor and no reflection involves parameterizing the transformation \mathbf{T}_s as,

$$\mathbf{U} = \mathbf{T}_s \mathbf{X} = \begin{bmatrix} a_1 & -a_2 & a_3 \\ a_2 & a_1 & a_4 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \\ \mathbf{1} \end{bmatrix}, \quad (7.5)$$

where a_1, a_2, a_3, a_4 are the 4 independent parameters that comprise \mathbf{T}_s . The parameters a_1 and a_2 collectively control the rotation and scale, while the negative sign prohibits reflecting the coordinates. The parameters a_3 and a_4 represent the translations. The error term may be formulated as,

$$E(\mathbf{T}_s) = \|\mathbf{T}_s(\mathbf{x}) - \mathbf{v}\|^2, \quad (7.6)$$

$$= \sum_i (a_1 x_i - a_2 y_i + a_3 - u_i)^2 + (a_2 x_i + a_1 y_i + a_4 - v_i)^2, \quad (7.7)$$

and differentiated with respect to each of the four parameters a_1, a_2, a_3 and a_4 . Setting the four resulting equations to zero provides the following matrix formulation,

$$\begin{bmatrix} \sum x_i^2 + y_i^2 & 0 & \sum x_i & \sum y_i \\ 0 & \sum x_i^2 + y_i^2 & -\sum y_i & \sum x_i \\ \sum x_i & -\sum y_i & \sum 1 & 0 \\ \sum y_i & \sum x_i & 0 & \sum 1 \end{bmatrix} \begin{bmatrix} a_1 \\ a_2 \\ a_3 \\ a_4 \end{bmatrix} = \begin{bmatrix} \sum u_i x_i + v_i y_i \\ \sum v_i x_i + u_i y_i \\ \sum u_i \\ \sum v_i \end{bmatrix}. \quad (7.8)$$

When the “Calculate Transform” button is clicked, HistoStitcher[©] solves the system in Equation 7.8 for $[a_1 \ a_2 \ a_3 \ a_4]^T$, yielding the optimal transformation \mathbf{T}_s . This transformation is then displayed in the “Transformation” box.

Isotropic Scaling allowing Reflection

Similar isotropic scaling constraints may be enforced while allowing for reflection. Such constraints may be required if a the mirror image of a slide is loaded. To select this configuration, the operator checks the boxes for both “Isotropic Scaling” and “Allow Reflection”. When reflection is allowed, but not enforced, we seek a transformation $\mathbf{T}_{p,s}$ using Procrustes analysis [156] to determine a constituent 2-by-2 orthonormal transformation matrix \mathbf{R} , a single scaling factor s and a translation vector $\mathbf{t} = [t_x; t_y]$. The major steps of the method employed by HistoStitcher[©] to solve for $\mathbf{T}_{p,s}$ are as follows:

1. Center \mathbf{x} and \mathbf{v} such that their mean corresponds to the origin, obtaining the following coordinate matrices,

$$\hat{\mathbf{X}} = [\hat{\mathbf{x}}; \hat{\mathbf{y}}], \quad (7.9)$$

$$\hat{\mathbf{U}} = [\hat{\mathbf{u}}; \hat{\mathbf{v}}], \quad (7.10)$$

where the elements of $\hat{\mathbf{x}}$ and $\hat{\mathbf{y}}$ are,

$$\hat{x}_i = x_i - \frac{1}{N} \sum_i^N x_i \text{ and } \hat{y}_i = y_i - \frac{1}{N} \sum_i^N y_i,$$

and the elements of $\hat{\mathbf{u}}$ and $\hat{\mathbf{v}}$ are,

$$\hat{u}_i = u_i - \frac{1}{N} \sum_i^N u_i \text{ and } \hat{v}_i = v_i - \frac{1}{N} \sum_i^N v_i.$$

2. The centered coordinates $\hat{\mathbf{X}}$ and $\hat{\mathbf{U}}$ are then scaled to unit norm by,

$$\tilde{\mathbf{X}} = \frac{\hat{\mathbf{X}}}{\|\hat{\mathbf{X}}\|}, \quad (7.11)$$

$$\tilde{\mathbf{U}} = \frac{\hat{\mathbf{U}}}{\|\hat{\mathbf{U}}\|}, \quad (7.12)$$

where $\|\cdot\|$ is the Frobenius matrix norm.

3. Compute the rotation matrix $\mathbf{R} = \mathbf{Q}\mathbf{P}^\top$ where $\mathbf{P}\mathbf{\Sigma}\mathbf{Q}^\top$ is the singular value decomposition (SVD) [157] of $\hat{\mathbf{U}}^\top \hat{\mathbf{X}}$. The SVD $\mathbf{P}\mathbf{\Sigma}\mathbf{Q}^\top$ is comprised of two orthogonal matrixes \mathbf{P} and \mathbf{Q} corresponding to the eigenvectors of $\hat{\mathbf{X}}^\top \hat{\mathbf{U}} \hat{\mathbf{U}}^\top \hat{\mathbf{X}}$ and $\hat{\mathbf{U}}^\top \hat{\mathbf{X}} \hat{\mathbf{X}}^\top \hat{\mathbf{U}}$ respectively, and a diagonal matrix $\mathbf{\Sigma}$ of the square roots of the eigenvalues.
4. Compute the scale factor $s = \text{trace}(\mathbf{\Sigma}) * \frac{\|\hat{\mathbf{X}}\|}{\|\hat{\mathbf{U}}\|}$.
5. Compute the translation vector $\mathbf{t} = \bar{\mathbf{U}} - s\bar{\mathbf{X}}\mathbf{R}$, where $\bar{\mathbf{X}} = [\bar{x}; \bar{y}]$ and $\bar{\mathbf{U}} = [\bar{u}; \bar{v}]$ are the mean coordinates of \mathbf{X} and \mathbf{U} .
6. Construct the homogeneous transformation matrix,

$$\mathbf{T}_{o,s} = \begin{bmatrix} s\mathbf{R}_{11} & s\mathbf{R}_{21} & t_x \\ s\mathbf{R}_{12} & s\mathbf{R}_{22} & t_y \\ 0 & 0 & 1 \end{bmatrix}. \quad (7.13)$$

When the ‘‘Calculate Transform’’ button is clicked, HistoStitcher[®] executes the above step to solve for $\mathbf{T}_{o,s}$, which is then displayed in the ‘‘Transformation’’ box.

Rotation and Translation: No Scaling

If the images were acquired at exactly the same magnification and stored at the same spatial resolution, no scaling is necessary. In such an event only a transformation comprising pure rotation and translation would be required to stitch the fragments. To set up this configuration, the operator unchecks both the boxes for “Allow Scaling” and “Allow Reflection”. As described in [156], the rotation matrix \mathbf{R} obtained via Procrustes analysis (steps 1-3 Section 7.3.3) is also used when no scaling factor is employed. The translation however is now calculated as $\mathbf{t} = \bar{\mathbf{U}} - \bar{\mathbf{X}}\mathbf{R}$. Finally, the homogeneous transformation matrix is constructed with the new translations,

$$\mathbf{T}_{o,n} = \begin{bmatrix} \mathbf{R}_{11} & \mathbf{R}_{21} & t_x \\ \mathbf{R}_{12} & \mathbf{R}_{22} & t_y \\ 0 & 0 & 1 \end{bmatrix}. \quad (7.14)$$

HistoStitcher[©] calculates $\mathbf{T}_{o,n}$ as described above when the “Calculate Transform” button is clicked.

Enforcing Reflection

For any of the configurations described above (anisotropic scale, isotropic scale, no scale, allowed reflection), HistoStitcher[©] provides a mechanism to reflect (horizontally and/or vertically) the histology fragment in the moving image shown on the right. The toggle buttons shown to the right of the moving image in the HistoStitcher[©] GUI are used to reflect both the image and any previously selected control points. A manually specified horizontal or vertical reflection is combined into the transformation by \mathbf{TM}_x or \mathbf{TM}_y , where,

$$\mathbf{M}_x = \begin{bmatrix} -1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \text{ and } \mathbf{M}_y = \begin{bmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 1 \end{bmatrix},$$

are the X and Y axis homogeneous reflection matrices. If both X and Y reflections are selected, a 180 degree rotation of the moving fragment is applied.

7.3.4 Transformation of Moving Image

On both images, background pixels not representing the tissue mounted on the slide are assigned a value α , a background intensity or color specified by a command line argument. The moving fragment image \mathcal{B} is transformed via \mathbf{T} determined from the control points as described in Section 7.3.3, obtaining $\beta = \mathbf{T} \circ \mathcal{B} = (f_\beta, C^\beta)$, defined by a new coordinate set $C^\beta = \{\mathbf{T}(c), \forall c \in C^B\}$ and intensity function $f_\beta(c)$ for each coordinate $c \in C^\beta$. Cubic spline interpolation on f_B is used to obtain the new image values $f_\beta(c)$ at each $c \in C^\beta$, while a value of α is assigned to coordinates originating outside the original image domain C^B .

7.3.5 Image Reconstruction

The stationary and transformed fragment images \mathcal{A} and β are combined to create the reassembled histology section contained within a larger image $\mathcal{S} = (f_S, C^S)$ by creating a continuous rectangular coordinate frame C^S that is large enough to encompass both C^A and C^β (i.e. $(C^A \cup C^\beta) \subset C^S$). The value $f_S(c)$ of each pixel $c \in C^S$ in \mathcal{S} is defined by,

$$(7.15) \quad f_S(c) = \begin{cases} f_A(c), & \text{if } c \in C^A \text{ and } (c \notin C^\beta \text{ or } f_B(c) = \alpha) \\ f_\beta(c), & \text{if } c \in C^\beta \text{ and } (c \notin C^A \text{ or } f_A(c) = \alpha) \\ \frac{f_A(c) + f_\beta(c)}{2}, & \text{if } (c \in C^A \text{ and } f_A(c) \neq \alpha) \text{ and } (c \in C^\beta \text{ and } f_B(c) \neq \alpha) \\ \alpha, & \text{if } c \notin C^A \text{ and } c \notin C^\beta \end{cases}$$

so that blending of f_A and f_β only occurs at pixels c representing tissue on both \mathcal{A} and β . Note that the blending of f_A and f_β will occur for few pixels in an accurately stitched pair of images, however it is necessary to address overlap that results from jagged features along the edges and other edge irregularities due to slicing and slide preparation. Thus, this area of blending will be minimal in a correct stitch and is only performed for regions where tissue deformation prevents perfect boundary alignment.

To facilitate this blending-with-transparency approach, a simple preprocessing tool

is included in HistoStitcher[©] to identify the boundary of the tissue on the slide and hence the background. Any value of α may be selected to represent the slide background on the stitched image, such as values representing white or black.

7.3.6 Considerations for High Resolution Images

As previously mentioned, operations on high resolution histology images require special handling to avoid exhausting computer system memory. For example, the 10x magnification scan of the single prostate histology quadrant in Figure 7.3(b) contains $30k \times 18k$ pixels, each pixel requiring 3 bytes to encode the RGB values, for a total of 1.6 GBytes of uncompressed data. Operations such as affine transformation on an image this size require excessive computation time and necessitate operating in tiles or chunks of the image when sufficient memory is not available to load the entire image. Therefore, HistoStitcher[©] has two important features to facilitate large images; these are (1) loading of any resolution within the image pyramid stored in an Aperio's SVS or multi-page TIFF file format and (2) down-sampling the two images for both display in the GUI and for previewing the stitch when the "Preview Overlay" button is pressed. By operating on a lower resolution pyramid level while placing control points and previewing the transformation, and by saving the obtained transformation, the full resolution images can be reassembled offline with the saved transformation using a system with more memory or using a routine that operates in blocks or tiles.

7.3.7 Reference Stitching: Manual Alignment of Fragments in Photoshop

In previous studies [5, 28, 112] that utilized reconstructed whole prostate histological sections for correlation with MR images, an expert pathologist performed stitching by manual manipulation of lower resolution quadrants in Photoshop. In this paper manual stitching of histology fragment images in Photoshop is compared to use of HistoStitcher[©]. Manipulation of the images in general purpose image editing software such as Photoshop involves performing any rotations, translations, and stretching (scaling) steps sequentially and thus independently. By contrast, HistoStitcher[©] simultaneously

determines the single transformation that contains the optimal sequence of rotations, scale, and translation. Further, to closely stitch the edges of two histology fragments in a general image editing program, it is also necessary to specify a transparency layer by carefully delineating the boundary of the tissue on the slides.

7.4 Criteria for Stitching Quality and Evaluation by Multiple Expert Consensus

The criteria for assessment of the quality of a reconstructed histology section is defined to reflect how well the resulting pathology image facilitates the disease annotation process. When a pathologist reviews a histopathologic section for prostate cancer, the pathologist first analyzes the tissue architecture at a low power magnification, at which a differential diagnosis is often made, and will zoom into higher power for confirmation. Thus, good stitching quality is most critical for the low power assessment where first impressions and often the diagnosis, which is of crucial importance in prostate cancer, are made. The disease annotation process involves inspecting salient anatomical structures (e.g. ducts, BPH nodules, urethra, capsule), which may span multiple quadrants. Therefore, with a well reconstructed section the assessment at low power is more efficient and accurate because a pathologist will not be required to perform complicated mental image transformations of the individual quadrants in order to follow the salient features across the cuts/boundaries.

In consort with our collaborating pathologists, we have defined stitch quality in terms of (1) continuity of tissue across the stitch boundaries, such that gaps and misalignment of anatomical structures are minimized, and (2) utility of the stitch for annotation of disease extent by a pathologist. A better stitch, as defined by this criterion, will facilitate a pathologist's ability to follow anatomical structures across the boundaries and, given the low to high power assessment approach employed by pathologists, perform annotation of disease extent more accurately and efficiently.

A quantitative 6 point scoring scheme, proposed by our collaborating and co-authoring pathologists, is employed to assess stitching quality. Under the 6 point

scoring scheme, each reconstructed section receives a score from 0 to 6, with up to 2 points awarded for alignment of each of the following three anatomical features:

1. Capsule: continuity across the quadrants
2. Urethra: ability to visualize in its entirety
3. Glands: boundary and structural preservation of glands and other histological instances across the cut

Each anatomical feature receives either 0 points for no alignment, 1 point for partial alignment, or 2 points for full alignment. Therefore, even with some remaining gaps, points are awarded for good alignment of the edges of the anatomical structures.

7.5 Results of HistoStitcher

7.5.1 Experiment 1: Reconstructing High Resolution Whole Prostate Sections from Quadrants

To demonstrate the operation of HistoStitcher[©] and evaluate both the accuracy of the reassembled histology sections and the efficiency of its use, we reconstruct high resolution whole histology sections of two different prostate specimens by stitching digitized slide images of the quadrants which compose each section. In the first prostate specimen, HistoStitcher's[©] operation is demonstrated in the context of the reassembly of a whole histology section from the very high resolution scans of the slides of the four fragments, or quadrants (Figures 7.3(a)-(d)). The slides were digitized using an Aperio slide scanner at 20x apparent magnification for approximately 0.5 microns per pixel. In HistoStitcher[©], image pyramid level 1 containing the 10x resolution images were loaded. The pixel dimensions of these four images shown in Figure 7.3(a)-(d) are approximately: 25k × 20k, 30k × 18k, 29k × 22k and 18k × 17k. To evaluate the efficacy of HistoStitcher[©] in terms of both the quality of the reassembled sections and the efficiency of stitching process, the same fragments were also stitched at a lower resolution (approximately 0.5x magnification or 20 microns per pixel) by an expert pathologist using Photoshop to manipulate the positions of the quadrants, as described in Section




Figure 7.4: (a) Photograph of specimen face prior to cutting into quadrants and mounting on slides. (b) Manually reassembled whole histology section obtained using Photoshop with low resolution images (final stitched image is 2,796 x 2,358 pixels). (c) HistoStitcher[®]-reassembled whole histology section from high resolution images (final stitched image is 48k x 41k pixels). Both stitches in (b) and (c) are validated by the photograph of the section taken prior to cutting. (d)-(g) Close up views of regions I-IV over the seams of the HistoStitcher[®]-reassembled image in (c). (h)-(k) Close up views of the regions nearest to I-IV on the manually-reassembled image in (b). The seams in (d)-(g) contain smaller gaps and better continuity of internal tissue structures compared to the seams in (h)-(k).

7.3.7. To establish a baseline for stitching quality, we also acquired an image of the gross pathology section with a digital camera prior to cutting the section into quadrants, and staining and mounting each quadrant on slides. While this photograph shown in Figure 7.4(a) is not useful for diagnostic purposes, it shows the original configuration of the quadrants prior to cutting.

As described in Section 7.3.1, the stitching of multiple fragments is cumulative and begins by stitching any two originally adjacent fragments. In this case, the quadrants in Figures 7.3(a) and (b) were stitched by placing corresponding pairs of control points along their common edge, automatic transformation of the moving image (the quadrant in Figure 7.3(a)), followed by reassembly of the larger image shown in Figure 7.3(e). Next, the quadrant in Figure 7.3(c) was stitched to the reassembled half in Figure 7.3(e) to generate the image in Figure 7.3(f). Finally, the quadrant in Figure 7.3(d) was stitched to the section in Figure 7.3(f) to generate the pseudo-whole mount section in Figure 7.3(g). The total time required to reassemble the final WMH section from the four quadrants was approximately 6 minutes. This included any time required to (i) load the images, (ii) select control points, (iii) preview the stitch, (iv) refine and/or add landmarks, (v) select appropriate transformation constraints, and (vi) perform the full resolution image transformation.

The reassembled whole-mount section generated using HistoStitcher[®] on the high resolution (10x apparent magnification) histology quadrants is shown in Figure 7.4(c), while the result of the manual stitching using Photoshop (see Section 7.3.7) is shown in Figure 7.4(b) for much lower resolution (0.5x magnification) images. To validate the accuracy of reassembled sections generated manually and using HistoStitcher[®], a block face photograph of the whole section prior to slicing into quadrants is shown in Figure 7.4(a). The HistoStitcher[®] result is shown in Figure 7.4(c) with four rectangular regions along the stitched boundaries highlighted. Close-up views of each of regions I-IV delineated in Figure 7.4(c) are shown in Figures 7.4(d)-(g), illustrating the continuity and smoothness of the stitched boundaries between the quadrants on the HistoStitcher[®]-reassembled section. For the manually stitched section in Figure 7.4(b), close-up views of approximately the same regions I-IV are shown in Figures 7.4(h)-(k), illustrating

significantly larger gaps between the adjacent quadrants and lack of continuity of internal structures compared with the HistoStitcher[©] result. The full 10x reconstruction had dimensions of 48k × 41k pixels (approximately 5.9GB) and was obtained using a 64-bit workstation with 8GB RAM. To perform stitching at resolutions higher than 10x magnification would require either a server machine with greater amounts of memory, or computational methods to process and save the stitched image in blocks, or tiles. Note however that if the parameterized linear transformations for each of the three stitching steps were saved (Save Points button), they may easily be reloaded into HistoStitcher[©] (via Load Points button), which automatically adjusts the translation for the resolution difference.

7.5.2 Experiment 2: Comparison with Manual Stitching of High Resolution Sections

On a second high resolution study, manual stitching was performed at 1x magnification (approximately 7MB per quadrant), and also attempted at approximately 4x apparent magnification (approximately 110MB per quadrant). HistoStitcher[©] was again used to reassemble the quadrants of this second study on 8x resolution images (approximately 420 MB per quadrant). For this study, we compare both the (a) feasibility of manually stitching very high resolution images without a tool such as HistoStitcher[©], and (b) the efficiency of the stitching process in terms of total processing time. In both of these experiments, we use the background mask obtained by HistoStitcher[©] to establish the transparency layer used in the manual stitching process. This speeds up the manual stitching process and provides for a more fair comparison of stitching quality between the two methods.

The results of attempted manual stitching using Photoshop at 4x apparent magnification, and successful stitching at 1x magnification are shown in Figures 7.5(a) and (b), while the HistoStitcher[©] result using 8x magnification images is shown in 7.5(c). The obvious failure of Photoshop-based manual stitching at the full resolution (Figure 7.5(a)) is a result of insufficient memory to accommodate the necessary transparency layers and all three color planes. Without the transparency layer, the quadrants could




Figure 7.5: Comparison of whole histology reassembly via Photoshop vs. HistoStitcher[®]. (a) Inadequate stitching is obtained using Photoshop on 4x resolution images (final image is 13975×13675 pixels), as computer system memory of 2GB was insufficient to introduce the transparency layers required to bring the images closer. (b) Stitching results using Photoshop on low resolution (about 1x magnification) images (final image was 2706×2244 pixels). (c) HistoStitcher[®] result on high resolution (8x magnification) images (final image is about $22k \times 18k$). (d)-(f) Zoomed regions I-III of manually stitched image (b), compared to (g)-(i) same regions of HistoStitcher[®] result (c), demonstrating both more contiguous stitching with minimal gaps and improved continuity of tissue structures across the stitch using HistoStitcher[®].

Stitching Method	Image Size	Time	Result	Output
Manipulation in Photoshop	Medium resolution (4x mag.)	56 min.	Partial failure due to memory requirements	Partially stitched image only
Manipulation in Photoshop	Low resolution (1x mag.)	4 min.	Success with some gaps, inconsistencies	Low-resolution stitched image only
HistoStitcher [©]	High resolution (8x mag.)	6 min.	Success with negligible gaps, overlap	All stitched images <i>and</i> fully-parameterized transformations

Table 7.2: Comparison of HistoStitcher[©] to manipulation via Photoshop for the task of reassembling a whole prostate histological section from four separate slide images of histology quadrants. HistoStitcher[©] is demonstrated to be fast, memory efficient, and capable of producing better quality stitches of very high resolution images.

Section	1	2	3	4	5	6
Photoshop	4 ± 1.0	3.3 ± 1.2	3.3 ± 0.6	4.3 ± 1.2	3.7 ± 0.6	4 ± 1.0
HistoStitcher [©]	5 ± 1.0	4.3 ± 1.2	5.3 ± 0.6	6 ± 0.0	4.3 ± 0.6	6 ± 0.0

Table 7.3: Average and standard deviation of scores for 6 sections reconstructed using both Photoshop and HistoStitcher[©].

not be brought any closer, as clearly shown in the resulting stitch in Figure 7.5(a). Therefore, the manual stitching task was performed again at a lower resolution of 1x apparent magnification. The reassembled section obtained by manipulation of these low resolution quadrants is shown in Figure 7.5(b). The machine used for manual stitching with Photoshop had 2GB of RAM, illustrating that this approach is clear memory inefficient. The high resolution reconstruction obtained using HistoStitcher[©] is shown in Figure 7.5(c). The 8x reconstruction the whole section had dimensions of about 22k x 18k pixels (approximately 1.2GB), and was obtained on a desktop computer with 6GB of RAM, although peak memory usage during stitching was less than 3GB, suggesting better memory efficiency.

A summary of these stitching results is described in Table 7.2, which also lists the operation times required to generate each result. The close-up views of the manually stitched section in Figure 7.5(b) are shown in Figures (d)-(f), while the same regions on the HistoStitcher[©] result in Figure 7.5(c) are shown in Figures 7.5(g)-(i). Continuity

of the gland architecture and tissue morphology is clearly visible across the stitched edges in Figures 7.5(g)-(i). The reconstructed section obtained using HistoStitcher[®] not only has smoother alignment of adjacent boundaries and better continuity of internal structures compared to the manually reconstructed section, but was obtained on the full resolution images in approximately the same amount of time required to perform a manual alignment on low resolution images. Using HistoStitcher[®], the reconstructed sections can be generated faster and/or more accurately than using an improvised approach in general purpose software, as evidenced by the results in Figure 7.5 and Table 7.2. Note in Table 7.2 the excessive time (56 minutes) required to operate on the medium resolution quadrants in Photoshop compared with the much shorter time (6 minutes) to combine the high resolution quadrants via HistoStitcher[®]. Only using drastically down-sized images was it possible to obtain a comparable stitch in a reasonable amount of time.

7.5.3 Experiment 3: Evaluation of Stitching Quality via Multiple Experts

On 6 additional high resolution studies, both HistoStitcher[®] and manual stitching were used to reassemble the quadrants at 4x resolution (approximately 160MB per quadrant). A different computer workstation with 8GB of RAM and an Intel Core 2 Quad CPU was used for manual stitching in this experiment, while the same desktop computer with 6GB of RAM was used to run HistoStitcher. The results of the manual Photoshop stitching of 3 studies are shown in Figures 7.6(a), 7.6(c), and 7.6(e). The results of HistoStitcher[®] for the same 3 studies are shown in Figures 7.6(b), 7.6(d), and 7.6(f). Three independent expert pathologists scored each section using the 6 point scoring system described in Section 7.4. The average scores for the six sections are presented in Table 7.3. Section numbers 6, 1, and 4 correspond to the three rows in Figure 7.6. In addition to the 6 point scoring system, all three experts consistently identified the reconstructions generated by HistoStitcher as the higher quality stitch for each of the 6 sections in this experiment.




Figure 7.6: Three prostate histology sections reconstructed using Photoshop (left column) and HistoStitcher[®] (right column). Stitching quality scores for these sections are given in Table 7.3 under section numbers 6 for (a) and (b), 1 for (c) and (d), and 4 for (e) and (f).

7.6 Concluding Remarks

Reassembly of whole histological sections from smaller tissue fragments is necessary to facilitate improved pathologist annotation, especially when pathologies span multiple fragments, and multimodal image fusion such as registration of WMH with MRI. To achieve accurate reconstruction of sections with variable tissue loss and uneven deformations along the incisions between fragments, an expert-guided interactive process is necessary. In this paper, we address this need for an efficient histology reassembly tool and present HistoStitcher[©], a graphical software package for combining images, that offers a powerful image alignment engine with flexible spatial transformation options. The program was demonstrated for the successful reconstruction of a whole histological section of the prostate from four “quadrants”, requiring only identification of pairs of anatomical landmarks along the boundaries via mouse clicks. Note that the reconstructed whole mount sections obtained by HistoStitcher[©] were generated by a naïve user, yet were obtained with greater accuracy and efficiency than the sections that were reconstructed via Photoshop by an expert pathologist. Using formalized criteria for stitching quality and a 6 point scoring scheme, which assesses the alignment and continuity of anatomical structures important for disease annotation, three independent expert pathologists evaluated the reconstructions of 6 prostate studies. Each of the reconstructed sections generated via HistoStitcher scored higher than the corresponding reconstructions generated by an expert pathologist using Photoshop. Further, the ability of HistoStitcher[©] to operate more efficiently on higher resolution images compared to Photoshop can be valuable in many applications. Although we demonstrated HistoStitcher[©] on prostate histology sections, the program can be applied to other applications such as reassembly of liver biopsy or breast lumpectomy histology fragments. Not only is HistoStitcher[©] efficient and flexible, with multiple options for constrained transformation, but the interface is intuitive enough to be used by relatively inexperienced users with minimal domain expertise.

Chapter 8

Conclusions

We have developed novel techniques for multimodal elastic image registration, and applied them to registration of whole-mount histology while multiprotocol MR images of the prostate. By establishing an accurate alignment of histology, upon which the ground truth for cancer extent and severity may be obtained, and MRI via image registration, we demonstrated the ability to establish a precise spatial extent of cancer on MRI. Our registration techniques were able to overcome the challenges of aligning images from modalities with dissimilar appearance (ex vivo digitized histopathology sections and in vivo multiprotocol MRI) and irregular and non-linear differences in shape. The COFEMI and COLLINARUS techniques utilizes multiple textural features within an information theoretic framework to improve affine and elastic registration, respectively, of multimodal images. We have shown quantitative improvement in registration accuracy over conventional mutual information (MI) similarity measures, with most significant improvements in cases with a image intensity artifacts including noise and MRI bias field inhomogeneity. The MACMI registration method provides a system for utilizing of multi-attribute images (e.g. multiprotocol MRI) simultaneously in the registration procedure. With the increasingly common acquisition of multiple imaging protocols, it is desirable to consider all available images when performing registration. We have demonstrated the use of MACMI in performing the alignment of 150 sets of *in vivo* T2-w MRI, DCE MRI and *ex vivo* whole mount histology images from 25 patient studies. In addition to the clinical prostate datasets, synthetic multiprotocol brain images were used to quantitatively demonstrated the performance improvements of both MACMI and COLLINARUS over conventional MI.

Using our image registration techniques, the spatial extent of cancer was mapped

via affine and non-linear spatial transformations from annotated *ex vivo* whole-mount histology onto both *ex vivo* and *in vivo* multiprotocol MRI datasets, as described in Chapters 2-5. The cancer maps on *ex vivo* MRI were used to train a novel supervised classifier for detection of cancer with accuracy greater than a classifier trained using cancer maps derived from manual delineation by experts, as described in Chapter 3. Preliminary studies by our group [101] investigating the development of a CAD system trained using the registration-derived *in vivo* cancer maps have demonstrated the efficacy of our registration techniques for mapping cancer extent from *ex vivo* histology to *in vivo* multiprotocol MRI. Future work will involve development and evaluation of more sophisticated CAD systems for *in vivo* multiparametric MRI data including diffusion weighted (DW) and spectroscopic imaging protocols (MRS). Thus, future work will also investigate the use of MACMI for inclusion of additional information channels with highly disparate spatial resolutions, such as MRS “metavoxels” that may span tissue elements up to 10 times the width of commonly acquired T2-w voxels.

We have also leveraged our elastic image registration and CAD methodologies to generate intensity-modulated radiation therapy (IMRT) dose plans providing increased dose to suspicious regions of the prostate, as described in Chapter 6. This use of image registration and prostate CAD to generate IMRT treatment plans with dose dose escalation to cancerous tissue demonstrates just one of the ways in which cancer treatment options can be improved using advanced image processing techniques. Future work should investigate the potential improvements to screening and diagnosis procedures via CAD and image-guided biopsy using existing clinical imaging protocols.

Vita

Jonathan Cecil Chappelow

- 2011** Ph.D. in Biomedical Engineering, Rutgers University
- 2002** B.S.E. in Biomedical Engineering, Duke University
- 1998** Graduated from Emmaus High School
-
- 2008-2011** Department of Defense Prostate Cancer Research Program Pre-doctoral Fellow, Department of Biomedical Engineering, Rutgers University
- 2007-2008** New Jersey Commission on Cancer Research Pre-doctoral Fellow, Department of Biomedical Engineering, Rutgers University
- 2005-2006** GAAN Graduate Fellow, Department of Biomedical Engineering, Rutgers University

References

- [1] Jonathan Chappelow, B. Nicholas Bloch, Neil Rofsky, Elizabeth Genega, Robert Lenkinski, William DeWolf, Satish Viswanath, and Anant Madabhushi. COLLI-NARUS: Collection of image-derived non-linear attributes for registration using splines. In *Medical Imaging: Image Processing*, volume 7259 of *Proc. SPIE*, 2009.
- [2] Jonathan Chappelow and Anant Madabhushi. Multi-attribute combined mutual information (MACMI): An image registration framework for leveraging multiple data channels. In *IEEE International Symposium on Biomedical Imaging*, 2010.
- [3] J. Chappelow, A. Madabhushi, et al. A combined feature ensemble based mutual information scheme for robust inter-modal, inter-protocol image registration. In *International Symposium on Biomedical Imaging*. IEEE, 2007.
- [4] Jonathan Chappelow, Anant Madabhushi, Mark Rosen, John Tomaszewski, and Michael Feldman. Multimodal image registration of *ex vivo* 4 Tesla MRI with whole mount histology for prostate cancer detection. In *SPIE*, volume 6512, pages 65121S1–8, 2007.
- [5] J. Chappelow, S. Viswanath, J. Monaco, M. Rosen, J. Tomaszewski, M. Feldman, and A. Madabhushi. Improving supervised classification accuracy using non-rigid multimodal image registration: detecting prostate cancer. In *Medical Imaging: Image Processing*, volume 6915 of *Proc. SPIE*, page 69150V, 2008.
- [6] Jonathan Chappelow, John E Tomaszewski, Michael Feldman, Natalie Shih, and Anant Madabhushi. Histostitcher : An interactive program for accurate and rapid reconstruction of digitized whole histological sections from tissue fragments. *Computerized Medical Imaging and Graphics*, in press, 2011.
- [7] Jonathan Chappelow, B. Nicolas Bloch, Neil Rofsky, Elizabeth Genega, Robert Lenkinski, William DeWolf, and Anant Madabhushi. Elastic registration of multimodal prostate mri and histology via multiattribute combined mutual information. *Medical Physics*, 38(4):2005–2018, 2011.
- [8] D. Skerl, D. Skerl, B. Likar, and F. Pernus. A protocol for evaluation of similarity measures for rigid registration. *IEEE Trans. Med. Imag.*, 25(6):779–791, 2006.
- [9] Baowei Fei, Andrew Wheaton, Zhenghong Lee, Jeffrey L Duerk, and David L Wilson. Automatic MR volume registration and its evaluation for the pelvis and prostate. *Phys Med Biol*, 47(5):823–838, Mar 2002.
- [10] B. Porter et al. Histology and ultrasound fusion of excised prostate tissue using surface registration. In *Ultrasonics Symposium 2001*, pages 1473–1476. IEEE, 2001.

- [11] A. Bharatha, M. Hirose, et al. Evaluation of three-dimensional finite element-based deformable registration of pre- and intraoperative prostate imaging. *Medical Physics*, 28(12):2551–2560, 2001.
- [12] Zhenghong Lee, D. Bruce Sodee, Martin Resnick, and Gregory T MacLennan. Multimodal and 3D imaging of prostate cancer. *Comput Med Imaging Graph*, 29(6):477–486, Sep 2005.
- [13] Tahier Mansoori, Gernot Plank, Rebecca Burton, Jrgen Schneider, Peter Kohl, David Gavaghan, and Vicente Grau. An iterative method for registration of high-resolution cardiac histoanatomical and MRI images. In *ISBI*, pages 572–575. IEEE, 2007.
- [14] Anant Madabhushi, Peng Yang, Mark Rosen, and Susan Weinstein. Distinguishing lesions from posterior acoustic shadowing in breast ultrasound via non-linear dimensionality reduction. *Conf Proc IEEE Eng Med Biol Soc*, 1:3070–3073, 2006.
- [15] Anant Madabhushi, Jayaram K. Udupa, and Andre Souza. Generalized scale: Theory, algorithms, and application to image inhomogeneity correction. *Computer Vision and Image Understanding*, 101(2):100–121, February 2006.
- [16] Anant Madabhushi and Jayaram K Udupa. New methods of mr image intensity standardization via generalized scale. *Med Phys*, 33(9):3426–3434, Sep 2006.
- [17] D. Rueckert, L. I. Sonoda, C. Hayes, D. L. Hill, M. O. Leach, and D. J. Hawkes. Nonrigid registration using free-form deformations: application to breast mr images. *IEEE Trans Med Imaging*, 18(8):712–721, Aug 1999.
- [18] Jennifer M. Hensel, Cynthia Mnard, Peter W.M. Chung, Michael F. Milosevic, Anna Kirilova, Joanne L. Moseley, Masoom A. Haider, and Kristy K. Brock. Development of multiorgan finite element-based prostate deformation model enabling registration of endorectal coil magnetic resonance imaging for radiotherapy planning. *International Journal of Radiation Oncology*Biophysics*, 68(5):1522 – 1528, 2007.
- [19] Lawrence S Taylor, Brian C Porter, Gyongyi Nadasdy, P. Anthony di Sant’Agnese, David Pasternack, Zhe Wu, Raymond B Baggs, Deborah J Rubens, and Kevin J Parker. Three-dimensional registration of prostate images from histology and ultrasound. *Ultrasound Med Biol*, 30(2):161–168, Feb 2004.
- [20] Mark LaSpina and Gabriel P Haas. Update on the diagnosis and management of prostate cancer. *Can J Urol*, 15:3–13; discussion 13, Aug 2008.
- [21] W. J. Catalona, D. S. Smith, T. L. Ratliff, K. M. Dodds, D. E. Coplen, J. J. Yuan, J. A. Petros, and G. L. Andriole. Measurement of prostate-specific antigen in serum as a screening test for prostate cancer. *N Engl J Med*, 324(17):1156–1161, Apr 1991.
- [22] M. D. Rifkin, E. A. Zerhouni, C. A. Gatsonis, L. E. Quint, D. M. Paushter, J. I. Epstein, U. Hamper, P. C. Walsh, and B. J. McNeil. Comparison of magnetic resonance imaging and ultrasonography in staging early prostate cancer. results of a multi-institutional cooperative trial. *N Engl J Med*, 323(10):621–626, Sep 1990.

- [23] Dirk Beyersdorff, Matthias Taupitz, Bjoern Winkelmann, Thomas Fischer, Severin Lenk, Stefan A Loening, and Bernd Hamm. Patients with a history of elevated prostate-specific antigen levels and negative transrectal US-guided quadrant or sextant biopsy results: value of MR imaging. *Radiology*, 224(3):701–706, Sep 2002.
- [24] Hee-Won Kim, David L Buckley, David M Peterson, G. Randy Duensing, Jim Caserta, Jeffrey Fitzsimmons, and Stephen J Blackband. In vivo prostate magnetic resonance imaging and magnetic resonance spectroscopy at 3 tesla using a transceive pelvic phased array coil: preliminary results. *Invest Radiol*, 38(7):443–451, Jul 2003.
- [25] Jacob Sosna, Neil M Rofsky, Sandra M Gaston, William C DeWolf, and Robert E Lenkinski. Determinations of prostate volume at 3-Tesla using an external phased array coil: comparison to pathologic specimens. *Acad Radiol*, 10(8):846–853, Aug 2003.
- [26] Anant Madabhushi, Michael D. Feldman, Dimitris N. Metaxas, Deborah Chute, and John E. Tomaszewski. A novel stochastic combination of 3D texture features for automated segmentation of prostatic adenocarcinoma from high resolution MRI. In *MICCAI (1)*, pages 581–591, 2003.
- [27] A. Madabhushi, M. Feldman, D. Metaxas, D. Chute, and J. Tomaszewski. Optimal feature combination for automated segmentation of prostatic adenocarcinoma from high resolution mri. In *Proc. 25th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, volume 1, pages 614–617, 17–21 Sept. 2003.
- [28] Anant Madabhushi, Michael D Feldman, Dimitris N Metaxas, John Tomaszewski, and Deborah Chute. Automated detection of prostatic adenocarcinoma from high-resolution *ex vivo* MRI. *IEEE Trans. Med. Imag.*, 24:1611–1625, December 2005.
- [29] A. Madabhushi, J. Shi, et al. Graph embedding to improve supervised classification and novel class detection: application to prostate cancer. In James S. Duncan and Guido Gerig, editors, *Medical Image Computing and Computer-Assisted Intervention 2005*, pages 729–737, 2005.
- [30] A Madabhushi, M Rosen, J Tomaszewski, and M Feldman. Eliminating mislabeled training instances: Detecting prostate cancer from high resolution mri. In *MICCAI Workshop on Medical Image Processing in Oncology*, pages 24–31, 2006.
- [31] J. S. Sánchez, R. Barandela, A. I. Marqués, R. Alejo, and J. Badenas. Analysis of new techniques to obtain quality training sets. *Pattern Recogn. Lett.*, 24(7):1015–1022, 2003.
- [32] Carla E. Brodley and Mark A. Friedl. Identifying mislabeled training data. *J. Artif. Intell. Res. (JAIR)*, 11:131–167, 1999.
- [33] S.K. Warfield, K.H. Zou, et al. Simultaneous truth and performance level estimation (staple): an algorithm for the validation of image segmentation. *IEEE Trans. Med. Imag.*, 23(7):903–921, 2004.

- [34] Hong Zhang, Zhen Zhang, and A. W. Partin. Neural network based systems for prostate cancer stage prediction. In *Proc. IEEE-INNS-ENNS International Joint Conference on Neural Networks IJCNN 2000*, volume 3, pages 659–662 vol.3, 2000.
- [35] O. Basset, Z. Sun, J. L. Mestas, and G. Gimenez. Texture analysis of ultrasonic images of the prostate by means of co-occurrence matrices. *Ultrason Imaging*, 15(3):218–237, Jul 1993.
- [36] A. E. Wefer, H. Hricak, D. B. Vigneron, F. V. Coakley, Y. Lu, J. Wefer, U. Mueller-Lisse, P. R. Carroll, and J. Kurhanewicz. Sextant localization of prostate cancer: comparison of sextant biopsy, magnetic resonance imaging and magnetic resonance spectroscopic imaging with step section histology. *J Urol*, 164(2):400–404, Aug 2000.
- [37] Jurgen J Ftterer, Tom W J Scheenen, Henkjan J Huisman, Dennis W J Klomp, Ferdi A van Dorsten, Christina A Hulsbergen van de Kaa, J. Alfred Witjes, Arend Heerschap, and Jelle O Barentsz. Initial experience of 3 tesla endorectal coil magnetic resonance imaging and 1h-spectroscopic imaging of the prostate. *Invest Radiol*, 39(11):671–680, Nov 2004.
- [38] F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens. Multi-modality image registration by maximization of mutual information. *IEEE Trans Med Imaging*, 16(2):187–198, Apr 1997.
- [39] Josien P W Pluim, J. B Antoine Maintz, and Max A Viergever. Mutual-information-based registration of medical images: a survey. *IEEE Trans Med Imag*, 22(8):986–1004, 2003.
- [40] Xiaolei Huang, Nikos Paragios, and Dimitris N Metaxas. Shape registration in implicit spaces using information theory and free form deformations. *IEEE Trans Pattern Anal Mach Intell*, 28(8):1303–1318, Aug 2006.
- [41] S. Lee, G. Wolberg, and S. Y. Shin. Scattered data interpolation with multilevel b-splines. *IEEE Transactions on Visualization and Computer Graphics*, 3(3):228–244, July–Sept. 1997.
- [42] Dinggang Shen and Christos Davatzikos. Hammer: hierarchical attribute matching mechanism for elastic registration. *IEEE Trans Med Imaging*, 21(11):1421–1439, Nov 2002.
- [43] B. Kim, J. L. Boes, K. A. Frey, and C. R. Meyer. Mutual information for automated unwarping of rat brain autoradiographs. *Neuroimage*, 5(1):31–40, Jan 1997.
- [44] Jochen F Krcker, Gerald L LeCarpentier, J. Brian Fowlkes, and Paul L Carson. Rapid elastic image registration for 3-d ultrasound. *IEEE Trans Med Imaging*, 21(11):1384–1394, Nov 2002.
- [45] A. Moskalik, P. Carson, et al. 3D registration of ultrasound with histology in the prostate. In *Ultrasonics Symposium 2001*, pages 1397–1400. IEEE, 2001.

- [46] A. Rangarajan, H. Chui, and J. S. Duncan. Rigid point feature registration using mutual information. *Med Image Anal*, 3(4):425–440, Dec 1999.
- [47] Haili Chui and Anand Rangarajan. A new point matching algorithm for non-rigid registration. *Comput. Vis. Image Underst.*, 89(2-3):114–141, 2003.
- [48] Xiaolei Huang, Yiyong Sun, D. Metaxas, F. Sauer, and Chenyang Xu. Hybrid image registration based on configural matching of scale-invariant salient region features. In *Proc. Conference on Computer Vision and Pattern Recognition Workshop CVPRW '04*, pages 167–167, 2004.
- [49] A. du Bois d’Aische, M. De Craene, S. Haker, N. Weisenfeld, C. Tempany, B. Macq, and S. Wareld. Improved non-rigid registration of prostate mri. In *MICCAI Med Image Comput Comput Assist Interv. MICCAI 2004*, 09 2004.
- [50] B. C. Porter, D. J. Rubens, J. G. Strang, J. Smith, S. Totterman, and K. J. Parker. Three-dimensional registration and fusion of ultrasound and MRI using major vessels as fiducial markers. *IEEE Trans Med Imaging*, 20(4):354–359, Apr 2001.
- [51] Ron Alterovitz, Ken Goldberg, Jean Pouliot, I-Chow Joe Hsu, Yongbok Kim, Susan Moyher Noworolski, and John Kurhanewicz. Registration of mr prostate images with biomechanical modeling and nonlinear parameter estimation. *Med Phys*, 33(2):446–454, Feb 2006.
- [52] Mark Foskey, Brad Davis, Lav Goyal, Sha Chang, Ed Chaney, Nathalie Strehl, Sandrine Tomei, Julian Rosenman, and Sarang Joshi. Large deformation three-dimensional image registration in image-guided radiation therapy. *Phys Med Biol*, 50(24):5869–5892, Dec 2005.
- [53] J. Chappelow, A. Madabhushi, M. Rosen, J. Tomaszewski, and M. Feldman. A combined feature ensemble based mutual information scheme for robust inter-modal, inter-protocol image registration. In *Biomedical Imaging: From Nano to Macro, 2007. ISBI 2007. 4th IEEE International Symposium on*, pages 644–647, 12-15 April 2007.
- [54] Yujun Guo, Radhika Sivaramakrishna, Cheng-Chang Lu, Jasjit S Suri, and Swamy Laxminarayan. Breast image registration techniques: a survey. *Med Biol Eng Comput*, 44(1-2):15–26, Mar 2006.
- [55] W. M. Wells, P. Viola, H. Atsumi, S. Nakajima, and R. Kikinis. Multi-modal volume registration by maximization of mutual information. *Med Image Anal*, 1(1):35–51, Mar 1996.
- [56] Colin Studholme, Derek L. G. Hill, and David J. Hawkes. An overlap invariant entropy measure of 3d medical image alignment. *Pattern Recognition*, 32(1):71–86, 1999.
- [57] J.P.W. Pluim, J.B.A. Maintz, and Max A. Viergever. Image registration by maximization of combined mutual information and gradient information. *IEEE Trans. Med. Imag.*, 19:809–814, August 2000.

- [58] C. Studholme, D.L.G. Hill, and D. J. Hawkes. Incorporating connected region labelling into automatic image registration using mutual information. In *Math. Methods in Biomed. Image Analysis*, volume 3979, pages 23–31, 1996.
- [59] D. Rueckert, M.J. Clarkson, and D. J. Hawkes. Non-rigid registration using higher-order mutual information. In *Medical Imaging: Image Processing*, volume 3979 of *Proc. SPIE*, pages 438–447, 2000.
- [60] J. Liu and J. Tian. Multi-modal medical image registration based on adaptive combination of intensity and gradient field mutual information. In *IEEE EMBS*, volume 1, pages 1429–1433, 2006.
- [61] J.L. Boes and C.R. Meyer. Multi-variate mutual information for registration. In *MICCAI*, volume 1679, pages 606–612, 1999.
- [62] A. Madabhushi and D.N. Metaxas. Combining low-, high-level and empirical domain knowledge for automated segmentation of ultrasonic breast lesions. *IEEE Trans. Med. Imag.*, 22(2):155–169, Feb. 2003.
- [63] Uros Vovk, Franjo Pernus, and Bostjan Likar. Mri intensity inhomogeneity correction by combining intensity and spatial information. *Phys Med Biol*, 49(17):4119–4133, Sep 2004.
- [64] Guorong Wu, Feihu Qi, and Dinggang Shen. Learning-based deformable registration of MR brain images. *IEEE Trans Med Imaging*, 25(9):1145–1157, Sep 2006.
- [65] G. Wu, F. Qi, and D. Shen. Learning best features and deformation statistics for hierarchical registration of MR brain images. *Inf Process Med Imaging*, 20:160–171, 2007.
- [66] G. D. Tourassi, E. D. Frederick, M. K. Markey, and C. E. Floyd. Application of the mutual information criterion for feature selection in computer-aided diagnosis. *Med Phys*, 28(12):2394–2402, Dec 2001.
- [67] Hanchuan Peng, Fuhui Long, and Chris Ding. Feature selection based on mutual information: criteria of max-dependency, max-relevance, and min-redundancy. *IEEE Trans Pattern Anal Mach Intell*, 27(8):1226–1238, Aug 2005.
- [68] C. Grova, P. Jannin, et al. Evaluation of registration of ictal SPECT/MRI data using statistical similarity methods. In *MICCAI*, pages 687–695, 2004.
- [69] J.M. Fitzpatrick, J.B. West, and Jr. Maurer, C.R. Predicting error in rigid-body point-based registration. *IEEE Trans. Med. Imag.*, 17(5):694–702, Oct. 1998.
- [70] R K S Kwan, A C Evans, and G B Pike. Mri simulation-based evaluation of image-processing and classification methods. *IEEE Transactions On Medical Imaging*, 18(11):1085–1097, 1999.
- [71] Darko Skerl, Bostjan Likar, and Franjo Pernus. A protocol for evaluation of similarity measures for non-rigid registration. *Med Image Anal*, 12(1):42–54, Feb 2008.

- [72] C. E. Shannon. A mathematical theory of communication. *Bell Systems Technical Journal*, 27:623–656, 1948.
- [73] W. McGill. Multivariate information transmission. *IRE Professional Group on Information Theory*, 4(4):93–111, September 1954.
- [74] R. W. Yeung. A new outlook on shannon’s information measures. *IEEE Trans. Inf. Theory*, 37(3):466–474, 1991.
- [75] H. Matsuda. Physical nature of higher-order mutual information: Intrinsic correlations and frustration. *Phys. Rev. E*, 62(3):3096–3102, Sep 2000.
- [76] Dana Harry Ballard and Christopher M. Brown. *Computer Vision*. Prentice Hall Professional Technical Reference, 1982.
- [77] S. y. Lu, J.E. Hernandez, et al. Texture segmentation by clustering of gabor feature vectors. In *International Joint Conference on Neural Networks, 1991*, pages 683–688, 1991.
- [78] Robert M. Haralick, K. Shanmugam, and Its’Hak Dinstein. Textural features for image classification. *IEEE Transactions on Systems, Man and Cybernetics*, 3(6):610–621, Nov. 1973.
- [79] Jeffrey C. Lagarias, James A. Reeds, Margaret H. Wright, and Paul E. Wright. Convergence properties of the Nelder–Mead simplex method in low dimensions. *SIAM J. on Optimization*, 9(1):112–147, 1998.
- [80] P. Rogelj, S. Kovacic, et al. Validation of a nonrigid registration algorithm for multimodal data. In *SPIE*, pages 299–307, 2002.
- [81] D.P. Huttenlocher, G.A. Klanderman, and W.J. Rucklidge. Comparing images using the hausdorff distance. *IEEE Trans. Pattern Anal. Mach. Intell.*, 15(9):850–863, Sept. 1993.
- [82] D. L. Collins, A. P. Zijdenbos, V. Kollokian, J. G. Sled, N. J. Kabani, C. J. Holmes, and A. C. Evans. Design and construction of a realistic digital brain phantom. *IEEE Trans Med Imag*, 17(3):463–468, 1998.
- [83] Yosi Keller and Amir Averbuch. Multisensor image registration via implicit similarity. *IEEE Trans Pattern Anal Mach Intell*, 28(5):794–801, May 2006.
- [84] J. C. Mazziotta, A. W. Toga, A. Evans, P. Fox, and J. Lancaster. A probabilistic atlas of the human brain: theory and rationale for its development. the international consortium for brain mapping (icbm). *Neuroimage*, 2(2):89–101, Jun 1995.
- [85] Peter Rogelj, Stanislav Kovačič, and James C. Gee. Point similarity measures for non-rigid registration of multi-modal data. *Comput. Vis. Image Underst.*, 92(1):112–140, 2003.
- [86] E. Oubel, C. Tobon-Gomez, A. O. Hero, and A. F. Frangi. Myocardial motion estimation in tagged MR sequences by using alphaMI-based non rigid registration. *Med Image Comput Comput Assist Interv Int Conf Med Image Comput Comput Assist Interv*, 8(Pt 2):271–278, 2005.

- [87] Huzefa Neemuchwala, Alfred Hero, and Paul Carson. Image matching using alpha-entropy measures and entropic graphs. *Signal Process.*, 85(2):277–296, 2005.
- [88] Fabrice Muhlenbach, Stéphane Lallich, and Djamel A. Zighed. Identifying and handling mislabelled instances. *J. Intell. Inf. Syst.*, 22(1):89–109, 2004.
- [89] David Mattes, David R. Haynor, Hubert Vesselle, Thomas K. Lewellyn, and William Eubank. Nonrigid multimodality image registration. In Milan Sonka and Kenneth M. Hanson, editors, *SPIE*, volume 4322, pages 1609–1620. SPIE, 2001.
- [90] M. Ferrant, A. Nabavi, B. Macq, F. A. Jolesz, R. Kikinis, and S. K. Warfield. Registration of 3-D intraoperative MR images of the brain using a finite-element biomechanical model. *IEEE Trans Med Imaging*, 20(12):1384–1397, Dec 2001.
- [91] C. Studholme, D. L. Hill, and D. J. Hawkes. Automated three-dimensional registration of magnetic resonance and positron emission tomography brain images by multiresolution optimization of voxel similarity measures. *Med Phys*, 24(1):25–35, Jan 1997.
- [92] F. L. Bookstein. Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Trans. Pattern Anal. Mach. Intell.*, 11(6):567–585, 1989.
- [93] J. F. Cardoso. High-order contrasts for independent component analysis. *Neural Comput*, 11(1):157–192, Jan 1999.
- [94] Harald Stgbauer, Alexander Kraskov, Sergey A Astakhov, and Peter Grassberger. Least-dependent-component analysis based on mutual information. *Phys Rev E Stat Nonlin Soft Matter Phys*, 70(6 Pt 2):066123, Dec 2004.
- [95] Seungyong Lee, G. Wolberg, Kyung-Yong Chwa, and Sung Yong Shin. Image metamorphosis with scattered feature constraints. *IEEE Transactions on Visualization and Computer Graphics*, 2(4):337–354, Dec. 1996.
- [96] H. Degani, V. Gusis, et al. Mapping pathophysiological features of breast tumours by MRI at high spatial resolution. *Nature Medicine*, 3(7):780–782, 1997.
- [97] B. Nicolas Bloch, Edna Furman-Haran, Thomas H Helbich, Robert E Lenkinski, Hadassa Degani, Christian Kratzik, Martin Susani, Andrea Haitel, Silvia Jaromi, Long Ngo, and Neil M Rofsky. Prostate cancer: accurate determination of extracapsular extension with high-spatial-resolution dynamic contrast-enhanced and T2-weighted MR imaging—initial results. *Radiology*, 245(1):176–185, Oct 2007.
- [98] John Kurhanewicz, Daniel Vigneron, Peter Carroll, and Fergus Coakley. Multiparametric magnetic resonance imaging in prostate cancer: present and future. *Curr Opin Urol*, 18(1):71–77, Jan 2008.
- [99] Caroline M A Hoeks, Jurgen J Fütterer, Diederik M Somford, Inge M van Oort, Henkjan Huisman, and Jelle O Barentsz. Multiparametric MRI for prostate cancer screening. *Ned Tijdschr Geneesk*, 153:B487, 2009.

- [100] Deanna L Langer, Theodorus H van der Kwast, Andrew J Evans, John Trachtenberg, Brian C Wilson, and Masoom A Haider. Prostate cancer detection with multi-parametric MRI: logistic regression analysis of quantitative T2, diffusion-weighted imaging, and dynamic contrast-enhanced MRI. *J Magn Reson Imaging*, 30(2):327–334, Aug 2009.
- [101] Satish Viswanath, B. Nicolas Bloch, Mark Rosen, Jonathan Chappelow, and Robert Toth. Integrating structural and functional imaging for computer assisted detection of prostate cancer on multi-protocol *in vivo* 3 Tesla MRI. In *Medical Imaging: Image Processing*, volume 7260 of *Proc. SPIE*, 2009.
- [102] Pallavi Tiwari, Mark Rosen, and Anant Madabhushi. A hierarchical spectral clustering and nonlinear dimensionality reduction scheme for detection of prostate cancer from magnetic resonance spectroscopy (MRS). *Med Phys*, 36(9):3927–3939, Sep 2009.
- [103] Masoom A Haider, Theodorus H van der Kwast, Jeff Tanguay, Andrew J Evans, Ali-Tahir Hashmi, Gina Lockwood, and John Trachtenberg. Combined T2-weighted and diffusion-weighted MRI for localization of prostate cancer. *AJR Am J Roentgenol*, 189(2):323–328, Aug 2007.
- [104] Piotr Kozlowski, Silvia D Chang, Edward C Jones, Kenneth W Berean, Henry Chen, and S. Larry Goldenberg. Combined diffusion-weighted and dynamic contrast-enhanced MRI for prostate cancer diagnosis—correlation with biopsy and histopathology. *J Magn Reson Imaging*, 24(1):108–113, Jul 2006.
- [105] Pallavi Tiwari, John Kurhanewicz, Mark Rosen, and Anant Madabhushi. Semi supervised multi kernel (SeSMiK) graph embedding: identifying aggressive prostate cancer via magnetic resonance imaging and spectroscopy. *Med Image Comput Comput Assist Interv*, 13(Pt 3):666–673, 2010.
- [106] C. Bartolozzi, I. Menchi, R. Lencioni, S. Serni, A. Lapini, G. Barbanti, A. Bozza, A. Amorosi, A. Manganelli, and M. Carini. Local staging of prostate carcinoma with endorectal coil MRI: correlation with whole-mount radical prostatectomy specimens. *Eur Radiol*, 6(3):339–345, 1996.
- [107] Pieter C Vos, Thomas Hambroek, Christina A Hulsbergen van de Kaa, Jurgen J Fütterer, Jelle O Barentsz, and Henkjan J Huisman. Computerized analysis of prostate lesions in the peripheral zone using dynamic contrast enhanced MRI. *Med Phys*, 35(3):888–899, Mar 2008.
- [108] Sedat Ozer, Deanna L Langer, Xin Liu, Masoom A Haider, Theodorus H van der Kwast, Andrew J Evans, Yongyi Yang, Miles N Wernick, and Imam S Yetik. Supervised and unsupervised methods for prostate cancer segmentation with multi-spectral MRI. *Med Phys*, 37(4):1873–1883, Apr 2010.
- [109] Xin Liu, D. L. Langer, M. A. Haider, Y. Yang, M. N. Wernick, and I. S. Yetik. Prostate cancer segmentation with simultaneous estimation of markov random field parameters and class. *IEEE Transactions on Medical Imaging*, 28(6):906–915, 2009.

- [110] Pieter C Vos, Thomas Hambroek, Jelle O Barentsz, and Henkjan J Huisman. Computer-assisted analysis of peripheral zone prostate lesions using T2-weighted and dynamic contrast enhanced T1-weighted MRI. *Phys Med Biol*, 55(6):1719–1734, Mar 2010.
- [111] Hyunjin Park, Morand R Piert, Asra Khan, Rajal Shah, Hero Hussain, Javed Siddiqui, Thomas L Chenevert, and Charles R Meyer. Registration methodology for histological sections and *in vivo* imaging of human prostate. *Acad Radiol*, 15(8):1027–1039, Aug 2008.
- [112] Yangming Ou, Dinggang Shen, M. Feldman, J. Tomaszewski, and C. Davatzikos. Non-rigid registration between histological and MR images of the prostate: A joint segmentation and registration framework. In *Computer Vision and Pattern Recognition Workshops*, pages 125–132, Los Alamitos, CA, USA, 2009. IEEE Computer Society.
- [113] B. Nicolas Bloch, Robert E Lenkinski, and Neil M Rofsky. The role of magnetic resonance imaging (MRI) in prostate cancer imaging and staging at 1.5 and 3 Tesla: the Beth Israel Deaconess Medical Center (BIDMC) approach. *Cancer Biomark*, 4(4-5):251–262, 2008.
- [114] John Kurhanewicz, Mark G Swanson, Sarah J Nelson, and Daniel B Vigneron. Combined magnetic resonance imaging and spectroscopic imaging approach to molecular imaging of prostate cancer. *J Magn Reson Imaging*, 16(4):451–463, Oct 2002.
- [115] Sheng Xu, Jochen Kruecker, Baris Turkbey, Neil Glossop, Anurag K. Singh, Peter Choyke, Peter Pinto, and Bradford J. Wood. Real-time MRI-TRUS fusion for guidance of targeted prostate biopsies. *Computer Aided Surgery*, 13(5):255–264, 2008.
- [116] Anurag K. Singh, Jochen Kruecker, Sheng Xu, Neil Glossop, Peter Guion, Karen Ullman, Peter L. Choyke, and Bradford J. Wood. Initial clinical experience with real-time transrectal ultrasonography-magnetic resonance imaging fusion-guided prostate biopsy. *BJU International*, 101:841–845(5), 2008.
- [117] C. C. Parker, A. Damyanovich, T. Haycocks, M. Haider, A. Bayley, and C. N. Catton. Magnetic resonance imaging in the radiation treatment planning of localized prostate cancer using intra-prostatic fiducial markers for computed tomography co-registration. *Radiother Oncol*, 66(2):217–224, Feb 2003.
- [118] Henkjan J Huisman, Jurgen J Ftterer, Emile N J T van Lin, Arjan Welmers, Tom W J Scheenen, Jorn A van Dalen, Andries G Visser, J. A. Witjes, and Jelle O Barentsz. Prostate cancer: precision of integrating functional MR imaging with radiation therapy treatment by using fiducial gold markers. *Radiology*, 236(1):311–317, Jul 2005.
- [119] Sandra Vidakovic, Hans S. Jans, Abe Alexander, and Ron S. Sloboda. Post-implant computed tomographymagnetic resonance prostate image registration using feature line parallelization and normalized mutual information. *J App. Clin. Med. Phys.*, 8(1):21–32, 2007.

- [120] Hyunjin Park, S. Kwee, G.P. Thibault, R. Stack, I.A. Sesterhenn, K. Potter, and C.R. Meyer. Registration methods for histological slides and *ex vivo* MRI of prostate. In *Nuclear Science Symposium Conference Record, 2007. NSS '07. IEEE*, volume 4, pages 3102–3104, Oct. 26 2007–Nov. 3 2007.
- [121] Charles R Meyer, Bradford A Moffat, Kyle K Kuszpit, P. L. Bland, T. L. Chenevert, A. Rehemtulla, and B. D. Ross. A methodology for registration of a histological slide and *in vivo* MRI volume based on optimizing mutual information. *Mol Imaging*, 5(1):16–23, 2006.
- [122] A. Bharatha, M. Hirose, N. Hata, S. K. Warfield, M. Ferrant, K. H. Zou, E. Suarez-Santana, J. Ruiz-Alzola, A. D’Amico, R. A. Cormack, R. Kikinis, F. A. Jolesz, and C. M. Tempany. Evaluation of three-dimensional finite element-based deformable registration of pre- and intraoperative prostate imaging. *Med Phys*, 28(12):2551–2560, Dec 2001.
- [123] Samuel B. Park, Frank C. Rhee, James I. Monroe, and Jason W. Sohn. Spatially weighted mutual information image registration for image guided radiation therapy. *Medical Physics*, 37(9):4590–4601, 2010.
- [124] Yiqiang Zhan, Yangming Ou, Michael Feldman, John Tomaszewski, and Christos Davatzikos. Registering histologic and MR images of prostate for image-based cancer detection. *Academic Radiology*, 14(11):1367–1381, 2007.
- [125] K. K. Bhatia, J. V. Hajnal, B. K. Puri, A. D. Edwards, and D. Rueckert. Consistent groupwise non-rigid registration for atlas construction. In *ISBI*, pages 908–911, 2004.
- [126] C. Studholme. Simultaneous population based image alignment for template free spatial normalisation of brain anatomy. In *WBIR*, pages 81–90, 2003.
- [127] T. F. Cootes, S. Marsland, C.J. Twining, K. Smith, and C.J. Taylor. Groupwise diffeomorphic non-rigid registration for automatic model building. In *ECCV*, pages 316–327, 2004.
- [128] S. Balci, P. Golland, M. Shenton, and W. Wells. Free-form B-spline deformation model for groupwise registration. In *MICCAI*, 10 2007.
- [129] M. Staring, U. A. van der Heide, S. Klein, M. A. Viergever, and J. Pluim. Registration of cervical mri using multifeature mutual information. *IEEE Transactions on Medical Imaging*, 28(9):1412–1421, 2009.
- [130] Gaoyu Xiao, B. Nicolas Bloch, Jonathan Chappelow, Elizabeth M. Genega, Neil M. Rofsky, Robert E. Lenkinski, John Tomaszewski, Michael D. Feldman, Mark Rosen, and Anant Madabhushi. Determining histology-MRI slice correspondences for defining MRI-based disease signatures of prostate cancer. *Comput Med Imaging Graph*, in press, 2011.
- [131] V. Nicolas, D. Beyersdorff, U. Mueller-Lise, W. Pennekamp, and C. Heyer. *MR Imaging of the Abdomen and Pelvis*, chapter The Prostate and Seminal Vesicles, pages 220–240. Thieme, 2002.

- [132] M. L. Schiebler, J. E. Tomaszewski, M. Bezzi, H. M. Pollack, H. Y. Kressel, E. K. Cohen, H. G. Altman, W. B. Geftter, A. J. Wein, and L. Axel. Prostatic carcinoma and benign prostatic hyperplasia: correlation of high-resolution mr and histopathologic findings. *Radiology*, 172(1):131–137, Jul 1989.
- [133] J. F. Platt, R. L. Bree, and R. E. Schwab. The accuracy of CT in the staging of carcinoma of the prostate. *AJR Am J Roentgenol*, 149(2):315–318, Aug 1987.
- [134] C. Rasch, I. Barillot, P. Remeijer, A. Touw, M. van Herk, and J. V. Lebesque. Definition of the prostate in CT and MRI: a multi-observer study. *Int J Radiat Oncol Biol Phys*, 43(1):57–66, Jan 1999.
- [135] Pallavi Tiwari, Mark Rosen, Galen Reed, John Kurhanewicz, and Anant Madabhushi. Spectral embedding based probabilistic boosting tree (sceptre): Classifying high dimensional heterogeneous biomedical data. In *MICCAI '09: Proceedings of the 12th International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 844–851, Berlin, Heidelberg, 2009. Springer-Verlag.
- [136] Barby Pickett, Eric Vigneault, John Kurhanewicz, Lynn Verhey, and Mack Roach. Static field intensity modulation to treat a dominant intra-prostatic lesion to 90 Gy compared to seven field 3-dimensional radiotherapy. *Int J Rad Onc Biol Phys*, 44(4):921–929, 1999.
- [137] M. S. Cohen, R. M. DuBois, and M. M. Zeineh. Rapid and effective correction of RF inhomogeneity for high field magnetic resonance imaging. *Hum Brain Mapp*, 10(4):204–211, Aug 2000.
- [138] Jianbo Shi and J. Malik. Normalized cuts and image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 22(8):888–905, August 2000.
- [139] George Lee, Carlos Rodriguez, and Anant Madabhushi. Investigating the efficacy of nonlinear dimensionality reduction schemes in classifying gene and protein expression studies. *IEEE/ACM Trans Comput Biol Bioinform*, 5(3):368–384, 2008.
- [140] Robert Toth, Pallavi Tiwari, Mark Rosen, Arjun Kalyanpur, Sona Pungavkar, and Anant Madabhushi. A multi-modal prostate segmentation scheme by combining spectral clustering and active shape models. In *SPIE: Medical Imaging*, 2008.
- [141] James P Monaco, John E Tomaszewski, Michael D Feldman, Ian Hagemann, Mehdi Moradi, Parvin Mousavi, Alexander Boag, Chris Davidson, Purang Abolmaesumi, and Anant Madabhushi. High-throughput detection of prostate cancer in histological sections using probabilistic pairwise markov models. *Med Image Anal*, accepted, Apr 2010.
- [142] S Doyle, J Tomaszewski, M Feldman, and A Madabhushi. Hierarchical boosted bayesian ensemble for prostate cancer detection from digitized histopathology. *IEEE Transactions on Biomedical Engineering*, 57(7):accepted, 2010.
- [143] Anant Madabhushi. Digital pathology image analysis: opportunities and challenges. *Imaging in Medicine*, 1(1):7–10, October 2009.

- [144] M.A. Jacobs, J.P. Windham, et al. Registration and warping of magnetic resonance images to histological sections. *Med. Phys.*, 26(8):1568–1578, 1999.
- [145] Jiaya Jia and Chi-Keung Tang. Image stitching using structure deformation. *IEEE Trans Pattern Anal Mach Intell*, 30(4):617–631, Apr 2008.
- [146] Simon T Suen, Edmund Y Lam, and Kenneth K Wong. Photographic stitching with optimized object and color matching based on image derivatives. *Opt Express*, 15(12):7689–7696, Jun 2007.
- [147] Assaf Zomet, Anat Levin, Shmuel Peleg, and Yair Weiss. Seamless image stitching by minimizing false edges. *IEEE Trans Image Process*, 15(4):969–977, Apr 2006.
- [148] Vladan Rankov, Rosalind J. Locke, Richard J. Edens, Paul R. Barber, and Borivoj Vojnovic. An algorithm for image stitching and blending. In Jose-Angel Conchello, Carol J. Cogswell, and Tony Wilson, editors, *SPIE: Three-Dimensional and Multidimensional Microscopy: Image Acquisition and Processing XII*, volume 5701, pages 190–199. SPIE, 2005.
- [149] S. K. Chow, H. Hakozaki, D. L. Price, N. A B MacLean, T. J. Deerinck, J. C. Bouwer, M. E. Martone, S. T. Peltier, and M. H. Ellisman. Automated microscopy system for mosaic acquisition and processing. *J Microsc*, 222(Pt 2):76–84, May 2006.
- [150] Bin Ma, Timo Zimmermann, Manfred Rohde, Simon Winkelbach, Feng He, Werner Lindenmaier, and Kurt E J Dittmar. Use of autostitch for automatic stitching of microscope images. *Micron*, 38(5):492–499, 2007.
- [151] Dirk Steckhan, Tobias Bergen, Thomas Wittenberg, and Stephan Rupp. Efficient large scale image stitching for virtual microscopy. *Conf Proc IEEE Eng Med Biol Soc*, 2008:4019–4023, 2008.
- [152] H. C. da Gama Leitao and J. Stolfi. A multiscale method for the reassembly of two-dimensional fragmented objects. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(9):1239–1251, 2002.
- [153] Efthymia Tsamoura and Ioannis Pitas. Automatic color based reassembly of fragmented images and paintings. *IEEE Transactions on Image Processing*, 19(3):680–690, 2010.
- [154] F. Amigoni, S. Gazzani, and S. Podico. A method for reassembling fragments in image reconstruction. In *Proc. Int. Conf. Image Processing ICIP 2003*, volume 3, pages 581–4, 2003.
- [155] M. S. Sagioglu and A. Ercil. A texture based matching approach for automated assembly of puzzles. In *Proc. 18th Int. Conf. Pattern Recognition ICPR 2006*, volume 3, pages 1036–1041, 2006.
- [156] John C. Gower and Garmt B Dijkstrahuis. *Procrustes Problems*. Oxford University Press, 2004.
- [157] Lloyd N. Trefethen and David Bau. *Numerical linear algebra*. Society for Industrial and Applied Mathematics, 1997.