A 3D ULTRASOUND-BASED TRACKING SYSTEM FOR PROSTATE BIOPSY DISTRIBUTION QUALITY INSURANCE AND GUIDANCE.

Michael Baumann

INSTITUT POLYTECHNIQUE DE GRENOBLE

N° attribué par la bibliothèque
|__|__|__|__|__|__|__|__|__|__|

THESE

pour obtenir le grade de

DOCTEUR DE L'IP GRENOBLE

Spécialité : Imagerie, Vision et Robotique

préparée au laboratoire TIMC/GMCAO
dans le cadre de l'Ecole Doctorale Mathématiques, Science et Technologies de l'Information, Informatique

présentée et soutenue publiquement

par

MICHAEL BAUMANN

le 26 mai 2008

A 3D ULTRASOUND-BASED TRACKING SYSTEM FOR PROSTATE BIOPSY DISTRIBUTION QUALITY INSURANCE AND GUIDANCE.

-

SYSTÈME DE SUIVI BASÉ SUR L'ECHOGRAPHIE 3D POUR L'ASSURANCE DE LA QUALITÉ DE LA DISTRIBUTION DES BIOPSIES DE LA PROSTATE ET LE GUIDAGE DU GESTE.

Directeurs de thèse

Jocelyne Troccaz et Vincent Daanen

Jury

M.	Augustin Lux	, Président
Mme.	Odile Berger	, Rapporteur
M.	David Hawkes	, Rapporteur
Mme.	Jocelyne Troccaz	, Directrice de thèse
M.	Vincent Daanen	, Co-encadrant
M.	Purang Abolmaesumi	, Examinateur
M.	Emmanuel Chartier-Kastler	, Examinateur

To my parents.

Acknowledgements

The last three years represent one of the most interesting and stimulating periods of my life. I am very grateful for the time I could spend with the GMCAO group of the TIMC laboratory, and for all the intellectual, cultural and personal experiences that I encountered there. I particularly appreciated the team spirit in the GMCAO group, which is in my opinion the group's key value to achieve excellence. A lot of gratitude from my side goes to Joce who was a very supportive guide and mentor during both tough and victorious times. Also, the support and friendship of Vince and Tonio were important factors for the success of this thesis. A great 'thanks' goes to all the clinicians with whom I collaborated during this work, in particular to Pierre Mozer who had, together with Joce, the original idea to work on computer-assistance systems for prostate biopsies, but also to Stéphane Bart, Grégoire Chevreau, Jean-Luc Descottes, Jean-Alexandre Long and all the other persons from the urology services at the CHU Grenoble and the Pitié-Salpétrière in Paris who participated in the project. I would also like to thank the 'registration crew', composed of Marek and Sébastien (and me), which formed sort of a think tank for all kind of imaginable and unimaginable registration problems. Thanks also to my interns Kevin, Marc and Nils who carried out the parts of the project for which I would have never had the time to do it on my own. Thanks to my office mates Anne and Céline for their kind support when I had some hard times with this work. Thanks to Philippe, Yohan, Manu and Laurent for their instructive comments and suggestions on this work. Thanks to Guy, Corinne, Olivier, Angélique and Chantal for keeping the lab running smoothly. Thanks to Fabien, Joe, Anna, Sandrine, Adriana, Christophe, Yannick, Claudio, Juan-Carlos, Maud, Nabil, Jonathan, François, Olivier, Ziad, Jocelin, Emilie, Laure, Oriane, Ariane, Louise, Nikolai, Pilar, Pascal, Anne-Thérèse, Sandrine, Bruno, Chantal and everybody else who I forgot to mention (forgive me!) for the nice time that we spent together.

Thanks also to my roommates Nadji and Jan for the great time that we spent together. Big hugs for my brother Christian, his wife Erika and my mother who helped me out a lot during the last three years, and in particular during the preparation of the vitae of this work. Finally thanks to my father who encouraged me to start this thesis.

Contents

S-1 Symbol Table

$[\cdot]$	functional parameter
$(\cdot)$	function parameter
$\mathcal{F}$	functional
F	function
$\wedge$	logical "and"
$\circ$	mathematical composition
$\Omega[I] \subset \mathbb{R}^3$	non-masked part of the image I
$\Omega = \Omega[R, T; \varphi]$	image overlap $\Omega[R] \cap \varphi(\Omega[T])$
$\Omega^{\mathbb{N}} = \Omega \cap \mathbb{N}^3$	discrete image overlap
$\Psi \subset \mathbb{R}^3$	non-rigid registration domain
$\Psi^{\mathbb{N}} = \Psi \cap \mathbb{N}^3$	discrete non-rigid registration domain
$I : \Omega[I] \to \mathbb{R}$	monochromatic image
$R : \Omega[R] \to \mathbb{R}$	reference image
$T : \Omega[T] \to \mathbb{R}$	template image
$\varphi : \mathbb{R}^3 \to \mathbb{R}^3$	image transformation
$\hat{\varphi}$	physical solution of the registration problem
φ^*	numerical solution of the registration problem
$u : \mathbb{R}^3 \to \mathbb{R}^3$	displacement vector field, $\varphi(x) = x + u(x)$
$\hat{u}$	physical solution of the registration problem (displacements)
u^*	numerical solution of the registration problem (displacements)
$\Theta \subset \mathbb{R}^6$	rigid transformation parameter domain
$\vartheta \in \Theta$	transformation parameter
$\mathcal{D}[R, T; \varphi]$	image distance measure
$\mathcal{R}[u]$	regularization term
$M : \mathbb{R}^3 \to \mathbb{R}^3$	rotation matrix
o	translation vector
$K_i : \mathbb{R}^3 \to [0, 1]$	restriction kernel
Pr_{Surf}	prostate surface
FP_{Rect}	rectal probe fix-point
O_{US}	ultrasound transducer origin
C_{Pro}	prostate center (bounding box)
$\alpha_i \in \mathbb{R}^+$	scaling parameter

Chapitre F

Synthèse (Français)

F-1 Introduction

F-1.1 Contexte

Cette thèse présente le principe de la mise en oeuvre et l'évaluation d'un système permettant le suivi des bougés et déformations de la prostate lors de biopsies. L'approche repose sur le recalage rigide et élastique de volumes échographiques 3D acquis avant et pendant les biopsies.

F-1.2 Biopsies de la Prostate

La Prostate

La prostate est une glande du système reproductif masculin, qui est située en dessous de la vessie et en face du rectum. Elle entoure l'urèthre à l'endroit de sa connection avec la vessie. La fonction principale de la prostate est la production du fluide séminal. Petite à la naissance et pendant l'enfance, la prostate pousse considérablement à partir de la puberté pour atteindre un diamètre moyen de 2 cm environ à l'âge de 25 ans. La prostate recommence à croitre lentement à partir de l'âge d'environ 50 ans. Cette deuxième étape de croissance s'avère souvent problématique, car elle est fréquemment accompagnée de pathologies telle que l'hyperplasie bénigne et le cancer de la prostate.

Cancer

En Europe et aux Etats-Unis, le cancer de la prostate est le cancer le plus fréquent chez l'homme et la deuxième cause de décès liés à un cancer. Plusieurs centaines milliers de nouveaux cas sont diagnostiqués dans les deux continents chaque année, et environ 100.000 décès lui sont attribués. Plus de trois quarts des cancers de la prostate sont dépistés chez des hommes âgés de plus de 65 ans, et on peut observer cette pathologie seulement très rarement chez des sujets plus jeunes que 40 ans. Comme tout type de cancer, les tumeurs de la prostate peuvent être bénins ou malins. Seulement les tumeurs malignes se transforment en un cancer de la prostate. Les tumeurs malignes peuvent produire des métastases si des cellules cancéreuses se séparent de la tumeur et entrent dans le circuit sanguin. Le cancer de la prostate est très souvent multifocal à cause de la très forte vascularisation du tissu prostatique.

Dépistage du Cancer

A l'heure actuelle il y a trois outils de dépistage du cancer de la prostate. La première est le toucher rectal qui vise à détecter des foyers cancéreux par palpation.

Cette méthode permet de détecter uniquement des cancers dans des stades avancés et qui sont situés près de la paroi rectale. La sensibilité est donc faible et varie fortement en fonction du clinicien et de son expérience. En outre, le toucher rectal est peu spécifique.

La deuxième méthode de détection est biologique et consiste en l'identification du taux de l'antigène spécifique de la prostate (PSA) dans le sang. Un taux de PSA élevé est un indicateur d'une pathologie de la prostate, mais il n'est pas spécifique du cancer de la prostate. Le dépistage du cancer de la prostate par le taux de PSA fait preuve d'une meilleure sensibilité que le toucher rectal, mais il est également peu spécifique.

Afin de valider l'hypothèse du cancer de la prostate, le clinicien est donc obligé d'avoir recours à une méthode plus invasive, les biopsies de la prostate. Les échantillons prélevés lors d'une série de biopsie sont soumis à un examen histologique qui permet d'atteindre une spécificité élevée. Malheureusement, la sensibilité des biopsies de la prostate est relativement faible et il y a donc un risque élevé de faux négatifs.

Biopsies de la Prostate

Les biopsies de la prostate sont prélevées soit par voie transrectale, soit par voie transpérinéale. Dans les deux cas, le geste s'effectue sous contrôle échographique. Les images échographiques sont acquises avec une sonde endorectale. Dans le cadre des biopsies transrectales, un guide d'aiguille qui est rigidement fixé sur la sonde échographique permet d'aligner la trajectoire de l'aiguille avec le champ de vue échographique. La trajectoire est affichée sur l'image échographique de contrôle, ce qui permet au clinicien de choisir les tissus qu'il désire échantillonner. Il est important de noter que la sonde échographique joue donc un double rôle, car elle est non seulement utilisée pour obtenir une image de la glande, mais également pour placer l'aiguille. L'inconvénient de cette approche est qu'il est difficile pour le clinicien de placer les aiguilles de biopsie avec une bonne précision, dû aux faits que l'échographie fournit très peu d'information structurelle à l'intérieur de la glande et que l'image est seulement bi-dimensionnelle. En particulier, le cancer n'est pas visible dans les images échographiques dans la plupart des cas. Il s'y ajoute la problématique des bougés de la prostate pendant l'intervention causés par les mouvements non seulement de la sonde, mais également du patient qui n'est en règle générale pas sous anésthésie générale.

Les biopsies par voie transpérinéales sont plus invasives et requièrent une anésthésie générale. Pour cette raison seulement peu de cliniciens prélèvent des biopsies de la prostate par cette voie d'accès, malgré le fait que ce procédé permet d'obtenir une meilleure précision de placement. En conséquence nous allons nous focaliser sur les biopsies transrectales dans le cadre de ce travail.

Bougés de la Prostate

On peut distinguer plusieurs types de bougés de la prostate pendant le prélèvement de biopsies de la prostate. Le premier type est lié aux bougés du patient qui sont souvent provoqués par une douleur aigüe lors de la ponction. Il se manifeste généralement par une contraction des muscles du bassin et provoque des mouvements et déformations de la glande par rapport à des tissus environants, mais aussi un mouvement par rapport à la salle opératoire.

Le deuxième type de bougés est induit par les bougés de la sonde transrectale lors du placement de l'aiguille, dû aux contraintes que la tête de la sonde doit être en contact avec la paroi rectale devant la prostate et qu'une légère pression doit être appliquée pour qu'une image échographique de la prostate puisse être acquise. Lors du placement de l'aiguille, la prostate peut donc bouger et se déformer avec les bougés de la sonde.

Le troisième type de bougés est causé par les organes environnants, notamment les organes respiratoires et digestifs. Ce type de mouvements posant des problèmes pour certains traitements du cancer de la prostate, il est d'une importance moindre pour les gestes diagnostiques.

Sur une plus longue période, les tissus prostatiques peuvent finalement bouger à cause de changements internes, notamment suite à la croissance d'un cancer ou la décroissance de la glande à cause d'un traitement hormonal.

Objectif : Suivi des Tissus Prostatiques

L'objectif scientifique de ce travail est de concevoir un système permettant de suivre les bougés des tissus prostatiques pendant l'intervention et entre différentes séances de biopsies. Un tel système permettrait d'atteindre plusieurs objectifs cliniques : tout d'abord il sera possible de mettre en correspondance un planning d'intervention pré-opératoire avec les images de contrôle per-opératoires. Ensuite, le système permettrait la mise en correspondance des positions des échantillons prélevés par rapport à un référentiel anatomique commun pour des analyses de qualité per- et post-opératoires. Finalement, une visualisation spatiale précise des résultats de l'analyse histologique des tissus serait réalisable et pourra conduire à des traitements focalisés.

Plusieurs approches de suivi ont été proposées dans la littérature pour atteindre certains de ces objectifs cliniques. La première classe d'approches consiste à effectuer les biopsies sous guidage IRM afin de pouvoir viser des lésions suspectes visibles sous IRM. Ces approches sont actuellement encore relativement peu robustes dû au fait que l'acquisition d'images de contrôle IRM de qualité acceptable ne peut se faire en temps réel avec la technologie actuelle. Le patient peut donc avoir bougé entre le moment d'acquisition de l'image et la ponction. En outre, le

coût de l'imagerie IRM étant très élevé par rapport à l'échographie, il est peu probable que cette approche devienne un standard clinique pour une première série de biopsies.

D'autres approches sont basées sur la mise en correspondance des images échographiques de contrôle avec une volume échographique qui sert de référence anatomique. Ce volume de référence est typiquement acquis quelques minutes avant l'intervention. Les premières approches se limitaient à établir la correspondance uniquement avec un système de suivi stéréotaxique, magnétique ou mécanique par rapport à un référentiel non anatomique, typiquement défini par le système de suivi. Ces approches n'étaient pas capable d'identifier et de compenser les bougés de l'organe par rapport à ce référentiel pendant l'intervention.

Pour cette raison des approches plus récentes rajoutent à la localisation par rapport à ce référentiel externe une étape de mise en correspondance avec des techniques de traitement d'image. Cette étape consiste essentiellement en l'optimisation d'une mesure de similarité entre la partie du volume de référence correspondant à l'image interventionelle pour une transformation donnée, et l'image interventionelle. Les techniques d'optimisation utilisées sont locales et requièrent que la transformation fournie par le système de suivi externe soit suffissamment proche de la solution physique, sinon le recalage par comparaison d'images fournit une transformation correspondant potentiellement à un faux minimum local. Or, cette condition n'est souvent plus satisfaite si le patient a bougé son bassin pendant l'intervention, ce qui contraint le clinicien à réacquérir une nouvelle image de référence. Les biopsies déjà prélevées ne pourront pas être mises en correspondence avec les biopsies suivantes.

F-1.3 Suivi par Recalage d'Images

Ce travail se situe dans le cadre des systèmes de suivi échographiques. Une approche originale sans utilisation de systèmes de localisation externe et purement basée sur des techniques de recalage d'images par minimisation d'une mesure de similarité est mise en oeuvre. Dans le cadre de ce travail, l'échographie 3D sera choisie en tant que modalité d'imagerie, afin d'avoir suffisamment d'information spatiale pour réaliser un recalage robuste et précis. En particulier, cette modalité d'imagerie permet d'estimer et de compenser des déformations de la prostate. L'échographie 3D est une technologie récente qui est très probablement déstinée à remplacer les échographes 2D dans un avenir proche.

La difficulté de cette approche consiste en l'estimation efficace des transformations par comparaison d'images afin d'obtenir un suivi temps-réel. Le temps-réel est nécessaire pour réaliser un système de guidage du geste vers des cibles non visibles sur l'image échographique, notamment des lésions suspectes identifiées sur des images IRM, des cibles provenant d'atlas statistiques de la répartition des

tumeurs prostatiques ou des zones échantillonées lors d'une série précédente de biopsies.

F-1.4 Décomposition du Problème

L'estimation de la transformation entre l'image de référence anatomique et une image interventionnelle est d'autant plus complexe et donc couteuse au niveau du temps de calcul que le nombre de degrés de liberté du modèle de transformation augmente. Une stratégie de décomposition de l'estimation en plusieurs étapes est donc utilisé pour affiner le recalage au fur et à mesure.

Trois niveaux de finesse sont mis en oeuvre : tout d'abord, une estimation grossière de la transformation est effectuée en utilisant un modèle cinématique des déplacements de la sonde dans le rectum. Cette première estimation est mise en oeuvre avec des techniques d'optimisation globale, nous permettant d'être indépendant d'un système de localisation externe. Le résultat de cette analyse globale est par la suite affiné par une étape d'optimisation locale d'une mesure de similarité appliquée sur un modèle de transformation rigide. Finalement, le résultat de l'optimisation rigide est utilisé en tant que point de départ pour une étape d'estimation des déformations élastiques.

L'étape d'optimisation sur un espace de transformation défini par un modèle cinématique et l'étape d'optimisation sur l'espace de transformations rigide sont techniquement très similaires et sont discutées dans le chapitre "Recalage Rigide". L'étape d'estimation de la transformation sur un espace de transformation élastique requièrt une approche différente, à cause du grand nombre de degrés de liberté de l'espace de transformation. Nous avons opté pour une approche d'optimisation non-paramétrique et variationelle qui est discutée dans le chapitre "Recalage Elastique".

F-2 Recalage Rigide

F-2.1 Introduction

Dans ce chapitre une méthode de recalage d'images paramétrique est proposée. Elle permet de trouver rapidement une estimation initiale de la transformation rigide grâce à un modèle cinématique qui modélise les contraintes des déplacements de la sonde lors d'une biopsie de la prostate. L'espace réduit correspond à un sous-ensemble de l'espace des transformations rigides. L'objectif de la recherche sur ce sous-ensemble est de trouver une transformation suffisamment proche de la transformation rigide physique que nous désirons estimer, c'est-à-dire une transformation qui permet de converger vers la transformation physique si des techniques

classiques d'optimisation locale telles que les méthodes de Newton, de descente de gradient ou de Powell-Brent sont appliquées.

Afin d'obtenir une meilleure résolution spatiale du problème, l'approche proposée se base sur l'imagerie échographique 3D. Ceci robustifie les étapes de recalage rigides et rend possible le recalage élastique.

F-2.2 Formulation Paramétrique du Problème d'Optimisation

Le problème de recalage d'image peut être formulé en tant que problème d'optimisation d'une fonctionnelle de coût $\mathcal{D}$ qui dépend d'une image de référence R, d'une image flottante T et d'un modèle de transformation φ. Dans sa formulation générale, nous obtenons donc

$$\varphi^* = \arg\min_{\varphi} \mathcal{D}[R, T \circ \varphi]. \tag{F.1}$$

Pour les problèmes avec peu de degrés de liberté il est préférable d'utiliser une formulation paramétrique du problème d'optimisation. Le modèle de transformation φ dépend donc d'un ensemble de paramètres ϑ. Nous cherchons les paramètres ϑ^* qui minimisent $\mathcal{D}$.

F-2.3 Approche Multirésolution

Afin de réduire le temps de calcul du processus d'optimisation, la majeure partie de l'algorithme est effectuée sur des niveaux de résolution grossières. Pour réaliser ceci avec une robustesse améliorée, des techniques de minimisation de perte d'information lors du calcul des images de résolution grossière sont mises en oeuvre. Une perte d'information peut arriver à cause de la forme particulière du champs de vue de l'échographe : pour calculer un voxel du niveau de résolution $n+1$, huit voxels du niveau de résolution n sont nécessaires. Or, sur les bords du masque, le nombre de voxels disponible est en général inférieur à huit.

Au lieu de créer un voxel masqué pour cette configuration nous utilisons la moyenne des voxel disponibles si le nombre de voxels disponibles au niveau n est supérieur ou égal à la moitié des voxels nécessaires. Le voxel au niveau $n+1$ est masqué seulement si moins de la moitié des voxels sont disponibles. La technique décrite est également appliquée pour tout type de filtre tel que le calcul du gradient d'images ou l'interpolation trilinéaire.

F-2.4 Fonctionnelle de coût

La fonctionnelle de coût doit être capable de fournir une mesure statistiquement robuste de la similarité de deux images échographiques, en tenant compte des variations des caractéristiques d'images qui sont typiques pour les images échographiques. On peut distinguer les mesures de coût définies sur des éléments géométriques identifiables dans l'image ou les mesures basées sur une comparaison directe des intensités des voxels qui correspondent dans les deux images pour une transformation donnée. Le segmentation automatique de structures géometriques dans des images échographiques étant difficile à mettre en oeuvre, le deuxième type de mesure de similarité a été choisi dans le cadre de ce travail.

Un ensemble de mesures basées sur la comparaison d'intensités, ces mesures étant souvent également dénominées iconiques, a été évalué sur des images échographiques de la prostate. Les mesures évaluées peuvent être associées à trois classes, la prémière comprenant les mesures qui supposent un lien fonctionnel entre les intensités dans les deux images. Il s'agit notamment des mesures très classiques de la somme des différences absolues (SAD), des sommes des différences au carré (SSD), du coéfficient de corrélation (CC) et du rapport de corrélation (CR). Ces mesures représentent des estimateurs de la probabilité maximale pour du bruit Gaussien. Une deuxième classe de mesures fait également l'hypothèse d'un lien fonctionnel entre les intensités, mais utilise un modèle de bruit de Rayleigh, qui est spécifique aux images échographiques. Le représentant de cette classe qui a été évalué dans le cadre de ce travail est la mesure de Cohen Dienstein 2 (CD2), qui a l'avantage de considérer le fait que les intensités des images échographiques sont compressées avec un logarithme binaire par le dispositif échographique. Finalement, la troisième classe de mesures suppose uniquement un lien statistique entre les intensités des images, la mesure étant donc non-paramétrique, done le représentant le plus utilisé est sans dout l'information mutuelle normalisée (NMI).

L'évaluation montre que la mesure basée sur un modèle de bruit Rayleigh n'apporte pas une amélioration de la robustesse du processus d'optimisation, mais augmente considérablement la complexité du calcul numérique. Les mesures SAD et SSD font l'hypothèse que le lien fonctionnel des intensités entre les deux images correspond à l'identité, ce qui s'est avéré être une hypothèse trop simpliste. Ceci est dû aux variations des intensités des images en fonction de la pression appliquée à la sonde, de l'orientation et la position du champ de vue de la sonde et de la possibilité que l'opérateur peut changer le gain de la sonde entre les acquisitions. De l'autre côté, les mesures qui permettent des relations plus complexes telles que le rapport de corrélation et l'information mutuelle normalisée requièrent un grand nombre de voxels pour obtenir des résultats qui sont statistiquement significatives. Ceci pose des problèmes pour des approches multi-résolution, ou le nombre total de voxels peut être très faible à des niveaux de résolution grossiers. Le rapport de corrélation,

faisant l'hypothèse d'un lien linéaire entre les intensités des images, s'avère être le meilleur compromis entre généralité fonctionnelle et robustesse statistique pour le problème étudié. Pour robustifier la mesure, le CC est non seulement appliquée sur l'image brute, mais également sur l'image contenant la norme du gradient.

F-2.5 Images Panorama

En règle générale, le champs de vue pyramidal de la sonde échographique 3D ne peut couvrir l'intégralité de la glande. Pour cette raison, une image panoramique est calculée à partir d'un ensemble d'acquisitions de différentes parties de la prostate. Ceci permet d'avoir une image intégrale de l'organe, limitant ainsi les échecs ultérieurs de recalage dus au fait que les deux images contiennent seulement partiellement la prostate.

F-2.6 Filtrage des Déformations

Un recalage guidé par un modèle de transformation rigide recquiert que la fonctionnelle de coût soit aussi invariante que possible vis à vis des déformations présentes dans l'image. Il est donc nécessaire que les déformations soient locales, sinon il est impossible de décomposer le problème. Heureusement, les principales déformations de la prostate sont causées par la tête de la sonde lors du positionnement de l'aiguille. En plus, ces déformations sont relativement locales et se limitent à un voisinage autour de la tête de la sonde. Il est donc possible d'augmenter l'invariance du recalage rigide vis à vis des déformations en masquant la zone proche de la tête où les déformations se manifestent le plus. En pratique, tous les voxels plus proches qu'un seuil expérimentalement choisi sont éliminés.

F-2.7 Filtrage de la Vessie

La vessie, située à côté de la prostate, apparaît en tant que zone noire dans les images échographiques, du au fait qu'elle contient essentiellement de l'eau qui transporte des ondes sonores pratiquement sans atténuation et sans réfraction ou réflexion. En conséquence, la différence d'intensité entre la vessie et d'autres tissus dans l'image est très prononcé, ce qui fait que le coéfficient de corrélation pénalise fortement un mauvais alignement de la vessie. En même temps, il n'y a aucune structure visible à l'intérieur de la vessie et la vessie est relativement symmétrique. Ceci fait que la fonctionnelle de coût n'est que peu discriminative pour la détermination de la rotation. En d'autres mots, elle est caractérisée par de faux minima très prononcés si la vessie est présente dans les deux images. Pour cette raison nous éliminons la vessie au moins dans une des images que nous visons à aligner.

F-2.8 Modèle Cinématique

Toutes les mesures de distance d'images connues sont seulement localement unimodales autour de la solution physique du problème. Pour pouvoir appliquer des méthodes d'optimisation locales il est donc nécessaire de trouver un point à l'intérieur de la region unimodale correspondant au minimum global recherché[1]. Pour trouver un point à l'intérieur de cette region il faut avoir recours à des méthodes d'optimisation globales. Or, l'espace de transformation rigide est bien trop large pour effectuer une optimisation globale avec un temps de calcul raisonable pour un système de tracking.

Il convient donc d'effectuer la recherche globale dans un sous-ensemble des transformations rigides avec une faible dimensionalité qui regroupe des transformations ayant une forte probabilité de se manifester pendant une intervention. Dans l'approche présentée, ce sous-ensemble est déterminé à l'aide d'un modèle cinématique de la sonde qui rend compte des contraintes physiques exercées sur la sonde par le rectum et par le processus de création d'image.

Le point de départ sont les trois observations suivantes : Premièrement, les muscles du sphincter anal définissent une zone de pivot pour la sonde qui peut être approximée par un point fixe. Deuxièmement, pour la création de l'image il est indispensable que la tête de la sonde soit en contact avec la paroi rectale en face de la prostate, sinon la prostate ne peut être visible dans l'image. Finalement, les rotations les plus importantes et les plus difficiles à estimer se font autour de l'axe longitudinal de la sonde.

Le modèle cinématique modélise ces observations à l'aide d'une approximation grossière de la surface de la prostate qui peut être décrite avec un modèle bi-dimensionnel et à l'aide d'un point qui approxime le sphincter anal. Le modèle admet par la suite uniquement des transformations pour lesquelles l'axe longitudinale de la sonde passe par le point fixe et pour lesquelles le centre du transducteur se situe sur la surface. Cet espace peut être décrit avec trois paramètres, deux pour choisir le point sur la surface de la prostate, et un qui modélise la rotation de la sonde autour de l'axe longitudinale. Il est donc possible d'effectuer la recherche globale sur un espace tri-dimensionnel qui est sufissamment petit pour mettre en oeuvre une recherche exhaustive.

F-2.9 Expérimentations et Résultats

Le taux de succès du recalage rigide a été évalué sur des données recueillies sur 47 patients, consistant en 785 recalages sur un ensemble de 926 volumes. Les résultats des recalages ont été analysés visuellement par des experts et classifiés

[1]Notez qu'on fait ici l'hypothèse que le minimum global correspond à la transformation physique recherchée, ce qui n'est pas forcément vérifié.

en "succès" et "échec". 27 recalages ont été classifiés en tant qu'échecs, ce qui résulte en un taux de succès de 96.5%. La précision et la reproductibilité ont été testées avec un jeu de test issues de prélèvements sur 14 patients. La précision a été mesurée à partir de segmentations de microstructures dans l'image (notamment des calcifications). La distance RMS de ces structures a été évaluée à 1.41 mm avec un maximum de 3.84 mm après recalage. Le temps de calcul moyen était de 6.5s.

F-3 Recalage Elastique

F-3.1 Introduction

L'algorithme de recalage rigide aligne essentiellement la partie de la capsule de la prostate visible dans les deux images qui est le plus éloigné de la tête de la sonde. A proximité de la sonde on peut observer des déformations plus ou moins importantes qui ne peuvent être corrigées par un recalage rigide. Pour corriger ces déformations, une troisième étape de recalage est mise en oeuvre, en partant du résultat du recalage rigide.

F-3.2 Formulation du Problème d'Optimisation

Pour l'estimation des déformations nous utilisons des champs de déformation denses, c'est-à-dire des champs de vecteurs $u : \mathbb{R}^3 \rightarrow \mathbb{R}^3$. Le problème de minimisation est formulé de manière suivante

$$\begin{aligned} u^* &= \underset{u}{\operatorname{argmin}} \left(\mathcal{R}[u] + \mathcal{D}_{SSD}[R, T; u] + \mathcal{D}_{Probe}[R, T; u] + \mathcal{D}_{Cons}[u, v]\right) \\ v^* &= \underset{v}{\operatorname{argmin}} \left(\mathcal{R}[v] + \mathcal{D}_{SSD}[T, R; v] + \mathcal{D}_{Probe}[T, R; v] + \mathcal{D}_{Cons}[v, u]\right), \quad \text{(F.2)} \end{aligned}$$

où u correspond à la transformation de l'image R vers l'image T, v à la transformation de T vers R, $\mathcal{R}$ correspond à un régularisateur, $\mathcal{D}_{SSD}$ correspond à une fonctionnelle de distance entre images, en l'occurrence basés sur le SDD, $\mathcal{D}_{Probe}$ à une fonctionnelle dérivée d'un modèle d'insertion de sonde, et $\mathcal{D}_{Cons}$ à une contrainte qui pénalise les inconsistences avec l'inverse par le terme $||u \circ v - Id||$. L'optimisation de ce problème est effectuée de manière itérative et alternée.

F-3.3 Regularisation

L'estimation des déformations basée sur la comparaison d'images est un problème mal posé car on cherche à trouver une projection spatiale des voxel qui

minimise la différence entre les intensités des voxels correspondants des deux images après projection. Or, il y a en règle générale une multitude de solutions à ce problème. Il faut donc injecter des contraintes supplémentaires dans le processus d'optimisation afin de réduire le nombre de solutions non-physiques. Un régularisateur très approprié en présence de déformations surtout élastiques est le potentiel élastique linéaire.

Le problème d'optimisation est résolu avec une approche variationnelle, en dérivant les équations d'Euler-Lagrange du problème à partir de sa dérivée de Gâteaux. Le système d'équations de dérivées partielles résultant est résolu à l'aide de l'algorithme itératif par relaxation de Gauss-Seidel, en le combinant avec la stratégie multirésolution "Full Multigrid". L'algorithme d'optimisation final correspond à un algorithme de descente de gradient sur le problème.

F-3.4 Forces issues de la Comparaison d'Images

Il est nécessaire de calculer les forces dérivées des différences entre les images localement afin de pouvoir estimer les déformations locales. La mesure utilisée pour le recalage rigide, le coefficient de corrélation, doit être calculé sur un certain nombre de voxels pour obtenir un résultat qui soit statistiquement significatif. Ceci limite son utilisabilité dans le cadre d'une approche multirésolution, car il doit être calculé sur une grande partie de l'image pour les résolutions grossières et il ne correspond donc plus à une mesure locale. En outre, les effets des images partielles liés à la forme du champ de vue échographique auraient un impact bien plus considérable sur la qualité du recalage.

Il est donc préférable d'utiliser une mesure plus simple telle que le SSD. Le SSD fait par contre l'hypothèse que les deux images à comparer sont identiques après application de la transformation physique. Cette hypothèse n'est pas généralement vérifiée dans le cadre de l'échographie transrectale de la prostate, notamment à cause des différences d'intensités dues à des variations de la pression de la tête de la sonde et à des variations de l'angle de vue. Ce problème est atténué par l'application d'un filtre passe-haut des différences d'intensité afin d'éliminer les variations locales de faible fréquence.

Au lieu d'utiliser directement la dérivée de Gâteaux en tant que force issue de la comparaison des images, seule la direction de la dérivée est retenue pour le calcul de la force. L'amplitude de la force est déterminée en effectuant une minimisation unidimensionnelle en direction du gradient. Ceci permet d'accélérer considérablement la convergence du processus de minimisation itératif.

F-3.5 Forces issues de la Sonde

Les déformations de la prostate sont les plus prononcées autour de la tête de la sonde lors d'une série de biopsies de la prostate transrectale. En même temps, seule peu d'information visuelle est disponible à proximité de la tête de la sonde, ce qui entrave la convergence du problème vers la solution physique, car aucune force issue de la comparaison des intensités ne peut être calculée à la position de la sonde. Pour cette raison, un modèle biomécanique des bougés des tissus prostatiques suite à des mouvements de la sonde a été conçu. Il tente à prédire des bougés de tissu probables à partir de la différence de la position de la sonde dans l'image de référence et l'image flottante après recalage rigide.

F-3.6 Forces issues de la Consistance de l'Inverse

Il est bien connu que le processus de recalage déformable de R vers T, et de T vers R n'est que rarement symétrique, car il s'agit d'un problème mal posé. Il est possible d'exploiter cette asymétrie pour améliorer la qualité du recalage. L'idée est d'effectuer une optimisation alternée du champ de déformation de R vers T, dénommé u, et du champ de T vers R, dénommé v. Dans l'idéal, la norme

$$\int_{\Omega} ||\varphi_u \circ \varphi_v(x) - x||\ dx \tag{F.3}$$

correspond à zéro à la fin du recalage. Nous allons utiliser le gradient de cette norme en tant que force supplémentaire dans le processus d'optimisation itératif. Ainsi, les deux estimations sont couplées et peuvent se corriger mutuellement. Des expérimentations ont montré que le processus d'optimisation converge plus rapidement si cette technique est appliquée.

F-3.7 Expérimentations et Résultats

Le recalage élastique a été évalué avec le même jeu de données que celui utilisé pour évaluer le recalage rigide. La précision RMS du recalage s'améliore à 1.1 mm et l'erreur maximale à 2.93 mm. Le recalage élastique permet donc d'améliorer la précision du recalage rigide d'environ 30%. L'intégration du modèle biomécanique des déplacements de tissu dus aux bougés de la sonde fournit de bons résultats en général, mais dans certains cas il dégrade les résultats. Il est éventuellement nécessaire de recaler avec et sans modèle afin de pouvoir filtrer les cas d'échec. Le temps de calcul moyen du recalage total (rigide plus élastique) est inférieur à 17 secondes.

F-4 Applications Cliniques

Une étude clinique a été effectuée avec l'objectif de mettre en évidence la difficulté d'atteindre avec précision une cible sous contrôle échographique 2D. Lors du geste actuel il y a une double source d'imprécision : d'un côté, les cibles sont déterminées à partir d'un protocole systématique. Le protocole est défini sur une représentation bidimensionnelle de la prostate qui doit être mentalement projetée sur l'anatomie du patient par le clinicien lors de l'intervention. De l'autre côté, l'imagerie échographique ne fournit que très peu d'information structurelle de la glande, il est par exemple difficile d'identifier l'urètre et l'anatomie zonale de la prostate. Ce dernier point est accentué par le fait que la prostate est relativement symétrique et que le contrôle du geste est effectué avec une image bidimensionnelle seulement. En conséquence, il est quasimment impossible d'attribuer un échantillon prélevé à une zone anatomique avec une précision satisfaisante après l'intervention. Il y a un risque élevé que l'échantillonnage soit effectué de manière non-régulière et qu'il y ait de larges zones non-échantillonées.

Nous avons pu mettre en évidence ces hypothèses en effectuant un nombre statistiquement significatif de séries de biopsies sous guidage échographique 2D, mais en acquiérant en plus une image de contrôle tridimensionnelle après chaque prélèvement. Après recalage des images de contrôle nous étions capable d'estimer la précision qu'on peut atteindre avec la procédure classique. La première observation faite est qu'il y a des secteurs qui sont significativement souséchantillonés, en particulier des secteurs de l'apex et de la base. En outre, une courbe d'apprentissage statistiquement significative a été mise en évidence chez un des cliniciens, celui-ci ayant eu un retour visuel sur ses performances après les interventions.

F-5 Discussion et Conclusion

En conclusion on peut constater que le système proposé permet d'estimer les transformations physiques entre une image de référence et un flux d'images interventionnelles. La précision du système est suffissante pour les interventions diagnostiques, et il est également envisageable d'améliorer la précision de certains traitements du cancer de la prostate avec ce système. Le temps de calcul de l'algorithme est suffissamment rapide qu'une application temps-réel semble envisageable en ayant recours à des processeurs spécialisés du domaine des cartes graphiques.

Un point ouvert reste la validation automatique du recalage. Pour l'instant le recalage doie encore être validé visuellement, ce qui représente un inconvénient majeur pour des applications qui visent à guider le clinicien vers une cible hypoéchogène ou virtuelle. Ce problème sera adressé dans un avenir proche.

Chapter 1

Introduction

1.1 Computer-assisted Medical Interventions

During the last three decades, computer-based assistance systems for medical interventions have become an integral part of modern medicine. The objective of computer-assisted medical interventions is to improve the quality of traditional diagnosis and therapies, and to open ways for novel types of interventions. The role of the computer is to gather, process and interpret the abundant information provided by medical imaging and sensor systems, and to transform it into a condensed, stringent and human-readable form. The objective is to facilitate, improve or even drive the decision finding process for diagnosis, treatment planning and therapy. In the domain of medical robotics, the assistance goes even beyond pure information processing and extends to the partial or integral execution of the intervention, with various degrees and forms of human control.

1.1.1 Objectives

Hence, a main objective of computer assistance for medical interventions is to improve the quality of diagnosis and therapy by providing the clinician tools that facilitate or allow the analysis of information provided by medical cognition systems like imaging devices or sensors, in order to assist the diagnostical or therapeutic gesture.

Modern imaging technologies with their capacity to provide two- or three-dimensional images of human tissues have considerably expanded the medical options and made many novel types of intervention possible. However, images of different modalities are often used for diagnosis, planning and therapy, they are taken under different conditions and from different positions, and the imaged body also changes in position and time, which makes it difficult for many interventions to establish correspondences between anatomical features. This is accentuated by the fact that the human perception is not particularly well suited for the reception and interpretation of dense three-dimensional tissue images, which are quite different from the images that we are used to see. The goal is hence to identify the information that is relevant for the intervention, to isolate it and to present it in a form that is easily accessible for a human operator, or to use it to control a robotic medical system.

The clinical usefulness and applicability of such a tool depend on multiple criteria, which can be divided into health-oriented, operator-oriented and resource-oriented criteria. Computer-assistance systems aim to improve the health-oriented criteria, which consist of

- an easier and less ambiguous identification of the pathology and the optimal therapy,

- an increased accuracy of the diagnosis, planning and/or gesture,
- less inter-operator variability of the quality of diagnosis, planning and/or therapy,
- less side-effects and failure rates,
- more patient safety,
- and a reduced invasiveness of the intervention.

The operator- and resource-oriented criteria also have a decisive impact on the clinical acceptance of a novel tool, which is driven by the reimbursement politics of the social security and the practical convenience of the tool for the clinician. These criteria include

- the time necessary to perform the clinical action, whether by automation or by a facilitated access to relevant information,
- the ergonomy of the system, in particular of the software interface,
- the impact of the tool on the clinical work-flow and logistics, in particular concerning sterilization and obstruction in the operating room,
- and the cost of a system.

The latter criteria are often neglected in scientific discussions about medical computer-assistance, which is the reason why many proposed systems have difficulties to find their way into the operating room. We prefer an approach which considers the totality of the enumerated points. This "holistic" approach drove our decisions at all stages of this work.

1.1.2 Methodology

The methodology of computer-assisted medical interventions can be divided into four distinct, but dependent parts:

1. Information Acquisition

The first part consists of the acquisition of digital information about the problem, typically provided by medical imaging devices, external tracking systems, tactile sensors and pressure sensors or by the clinician. The information is filtered and modeled in order to obtain a problem-focussed and dense representation. The information can consist in organ position information, instrument position information, intervention planning information and others.

2. Establishment of Correspondence

The second part consists of the representation of the different information in a single reference frame, also called registration, which makes it possible to fuse and super-impose different types and layers of information. Depending on the application, the correspondence can be established with the help of the clinician or it can be automatically computed.

3. Reasoning

Decision making and reasoning based on the combined information represents the third part of a computer-assisted intervention. This step is either performed by the clinician based on a synthetic visualization of the combined information, or, in the case of full automation, by the computer system itself. Simulation of the intervention can also be a powerful tool at this stage.

4. Action

The final step consists in the clinical action, executed by the clinician or a robotic system, where hybrid approaches with the clinician being assisted by a robot are frequent. The action follows the decisions taken during the reasoning step.

1.1.3 Establishment of Correspondence: Tracking Systems

The correspondence between anatomical features modeled or imaged in different reference spaces, the main research theme of this work, is established with tracking systems. The most commonly used tracking systems are based on external reference tracking, on implanted fiducial tracking and on image processing, where the three methods can be used in mixed forms.

1. **External Reference Tracking:** Tracking of external references that are rigidly attached to the intervention target is widely used in computer assisted medical interventions. Optical systems follow tracking references with a stereo-vision camera, detect them in the images and derive their spatial position from the 2D geometry of the references in the stereo images. Stereo-vision based tracking is fast, robust and accurate, but it is problematic and often very invasive to fix visual references on subcutaneous targets, and these systems can considerably constrain the kinematic liberty of the operator during the intervention due to visual occlusion problems. Magnetic systems track distortions of an artificial magnetic field caused by small ferromagnetic probes which serve as reference. Magnetic systems are less accurate than optical systems and their accuracy is sensitive to the presence of ferromagnetic

and electrical objects in the operating room, but they are smaller than visual references and they are not exposed to the occlusion problem. They can hence be introduced into the human body, which is a considerable advantage compared to optical systems.

2. **Artificial Fiducial Tracking:** The second class of tracking methods relies on artificial fiducials that are introduced into the human body, either by implantation or, in the case of contrast agents, by injection into the metabolism. The fiducials have excellent visibility characteristics for the image modality with which they are used. Correspondence between different images of the target organ can be computed on the segmented fiducials. Typical fiducials are gold markers for CT-scans. Magnetic fiducials have also been used in combination with magnetic tracking systems. The inconvenients of artificial fiducial tracking are its invasiveness and the overhead that their implantation or injection into the body adds to the clinical protocol.

3. **Image-based Tracking:** Image-based tracking could also be called anatomical fiducial tracking. The idea is to identify visual correspondences between anatomical elements in different images of the intervention target, where the granularity of the tracked fiducials can vary from segmentation of the entire target to statistical voxel-intensity comparison. Image-based tracking is less invasive than external reference tracking and artificial fiducial tracking, but it is often challenging to find robust and fast registration algorithms for a particular intervention.

1.1.4 Clinical Protocol

The methodological steps presented in Sec. 1.1.2 have to be integrated into a clinical protocol, with a degree of transparency that depends on the interaction with the operator at each step. Typically, the clinical protocol consists of four temporal phases:

1. Pre-operative phase

This phase includes the acquisition of diagnostic data and the planning of the intervention, where the planning itself can already be carried out with computer assistance, and hence integrate all methodological steps enumerated in Sec. 1.1.2.

2. Intra-operative preparation

The intra-operative preparation phase consists in the establishment of the correspondence between the pre-operative planning and the tracking system that is

used to measure the intra-operative position of the patient anatomy.

3. Intervention

During the intervention, the tracking system continuously updates the position of the patient anatomy and the instruments with respect to the previously registered planning. This allows for a combined representation of the pre-operative planning with the anatomy and the instrument, which makes it possible to improve the intra-operative decision process and the accuracy of the action.

4. Post-operative validation

Computer-assistance systems can perform the post-operative validation of the intervention. They can also provide rich information for post-interventional analysis and patient or cohort statistics.

1.1.5 Establishment of Correspondence of Soft Tissues

While the correspondence problem is considered solved for interventions on hard tissue (bones), subcutaneous soft tissue tracking is still an active research domain. External tracking systems are often only of limited use when soft tissue tracking is required, since the tracking references can rarely be attached directly on the tissue without increasing the invasiveness of the procedure. Furthermore, these systems can only measure the rigid transformations of the target. Soft tissues, however, often get significantly deformed during interventions. Artificial fiducials are also invasive and large quantities of fiducials are required for deformation estimation.

Image-based tracking of anatomical fiducials poses different problems, depending on the image modality. Magnetic resonance imaging (MRI) can provide detailed images of the target tissues but limits the access to the patient, the range of instruments that can be used, and high quality MR imaging is time-consuming, which leads to low update rates of the images. Radiography-based imaging systems also yield sufficient visual information for a large range of soft tissues, but they irradiate the patient and often also the operator. Moreover, soft tissue imaging with radiographies most often requires the injection of contrast agents, and real-time computed tomography (CT) systems provide blurred images of low resolution. Ultrasound-images are finally real-time and image acquisition is cheap compared to MRI and CT scanners, but the detail and signal to noise levels of the images are low. Many diagnostic and therapeutic interventions are nevertheless performed under ultrasound control, which makes this modality an interesting candidate for soft tissue tracking.

1.2 Prostate Interventions

The main objective of this work is the conception of an ultrasound-based tracking system for the prostate. The clinical target applications are prostate interventions that are carried out under endorectal ultrasound control. The focus lies on needle punctures of the prostate, and in particular on prostate biopsy acquisitions.

1.2.1 The Targeting Problem

It is challenging for the clinician to perform minimally invasive interventions on subcutaneous targets under ultrasound control. For example in the context of prostate biopsies, brachytherapy and cryotherapy, dozens of needles have to be placed according to a pattern defined on a pre-operative planning. Pathologic tissues are in general isoechogenic, i.e. invisible in the ultrasound images, and interpretation of the poor anatomical information provided by ultrasound imaging requires a certain level of expertise for accurate needle placement. Worse, the target moves and gets deformed during the intervention, either by patient movements, by the respiratory cycle or by the needle placement procedure itself. The clinician hence needs to update the planned target positions with respect to tissue motion. This mental adjustment of the planning to intra-operative target modifications can be challenging and error-prone.

1.2.2 The Localization Problem

The localization problem is closely related to the targeting problem and concerns in particular prostate biopsies: since the clinician cannot derive and memorize the needle position in the gland with high accuracy on intra-interventional ultrasound images, and since the biopsy needle trajectories are in general not any more visible on ultrasound images once the needle is removed, the clinician has only a vague idea about the location of the trajectories once the needle is removed. This is in particular embarrassing in the context of prostate biopsies, where exact knowledge about the localization of the samples could be very useful for diagnosis and treatment planning.

1.2.3 The Inter-modality Target Mapping Problem

While carcinoma are in general isoechogenic, MR imaging is more sensitive to cancerous tissues. It is however technically difficult to perform intervention under MR control since MR images cannot be acquired in real-time with an acceptable resolution. Also, due to the dissimilarities between MR and ultrasound images, it is challenging for the clinician to map suspicious lesions detected in MR images

mentally into real-time ultrasound images. It is hence almost impossible to target suspicious lesions directly, which is one of the reasons why biopsies are acquired with a systematic protocol defined on statistical cancer distributions.

1.2.4 Critical Structures

While prostate cancer can be cured for the majority of cases, current prostate cancer therapies exposes the patient to high risks of severe side-effects, in particular incontinence and impotence, which can significantly reduce his quality of life. While it is possible to observe a tendency towards less side-effects with increasing application of less-invasive treatments like brachytherapy, the fact that always the whole gland is treated because of missing information about the exact cancer localization makes it still necessary to treat glandular tissues very close to critical structures. These are in particular the urethra, the nervous tissues on the outer prostate capsule and the rectal wall. In the context of biopsies, most structural damages concern the urethra.

1.3 Research Themes

1.3.1 Prostate Tracking

The main focus of this thesis lies on the establishment of the correspondence between prostate images acquired with transrectal 3D ultrasound. The objectives are to detect the exact position of the biopsy trajectories and to guide the clinician accurately towards a given target. The proposed method uses voxel intensity based image comparison techniques and a priori models of the endorectal probe movement during the intervention. The principal objective of the tracking system is to enable computer-assisted medical interventions on the prostate that optimize the criteria enumerated in Sec. 1.1.1. We will hence not only consider the technical aspects of the tracking problem, but also the general context of prostate interventions under ultrasound control, and propose several computer-assisted medical applications that can be realized with the proposed system.

1.3.2 Prostate Biopsies

Prostate biopsy acquisition serves as case study for an application of the prostate tracking system. Image-based prostate tracking during biopsy acquisition is challenging when compared to other prostate interventions under ultrasound control, since the probe is used as guide for needle positioning. Hence, the position of the ultrasound view cone of the probe varies considerably, and the probe moves and

deforms the gland during the intervention. Moreover, the patient is not under total anesthesia and moves for example when he feels pain during needle insertion. All these factors make prostate tracking during biopsy acquisition relatively challenging. Adaptation of the tracking system for different applications should be relatively straight-forward.

1.4 Context of this Work

This work has been carried out within the GMCAO[1] group of the TIMC[2] laboratory, which has been specializing in computer-assisted medical interventions and medical robotics for more than twenty years now. This allowed me to benefit from a rich experience in this domain and to avoid many pit-falls.

A particular strength of the laboratory are the many close collaborations with clinical partners on the one hand, and industrial partners on the other hand. During this work I had direct access to the urology departments of the university hospital of Grenoble and the Pitié-Salpétrière hospital in Paris, and in particular the collaboration with the Pitié-Salpétrière hospital was of inestimable value for this project. I would in particular like to mention the very inspiring exchanges that I had with Pierre Mozer, MD, PhD, who followed this work closely and with great enthusiasm. I would also like to mention Grégoire Chevreau, Jean-Claude Bousquet, MD, and Stéphane Bart, MD, who participated in the acquisition of the patient data that were used for this study.

In addition, this thesis had from the beginning on an industrial focus with the objective to conceive a prototype for a commercial product. The idea that this work would potentially be commercialized and hence used by a large number of clinicians was an important supplementary motivation during the last three years. The prostate tracking project was started with the research department of the company Praxim, La Tronche, under the direction of Antoine Leroy. At the end of the first year, Antoine created his own company, Koelis SA, as a spin off of the urology activities from Praxim, and I followed him into this adventure. Together, we did not only work on ultrasound-based prostate tracking, but we also conceived a computer guidance system for renal punctures under fluoroscopy control (which is out of the scope of this thesis).

This work was also part of collaborative projects financed by the French ministry of research: From 2004 to 2006 we benefited from the "Programme Hospitalier de Recherche Clinique - Prostate-Echo" and from 2005 to 2007 we were supported

[1]GMCAO = "Gestes Médico-Chirurgicaux Assistés par Ordinateurs" = Computer-Assisted Medical Interventions.

[2]TIMC = "Techniques de l'Ingénierie Medicale et de la Complexité" = Methods in Medical Engineering and Complexity.

by the Agence Nationale de Recherche (ANR) within the national Health Technology Program and the "Surgétique Minimalement Invasive" (SMI) project.

Chapter 2

A Tracking System for Prostate Biopsy Acquisition

Abstract

In this chapter we first discuss the clinical context of this work: the anatomy and the function of the prostate are introduced, followed by a brief overview on prostate cancer and prostatic cancer distribution, and also the current standard approaches in diagnosis and therapy. Crucial for the understanding of this work is the discussion on prostate motion and deformation that occur during therapeutic and diagnostical interventions. In the second section, prostate biopsies under 2D transrectal ultrasound, the case study of this work, is presented in detail, followed by a discussion on the insufficiencies of this method. Prostate biopsy tracking and guidance is presented as a solution to these insufficiencies in the following section, where the state of the art in this domain is presented. Finally, we give an overview on the method that we propose in this work to solve the tracking and guidance problem, a 3D ultrasound-based elastic prostate tracking system.

2.1 The Prostate

2.1.1 Anatomy and Function

The prostate is a gland of the male reproductive system, which is located below the bladder and in front of the rectum. It surrounds the urethra at the point where it connects to the bladder. The prostate is made up of approximatively 30 percent of fibromuscular stroma while the rest consists of glandular tissue, see Fig. 2.1.

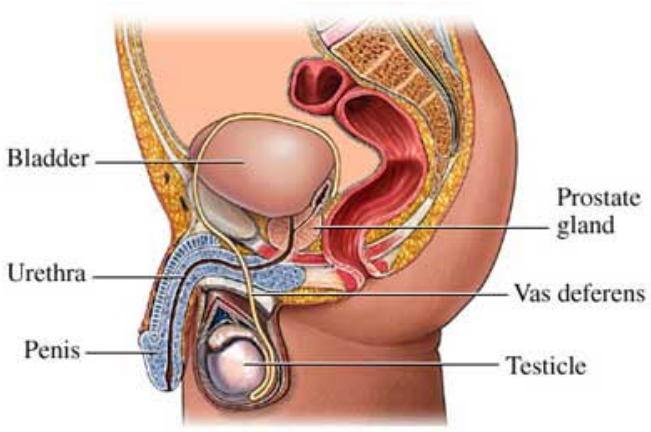

Figure 2.1: Prostate anatomy. [Image found at www.liv.ax.uk]

The main function of the prostate gland is to add seminal fluid to the semen. The seminal vesicle provide nutrients for the semen. The urethra carries urine from the bladder out to the penis. Tiny at birth and during childhood, the prostate gland doubles its size from the beginning of puberty until the age of about 25 years. The average diameter of a healthy prostate in adult men is 2 cm. After the age of 25 prostate growth almost stops completely. It slowly begins to grow again when men reach their early fifties. This second stage of growth is often problematic because it may lead to benign prostatic hyperplasia which may cause incontinence because of the urethra is squeezed by the growing tissues. Moreover, it is from this moment on that most prostate tumors begin to develop. Prostate growth is stimulated by the testosterone hormone.

2.1.2 Cancer

In the United States of America, prostate cancer is the most frequent cancer and the second cause of cancer-deaths for men. For 2007, Jemal et al. estimate [62] 218,890 new prostate cancer cases, and 27,050 deaths caused by prostate cancer. For 2006, a total number of 345,900 new prostate cancer cases, and 87,400 prostate

cancer deaths have been reported by Ferlay et al. [46] for the European Union (EU25). Today, one out of eight men risks the discovery of prostate cancer during his life. Prostate cancer represents 25 percent of all newly detected male cancers.

Prostate cancer is a in general slow-growing disease that mostly affects older men. In fact, more than 75 percent of all prostate cancers are found in patients older than 65 and it occurs rarely in subjects younger than 40. As any type of neoplasm, prostate tumors can be benign or malign. Only malign tumors lead to prostate cancer. In its early stage, prostate tumors require testosterone to grow and survive. Malign tumors can lead to metastasis when cancerous cells break away and enter the bloodstream or the lymphatic system. In 85 percent of all cases, prostate cancer is multifocal, i.e. tumors spread easily inside the gland due to its highly vascularized composition [51].

When prostate cancer grows it can affect surrounding non prostatic tissue. After capsular penetration, the tumor grows into the periprostatic fat, most often in the posterior-lateral part of the gland from where it moves towards the seminal vesicles via contact. The rectal part is only affected in rare cases, typically for unusually large tumors [51].

2.1.3 Zonal Anatomy and Cancer Distribution

Prostate tumors are more likely to be found in some regions of the gland than in others. The most widely accepted zonal model of the prostate was conceived by McNeal [81–83] in which the urethra divides the prostate into an anteriorly located fibromuscular stroma and posteriorly located glandular tissue. The glandular tissue is further subdivided into a peripheral zone (PZ), a central zone (CZ), and a transition zone (TZ), see Fig. 2.2. The peripheral zone comprises 70 percent of the glandular tissue and approximately 68 percent of prostatic cancers arise from this zone. The central zone comprises 25 percent of the glandular tissue and 8 percent of prostate cancers arise from this zone. The transition zone comprises 5-10 percent of the glandular tissue in young men but exhibits significant growth with aging. 24 percent of prostatic carcinoma arise in this area [41].

2.1.4 Diagnosis

If not treated, the mortality of prostate cancer is 75 percent for men younger than 65 [126]. Systematic prostate cancer screening is therefore essential to detect prostate cancer in early stages. Prostate cancer detection is nowadays carried out with a combination of the following methods:

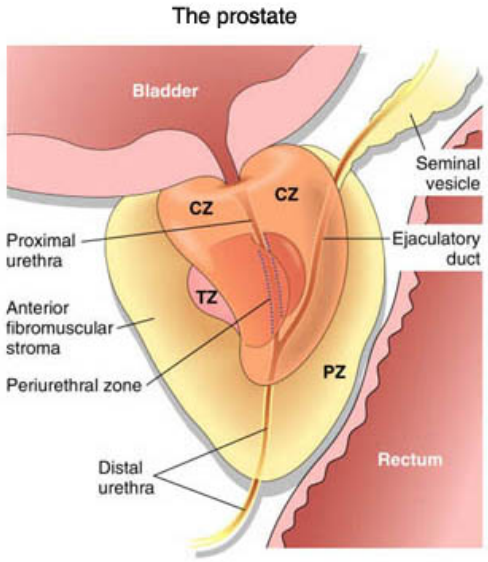

Figure 2.2: Zonal anatomy of the prostate.

Biological: Prostate Specific Antigene Screening

The prostate specific antigene (PSA) level raises in presence of pathologies in the prostate. The antigene is a protein produced by the cells of the prostate gland that can be used as a biological tumor marker. However, the antigene is not specific for cancer and also raises e.g. in presence of prostatitis or benign prostatic hyperplasia. Moreover, some men have naturally high PSA levels and some rare prostate cancers do not provoke a raise of the antigene level. Keetch et al. reported a positive predictive value of only 25 percent [67]. The cancer hypothesis hence needs additional confirmation in presence of high antigene levels. The currently applied threshold beyond which the cancer hypothesis is raised corresponds to 4ng/mL [126].

Clinical: Digital Rectal Exams

Digital rectal exams (DRE) aim to detect suspicious nodes by rectal palpation, see Fig. 2.3. This technique is still imperative because 10 percent of prostate cancers can be detected this way while the PSA level is below the commonly applied threshold of 4 ng/mL [126]. Early stage tumors can only be detected with DRE when cancer growth started in the peripheral zone of the gland.

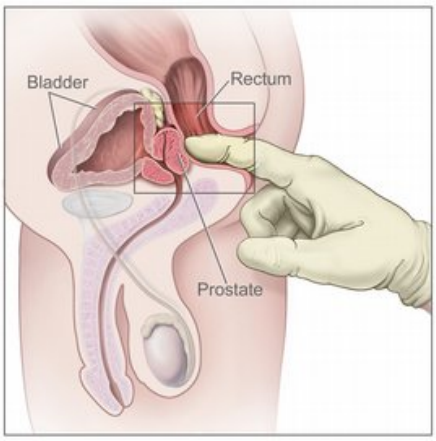

Figure 2.3: Digital rectal exam. [Image found at www.fccc.edu]

Histological: Prostate Biopsies

The only method to confirm the prostate cancer hypothesis consists in the histological analysis of tissue samples acquired from the gland. While PSA screening has a high cancer sensitivity but a low specificity, prostate biopsies yield highly specific result, but the sensitivity is only between 60 and 80 percent [67], depending on the biopsy acquisition pattern and the puncture path. In the case of persistent high PSA levels and negative biopsy results, the clinician faces a dilemma when all other explanations of the PSA level have been ruled out. In these cases, the biopsy series is repeated after several months. Each year, more than two million biopsies are performed in the United States and the European Union together.

2.1.5 Transrectal vs. Transperineal Biopsies

In the previous section we have introduced biopsies as the only method that allows to confirm the cancer hypothesis. Prostate biopsies are in the focus of this work, and we therefore introduce the different biopsy acquisition techniques, which are characterized by the access path to the gland. Two puncture paths have been reported for prostate tissue sample acquisition, the transrectal and the transperineal. The transrectal path is shorter than the transperineal and transrectal biopsies are in consequence well tolerated by patients, hence requiring only minimal anesthesia. There is, however, a relatively high risk of infection, which is the reason why the transrectal path is not considered for therapy. The transperineal path is more invasive but provides a better access to the peripheral zone of the prostate. Mohan et al. have estimated with simulations that 98.5 percent of the peripheral zone can be accessed via the transperineal path, while only 64.9 percent can be reached by

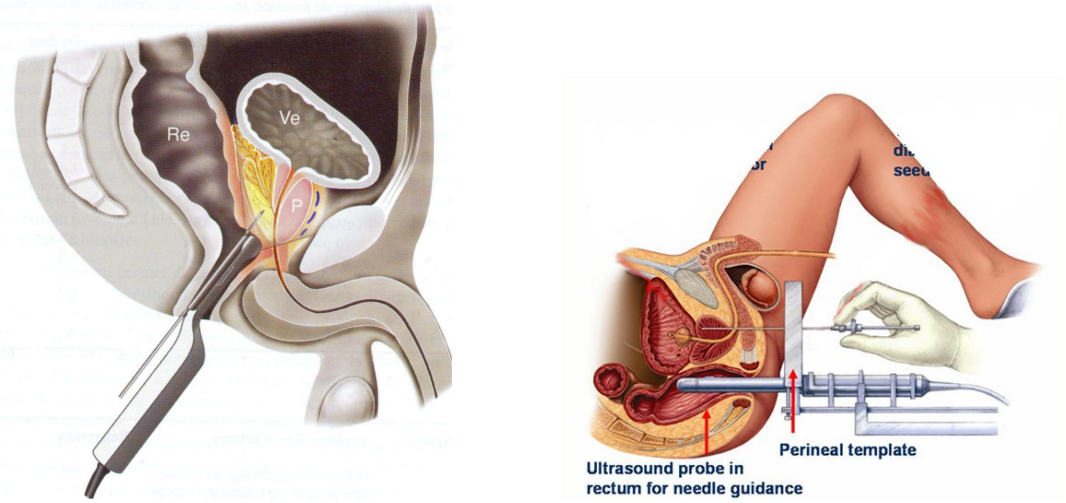

(a) (b)

Figure 2.4: Biopsy puncture paths. Fig. (a) illustrates the transrectal puncture path [125], whereas Fig. (b) illustrates the transperineal access with a needle template [Image found at www.prostatebrachytherapyinfo.net].

transrectal access. Both paths are illustrated in Fig. 2.4. The transperineal path is the preferred access for therapy.

Transrectal access represents currently the standard procedure for prostate biopsies, due to its simplicity. Only in rare cases, mainly when the rectum is not accessible in its natural way for example due to an anal stenosys or an abdomino-perineal amputation, transperineal biopsies are scheduled.

2.1.6 Therapy

PSA screening made it possible to detect cancer in early stages, where it can be successfully treated. Nowadays, the entire gland is treated when the cancer hypothesis is confirmed, mainly due to the fact that prostate cancer is multi-focal, since it can spread easily in the highly vascularized prostate tissues. Moreover, prostate tumors cannot be accurately located with current imaging technologies, despite the recent progresses that have been made in the domain of spectroscopic magnetic resonance image diagnosis and in ultrasound elastography.

Prostate cancer therapy is today a "nothing or all" decision: at early cancer stages, hormone treatment that prevents the production of testosterone can stop prostate cancer growth. When the cancer gets significant, the treatments are radical, i.e. the entire gland is treated. Surgery is the classic type of intervention, where radical prostatectomy, which involves removal of the prostate gland, seminal

vesicles and nearby lymph nodes is the most commonly exercised treatment. Less invasive treatments are therapies based on radiation. In conformal radiotherapy the gland is radiated from an external ray source until the tissues are destroyed. In brachytherapy, radioactive seeds (for example iridium seeds) are implanted into the prostate, which reduces the irradiation of healthy tissues. Thermal ablation is a heat-based cancer treatment where tissues are cooked with radiofrequency waves, microwaves or high intensity focused ultrasound (HIFU). The goal of cryotherapy is to eradicate prostate cancer by freezing the gland.

Radical therapy involves significant risks like incontinence, impotence and rectal problems for the patient. Damaging of structures surrounding the gland due to therapy can be observed in up to half of the men [1]. With the current tendency of screening with lower PSA thresholds, men are being diagnosed at a younger age with lower-risk disease. When the cancer hypothesis is confirmed, the patient often faces a therapeutic dilemma: recent studies show that about half of the detected carcinoma do not have a negative impact on the patient health [31]. A choice between two extremes has to be made: active surveillance, which ensures the preservation of the genitourinary functions at the expense of the psychological burden of intensive surveillance, and radical therapy, with all its potential side-effects.

For low PSA levels, the detected cancers are early-stage and hence often unifocal or unilateral. Ahmed et al. pose in [1] the question whether a paradigm shift towards focal therapy is appropriate to reduce side-effects and to increase the quality of life of prostate cancer patients. Focal treatments require, however, and relies on the accurate localization of malignant foci so that they can be treated with an adequate margin.

2.1.7 Prostate Motion and Deformation

Prostate motion and deformation poses a challenge for many diagnostic and therapeutic interventions. While the motion of the prostate is less pronounced than that of upper organs due to its distal location from the diaphragm and the support provided by the pelvis, it can nevertheless move and get deformed for several reasons. Malone et al. report prostate shifts of several millimeters due to respiratory motion [80].

Moreover, motion depends on the patient position and is larger in ventral decubitus position than in the supine position [123]. Herk et al. observed rotations of up to 4° [123] and Artignan et al. observed movements of 3 mm during a 20 minute radiotherapy session caused by rectal motility [5]. Hirose et al. observed that probe pressure and changes of the patient position lead to considerable gland deformations [57].

The insertion of the needle may also modify the form of the gland, depending

on the access path, the size and type of the needle and whether the needle is inserted with a spring needle gun or manually. Rotations of up to 13.8 degrees have been reported by Lagerburg et al. in brachytherapy when no locking needles were used [71]. In the same study, rotations up to 10.2 degrees have been observed even with locking needles. Manual insertion, for instance in brachytherapy, leads to more considerable prostate motion than when using a spring needle gun, since needle shots lead to a fast perforation of the gland.

Indications about the impact of probe movements and pressure can be found in the works of Beard et al., Padhani et al. and Balter et al. [8, 11, 93] who observed displacements up to 1.5-2.2 cm with respect to anatomical references. End-fire endorectal ultrasound probes, typically used for transrectal biopsy acquisition, can easily move and deform the prostate in a significant manner. Endorectal MRI coils may cause similar movements and deformations.

Patient movements cause additional problems when the organ needs to be immobile with respect to a reference frame during diagnosis or therapy. This is for example the case for radiotherapy treatment of the prostate, where the patient is not under total anesthesia and thus has to be immobilized to prevent accidental irradiation of healthy tissues. Patient fixation represents, however, a significant increase in patient discomfort. While this is an acceptable price to pay for treatment security, it is much more difficult to justify in the context of diagnosis.

2.2 Transrectal Biopsies under 2D Ultrasound Control

In this section the transrectal biopsy acquisition technique under 2D ultrasound control is presented, which represents the current clinical standard. This method is subject to several insufficiencies, enumerated at the end of this section, which are addressed by the solution that we propose with this work.

2.2.1 Methodology

Transrectal prostate biopsy acquisition is commonly performed under 2D ultrasound control. Most prostate tumors are isoechogenic and hence not visible on ultrasound images, which has the consequence that ultrasound imaging is not suited for tumor targeting. Hence, the role of ultrasound control is limited to the localization of biopsy sites identified either on images of different modalities or on a systematic biopsy distribution protocol. Modern 2D ultrasound probes can provide transverse and longitudinal images of the gland in real-time, see Fig. 2.5.

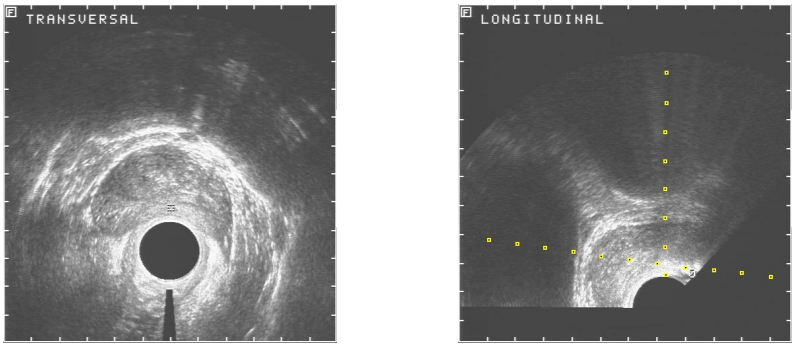

(a) (b)

Figure 2.5: 2D ultrasound planes. Fig. (a) depicts the transverse plane, and Fig. (b) shows the longitudinal plane.

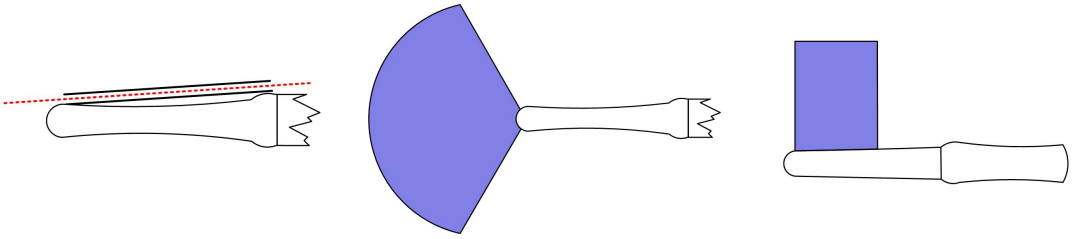

(a) (b) (c)

Figure 2.6: Endorectal ultrasound probes. Fig. (a) depicts a probe equipped with a needle guide, Fig. (b) illustrates the ultrasound beam of an end-fire probe, whereas Fig. (c) depicts the beam of a linear lateral-fire probe.

A particularity of ultrasound-guided prostate biopsy acquisition is the fact that the probe not only serves as imaging device, but also as needle guide. A small metal tube which determines the trajectory of the needle with respect to the ultrasound cone is rigidly attached to the probe prior to the intervention, see Fig. 2.6 (a). The axis of the tube is aligned with the longitudinal image plane of the probe, i.e. the needle trajectory lies in the plane of the longitudinal image. The double role of the probe is the main reason why end-fire probes are used for prostate biopsy acquisition (see Fig.2.6 (b)), in contrast to the linear lateral probes used for example in brachytherapy (see Fig. 2.6 (c)). Moreover, ultrasound devices for prostate biopsy acquisition visualize the hypothetical puncture trajectory, see the small dots in in Fig. 2.5 (b). Targeting thus consists in the alignment of the visualized trajectory with the ultrasound target. Spring needle guns are typically used for transrectal biopsy acquisition. The sample acquisition loop is summarized in Fig. 2.7.

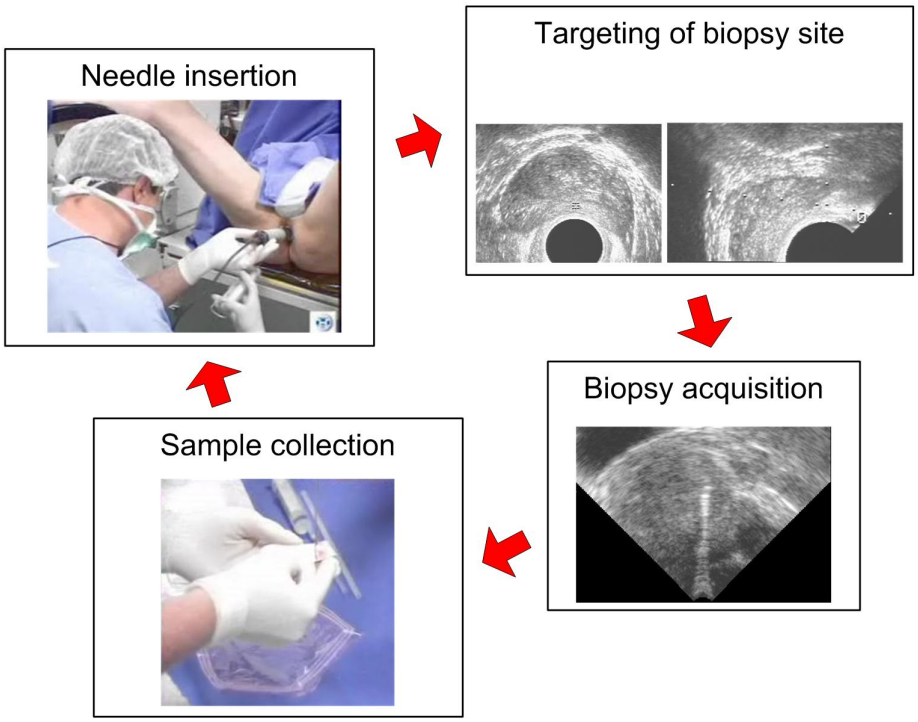

Figure 2.7: Biopsy acquisition loop. After needle insertion, the target is identified on the transverse and the longitudinal ultrasound planes. After biopsy acquisition the tissue sample is removed from the needle and put into a receptacle. The procedure is repeated until all planned sites have been biopsied.

2.2.2 Biopsy Target

The shape and the size of prostate cancer tumors vary significantly. A study on 1,832 radical prostatectomy specimens, performed by Ohori et al. [91], showed tumor volumes varying from 0.04 cm^3 to 2.13 cm^3. A tumor is considered clinically significant when it contains more than 0.5 cm^3 of tissue with a Gleason score of 7 or more [44, 110]. The Gleason score is a histological measure based on tissue texture analysis. Low scores from 1 to 5 indicate benign tumors while high scores from 6 to 10 indicate malign tumors. After considering that the tissue volume shrinks by approximately 20 percent after resection, the target tumor size can be approximated by a sphere of 5.2 mm radius. Even though the shapes of prostate tumors are irregular, this sphere size sets the upper limit of the biopsy site distances.

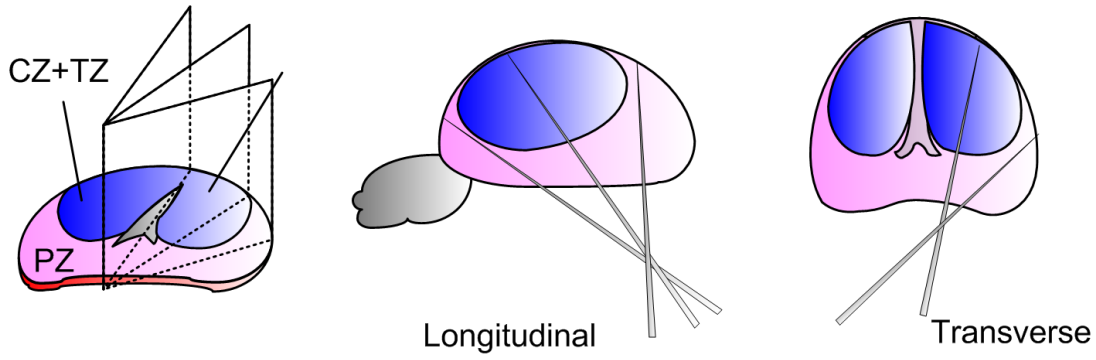

Figure 2.8: 12-core protocol. The Fig. illustrates the six targets in the right lobe; the targets in the left lobe are symmetrically defined. Most parts of the trajectories lie in the peripheral zone. The center of the prostate is not explicitly targeted to avoid damaging of the urethra.

2.2.3 Systematic Protocol

The lack of tumor sensitivity of ultrasound images makes it necessary to define the biopsy sites with statistical approaches. Nowadays, the most commonly accepted biopsy acquisition protocol is the so-called "twelve core" protocol, which has been shown to represent a reasonable compromise between tumor sensitivity maximization and hazard risk minimization. In this protocol, the biopsy sites are regularly distributed over the gland such that the statistical cancer distribution in the different prostate zones is taken into account, i.e. most of the trajectories aim the peripheral zone, see Fig. 2.8. Zones near the urethra are not targeted to avoid damaging of the urinal tract. In clinical practice, the 12-core protocol is often schematically defined in the coronal plane, see Fig. 2.9.

2.2.4 Tumor-atlas based biopsies

An alternative to the systematic biopsy protocol was introduced by Shen et al., who created a tumor atlas from digitized step-sectioned whole-mounted radical prostatectomy specimens with clinically localized cancers [108]. The digitized volumes were registered using non-linear registration in order to construct a statistical atlas of the cancer distribution. The cancer distribution was then used for an optimized target definition with the respect of the systematic protocol. A statistical predictive model was developed to optimize the biopsy sites, and experiments showed that in the samples used to build the atlas, the proposed method achieved a sensitivity of 99 percent with only 7 needles. In practice, it is challenging to implement this protocol since it is difficult to target the biopsy sites that the method proposes

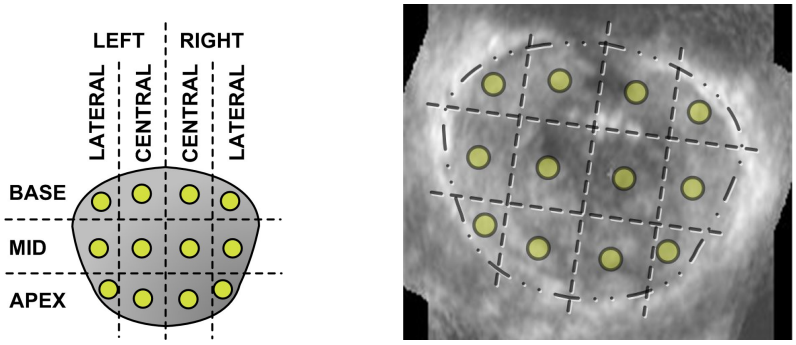

(a) (b)

Figure 2.9: 12-core protocol. Fig. (a) and (b) illustrate the clinical planning of the 12-core protocol, defined in the coronal plane of the gland.

under 2D ultrasound control.

2.2.5 Problems and Insufficiencies

Prostate biopsy acquisition under 2D ultrasound control suffers from several insufficiencies. First, targeting under 2D control reduces sample localization accuracy due to a lack of information in z direction. In particular for novice practitioners it is difficult to explore the gland regularly. This problem is accentuated by the low amount of structural information provided by the ultrasound modality and the highly symmetric shape of the prostate. Often, only the outer membranes of the gland can be clearly identified on the ultrasound images.

Due to the schematic target definition and the lack of localization accuracy, the exact localization of the tissue samples is unknown after the intervention. As a consequence it is difficult to map the localization of tumorous tissues identified via histological analysis back into the prostate. This is one reason that makes it difficult to reuse the histological results for localized treatment, in addition to the multifocal nature of mid- and late-stage prostate cancer.

A related issue arises when the biopsy series has to be repeated in the case of persisting high PSA levels. In this situation it is necessary to apply the systematic protocol again, without having the possibility to identify the sites that have already been sampled during the previous biopsy series. Finally, it is very difficult to target tissues identified as suspicious in images acquired with non-ultrasound imaging modalities, e.g. with spectroscopic MRI (MRSI) and standard MRI. If a clinician wants to target these lesions, he has to mentally map a 3D MRI target on the 2D ultrasound planes. We do not know of a clinical study where this approach has been investigated, which indicates the difficulties inherent in such a task.

2.3 Prostate Biopsy Tracking and Guidance

2.3.1 Interest

A prostate biopsy tracking system could provide solutions to the presented problems. A system capable of providing the exact location of the needle trajectory at a given instance during the intervention, and with respect to a tracking space that has yet to be defined, would enable several interesting approaches:

1. **Post-interventional Quality Control.** First of all, knowledge about the exact sample distribution in the gland would make it possible to perform a post-interventional quality control. In case when a large zone has remained unsampled, the clinician can acquire additional samples to ensure more reliable diagnostical results.

2. **Post-interventional Cancer Distribution Analysis.** Moreover, it would be possible to establish a precise cancer distribution map, which would eventually make it feasible to envision focal treatments, in particular at early cancer stages, when the probability of the presence of metastasis is still low.

3. **Intra-Interventional Positioning Control.** Intra-interventional visual feedback about the localization of the already acquired samples would allow to avoid multiple sampling of prostate sites, and the samples could be positioned more regularly.

4. **Previous Biopsy Site Navigation.** In analogy to the intra-interventional positioning control, a tracking system could be used to visualize the sites that have already been sampled during a previous biopsy series. This would make it possible to focus on unexplored zones during a repeated biopsy series.

5. **Statistical Target Navigation.** Moreover, a tracking system could guide the clinician to previously planned sites by super-imposing the targets onto the ultrasound image.

6. **Non-ultrasound Modality Target Navigation.** Finally, it could be interesting to project images acquired with different modalities into the tracking space in order to provide ultrasound-based targeting of non-ultrasound targets.

Several research groups have investigated the insufficiencies discussed in 2.2.5. In this section we present the state of the art in prostate biopsy tracking.

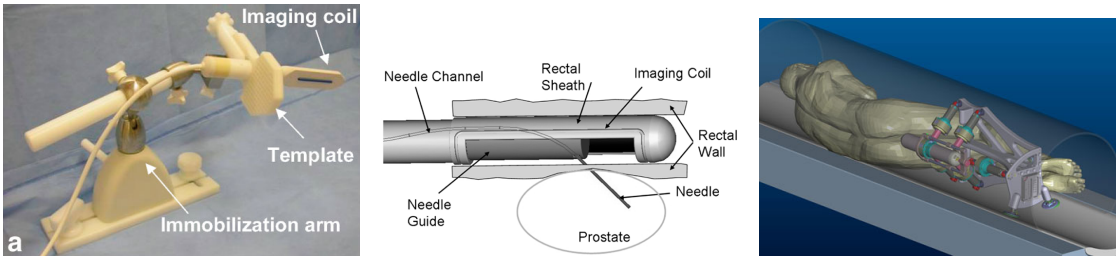

(a) (b) (c)

Figure 2.10: MRI-guided approaches. Fig. (a) illustrates the transperineal, template-based approach of Susil et al [116], Fig. (b) depicts the transrectal end-effector based approach of Krieger et al. [70], and Fig. (c) shows Stoianovici's robot [113]

2.3.2 MRI-guided Approaches

MRI-guided solutions seek to exploit the superior prostate cancer sensitivity of MR images with respect to ultrasound. Suspicious lesions are identified on the MR images and corresponding samples are acquired using a needle device that is aligned with the MRI coordinate space. Hata et al. use real-time fast gradient-recalled echo MR images to track transperineal needle insertion in real-time in an open MRI scanner [56]. Susil et al. perform transperineal biopsies with a system based on a lockable positioning arm equipped with a needle template, which is fixed to the endorectal MR imaging coil, see Fig. 2.10 (a). The template is calibrated with the imaging space, thus allowing to map image coordinates to the template. Biopsies are acquired in a closed 1.5 T MR scanner [116]. Beyersdorff et al. propose a system that provides transrectal access. Alignment of the needle with the MR images is performed with passive fiducial markers that are automatically detected in the MR images [14]. This approach works only with surface antennas, and alignment has to be performed repeatedly during needle positioning, which requires repeated high resolution volume acquisition and thus substantial acquisition times. Krieger et al. propose a manipulator mounted for transrectal biopsy acquisition equipped with an active tracking fiducial and mounted on the MR imaging coil, see Fig. 2.10 (b). Active fiducials can be detected in low-resolution and hence low-quality MR images which makes real-time tracking possible [70]. In [69] Krieger adds a needle system to the end-effector. Stoianovici et al. report in [113] a pneumatic, fully actuated robot with remote control, which can be used for both transrectal and transperineal access, see also Fig. 2.10 (c). The robot is fully encoded, and hence the alignment with the scanner space needs to be carried out only once, before the intervention. Alignment is performed by registration of a passive marker that is attached on the end-effector.

MRI-based approaches can provide precise targeting of suspicious lesions identified in MR images. The presented methods have, however, only limited capacities to deal with patient and organ movements. The presented manipulators and robots are all registered with the MRI reference, not with the organ. Both the patient and the organ can move considerably between the image acquisition, target identification and biopsy acquisition. The real-time manipulator control on low resolution MR images presented by Krieger is based on a 2D image stream and it is questionable if it is possible to use this stream for patient movement compensation. Its purpose is thus limited to the tracking of manipulator movements relative to the MRI frame. A cautious clinician can theoretically detect patient movements on the stream and eventually restart the procedure. However, repeated high-resolution MR image acquisition and target identification prolongate considerably the intervention time, which leads to an increased patient discomfort, higher costs and additional staff requirements.

Moreover, MR imaging is very costly when compared to low-cost ultrasound, and the additional hardware required for the presented MR-based approaches further increase the interventional cost. In the United States, more than 1,000,000 biopsies have been carried out in 2006, and similar numbers have been reported for the European Union. It seems very unlikely that hospitals in both continents will have the resources to let benefit a majority of the concerned patients from these techniques.

Further problems arise in the context of cleaning and sterilization of the presented devices, difficulties to make the solutions work for obese patients due to the limited work-space in the scanner, the requirement of special tables that allow fixation of the devices, additional staff training for competent handling of the innovative techniques.

To conclude, discussions with clinicians about the prospects of MR-guidance often led to the conclusion that the existing systems are "far from the clinical reality" when used for diagnosis. The reported solutions are, however, very relevant in the context of therapy, where the organizational and financial overhead that they introduce is fully justified by the considerable increase in treatment accuracy that they can provide, at the condition that the organ movement problem is solved.

2.3.3 Tracking-system based Approaches

In 2004, the TIMC laboratory performed a small clinical study, the ProNav project, which had the objective to investigate biopsy tracking based on external localization, and to collect data that allows to estimate organ movements during biopsy acquisition. The approach consisted in adding an optical tracking camera to the clinical setup which tracks a reference fixed on the probe (see Fig. 2.11 (a)+(b)), and a reference fixed on the pelvis of the patient, see Bucki et al. in [28]. The

(a) (b) (c)

Figure 2.11: ProNav project. Fig. (a) illustrates the infrared stereo tracking system, Fig. (b) shows the endorectal ultrasound probe with the tracking reference and Fig. (3) illustrates the prostate containing the biopsy axes (violet)

ultrasound beam was calibrated with respect to the tracking reference fixed on the probe. Before the intervention, a free-hand 3D ultrasound volume of the prostate was acquired. The ultrasound image sequences were recorded during the intervention, together with the corresponding spatial transformation. After the intervention, the biopsy needles were manually segmented in the images and projected into the free-hand volume, see Fig. 2.11 (c). The main observation of this study was, that the projected trajectories laid sometimes several centimeters next to the prostate, whereas the histological examination clearly showed that the sample contained only prostate tissues. This was not due to system imprecisions, but to the difficulties to compensate for patient movements, and to the fact that the gland moves and gets deformed during the intervention. We concluded that tracking alone is not sufficient, and that registration-based approaches are the preferred way to handle ultrasound-based biopsy tracking.

2.3.4 Tracking and Registration-based Approaches

Xu et al. propose in [130] a system based on a magnetic localization system which is calibrated with a 2D ultrasound probe, and on slice to volume ultrasound image registration. The free-hand reference volume is acquired before the intervention and serves to define the organ space. During the intervention, the ultrasound images are rigidly registered with the reference volume, using the spatial transformation provided by the tracking system as initial pose estimation. The principal objective of this work is to allow MRI-ultrasound guidance via multi-modal image fusion. A pre-operative MR image of the prostate is rigidly registered with the free-hand ultrasound volume before the intervention. The multi-modal registration requires manual segmentation of the prostate in both volumes prior to the intervention. During the intervention, the MR image of the prostate is visually projected into the 2D ultrasound image stream, which makes it possible to target

lesions detected in MR images. 2D-3D registration is carried out in 3 seconds using multiple images of the stream to achieve an increase in robustness.

The presented solution can correct for organ movements and hence should represent less intrinsic risks for the patient than the MR-guided procedures. Its accuracy is however compromised by the underlying image processing chain and the rigid transformation assumption: first, the prostate moves and is significantly deformed during endorectal image acquisition due to probe pressure. The assumption of rigid deformation between MR image acquisition and free-hand volume acquisition is questionable. Second, the free-hand volume acquisition process itself is biased due to the probe movements required to obtain an integral volume of the prostate. Moreover, the patient can move during the free-hand acquisition process, which can neither be detected nor corrected by the system. Third, simultaneous registration of several 2D images that have been acquired at different moments with a 3D volume introduces additional biases when patient movements and organ deformations occur meanwhile. Fourth, the system can lose track when the patient moves his pelvis considerably, i.e. more than 2 cm. In that case, the delivered sensor positions will not correspond any more to the sensor positions with which the reference volume was acquired. Fifth, 2D planes do not contain enough information for robust registration of images acquired during lateral biopsy acquisition, a fact that we demonstrate in Sec. 3.9.2 where we evaluate an approach based on the registration of three orthogonal planes with a volume. Sixth, the presented method ignores deformations at all stages of the processing chain. Gland deformations caused by probe movements are however inevitable when performing transrectal prostate biopsies. Xu et al. have not reported the accuracy and the robustness of their slice-to-volume tracking system, but we believe from our experiences that it is difficult to track lateral biopsies with this approach, and we also believe that the average accuracy of the ultrasound tracking system is not better than 4 to 6 mm, which is not sufficient when considering that the diameter of the average prostate is only 2 cm.

2.4 Proposed Method: 3D Ultrasound Elastic Prostate Tracking

In this work we propose a 3D ultrasound based elastic prostate tracking approach that makes advantage of the abundant information available in ultrasound volumes.

(a) (b) (c)

Figure 2.12: 3D ultrasound. Fig. (a) illustrates the longitudinal, Fig. (b) the transverse and Fig. (c) the coronal planes of a prostate. The illustrated planes have all been reconstructed from the same 3D volume.

2.4.1 3D Ultrasound

Recently, transducer-swept endorectal ultrasound probes have become available that are capable of acquiring 3 to 5 3D volumes of the prostate per second. Moreover, the ultrasound device industry currently spends considerable resources for the development of miniaturized 2D transducer arrays capable of achieving even higher acquisition rates. Simultaneously, most ultrasound manufactures investigate the novel CMUT transducer technology, which could replace the currently used discrete piezo transducer arrays by a single analogous transducer element capable of volume acquisition with considerably higher resolutions. Fig. 2.12 (a)+(b)+(c) show the longitudinal, transverse and coronal planes of a prostate, all reconstructed from a single 3D ultrasound volume.

2.4.2 Registration-based Approach

A purely registration-based approach is proposed that does not require additional information about the probe position originating from external tracking systems. The system estimates the rigid and the elastic transformations of the prostate that occur during the intervention. This is achieved with a priori models of endorectal probe movements during biopsy acquisition and of probe-induced deformations. Thanks to the model-based estimation, the presented approach is insensitive to patient movements, which is a considerable advantage when compared to existing approaches. Also, to our knowledge this is the first prostate tracking approach that considers elastic deformations. The tracking reference space is defined by a volume denoted anchor volume, acquired before the intervention, and with which all subsequently acquired volumes are registered.

2.4.3 Addressed Problems

The presented method is general enough to be used for all applications enumerated in Sec. 2.3.1, e.g. it could be used to solve the correspondence problem for post-interventional biopsy controls, post-interventional cancer distribution analysis, intra-interventional positioning control, repeated biopsy navigation and statistical target navigation. Moreover, it is feasible to realize MRI-based targeting, at the condition that a MRI to ultrasound volume registration technique is available, a topic that is beyond the scope of this thesis.

2.4.4 Principal Advantages of our Approach

Compared to existing biopsy tracking systems, the method that we have developed yields several important advantages:

1. We barely modify the clinical standard approach. The main differences consist in the replacement of the 2D ultrasound device by a 3D ultrasound device, and in the acquisition of an anchor volume before the intervention.
2. We solve the organ movement problem *and* the organ deformation problem such that we can guarantee an average root mean square (RMS) tracking error smaller than 1 mm for all tissues inside the prostate.
3. Our approach does not require patient fixation, since we can correct even very extreme patient motion.
4. We do not require additional tracking hardware, which means that we do not introduce any additional material and logistical overhead into the current procedure.
5. Our approach is potentially cheaper than MRI-guided approaches and external-tracking based approaches, both concerning cost of hardware and staff requirements.
6. We do not introduce additional sterilization issues.

2.4.5 Principal Scientific Contributions

From a technical point of view, the principal contributions of this work are

1. A comprehensive strategy is developed for information loss containment in multi-resolution approaches, which makes it possible to estimate most of the registration problem on very coarse levels. Information loss handling is the core concept that makes it feasible to envision real-time tracking with a volume to volume registration approaches.

2. An a priori probe movement model is presented that incorporates endorectal probe cinematics and transrectal ultrasound prostate image creation constraints. The model makes it possible to identify rapidly initial approximations of the rigid transformation between the gland in different images, independent of patient and probe movements. Hence, the proposed method definitely solves the problem of patient movements with respect to the tracking reference frame and does not require external tracking systems.

3. A panorama anchor image acquisition protocol is proposed in order to obtain a high quality tracking reference space that contains the complete prostate. This allows to register even very partial prostate images, typically arising during lateral biopsy acquisition.

4. In addition to the rigid probe movement model, a priori models on the probe-related deformations are conceived to improve the success-rate of the rigid pre-registration. A similar model is used to combine image-based elastic registration with a probe insertion simulation, which makes it possible to estimate probe-related deformations with very good accuracy.

5. A novel line search scheme is proposed for the computation of the image-based forces, which replaces a force scaling parameter that is difficult to determine.

6. A novel intensity homogenization approach is presented in order to make registration of dissimilar volumes of the same prostate possible.

Many additional smaller contributions are proposed concerning local optimization, threshold determination and optimization robustification.

2.4.6 Principal Clinical Contributions

Fortunately, this work could be carried out in very close collaboration with the urology department of the Pitié-Salpétrière hospital in Paris, where we were supported by a very motivated and encouraging clinical staff. Together, we evaluated from the beginning on the clinical aspects and prospects of the project. This collaboration made it possible to put the tracking system into practice already at an early stage of the project for a quality analysis of biopsy distributions in function of the experience of the practitioner.

In this work present the following clinical contributions:

1. We established biopsy maps for more than 50 patients which made it possible for the first time to visualize and to control the quality of the biopsy distribution after the intervention.

2. A statistical analysis was performed on these maps showing clearly that some zones of the gland are not sufficiently sampled under standard 2D ultrasound control. We were also capable of showing that the biopsies are more accurately placed with increasing experience of the practitioner.

2.5 Outline of this Work

2.5.1 Rigid Registration

In Chap. 3, the tracking problem is formulated as a registration problem, which is separated into a rigid registration part and an elastic registration part. We propose a rigid ultrasound volume registration system capable of estimating large transformations based on a rectal probe movement model. Furthermore, loss-containing image processing methods are introduced that make it possible to perform registration on very coarse levels of a multi-resolution pyramid of the volumes to be registered, making registration very efficient. The robustness of the proposed method is improved by filtering the bladder and probe-related deformations. Several improvements for classical local optimization concepts are presented. We further introduce a clinical protocol for acquisition of panorama reference volumes of the prostate. The presented algorithm is finally validated on patient data.

2.5.2 Elastic Registration

In Chap. 4 we introduce a fast and robust elastic registration framework, capable of estimating elastic deformations of the prostate. A cost function is proposed for the elastic registration problem that integrates image-based forces, probe-based forces and inverse-consistency constraints, i.e. our approach combines image-based registration methods with a bio-mechanical simulation. After introduction of a solution scheme for the problem, an analysis of its convergence properties is carried out, which leads to a timing parameter that depends on the tissue elasticity parameters. This makes it possible to set the tissue compressibility parameter explicitly without impacting the convergence rate of the estimation process. The force terms used for registration are analyzed in detail, starting with the image-based forces over the model-derived probe forces to a force term derived from the inverse consistency constraints. Moreover, a method is proposed for the intensity homogenization of a pair of ultrasound images that makes it possible to register images containing local dissimilarities. Again, the elastic registration framework is validated on patient data.

2.5.3 Clinical Application

In Chap. 5 the first clinical application of our system is presented, which consists in a prostate biopsy distribution study depending on operator experience. Biopsy needles were segmented in a post-processing step, which made it possible to project the sample locations into the anchor volume and to reconstruct hence the sample distributions for the patients who participated in the study. We had one patient who underwent a second biopsy series. Both series were fused using inter-series reference image registration to obtain the distribution of the biopsies acquired during both interventions in a single anchor volume. A statistical analysis of the biopsy distribution reconstructions shows clearly that some zones are systematically under-sampled. We were also able to show an improvement in biopsy distribution quality with augmenting experience of the clinician.

2.5.4 Discussion and Conclusion

In Chap. 6 we discuss the prospects of the presented system, in particular its potential clinical applications, and also its limitations. We describe the open issues and outline solution schemes that will be investigated in the future.

Chapter 3

Rigid Registration

Abstract

Soft-tissue tracking is still a challenging issue for many computer-assisted medical interventions. With the recent transducer-swept real-time 3D ultrasound probes it became feasible to conceive robust soft tissue tracking systems that are not only capable of identifying the rigid part of the target organ transformation, but also its non-rigid parts. Image-based tracking is mainly a registration problem for a stream of images. In this chapter we will briefly discuss a general registration framework for image-based tracking. Then, a rigid pre-registration system will be conceived that makes it possible to perform tracking without requiring additional external tracking systems like e.g. magnetic or optical tracking systems. The main parts of this system are a novel multi-resolution approach capable of handling the gain in entropy, i.e. loss of information, present in standard multi-resolution methods. Also, a similarity measure optimized for global search problems and three practical improvements to the Powell-Brent algorithm are presented. Moreover, panorama images of the prostate will be created to ensure that the entire gland is visible in the anchor image. A pre-processing strategy is presented that consists in bladder and probe-related strong local deformation masking. Finally, we present a method that combines purely image-based registration with the exploration of a priori models of the registration problem and conceive an efficient global search strategy based on a kinematic model of rectal probe movements during prostate image acquisition. The method was tested on 237 prostate volume pairs acquired from 14 different patients. The rigid pre-registration algorithm converged correctly in 96.7% of all cases with an accuracy of 1.41mm (RMS) and 3.84mm (max).

3.1 Introduction

3.1.1 Problem Positioning

The main objective of this work is to conceive a soft tissue tracking system for computer-assisted medical interventions on the prostate. Soft-tissue tracking is still an unresolved problem for many applications. Classical approaches that rely on marker or sensor tracking can in general not compensate organ displacements relative to the tracking reference, since the latter cannot be "fixed" to the organ, or only by increasing the invasiveness of the intervention. Moreover, it is impossible to estimate organ deformations occurring during intervention. Causes for deformation are manifold, the most frequent being patient movements, the respiratory and cardiac cycles and instrument pressure. Patient movements are a concern in particular for diagnostics, when no total anesthesia is applied.

With real-time 3D ultrasound it became possible to observe large parts of a given target continuously with a non-invasive imaging modality. The idea of using 3D ultrasound for soft tissue tracking thus seems to be an extraordinarily promising approach to solve the soft-tissue tracking problem for a large number of applications. This is in particular true for medical interventions on the prostate, since many of them are already performed under ultrasound control, e.g. seed implantation for brachytherapy, prostate biopsy acquisition or high intensity focal ultrasound (HIFU) treatments. The focus of this work will lie on prostate biopsy acquisition, which is, despite of the relative simplicity of the intervention itself, a challenging task for a soft tissue tracking system: the patient is not under anesthesia and can thus move, and the probe is not only used as imaging device, but also as a guide for needle positioning. This leads to large organ transformations during prostate biopsy acquisition when compared to brachytherapy or HIFU treatments. In this chapter we will first briefly present a general image-based tracking framework, and then focus on the estimation of the rigid transformation between the prostate in different ultrasound volumes.

3.1.2 3D Ultrasound-based Tracking

Several authors have taken advantage of the abundant information in 3D ultrasound images for tracking purposes: Ding and Fenster present in [40] a real-time algorithm based on the Hough transform capable of detecting linear instruments in ultrasonic volumes. Novotny et al. presented in [89, 90] a 3D ultrasound-based algorithm that detects linear instruments in real-time, based on Ding's algorithm and using a GPU implementation. A comparison of methods for tool localization in 3D ultrasound data is presented by Barva et al. in [10]. Instrument detection approaches all rely on detection of linear targets using a Hough transform based

segmentation approach. This approach cannot be applied for prostate tracking, since there is no similar algorithm available that rapidly and robustly segments the prostate in the volumes.

Relatively few investigations involving 3D ultrasound image based tracking of soft tissues have been reported. Morsy et al. analyzed in [87] a correlation search scheme to estimate rigid tissue motion in a stream of 3D ultrasound images. In the context of respiratory gated radiation treatment, Sawada et al. acquire a localized 3D ultrasound reference image of the liver in breath-hold state and register it rigidly with the treatment planning CT volume [105]. During therapy, localized US slices of the liver are continuously compared with the reference volume using image correlation to retrieve the planning position of the liver. Huang et al. register real-time 3D ultrasound images of the beating heart with a set of 4-D MR images covering the entire cardiac cycle [60]. An optical tracking system is used to initialize the spatial registration process while the ECG signal serves for temporal alignment. The authors achieve precise rigid registration in an overall computation time of 1 second with a mutual information based rigid registration algorithm. In both studies relative rigid movements between probe and target organ are limited to respiratory/cardiac cycle caused movements and are thus both predictable and repeatable up to a certain extent. The presented 3D ultrasound-based soft tissue tracking models have in common that they only estimate the rigid part of the organ transformation.

Krücker et al. propose a more general approach based on a two-stage, sub-volume based deformable registration algorithm for 3D ultrasound volumes [68]. In a first step, the affine transformation of the entire volume is computed, then, the volume is iteratively divided into sub-volumes which are rigidly registered to obtain a local estimation of the transformation. Thin plate splines are used to interpolate the transformation. The proposed algorithm executes in less than 5 minutes on a standard PC (2002) and achieves an accuracy between 1 to 3 voxel on artificially deformed images, dependent on the number of sub-volumes. Foroughi et al. [50] propose a rapid feature-based deformable registration approach for 3D ultrasound. Features are detected with an importance function defined on the image intensity, the gradient magnitude and the Laplacian of the Gaussian. The deformation field in between the feature points is interpolated with the sum of the distance-weighted displacements of all features. Non-rigid registration can be performed in 430 seconds with this approach on a standard PC (2006). Zikic et al. [133] present an approach for deformable registration of 3D ultrasound volumes of the liver, based on a variational formulation of the registration problem. Application of full multi-grid techniques to solve the Euler-Lagrange equations derived from the cost function makes their 3D registration approach remarkably fast compared to other deformable registration approaches, with an average computation

time of only 27 seconds for rectilinear volumes with 257^3 voxel on a standard PC (2006). Zikic and Foroughi do not treat the rigid pre-registration problem, which is, however, necessary to obtain robust results with the proposed methods. Rigid pre-registration further increases the reported computation times.

We can conclude that 3D ultrasound-based organ tracking can be formulated as a registration problem. There exist efficient approaches to estimate both the rigid and non-rigid part of the physical transformation. The presented solutions are, however, not yet fast enough to be used in a tracking system, and it is difficult to estimate their robustness in presence of noise and partial occlusion or partial visibility of the target. Furthermore, few information is given on the assumptions that the authors made on the probe movements between image acquisitions, in particular for the deformable registration methods. In the context of prostate biopsies, probe movements cannot be constrained without impacting the quality of the clinical intervention, since the needle guide is attached to the probe.

In this work we will try to solve the prostate tracking problem with a registration approach. We will try integrate a maximum of a priori information on the physical organ transformation to obtain a system which is sufficiently robust and to accelerate the estimation process as much as possible.

3.1.3 Tracking as a Registration Problem

The purpose of a tracking system is to provide the transformation between an object in reference space and the same object in tracking space. The reference space is fixed, while the tracking space corresponds to the current position of the device that is used to obtain positional information. In the context of transrectal ultrasound based prostate tracking, the tracking space is defined by the probe position at a given moment. The reference space will be defined by a volume that will be denominated anchor volume[1], while the tracking space will be defined by a volume called tracking volume. The tracking target itself is implicitly embedded via the image information in both the anchor volume and the tracking volume. See also Tab. 3.1 for a summary of the terminology that we just introduced.

Note that the spatial transformation between the ultrasound beam cones of the anchor and the tracking volume does not necessarily correspond to the spatial transformation between the tracking target in the anchor and the tracking volume. Recall from Sec. 2.1.7 that the tracking target can change its position, orientation and form during the intervention due to patient movements, respiratory motion and endorectal instrument motion. When considering longer time intervals between image acquisitions, target differences can be caused by changes in bladder

[1]This denomination was chosen instead of *reference volume* to avoid confusion with the same term commonly used in the context of image registration. In registration terminology, the anchor volume in our approach actually corresponds to the *template volume*.

Reference Space	Global referential for correspondence establishment
Tracking Space	Referential of a tracked object
Anchor Volume	Dedicated ultrasound volume that defines the reference space in our approach
Tracking Volume	Ultrasound volume that defines the tracking space in our approach
Template Volume	Floating volume during registration, corresponds to the anchor volume in our approach
Reference Volume	Fixed volume during registration, corresponds to the tracking volume in our approach

Table 3.1: Terminology.

and rectal filling, tumor growth, hormonal treatments and by other long-term variations in the human body.

These variations of the target object in tracking space are the reason why probe tracking alone using e.g. a magnetic or optical rigid transformation tracking system is not sufficient. A solution to this problem is to exploit the information available in the images using image-based registration algorithms. The basic idea of image-based registration is to estimate the transformation by comparing an image of the object acquired in tracking space with an image of the object acquired in reference space. This estimation process, also called image registration, can be formalized in its simplest form as the minimization problem

$$\varphi^* = \arg\min_{\varphi} \mathcal{D}[R, T \circ \varphi], \tag{3.1}$$

where $\mathcal{D} : \Omega \to \mathbb{R}$ is a function that measures a distance between the reference image $R : \mathbb{R}^3 \to \mathbb{R}$ and the template image $T : \mathbb{R}^3 \to \mathbb{R}$, and $\varphi : \mathbb{R}^3 \to \mathbb{R}^3$ is a voxel mapping function.

Note that we distinguish functional parameters, which will be put into brackets, from normal variables, which will be put into parentheses', e.g. $\mathcal{D}[R, T; \varphi]$ designates that $\mathcal{D}$ depends on the functionals R, T and φ, whereas $\varphi(\vartheta)$ designates that the functional φ depends on the parameter vector ϑ. The semicolon is a syntactical aid to group semantically related functionals.

3.1.4 Parametric image registration

In this chapter we will separate the registration problem into a rigid and a non-rigid part, and then focus on rigid registration. The non-linear part of φ will be

estimated in Chap. 4. To estimate the 6-dimensional rigid part of φ we use a parametric formulation of the registration problem, which is the preferred choice for transformation spaces with few degrees of freedom.

In the parametric framework, the transformation φ depends on the parameter vector $\vartheta \in \Theta$, where $\Theta \subseteq \mathbb{R}^d$ is the parameter space of dimension d. We can thus reformulate Problem 3.1 as

$$\vartheta^* = \arg\min_{\vartheta \in \Theta} \mathcal{D}[R, T \circ \varphi(\vartheta)]. \tag{3.2}$$

In our approach, T defines the reference space, and R defines the tracking space. We further denominate with $\hat{\varphi}$ the real physical transformation between the target in the reference image and the target in the template image, and the parameter set that generates $\hat{\varphi}$ will be called $\hat{\vartheta}$. Another notation that will be extensively used is

$$T_\varphi := T \circ \varphi. \tag{3.3}$$

3.1.5 Problem Separation into Rigid and Non-rigid Registration

For many medical image registration problems it is convenient to perform a linear, for example rigid or affine, pre-registration before estimating the nonlinear parts of the transformation between the objects in the reference image and the objects in the target image, in particular when the affine part of the transformation is large. The reason is that most efficient nonlinear registration algorithms only converge to the desired solution when most of the image elements lie inside the convergence range of $\hat{\varphi}$ for the used distance measure. Recall that the ultrasound probe not only serves as an imaging device in the context of prostate biopsies, but also as a guide to position the biopsy needle. For the needle to be positioned properly it is necessary to turn the probe by 180° around its axis when changing from one lobe to the other. Also, tilting up to 30° degrees, and probe head displacement on the gland membrane of a dozen of millimeters can frequently be observed. These are large affine displacements that are unlikely to be estimated robustly with a non-linear algorithm. This fact becomes more evident when considering that image-based registration is an inherently ill-posed problem: in general there exists a multitude of mappings φ such that $R \equiv T_\varphi$, but most of these mappings do not have any physical meaning. As a consequence, non-linear registration without affine pre-registration gets easily trapped in the local minima of an undesired solution. For these reasons it is preferable to divide registration into a rigid pre-registration step followed by a finalizing non-linear registration step. The feasibility of this approach depends on the size of the capture range of $\hat{\vartheta}$ when φ models a rigid transformation in Problem 3.2. The capture range not only

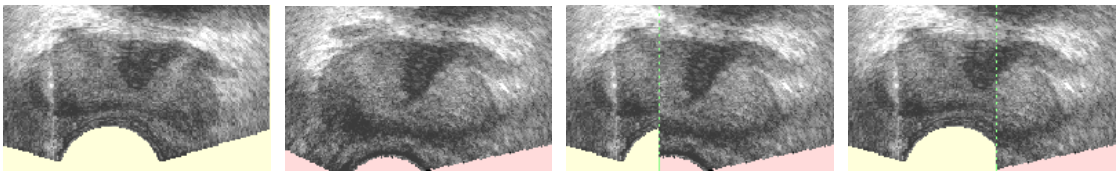

(a) (b) (c) (d)

Figure 3.1: Deformation study. Fig. (a) shows a deformed prostate in a volume acquired during biopsy acquisition (note the needle trajectory in the left side of the image), (b) shows the registered template image, in (c) and (d) the reference (right) and the template (left) are super-imposed to illustrate the differences in tissue compression. Note that deformations are strong near the probe head, whereas the upper gland and in particular the upper membranes are barely deformed.

depends on the distance measure $\mathcal{D}$, but also on the strength of the non-linear part of $\hat{\vartheta}$. When the non-linear part increases, the capture range gets smaller and may eventually even vanish for a rigid transformation model. In the context of transrectal ultrasound prostate images, this might happen when the operator presses the probe hard against the prostate, which would lead to a considerable squeezing of the gland. In a different context a distance measure might only yield interpretable results when applied locally, e.g. when the target has been bended by 90 degrees. Fortunately, the prostate deformations that can be observed in biopsy images are significant, but limited enough for a separation of the registration problem. The most flagrant deformations are caused by the probe head and they are limited to surrounding tissues. In particular the membranes far from the probe head are rarely considerably deformed, see Fig. 3.1. Zones of strong deformations are thus local and, since the probe head position with respect to the volume is known, their positions are reliably predictable.

In the remainder of this chapter we will focus on the rigid registration part of this thesis, and the non-rigid registration method will be discussed in the next chapter.

3.1.6 The Image Overlap Ω

Before presenting the state of the art of image-based cost functions it is necessary to define the image overlap Ω. Let the image domain $\Omega[I] \subset \mathbb{R}^3$ be the subset of $\mathbb{R}^3$ on which the image I is defined, i.e. the part of the image containing information. In this work we will assume that the ultrasound images delivered by the ultrasound hardware have been projected into a cartesian coordinate space. Hence, $\Omega[I]$ contains the Cartesian voxel that lie inside the ultrasound beam. In practice, $\Omega[I]$ is either implemented with a separate image mask structure or by a

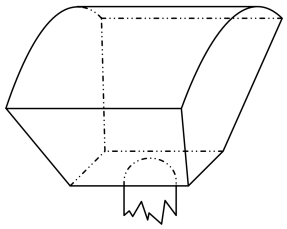

Figure 3.2: Image domain. The image domain $\Omega[I]$ of an image I corresponds to the ultrasound beam cone.

particular intensity value that is interpreted as mask value. Fig. 3.2 illustrates a typical image domain of 3D transrectal ultrasound images.

To simplify discussion we will make abstraction of image scaling and hence assume that volume pairs are always identically scaled. Then, the image overlap $\Omega := \Omega[R, T, \varphi]$ can be defined as $\Omega[R] \cap \varphi(\Omega[T])$. The grid points on which we will evaluate the distance measure are given by the set $\Omega^{\mathbb{N}} := \Omega \cap \mathbb{N}^3$. Note that the grid-points are voxel-centered. We define the registration problem in the Euler frame, i.e. we evaluate the similarity of image points of R and T_φ at the grid points $\Omega^{\mathbb{N}}$.

3.1.7 Feature-based Cost Functions

It is possible to define metrics on image features identified with an extraction algorithm. Barrow et al. [9] and Borgefors [21] base their feature matching algorithm on the chamfer distance. Besl and McKay [13] use the distance between closest feature points as distance measure. Feature identification can be target-oriented, i.e. parts or the totality of the target object is extracted e.g. using a segmentation approach, or invariant-descriptor oriented, i.e. the detection focuses on transformation-invariant elements in the images. In general, target extraction based methods are rather large-scale approaches that identify only few features like e.g. surfaces, while invariant descriptor based approaches rather yield rich sets of features. Most often, the invariant descriptors consist of invariant moments.

The attribute-vector approach proposed by Shen et al. represents a compromise between purely feature-based approaches and purely intensity-based approaches [107]. Shen's attribute vectors are composed of intensity information, contour information and geometric moment invariants. Both the contour information and the geometric moment invariants are rather basic and locally defined feature models that lead to a rich set of features, in contrast to approaches that rely for example on the segmentation of the entire organ. Shen's distance measure still represents the state of the art for MRI to MRI registration of brain images.

The attribute vector of Foroughi et al. used in [49, 50] is composed of even more locally defined features than Shen's, which leads to an abundant amount of information. For this reason, the authors introduced an importance function to filter out attribute combinations that seem less important than others. This approach was applied on ultrasound images of the liver and yielded good results.

Note that feature-based methods entirely rely on the quality of the feature extraction process, whether it is manual or automatic. When fully automatic registration is desired, the applicability of the extraction methods depend on the target visibility, the image modality and the type and quantity of noise in the images. Registration accuracy of feature-based method is in general determined by the detail level of the extracted features. When features can be robustly identified with a satisfying detail level they yield often better results than intensity-based approaches. In case when the images are very noisy, and when the features cannot be identified with a high degree of accuracy, the intensity-based approaches, which we will present in the next section, perform better.

3.1.8 Intensity-based Cost Functions

The simplest image distance measure $\mathcal{D}$ is the sum of absolute distances between image intensity pairs, i.e.

$$\mathcal{D}_{SAD}[R, T; \varphi] = \frac{1}{|\Omega^{\mathbb{N}}|} \sum_{j \in \Omega^{\mathbb{N}}} |R(j) - T_{\varphi}(j)| \,. \tag{3.4}$$

Very similar to this measure is the sum of squared distances

$$\mathcal{D}_{SSD}[R, T; \varphi] = \sqrt{\frac{1}{|\Omega^{\mathbb{N}}|} \sum_{j \in \Omega^{\mathbb{N}}} (R(j) - T_{\varphi}(j))^2}. \tag{3.5}$$

$\mathcal{D}_{SSD}$ finds a least square fit of the images which is optimal for many applications. Note that intensity difference squaring overweights outliers, which can be a disadvantage in image registration for instance in the presence of local noise or occlusion. Both $\mathcal{D}_{SAD}$ and $\mathcal{D}_{SSD}$ rely on the identity hypothesis, i.e. $R \equiv T_{\hat{\varphi}}$. In cases where the identity hypothesis is only weakly verified, more complex correlation models can be used. If the functional relationship between the intensities of R and T is only affine, i.e. $R \equiv f(T_{\hat{\varphi}})$ with $f = ax + b$, it is preferable to use the Pearson correlation

$$\mathcal{D}_{CC}[R, T; \varphi] \;=\; \frac{\mathrm{Cov}_{\Omega^{\mathbb{N}}}(R, T_{\varphi})}{\sqrt{\mathrm{Var}_{\Omega^{\mathbb{N}}}(R)}\sqrt{\mathrm{Var}_{\Omega^{\mathbb{N}}}(T_{\varphi})}}, \tag{3.6}$$

where the variance and the covariance are computed on Ω. When f is unknown, it is possible to use the correlation ratio, which estimates the optimal non-linear

least-square approximation $\tilde{f}$ of R in terms of T_φ, i.e.

$$\mathcal{D}_{CR}[R,T;\varphi] = 1 - \frac{\mathrm{Var}_{\Omega^{\mathbb{N}}}(R - \tilde{f}(T_\varphi))}{\mathrm{Var}_{\Omega^{\mathbb{N}}}(R)}. \tag{3.7}$$

or

$$\mathcal{D}_{CR}[R,T;\varphi] = 1 - \frac{\mathrm{Var}_{\Omega^{\mathbb{N}}}(T_\varphi - \tilde{f}(R))}{\mathrm{Var}_{\Omega^{\mathbb{N}}}(T_\varphi)}. \tag{3.8}$$

The correlation ratio was introduced by Roche et al. in [101], where one can also find the modeling of $\tilde{f}$. It is important to realize that the correlation is asymmetric: in fact, in Eqn. 3.7 the template image predicts the reference image, while in Eqn. 3.8 the reference image predicts the template image. It is always preferable to chose the image of the highest quality as the predicting image. When replacing $\tilde{f}$ by a linear mapping, one ends up with the Pearson correlation. In general, if the mapping function is known, the corresponding correlation-based distance measure can be found by replacing $\tilde{f}$ with this function.

The widely used mutual information distance measure, first applied in the context of image registration by Collignon et al. in [38] and by Viola et al. in [127], does not make any functional assumption on the relation between the intensities of R and the intensities of T, i.e. it is a purely statistical measure. Mutual information is defined on the marginal intensity histograms $G[I] : \mathbb{N} \to \mathbb{N}, I = R, T_\varphi$, and the joint intensity histogram $G[R, T_\varphi] : \mathbb{N}^2 \to \mathbb{N}$, all computed on the image overlap $\Omega^{\mathbb{N}}$. Then

$$\begin{aligned} H[I] &= -\sum_l G[I](l) \log G[I](l) \\ H[I|J] &= -\sum_l \sum_m G[I,J](l,m) \log G[I,J](l,m) \\ \mathcal{D}_{NMI}[R,T;\varphi] &= \frac{H[R] + H[T_\varphi]}{H[R|T_\varphi]}. \end{aligned} \tag{3.9}$$

This is the normalized formulation of mutual information, proposed by Studholme et al. in [115], which is more robust with respect to changes in the amount of image overlap than the original version.

Ultrasound specific distance measures derived from a Rayleigh noise model instead of a Gaussian noise model have been pioneered by Strintzis and Kokkinidis [114] with

$$\mathcal{D}_{SK}[R,T;\varphi] = \sum_{x \in \Omega} \left(\log(R(x)) + \frac{(R(x) - T(\varphi(x)))^2}{T(\varphi(x))} \right). \tag{3.10}$$

Cohen and Dinstein propose an improvement to Strintzis' measure by taking into account that ultrasound images are log compressed by the ultrasound machine manufacturers for visualization purposes [37], where

$$\mathcal{D}_{CD2}[R,T;\varphi] = \sum_{x\in\Omega} \left(\log(R(x)) - \log(T(\varphi(x))) - \log(\exp^{2(R(x)-T(\varphi(x)))}+1)\right) \tag{3.11}$$

This measure makes the simplifying assumption that the noise distribution is identical in both images, which is in general not true due to ultrasound attenuation and tissue variability [61]. Both ultrasound methods are maximum likelihood estimators for Rayleigh noise models.

As a general rule, the choice of a distance measure should be made in function of the a priori knowledge on the form of the intensity mapping f between the reference and the template image modality. When no functional assumption can be made, statistical measures are the only choice left to the user. However, when a precise functional model is available, it is preferable to use the most restrictive correlation model that still allows to estimate the function, thus incorporating a maximum of a priori information into the registration process. The reason is that the given metrics require increasing amounts of samples with increasing complexity of the estimated function f to yield statistically significant results. This is a concern when the image overlap Ω is relatively small, e.g. in the context of image mosaicing. Wachinger et al. add multiplicative weights in function of the size of the overlap to the measure in order to penalize small overlaps [128].

A given distance metric $\mathcal{D}[R,T;\varphi]$ only exploits a small part of the information available in both images. It can therefore be convenient to combine multiple measures in order to get more reliable estimates. Shen et al. apply in [107] the energy function on different aspects of the template and the reference images, more precisely on intensity information, edge information and geometric moment invariants, which they store in a so-called attribute vector. For each aspect of the image a cost function is evaluated and summed up to form a global energy term, i.e.

$$\mathcal{D} = \sum_i \alpha_i \mathcal{D}_i. \tag{3.12}$$

The sum of the measures corresponds to the maximum likelihood estimator when a Gaussian intensity distribution is assumed, see e.g. [100, 127, 128]. Foroughi et al. use attribute vectors to combine the intensity information with the magnitude of the gradient and the Laplacian of the Gaussian of the images [49, 50]. Haber et al. discard the intensity aspect entirely and define their distance measure only on the gradient [55].

3.1.9 Parametric Optimization

Minimization of the parametric Problem 3.1 is not trivial for prostate images acquired during a prostate biopsy procedure. Recall from Sec. 2.2 that the probe is used for needle positioning. Most clinicians perform a rotation of 180° around the probe axis when switching from targets of the left lobe to targets on the right lobe. Further, the probe needs to be tilted up to 30° for lateral biopsy sites. The position of the probe head contact with the prostate membrane can vary for several centimeters during biopsy. The translations and deformations of the prostate caused by probe and patient movements add to the large transformations of the acquisition space. As a consequence, the rigid part of the physical transformation $\hat{\varphi}$ that we try to estimate does in general not lie inside the capture range of the distance measure, where the capture range can be informally defined as the set of transformations for which local optimization converges to the physically correct transformation. The capture range varies strongly with noise, image occlusion, and in particular with the size of the overlapping parts of the images to be registered. The latter in particular applies for lateral prostate biopsies, on which often only 20 to 40 percent of the gland are visible. For these cases, the capture range reduces dramatically, and fast local optimization methods cannot be applied. On the other hand, the computational burden of global optimization in a 6 dimensional rigid transformation space is prohibitive for tracking tasks. To conclude, it is neither possible to rely on local optimization alone, nor is it possible to employ standard global optimization techniques.

Xu et al. [130] add a magnetic tracking system to the probe, which makes it possible to identify the probe position and orientation with good accuracy. The probe position is then used as initial guess for an image registration that identifies prostate movements relative to the acquisition space. The draw-backs of this approach are the additional hardware required near the patient, which can be cumbersome and leads to a longer preparation of the intervention. Also, interferences originating e.g. from metallic objects can distort the magnetic field which leads to erroneous transformations. Further, this approach requires calibration of the ultrasound image space of the probe with the sensor. If the sensor is not permanently fixed on the probe, this calibration step has to be carried out before every intervention, which can lead to significant additional setup time. This step may even double or triple the biopsy procedure time, which is a severe issue concerning clinical acceptance. Note finally that this approach does not solve the capture range problem, e.g. the combination of lateral biopsies with patient movements remains problematic.

Image-based solutions have been proposed by Gueziec et al. in [54] and Eadie et al. in [43] who both propose to reduce the intra-interventional computation time of global search by pre-computing a feature-based index hash table. During

intervention, similarity evaluation is replaced by computation of the geometric index followed by a fast database look-up. Concerning US image tracking this approach has the disadvantage of relying on feature extraction, which often lacks robustness when confronted with partial target images, speckle and US shadows. Also, they cannot reduce the complexity of the optimization problem and pre-computation time is not negligible.

Local optimization of Problem 3.1 in the space of rigid transformations can be carried out with gradient descent techniques, Newton-like methods, Levenberg-Marquardt, Simplex, Powell-Brent and others, where the latter two do not require the computation of the gradient. Maes et al. have shown in [79] that the performance of the different methods does not vary significantly in terms of computation time and accuracy. It seems that threshold choices, e.g. the initial step width, have a greater impact on performance than the choice of the algorithm itself. For this reason we have somewhat arbitrarily chosen the Powell-Brent scheme [24] that we will discuss in more detail in a dedicated section.

3.1.10 Outline of this Chapter

We have chosen an intensity-based registration approach of the rigid part of the registration problem. We will use a multi-resolution approach to accelerate the estimation and an attribute-vector based similarity-measure. An a priori model of the endorectal probe movements is used to find a point in the capture range of $\hat{\varphi}$. We propose a multi-step search algorithm consisting of a global pre-search and several local refinements on coarse levels, followed by a final local registration on finer levels.

In Sec. 3.3 we will first discuss a novel multi-resolution approach that is capable of handling the loss of information present in most classical multi-resolution systems, by introducing a loss-preventing inter-grid transfer operator and a loss-preventing tri-linear interpolation method. This will make it possible to carry out most of the rigid pre-registration on very coarse levels, thus obtaining an extremely efficient *and* robust rigid volume to volume registration approach.

In Sec. 3.4, an attribute-vector based distance measure will be presented that takes advantage of the capacity of the correlation coefficient to detect inverse correlations.

In Sec. 3.5 we will improve the capture range of $\hat{\varphi}$ with panorama images of the prostate which ensures that the anchor image contains the entire gland. Masking of the bladder and strong probe-related local deformations will further increase the capture range.

In Sec. 3.6 we will discuss several improvements of the local search algorithm of our choice, the Powell-Brent algorithm, for which we formulate a clinically meaningful termination criterion and integrate a resolution-switch for multi-resolution

registration. Moreover, we improve the robustness of the algorithm by detecting violations of the bracketing invariant on which Brent's line search is built.

In Sec. 3.7 we will conceive an efficient global search algorithm based on a kinematic model of rectal probe movements during prostate image acquisition.

In Sec. 3.9 the method is tested off-line on 237 prostate volumes acquired from 14 different patients.

3.2 Rigid Prostate Tracking Framework

3.2.1 Registration Objectives

In the context of prostate biopsies we not only want to register the volumes acquired during an intervention with the corresponding anchor image, we would also like to register the anchor volumes of different biopsy series performed on the same patient. Biopsy series are repeated when no tumorous tissue has been found in the biopsy samples of the last series, and when a persistent abnormally high prostate specific antigene level can be observed. The time lapse between the series can go from several months to one or two years. The prostate can significantly change during such long periods, in particular in presence of pathologies or treatments, but variations can also come from different ultrasound settings, different bladder filling and others. We do not expect to find an algorithm that will work in all cases, but we would also like to achieve good results for inter-series anchor image registration in the majority of cases with an automated approach.

A second important objective is near real-time registration, that can potentially be used during intervention after application of common optimization techniques like algorithm parallelization on specialized hardware. This requirement rules out many registration techniques due to their computational burden. It is for instance impossible to use general global optimization methods to solve the rigid registration problem.

Recall from Sec. 2.2.2 that tumors are considered significant when they have a volume superior to 0.5 cm^3. When the cancer has the shape of a sphere, the radius of that would be 5.2 mm. Tracking accuracy defines an upper limit to the accuracy of the clinical gesture when the system is used to for computer-assisted interventions. We therefore seek sub-millimeter (¡1 mm) root mean square accuracy for tissue registration inside the prostate, which should be sufficient for most clinical applications.

Registration robustness for intra-series registrations is essential when the algorithm is used for intra-interventional tracking, where permanent registration validation is time-consuming and potentially not feasible. Criteria are required that make it possible to automatically distinguish mis-registrations from proper

alignments.

3.2.2 Intensity-based approach

Recall from Sec. 3.1.7 that robust feature detection is a challenging process, in particular in the context of noisy ultrasound images with varying image characteristics. Also, feature detection algorithm often identify only sparse sets of features, which makes registration more sensitive to outliers and hence can have a negative impact on accuracy. This problem is accentuated when the overlap Ω is small, i.e. if either one or both of the reference and template images contain only part of the tracking target. This is for example the case for lateral biopsies, where the probe has to be positioned such that sometimes only about 20 percent of the gland are visible in the reference image.

Another difficulty consists in the rather symmetric shape of the prostate, which risks to render feature-based methods relatively insensitive to the rotational part of $\hat{\varphi}$, due to the fact that surface extraction algorithms discard information on the local properties of the surface, depending on the used a priori model of the surface.

Intensity-based problems do not rely on an a priori model and thus discard less information. Operating on a richer set of information leads to more robust and more accurate registration. For these reasons, organ-extraction based approaches have not been analyzed in the context of this work, and the focus was set on intensity-based and attribute-vector approaches.

3.2.3 The Rigid Transformation Model

We have chosen to limit our approach to rigid rather than affine registration since the pure affine part of the transformation $\hat{\varphi}$ is small and can be approximated more efficiently with non-linear registration approaches. Note that almost all the concepts that we will develop in this chapter can be easily extended to affine registration with nine or twelve degrees of freedom.

Rigid transformations in 3D space can be parameterized with a displacement vector $o \in \mathbb{R}^3$ and three Euler angles $\omega \in [-\pi, +\pi]^3$. We can construct a rotation matrix from ω with

$$M(\omega) = \begin{bmatrix} \cos\omega_3 & \sin\omega_3 & 0 \\ -\sin\omega_3 & \cos\omega_3 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos\omega_2 & \sin\omega_2 \\ 0 & -\sin\omega_2 & \cos\omega_2 \end{bmatrix} \begin{bmatrix} \cos\omega_1 & \sin\omega_1 & 0 \\ -\sin\omega_1 & \cos\omega_1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \tag{3.13}$$

and set $\varphi(\omega, o; x) = M(\omega)x + o$ for a point $x \in \mathbb{R}^3$. The parameter space Θ is thus 6-dimensional, and the parameter vector can be defined as $\vartheta = [o, \omega]$. If affine registration is required, the symmetric and positive definite rotation matrix M has to be replaced by a general affine matrix.

3.2.4 Interpolation

An interpolation operator is required for points $\varphi(x)$ for $x \in \Omega[T]$ that do not lie in $\mathbb{N}^3$. The most commonly used method is tri-linear interpolation, which represents a reasonable compromise between computational burden and interpolation quality. Note that the low pass filtering effect of linear interpolation can lead to oscillations in the distance measure $\mathcal{D}$ [121]. This is caused by the uniform or periodic interpolation error repartition on the points $Mx + o$ when both M is the identity matrix and the template is shifted along one of the three coordinate axis. A simple and effective way to avoid this effect in virtually all situations is to start registration always with a small rotational part in the initial guess.

3.2.5 Anchor Volume

A dedicated 3D volume of the prostate, the anchor volume, will be used to define the reference space. All volumes acquired during the intervention will be registered with respect to the anchor volume. For intra-interventional tracking, the anchor volume has to be acquired just before the intervention. For post-interventional biopsy distribution analysis, any volume can be chosen as anchor volume in theory. In practice, the anchor volume will have to fulfill several quality criteria, as we will see later.

3.2.6 Bounding box

Several parts of the presented algorithm will need approximate knowledge about the prostate location in the anchor volume. For this reason, the user is asked to provide an axis-aligned bounding box that makes it possible to delimit the prostate, and to identify approximately the prostate center, defined as the bounding box center.

3.3 Multi-Resolution and Information-Loss

Optimization on coarse resolution levels of a Gaussian image pyramid yields two important advantages: first, coarse levels are statistical aggregates of finer resolution levels and average out high frequency differences of the original image. This can lead to a more robust registration process, in particular in the presence of small local deformations in image regions with high frequency intensity patterns of large amplitudes. Second, optimization can be accelerated by several orders of magnitude. Once optimization on the coarsest level is terminated, the solution will be refined on denser levels, but from significantly better starting points, hence requiring less time consuming iterations on finer grids.

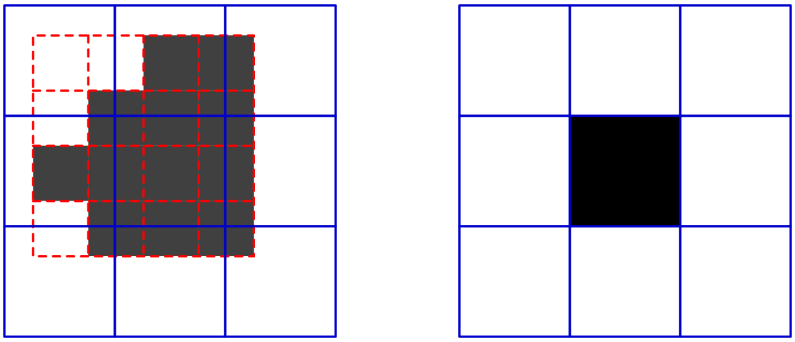

(a) (b)

Figure 3.3: Loss of information. Fig. (a) shows the fine grid voxel pattern (red), containing out-masked voxel (white) and information-containing voxel (gray). The blue grid illustrates the voxel-centered standard coarse grid overlay. Fig. (b) illustrates the resulting coarse grid data mask when the coarse grid voxels are masked out if at least one of the corresponding fine-grid voxels is masked out. In that case, only the center voxel contains data, all data in the border voxels of the fine grid is lost.

A coarse level voxel is in general built from s^d voxel of the finer level, where s is the inter-grid scaling, and d is the image dimension. We will in the following discuss the case $s = 2$ and $d = 3$, but the discussions in this section also hold for different values. A major problem of multi-resolution approaches is the loss of information inherent to the coarse grid construction process, which occurs when the number of voxels of the fine grid available to construct a coarse grid voxel is smaller than s^d (See also Fig. 3.3). While this is not a concern on fine levels, loss of information grows exponentially with respect to the number of grid levels, and thus represents a severe limit to extreme multi-resolution approaches. Similar problems arise for interpolation on coarse levels, which can create aliasing effects on the distance measure if not properly handled, and filtering, for instance for gradient computation. In this section we will present modifications of the standard inter-grid restriction operator that significantly reduce the loss of information. Moreover, the developed concepts will also be applied to linear interpolation and to gradient computation. The novel loss-containing multi-resolution framework makes it possible to perform registration robustly on extremely sparse levels, we will for instance carry out large parts of the non-linear estimation process presented in the following chapter on a $7 \times 7 \times 7$ grid only, constructed recursively from a $199 \times 199 \times 199$ grid. This makes extremely efficient volume to volume registration feasible, which is an important step towards real-time, volume to volume registration based organ tracking.

3.3.1 Inter-grid Restriction: the Fifty Percent Rule

It is necessary to define a grid transfer operator to build a resolution pyramid. We will first fix the inter-grid scaling factor to two. To be able to distinguish different grid resolutions, we will introduce the notation $\Omega^{h\mathbb{N}}, h \in \mathbb{N} \setminus \{0\}$ where h denotes the current grid scaling with respect to $\Omega^{\mathbb{N}}$. The inter-grid transfer operator, also called restriction operator, maps the intensity values of the grid $\Omega^{h\mathbb{N}}$ to the grid $\Omega^{2h\mathbb{N}}$ at the grid points of the coarse grid. The restriction operator can be defined as a convolution on the fine grid, using for example a Gaussian kernel. We will use the grid transfer operator kernel

$$K_1(j) = \begin{cases} 1/8, & j_1 = j_2 = j_3 = 0 \\ 1/16, & |j_1| + |j_2| + |j_3| = 1 \wedge |j_i| \leq 1, i = 1,2,3 \\ 1/32, & |j_1| + |j_2| + |j_3| = 2 \wedge |j_i| \leq 1, i = 1,2,3 \\ 1/64, & |j_1| + |j_2| + |j_3| = 3 \wedge |j_i| \leq 1, i = 1,2,3 \\ 0, & else \end{cases}, \tag{3.14}$$

where $j \in \mathbb{N}^3$. The inner grid points of an image I on level $2h$ can be computed with the discrete convolution $I^{2h}(j) = \sum_{k \in \mathbb{N}^3} K_1(k) I^h(2j+k)$. This voxel-centered restriction operator is also known as full-weighting restriction. It has optimal error smoothing properties when used in a multi-grid solver, see e.g. Briggs et al. [25], which will be an advantage for non-linear registration.

In the presented form, the kernel operator $K_1(j)$ cannot be applied on the border domain $\partial\Omega^{h\mathbb{N}}$, which leads to successive border cropping with every new resolution level and thus growing masks and shrinking domains. Fig. 3.4 (a) illustrates the two-dimensional case of this problem. To contain this gain in entropy[2] we apply the kernel on all available voxels on the given border position and obtain the intensity I, sum the kernel weights that correspond to these voxel in the variable θ, and finally check if $\theta \geq \frac{1}{2}$. If this condition is fulfilled, the coarse grid voxel is set to the intensity I/θ, else it is masked out. The result of this rule, which will be called fifty percent rule in the following, is illustrated in Fig. 3.4 (b).

3.3.2 Inter-grid Restriction: the Shift Rule

Unfortunately, the result is still not satisfactory, since the upper border of a dimension is not covered when the corresponding image resolution is even. We therefore introduce voxel shifts of $\frac{1}{2}h$ between grids for even resolutions, and the restriction

[2] When interpreting the intensities of the cropped border voxels as random noise, the information loss caused by the restriction operator can be measured with the Shannon entropy, also known as information entropy. Increasing entropy means loss of information.

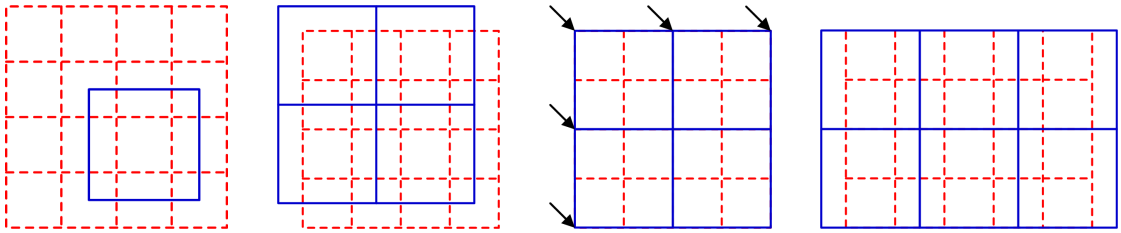

Figure 3.4: Restriction and information loss. In each figure, the solid blue frames correspond to the coarse-grid voxel resulting from the restriction, whereas the dotted, red lines correspond to the voxel of the fine grid. In Fig. (a), large amounts of information of the fine grid are lost on the coarse grid, since only the center voxel can be built from an integral set of fine grid voxel. In Fig. (b), the information of the lower borders in each dimensions has been recovered, but not of the upper borders. In Fig. (c), all information has been recoverd. Fig. (d) shows the case of a mixed pair-impair fine grid.

kernel

$$K_2(j) = \begin{cases} 1/8, & 0 \leq j_i \leq 1, i = 1, ..., 3 \\ 0, & else \end{cases} . \tag{3.15}$$

The grid resulting from the application of this kernel together with a shift is illustrated in Fig. 3.4 (c). The kernels for mixed pair and impair resolutions for the three dimensions will not be given explicitly, just note that the corresponding kernels are mixed forms of K_1 and K_2. An example of the a mixed form is illustrated in Fig. 3.4 (d). Note that the one has to take care to incorporate the grid shift into the transformation φ. Note further that both kernels are Gaussian.

An illustration of the fifty percent rule and the shift rule applied on patient data is given in Fig. 3.5. It can be observed that with information loss containment techniques the structural information of the image is preserved, and the integral prostate membrane is identifiable, whereas without loss containment, the available information shrinks to a small block compared to the original volume.

3.3.3 Masks and Information Loss

A gain in entropy can also be observed at mask borders of an image I, i.e. on $\partial\Omega[I]$. Endorectal prostate images are particular in the sense that the object of interest, the prostate, is very close to the probe head during image acquisition. Hence, tissue near the probe head is very close to the mask border and thus threatened by information loss during restriction. Worse, since the ultrasound beam is quite narrow near the transducer, it is in general impossible to capture the entire prostate

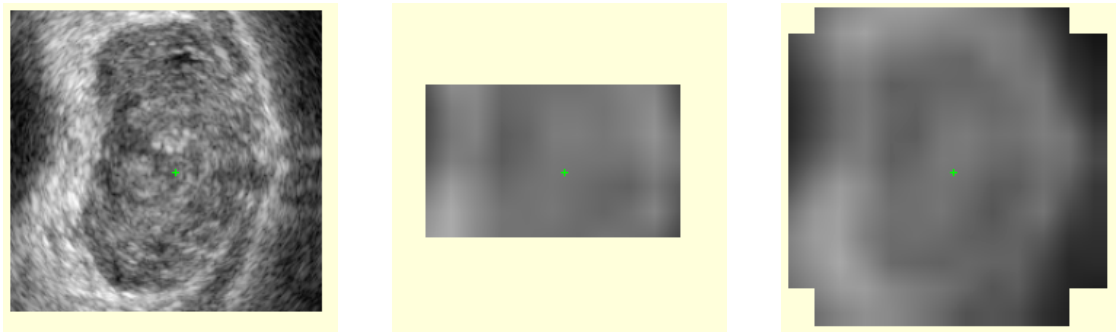

(a) (b) (c)

Figure 3.5: Restriction and information loss. Fig. (a) shows the original image with a resolution of $199 \times 199 \times 199$, Fig. (b) illustrates the four times downsampled image of resolution $13 \times 13 \times 13$ without information loss containment, and Fig. (c) illustrates the same image when loss-containing downsampling was used.

and mask borders hence lie even inside the target. Target information loss is thus a problem from the first restriction on, and rapidly leads to barely recognizable fragments of the original images. The kernel weighting mechanism introduced for border information loss handling, together with the 50 percent rule to decide whether a coarse grid voxel is masked out or constructed from partial data, also works fine for mask borders.

3.3.4 Interpolation and Information Loss

Information loss at mask and image borders is also a concern for tri-linear interpolation of the template image. Standard tri-linear interpolation on coarse level borders leads to aliasing effects and thus to smaller overlaps (see Fig. 3.6 (b)), which in consequence leads to rougher graphs of the distance measure. We therefore also modified the tri-linear interpolation operator such that the interpolation weights are summed up and checked against the 50 percent threshold. If the cumulated weight is superior to the threshold, it will be used to normalize the sum of the weighted intensities and the result will be returned as intensity value, else the mask value is returned. The result is illustrated in Fig. 3.6 (c).

3.3.5 Gradient Operator and Information Loss

In general, information loss is a concern for all filters with a kernel size larger than one voxel. In this work we will make excessive use of gradient computation on all resolution levels, both for rigid and non-linear transformation estimation. The

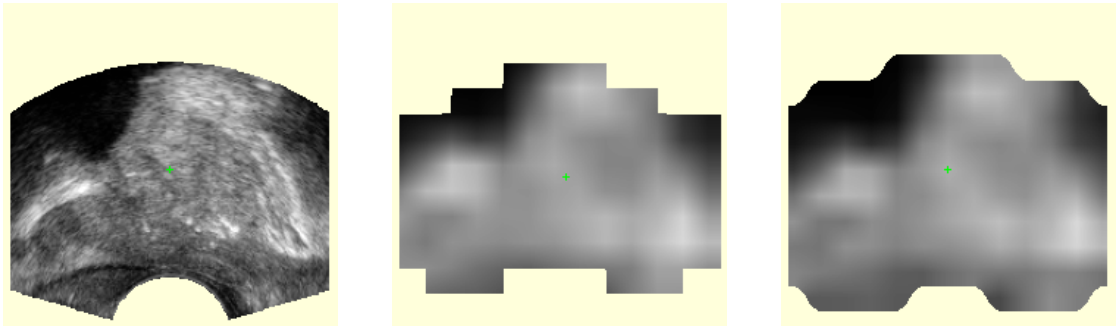

(a) (b) (c)

Figure 3.6: Interpolation and information loss. Fig. (a) shows the original image of resolution $199 \times 199 \times 199$, Fig. (b) illustrates the result of linear interpolation without using the fifty percent rule (image resolution $13 \times 13 \times 13$), and Fig. (c) illustrates the result of linear interpolation when the fifty percent rule is applied. The original image depicts a sagittal cut through the prostate.

components of the discretized gradient

$$\nabla I \approx (\frac{\Delta}{\Delta x_1}, \frac{\Delta}{\Delta x_2}, \frac{\Delta}{\Delta x_3})I \tag{3.16}$$

of an image I can be computed by applying the filter kernel $[\frac{1}{2}, 0, \frac{-1}{2}]$ successively in the directions $x_i (i = 1, 2, 3)$. This filter operator yields in general noisy results, since it is applied only on very few neighbors. The robustness of this kernel can be improved by applying the Gaussian filter operator $[\frac{1}{4}, \frac{1}{2}, \frac{1}{4}]$ in all directions prior to gradient filtering. Now, the gradient is computed on all 27 neighbors of a given voxel, each voxel being distance-weighted with the Euclidean norm. The weighting has been chosen such that the computation domain of the gradient can be interpreted as a sphere with the radius of one.

Weight counting and the fifty percent rule will again be used to reduce the impact of information loss on coarse levels, but only on the Gaussian filter. If one of the two voxels required to compute the partial derivative is missing, we will use the center voxel value instead, and divide the result by 2. This corresponds to the best possible approximation of the gradient for the given information. When both voxels are masked out, the mask value is returned. Note that this method introduces also an information barycenter shift corresponding to half the voxel side length in the direction of the partial derivative.

3.3.6 Discussion and Conclusion

The information loss problem imposes an upper limit to the number of grids that can be used for multi-resolution registration approaches. Border and mask cropping can only be contained by extrapolation of fine grid data on the coarse grid. In consequence the information at the mask and volume borders on the coarse grid is noisy because high frequencies are not attenuated, and the barycenter of the information used to create the voxel does not lie on the grid point. The information barycenter shift and a growing volume for impair fine grid resolutions resulting from the extrapolation introduce biases in the image that can lead to mis-registration. Worse, the bias grows exponentially with the number of levels.

Compared to information loss, the bias is, however, the smaller problem. On our test data set, robust rigid registration is possible with 5 resolution levels, where the resolution of the finest level is $199 \times 199 \times 199$, and the resolution of the coarsest level is $13 \times 13 \times 13$. When passing from the 5th to the 6th level, registration fails for a large number of volume pairs. The non-rigid registration algorithm that will be discussed in the next chapter robustly operates with 6 resolution levels, the coarsest level containing $7 \times 7 \times 7$ voxel. Without information loss handling, rigid registration starts to fail on the fourth level, and becomes completely instable on the fifth level. Note that the computation time, in particular for global search, is reduced by the factor of eight with each additional level.

A weak point of our information loss containment method is that the weights ω accumulated on border grid points with values between 50 and 100 percent are not recorded, which inevitably leads to incorrect weighting at coarse levels. A better approach is to memorize all weights in a mask structure, and to consider them when applying the fifty percent rule. The computational and the resource overhead of this approach can, however, pose additional problems.

3.4 Distance Measure

In this section we will discuss the choice of the similarity measure $\mathcal{D}$ for Problem 3.1. A comparative analysis of correlation-based measures will be carried out on different registration problems. The main objective is to choose a measure that produces few local minima and yields large and well-developed capture ranges around $\hat{\varphi}$. Also, the global minimum should correspond to $\hat{\varphi}$. Finally, the measure should also perform reliably on coarse levels.

3.4.1 Preliminary Considerations

The most commonly used similarity measures for monomodal registration are $\mathcal{D}_{SSD}$ and $\mathcal{D}_{CC}$. The correlation coefficient has the advantage that it can handle linear

variations in the image intensity distribution. Changes in the intensity distribution typically arise when the clinician modifies the ultrasound gain and ultrasound focus. The difference in the gain Δi is

$$\Delta i = 20 \log_{10} \frac{P_1}{P_2}, \tag{3.17}$$

where P_1 and P_2 are the sound pressures of two different acquisitions. When applying the distance law for sound pressure, which states that the sound pressure is inverse-proportional to the distance r of a punctual sound source, we get that Δi is independent of the image position. Since acoustic attenuation characteristics of tissue are independent of sound pressure, the new intensity i' at a given point is

$$i' = \Delta i + i. \tag{3.18}$$

When considering log-compression of the intensity value, the mapping is

$$i' = \alpha \log(i + \Delta i) + \beta, \tag{3.19}$$

for $\alpha, \beta \in \mathbb{R}$.

The ultrasound-specific speckle noise has a Rayleigh distribution, which is a multiplicative noise model. Some authors report registration improvements when using similarity measures based on Rayleigh noise based maximum likelihood models (see e.g. [23,37]) instead of Gaussian maximum likelihood estimators like for example $\mathcal{D}_{SSD}$. However, the presented measures $\mathcal{D}_{CD2}$ and $\mathcal{D}_{SK}$ all use the identity as intensity mapping function. Also, Rayleigh noise is not the only noise present in the images, we also have structural noise like ultrasound shadows caused by small air bubbles in the contact gel or arbitrary information due to partial probe contact and, of course, additive Gaussian noise from the image processing of the ultrasound hardware. Further, speckle noise textures can be used to characterize different types of tissue, i.e. speckle transports useful information on micro-structures of the tissue.

3.4.2 Comparative Study on Identical Images

We have carried out a comparative analysis of the $\mathcal{D}_{SAD}$, $\mathcal{D}_{SSD}$, $\mathcal{D}_{CD2}$, $\mathcal{D}_{CC}$ and $\mathcal{D}_{CR}$ on different types of images, and at different distances from $\hat{\varphi}$. The first comparison has been performed on a transrectal ultrasound image of the prostate that was used as reference *and* as template image, i.e. the identity assumption $R \equiv T_{\hat{\varphi}}$ holds. A translation of -40 mm to +40 mm in x-direction was applied, and the similarity measure was evaluated every 2 mm. Note that the choice of the direction is arbitrary, and that the choice of different directions leads to similar results. Note also that evaluation in a single direction is somehow superficial. It

is, however, methodologically difficult to compare the capture range of different similarity measures in more than one dimension. In our opinion, a comparison of similarity measures should be carried out on a large set of patient data together with a fiducial segmentation based validation, but this would have taken too much time in the context of this thesis.

The voxel side length was 0.33 mm, and the image resolution $199 \times 199 \times 199$ voxels. The results are given in Fig. 3.7. The results of the different measures have been multiplied by -1 to be usable for minimization. Also, they were normalized such that they fit between the interval [0,1], and their offset has been adapted to allow for better visualization. We modified the correlation coefficient such that it does not accept inverse correlations, since we got several times stuck with it on inversely correlated images during global optimization. Positions that lead to inverse correlations are set to the maximum, i.e.

$$\mathcal{D}_{CC+}[R,T;\varphi] = \begin{cases} 1 & ,\mathcal{D}[R,T;\varphi] \leq 0 \\ 1-\mathcal{D}[R,T;\varphi] & ,\text{else} \end{cases}. \tag{3.20}$$

The focus of this experiment lies on the size and form of the capture range of the different measures, and the position of the global minimum. In Fig. 3.7 we can observe that almost all measures have a similar graph with barely varying capture ranges. $\mathcal{D}_{CD2}$ is contrasting with its large false capture range from -23 to -40, while it performs better than the other measures from 40 to 29.

3.4.3 Comparative Study on Intra-series Images

Since the previous test is not very realistic, we applied the identical test on two different - but still almost identical - volumes acquired during the same biopsy series. Alignment was carried out using rigid registration. The aligned position corresponds to point zero in Fig. 3.8. $\mathcal{D}_{SSD}$ and $\mathcal{D}_{SAD}$ have a similar capture range than in the previous test, while $\mathcal{D}_{CC+}$ and $\mathcal{D}_{CR}$ have wider capture range than in the previous tests. $\mathcal{D}_{CD2}$ has, again, an asymmetric graph with a strong false capture range from -23 to -40.

3.4.4 Comparative Study on Inter-series Images

A third test was carried out on the anchor volumes of two biopsy series of the same patient. The time lag between both series was 14 months. The speckle patterns of both volumes are different, in particular inside the prostate. Also, the overall intensity is clearly different. Instead of translating the template volume, we now rotate it from -90 to +90 degrees around the probe axis. Again, both volumes were rigidly aligned, and an artificial displacement of 7 mm and a probe tilting

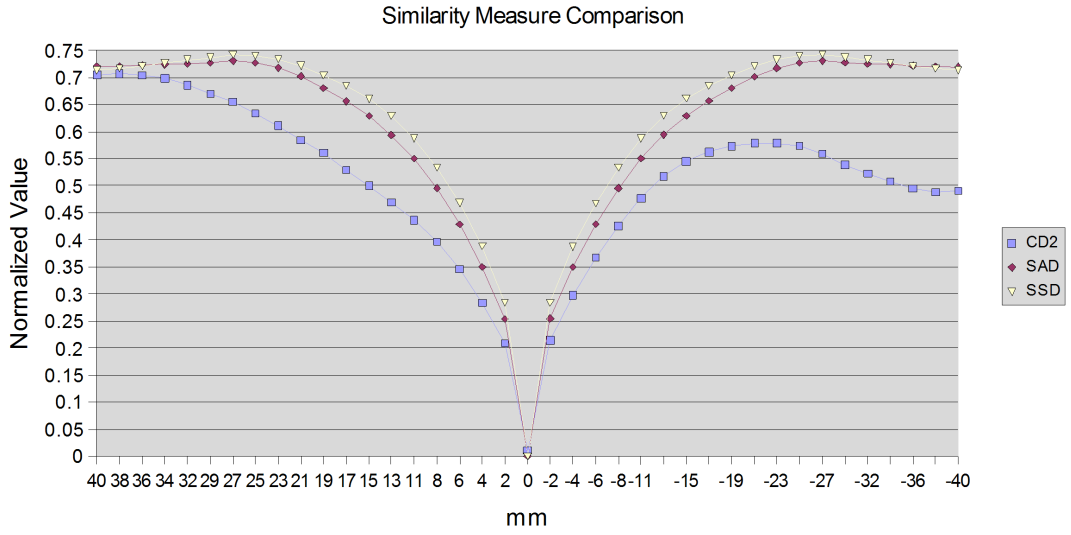

(a)

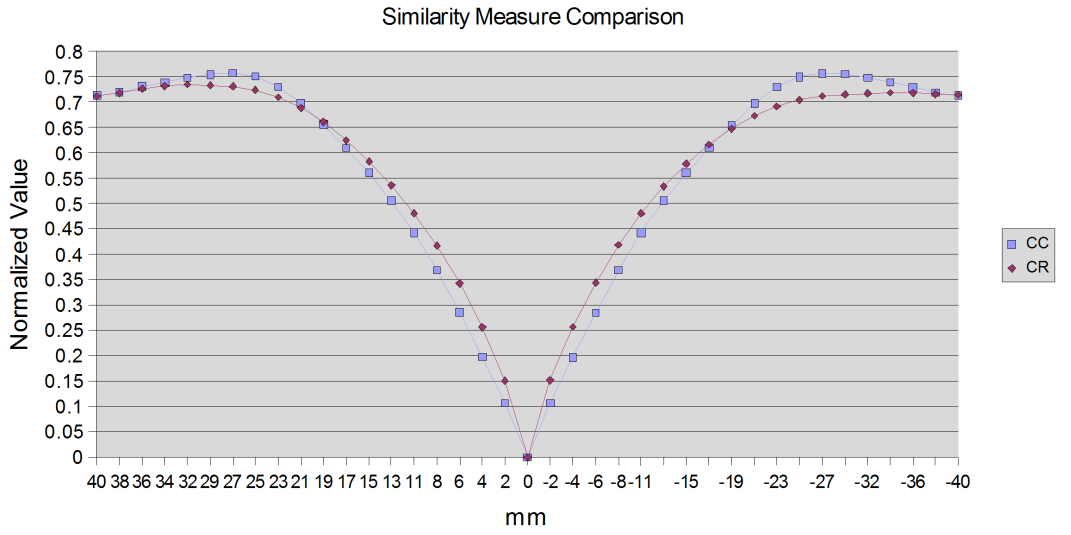

(b)

Figure 3.7: Similarity measure evaluation (single image comparison). Fig. (a) shows the graphs of $\mathcal{D}_{SAD}, \mathcal{D}_{SSD}$ and $\mathcal{D}_{CD2}$ for a -40 mm to +40 mm translation along the x axis for two identical volumes. The point zero corresponds to the aligned position. Fig. (b) shows the graphs of $\mathcal{D}_{CC+}$ and $\mathcal{D}_{CR}$.

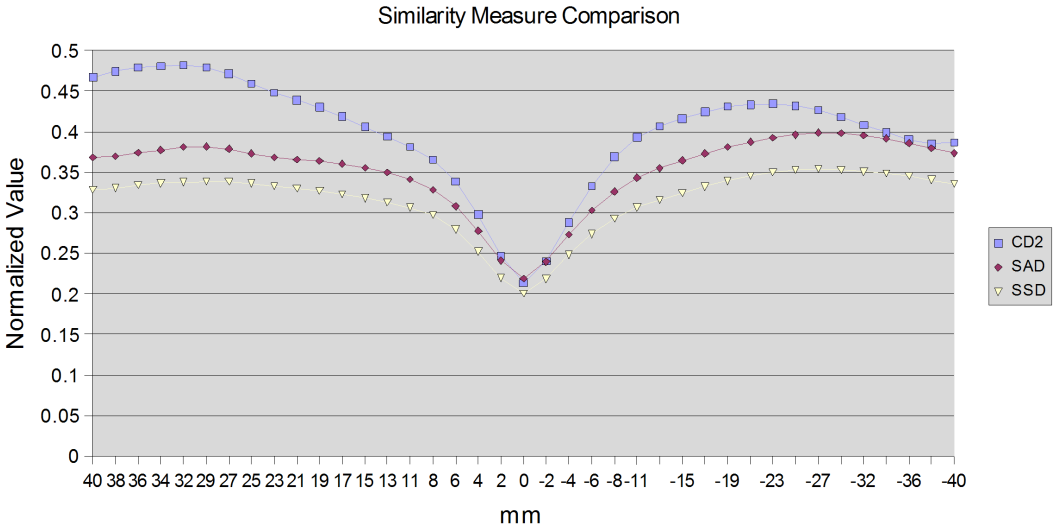

(a)

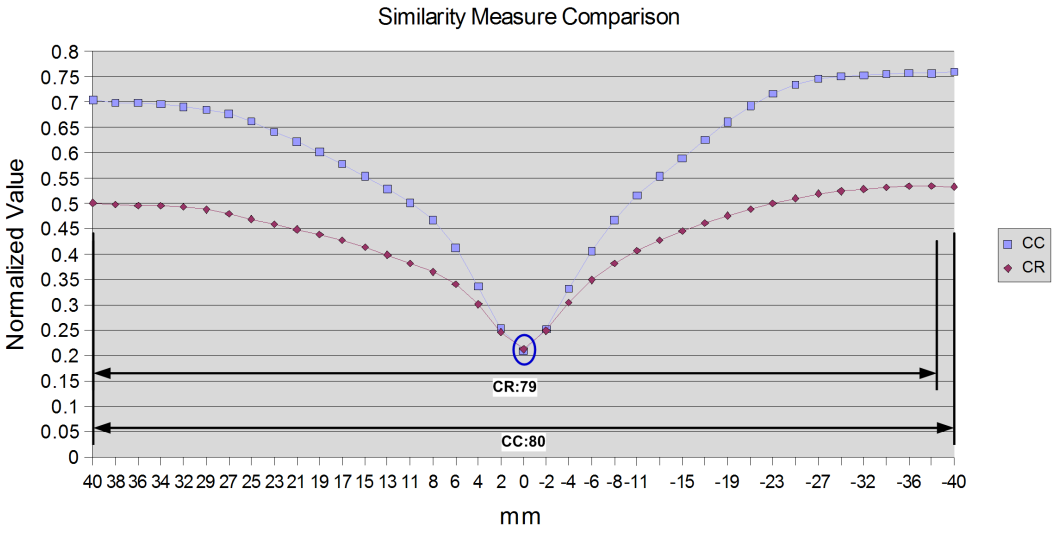

(b)

Figure 3.8: Similarity Measure Evaluation (intra-series comparison). Fig. (a) shows the graphs of $\mathcal{D}_{SAD}, \mathcal{D}_{SSD}$ and $\mathcal{D}_{CD2}$ for a -40 mm to +40 mm translation along the x axis for two different volumes. Alignment (point zero) was carried out using rigid registration. Fig. (b) shows the graphs of $\mathcal{D}_{CC+}$ and $\mathcal{D}_{CR}$.

of 10 degrees was added to the point zero. In reality, the initial displacement can be as far as 20 mm and probe tilting of 30 degrees has been observed. The results of this configuration are given in Fig. 3.9, and we can observe that $\mathcal{D}_{SAD}$ and $\mathcal{D}_{CD2}$ produce false local minima, and the capture ranges of all measures have flattened out. $\mathcal{D}_{CC+}$ has a smaller capture range than the simple correlation models, but the global minimum is at the correct position. The capture range of $\mathcal{D}_{CR}$ is considerably flatter when compared with the other measures. It produces a false local minimum in the zone where the correlation coefficient produces inverse correlations. Its minimum is correctly positioned. $\mathcal{D}_{SSD}$ clearly has the largest capture range, but it is relatively flat.

3.4.5 Comparative Study on Coarse Levels

Before passing to the next set of tests, we will eliminate the measure $\mathcal{D}_{CD2}$, first, because its performance is not convincing. The tests show clearly that Rayleigh models do not improve registration success or accuracy; in contrary, $\mathcal{D}_{CD2}$ rather underperforms its peers. Also, the differences in the speckle pattern in the previous test did not represent a challenge for the Gaussian noise based correlation models. The second reason is that we will make a test on the gradient magnitude image, where we will have intensity values of zero, which would have required to adapt the measure. Before doing this, however, we perform the same test on the 4th resolution level, i.e. the image resolution is now $13 \times 13 \times 13$, everything else being identical than in the previous test, but without $\mathcal{D}_{CD2}$. The results are given in Fig. 3.10. Both $\mathcal{D}_{SAD}$ and $\mathcal{D}_{SSD}$ still have wide capture ranges, but they flattened out further: Both measures yield almost identical values near the global minimum. The global minimum of $\mathcal{D}_{SAD}$ is at the correct position again, but it still has a false local minimum 24 degrees. The correlation coefficient still has a nicely shaped capture range, without excessive flattening. The minimum is at the correct position and is well pronounced. Its capture range, however, is smaller than the ranges of $\mathcal{D}_{SAD}$ and $\mathcal{D}_{SSD}$. The correlation ratio finally breaks down due to an insufficient number of statistical samples. Similarity measures with no particular a priori assumption on the intensity mapping like the correlation ratio and also mutual information require a much larger number of statistical samples to produce significant results. They are therefore less appropriate for extreme multi-resolution approaches than less flexible measures.

In the next test we evaluate the similarity measures on both the original images, and then the gradient magnitude images, in analogy to the Hammer approach [107]. The measures are combined by multiplication. If the correlation coefficient yields an inverse correlation for one image pair, the entire result is set to the maximum. The results of this test are given in Fig. 3.11. The capture ranges of $\mathcal{D}_{SAD}$ and $\mathcal{D}_{SSD}$ are less flat than in the previous test, but it is still difficult to clearly identify

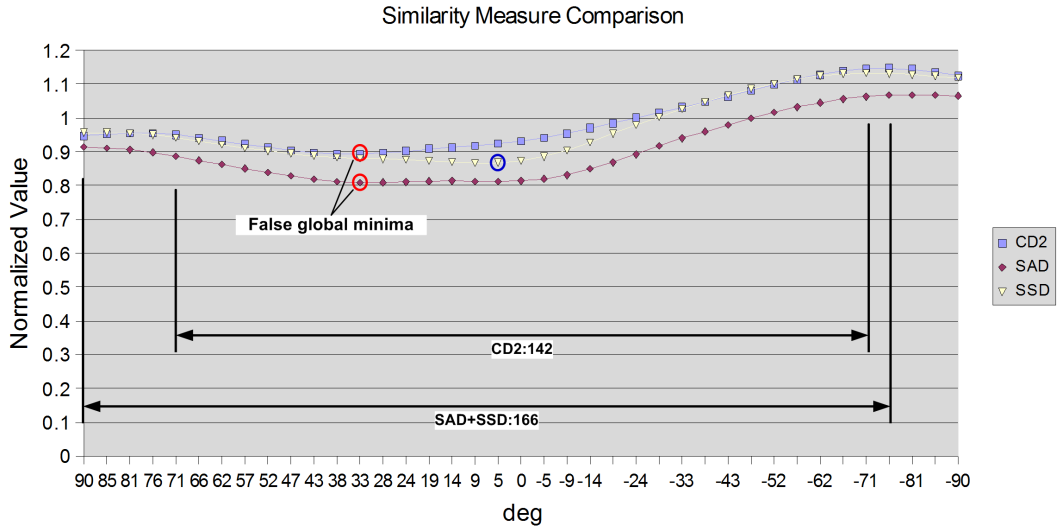

(a)

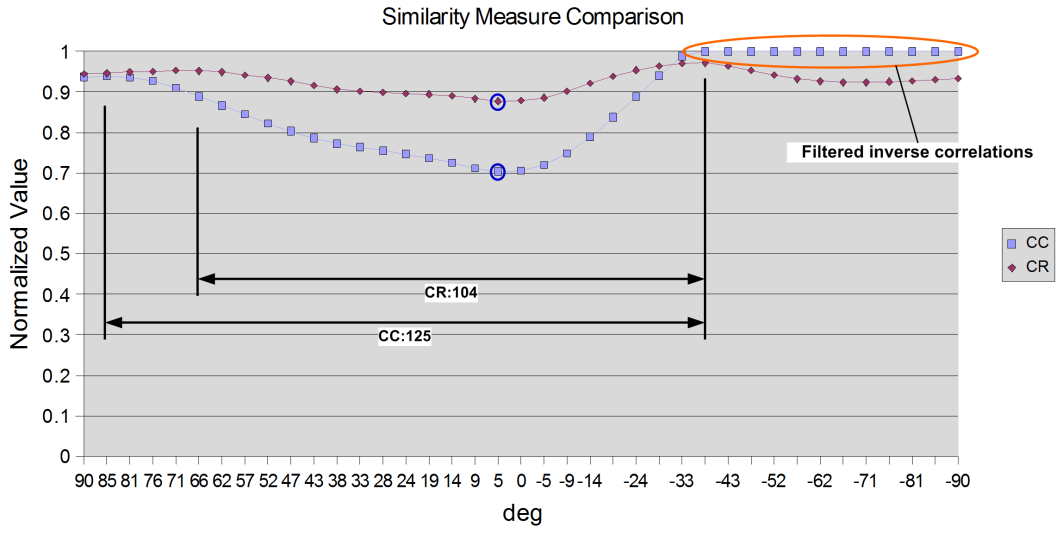

(b)

Figure 3.9: Similarity Measure Evaluation (inter-series comparison). Fig. (a) shows the graphs of $\mathcal{D}_{SAD}, \mathcal{D}_{SSD}$ and $\mathcal{D}_{CD2}$ for a -90° to +90° rotation for two different volumes. Alignment of point zero was carried out using rigid registration, and then a rigid offset of 7 mm, and a tilting of 10 degrees was added. Fig. (b) shows the graphs of $\mathcal{D}_{CC+}$ and $\mathcal{D}_{CR}$.

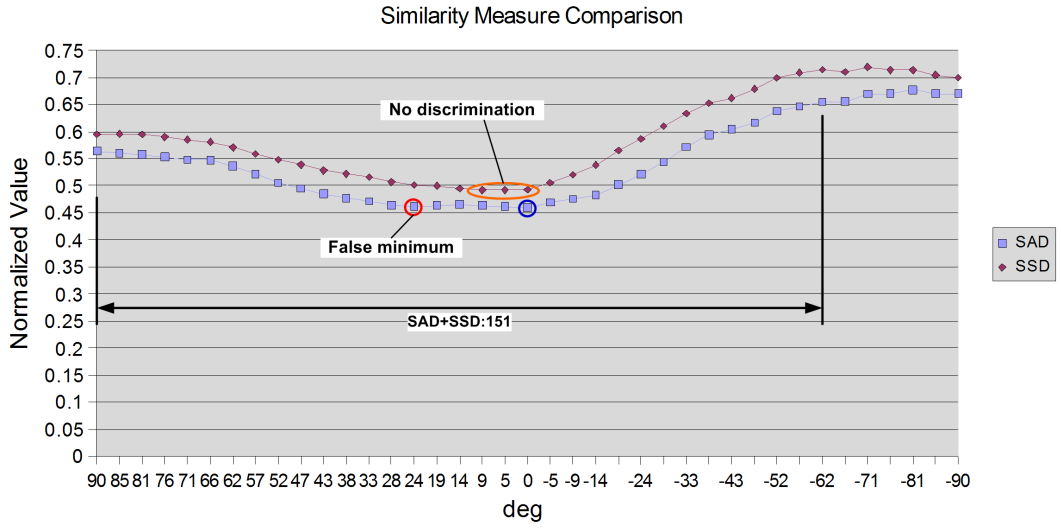

(a)

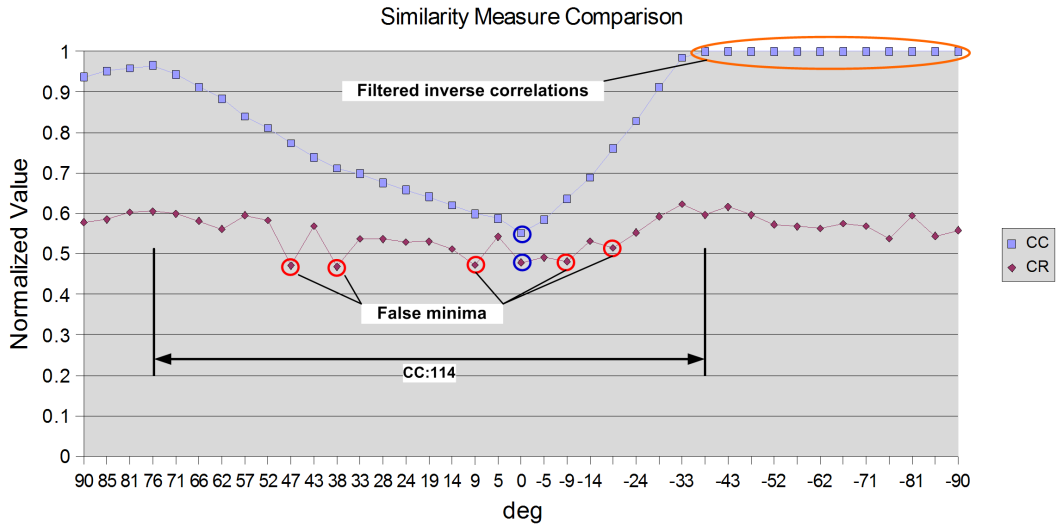

(b)

Figure 3.10: Similarity Measure Evaluation (inter series, coarse level comparison). Fig. (a) shows the graphs of $\mathcal{D}_{SAD}$, $\mathcal{D}_{SSD}$ and $\mathcal{D}_{CD2}$ for a -90° to +90° rotation for two different volumes. Alignment of point zero was carried out using rigid registration, and then a rigid offset of 7 mm, and a tilting of 10 degrees was added. The function evaluations were performed on the fourth level of the resolution pyramid. Fig. (b) shows the graphs of $\mathcal{D}_{CC+}$ and $\mathcal{D}_{CR}$.

the global minimum. On the other hand, the capture range of both measures shrinked slightly due to a new local minimum at 81 degrees. Both measures still have problems at 24 degrees with a local minimum. The correlation ratio is broken again. The local minimum of the correlation coefficient from 90 to 76 degrees in the previous test has vanished due to inverse correlation detection. Finally, not a single false minimum is present anymore in the graph of the correlation coefficient.

3.4.6 Discussion and Conclusion

To summarize, in order to maximize and stabilize the capture range of the similarity measure on coarse resolutions, we opted to use the correlation coefficient in combination with inverse correlation elimination. Together with attribute vectors built from the intensity and the gradient magnitude images, the number of false local minima can be considerably reduced. This is particularly important for global search algorithms, which need to evaluate the similarity measure also at points far from $\hat{\varphi}$. Strong local minima at these points are frequent, and they can often have large capture ranges. When, by misfortune, the global search algorithm does not get close enough to $\hat{\varphi}$ during its exploration, it is possible that the wrong capture range will be chosen. This happens in particular when only a small part of the prostate is visible in the reference image. The correlation coefficient not only produces few local minima, but also has a more pronounced and steeper capture range near $\hat{\varphi}$ when compared to the other measures that we evaluated, probably due to the fact that the linear model better handles the intensity distribution modifications that may occur between acquisitions.

Hence, the final similarity measure is

$$\mathcal{D} = \mathcal{D}_{CC+}[R, T_\varphi] \cdot \mathcal{D}_{CC+}[||\nabla R||, ||\nabla T_\varphi||]. \tag{3.21}$$

We use a multiplicative combination of the similarity measures and hence leave the grounds of maximum likelihood estimation. In the maximum likelihood framework with a Gaussian noise assumption, the similarity measures would have to be summed. The justification of the multiplicative approach is rather intuitive than mathematical: we think that local minima are more efficiently eliminated by multiplication than by addition. In practice we achieved good results with this approach, however, we did not properly validate it against the maximum likelihood estimator. This should be carried out in the near future.

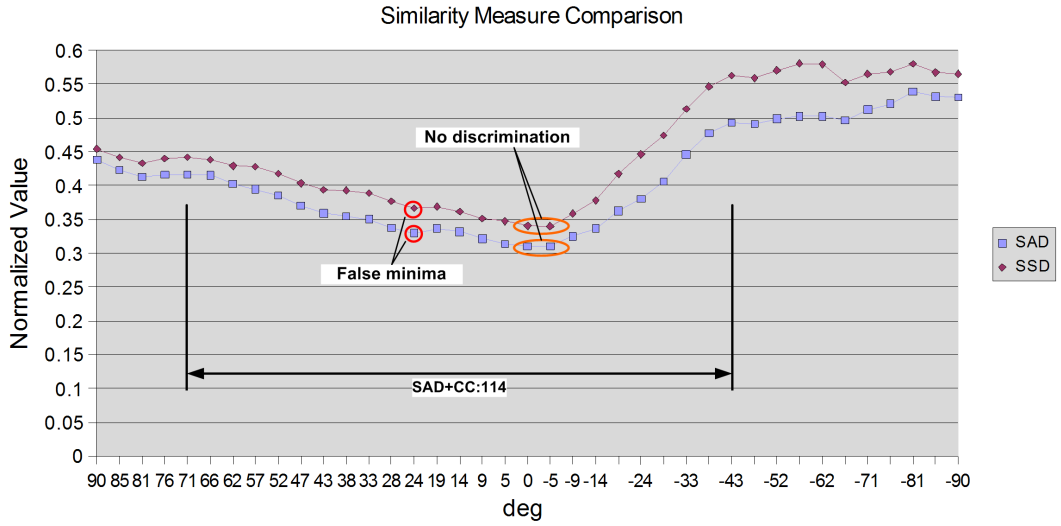

(a)

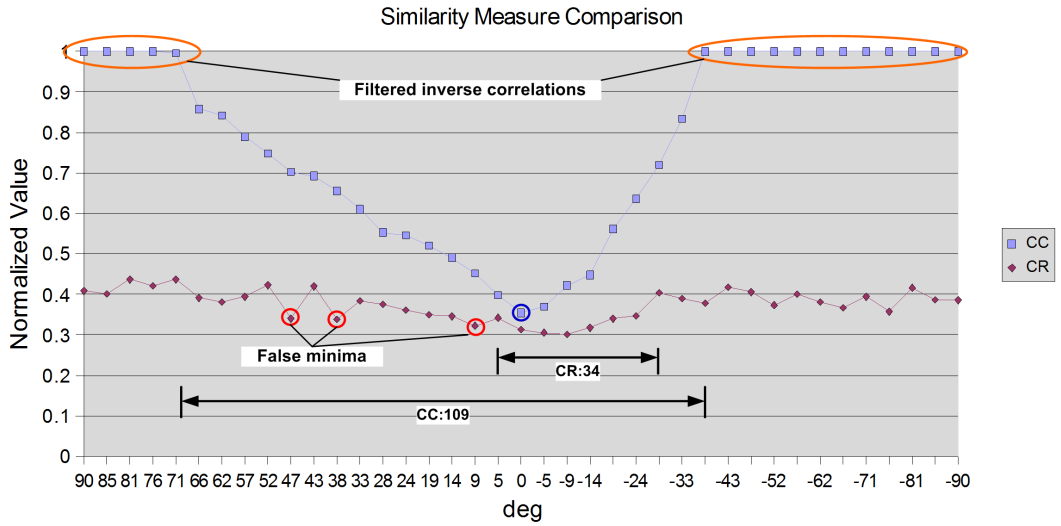

(b)

Figure 3.11: Similarity Measure Evaluation (inter-series, coarse level, intensity+gradient magnitude images). Fig. (a) shows the graphs of $\mathcal{D}_{SAD}, \mathcal{D}_{SSD}$ and $\mathcal{D}_{CD2}$ for a -90° to +90° rotation for two different volumes. Alignment of point zero was carried out using rigid registration, and then a rigid offset of 7 mm, and a tilting of 10 degrees was added. The function evaluations were performed on the fourth level of the resolution pyramid of the original images, and on the gradient magnitude images. Fig. (b) shows the graphs of $\mathcal{D}_{CC+}$ and $\mathcal{D}_{CR}$.

3.5 Prostate Tracking Specific Capture Range Improvements

After having chosen the similarity measure, we will eliminate several pit-falls that are specific to transrectal ultrasound based prostate tracking. First, the partial prostate visibility problem, due to the shape of the view cone of the used ultrasound probe, is solved using panorama anchor volumes. The similarity measure graphs will be further improved by filtering probe-related deformations, making advantage of a priori knowledge about the probe head position relative to the volume. Finally, most of the bladder will be removed from the volumes, since the very different intensities of liquids in the bladder strongly influence the similarity measure.

3.5.1 Panorama Anchor Volume

It is well known from signal processing that it is possible to identify patterns robustly even in very noisy signals. However, if the pattern itself is noisy or incomplete, automated detection can degrade considerably. The shape of the view cone of the ultrasound probe that we were using, a General Electric RIC5-9, resembles a pyramid with a rectangular base. Unfortunately, one side of this rectangle is so small that it is impossible to acquire an integral volume of a typical prostate, independent of the chosen ultrasound depths (see Fig. 3.12 (a)+(b)). Hence, large parts of the capsule of the gland is not visible. This problem is common to all currently available 3D ultrasound probes. Partial images can lead to registration failures in particular for lateral biopsy volumes, which contain often only 30 to 40 percent of the prostate. As a result, the overlap Ω of both images can get very small.

We therefore acquire three partial prostate volumes using the following protocol: the operator first acquires the anchor volume with the prostate being centered in the US beam. The ultrasound depth should be set such that all membranes are clearly visible, but it should not be too deep either, because the prostate would get too small in the volume. Then, two additional images with rotations of plus and minus 60 degrees around the principal axis of the probe are taken (see Fig. 3.12 (c) and see an example in Fig. 3.13), following the same guidelines as for anchor image acquisition. Care is taken to avoid deformation and US shadows. The two additional images contain the parts that are not visible in the anchor volume. Since all three volumes are acquired within a short time lapse and since all contain large parts of the prostate, it is relatively unproblematic to register and to fuse them into a panorama anchor volume. Image fusion is performed by intensity averaging.

Wachinger et al. presented recently a mosaicing algorithm based on multivariate formulation of the registration problem that makes it possible to register all

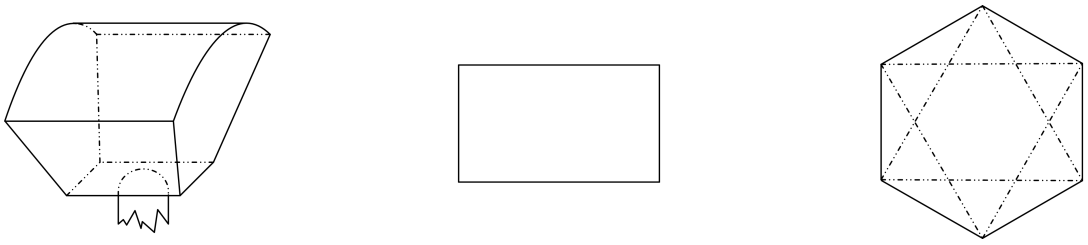

(a) (b) (c)

Figure 3.12: Panorama construction. Fig. (a) illustrates the view cone of a 3D acquisition, Fig. (b) illustrates the form of the beam in the plane orthogonal to the probe axis, and Fig. (c) illustrates the covered area in the same axial plane after fusion of the anchor image with the -60° and +60° images.

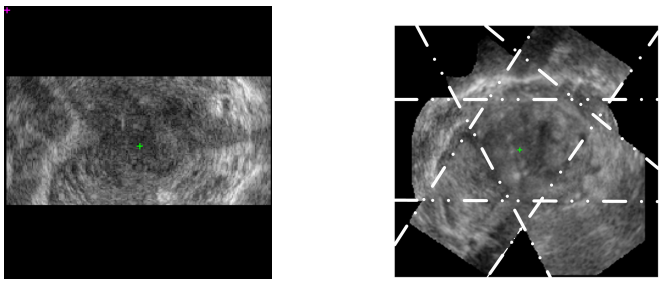

(a) (b)

Figure 3.13: Panorama construction. Fig. (a) shows an axial cut of the anchor image near the probe head. Fig. (b) gives an example of a panorama image.

images simultaneously [128]. This leads to a more consistent overall registration and improves the quality of the panorama anchor image. In our application, the overlaps of the images that form the panorama are relatively large, and in our tests it was rarely possible to identify discrepancies between the registrations of the "+60" and the "-60" images with the anchor image. Nevertheless, this is an interesting approach, and we are looking forward to investigate it in the near future.

3.5.2 Deformation Filtering

We stated in the introduction that the transrectal ultrasound prostate image registration problem is separable into a rigid and a non-rigid intensity-based registration part. We also stated that representing the registration problem on coarse grids of a Gaussian resolution pyramid averages small deformations out. It is possible to further decrease the dependence of the rigid registration algorithm on the amount

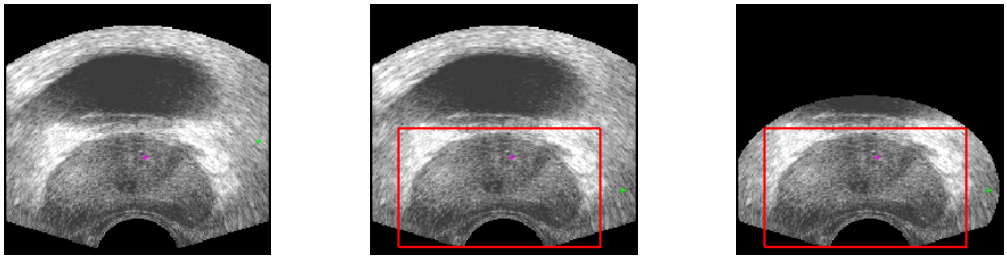

(a) (b) (c)

Figure 3.14: Bladder filtering. Fig. (a) shows a volume where the bladder is almost as large as the prostate. Fig. (b) shows the bounding box placement, and Fig. (c) finally illustrates the ellipsoidal cut performed to remove all unnecessary information.

of deformation present in the physical transformation $\hat{\varphi}$ by incorporating knowledge on the location of the main source of deformation, the probe head. Recall Fig. 3.1 that illustrates typical gland deformations that can be observed during biopsy acquisition. Note the squeezing of the prostate, and note also that the deformation is strongest near the probe head. Our solution consists in out-masking voxels that are closer to the probe head than a threshold of several millimeters. This method also allows to filter near-field effects. The threshold can be chosen in function of the prostate diameter in axial direction, which can be approximated with the bounding box. We made good experiences when using ten percent of the diameter. A similar approach was used by Blackall et al. in [18].

3.5.3 Bladder Filtering

The bladder is a strong attractor due to the fact that bladder intensities are very dissimilar to tissue intensities. If the bladder is filled has at least the same volume as the prostate, the similarity measure overweights the bladder in both the intensity and the magnitude of gradient image. For this reason we use the bounding box to cut out the prostate in the anchor image, thus reducing the influence not only of the bladder, but of all surrounding tissues. This is performed by cutting an ellipsoid defined by the bounding box plus a margin of 2^{n-1} voxel, where n is the number of pyramid levels used for registration, see also Fig. 3.14.

3.5.4 Discussion and Conclusion

In this section we have presented several problem-specific techniques that increase the convergence range of local optimization around the physical solution $\hat{\varphi}$, and which decrease the number of false minima in the graph of the similarity mea-

sure. Panorama anchor images increase considerably the stability of lateral biopsy volume registration. In a similar extent, bladder and deformation filtering considerably increases the robustness of the registration. In particular without bladder filtering, mis-registrations occur even for apparently "easy" volume pairs, because global optimization can get stuck in the capture ranges of strong bladder-related false local minima. To conclude, focusing on the registration target is a strategy that pays well in the context of endorectal prostate image tracking.

3.6 Local Parametric Optimization

We already mentioned in Sec. 3.1.9 that the most commonly used parametric optimization schemes do not differ significantly in terms of accuracy and computational burden. Following this observation, we have opted to use the Powell-Brent optimization scheme which has the advantage of not requiring derivatives of the cost function. This makes it easier to add complex similarity measures, for which derivation is not straightforward.

In this section we will adapt the Powell-Brent optimization scheme to fit our needs: first, we will define a termination criterion that makes more sense for registration than the original one that is defined on the machine precision. Second, we will adapt the Powell-Brent to work better with multi-resolution approaches. And finally we will discuss the δ-unimodality assumption which is the base of Brent's line search, and we will propose a small improvement that renders Brent's algorithm more robust. All discussions in this section are based on the implementation of the Powell-Brent algorithm discussed in [99].

3.6.1 Termination Criterion

There are two possibilities to define a termination criterion for a local optimization scheme: the first is based on the difference of the cost function $\Delta\mathcal{D}$ between two consecutive evaluations, while the second is based on the step size Δs carried out by the optimization algorithm between two consecutive position updates, where $\Delta s = s_{n-1} - s_n$, with $s_i \in R^d$, d corresponding to the number of unknowns of the optimization problem. The direction-set part of the Powell-Brent algorithm [24] uses $||\Delta s||$ to decide whether optimization is finished or not, and uses a threshold near the machine precision. The line-search part of the algorithm terminates when the step difference $\Delta t \in \mathbb{R}$ on the investigated line in the d-dimensional search space is smaller than a threshold near machine precision.

Thresholds derived from machine precision are clearly no sound choices in the context of image registration, where the achievable accuracy is determined by the image resolution, the voxel size and the noise present in the images. We have

chosen to base the termination criterion on a translational threshold δ_t and a rotational threshold δ_α, which is verified if

$$||A(\omega_{n-1}) - A(\omega_n)|| < \sqrt{4(1-\cos(\delta_\alpha))} \wedge ||o_{n-1} - o_n|| < \delta_t \tag{3.22}$$

where n is the current iteration. The matrix norm $||\cdot||$ is the Frobenius norm $||A|| = \sqrt{\sum_{i,j} a_{ij}^2}$. The Frobenius norm of the difference of two rotation matrices is related to the rotation angle α between both matrices by the identity $||A_1 - A_2|| = \sqrt{4(1-\cos(\alpha))}$, which can be shown[3] with the equalities $\mathrm{tr}(A_1 A_2^{-1}) = \mathrm{tr}(A_1 A_2^T) = 1 + 2\cos\alpha$ and $||A_1 - A_2||^2 = 6 - 2\mathrm{tr}(A_1 A_2^T)$.

The termination function T returns true, when both the rotational and translational changes are below the thresholds. This makes it possible to define the angular and rotational precision separately. Also, the precision can be set to reasonable values in function of the voxel sizes. The elegance of using the Frobenius norm is that it can also be computed for general affine matrices, which makes it possible to define an affine threshold as its angle equivalent. This is an interesting feature for affine registration systems. Note also that we apply this criterion for both the direction set and the line search parts of the Powell-Brent algorithm. Both thresholds are set to clinically meaningful values. In general we require a translational accuracy of 0.1 mm and a rotational accuracy of 0.25 degrees, which is largely sufficient for a prostate tracking system.

3.6.2 Powell-Brent and Multi-resolution

Maes et al. have shown with their local optimization algorithm comparison [79] that the Powell-Brent algorithm does not scale very well with multi-resolution approaches. When switching the resolutions outside the algorithm, the algorithm has to be restarted and needs to iterate at least twice over the entire direction set before the termination criterion can be checked for the first time. Changing the resolution in the Powell iteration of the algorithm is a solution to this problem. In our system, the resolution is switched to the next finer level every time when the termination criterion is satisfied on the coarse level. The algorithm exits when the termination criterion is fulfilled on the finest level. If the estimation does not change significantly on a given level, the resolution switch is performed immediately.

3.6.3 Brent's Invariant

A final improvement concerns Brent's line search. Brent's method examines a bounded section of a line in the search space using alternatively quadratic and

[3] I would like to thank my colleague Ziad Aramouni for the proof of this theorem.

linear approximations. The entire algorithm is built on the assumption that the objective function is δ-unimodal on the segment. Brent defines δ-unimodality as follows:

Definition 1. The objective function f is unimodal on $[x_1, x_2]$ if, for some (unique) $x^* \in [x_1, x_2]$, either f is strictly monotonic decreasing on $[x_1, x^*)$ and strictly monotonic increasing on $[x^*, x_2]$, or f is strictly monotonic decreasing on $[x_1, x^*]$ and strictly monotonic increasing on $(x^*, x_2]$.

Definition 2. Let I be an interval and f a real-valued function on I. We say that f is *strictly δ-monotonic increasing on I* if, for all $x_1, x_2 \in I$,

$$x_1 + \delta < x_2 \Rightarrow f(x_1) < f(x_2). \tag{3.23}$$

Strictly δ-monotonic decreasing functions are defined in the obvious way.

Definition 3. Let I be an interval and f a real-valued function on I. We say that f is *δ-unimodal on I* for a $\delta \in \mathbb{R}^+$ if, for all $x_1, x_2, x_3 \in I$,

$$\begin{aligned} x_1 + \delta < x_2 \wedge x_2 + \delta < x_3 \quad \Rightarrow \quad & (f(x_1) \leq f(x_2) \Rightarrow f(x_2) < f(x_3)) \\ & \wedge (f(x_2) \geq f(x_3) \Rightarrow f(x_1) > f(x_2)) \,. \end{aligned} \tag{3.24}$$

δ-unimodality is equivalent to unimodality for $\delta = 0$.

Brent introduced the concept of δ-modality to cope with the limits of machine precision. In his framework, δ is the minimum distance between points at which the objective function can be evaluated without getting stuck in a local minimum that is due to limited machine precision. Brent's algorithm assumes that δ is always set such that the invariant $f(x_2) < f(x_1)$ and $f(x_2) < f(x_3)$ always holds for $x_1 < x_2$ and $x_2 < x_3$, and an objective function f. We denote this invariant as "Brent's invariant".

In the context of image registration it is unfortunately almost impossible to find a δ for which all line-segments in the search space are δ-unimodal on the cost function without losing most of the accuracy that the registration process could achieve. Noise, interpolation artifacts, tissue changes between image acquisition and many other reasons are the reason why there will always be local minima with convergence ranges of unpredictable size. Unfortunately, despite of its reputation of being a particularly robust local optimization algorithm, the parabolic line search often has "hiccups" when Brent's invariant is violated, and optimization can even terminate with a position at which the similarity criterion is worse than at the start position. This is unacceptable in the context of a medical application, which is the reason why we decided to detect situations where Brent's invariant is not fulfilled.

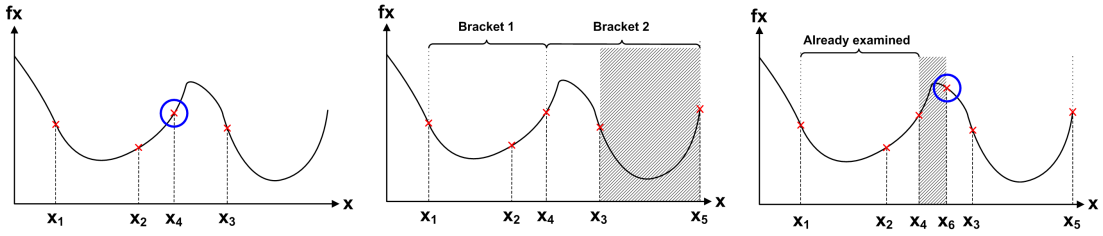

Figure 3.15: Brent's invariant. Fig. (a) shows the algorithm a violation of Brent's invariant. Normally the algorithm would ignore this violation and would continue with the bracket $[x_1, x_2, x_4]$, thus getting stuck in a local minimum. Fig. (b) illustrates the creation of a second bracket $[x_4, x_3, x_5]$. Fig. (c) illustrates a case where the algorithm has to take care of not creating a novel bracket that re-examines the bracket $[x_1, x_2, x_4]$.

A detailed discussion of all the necessary modifications would be long and technical, and would hence shift the focus of this work to a marginal problem. We will therefore only motivate the main ideas of the modification. First of all, at each update of the bracket $[x_1, x_2, x_3]$ that defines the currently examined line section, we check if Brent's invariant is violated for some new evaluation x_4, see Fig. 3.15 (a). If this is the case, we break the bracket into two by finding novel outer borders $x_5 < x_1$ if x_4 lies between x_1 and x_2, or $x_5 > x_3$ if x_4 lies between x_2 and x_3. Then, Brent's routine is called recursively on the first bracket, while optimization of the second bracket continues as usual (see Fig. 3.15 (b)). The bracket expansion algorithm has to ensure to find new borders on sections that have not yet been analyzed, it is therefore necessary to memorize previous function evaluations (see Fig. 3.15 (c)). Since the bracket algorithm will find a minimum on each split bracket, the minima also need to be memorized.

The modified Brent algorithm guarantees that the result of the optimization process is not worse than the start position. Also, many false local minima are detected and avoided. This improves convergence rate and the overall stability of the registration system. Also, less evaluations of the similarity measure are necessary.

3.6.4 Discussion and Conclusion

In this section we have adapted the Powell-Brent algorithm to the medical image registration problem. We also introduced a novel termination criterion that makes it possible to define explicitly the desired angular and positional accuracy. Having defined the rotational part with the Frobenius norm, it is possible to use the cri-

terion also for affine registration, in which case the threshold can be interpreted as an angle-equivalent distance. Integration of the resolution switch into the algorithm significantly reduces the number of iterations on fine levels, which makes the Powell-Brent algorithm competitive for multi-resolution approaches. Finally, verification of Brent's invariant guarantees that the result of the optimization process is not worse than the start position. Also, many small local minima are detected and avoided, which improves the convergence rate and the overall stability of the registration system.

3.7 Global Optimization

The role of global optimization in our approach is to find a point inside the capture range of the solution $\hat{\varphi}$. Due to the large transformation space and cropped prostate images in particular for lateral biopsies, registration failed for about 65 percent of the volumes that were acquired during biopsy when using local optimization alone. Unfortunately, global optimization on the entire six-dimensional search space, whether by evaluation of a regular grid of points in the search space or by more sophisticated methods like simulated annealing or genetic optimization, would increase computation time by several orders of magnitude. In this section we will discuss a fast global optimization method for registration of endorectal prostate images, based on a kinematic model of probe movements during transrectal ultrasound prostate image acquisition. Then, the global search approach will be rendered more robust by penalizing small volume overlaps during the optimization process.

3.7.1 Endorectal Probe Kinematics

The most innovative part of our rigid registration algorithm consists in a search space reduction based on a kinematic model of probe movements considering rectal and image formation constraints. The model holds for end-fire ultrasound probes, typically used for prostate biopsy acquisition. The main idea consists in combining pure image-based registration with a priori models of the concrete registration problem. When the physical transformation $\hat{\varphi}$ between the prostates in R and T is large, then it is necessary to find a point inside the capture range of the similarity measure before a local optimization can be tempted. This is not trivial: regular sampling of a 6-D rigid transformation space using a very sparse grid size of 10 already requires 10^6 function evaluations, which results in an unacceptable computational burden. The physical constraints exerted by the rectum on probe movements, and the fact that the probe head always remains in contact with the thin rectal wall at the prostate location lead to the following assumptions:

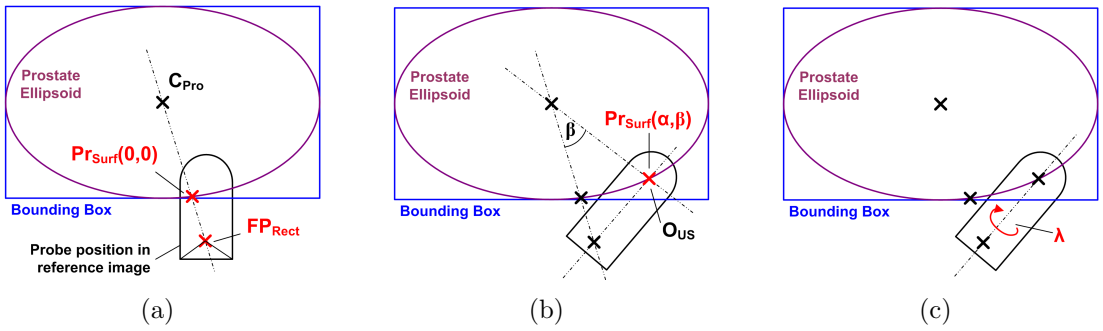

Figure 3.16: Mechanical probe movement model in 2D: (a) shows the computation of the search model surface origin $PR_{Surf}(0,0)$ from the prostate center C_{Pro} and the (hypothetical) rectal probe fix point FP_{Rect}. In (b), a 2D polar parameterization is used to determine a surface point $PR_{Surf}(\alpha,\beta)$. The probe is then rotated and translated such that its US origin O_{US} coincides with $PR_{Surf}(\alpha,\beta)$. In (c), the probe is rotated around its principal axis by an angle λ.

1. the probe head is always in contact with the prostate capsule,

2. the most important rotations occur around the principal axis of the probe, and

3. all other rotations have a rotation point that can be approximated by a unique fix point FP_{rect} in the rectum.

Note that the first point is always verified, else the captured volume does not contain image information of the prostate. The other two points depend on patient physiognomies and on the clinician.

With these assumptions it is possible to define a probe movement model based on a prostate surface approximation, the probe position in the US image (which is known) and a rotational fix point in the rectum. As shown in Fig. 3.16(a), the prostate surface is approximated by a bounding-box aligned ellipsoid. The ellipsoid is modeled using a 2D polar parameterization $PR_{Surf(\alpha,\beta)}$. The origin $PR_{Surf(0,0)}$ of the parameterization corresponds to the intersection of the line from the prostate center C_{Pro} to FP_{Rect}. As illustrated in Fig. 3.16(b), $PR_{Surf(\alpha,\beta)}$ implements assumption 1) by determining plausible US transducer positions on the prostate surface. Assumption 3) is satisfied by requiring that the principal probe axis must always pass through FP_{Rect}. Finally, a rotation about the principal probe axis implements assumption 2) and thus adds a third DOF (See Fig. 3.16(c)).

The 3D subspace defined by the probe movement model is systematically explored using equidistant steps. Since the exploration grid points do not change

during an intervention, it is possible to pre-compute and to store all reconstructions $T \circ \varphi(\omega, o)$ of the anchor image necessary for the evaluation of the intensity measure. Hence, the T_φ have to be computed only once, after construction of the panorama anchor image, which considerably accelerates the global search. The rotational space around the principal axis of the probe is unconstrained (360°), while tilting ranges are limited to the maximum value determined on test data, plus a security margin. The number of steps per dimension were also experimentally determined.

3.7.2 Statistical Significance of the Similarity Measure

A problem inherent to registration is the lack of statistical significance of the similarity measure when the overlap Ω gets small. This is accentuated on coarse resolution levels, where the number of grid points, i.e. statistical samples, is reduced by the factor 8^{h-1}, where h is the number of the resolution level. Unfortunately, the correlation coefficient is a good measure of the strength of a correlation, but it cannot be used to decide whether the correlation is significant or not. This is due to the fact that the coefficient is ignorant of the individual distributions of the intensities in both images (see also the discussion on linear correlation in [99]).

Hence, it is impossible to distinguish significant from non-significant results with our similarity measure. We therefore set the similarity measure to the maximum if the overlap is smaller than a pre-defined percentage of the minimum of the total number of grid-points in the reference image, and the total number of grid-points in the template image. The chosen threshold is 30 percent, which was experimentally derived by identifying analysis of the failure cases : we augmented the threshold until no failure caused by small overlaps was left. The chosen threshold is a conservative threshold, that ensures that almost all false local minima created by overly small overlaps are eliminated. On the other side, it is not any more possible to register images that contain less than approximately thirty percent of the prostate.

3.7.3 Discussion and Conclusion

A probe movement model was conceived that reduces the original six-dimensional global search problem into a three-dimensional problem, which can be explored with a reasonable number of evaluations. A future project consists in fitting the model to the patient data acquired during the last months. This can be carried out by optimization of an objective function that measures the average distance of the grid points evaluated during global search in function of the model parameters to the registration results obtained on the training set.

Moreover, a minimum volume overlap was defined to ensure the statistical significance of the similarity measure evaluations, hence eliminating large false local minima that could undermine global optimization. Similarity measures have been reported for which the statistical properties are known [66]. They could be used to decide whether a result is significant or not and thus replace the threshold, in order to allow for smaller overlaps.

3.8 Putting Everything Together

In this section we will put the different parts of the registration process together.

3.8.1 Pre-processing of the Anchor Volume and Caching

Fig. 3.17 illustrates the pre-processing that is performed on the anchor volume once before the first registration. First, the bladder is removed. Then, if non-linear registration is desired, an anchor pyramid for non-linear registration is constructed. In a next step, tissue near the probe head is removed and the anchor pyramid for rigid registration is constructed. Finally, we make advantage of the fact that the global grid search on the kinematic probe movement model evaluates always the same set of transformations $\{\varphi_i\}, i = 1, ..., N$. As long as the voxel side lengths and the volume resolutions of the tracking images do not change due to ultrasound depth modifications or ultrasound bounding box modifications, it is possible to pre-compute and to cache the images $T \circ \varphi(\Theta_i)$ in the tracking space Ω. When caching is used, it is possible to perform the global search without requiring tri-linear interpolation of the anchor images, which represents more than 90 percent of the computational burden.

3.8.2 Multi-step Registration

Fig. 3.18 also illustrates the registration chain which consists essentially of three steps: global registration on coarse levels, local registration on coarse levels, and either rigid or non-linear refinement on both coarse and finer levels. The choice of the finest level depends on the desired final accuracy and on the timing constraints.

3.9 Experiments and Results

3.9.1 Experimental Setup

Unfortunately, most currently available 3D ultrasound systems do not provide real-time access to volume data. They can, however, visualize two or three orthogonal

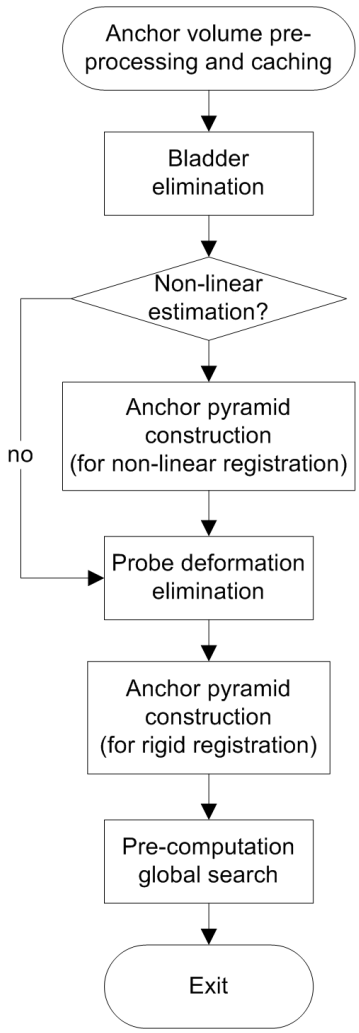

Figure 3.17: Pre-processing flow chart. The figure illustrates the different steps of the anchor image pre-processing and caching flow.

2D (O2D) slices inside the field of view of the probe in real-time, see also Fig. 3.19. These slices can be captured using a frame-grabber and used for registration with a previously acquired reference volume, that will define the tracking space [60,105]. Note that compared to 2D US images, O2D planes deliver significantly more spatial information, which potentially makes 3D to O2D registration more robust than 3D to 2D registration. In this study, we evaluated both 3D to 3D and 3D to O2D registration for rigid image-based tracking.

The rigid registration method was validated on 237 3D images of the prostate acquired during biopsy of 14 different patients. The imaging device was a General Electric 3D US Voluson 730 equipped with a volume-swept endorectal probe (GE

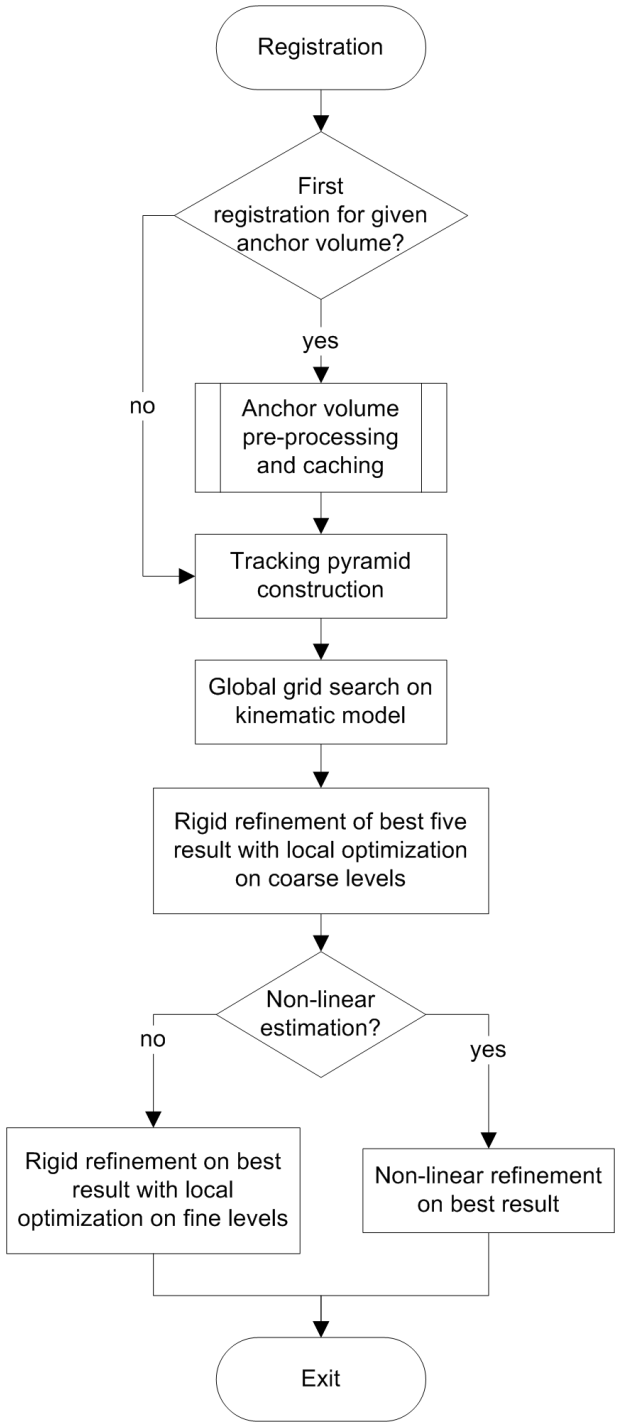

Figure 3.18: Registration flow chart. The figure illustrates the different steps of the registration flow.

RIC5-9). All images, except the images used for panorama image creation, have been acquired immediately after a needle biopsy acquisition. All registrations have been carried out in a post-processing step. The O2D images used in the tests are not frame-grabbed but reconstructed from 3D images. The image resolution is 199^3. The voxel side lengths vary from 0.33 mm to 0.47 mm. A five level resolution pyramid was used for volume to volume registration, whereas for volume to O2D registration only four levels were used, taking into account the smaller amount of available voxels. The final rigid multi-level search is carried out from the coarsest to the third-finest level for volume to volume, and to the second-finest level for volume to O2D registration. A total of 12960 grid points on the endorectal probe kinematics model were explored for each registration. The test was carried out on a Pentium 4 with 3GHz.

To measure reproducibility and registration success, 10 registrations have been carried out for each volume pair from slightly different initial points by adding noise of 2 mm and 2 degrees. This yields 10 transformations T_i that approximate the unknown rigid transformation between the prostate in both volumes. The average transformation $\overline{T}$ of the T_i has been computed with the method presented in [53]. The Euclidean distance error $\epsilon_E^i = ||T_i \cdot C - \overline{T} \cdot C||$, with C being the image center, and the angular error ϵ_A^i, which corresponds to the rotation angle of $T_i^{-1} \cdot \overline{T}$, are used to compute the root mean square (RMS) errors ϵ_E and ϵ_A. A registration is considered correct if $\epsilon_E < 2.0$ mm and $\epsilon_A < 5$ degrees, and if the result $\overline{T}$ is visually satisfactory when both volumes are superimposed in a composite image (See Fig. 3.20(c)).

Reconstruction accuracy evaluation was more difficult to implement since there is no straightforward gold standard. In some images, the needle trajectories from previous biopsies were still visible. In these cases, the trajectories were manually segmented, and the angular error between corresponding needle trajectories were used to evaluate rotational accuracy. Also, some patients had significant and clearly visible calcifications inside the prostate. The distances between segmented fiducials were used as target registration error for translational accuracy. This error measure is close to a gold standard, it is only biased by segmentation and ultrasound image formation related errors.

3.9.2 Results

The results of the 3D to 3D evaluation are given in Tab. 3.2, and the results of the 3D to O2D evaluation are given in Tab. 3.3. Fig. 3.20 illustrates a registered volume pair.

The overhead introduced by the systematic model-based exploration accounts for about 25% of 3D-3D, and for 35% of 3D-O2D registration time. The five optimizations on the coarsest level account for about 10% in 3D-3D, and for 20%

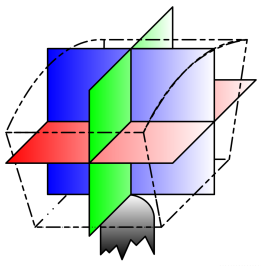

Figure 3.19: O2D slices. The Fig. illustrates three orthogonal slices that can be visualized in real-time by modern 3D ultrasound devices.

	result	pairs
Registration success	96.7%	237
Average computation time	6.5 s	237
Angular precision ϵ_A (reproducibility, RMS)	1.75 deg	229
Euclidean precision ϵ_E (reproducibility, RMS)	0.62 mm	229
Needle trajectory reconstruction (RMS)	4.72 deg	10
Needle trajectory reconstruction (max)	10.04 deg	10
Calcification reconstruction (RMS)	1.41 mm	189
Calcification reconstruction (max)	3.84 mm	189

Table 3.2: 3D-3D registration results. The third column indicates the number of registrations for which the given test could be evaluated.

	result	pairs
Registration success	87.7%	237
Average computation time	2.3s	237
Angular precision ϵ_A (reproducibility, RMS)	1.71deg	208
Euclidean precision ϵ_E (reproducibility, RMS)	0.47mm	208
Needle trajectory reconstruction (RMS)	4.74deg	9
Needle trajectory reconstruction (max)	10.5deg	9
Calcification reconstruction (RMS)	1.37mm	181
Calcification reconstruction (max)	4.30mm	181

Table 3.3: O2D-3D registration results. The third column indicates the number of registrations for which the given test could be evaluated.

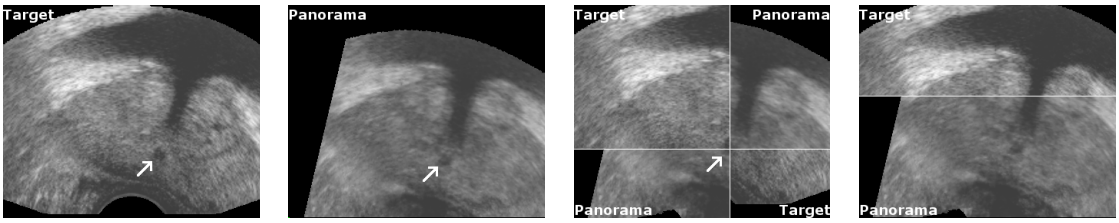

(a) (b) (c) (d)

Figure 3.20: Registration accuracy. (a) shows the target image, (b) the aligned panorama image, in (c) both volumes are superimposed to illustrate registration accuracy for the urethra (arrow), and (d) illustrates the registration accuracy in the upper gland.

in 3D-O2D. Panorama image pre-processing and pre-computation of the images for systematic exploration are performed before the intervention and require about one minute of computation.

3.9.3 Failure Case Study

A qualitative assessment of the failure cases was also performed. It consisted in a classification of the failure reasons. We created four classes that distinguish between failures due to

1. inadequate ultrasound depth, leading either to extremely small or only partially visible prostate images, see Fig. 3.21) (a),

2. anchor volumes of low quality either due to noise or because they contain only a fraction of the target, which lead for example to the registration failure of the tracking volume illustrated in Fig. 3.21 (b), despite its good visual quality,

3. noise, see Fig. 3.21 (c) and (d),

4. and tracking images that contain only small parts of the prostate.

785 registrations on a database of 926 volumes acquired on 47 patients have been analyzed, and 27 registrations have been identified as mis-registrations. Only volume to volume registration was evaluated. The result of the classification is given in Tab. 3.4.

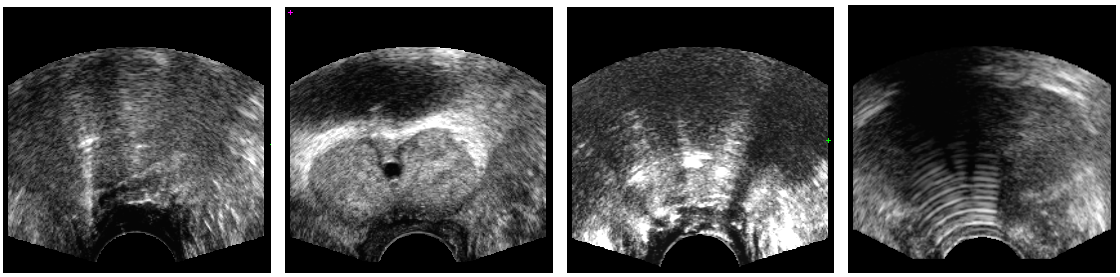

(a) (b) (c) (d)

Figure 3.21: Failure study. Registration of tracking volume (a) failed due to inadequate adjustment of the ultrasound depth. Tracking volume (b) could not be registered, despite its perfect quality, because the imaged part of the prostate was only partially present in the anchor volume. Tracking volume (c) could not be registered because almost no information was available in the volume, probably due to problems with the contact gel. Tracking volume (d) shows a case of partial probe contact, which also broke registration.

	Number of failures
Inadequate ultrasound depth	14
Low quality anchor volumes	4
Noise	3
Partial tracking images	6

Table 3.4: Failure study. Qualitative classification of the failure cases.

3.10 Discussion and Conclusion

We presented a fast and robust rigid registration framework for TRUS prostate images in a context of unconstrained patient movements, of only anatomy-constrained probe movements and of probe-induced prostate displacements. The algorithm yields reproducible results and acceptable accuracy for both 3D-3D and 3D-O2D registration.

The success-rate of 3D-3D registration is satisfactory, since all failures were either due to significant US shadows caused by only partial contact of the probe head with the rectal wall or by air bubbles in the US contact gel, or to an insufficient US depth with the result that the gland membrane is not visible on the images. In these cases the similarity measure fails because of missing information in the image, and there exists probably no algorithmic remedy. Additional failure causes can be observed for 3D-O2D registration, in particular for very small prostates, for

which the coronal plane does not contain prostate tissues. 3D-O2D registration is also more sensitive to low image quality, e.g. low contrast, too large deformations and to partial prostate images (for which often only one out of the three planes contains prostate tissues). Note that the presented algorithm is not very sensible to bounding box placement precision.

Computation time of local searches could be accelerated using the GPU for image reconstruction (which corresponds to approximatively 95% of the computational burden of a similarity measure evaluation), while further optimization of the systematic exploration would require parallelization of the evaluations.

The presented algorithm accurately registers in particular the prostate membranes that are distant to the probe head, and the urethra. The relatively high angular RMS error observed with the needle reconstruction study can be explained with probe-related local deformations that are particularly strong at the needle entry point.

3.10.1 Automatic Mis-registration Detection

A crucial part of a fully automatic image-based tracking system is its capacity to detect registration failures. Failures will always occur, most frequently they are caused by insufficient visibility of the target. Visibility problems include partial visibility of the target in the image, and noise-related visibility issues.

The fact that there exists no test for the significance of the function value of a parametric correlation based similarity measure is a serious disadvantage for the correlation coefficient measure. Its results strongly depend on the image pair to be registered, and it is very difficult to define a threshold that can separate registration results into the classes 'correctly registered' and 'mis-registered'.

A possible solution to this problem is to use distinct similarity measures for registration and for automatic validation. A similarity measure successfully employed for image registration with known statistical properties is for instance the increment sign correlation measure [66]. While we think that this measure is less appropriate for our registration problem, it seems to be a good candidate for the validation problem, since its significance can be measured with Student's distribution. Unfortunately, we were not yet able to investigate this promising approach. Remarkably few authors have investigated this problem, at least in the context of medical image registration.

Albeit being part of the main objectives of the prostate tracking project, automatic validation is hence still unsolved and requires further analysis. To conclude, the current version of the registration system relies on manual validation by the clinician or his assistant, either during the intervention or at a post-processing step.

3.10.2 Operator-dependency

Incomplete anchor volumes, either caused by the acquisition of inadequate volumes for panorama anchor image construction or by incorrect ultrasound depth settings, are the main causes of registration failures. Despite the fact that the clinicians have been instructed to acquire 3 images containing a maximum of the gland with rotations of 0, +60 and -60 degrees around the probe axis, we received only very few image sets for which these instructions have been properly respected. In consequence, the organ coverage of the panorama images is often not satisfactory. To reduce the number of panorama image related failures we are currently developing a software that guides the clinician during panorama acquisition, and a clinical protocol was developed that allows for rapid and reliable panorama acquisition.

The clinical protocol conceived for panorama anchor image acquisition is illustrated in Fig. 3.22. It comprises two validation steps: first, after registration of the volumes $+60^{\circ}$ and -60°, respectively, with the volume 0°, a check is performed whether the registration result corresponds to a rotation of $60 \pm \alpha$ degrees about the probe axis with respect to the anchor image. The threshold α for distinction between correct and incorrect acquisitions have still to be defined. They should represent a compromise between acquisition comfort for the clinician and quality requirements. Validation of the rotation angle and axis makes it possible to identify not only most of the operator-related acquisition errors, but also most of the panorama registration errors. A final visual validation is performed after the fusion of the three volumes. The software will be ready for our next biopsy campaign.

Issues can also arise with inadequately defined prostate bounding boxes. Fortunately, the algorithm is remarkably tolerant with respect to the bounding box settings. In practice we did not observe bounding-box related registration failures for boxes containing most of the prostate and for which the center of the box corresponds roughly to the prostate center. With the clinical protocol for panorama image acquisition, aberrant bounding-box settings are detected at the visual validation step, where the bounding-box based algorithm for bladder removal is already applied. Misplaced boxes lead to removal of prostatic tissues, which should be immediately detected by visual inspection.

The necessity of manual validation of the registrations introduces additional dependencies on the operator into the system. The principal draw-back of thorough manual validation is that it requires interaction with the interface of the tracking system. This is almost impossible during prostate biopsy acquisition, due to the fact that the clinician has to hold the probe and is hence constrained in his interactions with the system. This makes it necessary to allow for multiple acquisitions without immediate validation, and we believe that in practice, manual validation will most often be performed after the intervention. If the system deliv-

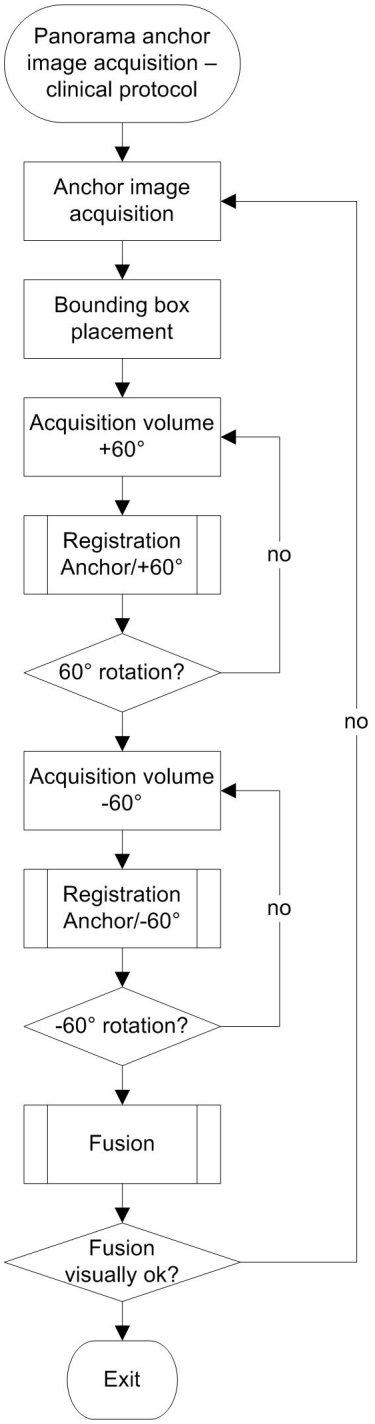

Figure 3.22: Clinical protocol. The clinical protocol for panorama anchor image acquisition.

ers intra-interventional feed-back it is hence necessary to communicate clearly that the performed registrations are not validated and that the tracking transformations could be wrong. It then depends on the common sense of the operator up to which degree he is confident about the information provided by the system. Concerning the validation task itself, the algorithm has fortunately the nice property of being either pretty accurate or completely wrong, i.e. the anatomical structures clearly do not correspond when registration fails. This makes visual validation relatively fail-proof and rapid. All but one of the 27 mis-registrations that have been observed during our patient study were immediately identified based on a visualization of the super-imposition of the registered volumes on three orthogonal planes. One mis-registration was only identified after fiducial-segmentation, where we noticed that the distance between the fiducials was of several centimeters. The transformation between $\hat{\varphi}$ and φ^* was a 180 degree rotation around the probe axis, and it was difficult for the operator to identify the mis-registration due to a surprisingly perfect symmetry of the concerned prostate.

While we think that the novel software will considerably reduce protocol-related failures with our novel acquisition software, the adequacy of the ultrasound depth settings cannot be easily verified. A possibility is to check the distance of the bounding box in the anchor image to the image side walls. This approach might however cause unintended behavior in the sense that the clinician might develop the habit of always defining the box a little bit smaller than required after having experienced several anchor image rejections. Our current approach is to visualize prostate image examples acquired with correct settings and examples acquired with incorrect settings. This measure also increases the transparency of registration failure causes for the clinician. In addition to ultrasound setting related failure case we also provide examples of noise that can cause registration failures, e.g. partial probe contacts.

3.10.3 Rigid Registration and Real-Time

Rigid registration is performed in several seconds on a 3Ghz Pentium 4 PC. Most of the computational burden consists of image convolutions. Considering the fact that graphic processing unit manufacturers have begun to use their skills to build highly parallel processing architectures for financial and scientific computing, we think that real-time in the sense of a tracking frequency of 20Hz is achievable. Recently, the first TFlop processing units have been presented at prices that easily compete with optical or magnetic tracking systems. This corresponds to a theoretical acceleration of approximately the factor 200, when compared to our test PC.

3.10.4 Global Search vs. External Tracking

Recall from Sec. 2.3.4 that Xu et al. [130] use a magnetic tracking system for their transrectal ultrasound based prostate biopsy guidance system. In this framework, the local registration algorithm is not initialized with the results of a global search, but with the current probe position delivered by the tracking system. This approach has several draw-backs: first, it is difficult to compensate for patient movements, due to the fact that it is neither straight-forward to detect, nor to measure them. If there is an undetected movement of the prostate of, say, 2 or 3 cm, the tracking system breaks. It is possible to fix the patient, but this would only considerably reduce patient comfort. We have tried to solve this problem with a second tracking reference fixed near the pubic arch of the patient, but the results were not satisfying.

It is, however, possible to use global search combined with external tracking. Global search can be used to correct for patient motion, while the tracking system will increase robustness in presence of target visibility problems. Note that introduction of a tracking system leads to additional issues like ultrasound probe calibration, instabilities due to magnetic interferences and cumbersome additional hardware in the operating room. In App. A we will present a calibration system for 3D ultrasound probes with respect to a tracking reference. We developed the calibration system before investigating the kinematic model.

We nevertheless think that global optimization alone will work out fine. A powerful improvement can be made in the context of real-time tracking: when the tracking frequency is high enough, the result of the previous registration can be injected into the registration process. More precisely, it has to be added to the list of transformations identified with the global search algorithm. This would considerably improve tracking robustness, and the clinician can "guide" the algorithm by slowly moving the probe towards problematic zones.

3.10.5 Conclusion

To conclude, we think that the presented rigid registration system is robust enough for computer-assisted medical interventions, at the condition that manual registration validation can be carried out either during the intervention or at a post-processing stage. To go beyond this limit it is necessary to find a method capable of validating the registration automatically.

Chapter 4

Elastic Registration

Abstract

In this chapter we propose a non-rigid registration framework with linear elastic regularization, image-based forces derived from the SSD similarity measure, probe forces derived from a priori models, and forces that ensure inverse-consistent estimates. Both the forward and the backward transformation are computed simultaneously. We present a novel approach to define the tissue compressibility parameter without impacting the convergence rate of the registration process. We propose also a line search scheme for image force computation to obtain more accurate estimates and faster convergence. The probe forces are derived from assumptions on the hypothetical undeformed tissue state and from the knowledge about the probe positions acquired with the rigid pre-registration. Integration of the probe forces can be interpreted as a fusion of a biomechanical simulation with the image-based registration process. Finally, a novel local intensity homogenization approach is proposed, which makes the SSD force term applicable to image pairs for which the identity assumption does not hold. The presented methods are all individually tested, and a final test is performed on [TODO]

4.1 Introduction

4.1.1 Problem Positioning

In the previous chapter we have described a fast, robust and accurate method to estimate the rigid transformation between two ultrasound volumes of the prostate. Starting from this linear approximation we now estimate the non-linear part of the transformation.

Recall that the prostate is mainly deformed by the head of the ultrasound probe, due to the fact that the clinician has to press the probe onto the gland to obtain a good contact of the transducer with the tissue. For rigid registration, a mask was used to eliminate tissues near the probe head. A principal objective of non-linear registration consists in the estimation of the deformation of these tissues. To a lesser extent, patient movements, bladder filling and breathing cause additional deformations, which is also addressed in this chapter.

For an exact modeling of these deformations we would need knowledge about the tissue properties of the prostate and the forces applied on the tissue. However, we do not know about a practical way to obtain these informations before or during intervention, which limits the possibilities to estimate deformations with a biomechanical tissue deformation model.

We can, however, 1) derive voxel displacements for identical structures in the ultrasound prostate volumes, we have 2) knowledge about the probe positions in both volumes, and we can assume that 3) the non-linear transformation is an elastic deformation. From the latter assumption follows that we are looking for a diffeomorphism, which means that the transformation is invertible, and that both the transformation and its inverse are differentiable. In other words, we are looking for a transformation that does not fold space and does not have discontinuousities. Note, however, that the assumption of pure elasticity is only valid as long as the two images to register show identical tissues, e.g. no tumor has grown between acquisitions and no surgical incision has been performed. In the context of prostate biopsies, the removal of tissue during the intervention and the resulting development of hematomas locally invalidate this hypothesis. These effects are fortunately local enough for the elasticity assumption to hold for strong probe and patient-movement caused deformations. When registering images of the same patient that are acquired with a considerable time lag e.g. images acquired during different biopsy series, the assumption is, however, more questionable.

The identification of elastic tissue deformations between two ultrasound acquisitions of the gland is far from being trivial. Non-linear transformations have several orders of magnitude more degrees of freedom than the global rigid transformation model, which, at first glance, seems to be a problem impossible to solve within a time frame compatible with the tracking objective. Worse, the non-linear

registration problem is ill-posed in the sense that it is almost always possible, and even trivial, to find mappings $\varphi : \mathbb{R}^3 \to \mathbb{R}^3$ from the reference image R to the template image T such that $R \equiv T \circ \varphi$. The transformation φ, however, is in general absolutely meaningless.

4.1.2 Parametric Approaches

Many authors extended the parametric linear registration approach 3.2 by introducing more complex parametric deformation models. Parametric models are built with a set of basis functions, the most widely used being bivariate polynomials, which include linear transformations, and radial basis functions.

Little et al. [78] and Fornefett et al. [48] propose radial basis functions. They define the transformation function $u : \mathbb{R}^3 \to \mathbb{R}^3$ is defined as

$$u(x) = \Phi_s(x) + R_s(x), \tag{4.1}$$

where Φ_s is a linear combination of monomials, and R_s is a linear combination of radial basis functions of the type

$$R_s(x) = \sum_{i=1}^{n} \alpha_i R(x, x_i), \tag{4.2}$$

where the x_i originate from landmark constraints $u(x_i) = q_i$, R is the chosen radial basis function, and α_i are scaling parameters.

A variety of radial basis functions have been proposed for non-linear image registration, see e.g. Bookstein et al. [20] or Evans et al. [45] for an example of thin plate splines, Little et al. [78] for multiquadrics and Arad et al. [4] for Gaussian basis functions. Rueckert et al. [104] build their parametric approach on B-splines of the form

$$u(x) = \sum_{j_1=0}^{3} \sum_{j_2=0}^{3} \sum_{j_3=0}^{3} B_{j_1}(y_1) B_{j_2}(y_2) B_{j_3}(y_3) \Phi(z_1, z_2, z_3), \tag{4.3}$$

where $y_i = \frac{x_i}{N_i} - \lfloor \frac{x_i}{N_i} \rfloor$, $z_i = \lfloor \frac{x_i}{N_i} \rfloor - 1 + j_i$ and the $N_i \in \mathbb{N}$ are the number of control points in direction i. In contrast to thin-plate splines, B-splines are locally controlled, which makes them computationally more efficient.

Blackall et al. propose a model-based regularization for registration of CT images of the liver with free-hand ultrasound images [18]. The model predicts the rigid and non-rigid transformations of the liver with respect to the respiratory cycle, and it interacts on the B-spline control points of the parametric registration framework. The model reduces the degrees of freedom of the non-rigid registration to a single parameter, i.e. to the position in the respiratory cycle.

Parametric registration is robust, flexible and can be used for a large variety of problems, but nevertheless lost a lot of its popularity in the medical image registration community in recent years. This is mainly due to the facts that the basis functions are more or less arbitrarily chosen and cannot provide a physically meaningful transformation. Moreover, optimization on parametric models can become inefficient, depending on the complexity of the chosen basis function. In the previous chapter we have seen that the latter point is already a big challenge when estimating only 6 DOFs. When using parametric motion models to reduce the dimensionality of the search space without further refinement, the registration quality depends entirely on how accurately the model can describe the organ motion, and hence on the intra- and inter-subject variability of the motion.

4.1.3 Non-parametric Approaches

Non-parametric registration methods model the mapping between two images as a function $\varphi : \mathbb{R}^d \to \mathbb{R}^d$, where d is the image dimension, which we assume to be 3 in the remainder of this chapter. The goal is to find a φ such that, for a reference image R and a template image T, $T \circ \varphi$ is similar to R. It is customary to decompose φ into the identity part $x = (x_1, x_2, x_3)$ and the displacement part $u : \mathbb{R}^3 \to \mathbb{R}^3$, i.e.

$$\varphi_u(x) := x + u(x), \tag{4.4}$$

with $u(x) = (u_1(x), u_2(x), u_3(x))$.

To increase the readability of this chapter, we further introduce the abbreviation

$$T_u(x) := T \circ \varphi_u(x) = T(x + u(x)). \tag{4.5}$$

Finally, we make use of the notation ∂_{x_i} for the partial differential operator $\frac{\partial}{\partial x_i}$, and $\partial_{x_i x_j}$ for $\frac{\partial^2}{\partial x_i x_j}$.

Taking advantage of the diffeomorphism assumption motivated in Sec. 4.1.1, we can further expect that the inverse φ_u^{-1} exists and that both φ_u and φ_u^{-1} are differentiable.

Given an image-based distance measures $\mathcal{D}[R, T; u] : \mathbb{R}^3 \to \mathbb{R}$ the fundamental non-parametric registration framework is

$$u^* = \arg\min_u \mathcal{D}[R, T; u]. \tag{4.6}$$

This is a variational formulation of the registration problem that can be analyzed with the calculus of variations.

Without further constraints it is not likely that the result of a registration based on Problem 4.6 will be satisfying, since far too many false local minima are present in typical medical images. It is therefore necessary to add an a priori model of the

displacement in order to eliminate most of the obviously false solution. Joshi et al. view in [63, 65] the feature registration problem as a Dirichlet problem, i.e.

$$u^* = \arg\min_u \mathcal{R}[u] \tag{4.7}$$

with

$$\mathcal{R}[u] = \int_\Omega \|\mathcal{L}u(x)\|^2 \, dx, \tag{4.8}$$

which is subject to the explicit constraints $u(p_i) = q_i - p_i$ for $i = 1, ..., M$, where the q_i and p_i represent the features detected respectively in both images. The operator $\mathcal{L}$ denotes a symmetric linear differential operator and is used to interpolate the displacement field at positions where no features are available.

Christensen et al. combine in [34, 36] multiple cost functions with the regularization term to obtain a framework of the form

$$u^* = \arg\min_u \Big(\sum_i \alpha_i \mathcal{D}_i[R, T; u] + \mathcal{R}[u] \Big), \tag{4.9}$$

They define the $\mathcal{D}_i$ both on landmarks and voxel-based image similarity measures.

Cachier et al. and Vercauteren et al. propose in [29, 124] a slightly different registration framework, which uses two distinct deformation fields, u and v, for regularization and optimization on the similarity measure, respectively, and couple them with an additional distance term:

$$u^* = \arg\min_{u,v} \left(\alpha_D \mathcal{D}[R, T; u] + \alpha_E \mathcal{E}[u, v] + \alpha_R \mathcal{R}[v] \right), \tag{4.10}$$

where the minimization process is implemented with an iterative scheme that alternates between the minimization of the terms that depend on u and the terms that depend on v. The motivation of this approach is the observation that most known regularization models are no perfect description of the physical deformation process, e.g. the material parameters of imaged bodies can vary and most image modalities are noisy. Using a coupling term, regularization is weaker and more easily allows for local deviations from the model. In [29], the deviations of the regularization process from the real physical process are modeled as Gaussian noise by defining the coupling term as $\mathcal{E} = ||u - v||^2_{L_2}$.

4.1.4 Inverse-consistent Frameworks

Computer-assisted medical interventions often require both the forward and backward registration from the reference to a template image. In the context of prostate biopsies, a primary objective is to project the biopsy needles into the anchor volume in order to allow biopsy distribution analysis. In the context of computer-guidance,

it is necessary to project e.g. a planning or previously biopsied sites in real-time into the tracking volume, which makes the computation of the inverse projection necessary.

Computation of the inverse transformation from u after registration represents almost the same computational burden as the registration process itself. When considering that the inverse registration problem that results from swapping the reference and the template image consists in the estimation of the problem from the opposite direction, we observe that the inverse registration, also called backward registration, does not produce the same intermediate and final results than the forward registration process. The asymmetry of the forward and backward registration leads to inconsistent forward and backward estimates in the sense that $u \circ u^{-1}$ do in general not correspond to the identity transformation. The advantage of the asymmetry is, however, that it can be exploited to improve the registration process by combining both estimations.

Christensen and Cachier propose in [30,32] to introduce an inverse consistency constraint with the cost function

$$\mathcal{D}_{cons}[R, T; u] := \mathcal{D}[R, T; u] + \mathcal{D}[T, R; u^{-1}]. \tag{4.11}$$

This approach combines the distance measures of the forward registration problem and the backward registration problem. The inverse u^{-1} is computed directly from u in this approach.

Rogelj et al. present in [102] a technique that they call symmetric inverse computation, which consists in computing the forces $\bar{f}_f$ of the forward registration problem and the forces $\bar{f}_b$ of the backward registration problem simultaneously. They define the novel forces $f_f = \frac{1}{2}(\bar{f}_f - \bar{f}_b)$ and $f_b = -f_f$ for the forward and backward registration, respectively. This approach is justified with Newton's third law of motion, which states that for every action, there is an equal and opposite reaction. This approach can be useful in the presence of small deformations. In the presence of larger deformations, the time-inverse character of the forward and the backward registration processes has the consequence that the forces are added at the "wrong moment" and in consequence to the wrong grid point. This can lead to slow registration convergence or even mis-registrations.

Both presented approaches are blind to the discretization errors present in the field approximations u and u^{-1}, i.e. when $u^n \circ (u^n)^{-1}$ does not correspond exactly to the identity at iteration n, the error will never be corrected. Rogelj et al. reduces this effect using more precise estimates based on the local Jacobian of the displacement fields when computing the force terms.

In [33], Christensen allows the forward transformation u and the backward transformation v to be independent by adding two extra terms to the objective

function, i.e.

$$u^*, v^* = \arg\min_{u,v}\Big(\mathcal{D}[R,T;u] + \mathcal{D}[T,R;v] + \|u - v^{-1}\|^2_{L_2} + \|u^{-1} - v\|^2_{L_2}\Big). \quad (4.12)$$

Christensen solves this problem with an alternating optimization of u and v, while Leow et al. optimize the problem directly [75]. Again, this formulation of the problem requires the computation of u^{-1} and v^{-1} at each iteration. The advantage of this approach is that the inverse consistency forces are only indirectly coupled, which reduces the risk of mis-registration due to the fact that forward and backward registrations are time-inverse.

Zhang at al. propose in [132] also a decoupled formulation of the inverse consistency constraint, based on the composition of the forward and the backward transformation. They further separate the forward and backward objective functions, i.e.

$$\begin{aligned} u^* &= \arg\min_{u}\Big(\|\varphi_u \circ \varphi_v - Id\| + \sum_i \alpha_i \mathcal{D}_i[R,T;u] + \mathcal{R}[u]\Big) \\ v^* &= \arg\min_{v}\Big(\|\varphi_v \circ \varphi_u - Id\| + \sum_i \alpha_i \mathcal{D}_i[T,R;v] + \mathcal{R}[v]\Big). \end{aligned} \quad (4.13)$$

The forward and the backward transformation are alternately optimized in an iterative process. The advantage of this approach is that the inverse consistency cost $\|u \circ v - Id\|$ can correct discretization errors from previous iterations, and that it does not require the computation of u^{-1} and v^{-1}. Note that this approach cannot, however, ensure the diffeomorphism property on the estimated fields.

A model capable of handling time-inverse flows was presented by Avants et al. in [6] that is defined on a forward flow $\Phi^1_{0,t}$ and a backward flow $\Phi^2_{1,t}$, where the time indexing t of Φ^1 is opposite to the indexing of Φ^2. The cost function is proposed to be

$$\mathcal{D}_{cons} = \int_\Omega \int_0^{0.5} \|u\|^2_{L_2} + \|v\|^2_{L_2} + \lambda \|R_{\Phi^1_{t,0}} - T_{\Phi^2_{1-t,1}}\|^2_{L^2} dt\ dx. \quad (4.14)$$

This approach can guarantee that the inverse constraint is sub-voxel accurate everywhere, which cannot be guaranteed with Christensen's or Zhang's approach. Flow integration, however, increases registration time by an order of magnitude, which renders it prohibitive for a tracking system.

4.1.5 Regularization

In this section we present the most popular regularization models for non-linear registration, beginning with the linear elasticity tensor, first applied to image registration by Broit in [27] and introduced into the medical image processing domain

by Bajcsy and Kovacic in [7]. It is derived from the linearized elastic potential of the displacement u [86]

$$\mathcal{R}_E[u] = \int_\Omega \frac{\mu}{4} \sum_{j,k=1}^{3} \left(\partial_{x_j} u_k + \partial_{x_k} u_j\right)^2 + \frac{\lambda}{2} (\text{div } u)^2 \, dx. \tag{4.15}$$

where λ and μ are the Lamé coefficients. The advantage of the linear elasticity model is that it provides a well-defined physical approximation of small elastic deformations of tissues in the human body. It is, however, not appropriate in the presence of strong, highly non-linear elastic deformations. Also, the constants λ and μ have to be parameterized in function of u when the material properties vary significantly for the bodies present in the images to register. Even for uniform materials the choice of λ and μ can be challenging since they are often unknown, and since the forces used for image registrations are estimations of image element displacements, and not physical forces. Note that this last point is true for all physically motivated regularization approaches.

Christensen [35] uses the fluid continuum as a regularization model. This model is linked to the linearized elastic potential by the time derivative $\mathcal{R}_F[u] = \mathcal{R}_E[\partial_t u]$. In other words, the regularization is rather carried out on the velocity field $\partial_t u$ than on the displacement field u. It follows that fluid regularization looses its strength when u approaches a steady state solution, which makes this approach more flexible and less constraining than the linear elasticity model. Bro-Nielsen [26] was the first to provide a fast solution scheme for the fluid model by exploiting the linearity of its partial derivative equations. In principle, it is possible to deform any template image to any reference image with fluid registration, at the condition that the gray-level scales of both images are identical: Modersitzki demonstrates in [86] the registration of a radiography of a hand with a black disk under fluid regularization. Hence, the fluid model is most appropriate when registering viscous bodies like liquids, but it is a questionable physical model for soft tissues in the human body.

Another class of regularizers is the diffusion scheme, which is based on the regularization term

$$\mathcal{R}_D[u] = \tfrac{1}{2} \sum_{i=1}^{3} \int_\Omega \|\nabla u_i\|_{L_2}^2 dx \tag{4.16}$$

cf. for instance [3, 47, 118]. The motivation for diffusion registration is to privilege smooth deformation fields, since smoothness is a powerful a priori assumption on the displacement field which is valid for many clinical applications. It has been shown that the widely used algorithm of Thirion [119] can be interpreted as a diffusion scheme [86, 124].

Another widely used regularization model is the thin plate model, first investigated by Duchon in [42] and Bookstein in [20]. It is based on the second order

regularization

$$\begin{aligned}\mathcal{R}_{TPS}[u] &= \int_\Omega \Big((\partial_{x_1x_1}u)^2 + (\partial_{x_2x_2}u)^2 + (\partial_{x_3x_3}u)^2 \\ &+ 2\left(\partial_{x_1x_2}u\right)^2 + 2\left(\partial_{x_1x_3}u\right)^2 + 2\left(\partial_{x_2x_3}u\right)^2 \Big)\, dx. \end{aligned} \tag{4.17}$$

The thin plate regularization models the bending of a thin plate of metal and yields good approximations to the bending behavior of membranes.

Until now the presented regularizers have been derived from physical models. It is also possible to follow a Bayesian approach for deformation field regularization, where the optimal field may be derived on basis of the posterior probability

$$p(u|R,T) = \frac{p(R,T|u)p(u)}{p(R,T)} \tag{4.18}$$

and the regularization energy

$$\mathcal{R}_B[u] = \int_\Omega -\log p(R,T|u)(x) - \log p(u)(x)dx. \tag{4.19}$$

The probability density functions are retrieved from a statistical atlas that has to be build from a representative set of learning data. Bayesian models are in particular interesting when the image-based distance terms $\mathcal{D}_i$ are only weak measures of the structural image identity, which is often the case e.g. in the context of registration of images acquired with very dissimilar modalities. For these applications, a probability density function can incorporate significantly more detailed a priori information into the registration process than the physical models discussed earlier. The registration quality, however, heavily depends on the quality of the statistical atlas. An example for statistical atlas construction and registration for the prostate can be found in [108].

4.1.6 Image-based Forces

We now have discussed registration frameworks and regularization models, but the most important part is still missing - the forces that move the image elements under the control of the regularization model. Unfortunately, we have to leave physical grounds from now on, since it is impossible to derive the forces acting on a system at image acquisition from the image. What we can estimate, however, are displacements between image elements, identified with a similarity measure.

Similarity measures have already been discussed in Sec. 3.1.8, namely sums of squared distances, normalized mutual information, Pearson correlation coefficient and the correlation ratio. In contrast to affine registration, where a global similarity criterion was required, we now need a local criterion that can be applied

to each image element. A global similarity measure $\mathcal{D}(R, T \circ u)$ can be localized around a point x by restricting its computation to a neighborhood $\Omega(x)$ of x. The simplest neighborhood is the block neighborhood, i.e.

$$\Omega^r_{Block}(x) = \{y \in \mathbb{R}^3 | \; ||x - y||_\infty \leq r\} \tag{4.20}$$

where $r \in \mathbb{R}$ is the radius of the neighborhood, and $|| \cdot ||_\infty$ is the maximum norm. A mathematically more accurate neighborhood in the sphere neighborhood, which is defined on the Euclidean norm, i.e.

$$\Omega^r_{Sphere}(x) = \{y \in \mathbb{R}^3 | \; ||x - y||_{L_2} \leq r\} \tag{4.21}$$

It is also possible to weight the contribution of each element to the energy in function of its distance to x. The interested reader can for instance easily integrate window functions in the framework developed by Valadez in [122], which is a little bit more general than ours.

What matters to us is in particular the window size required for a given similarity measure to produce statistically significant results. We have seen in Sec. 3.4.5 that the correlation ratio does not deliver stable results on coarse levels, even when computed on the *entire* image. The Pearson correlation that we used as a global measure for rigid registration is a rather poor statistical measure for deciding whether an observed correlation is significant, since it is ignorant of the individual intensity distributions in the images R and T. In other words, we either have to choose rather large values for r with, say, $|\Omega^r| > 1000$ voxel. This, however, would drag down registration performance to an unacceptable level for an application that requires intra-interventional registration results. The alternative would be to accept that the Pearson correlation computed on a small neighborhood does not always yield correct results, paying the performance improvement with an additional source of noise. Another solution is the block matching algorithm, where the transformation kernel is rather modeled as a rigid or affine transformation than a simple displacement vector, and for which the richer local transformation model allows for a sparser evaluation of the energy term on the grid, thus leaving more computation time for the convolution of the energy function with the image, see e.g. [73]. Block matching is, however, still not competitive in terms of computational efficiency, and interpolation of the transformation at locations between evaluated grid points can be a subtle issue.

Identical problems arise for other more sophisticated similarity measures, e.g. normalized mutual information. These considerations led us quickly to the decision to use the widely used and rather simple sums of squared distances $\mathcal{S}_{SSD}[R, T; u]$ as similarity measure. Recall that this measure evaluates the identity assumption $R \equiv T$, which is not quite verified for the transrectal ultrasound prostate image registration problem, but we save this discussion for Sec. 4.7. The advantage of

SSD is that we can choose the window size as small as 0.5, which means that only a single pair of image elements is considered after discretization. Note that even for this smallest possible neighborhood SSD delivers *always* the correct answer to the question whether the image is identical at a given point or not.

The distance $\mathcal{D}$ in function of R, T and u is defined as follows

$$\mathcal{D}_{SSD}[R, T; u] := \alpha_{SSD} \tfrac{1}{2} \int_{\Psi} \|R(x) - T_u(x)\|^2_{\mathbb{R}^3} \, dx, \tag{4.22}$$

where $\Psi \subset \mathbb{R}^3$ and $\alpha_{SSD} \in \mathbb{R}$ is a scaling parameter that is necessary in the framework 4.9 to balance the distance measure with respect to the regularization.

4.1.7 Feature-based Forces

The cost functions $\mathcal{D}_i$ in framework 4.9 can also be a combination of image-based and feature-based distance measures, see e.g. Christensen et al. in [36], or purely feature-based distances. In general, the feature-based forces are defined as the Euclidean distances between feature pairs. Recall from Sec. 3.1.7 that feature-based methods require a detection algorithm capable of identifying feature pairs when fully automated registration is required. However, in particular in the context ultrasound image registration, manual segmentation is still prevalent in most works. The automatic attribute-vector based solution of Foroughi et al. [50] for elastic registration of 3D liver ultrasound images has already been discussed in Sec. 3.1.7 and Sec. 3.1.8.

An original approach is proposed by Penney et al. in [95] for the registration of CT images of the liver with freehand 3D ultrasound images. In both modalities, a statistical approach is used to determine the probability of the presence of a feature at a given point. The driving forces are derived from the probability density maps of the CT and the ultrasound image. The probability distribution function was computed on a training data set, and is defined on the image intensities for the CT data, and on the intensity plus the intensity changes in ultrasound scan line direction for the ultrasound images. This approach can be viewed as a statistical segmentation process. The probabilities maps are interpreted as images and registered with the SSD distance measure. Unfortunately, it is difficult to use this approach on transrectal ultrasound images of the prostate, since they do not contain easily identifiable features like the hepatic vessels segmented by Penney.

4.1.8 Deformable Prostate Registration

In this section a little overview is given on the registration approaches proposed for non-linear prostate registration. Bharatha et al. present in [15] a segmentation-based non-linear registration system for pretreatment MR prostate images and

intraoperative MR imaging during brachytherapy. The shape of the prostate and its substructure are manually segmented and represented in a tetrahedral finite element mesh. The forces were computed from the mesh distances, and the linear elastic model was used as regularizer.

For estimation of endorectal probe-related prostate deformations, Alterovitz et al. discuss in [2] a method for registration of MR prostate images and spectroscopic MR images acquired with an endorectal probe to MR images acquired during radiotherapy with a surface probe. Here, large parts of the image dissimilarities are due to an object that is present in one image and absent in the second. Alterovitz proposes a finite element method to estimate the deformation of the prostate and surrounding tissues due to displacements and forces resulting from the endorectal probe. The finite element mesh is obtained with a Delaunay triangularization from a manual segmentation of the prostate, the rectum and bones. Here, the forces that drive the registration are entirely model-based and depend on the accuracy of the tissue parameters used during estimation. The parameters of the model are optimized during the iterative deformation estimation on a criterion that depends on the overlap of the segmentation of the prostate.

Zhan et al. use the shape of the prostate to define landmarks in histological prostate samples and in MR images of the same prostate to perform registration [131]. "Blob-like" structures in the inner gland are used to define additional landmarks. The forces are derived from the landmark distance and the mutual information similarity measure. The regularization model is not explicitly specified, the authors say that thin-plate splines or positive definite kernels can be used.

Shao et al. propose a deformable registration method for fusion of MRI/MRSI information with transrectal ultrasound images [106]. The approach relies on the manual segmentation of the MR images. The forces are derived from the gradient image of the ultrasound image after regularization multiplied with the surface normal of the MRI segmentation after regularization, i.e.

$$f(T_g) = \frac{1}{N} \sum_{k=1}^{N} \langle G\left(T_g(x_k)\right) \cdot T_g(n_k) \rangle_{L_2}\,, \tag{4.23}$$

where the T_g is the transformation parameter, x_k are the coordinate of the k-th sampled point on the surface, n_k is the surface normal at this point, and G is the gradient image. The thin-plate splines model was used for regularization in a parametric formulation.

Lian et al. propose a non-linear registration method for coil-based MR/MRS images with CT images [76]. Again, the thin plate splines model was used in a parametric formulation for regularization. The control points were set on expert-segmented contours.

The presented deformable prostate registration methods rely all on organ segmentation for elastic prostate registration. Automated organ segmentation is difficult to implement on ultrasound images, which is the reason why most authors use manual segmentation approaches. Manual segmentation, however, cannot be used in a tracking system for computer-assisted medical interventions.

4.1.9 Outline of this Chapter

The registration framework that we discuss in this work incorporates image-based forces, probe-related forces and inverse-consistency forces, and estimates the forward and the backward transformation of the registration problem simultaneously. It is of the form

$$\begin{aligned} u^* &= \underset{u}{\operatorname{argmin}} \left(\mathcal{R}[u] + \mathcal{D}_{SSD}[R, T; u] + \mathcal{D}_{Probe}[R, T; u] + \mathcal{D}_{Cons}[u, v] \right) \\ v^* &= \underset{v}{\operatorname{argmin}} \left(\mathcal{R}[v] + \mathcal{D}_{SSD}[T, R; v] + \mathcal{D}_{Probe}[T, R; v] + \mathcal{D}_{Cons}[v, u] \right) \end{aligned} \quad (4.24)$$

where u is the forward transformation and v is the backward transformation.

The framework is constructed step by step starting with a discussion about an efficient and stable solution scheme for the linear elastic regularization model in Sec. 4.2. We phrase the Euler-Lagrange equations of Problem 4.9 as a time-dependent diffusion problem and derive an iterative solver, that is used as relaxation method in a full multigrid optimization strategy. Finally, a von Neumann stability analysis is carried out on the resulting discretization of the problem to derive a stability criterion.

In Sec. 4.3 the stability criterion is used to derive a novel approach to define the Lamé tissue elasticity parameters, in function of Poisson's coefficient and the window size of the image-based similarity measure. This makes it possible to achieve optimal convergence rates for the iterative solution scheme with respect to the tissue parameters.

In Sec. 4.4 the convergence rate of the iterative optimization is improved by a novel line search scheme in direction of the local image force $f_{SSD}(x)$. The line search provides exact solutions for the local intensity measure minimization problem inside an interval derived from the window size of the similarity measure.

In Sec. 4.5 we present an original contribution of this work, which consists in the integration of a priori knowledge about the mechanical deformation process into the registration framework. We make advantage of the information about the probe positions, known from the rigid pre-registration, to derive probe-related forces. In other words, this approach fuses a biomechanical deformation simulation with the image-based registration process.

In Sec. 4.6 the registration framework is formulated as a simultaneous forward and backward transformation estimation process to incorporate Zhang's inverse consistent framework 4.13.

In Sec. 4.7 we derive an intensity homogenization algorithm from the local intensity mapping model $I'(x) = I(x) + a$, i.e. we correct the image pairs for local intensity shifts. This is important when the image pairs are dissimilar due to noise.

In Sec. 4.8 we validate the algorithm on patient data.

4.2 Registration Framework

In the previous section we gave an overview over the principal ingredients that are necessary to build a non-linear, intensity-based registration system based on a variational approach. In this section we chose the general framework and the regularization model. Then, an iterative solution method for the Euler-Lagrange equations of the framework is discussed. A von Neumann analysis is performed on the resulting discretized solution scheme. Finally, the iterative scheme is embedded into a multigrid optimization strategy.

The principal criteria that drive our choices are: 1) computational efficiency, i.e. the computational burden of non-linear registration should be comparable to the computational burden of the rigid registration approach, 2) sub-millimeter accuracy should be achieved, and 3) the registration should be robust even for image pairs for which the identity assumption $R \equiv T_{\hat{u}}$ does not hold.

We start our discussion with Christensen's framework 4.9 and propose a solution scheme for the linear elastic regularization model. In Sec. 4.4 and Sec. 4.5 we successively add the constraints $\mathcal{D}_{SSD}$ and $\mathcal{D}_{probe}$, and extend it to Zhang's framework 4.13 when we introduce the inverse consistency constraints $\mathcal{D}_{cons}$ in Sec. 4.6.

Christensen's framework provides a stronger regularization than Cachier's framework 4.10. As a consequence, we will not benefit from the capacity of Cachier's framework to handle deviations of the physical transformation $\hat{\varphi}$ from the transformation predicted by the regularization model. This is for example the case when the tissue elasticity parameters vary in the image. A stronger regularization, however, performs better in zones with a low signal-to-noise ratio, which are frequent in typical ultrasound images.

4.2.1 Problem Domain Ψ

In Sec. 3.1.6 we have defined the image overlap Ω as the domain on which the rigid registration problem is defined. In this chapter we embed Ω into a larger domain $\Psi \subset \mathbb{R}^3$, which is chosen to be a parallelepiped. The displacement field

u is thus a mapping from Ψ into $\mathbb{R}^3$. This type of embedding was chosen to reduce the impact of boundary conditions on the part of u that is defined on Ω. Again, the estimation process is accelerated with a multi-resolution method. The choice of a problem domain larger than Ω ensures that loss of information in the displacement field, caused by inter-grid transfers, occurs at some distance to the image borders. Furthermore, the impact of boundary conditions on the registration result is reduced.

4.2.2 Boundary Conditions

The most commonly used boundary conditions are Dirichlet boundaries with $u(x) = 0 \, \forall x \in \partial\Psi$. With Dirichlet boundary conditions the border walls of the volume are fixed in space. This can be rather restrictive in case when there are significant deformations near the border. This is very relevant in the biopsy application, where an external object, the probe, enters the volume from the side walls and pushes tissue away.

We therefore opted to fix the edges of the volume and to implement bending side walls with the boundary conditions

$$
\begin{aligned}
&u_1(x_1, 0, x_3) = u_1(x_1, N_2 - 1, x_3) = u_1(x_1, x_2, 0) = u_1(x_1, x_2, N_3 - 1) = 0 \\
&u_2(x_1, x_2, 0) = u_2(x_1, x_2, N_3 - 1) = u_2(0, x_2, x_3) = u_2(N_1 - 1, x_2, x_3) = 0 \\
&u_3(0, x_2, x_3) = u_3(N_1 - 1, x_2, x_3) = u_3(x_1, 0, x_3) = u_3(x_1, N_2 - 1, x_3) = 0 \\
&\partial_{x_1} u_1(0, x_2, x_3) = \partial_{x_1} u_1(N_i - 1, x_2, x_3) = 0 \\
&\partial_{x_2} u_2(x_1, 0, x_3) = \partial_{x_2} u_2(x_1, N_j - 1, x_3) = 0 \\
&\partial_{x_3} u_3(x_1, x_2, 0) = \partial_{x_3} u_3(x_1, x_2, N_k - 1) = 0,
\end{aligned}
$$

where $x_1 \in [0, N_i - 1], x_2 \in [0, N_j - 1], x_3 \in [0, N_k - 1]$.

4.2.3 Elastic Regularization

The deformation of the prostate is assumed to be mainly an elastic deformation, caused by probe pressure and patient movements. For this reason we use the linear elasticity operator to regularize the registration problem. Note, however, that we register an image of a deformed prostate with another image of the same prostate in another deformation state, i.e. we do not have a template image that contains the prostate in its undeformed state. This limitation leads to inexact energy potentials which can be an issue in presence of very strong and very dissimilar deformations in both images. Moreover, the linear elasticity model itself does not cope very well with large deformations. While these two points would be a serious concern for biomechanical simulations, the physical exactitude of the regularization model is less important in the context of registration, where the forces derived from the images can correct model errors to some extent.

A necessary condition for a function u to minimize Problem 4.9 is that the Gâteaux derivative of the functional

$$\mathcal{E}[R,T;u] = \sum_i \alpha_i \mathcal{D}_i[R,T;u] + \mathcal{R}[u] \tag{4.25}$$

vanishes for all perturbating functions $\psi : \mathbb{R}^3 \to \mathbb{R}^3$. The Gâteaux derivative of $\mathcal{E}$ with respect to u and in direction ψ is defined as

$$d\mathcal{E}[R,T;u,\psi] = \lim_{\tau \to 0} \frac{\mathcal{E}[R,T;u+\tau\psi] - \mathcal{E}[R,T;u]}{\tau}, \tag{4.26}$$

and hence,

$$d\mathcal{E}[R,T;u,\psi] = \sum_i \alpha_i d\mathcal{D}_i[R,T;u,\psi] + d\mathcal{R}[u,\psi]. \tag{4.27}$$

The Gâteaux derivative of the linearized elastic potential is

$$d\mathcal{R}_E[R,T;u,\psi] = \int_{\mathbb{R}^3} \langle \mu\Delta u + (\lambda+\mu)\nabla \text{div } u, \psi \rangle_{\mathbb{R}^3} \; dx, \tag{4.28}$$

see e.g. [86] for a proof. We thus get the Euler-Lagrange equations

$$\mu\Delta u + (\lambda+\mu)\nabla \text{div } u + \sum_i f_i[R,T;u] = 0, \tag{4.29}$$

where the so-called forces f_i are derived from the Gâteaux derivatives of the distance measures $\mathcal{D}_i$. The force term drives the registration process, while the gradient of the elastic potential smoothes the function u. The Eqns. 4.29 are also known as the Navier-Lamé equations.

With constant coefficients λ and μ, the Eqns. 4.29 describe uniform material, which represents another limitation of our approach. Note that Eqns. 4.29 transforms to Poisson's equation when setting the Lamé coefficient λ to $-\mu$.

The coefficients λ and μ are related to Young's modulus E and Poisson's ratio ν, i.e.

$$E = \frac{2\mu(2\mu+3\lambda)}{2\mu+2\lambda}, \quad \text{and} \quad \nu = \frac{\lambda}{2\mu+2\lambda}. \tag{4.30}$$

4.2.4 Solution Scheme and Discretization

Together with the bending boundary conditions defined in Sec. 4.25 we get an elliptic boundary value problem of the form

$$\mathcal{L}[u] = -f. \tag{4.31}$$

The existence of a solution for this boundary value problem has been proven in [86]. A classical approach to solve this type of partial differential equations if to introduce an artificial time t such that

$$\partial_t u = \mathcal{L}[u] + f, \tag{4.32}$$

see for instance [99] for a more detailed discussion. In other words, we formulate Problem 4.31 as a diffusion equation, where the initial displacement field u relaxes to an equilibruum solution as $t \to \infty$. All time derivatives vanish when approaching the equilibruum, therefore u is a solution to Problem 4.31. The term of the time derivative can be used to define the fixed point iteration

$$\frac{u^{n+1} - u^n}{\Delta t} = \mathcal{L}[u^n] + f, \tag{4.33}$$

which makes it possible to compute u^{n+1} from the previous solution u^n at iteration n, starting from an initial guess u^1 which is for example the identity transformation. The artificial timing parameter Δt is not physical and should rather be interpreted as a scaling parameter that controls the numerical step size at each iteration, and hence the convergence rate and the stability of the solution scheme.

The elastic partial differential operator $\mathcal{L}$ is of the form

$$\begin{aligned} \mathcal{L}[u] &= \mu\Delta u + (\lambda+\mu)\nabla \text{div } u \\ &= \mu\nabla^2 u + (\lambda+\mu)\nabla(\nabla \cdot u), \end{aligned} \tag{4.34}$$

and thus [86]

$$\begin{aligned} \mathcal{L}[u] = &\begin{pmatrix} (\lambda+2\mu)\partial_{x_1x_1}u_1 + \mu\partial_{x_2x_2}u_1 + \mu\partial_{x_3x_3}u_1 \\ \mu\partial_{x_1x_1}u_2 + (\lambda+2\mu)\partial_{x_2x_2}u_2 + \mu\partial_{x_3x_3}u_2 \\ \mu\partial_{x_1x_1}u_3 + \mu\partial_{x_2x_2}u_3 + (\lambda+2\mu)\partial_{x_3x_3}u_3 \end{pmatrix} \\ &+ (\lambda+\mu)\begin{pmatrix} \partial_{x_1x_2}u_2 + \partial_{x_1x_3}u_3 \\ \partial_{x_3x_2}u_1 + \partial_{x_2x_3}u_3 \\ \partial_{x_1x_3}u_1 + \partial_{x_2x_3}u_2 \end{pmatrix}. \end{aligned} \tag{4.35}$$

Problem 4.33 can be discretized with a central finite difference scheme, i.e. for a grid point $j \in \mathbb{N}^3$, a function $f : \mathbb{R}^3 \to \mathbb{R}$ and a vector field $u : \mathbb{R}^3 \to \mathbb{R}^3$ the discretization is

$$\begin{aligned} \partial_{x_lx_l} f(j) &= \tfrac{1}{h_l^2}\left(f(j+h_le_l) - 2f(j) + f(j-h_le_l)\right) + \mathcal{O}(h_l^2) \\ \partial_{x_lx_m} f(j) &= \tfrac{1}{4h_lh_m}\left(f(j+h_le_l+h_ke_m) - f(j-h_le_l+h_ke_m)\right) \\ &\quad + \tfrac{1}{4}\left(f(j-h_le_l-h_ke_m) - f(j+h_le_l-h_ke_m)\right) \\ &\quad + \mathcal{O}(h_l^2+h_m^2), \end{aligned} \tag{4.36}$$

where $e_1 = (1,0,0), e_2 = (0,1,0)$ and $e_3 = (0,0,1)$. When we assume that the images to register are isotropic, and if we operate in voxel space (i.e. $h_j = h_k = 1$), the finite difference scheme simplifies to

$$\begin{aligned} \partial_{x_l x_l} f(j) &= f(j+e_l) - 2f(j) + f(j-e_l) + \mathcal{O}(1), \\ \partial_{x_l x_m} f(j) &= \tfrac{1}{4}\Big(f(j+e_l+e_m) - f(j-e_l+e_m) \\ &\quad + f(j-e_l-e_m) - f(j+e_l-e_m) \Big) + \mathcal{O}(1) \end{aligned} \tag{4.37}$$

The central finite difference scheme will now be applied to Eqn. 4.33 and Eqn. 4.35. We only develop $\partial_t u_1$ explicitly, the developments of u_2 and u_3 are almost identical and follow immediately. With $j \in \mathbb{N}^3$ we get

$$\begin{aligned} &\frac{u_1^{n+1}(j) - u_1^n(j)}{\Delta t} = \\ &\quad (\lambda + 2\mu)\left(u_1^n(j+e_1) - 2u_1^n(j) + u^n(j-e_1)\right) \\ &\quad + \mu\left(u_1^n(j+e_2) - 2u_1^n(j) + u_1^n(j-e_2) + u_1^n(j+e_3) - 2u_1^n(j) + u_1^n(j-e_3)\right) \\ &\quad + \tfrac{1}{4}(\lambda+\mu)\Big(u_2^n(j+e_1+e_2) - u_2^n(j-e_1+e_2) - u_2^n(j+e_1-e_2) + u_2^n(j-e_1-e_2) \\ &\quad + u_3^n(j+e_1+e_3) - u_3^n(j-e_1+e_3) - u_3^n(j+e_1-e_3) + u_3^n(j-e_1-e_3) \Big) \\ &\quad + f_1. \end{aligned} \tag{4.38}$$

4.2.5 Stability Analysis

In this section we carry out a von Neumann or Fourier mode stability analysis on Eqn. 4.38. The von Neumann method makes it possible to investigate the stability of discretizations of linear and unbounded partial differential equations by analyzing the problem in the frequency domain. The von Neumann approach neglects the boundary conditions and thus represents only an approximation. The basic idea is that the solution of the unbounded problem can be represented as a superposition of wave functions with different wavelengths. A solution for a u_k is then

$$u_k^n(j) = \zeta^n e^{i\langle \omega, j \rangle}, \tag{4.39}$$

where $k = 1, ..., 3$, the wavelengths $\omega = (\omega_1, \omega_2, \omega_3) \in \mathbb{R}^3$ can have any real value, the $j = (j_1, j_2, j_3) \in \mathbb{N}^3$ is a grid point, the i is the complex $i = \sqrt{-1}$, ζ is a complex number and $n \in \mathbb{N}$ is finally the iteration number. After discarding the force term f, we can insert the solution into Eqn. 4.38, divide by u_x^n and simplify

by exploiting several trigonometric identities to finally get the amplification factor

$$\begin{aligned}\zeta(\omega) &= 1-\Delta t\Big(2\left((\lambda+2\mu)\sin^2\tfrac{\omega_1}{2}+\mu\sin^2\tfrac{\omega_2}{2}+\mu\sin^2\tfrac{\omega_3}{2}\right)\\ &\quad+(\lambda+\mu)\left(\sin^2\left(\tfrac{\omega_1+\omega_2}{2}\right)+\sin^2\left(\tfrac{\omega_1+\omega_3}{2}\right)\right)\\ &\quad+(\lambda+\mu)\left(\cos^2\left(\tfrac{\omega_1-\omega_2}{2}\right)+\cos^2\left(\tfrac{\omega_1-\omega_3}{2}\right)+2\right)\Big).\end{aligned} \tag{4.40}$$

The finite difference scheme is stable if $|\zeta|_{\mathcal{C}} \leq 1$ for all $\omega \in \mathbb{R}^3$. This leads us to a first stability criterion depending on Δt

$$\Delta t \leq \frac{1}{2(\lambda+4\mu)}. \tag{4.41}$$

We also get a stability criterion that is independent of Δt, i.e.

$$\lambda \leq -\mu. \tag{4.42}$$

Both stability criteria hold also for u_y and u_z. Now, after replacing Δt with

$$\Delta t_{max} = \frac{1}{2(\lambda+4\mu)} \tag{4.43}$$

in Eqn. 4.38 and solving for u_1^{n+1}, we obtain the well-known Jacobi method for solving partial differential equations, i.e.

$$\begin{aligned}&u_1^{n+1}(j) =\\ &\frac{1}{2(\lambda+4\mu)}\Big(f+(\lambda+2\mu)\left(u_1^n(j{+}e_1)+u_1^n(j{-}e_1)\right)\\ &+\mu\left(u_1^n(j{+}e_2)+u_1^n(j{-}e_2)+u_1^n(j{+}e_3)+u_1^n(j{-}e_3)\right)\Big)\\ &+\frac{\lambda+\mu}{8(\lambda{+}4\mu)}\Big(u_2^n(j{+}e_1{+}e_2)-u_2^n(j{-}e_1{+}e_2)-u_2^n(j{+}e_1{-}e_2)+u_2^n(j{-}e_1{-}e_2)\\ &+u_3^n(j{+}e_1{+}e_2)-u_3^n(j{-}e_1{+}e_2)-u_3^n(j{+}e_1{-}e_2)+u_3^n(j{-}e_1{-}e_2)\Big).\end{aligned} \tag{4.44}$$

4.2.6 Full Multigrid Solver

In the previous section we have derived the Jacobi solver for the Navier-Lamé equations. We now slightly modify the Jacobi scheme in order to get the Gauss-Seidel method, which has the advantage of not requiring additional storage capacities. To do this, we have to divide the grid points into two sets: the grid points $j \in \Psi^{\mathbb{N}}$ with $\sum_i j_i \mod 2 = 0$ represent the first set $\Psi_R^{\mathbb{N}}$, and the other half goes into the second set $\Psi_B^{\mathbb{N}}$. The points in the $\Psi_R^{\mathbb{N}}$ are called the red points, while the points in

$\Psi_B^{\mathbb{N}}$ are called the black points. $\Psi_R^{\mathbb{N}}$ and $\Psi_B^{\mathbb{N}}$ define a tri-dimensional checkerboard pattern.

Let us first apply the Jacobi scheme on the red grid points, and then carry out a second iteration on the black points. At the second iteration, 14 of the 26 direct neighbors of a grid point $j \in \Psi_B^{\mathbb{N}}$ have already been updated to $n+1$. This increases convergence speed by a factor of approximately two, and we do not need any longer to store the updated values u^{n+1} in a temporary buffer, all solutions being updated in-place. This strategy is also called the red-black Gauss-Seidel scheme.

To improve the convergence rate of the proposed solution scheme, we take advantage of multigrid techniques. A novel image resolution pyramid is constructed in analogy to the pyramid used for rigid pre-registration (see Sec. 3.3), but without elimination of the deformations near the probe head.

Recall from the last section that the solution of a linear and homogenuous boundary value problem at a point j is a composition of Fourier modes of the form $e^{i\langle\omega,j\rangle}$, where the wavelengths $\omega_i = \tilde{\omega}_i\pi/n$. Depending on the discretization granularity Δx, the approximation error is more rapidly reduced for some wavelengths than for others, in function of the fraction $\omega/\Delta x$. In other words, a red-black Gauss-Seidel iteration on level i is blind to high frequency errors with wavelengths $\omega << \Delta x$ and converges extremely slowly for low frequency errors with wavelengths $\omega >> \Delta x$.

The basic idea of the multigrid approaches is to solve the problem successively on different resolution levels starting from solutions computed on neighbored levels, i.e. for a level i either from level $i-1$ or from level $i+1$. The most classical multigrid approach is the V-cycle (Fig. 4.1, which consists in solving successively from fine to coarse levels, i.e. from 1 to n, and to iterate back to the finest level.

Note that many implementations of the multigrid scheme use an explicit solver at level n, in particular if n does not correspond to the coarsest possible level (which corresponds for a problem with an identical number of grid points N in each dimension to $\log_2(N)$). For an image registration problem driven by image-based forces that are computed locally within a limited window, this approach does not yield any particular advantage compared to an iterative relaxation scheme, since the latter benefits from the force update at each iteration.

$$n \geq \log_2(2\omega) \tag{4.45}$$

Note that not only the lower frequencies are more efficiently solved on coarser levels, but also that the computation time to solve the problem on level n reduces to $2^{-3(n-1)}$ times the computation time required to solve level 1. These two properties turn the multigrid approach into a very powerful optimization technique.

The most efficient multigrid scheme is the full multigrid strategy, which does

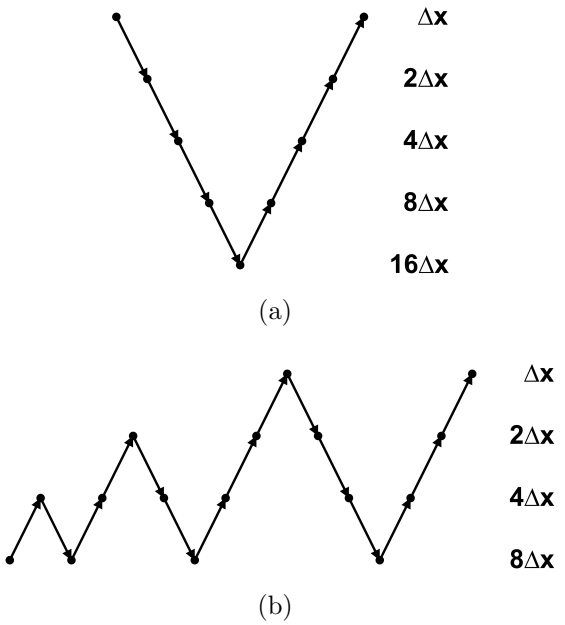

Figure 4.1: Schedule of grids for (a) the V-cycle strategy, and (b) the full multigrid strategy.

not start from the finest level, but from the coarsest level. The iteration scheme iterates from $n \to n-1 \to n \to n-1 \to n-2 \to n-1 \to n \to n-1...$ to 1, see Fig. 4.1.

The inter-grid transfer operator required to pass from level i to level $i+1$ is built in analogy to the restriction operator introduced in Sec. 3.3.1 and implements in particular the shift rule introduced in Sec. 3.3.2. This operator contains information loss at grid borders, and properly aligns low resolution versions of u with the corresponding low resolution versions of the images. For the coarse-to-fine transfer, also called prolongation, we use trilinear interpolation.

4.3 Elasticity and Timing Parameters

In Sec. 4.2.5 we have derived the timing parameter

$$\Delta t_{max} = \frac{1}{2(\lambda + 4\mu)},$$

which depends on the elasticity parameters λ and μ. It is important to note that Δt_{max} scales the length of the forces that are added to the displacement field at

each iteration, i.e.

$$\bar{f} = \frac{f}{\Delta t_{max}}.$$

In consequence, the convergence of the iterative estimation process degrades with increasing Δt_{max}.

In literature, the choice of λ and μ, which can also be expressed with Young's modulus E and Poisson's tissue compressibility coefficient ν, is often arbitrary. Most authors have chosen $\mu = 1$ and $\lambda = 0$, see e.g. [7, 27]. Modersitzki adapts the tissue elasticity parameters after the first iteration such that the convergence rate is maximized, in function of the norm of u^1 [86]. The draw-back of these approaches is that they abandon the capacity of the linear elastic regularization model to incorporate tissue elasticity properties.

We propose an alternative approach, which determines λ and μ in function of Poisson's coefficient and in function of the radius of the window on which the image-based force term is computed, which is 0.5 voxel. Hence, only Young's modulus is interpreted as a free parameter. The advantage of this approach is that it is possible to modify the tissue compressibility without modifying the convergence characteristics of the registration.

4.3.1 Choice of the Elasticity Parameters λ and μ

The choice of sound values for the Lamé coefficients λ and μ, or (almost) equivalently, for E and ν, is subject to some subleties. The first idea is to use good estimates of the average values for prostatic tissue. Souchon gives in [109] average values for the Young modulus of the prostate in function of the prostate pathology and the precompression state of the prostate. He finds that, when the prostate is precompressed by 2 percent, then the Young modulus varies from 36±9 kPa if the prostate is affected by a benign hyperplasia, to 100 ± 20 in the presence of cancer. In case that the precompression is 4 percent, Young's modulus varies from 36±11 to 212±32. The values for healthy prostates lie somewhere in between. In conclusion we can say that Young's modulus varies significantly in function of the pathology, but also in function of the compression state of the tissue. However, in the context of image-based prostate registration, the value of this information is limited: since we cannot determine the physical forces that have acted on the prostate when the image was acquired, the average Young modulus is of no interest to us.

Poisson's ratio ν describes the ratio between axial and transverse strain. Assume that a material is compressed in axial direction, then

$$\nu = \frac{\epsilon_{trans}}{\epsilon_{axial}} \tag{4.46}$$

where ϵ_{trans} is the transverse strain, and ϵ_{axial} is the axial strain. Almost all materials have Poisson coefficients between 0 and 0.5; only some composite materials have negative Poisson coefficients. A Poisson coefficient of 0 means that the material is highly compressible and does not extend its volume in directions orthogonal to compressing forces. A - physically impossible - coefficient of 0.5 corresponds to a completely incompressible material. In general, human tissue has a Poisson coefficient of about 0.49 to 0.499 [92], i.e. it is relatively incompressible due to the fact that it contains mainly liquids. During our experiments we could observe that tissue that surrounds the gland is considerably more compressible, in particular the rectal wall between the probe head and the prostate. In contrast to Young's modulus, Poisson's coefficient has the advantage that it can be used in the context of image-based registration that operates with voxel displacements rather than physical forces.

4.3.2 Timing, Stability and Convergence Rate

The driving forces of our registration problem are derived from image intensity displacements and are not physical. Moreover, the forces are computed inside a small window. Hence, they are only valid inside this limited domain and they are only estimates of the real displacement of limited quality. We therefore have to ensure that the contribution $\|f\|/\Delta t_{max}$ is smaller than the validity range of the image-based forces. Else, the optimization process, which is mainly a gradient descent, might jump outside the capture range of the physical transformation $\hat{u}$ that we seek to estimate. If this happens for too many grid points, optimization starts to diverge. The window on which the image-based forces and the inverse consistency forces are computed is just one voxel, or grid point. Hence, the radius of the window is 0,5 times the voxel side length. On the other hand, when Δt_{max} is too large, the contributions of the force terms to the displacement field u get extremely small, which can lead to very slow convergence rates. The choice of Δt_{max} is therefore crucial for the stability and the convergence rate of the registration problem.

We have chosen to establish the following important rule: each force of the registration framework has to be scaled such that its maximum length is one voxel side length, i.e.

$$\|f_s\| \leq 1 \quad \forall s \in \{\text{SSD, probe, cons}\} \tag{4.47}$$

This convention makes it possible to control the force contributions with a single parameter, i.e. with Δt_{max}. As a consequence, the different force terms used in our approach all contribute equally to the total force. The fine-tuning of the contribution of the forces depends on Δt_{max}, and a typical choice is for example 0.5, which would correspond to the window radius of the image-based force term.

Δt_{max} can be chosen freely when making advantage of the degree of freedom gained by abandoning the physical meaning of the Young modulus. Then, it is possible to define the Lamé coefficients λ and μ in dependency of Δt_{max} and ν using the Eqn. 4.43 and Eqn. 4.30, i.e.

$$\begin{aligned} \mu &= \frac{1}{\Delta t_{max}} \cdot \frac{2\nu - 1}{12\nu - 8} \\ \lambda &= \frac{1}{2\Delta t_{max}} - 4\mu. \end{aligned} \tag{4.48}$$

4.3.3 Experiments and results

We compared the performance of different timing and tissue compressibility parameters using a sample ultrasound image of the prostate on which we simulated a 15 mm probe insertion (see Fig. 4.2) using the elasticity parameters $\lambda = 0$ and $\mu = 1$ that model the tissue as completely compressible ($\nu = 0$). Moreover, Dirichlet boundary conditions were used for this simulation. Note that the simulated 15 mm probe insertion causes large deformations and represents a rather hard stress test for a solely image-based registration algorithm, due to the fact that there is only rudimentary information available where the deformations are strongest, i.e. at the location of the simulated probe head. However, this is exactly the type of deformation that can mainly be observed during transrectal prostate biopsy acquisition under transrectal ultrasound control. Nevertheless, we have never observed a 15 mm insertion, the simulation thus represents an extreme assumption. The warp $\hat{u}$ that generates the deformation is used to evaluate the RMS error between the estimated field u and $\hat{u}$, i.e.

$$\epsilon_{rms} = \sqrt{\frac{1}{N_i N_j N_k} \sum_{i=0}^{N_i-1} \sum_{j=0}^{N_j-1} \sum_{k=0}^{N_k-1} \|u(i,j,k) - \hat{u}(i,j,k)\|_{L_2}^2} \tag{4.49}$$

We also computed the decile values of the RMS error, since the global value can get heavily influenced by large outliers due to the quadratic weighting of the local distances. This is in particular true for the test transformation $\hat{u}$, where the largest displacement vectors point outside the image at the probe insertion point (recall that we work in the Euler frame). Furthermore, we do not give a maximum error for this particular test configuration, since it is relatively meaningless.

Decile error analysis allows to compare effects of the elastic registration on image elements that are more or less already aligned before registration (the most aligned elements lying in decile 1) and on unaligned elements (the most unaligned

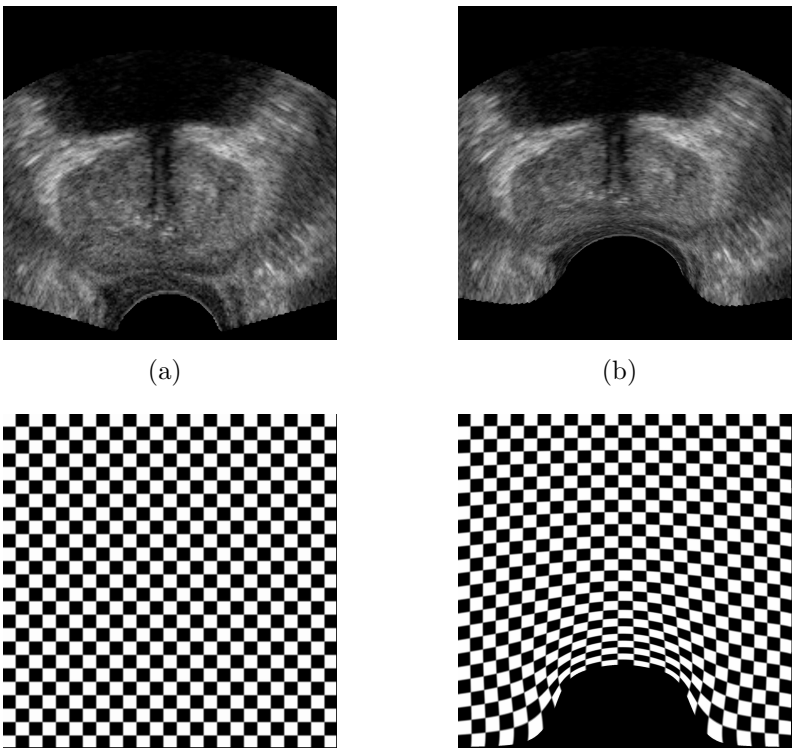

Figure 4.2: Experimental setup: test images. Fig. (a) shows the template T, Fig. (b) illustrates the reference R, created from T by simulating a probe insertion of 15 mm, and Fig. (c) depicts an undeformed volume filled with a 3D checkerboard pattern that corresponds to Fig. (a), and Fig. (d) illustrates the deformation $\hat{u}$ that generated R, applied on the checkerboard pattern of Fig. (c).

elements lying in decile 10). Note that elastic registration can locally degrade registration performance, either due to the local inadequacy of the elastic model or due to inadequate image, probe or inverse consistency forces. This is in particular true in regions that do not lie inside the image overlap Ω. In these regions registration is purely driven by the a priori models that we will introduce.

The most interesting intervals for most tests are the deciles 6 to 9, which represent the largest deformations of $\hat{u}$ in regions containing intensity information in both images. The displacements of the deciles 1 to 3 and 10 almost all lie in regions where no image data is available in at least one of the images during registration. These deciles allow to analyze the behavior of the regularization outside overlapping regions for different parameterizations, for instance in function of a given compressibility parameter ν.

For this particular test we have used the registration framework,

$$u^* = \arg\min_u \Big(\mathcal{D}_{SSD}[R, T; u] + \mathcal{R}[u] \Big), \tag{4.50}$$

where the computation of the corresponding forces f_{SSD} is discussed in Sec. 4.1.6.

The test was carried out using a six level resolution pyramid, and full multigrid optimization was carried out up to the second level. After each restriction/prolongation no more than two Gauss-Seidel smoothings were performed. The optimization was stopped when the finest level was reached for the first time. The termination criterion of the optimization is thus simply the number of iterations, and not a criterion defined on the convergence of the force terms. We hence do not expect perfect registrations as test results. The focus lies on the relative performance differences of the compared settings.

We carried out two test series: in the first, we tested the influence of the timing parameter Δt_{max} on the stability of the registration process. We have tested the parameters $\Delta t_{max} = 2, 1, 0.4, 0.25$, using a tissue compressibility parameter $\nu = 0,47$. Note that the tissue compressibility used for registration is different from the Poisson coefficient used to generate the deformation ($\nu = 0$), in an attempt to render the test more challenging. The results are given in Fig. 4.3 and Fig. 4.4. Gray zones in the 3D checker-board pattern indicate deformations in z-direction. Note that the transformations start to get smooth when approaching the discretization resolution of 0.5 voxel side length, while algorithm diverges for larger values. A decile error analysis is given in Fig. 4.5. For large timing parameters, the RMS distance of the estimated field u^* to $\hat{u}$ is larger than the RMS distance before registration. The convergence rate decreases with lower timings. We feel comfortable with a timing of 0.25 voxel side lengths, which seems to be a good compromise between stability and convergence speed.

The second test demonstrates the influence of different tissue compressibility constraints. Note that the test deformation has been created with a Poisson coefficient of 0, i.e. $\hat{u}$ contains many stretched and squeezed voxel. This test illustrates the impact of different compressibility constraints on image-based zones and on the probe head zone, where no forces are available, and where the elasticity model extrapolates the transformation. The results are given in Fig. 4.6 and the corresponding decile error analysis can be found in Fig. 4.7. Note the differences in the probe head area for the different settings. Note also that the registration diverges when the compressibility constraint is set to 0.4999, i.e. to nearly uncompressible (not illustrated).

An error analysis is given for all tested values in Fig. 4.7. Fig. 4.7 (a) shows the influence of the time step Δt_{max} on registration stability. We can observe that registration gets instable when Δt_{max} is superior to the window radius of

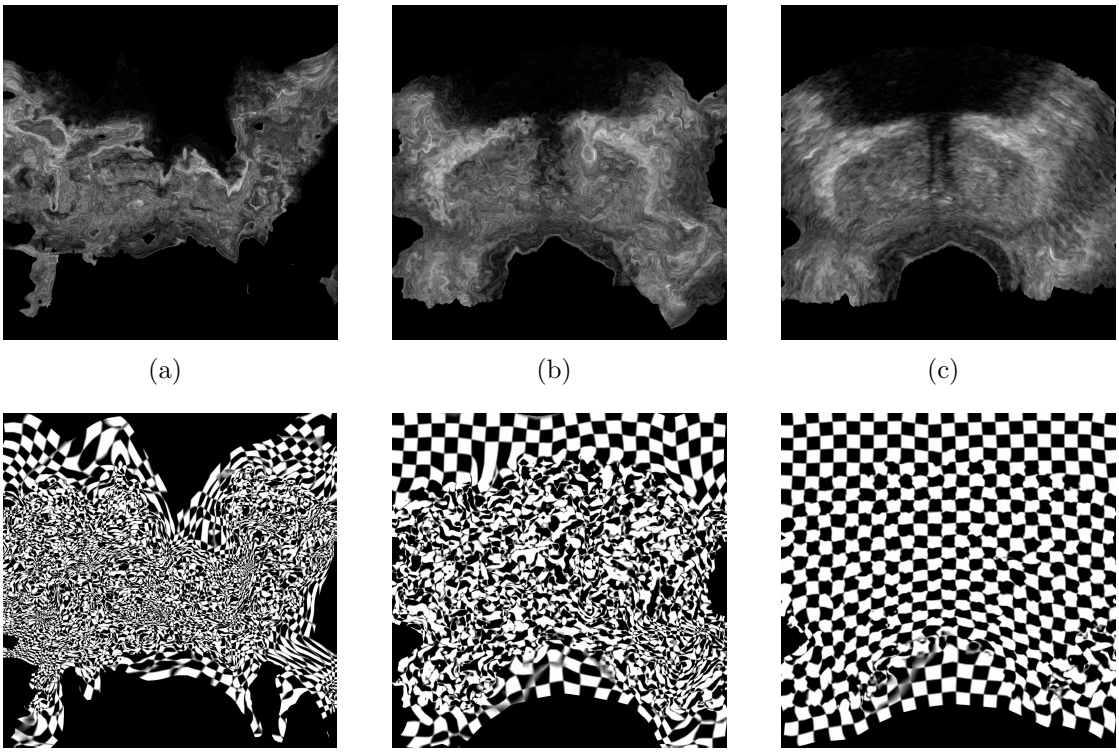

(a) (b) (c)

(d) (e) (f)

Figure 4.3: Timing parameter comparison I. All registrations were performed with $\nu = 0.47$. Fig. (d), (e) and (f) illustrate the deformations applied to a 3D checkerboard pattern. Fig. (a) and (d) were estimated with $\Delta t_{max} = 4.0$, Fig. (b) and (e) were estimated with $\Delta t_{max} = 2.0$, and Fig. (c) and (f) were estimated with $\Delta t_{max} = 1.0$.

the similarity measure, i.e. when it is greater than 0.5. Smooth, convergent registrations can be observed for values Δt_{max} below 0.5.

4.3.4 Discussion and Conclusion

The proposed parameter setting framework makes it possible to integrate a priori knowledge on tissue compressibility into the registration process. The advantage of our approach is that the tissue compressibility can be modified without impacting the convergence rate of the registration process.

The choice of the timing parameter Δt_{max} is crucial for robust registration with an acceptable convergence rate. We have found in our experiments that $\Delta t_{max} = 0.25$ is a sound compromise between robustness and convergence speed.

A general observation that can be made concerning tissue compressibility pa-

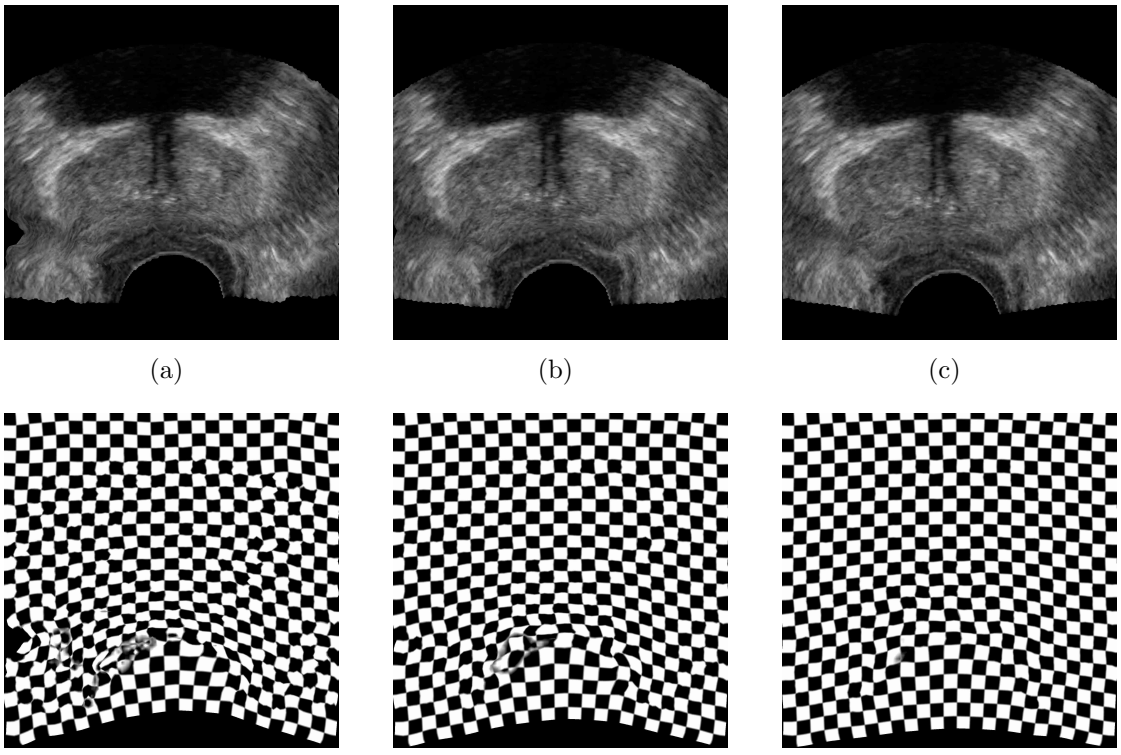

Figure 4.4: Timing parameter comparison II. All registrations were performed with $\nu = 0.47$. Fig. (d), (e) and (f) illustrate the deformations applied to a 3D checkerboard pattern. Fig. (a) and (d) were estimated with $\Delta t_{max} = 0.5$, Fig. (b) and (e) were estimated with $\Delta t_{max} = 0.25$, and Fig. (c) and (f) were estimated with $\Delta t_{max} = 0.125$.

rameterization is that low tissue compressibility, i.e. values for ν near 0.5, lead to a more restrictive regularization, see Fig. 4.7. This is due to the fact that low compressibility restricts the application of the "'easy solution"' when the image-based force field is inhomogeneous, which consists in stretching voxel between neighbored zones of contradictory forces. This might be a great advantage in the context of partially overlapping images, e.g. for lateral biopsy images containing only part of the prostate, since the deformation field is extrapolated to non-overlapping zones.

However, in the presence of large physical dissimilarities between tissues imaged in the reference and tissues images in the template image, e.g. for registration of panorama images of two different biopsy series acquired with a time lapse of several months, the low tissue compressibility model performs poorly and tends to diverge. When the algorithm tries to enforce the compressibility constraint between zones of local dissimilarities and zones of good similarity, it tends to push estimates away

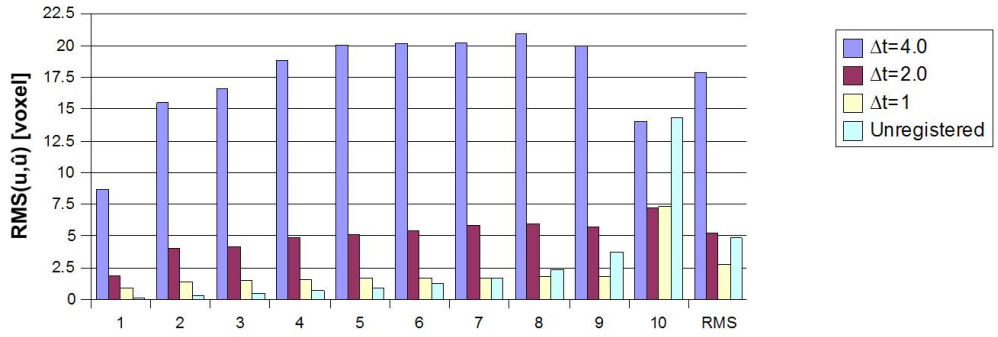

(a)

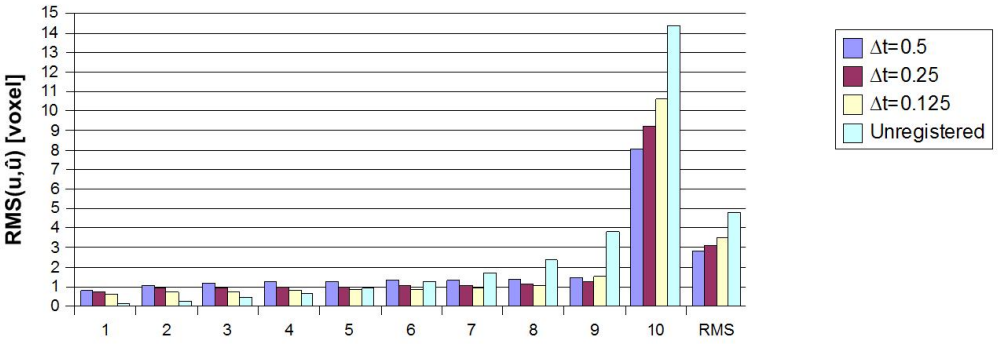

(b)

Figure 4.5: Timing parameter comparison. Fig. (a) and (b) illustrate the influence of the timing parameter Δt_{max} on the stability of the registration process. 1-10 are the decile RMS errors in voxel and RMS is the overall RMS error in voxel. The values of the unregistered transformation were computed with the identity transformation $u = Id$.

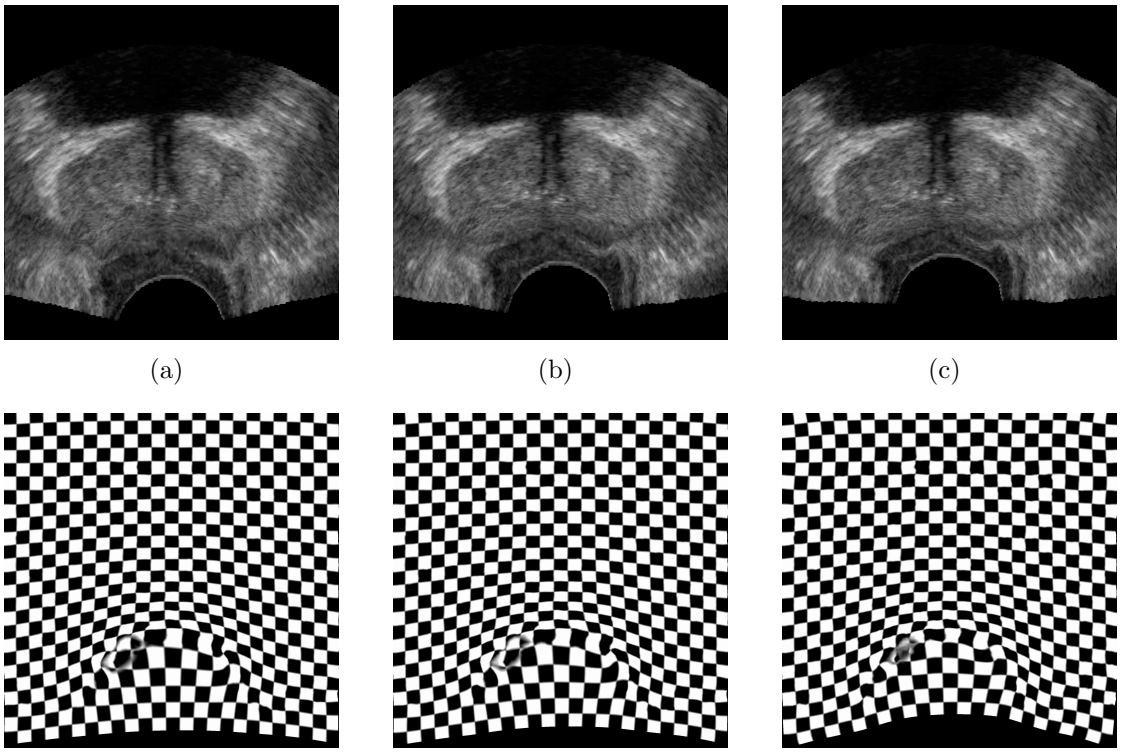

Figure 4.6: Tissue compressibility parameter comparison. Fig. (d), (e) and (f) illustrate the deformations applied to a 3D checkerboard pattern. Fig. (a) and (d) were estimated with $\nu = 0$, Fig. (b) and (e) were estimated with $\nu = 0.25$, and Fig. (c) and (f) were estimated with $\nu = 0.4$.

from the correct solution, and sometimes even outside the capture range of the image-based energy. Similar observations can be made in the presence of noise.

Less restrictive compressibility constraints like $\nu = 0.0$ reduce this problem by stretching or compressing voxel between zones of high image similarity and zones with flagrant dissimilarities. This is e.g. an excellent model for different bladder fillings or tumor growth in the template and the reference image. However, this is a poor model when field extrapolation is desired: without compressibility constraints, voxel at the overlap borders tend to get excessively stretched, and the sum of the norm of the displacement in $\Psi \setminus \Omega$ is close to zero.

We can conclude from these observations that it is not a good idea to use tissue compressibility constraints near the physical average of $\nu \approx 0,495$: without even thinking about physical tissue dissimilarities, speckle, ultrasound shadows and viewing angle related shifts of the ultrasound signal response intensities rep-

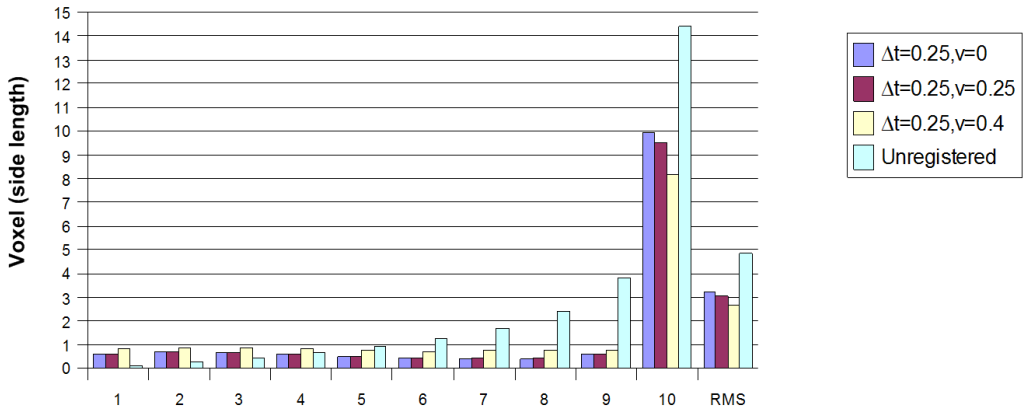

Figure 4.7: Tissue compressibility parameter comparison. The figure illustrates the influence of the compressibility parameter ν on registration. 1-10 are the decile RMS errors in voxel, and RMS is the overall RMS error in voxel. The values of the unregistered transformation were computed with the identity transformation $u = Id$.

resent caveats serious enough to spoil registration. We therefore recommend to use an intermediate value of $\nu \approx 0.3$ for images acquired with a time lag of several minutes at maximum, for which the probability of physical tissue changes between acquisitions is relatively limited, and for which field extrapolation in barely overlapping areas is of interest. Concerning greater time lags, it is preferable to use far less restrictive constraints. Fortunately, registration is carried out on image pairs containing both the entire prostate in the context of repeated biopsy series, i.e. we register the panorama images of two series. Extrapolation is thus less a concern. Be aware, however, that this choice can lead to significant voxel stretching or compression at locations where such a behavior is not justified.

4.4 Image-based forces

In this section we first derive the image-based forces from the Gâteaux derivative of the energy term $\mathcal{D}_{SSD}$. Then we discuss alternative force terms often used in image registration and carry out a comparative study. Finally, we propose a line search scheme that allows to eliminate the force scaling parameter α_{SSD} in framework 4.9.

4.4.1 Gâteaux derivative of $\mathcal{D}_{SSD}$

The Gâteaux derivative of $\mathcal{D}_{SSD}$ at point u with respect to the perturbing function $\psi : \Psi \to \mathbb{R}^3$ can be derived with the Taylor expansion of $T(x + u(x) + h\psi)$ at the expanding point $x + u(x)$,

$$T(x + u(x) + h\psi(x)) = T_u(x) + h\langle(\nabla T)(x + u(x)), \psi(x)\rangle_{\mathbb{R}^3} + \mathcal{O}(h^2).$$

Hence,

$$\begin{aligned}
&d\mathcal{D}_{SSD}[R, T; u, \psi] \\
&\quad = \lim_{h\to 0}\frac{1}{h}(\mathcal{D}_{SSD}[R, T; u + h\psi] + \mathcal{D}_{SSD}[R, T; u]) \\
&\quad = \lim_{h\to 0}\frac{1}{2h}\|T_u(x) + h\langle(\nabla T)_u(x), \psi(x)\rangle_{\mathbb{R}^3} + \mathcal{O}(h^2) - R(x)\|^2_{\mathbb{R}^3} \\
&\qquad - \|T_u(x) - R(x)\|^2_{\mathbb{R}^3}\ dx \\
&\quad = \int_\Psi \langle(T_u(x) - R(x))(\nabla T)_u(x), \psi(x)\rangle_{\mathbb{R}^d}\ dx.
\end{aligned}$$

We thus have to minimize the local forces $f_{SSD} : \Psi \to \mathbb{R}$,

$$f_{SSD}[R, T; u](x) := \begin{cases} \alpha_{SSD}(T_u(x) - R(x))(\nabla T)_u(x) & ,\ x \in \Omega \\ 0 & ,\ \text{else} \end{cases}. \tag{4.51}$$

Recall from Sec. 4.3.2 that the contribution of the force term to u at each iteration of the iterative solution process is determined by Δt_{max}. The contribution has to be smaller than the radius of the window on which $\mathcal{D}_{SSD}$ is computed, else the local solutions $u(x)$ can "jump" outside the capture range of the corresponding $\hat{u}(x)$ and the algorithm starts to diverge. To control the contribution of the forces, we established the rule that the length of the forces smaller than one voxel side length.

The Gâteaux derivative of $\mathcal{D}_{SSD}$ can be normalized to a maximum length of one for example with the local scaling factor

$$\alpha_{SSD}(x) = \frac{1}{\Delta I\|(\nabla T)_u(x)\|}, \tag{4.52}$$

where ΔI is the maximum possible intensity difference between R and T, computed on the entire images. Note that the scaling parameter depends on x. Using f_{SSD} in this form in Problem 4.33 corresponds to a gradient descent on the image gradient of $T(u(x) + x)$. A different normalization can be found for example in [124].

4.4.2 Alternative force terms

Many authors use the gradient $(R - T_u)\nabla R$ instead of $(T_u - R)(\nabla T)_u$ as descent direction [119, 133]. Other authors [12, 124] combine both gradients, i.e. they optimize in direction $\nabla R + (\nabla T)_u$. This can be justified by considering the Taylor expansion of $R(x + w(x))$ at the point x, i.e.

$$\begin{aligned} R_w(x) &= R(x + w(x)) \\ &\approx R(x) + \nabla R(x) \cdot w(x) + \frac{1}{2} w^T(x) \nabla^2 R(x) \cdot w(x) + \mathcal{O}(w(x)^3). \end{aligned} \tag{4.53}$$

An expansion of $\nabla R(x + w(x))$ at x yields

$$\nabla R(x + w(x)) = \nabla R(x) + \nabla^2 R(x) \cdot v(x) + \mathcal{O}(w(x)^3). \tag{4.54}$$

After plugging Eqn. 4.54 into Eqn. 4.53, we get

$$R(x) - R(x + w(x)) = -\frac{1}{2}\big(\nabla R(x) + \nabla R(x + w(x))\big) \cdot w(x) + \mathcal{O}(w(x)^3). \tag{4.55}$$

After discarding the third-order terms, this is a second order approximation of the minimization problem $R - R_w$ in a small neighborhood of x. Under the assumption $R \equiv T_{\hat{u}}$ we get $\varphi_u = \varphi_{\hat{u}} \circ \varphi_w$ and we obtain the approximization

$$R(x) - T_u(x) \approx -\frac{1}{2}\big(\nabla R(x) + \nabla T(x + u(x))\big) \cdot u. \tag{4.56}$$

In theory, the mean of the gradients ∇R and $(\nabla T)_u$ should thus yield a more robust guess of the minimization direction than $(\nabla T)_u$ alone.

Finally, Stefanescu et al. use ∇T_u instead of $(\nabla T)_u$ [112], derived from the displacement field composition scheme

$$\varphi_{u^{n+1}}(x) = \varphi_{u^n} \circ \varphi((\Delta u)^{n+1} + x), \tag{4.57}$$

where $(\Delta u)^{n+1}$ is the contribution of the forces to the displacement field at iteration $n + 1$. Stefanescu shows that in the presence of large deformations, the additive displacement scheme may become considerably less efficient or even invalid due to the fact that ∇T is computed only once. The Gâteaux derivative of the SSD cost function with the composition scheme is

$$d\mathcal{D}[R, T; u, \psi] = \int_\Psi \langle (T_u(x) - R(x)) \nabla T_u(x), \psi(x) \rangle_{\mathbb{R}^3} \, dx. \tag{4.58}$$

From Eqn. 4.57 we get the field update

$$u^{n+1}(x) = u^n((\Delta u)^{n+1}(x) + x) + (\Delta u)^{n+1}(x). \tag{4.59}$$

We will make now advantage of the fact that Δu is in our framework small by construction, with $\|\Delta u\|$ being smaller than Δt_{max}, which we normally choose to lie in the interval [0.25,0.125]. It is thus legitimate to use the following approximation, which makes it possible to integrate Stefanescu's gradient into our framework:

$$u^{n+1}(x) \approx u^n(x) + (\Delta u)^{n+1}(x). \tag{4.60}$$

This corresponds to the field update of the additive scheme that we use. Note that the error introduced by this approximation is corrected at the following iteration and hence does not cumulate, while the positive effect of the gradient ∇T_u in presence of large deformations is maintained.

4.4.3 Comparative study

We have compared the different gradients using the test configuration introduced in Sec. 4.3. Poisson's coefficient was set to 0.3, Δt_{max} to 0.25. For this particular test we performed 10 Gauss-Seidel smoothings instead of two at each multigrid iteration, and the error was computed on $\Omega[R, T; Id]$ instead of Ψ in order to focus on the gradient performances on image data.

Fig. 4.8 illustrates the decile analysis of the error and the overall RMS error. In the first decile there are only voxel that lie outside the image overlap, and the error in this decile is worse than before registration because of the regularization. The gradients perform almost identically, they all converge to the correct solution and have similar errors. ∇R and $\nabla R + (\nabla T)_u$ perform worse on large deformations (deciles 7-10) than $(\nabla T)_u$ and ∇T_u. ∇T_u performs best with an overall RMS error of 0.57, closely followed by $(\nabla T)_u$ with 0.63. ∇T_u is slightly less efficient than the other gradients on small deformations (deciles 2-5), probably due to the error introduced with the approximation 4.60, that we used to fit it into our additive scheme. We conclude from these observations that the choice of the gradient direction has less impact on the convergence rate than we expected. In the remainder of this work we will stick with ∇T_u since it yields the best overall results and performs best for large deformations.

4.4.4 Force Scaling vs. Line Search

The choice of $\alpha_{SSD}(x)$ in Eqn. 4.52 is not optimal, and it can be improved by analyzing the intensity distributions of the images. We were, however, neither happy with the convergence rate of the gradient descent, nor with the idea of optimizing α_{SSD} derived by additional image analysis. Instead, we preferred to carry out a line search in gradient direction, which revealed to be a surprisingly efficient optimization scheme. The line search is carried out by minimizing the

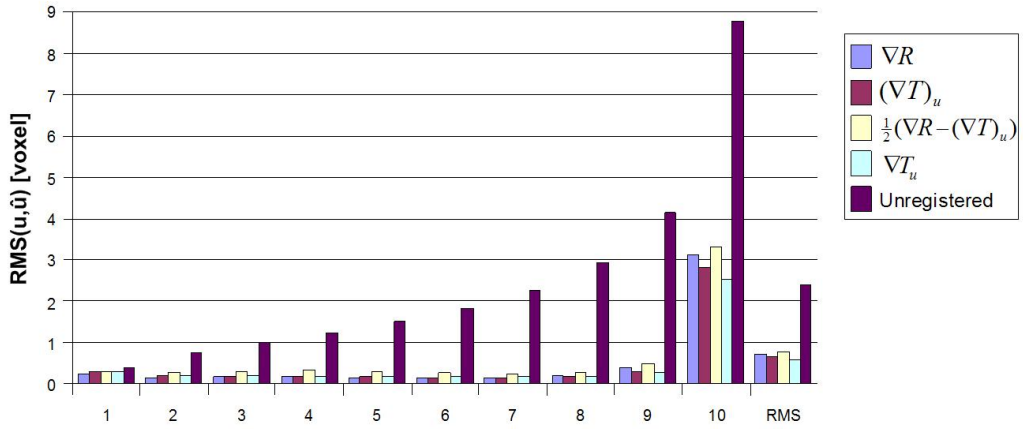

Figure 4.8: Gradient comparison. This plot illustrates the performances of the four most commonly used gradient terms. Note the the RMS error in this plot was computed on Ω only. 1-10 are the decile RMS errors in voxel and RMS is the overall RMS error in voxel. The values of the unregistered transformation were computed with the identity transformation $u = Id$.

difference $R - T_u$ in gradient direction on a line segment $d \in \mathbb{R}$ with

$$0 \leq d \leq 1, \tag{4.61}$$

such that

$$d^* = \arg\min_d \left(T(d \frac{R(x) - T_u(x)}{|R(x) - T_u(x)|\|\nabla T_u\|} \nabla T_u(x)) - R(x) \right). \tag{4.62}$$

Then we set the force to

$$f_{SSD}[R, T; u] = d^* \frac{R - T_u}{|R - T_u|\|\nabla T_u\|} \nabla T_u. \tag{4.63}$$

Since we use trilinear interpolation on the images, the minimization problem is linear on [0,1] in gradient direction. The implementation is hence straightforward and introduces only a negligible computational overhead. We tested this approach with the same test configuration as earlier in Fig. 4.4.2, the results are given in Fig. 4.10 and Fig. 4.9.

One can see the superior performance of the line search on large deformations in the reference image, i.e. the errors in the deciles 9 and 10 are considerably

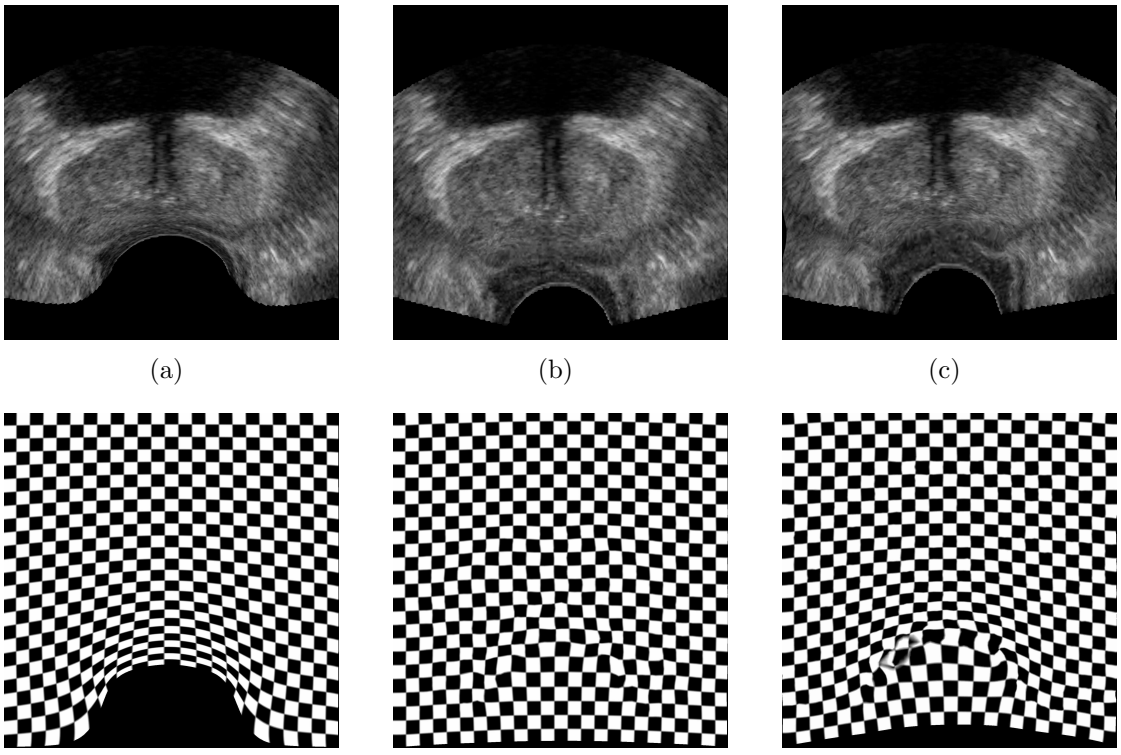

Figure 4.9: Gradient descent vs line search. (a)+(d) illustrate the reference image, (b)+(e) the result of the gradient descent and (c)+(f) the result of the line search. With an identical number of iterations, timing and elasticity parameters, the line search estimate is closer to the reference image than the gradient descent estimate.

more reduced than with gradient descent. Note that Δt_{max} serves as an upper limit for d and as a divisor of the force term. In other words, at each iteration we add the cropped, but exact, solution of Problem 4.62 to the displacement field instead of a damped solution. This leads to faster convergence, but reduces the local smoothness of the warp somehow.

4.4.5 Discussion and Conclusion

The comparative study of gradient terms has shown that it is preferable to stick with the Gâteaux derivative of $\mathcal{D}_{SSD}$ as search direction. At least in our case, the Gâteaux derivative yields considerably better results than the alternative terms, which can be explained with the strong deformations in the image.

The line search in gradient direction yields exact solutions of the linearized local search problem in a range of one voxel side length. This allows it to discard

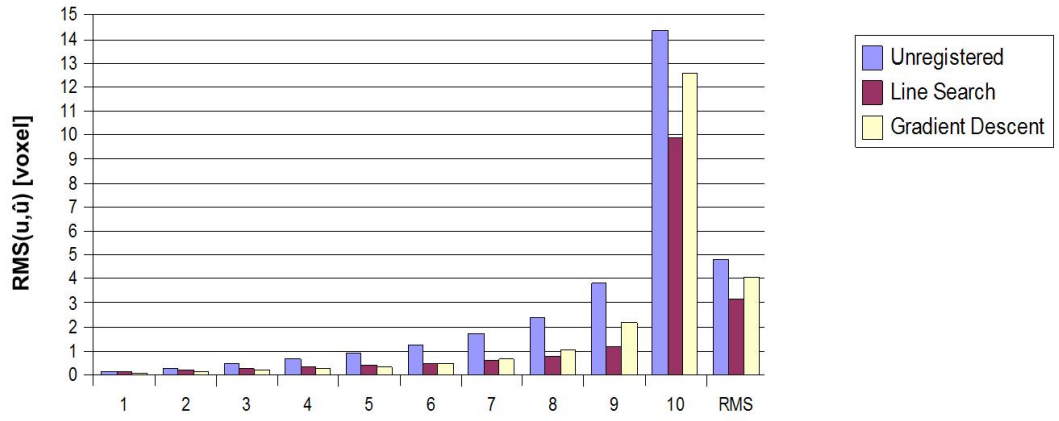

Figure 4.10: Line search. Comparison of the line search algorithm with the gradient descent scheme. 1-10 are the decile RMS errors in voxel and RMS is the overall RMS error in voxel. The values of the unregistered transformation were computed with the identity transformation $u = Id$.

the scaling parameter α_{SSD}, which is difficult to choose. Note that gradient scaling is in general a subtle issue in gradient descent schemes. The line search guarantees fast convergence rates in combination with estimations of high accuracy.

4.5 Probe-related Forces

When reconsidering the experiments of the last sections, one can observe that the error correction is not really satisfying in zones near the probe head where the deformation is strongest. The grid points corresponding to these zones can be found in the deciles 9 and 10 of the experiments in Sec. 4.3 and in Sec. 4.4. It would therefore be helpful if additional information could be used to improve convergence speed.

In this section we integrate a priori models on tissue displacements that occur during probe insertion in form of an additional cost function $\mathcal{D}_{Probe}$ into the registration framework 4.9. The approach is original in the sense that the probe forces are not be defined as a distance between image feature pairs, in contrast to the solutions proposed by e.g. [35,36,64,107]. Furthermore, the presented method does not rely on organ segmentation, in contrast to the deformable prostate registration methods presented in Sec. 4.1.8.

In our approach, the forces are derived from a priori assumptions on the probe

displacements before acquisition of the reference image, and before acquisition of the template image. They are formulated such that they support the image-driven registration process without dominating it, i.e. the role of the probe forces is to assist the image-based registration. We make advantage of the knowledge about the relative probe positions in reference space that were computed with the rigid pre-registration. We hence assume that the probe head radius and the probe head position in tracking space is known[1].

The presented method can be interpreted as the fusion of a bio-mechanical simulation with an image-based registration. The objective is to stabilize the registration process at strongly deformed image borders, where purely image-driven approaches produce poor results due to a lack of image information.

4.5.1 A Priori Assumptions on Probe Movements

For an end-fire endorectal ultrasound probe, two simplifying assumptions can be made on the probe displacements that occur before the acquisition of a prostate image. First, the depth of the probe insertion into the rectal tissues is equal or greater than the visible part of the probe in the image. Second, the trajectory of the probe displacement is linear in direction of the probe axis. Fig. 4.11 shows a typical image of the prostate acquired with an end-fire probe.

With these assumptions it is possible to approximate the probe-related tissue displacements in both the reference and the template images. Fig. 4.12 (a) illustrates the tissue displacements at the probe head location in the template image, Fig. 4.12 (b) illustrates the tissue displacements at the probe head location in the reference template, after projection into the template image space with the rigid transformation computed during rigid pre-registration. The relative tissue displacements from the template to the reference image, are thus the reference image displacement minus the template image displacements, see Fig. 4.12 (c).

4.5.2 Displacement Computation

The displacement computation is straightforward: one mainly has to compute line intersections with a sphere that represents the probe head. If we operate for a given image I in a space where the probe origin corresponds to the coordinate system origin O_I, and for a probe direction $d_I \in \mathbb{R}^3$, problem grid point $x \in \mathbb{R}^3$ and probe head radius r

$A \leftarrow 2\langle d_I, x \rangle_{\mathbb{R}^3}$

[1] In practice, we detect the probe head position and radius automatically and in real-time with an algorithm based on a recursive Hough transform. It is hence possible to change the ultrasound view cone during the intervention without impacting registration.

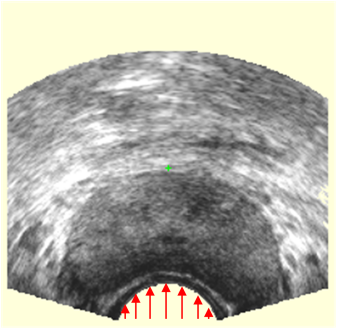

Figure 4.11: Probe assumptions. We assume that the probe displaced the tissues by the distances illustrated with the arrows, and that the displacements happened in the axial direction of the probe.

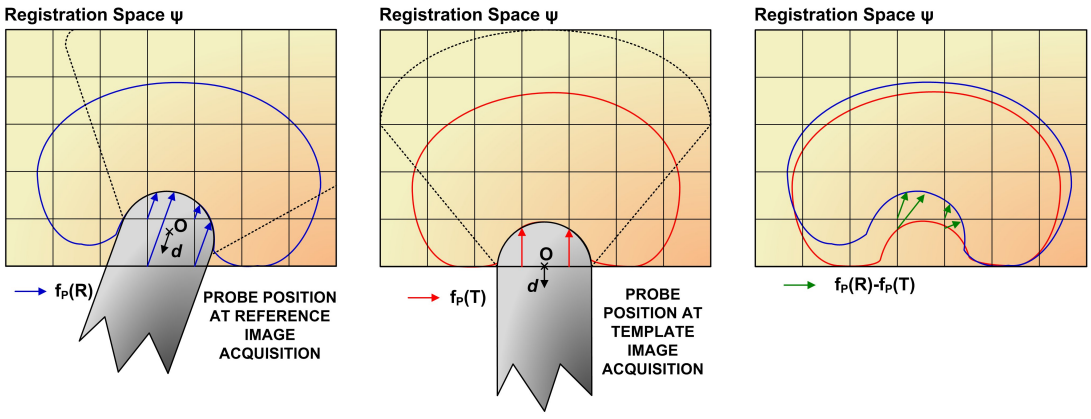

(a) (b) (c)

Figure 4.12: Probe-related tissue displacements. Fig. (a) shows the assumed tissue displacement trajectories caused by the probe at reference image acquisition, and (b) illustrates the assumed tissue displacement trajectories at template image acquisition. Fig. (c) shows the relative tissue displacements from the template image to the reference image, i.e. the displacements in the reference image minus the displacements in the template image.

$D \leftarrow A^2 - 4(\|x\| - r)\|d_I\|$
if $D \geq 0$ **then** {does x lie in the cylinder?}
  $A \leftarrow -(A + \sqrt{D})/(2\|d\|)$
  if $A \leq 0$ **then** {does x lie in front or inside the probe?}
    $F_p[I](x) \leftarrow Ad_I$
  end if
end if

The displacements $\Delta f_p := f_p[R](x) - f_p[T](x)$ are defined on the point $x + f_p[T](x)$, see Fig. 4.12 (c). It is thus necessary to distribute the displacements on the eight nearest neighbors of $x + f_p[T](x)$, linearly weighted by the distance of each neighbor from $x + f_p[T](x)$. This technique can be thought of as an inverse linear interpolation.

4.5.3 Force Capping

If we would set the probe force to Δf_p, it might become over-dominant with respect to the image-derived forces, which are capped to one voxel side length. Since the assumptions on which the probe forces are defined are an inexact model of reality, this is an undesired property. It is therefore preferable to cap the length of the probe forces also to one voxel side length, with the result that the probe forces can always be counter-balanced by image-driven forces during registration.

$$f_{probe} := \frac{\min\{1, \|\Delta f_p\|\}}{\|\Delta f_p\|}(\Delta f_p). \tag{4.64}$$

4.5.4 Experiments and Results

The probe forces have been tested within the framework

$$u^* = \arg\min_u \big(\mathcal{D}_{SSD}[R, T; u] + \mathcal{D}_{Probe}[R, T; u] + \mathcal{R}[u]\big), \tag{4.65}$$

where everything else is identical to the test configuration used in Sec. 4.4. The test results for the probe forces are given in Fig. 4.14 and in Fig. 4.13. One can observe that probe forces significantly reduce the error in the decile 10, where probe deformations are strongest. In all other deciles, the error is slightly worse than without probe forces. Nevertheless the overall RMS error is reduced from 3.19 to 1.21 when compared to registration without probe forces, and the error distribution is more regular.

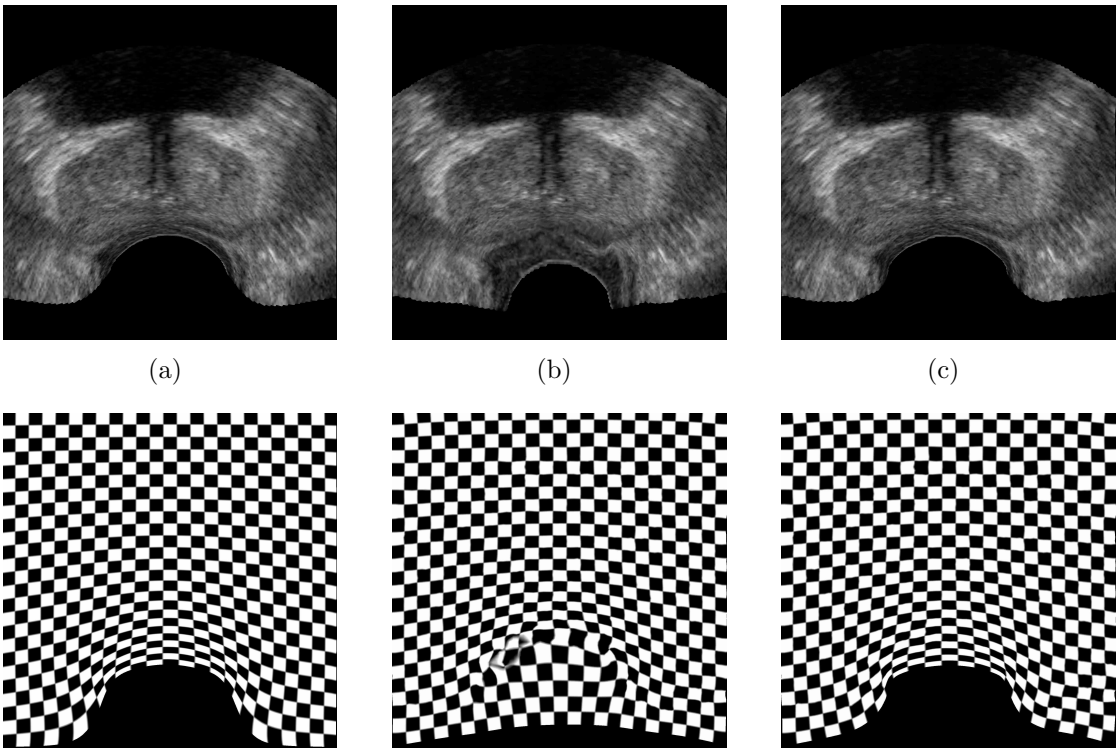

Figure 4.13: Probe forces. (a)+(d) illustrate the reference image, (b)+(e) the deformed template image after registration without probe forces and (c)+(f) the template image after registration with probe forces. The registration was carried out with a line search, $\nu = 0.3$ and $\Delta t_{max} = 0.25$.

4.5.5 Discussion and Conclusion

We found a way to incorporate approximative a priori knowledge about the deformation process into the registration framework, which makes it possible to register strong deformations that cannot be recovered with an image-based approach alone. The originality of our approach is that we derive the forces from assumptions on a virtual undeformed tissue state that is deformed during a probe insertion process, thus going beyond classical feature-pair based force terms. Our approach is a fusion of a bio-mechanical simulation of the major constraint applied to the prostate, the probe pressure, with an image-based registration process, with the limitation that we use displacements instead of physical forces to drive the bio-mechanical estimation process.

The presented approach can compensate for the absence of image information at the probe position in both images, where the relative tissue displacements be-

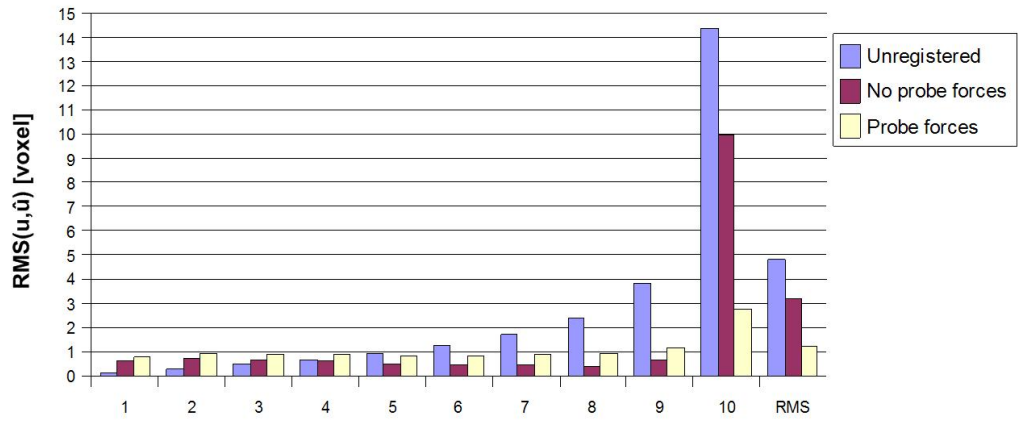

Figure 4.14: Probe forces. The figure illustrates the registration performance with and without usage of the probe forces. 1-10 are the decile RMS errors in voxel and RMS is the overall RMS error in voxel. The values of the unregistered transformation were computed with the identity transformation $u = Id$.

tween the reference and the template image are often strongest. It is modeled such that the bio-mechanical simulation does not over-dominate the estimation process and play a secondary role once the attraction basin of $\hat{\varphi}$ with respect to the image-based similarity measure is reached.

4.6 Consistent Inverse Computation

Independent forward and backward registration leads to transformations u and v which are not inverses of each other. We therefore couple the forward and backward registration process using an alternating optimization approach. Moreover, this approach makes it possible to profit from the fact that the forward and the backward registration are time-inverse, i.e. the problem is estimated from two directions. Coupling of both estimation processes hence leads to an improvement of the registration quality and can accelerate the convergence rate.

4.6.1 Gâteaux-Derivative of $\mathcal{D}_{cons}$

In our approach we have chosen the variational inverse consistency framework 4.13, which represents the best compromise between inverse consistency, compatibility with large deformations and computational burden for our application. At each

iteration of our multigrid algorithm, we first carry out a relaxation on the forward estimation problem, and then on the backward estimation problem. Additional coupling forces improve convergence and registration quality since the problem is estimated from two directions.

The forces are derived from the Gâteaux derivative of

$$\mathcal{D}_{cons}[u,v] = \int_\Psi \|\varphi_v \circ \varphi_u(x) - x\|^2_{\mathbb{R}^3}\ dx \tag{4.66}$$

at u in direction ψ. We consider the function $\varphi_v = v(x) + x$ and make advantage of the Taylor expansion of $\varphi_v(x + u(x) + h\psi(x))$ at the point $x + u(x)$, i.e.

$$\begin{aligned}
&\varphi_v(x + u(x) + h\psi(x)(x)) \\
&\quad = \varphi_v(x + u(x)) + h(\nabla\varphi_v)(x + u(x)) \cdot \psi(x) + \mathcal{O}(h^2) \\
&\quad = \varphi_v \circ \varphi_u(x) + h(\nabla\varphi_v) \circ \varphi_u(x) \cdot \psi(x) + \mathcal{O}(h^2)
\end{aligned} \tag{4.67}$$

Thus,

$$\begin{aligned}
&d\mathcal{D}_{cons}[u,v;\psi] \\
&\quad = \lim_{h\to 0}\frac{1}{h}\left(\mathcal{D}_{cons}[u + h\psi, v] - \mathcal{D}_{cons}[u,v]\right) \\
&\quad = \lim_{h\to 0}\frac{1}{h}\int_\Psi \|\varphi_v(x + u(x) + h\psi(x)) - x\|^2_{\mathbb{R}^3} - \|\varphi_v(x + u(x)) - x\|^2_{\mathbb{R}^3}\ dx \\
&\quad = \lim_{h\to 0}\frac{1}{h}\int_\Psi \|\varphi_v(x + u(x)) + h\nabla\varphi_v(x + u(x)) \cdot \psi(x) + \mathcal{O}(h^2) - x\|^2_{\mathbb{R}^3} \\
&\qquad\quad - \|\varphi_v(x + u(x)) - x\|^2_{\mathbb{R}^3}\ dx \\
&\quad = \lim_{h\to 0}\frac{1}{h}\int_\Psi h\langle\varphi_v(x + u(x)), \nabla\varphi_v(u(x) + x) \cdot \psi(x)\rangle \\
&\qquad\quad - h\langle x, \nabla\varphi_v(u(x) + x) \cdot \psi(x)\rangle + \mathcal{O}(h^2)\ dx \\
&\quad = \int_\Psi \langle(\varphi_v \circ \varphi_u(x) - x)^T(\nabla\varphi_v) \circ \varphi_u(x), \psi(x)\rangle\ dx.
\end{aligned} \tag{4.68}$$

Hence, the force term is

$$\hat{f}_{cons} = (\varphi_v \circ \varphi_u(x) - x)^T(\nabla\varphi_v) \circ \varphi_u, \tag{4.69}$$

where the forces are capped again such that their maximum length is 1, i.e.

$$f_{cons} := \frac{\min\{1, \|\hat{f}_{cons}\|\}}{\|\hat{f}_{cons}\|}\hat{f}_{cons} \tag{4.70}$$

Note that we set $\varphi_u \circ \varphi_v$ to Id if φ_u points outside Ψ, since φ_v is not defined there. We also tried out a more sophisticated force computation based on a line

search in direction $\hat{f}_{cons}$ in order to find an exact solutions in the interval $[0, 1]$, using a parabolic three point approximation scheme inspired by Brent's routine. The computational overhead, however, was massive due to the additional evaluations of $\varphi_u \circ \varphi_v$ involving trilinear interpolation of a displacement field. Fortunately, the proposed scheme 4.70 works well in practice and yields sensible improvements of the inverse consistency.

4.6.2 Experiments and Results

The performance of this approach is illustrated in Fig. 4.15, where the method was tested on our artificially deformed image, using the framework

$$\begin{aligned} u^* &= \arg\min_u \Big(\mathcal{D}_{SSD}[R, T; u] + \mathcal{D}_{probe}[R, T; u] + \mathcal{D}_{cons}[u, v] + \mathcal{R}[u] \Big) \\ v^* &= \arg\min_v \Big(\mathcal{D}_{SSD}[T, R; v] + \mathcal{D}_{probe}[T, R; v] + \mathcal{D}_{cons}[v, u] + \mathcal{R}[v] \Big) \end{aligned} \tag{4.71}$$

We measured the errors $\|\varphi_u \circ \varphi_v\|$ and $\|\varphi_v \circ \varphi_u\|$ for all φ_v pointing inside u, and for all φ_u pointing inside v, respectively. We were using the thresholds $\nu = 0,3$ and $\Delta t_{max} = 0.25$. We can observe that the inverse consistency error reduces significantly when adding f_{cons} to the registration forces (Fig. 4.15 (a)). Further, all decile errors indicate an improvement of registration accuracy in Fig. 4.15 (b) when using inverse consistency constraints.

4.6.3 Discussion and Conclusion

Inverse consistency constraints accelerate the registration process and lead to better estimates, since the problem is approximated from two directions. In addition, we get a backward transformation that is more consistent with the forward transformation then in the uncoupled approach.

4.7 Intensity homogenization

Until now we have tested our algorithms on a synthetic deformation $\hat{u}$, for which the identity assumption $R \equiv T_{\hat{u}}$ holds. This assumption is crucial for the validity of SSD-based image forces. Unfortunately, the identity assumption almost never holds for ultrasound images of the prostate. The principal reason is that the voxel intensity is strongest for interfaces orthogonal to the ultrasound wave direction. The intensity of the same tissue thus varies in function of the view angle. Another problem are changes in the ultrasound settings between acquisitions, which is typically the case for repeated biopsy series, where the time lapse between the

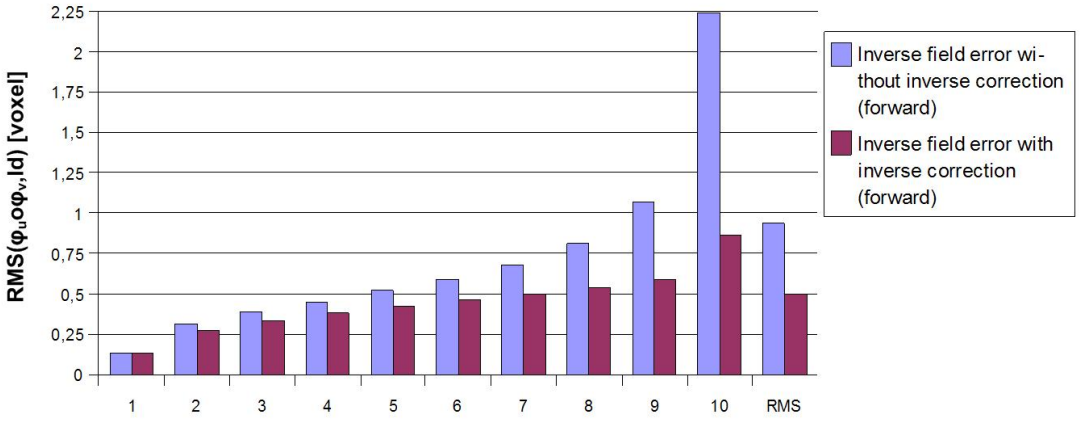

(a)

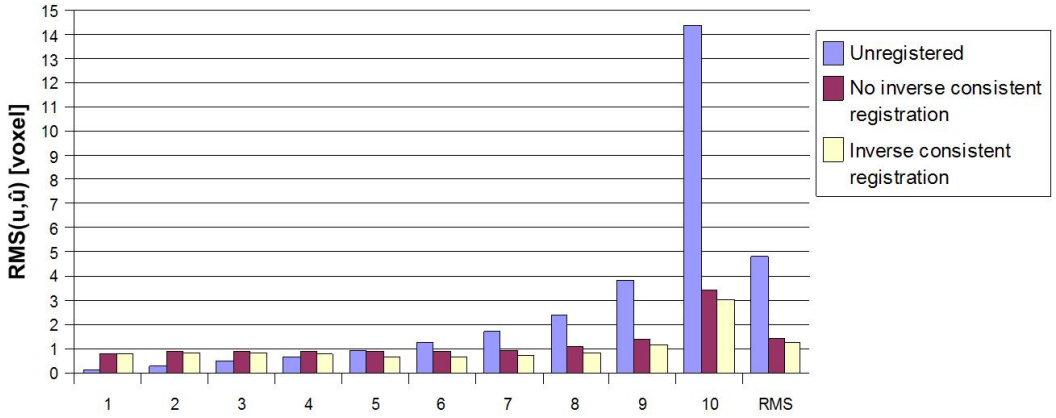

(b)

Figure 4.15: Inverse consistency forces. Fig (a) shows the decile analysis of the inverse consistency error $\|\varphi_v \circ \varphi_u - Id\|$ with and without inverse consistent forces, while Fig. (b) illustrates the convergence properties of the same configurations. 1-10 are the decile RMS errors in voxel and RMS is the overall RMS error in voxel. The values of the unregistered transformation were computed with the identity transformation $u = Id$.

acquisitions can be up to one year. Similar problems arise from noise like ultrasound shadows or speckle that can also vary in time. These variations between images lead to image distortions if not handled.

4.7.1 Local Intensity Shift Model

The obvious solution is to use a more sophisticated force term, e.g. the Pearson correlation. However, when using more complex correlation frameworks one has to carry out larger image convolutions for local force computation in order to get statistically significant results, which increases the computational burden of the registration process considerably. Moreover, the larger the convolution kernel, the less appropriate is the vector field representation of the transformation. This problem is well-known from block-matching, where the basis functions of φ are in general affine transformations.

Hence, we decided to use only a slightly more complex correlation model than the identity model, i.e.

$$R \equiv T_{\hat{u}} + b, \tag{4.72}$$

where $b : \mathbb{R}^3 \to \mathbb{R}$ models a local intensity shift. The shift can be locally estimated with a Gaussian convolution of $R - T_u$, i.e.

$$\begin{aligned} b_u^\sigma(x) &:= b^\sigma[R, T; u](x) \\ &= \int_\Omega (R(y+x) - T_u(y+x)) * \mathcal{G}_\sigma(y) dy \\ &= \int_\Omega R(y+x) * \mathcal{G}_\sigma(y) dy - \int_\Omega T_u(y+x) * \mathcal{G}_\sigma(y) dy, \end{aligned} \tag{4.73}$$

where σ^2 is the variance of the Gaussian. From the Gâteaux derivative of $\|R - T_u + b_u^\sigma\|$ at u we get the force term

$$f = (R - T_u - b_u^\sigma)(\nabla(T - b^\sigma))_u. \tag{4.74}$$

When using ∇T_u instead of $(\nabla T)_u$, the force term is

$$f = (R - T_u - b_u^\sigma)(\nabla T_u - (\nabla b^\sigma)_u). \tag{4.75}$$

4.7.2 Experiments and Results

In our experiments we have fixed the threshold to a discretized block neighborhood Ω^r_{sphere} with $r = 20$ and $\sigma = \frac{1}{3}r = \frac{20}{3}$. This choice was guided by the findings of the rigid-registration evaluation study, which has shown that inside the prostate, the maximum displacement error has an upper bound of approximately 3.5 mm. We

doubled this value, and multiplied it with the inverse of the most common voxel side length in our patient database, i.e. 0.35 mm. Considering that Term 4.75 is, in fact, a spatial high-pass filter on the image differences, we filter deformations with wavelengths greater than 7 mm. The registered images are the anchor volumes of two different biopsy series of the same patient. The images have many local dissimilarities, consisting in particular in speckle noise differences and local shadows. The Figs. 4.16 (a) and (d) show the result of the rigid registration. We can observe that the alignment of the prostatic structures is already of good quality, which means that the elastic registration should not find large deformations. In Figs. 4.16 (b) and (e) the registration was performed without homogenization. Without homogenization, the algorithm squeezes and distorts the upper half of the prostate (see the upper right quadrant of Fig. 4.16 (b)). The Figs. 4.16 (c) and (f) show the results with homogenization, which are significantly smoother.

An additional test was carried out with a Poisson coefficient of 0.47, i.e. the tissue was nearly incompressible. This leaves the regularization almost no room to compensate the local image dissimilarities with voxel squeezing and stretching. Due to the images differences, the registration process diverges completely for the chosen settings (see Fig. 4.17 (b)+(e)), despite the almost optimal rigid preregistration. When the local intensity homogenization is used, the registration process does not diverge any more (see Fig. 4.17 (c)+(f)). Further studies have to be done to analyze the impact of the homogenization on the convergence rate.

4.7.3 Discussion and Conclusion

The presented method represents a viable solution for the intensity homogenization problem, under the condition that an upper bound can be determined for the voxel displacements of $\hat{u}$, and that deformations are moderate. The latter requirement is, fortunately, attenuated by the fact that estimation of the most important deformations is assisted by the probe forces f_{probe}.

We are nevertheless not quite happy with this approach and we are looking forward to integrate different methods. An immediate idea is to carry out registration on both the original and the gradient magnitude image, in analogy to our rigid registration approach, knowing that the gradient image is free of the intensity shift bias and should therefore make registration more robust in the presence of local intensity shifts. Similar approaches have been presented by Foroughi et al. in [50], and by Wirtz et al. in [129]. Modersitzki et al. assume a multiplicative intensity bias $c \cdot T_u$ instead of an additive bias, and add $\|c\|$ as an additional energy term to the optimization framework [85]. These approaches will be investigated in the near future.

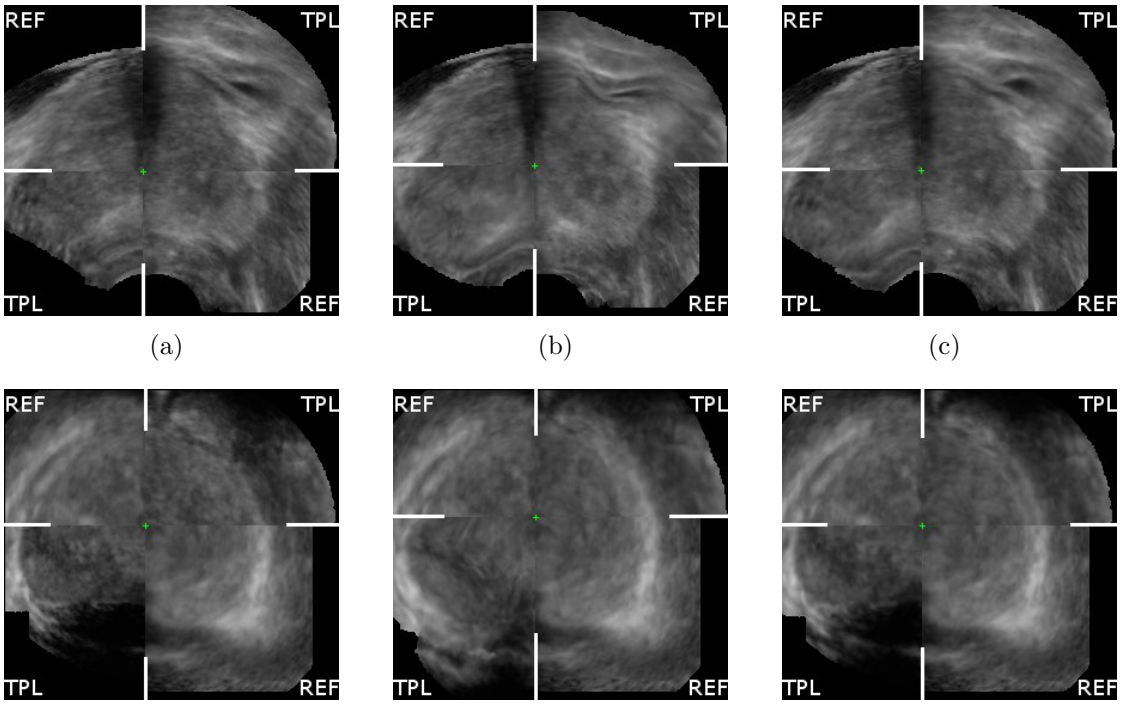

Figure 4.16: Intensity homogenization. The Fig. show the super-imposed reference (fixed) and template (moving) images. Fig. (a)+(d) show the rigid registration result (transversal and coronal cuts), Fig. (b)+(e) show the result of the registration without homogenization and (c)+(f) show the result of the registration with intensity homogenization. In Fig. (b) one can observe strong compressions in the upper right half that depicts the template image, whereas Fig. (c) is considerably smoother in this region, and the membranes are more accurately registered.

4.8 Experiments and Results

4.8.1 Accuracy Study

All elastic registration studies were carried out on a Pentium 4 single core processor with 3Ghz and 2GB of RAM. The images were acquired with the volume-swept GE RIC5-9 probe mounted on different GE ultrasound machines of the Voluson series. All images, except the images used for the creation of the panorama anchor image, were acquired immediately after a needle biopsy acquisition. All registrations have been carried out in a post-processing step. In contrast to the rigid registration study, we did not evaluate the elastic registration of O2D images with 3D anchor volumes, since we do not believe that 2D slices contain enough infor-

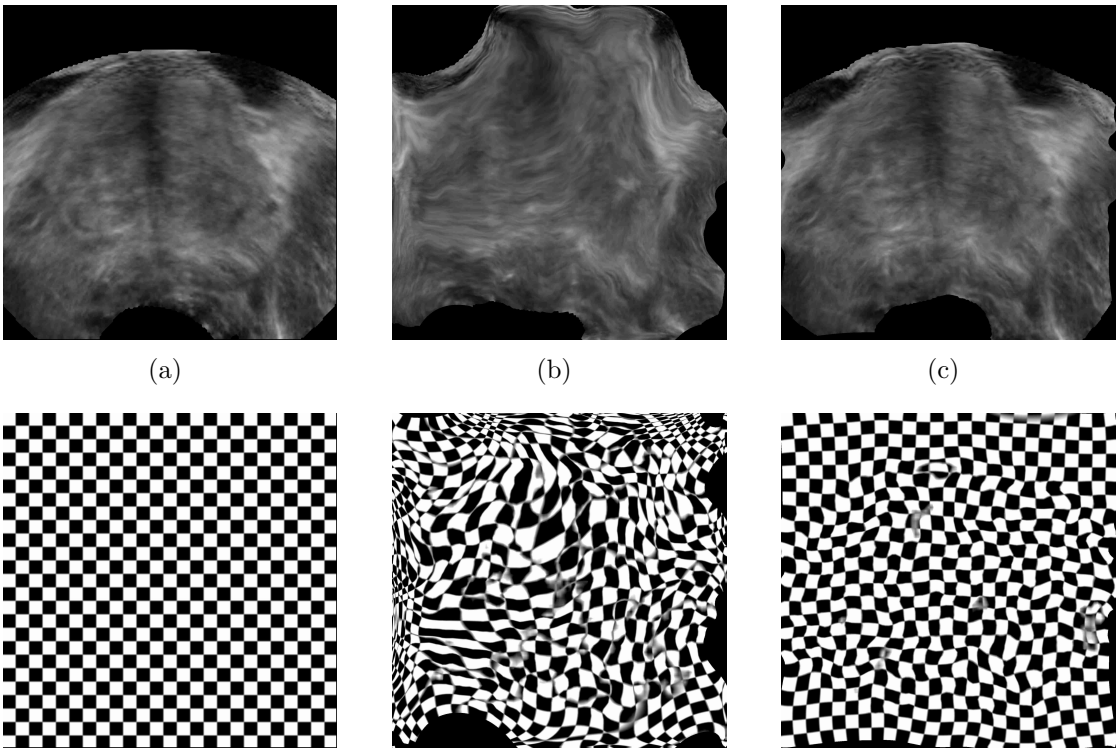

Figure 4.17: Intensity homogenization. Fig. (a) and (d) shows the template image before registration, (b) and (e) the registration result without intensity homogenization, (c) and (f) the registration result with intensity homogenization. The divergence was amplified by the fact that a Poisson coefficient of 0.47 was used for this test.

mation to achieve meaningful results with elastic registration, in particular if the image quality is not perfect.

The accuracy studies in this section were carried out on the 189 volume pairs that were used for the point fiducial reconstruction study of the rigid registration algorithm in Sec. 3.9. Elastic registration was executed starting from the results of the rigid registration accuracy study. The timing parameter Δt_{max} was set to 0.2, Poisson's coefficient was set to 0.3 and we chose 5 Gauss-Seidel smoothing steps per multi-grid iteration. We did not use a termination criterion, registration was stopped when the full multi-grid cycle was completely executed. All registrations were carried out on a six level resolution pyramid.

Four studies were carried out: first, the algorithm was executed on the resolution levels 6 ($7 \times 7 \times 7$) to 3 ($50 \times 50 \times 50$), without intensity homogenization and without probe forces (see Tab. 4.1 (b)). The same study was repeated with

		fiducial distance (RMS)	fiducial distance (max)	execution time (mean)
a)	rigid only	1.41 mm	3.84 mm	6.5 s
b)	6 - 3 wo homogenization and wo probe forces	1.81 mm	5.80 mm	11.7 s
c)	6 - 3 wo probe forces	1.09 mm	3.61 mm	15.5 s
d)	6 - 3 full	1.10 mm	2.93 mm	16.7 s
e)	6 - 2 full	1.03 mm	3.46 mm	38.7 s

Table 4.1: Registration accuracy. Line a) gives the results of the rigid registration on the test set, line b) the results of registration using SSD without intensity homogenization and without probe forces from the sixth-finest level ($7 \times 7 \times 7$) to the third-finest level ($50 \times 50 \times 50$). Line c) indicates the results when intensity homogenization is used without the probe force model, line d) the results when the full framework is executed, and line e) finally shows the result when optimization is carried out with the full framework from level 6 to level 2 ($100 \times 100 \times 100$).

intensity homogenization (see Tab. 4.1 (c)) and both intensity homogenization and probe forces (see Tab. 4.1 (d)). An additional study was performed using the resolution levels 6 to 2 ($100 \times 100 \times 100$), see Tab. 4.1 (e).

The first observation is that SSD without intensity homogenization degrades the accuracy achieved with rigid registration, with an increase of the RMS error from 1.41 mm to 1.81 mm, due to local image dissimilarities. The local intensity shift model performs well and can decrease the RMS error to 1.09 mm. The maximum error was also reduced, but this value has to be interpreted with care, since we noticed that it can fluctuate considerably and in an unpredictable manner even when the registration parameters are only slightly modified.

The performance of the probe insertion model was a little bit deceiving, since the overall RMS error is not improved when it is used. The model considerably improves the results for some patients, while it degrades accuracy for other patients. An explanation is the existence of image pairs that were acquired with considerably varying quantities of ultrasound gel in front of the transducer, a fact that we realized only during the final validation. In these cases, the distance between the probe heads computed with the rigid registration result can be quite large while the prostatic tissues near the probe head are almost perfectly registered, see Fig. 4.18 (a). The physical transformation of the matter between the capsule and the probe head does not correspond to an elastic deformation of tissue with low compressibility: between acquisitions, the gel is probably displaced on the probe

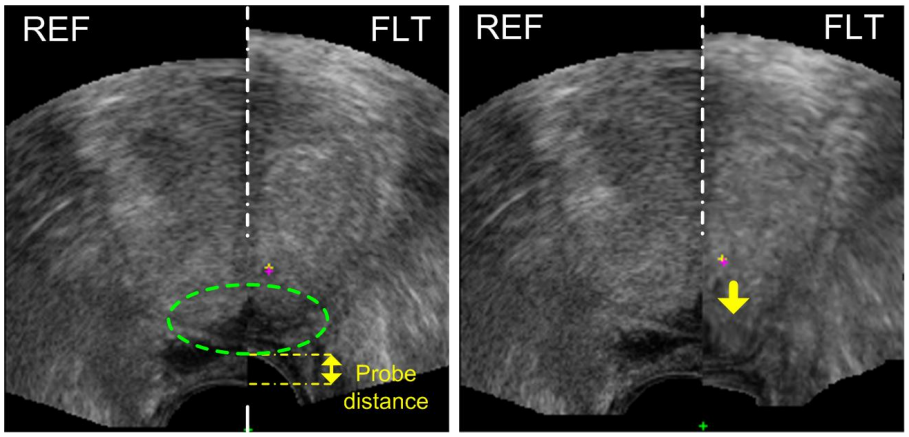

(a) (b)

Figure 4.18: Failure of probe insertion model. Fig. (a) shows the result of the rigid registration: the glandular structures near the probe head are almost perfectly aligned (green cercle), but the insertion depths of the probe heads considerably differ, probably due to a different quantity of contact gel in front the head during both acquisitions. The matter between the capsule and the probe head visible in the reference image has not been elastically squeezed in the floating image, but rather moved out of view. Linear elasticity is hence a grossly inaccurate model for the physical transformation of the matter in this region. Fig. (b) shows the result of the elastic registration: the floating image is elastically stretched and the capsule is torn downwards in a non-realistic way, leading to large mis-registrations of the lower part of the prostate. The image forces cannot counterbalance the stretching.

head and moves out of view. Hence, the elastic model is not valid in these image regions and the forces derived using the probe insertion model generate grossly inaccurate elastic deformations, as illustrated in Fig. 4.18 (b).

A further observation is that optimization on finer resolution levels does not yield significant improvements in accuracy. This can probably be explained with the unknown fiducial segmentation error. Due to this intrinsic error of the fiducial "gold standard", the RMS errors given in Tab. 4.1 overestimates the real registration RMS error.

4.9 Discussion and Conclusion

In this chapter we have presented an elastic registration algorithm that can efficiently estimate prostate deformations that occur typically during prostate biopsy acquisition. We have tried to incorporate a maximum of a priori knowledge into the framework without over-constraining the registration process. The principal

assumptions are elasticity of the deformation, presence of displacements caused by probe pressure, a local intensity shift model and inverse consistency. The combination of these models and assumptions in a single estimation process leads to a robust system with satisfying accuracy.

4.9.1 Biomechanical Simulation and Image-based Registration

The regularization framework 4.9 can be used for bio-mechanical simulations, when using physical forces, or for image registration, when using forces derived from image feature displacements. We have chosen to combine both approaches by introducing probe forces, derived from the relative probe positions computed with the rigid pre-registration, into the image-based registration process. Introduction of probe forces considerably improves the quality of the deformation estimation near the probe head.

The probe force model is, however, problematic when the quantity of ultrasound gel in front of the transducer differs considerably between both images. In that case, the elastic model is no longer valid in the probe head zone and the probe forces misguide the regularization, which leads to a local degradation of registration accuracy. This effect cannot be counterbalanced by the image-based forces. Unfortunately, we have no appealing solution to this problem.

4.9.2 Non-rigid registration and real-time

Simultaneous non-rigid forward and backward estimation of the tissue deformation can be carried out in about 10 seconds when the registration is stopped at the third-finest resolution level. The algorithm profits from the loss-containing intergrid transfer operators introduced for rigid multi-resolution registration: non-rigid registration can be robustly performed with up to 6 resolution levels when loss-containing techniques are used. Most computationally demanding parts of the presented algorithm consist in image convolutions. Image convolutions can be parallelized using stream processing techniques. This makes this algorithm a good candidate for integration into the novel hardware architectures recently presented by the major graphic chip manufacturers, which provide massive parallel stream computing capacities. We believe that a real-time implementation of this algorithm on modern graphic hardware is feasible, thus enabling real-time elastic deformation tracking.

4.9.3 Convergence rate vs. Smoothness

In Sec. 4.3 experiments were carried out to find the best compromise between the convergence rate of the optimization process and the smoothness of the solution u^*. The results of the experiments support the hypothesis that the norm of the force term should be smaller than the radius of the window on which it is computed. We start to feel comfortable with force lengths smaller than half the window size. When the force lengths are chosen too conservatively, the regularization dominates the forces, and the convergence rate deteriorates rapidly. From our experience, the forces in our framework should be allowed to have the maximum length of at least a quarter of the window size. The balancing between robustness and convergence remains, however, to some extent a subjective choice.

4.9.4 Diffeomorphism Property

We did not yet investigate if the diffeomorphism property of the estimate can be guaranteed within the presented framework, but this is an interesting problem. We think that the force limitation with respect to the window size ensures that the determinants of the local Jacobians of the estimated displacement fields do not get negative, but we still have to proof this property. Vercauteren et al. exploit in [124] the fact that diffeomorphisms form a Lie group and ensure the diffeomorphism property with algebraic considerations on the group exponential. It is possible that this approach can allow for larger force terms and hence faster convergence without a loss of smoothness.

4.9.5 Tissue Compressibility and Stability

In Sec. 4.3 we observed that registration gets instable with Poisson coefficients approaching the physical average for human tissue of 0.495, since local dissimilarities between the image pair cannot be compensated by voxel squeezing and stretching. This may lead to divergence of the iterative optimization process.

It seems interesting to investigate the effect of a lose coupling of the regularization energy with the image energy in analogy to the approach of Cachier et al. presented in Sec. 4.10. Such a model could be robust in the presence of noise, while still enforcing low tissue compressibility constraints with a sufficient degree of accuracy. To our knowledge, this approach has not yet been investigated for the linear elastic model. An alternative approach has been proposed by Modersitzki in [86], who applies the elastic regularization on the force field instead of the displacement field to obtain a less restrictive elastic regularization. We nevertheless think that the choice of a less restrictive compressibility constraint is a viable solution to the problem.

4.9.6 Intensity Mapping

We have proposed an intensity homogenization method that makes the assumption of local intensity shifts. In the context of intra-interventional tracking, this approach yields acceptable results as long as there are no significant ultrasound shadows or changes in the ultrasound parameters. In the context of inter-interventional registrations, the presented method lacks robustness, since the mapping model is too simple to yield good estimates in the presence of significant dissimilarities. Since inter-series registration is not subject to the real-time objective, it is possible to use more complex similarity measures. This approach will be investigated in the near future. We will also analyze additional image-based force terms, either feature based or based on different aspects of the images to be registered, e.g. by integrating force terms derived from the magnitude of the gradient image.

4.9.7 Panorama Anchor Image and Registration Quality

Even more than rigid registration, elastic registration strongly depends on the quality of the anchor image. Missing parts of the prostate or blurry panorama anchor images lead to oscillations and mis-registrations. It is absolutely necessary that the panorama image is of good quality. We hope that we can ensure this in the future with our novel anchor image acquisition protocol presented in Sec. 3.5.1.

4.9.8 Registration Accuracy

In the validation study, registration accuracy could be improved by 20 to 30 percent with elastic registration, depending on the parameterization. Elastic registration accuracy is probably understated due to the fiducial segmentation error, which makes it difficult to measure sub-millimeter accuracy and to compare fine-scale improvements of the algorithm. With a measured RMS error of 1.09 mm, the self-defined challenge to achieve sub-millimeter RMS accuracy was only missed by a small amount. Compared to the accuracy achievable with existing systems, one millimeter accuracy represents a considerable improvement, and it should be largely sufficient for most clinical applications that we target.

Chapter 5

Clinical Applications

Abstract

In this chapter we present a first clinical application of the tracking system. Biopsy maps have been created in anchor image space, which allows for post-interventional quality analysis. The distribution of the biopsies was statistically evaluated to analyze the targeting accuracy under 2D ultrasound control with respect to the systematic 12-core biopsy protocol. We also investigated 32 biopsy series performed by a single practitioner to answer the question whether a learning curve can be observed. We found that it is difficult to reach the targets accurately under 2D ultrasound control with statistical differences according to the regions of the prostate it concerns. We also found a statistically significant improvement in the targeting accuracy for the single practitioner when comparing the first 16 biopsies with the last 16 biopsies.

5.1 Introduction

Recall from Sec. 2.2 that prostate biopsies are traditionally performed under 2D transrectal ultrasound control according to a predefined pattern. The pattern is defined in the coronal plane and takes into account that most tumors are found in the peripheral zone of the prostate. It is based on the findings of statistical studies on the cancer distribution in radical prostatectomy specimens.

We present in this chapter a study on the needle placement accuracy under 2D transrectal ultrasound control. We were interested in the question whether it is possible to sample the gland according to the systematic biopsy pattern. This question has never been investigated before since the exact sample location could not be reconstructed after the acquisition. With the presented 3D ultrasound based tracking system, biopsy localization became possible, which made the study feasible.

5.2 Biopsy Acquisition: Clinical Protocol

The study was performed with a GE Voluson 730 and a RIC5-9 3D transrectal ultrasound probe. A 18-gauge needle with a cutting length of 23 mm was used. Before the intervention, a 3D anchor image of the prostate was acquired. Needle placement and the firing of the needle gun was performed as usual under 2D ultrasound control. Before retracting the needle containing the sample, the ultrasound device was switched into 3D mode, and a 3D image of the prostate with the needle inside was acquired. The clearly visible needles were manually segmented in a post-processing step and, after registration with the presented method, projected into the anchor volume to obtain the complete distribution of the samples. The biopsies were acquired following the 12-core protocol defined in Sec. 2.2.3, and the acquisitions were performed in a pre-defined order. This made it possible to identify for each 12-core target the corresponding needle path. The presented protocol acquired the permission of the ethical committee of the Pitié-Salpétrière hospital in Paris, and the biopsy acquisition protocol has been applied to 55 consenting patients so far. Four clinicians participated in the study, one trainee and three experts. 32 biopsies were acquired by a single expert clinician, which were separately analyzed to answer the question whether the biopsy distribution improves in time. The study was performed under the direction of Pierre Mozer, MD, PhD, clinician at the Pitié-Salpétrière Hospital in Paris. The statistical analysis was carried out by Pierre Mozer and Alexandre Moreau-Gaudry, MD, PhD, clinician at the Centre Hospitalier Universitaire de Grenoble, and member of the TIMC laboratory. In addition, 3D biopsy maps were computed for visual inspection, see Fig. 5.2.

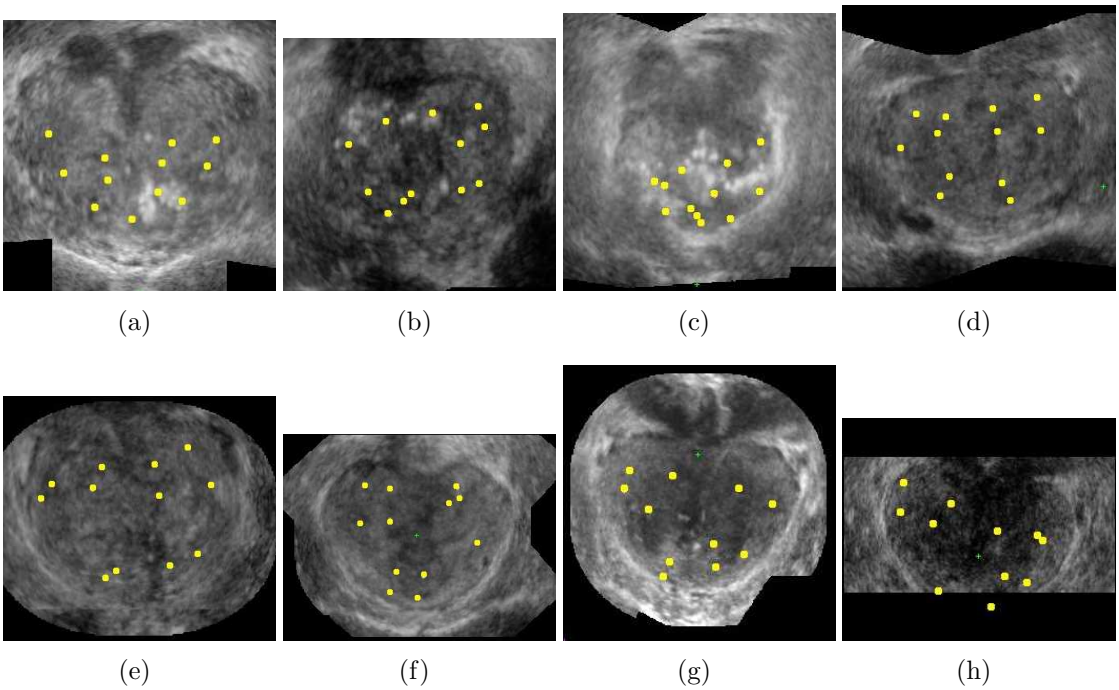

Figure 5.1: Coronal biopsy maps. The Fig. depict the coronal biopsy maps for 8 patients. While most maps correspond with good accuracy to the systematic scheme, the maps of the series (a) and (c) reveal a lack of sampling in the apex zone. In series (e), the central gland is not adequately sampled, and in series (c), (f) and (h), almost identical biopsy sites have been sampled multiple times.

5.3 Biopsy Maps

For a first quality control, the intersections of the trajectories with a manually chosen coronal plane of the anchor volume were computed. Fig. 5.1 illustrates the maps of 8 arbitrarily chosen patients, where 1 series was performed by a trainee (c), and the other 7 by three experts. In general, the biopsies are rather uniformly distributed and approximate the systematic protocol fairly well. The maps (a) and (c), however, reveal an undersampling of the base, and (e) an undersampling of the central gland. In the series (c), (f) and (h), almost identical sites have been biopsied multiple times. We can conclude from the visual investigation of the biopsy distributions that is difficult to target accurately under 2D ultrasound control.

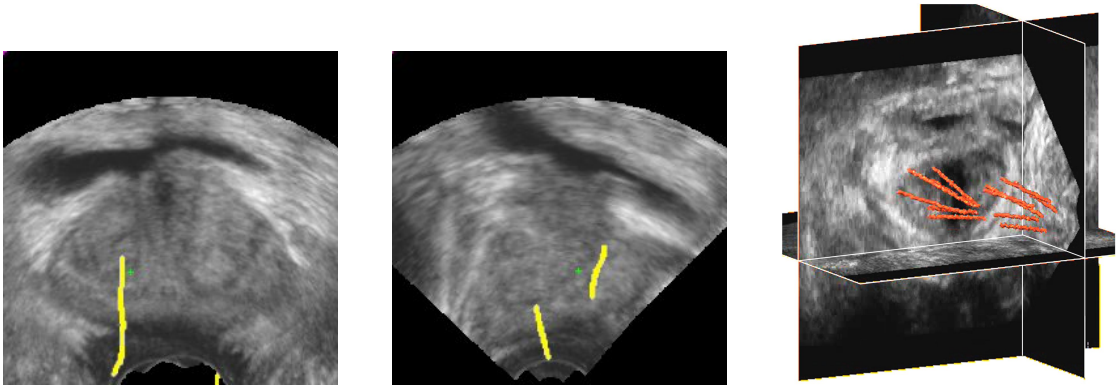

(a) (b) (c)

Figure 5.2: 3D biopsy maps. Fig. (a) shows a projected needle trajectory in the transversal plane, Fig. (b) shows sections of two trajectories in the longitudinal plane, and Fig. (c) shows a 3D visualization of the trajectories. Note that projected trajectories are deformed in the anchor images, which is due to the deformation correction.

5.4 3D Representation of the Systematic Protocol

A statistical study was performed on the accuracy with which the 12-core targets can be reached under 2D ultrasound control. Recall that the systematic 12-core protocol introduced in Sec. 2.2.3 is represented as a 2D scheme defined in the coronal plane of the prostate, see Fig. 2.9. To measure the length of the needle trajectories in each sector, we represented the protocol with 12 parallelepipeds, see Fig. 5.3. For each parallelepiped, we computed the intersection with the corresponding needle trajectory.

5.5 Targeting Accuracy

The first evaluation that we performed concerned the targeting accuracy. We considered a sector as successfully targeted, if the corresponding biopsy length in the sector was at least 3 mm. The results are given in Tab. 5.1. We can see that the percentage of biopsies that reached the base sectors BL and BC is only 62.5 percent and 64 percent. For the apex sectors AL and AC, which we combined for this particular study, the percentage is 34.4. The needle length quartile analysis in Fig. 5.4 shows that the needle lengths in the base sectors and in the apex sectors vary considerably. The results indicate that the apex and the base sectors are

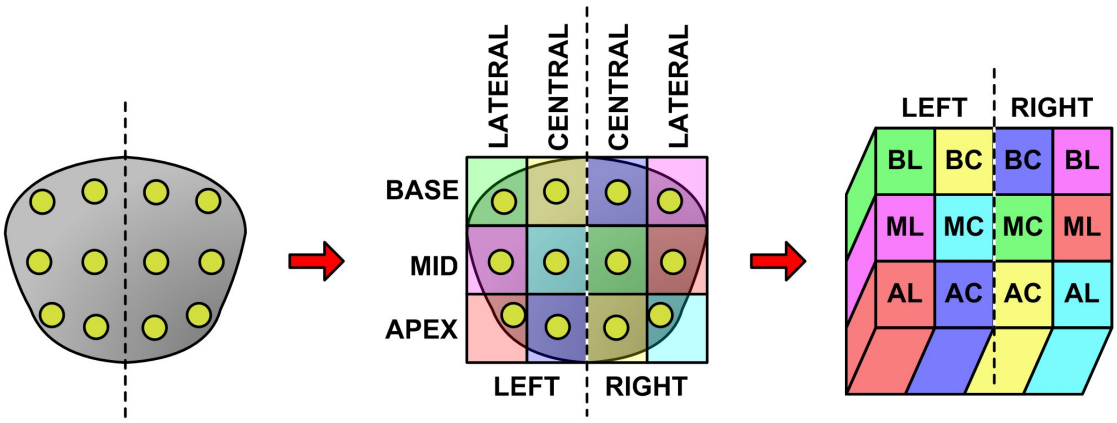

Figure 5.3: 3D sector representation of the 12-core protocol. The clinical 12-core protocol scheme is represented as a set of parallelepipeds, in order to make it possible to compute the needle lengths in each sector.

Sector	Total nb biopsies	Nb biopsies in sector	% biopsies in sector
BL	64	40	62.5 %
BC	64	41	64.0 %
ML	64	49	76.5 %
MC	64	60	93.8 %
AL+AC	119	78	65.5 %

Table 5.1: Targeting accuracy. The Tab. gives a first indication on the targeting accuracy with respect to the 12-core protocol under 2D ultrasound guidance. Targeting is considered successful when at least 3 mm of the sample lie in the sector. The AL and AC sectors have not been investigated separately.

significantly under-sampled. Note, however, that this study is biased due to the fact that the central mid-gland parallelepiped MC contains almost 100 percent of prostatic tissue, whereas the other sectors contain less tissue, due to the ellipsoidal shape of the prostate.

5.6 Operator Learning Curve

We also observed that the targeting skills of the operator improved from the beginning to the study (first 16 patients) to the end of the study (last 16 patients). Note that the operator received the coronal map of the biopsies after each intervention and thus had the possibility to auto-evaluate his performance. The Tab. 5.2

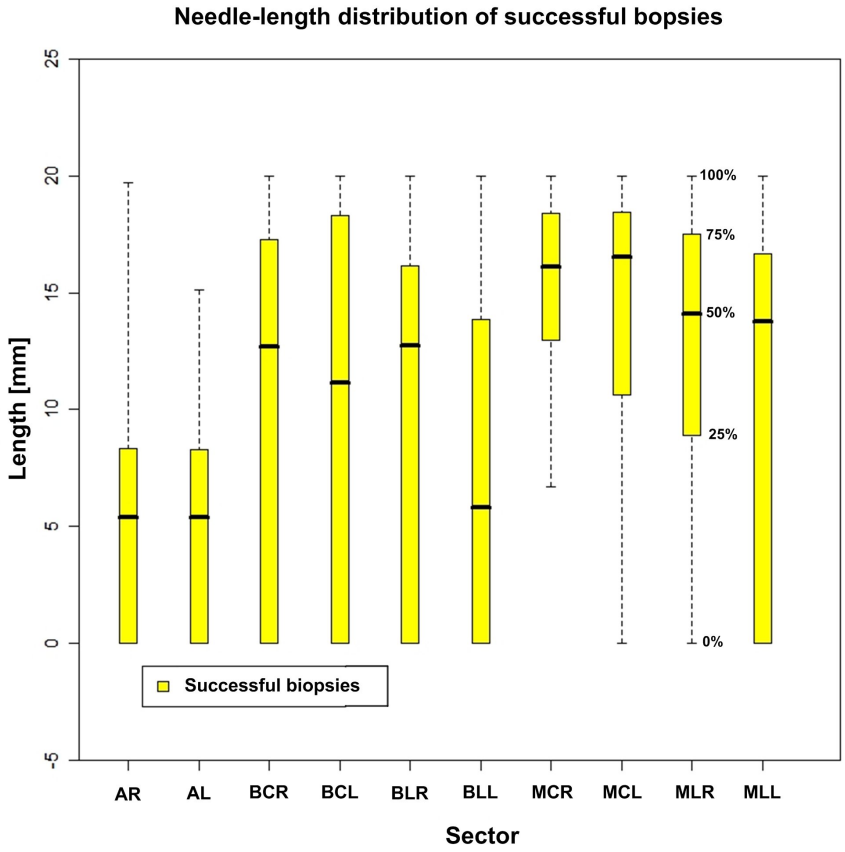

Figure 5.4: Needle lengths in the sectors. The Fig. indicates the mediane, minimum, maximum and inter-quartile lengths.

indicates a significant improvement in the median, mean, third quartile and maximum performance. In this evaluation, the lengths of the successful biopsies in each sector were added to obtain a single performance measure for each biopsy series. Tab. 5.3 shows the improvements per target sector, in particular for lateral targets.

5.7 Discussion and Conclusion

The 2D coronal maps are very well suited to provide visual feed-back about the sampling distribution during the intervention. The information is clear and does not lead to mis-interpretations. We are currently working on a software for clinical use, which will have the capacity of providing this type of feed-back about the biopsy distribution during the intervention. It is straight-forward to add targets

	Min	1st Qu.	Median	Mean	3rd Qu.	Max
16 first patients	36.60	72.73	90.38	91.69	109.70	130.50
16 last patients	31.5	101.8	121.0	110.0	137.1	155.8

Table 5.2: Operator learning curve. The table indicates the global needle length distribution for the 16 first patients and the 16 last patients. The needle lengths of the sectors were added to obtain a global performance measure. All lengths are in mm.

Sector	First 16 patients		Last 16 patients	
	Total nb biopsies	% biopsies in sector	Total nb biopsies	% biopsies in sector
BL	32	53.1	32	71.8
BC	32	56.2	32	71.8
ML	32	78.1	32	75.0
MC	32	96.8	32	90.6
AL+AC	59	55.9	60	75.0

Table 5.3: Operator learning curve. The table illustrates the targeting accuracy improvement between the first patient cohort and the second patient cohort.

to these maps for target navigation.

3D biopsy maps are less intuitive and are hence less appropriate to guide the gesture or to provide feed-back about acquired biopsies, since the clinician will not want to interact excessively with the software during the intervention. Interaction is, however, necessary for most 3D visualizations to ensure the correct positioning of the virtual camera, which needs, in addition, to be updated multiple times during the intervention. This is clearly not practical. 3D biopsy maps are, however, useful for therapy planning where the clinician can freely manipulate the virtual camera.

Visual inspection of the 2D coronal biopsy maps reveals that it is difficult to target the 12-core sectors accurately under 2D transrectal ultrasound control. Biopsy sites were sampled multiple times, and large zones have been left unexplored. This increases the probability of missing a tumor, i.e. the biopsy sensitivity is lower than it would be if the samples were more regularly placed. The statistical analysis of the biopsy trajectories that effectively reached the targeted sectors seem to confirm these findings. We are, however, aware of the bias introduced in the study by the 12-parallelepiped representation of the prostate. We are currently working on a software that allows us to define the 12-core protocol on a segmentation of

the prostate, which will lead to more accurate statistical evaluations.

The improvement of the operator performance is, however, unbiased and represents an interesting finding. The operator was already considered as an expert before the beginning of the study. The fact that he received after each intervention the corresponding biopsy map made it possible to evaluate his performance. It is therefore possible to formulate the hypothesis that the quality of the biopsy distribution can be improved even with post-interventional visual feed-back.

Chapter 6

General Conclusion and Perspectives

6.1 Prostate Tracking

The presented system provides the possibility to track prostate displacements and deformations with an RMS accuracy of 1.1 mm in 16.7 s. The average tracking success rate lies between 96 and 97 percent and should be further increased with the introduction of a quality control during panorama image acquisition.

We have developed strategies to reduce the loss of information in multi-resolution approaches that enables the estimation of the major part of the transformation $\hat{\varphi}$ on very coarse levels. Without information loss handling, it is difficult to envision a robust real-time tracking system based on deformable registration, in particular in the presence of complex image masks, as it is the case of transrectal ultrasound images of the prostate.

Moreover, a kinematic model of the probe movements under endorectal constraints was conceived that makes advantage of the fact that probe contact with the target organ is a necessary condition for prostate image acquisition. The model can be used to solve the patient movement problem completely and independently of external tracking systems, due to the fact that it defines a reduced subspace of the rigid search space that can be exhaustively explored in a very short time frame at low resolutions.

Furthermore, a panorama anchor image acquisition protocol was developed with the objective to ensure that the anchor image contains the entire prostate. A clinical protocol was proposed that can be used to verify the correct acquisition of the panorama image. Anchor images that contain the entire prostate make the registration of tracking images that contain only a small part of the prostate more robust.

The rigid registration was considerably improved with the elimination of zones of strong deformations near the probe head and of the bladder.

Concerning deformation estimation, a variational framework with a linear elastic regularization model was conceived that combines image-based forces, probe-based forces and inverse-consistency forces. An iterative solution scheme was conceived such that the balance between the elastic regularization and the force terms can be controlled with a single timing parameter. Moreover, we made advantage of the fact that Young's modulus has no physical meaning in an image-based registration framework and used it as free parameter. This makes it possible to define Poisson's compressibility parameter without impacting the convergence rate of the iteration process.

An a priori model on the tissue displacements during probe insertion is used in combination with the knowledge about the positions of the tracking and anchor image probe heads in the reference space, known from the rigid pre-registration. The model is used to simulate the probe insertion simultaneously to the image-based registration process, combining hence bio-mechanical simulation with image-

based registration. The model-based forces make it possible to estimate strong deformations near the probe head, where only few image information is available. The forces derived from the model are formulated such that they do not dominate, but rather assist the image-based forces.

A novel line search scheme is proposed for the computation of the image-based forces in order to replace a force scaling parameter that is difficult to determine. The line search yields exact solutions in direction of the used gradient, and hence improves the convergence rate and accuracy of the iterative estimation process.

Finally, a novel intensity homogenization approach is presented that makes it possible to register volumes of the same prostate that have local dissimilarities caused by noise, different ultrasound settings or changes of the tissues. The intensity homogenization approach models the dissimilarity as a local intensity shift and represents a high pass filter for the image dissimilarities.

The proposed system is robust and fast enough to achieve real-time tracking with a high degree of accuracy, at the condition that the automatic registration validation process is solved.

The clinical contributions of our work consist in the creation of biopsy maps for more than 50 patients which made it possible to visualize and control for the first time the quality of biopsy distributions that are acquired under 2D ultrasound control.

A statistical analysis was performed on these maps that shows clearly that some zones of the gland are not sufficiently sampled under standard 2D ultrasound control. We were also capable to show that the biopsies are more accurately placed with increasing experience of the practitioner.

An additional personal computer has to be installed in the operating room for execution of the tracking software. The multiplication of devices in the operating room has often negative impacts on the ergonomy of computer-assisted medical solutions, due to spatial constraints and the splitting of visual information on multiple devices. In the context of prostate tracking, a solution to this problem is the integration of the tracking system into the ultrasound hardware.

The system entirely solves the prostate motion problem, thanks to a priori models derived from assumptions on endorectal probe movements and on the necessary conditions for prostate image creation. It is hence neither necessary to fix the patient pelvis nor to use external tracking devices for probe movement identification. The latter is an important advantage of our system since external tracking systems can considerably constrain the clinical gesture (due to occlusion problems for optical tracking systems) or complicate the clinical setup (elimination of ferromagnetic interferences for magnetic systems, placement of the camera for optical tracking systems).

The system is robust enough to track central and lateral biopsy images reliably,

and the algorithm executes fast enough to be a candidate for real-time tracking, for example with an implementation on specialized hardware. Moreover, no additional sterilization issues arise with the 3D ultrasound tracking system.

The clinical standard procedure is only slightly modified with the proposed system. An important additional step is the placement of the bounding box, necessary to initialize the endorectal probe movement model, the probe deformation model and the bladder elimination algorithm. Note that the clinicians currently already define a 2D bounding box for prostate volume measurement. Introduction of a 3D bounding box is hence not a significant modification of the clinical routine and should easily be accepted by the clinicians.

A more time-consuming novel task is the manual validation of the registration, which is almost impossible to be performed during biopsy acquisition. It depends on the application whether manual validation is a viable solution or not.

The main modification that we introduce is the usage of 3D ultrasound instead of 2D ultrasound. Nowadays, there exist only few clinical applications that make advantage of the abundant information provided by 3D ultrasound, the most widely used being 3D fetus reconstruction for prenatal diagnosis. Ultrasound guided brachytherapy also moves towards 3D imaging with the increasing resolution of the mechanical steppers used for acquisition of integral images of the gland via 2D ultrasound sweeping. However, for most clinical interventions, 3D ultrasound imaging currently does not represent a revolutionary improvement, since 3D rendering of compact tissues does not yield easily interpretable results. In our opinion, the primary advantage of 3D ultrasound is that it makes reliable and nonrigid real-time organ tracking feasible. From a technical point of view, 3D ultrasound is barely more complex than 2D ultrasound and ultrasound should hence remain one of the most cost-effective medical imaging modalities. This makes ultrasound-based computer-assistance system very attractive when compared to MRI-based approaches whenever ultrasound images do provide enough information for a given target organ, in particular in the context of diagnosis.

6.2 Applications and Perspectives

In this section we discuss several applications that can be realized with 3D ultrasound-based prostate tracking.

6.2.1 Quality Control

Post-interventional sampling quality control is a straight-forward application of the tracking system which is not confronted with additional scientific challenges. It can be realized with the 2D and 3D biopsy maps that we introduced in Chap. 5.3, where

we think that 2D maps are best suited for this purpose. The maps make it possible to identify unsampled and multiply sampled zones. Detection of multiply sampled zones is important since multi-sampling can lead to false conclusions about the tumor size when the concerned samples contain cancerous tissues. Undersampling can be corrected during the intervention when the user has an immediate feed-back about the sampling locations. We are currently working on a software that provides this sort of feed-back.

6.2.2 Cancer Distribution Map

For diagnosis it is interesting to know the exact cancer distribution map. It can be obtained by establishing the spatial correspondence of the results of the histological biopsy analysis with the biopsy maps. From a scientific and engineering point of view, the implementation of such an application is straight-forward with the proposed tracking systems. From a logistical point of view, the digital patient data collected with the tracking system has to be managed and stored, and the histological results have to be fused with the biopsy maps, which adds a task to the clinical routine. Distribution maps of the histological information will not only incorporate the needle trajectory, but also the real length of the samples and the position of the cancer in the samples. This information may be very useful for further analysis.

6.2.3 Augmented Reality and Assistance

A primary objective of this work was to provide a tool which makes it possible to envision assistance of the clinician during the medical intervention. Currently, prostate biopsies are acquired almost "blindly": the operator can neither see the tumor, nor does he have an accurate idea about the current needle position with respect to previously sampled tissues and the planned targets. A very helpful assistance would consist in the real-time visualization of theses information in the ultrasound images that are consulted by the clinician for needle positioning. In other words, the ultrasound images would be augmented with additional information that could significantly improve the accuracy of the gesture.

Real-time assistance requires real-time access to ultrasound volumes, which is not provided by most ultrasound manufacturers, and a real-time implementation of the tracking system. While we think that these two points are minor obstacles, a third one represents a real scientific challenge: automatic registration validation. Manual validation performed by the user is still imaginable for quality control and the construction of cancer distribution maps, but it must be excluded as option for assistance systems. This interesting topic will be amongst the foci of our future research efforts.

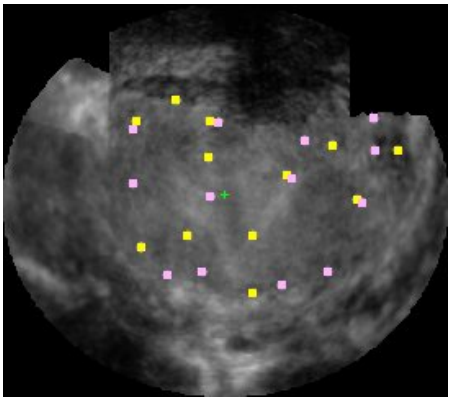

Figure 6.1: Repeated biopsy series. The yellow dots show the biopsy sites sampled during the first biopsy series, and the pink dots show the sites of the second series.

6.2.4 Repeated Biopsy Series

If the open issues of augmented reality can be solved, then it is feasible to project the sites of a previous biopsy series into the real-time ultrasound planes used by the clinician for targeting, after registration of the anchor volume of the previous biopsy series with the current anchor volume. This makes it possible to avoid re-sampling of tissues, and should hence considerably increase the sensitivity of repeated biopsy series. Another potential clinical application is to re-sample biopsies near a spot where cancer was detected with the previous series, in order to get a better idea about the size of the cancer. Fig. 6.1 illustrates two super-imposed biopsy series. One can observe that multiple sites have been biopsied twice, and that the regularity of the combined biopsy distribution is sub-optimal. Another feature could be to propose optimal needle positions in unsampled regions automatically.

6.2.5 Prostate Biopsies on Atlas-derived Targets

Shen et al. constructed in [108] a statistical tumor atlas from radical prostatectomy specimens with clinically localized cancers. He also developed a statistical predictive model for prostate cancer, which can theoretically achieve a sensitivity of 99 percent with only 7 needles. In practice, the problem is to find the targets in the ultrasound images. If it is feasible to register such an atlas with the gland in the anchor image, the predictive model could be used to define the cancer targets, and the biopsy sensitivity could be significantly reduced while the invasiveness of the procedure is decreased.

6.2.6 MRI Target Navigation

Another interesting application consists in MRI target navigation. The idea is to register an MRI or MRSI image of the prostate with the anchor image, and to super-impose the MRI-image on the tracking image stream during the intervention using the tracking estimates, see also Fig. 6.2. The MRI/MRSI to anchor image registration can be done off-line and is hence not time-critical. Further scientific investigations are, however, necessary to conceive an automatic or semi-automatic multi-modal registration algorithm. We are currently working on a project with the objective of MRI-ultrasound prostate image registration. MRI-targeted prostate biopsies could accompany or even replace the systematic protocol, and they should help to increase the sensitivity of prostate biopsies.

6.2.7 Focal Prostate Cancer Treatments

The current standard prostate biopsy procedure leaves the clinician with two uncertainties: first, the size and shape of the tumor can only be roughly identified, and second, the location of the samples itself is only very approximatively known. This makes it impossible to localize tumors precisely and hence poses great challenges to focal prostate cancer treatments, see also the conclusion of Ahmed et al. in [1]. Prostate tracking removes the uncertainty about the sample location, which leads to the interesting question whether this is sufficient to implement novel prostate therapy methods that are clinically convincing. An idea is to use statistical data about tumor sizes and shapes for example in function of the Gleason score of tumorous biopsy samples. This could be implemented with a statistical cancer atlas, for example with the cancer probability atlas of Shen et al. [108]. Recall from Sec. 2.1.6 that currently the treatment options in the presence of positive biopsies are either careful watching or radical therapy. While novel therapy methods seem to reduce the severe side-effects radical prostate therapy, they are still considerable, and the patient has to make a delicate choice. Focal therapy could represent an escape of this dilemma.

6.2.8 Brachytherapy

Until now we have only discussed prostate tracking for biopsy localization. Many therapeutic interventions are, however, also subject to the prostate motion problem, and they are performed under ultrasound control (see Fig. 6.3. In brachytherapy, irradiating seeds are implanted into the prostate via transperineal needle access, and the needle placement is controlled with lateral transrectal 2D ultrasound. With a single needle, multiple seeds can be placed on a line. Needle insertion leads to deformations and displacements of the prostate, and the prostate grows in vol-

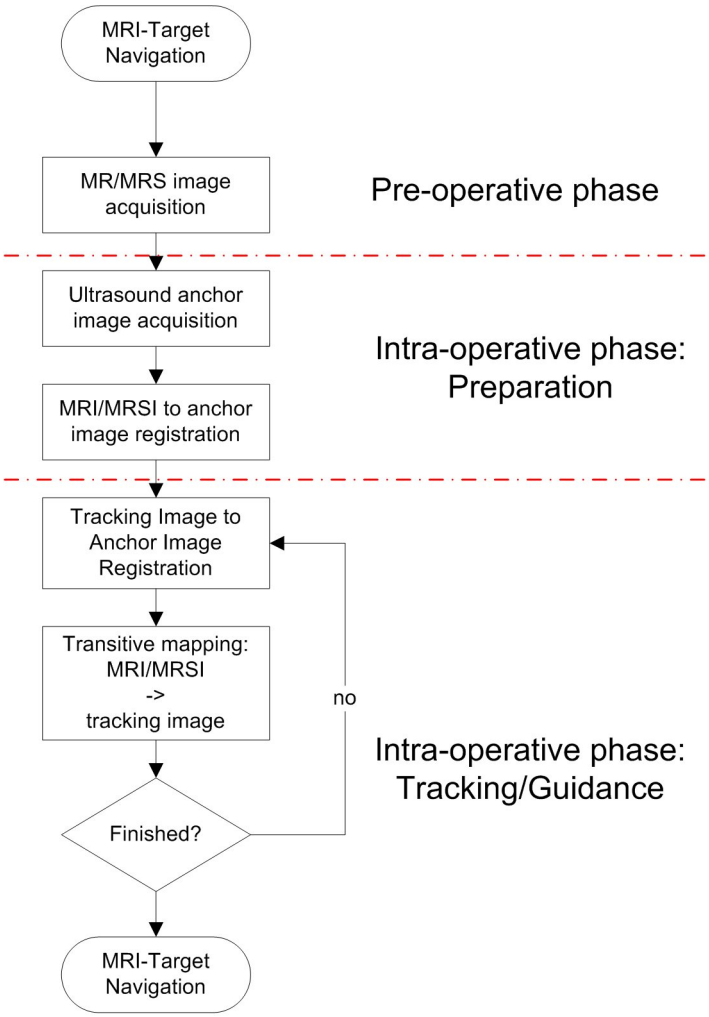

Figure 6.2: MRI target navigation. The flow chart illustrates the different steps of MRI/MRSI navigation based on ultrasound tracking. The MR images are registered with the anchor image. During the intervention, the MR to tracking image transformation is computed transitively.

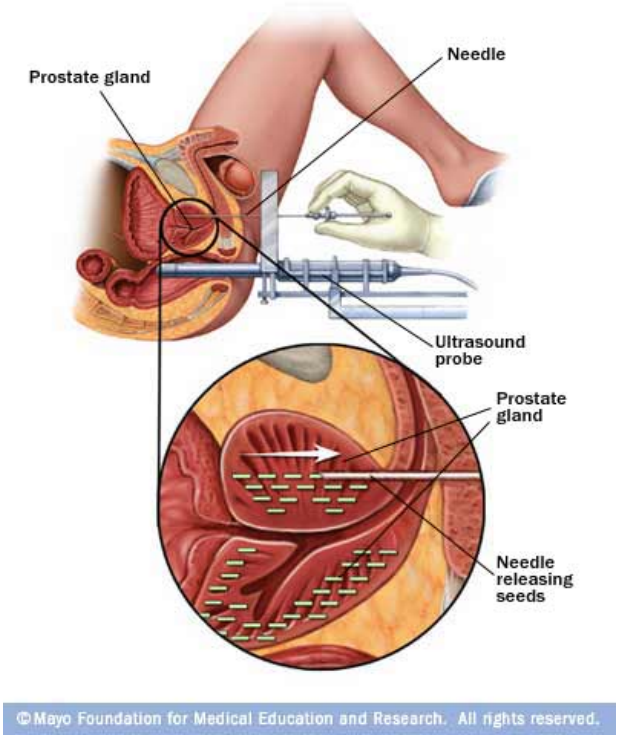

Figure 6.3: Brachytherapy. The Fig. illustrates the needle insertion with a needle template and seed placement. [Image found at www.content.revolutionhealth.com]

ume during the intervention due to the implanted seeds. This makes it necessary to update the seed distribution planning continuously during the intervention. The update is not trivial since the seeds are only barely visible on the ultrasound images, and their exact location is thus unknown. A prostate tracking system could identify the motion of the gland and communicate it to the planning software. The planning update could be carried out automatically. From a technical point of view, registration problems may arise due to needle visibility in the tracking images and due to the prostate volume growth caused by needle and seed insertions.

6.2.9 High Intensity Focused Ultrasound Ablation

Another candidate for prostate tracking could be HIFU prostate cancer ablation. HIFU is an ultrasound-guided intervention that cooks tissues with high intensity focused ultrasound. HIFU can burn non-prostatic tissues or critical structure like the rectal wall and the urethra when the prostate moves after the planning was defined. Prostate movements during HIFU therapy are mainly caused by

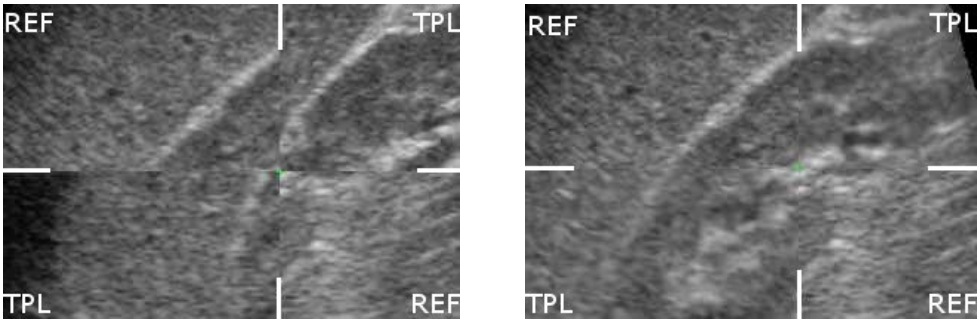

(a) (b)

Figure 6.4: Kidney registration. Fig. (a) depicts two kidney volumes before, and Fig. (b) the same volumes after elastic registration.

respiratory movements, which could be estimated with the tracking system. This would require the integration of a 3D ultrasound transducer into the HIFU system, which is potentially problematic due to the fact that it has to be integrated into the HIFU emitter.

6.2.10 Other Organs

3D ultrasound-based tracking is not limited to the prostate, registration algorithms have been presented for the liver [49, 133] and we made experiments with kidney registration that yielded encouraging results, see Fig. 6.4. The prostate-specific models will have to be replaced with models specific to the target organ. For many organs, it seems challenging to find a similarly efficient approach than the endorectal probe movement model for the prostate. In these cases it is probably preferable to work with an additional external tracking system to find a point inside the capture range of $\hat{\varphi}$. The kinematic model can, however, easily adapted to most applications that use endocavitary probes. This applies for most pelvic organs and many interventions in gynecology.

6.3 Conclusion

A large variety of potential applications for the prostate tracking system in the context of subcutaneous soft tissue interventions have been enumerated in the preceding sections. The main scientific issues for the implementation of some of these interventions are automatic validation of the tracking estimates and registration of the anchor image with images of different modalities or with a statistical cancer distribution atlas. We hope that we will be given the possibility to address these

open questions in the future, and that the outlined applications will find their way into the operating room one day.

Appendix

Appendix A

3D Ultrasound Calibration

Michael Baumann, Vincent Daanen, Antoine Leroy and Jocelyne Troccaz.
In proceedings of CVAMIA'06 - 2nd International Workshop on Computer Vision Approaches to Medical Image Analysis - ECCV'06, Graz, Austria, 2006.

Abstract

With the emergence of swept-volume ultrasound (US) probes, precise and almost real-time US volume imaging has become available. This offers many new opportunities for computer guided diagnosis and therapy, 3-D images containing significantly more information than 2-D slices. However, computer guidance often requires knowledge about the exact position of US voxels relative to a tracking reference, which can only be achieved through probe calibration. In this paper we present a 3-D US probe calibration system based on a membrane phantom. The calibration matrix is retrieved by detection of a membrane plane in a dozen of US acquisitions of the phantom. Plane detection is robustly performed with the 2-D Hough transformation. The feature extraction process is fully automated, calibration requires about 20 minutes and the calibration system can be used in a clinical context. The precision of the system was evaluated to a root mean square (RMS) distance error of 1.15mm and to an RMS angular error of 0.61°. The point reconstruction accuracy was evaluated to 0.9mm and the angular reconstruction accuracy to 1.79°.

1.1 Introduction

Until recently, 3-D ultrasound (US) volumes had to be manually reconstructed from a number of 2-D US slices acquired while slowly moving a 2-D probe over the target region. The so-called *3D freehand method* is time-consuming, imprecise and not usable for many clinical applications requiring real-time acquisition [103]. The emergence of 3-D swept-volume US probes solved most of the enumerated problems: a mechanical device capable of sweeping the 2-D crystal array of the probe over a target region makes it possible to acquire 3-D US volumes accurately and almost in real-time (1s to 4s per acquisition)[1].

These new capabilities open an entire new field of applications in the domain of computer guided medical interventions based on US imaging. One can imagine tool guidance systems that would operate with permanently updated US volumes, visualizing for instance slices at the tool tip position. More sophisticated applications could carry out target localization inside the volumes through real-time registration and segmentation techniques, thus allowing to match pre-operative planning with intra-operative data.

However, US-based guidance often requires knowledge about the position and orientation of the US volume in space. When using a tracking system this can be achieved by calibrating the US acquisition volume with a tracking reference fixed on the probe. Unfortunately it is virtually impossible to derive the calibration parameters directly from the geometry and parameterization of the probe. Almost all existing calibration systems rely therefore on statistical or segmentation-based object matching methods.

1.1.1 Calibration Methods Overview

A variety of techniques for 2-D US calibration was proposed in the literature; a comprehensive review being given in [84]. Calibration methods can be classified with respect to the target geometry they rely on. *Single-point target* methods identify a point, i.e. a bead, a calibrated pointer tip or a cross-wire, in the US image [88,97,111]. The difficulties consist in automatic geometry extraction in the US slice and US beam alignment with the phantom.

Multi-point target phantoms are extensions of the single-point bead or cross-wire phantoms. They consist of a number of point targets with precisely known coordinates in phantom space. Their geometric configuration makes it possible to derive the calibration parameters from the distances between the reconstructed

[1]However, most currently available systems don't yet provide a real-time data transfer interface for 3-D data. Nevertheless one can acquire so called "4-D" images of three orthogonal volume slices in real-time using a video-capture device. In the rest of this article we make abstraction of this restriction, hoping that it will disappear with the next generation of 3-D echographs.

intersection points visible in the 2-D US scan [74]. Compared to single-point phantoms they require less image acquisitions due to their more discriminative geometry, but share the phantom alignment and feature extraction problems.

Z-fiducial or *N-fiducial* phantoms address the alignment problem of point target methods. A calibration point is determined from the intersection points of a number of nylon strings with the US beam. This is possible due to a sufficiently discriminative wire geometry [39, 77, 94]. Fiducial methods are more robust than point target methods but the difficulties concerning fully automatic feature extraction subsist. Also, Z- or N-fiducial phantoms require a high manufacturing accuracy to achieve a satisfying calibration quality.

Wall phantom methods are based on detection of the intersecting line of a planar surface with the 2-D US beam. In [97], a water tank bottom is imaged for calibration. The authors of [117] address the reverberation and line thickness problems inherent of wall phantoms by using a membrane variant. Both phantoms have difficulties when confronted with steep angles between the US beam and the plane because they cause line intensity and line sharpness degradation [84]. The Cambridge phantom scans a rotating bar, thus creating a virtual plane, to solve these problems [97]. The advantage of plane phantoms lies in the robustness of the feature extraction process which can, as a consequence, be reliably and fully automated. The pitfall of this method lies in the non-discriminative phantom geometry which can result in underdetermined systems if the acquired calibration samples do not cover all degrees of freedom. This can be avoided by strictly respecting the acquisition protocols presented in [97, 120].

Registration Phantoms: the last class of calibration methods relies on surface or intensity based registration techniques and therefore has the advantage of being independent of phantom geometry. The only requirement on phantom shape is that its US image is sufficiently discriminative with respect to rotations and translations, which is true for non-symmetric phantoms. The lack of precision of registration algorithms is the major drawback of this approach. To our knowledge, only one study examined registration-based 2-D probe calibration, registering US slices with an MRI image of the phantom [19]. A 3-D approach is discussed in the next paragraph.

1.1.2 3-D Probe Calibration

Until today, only few studies about calibration of 3-D probes were carried out. Poon and Rohling [96] compared 3-D calibration based on a IXI-fiducial wire phantom, a pointer tip phantom and a cube phantom. The IXI wire phantom and the cube phantom methods require only one volume acquisition for calibration. The presented feature detection is semi-automatic. The best results yielded the IXI phantom with a mean error in reproducibility of 1.5 mm, a RMS error of

the point accuracy measure of 2.15 mm and a RMS error of the reconstruction accuracy by distance measure of 1.52 mm. Bouchet et al [22] examined Z-fiducial phantom and achieved a RMS point accuracy error of 1.1mm. Two variants of a surface registration based 3-D calibration method were presented by Lange and Eulenstein in [72]. The first one registers 3-D US images of the phantom with a geometric model derived from its CT scan. The second variant registers a number of US images of the phantom acquired from different positions. In both cases, surfaces are extracted manually. The authors claim that the latter approach could be fully automated. The CT variant performed best and yielded a RMS error in reproducibility precision of 1.8 mm and a RMS error in point accuracy of 2.0 mm. The ultrasound speed distortion problem is not addressed.

In this study we propose a 3-D US calibration method based on a single plane membrane phantom. A fast, precise and accurate 3-D feature extraction algorithm relying on the 2-D Hough transform is presented. In contrast to existing 3-D US calibration systems, the feature extraction process is fully automated. In the result section, precision and accuracy assessments are carried out using a specially designed validation phantom.

1.2 Materials and Methods

1.2.1 Acquisition Hardware

The acquisition hardware consists of a GE Voluson 730 Pro 3-D US scanner and a NDI Polaris optical tracker with a 0.25 mm RMS error (as communicated by NDI). The tracking system operates with wireless (passive) infrared-reflecting rigid bodies equipped with flat markers. The ultrasound volumes are acquired with a 5 to 9 Mhz two dimensional curved array probe (see Fig. A.1a). The piezo array of the probe is mounted on a mechanical device which is capable of sweeping regularly around its rotation axis within a predefined angular range. During the continuous sweeping process the US hardware reconstructs 3-D volumes from the series of acquired 2-D slices. The 3-D acquisition time ranges from 1s to 4s, depending mainly on sweep angle and axial acquisition depth. Images are digitally transferred using a proprietary software from GE Medical Systems named 4D View. The US scanner also communicates the voxel size. The scan converter assumes the speed of sound (SoS) in tissue to be 1540 m/s.

1.2.2 The Membrane Phantom

The calibration phantom being dedicated to a clinical context, ergonomics considerations had an important impact on its design. We use a variation of the wall

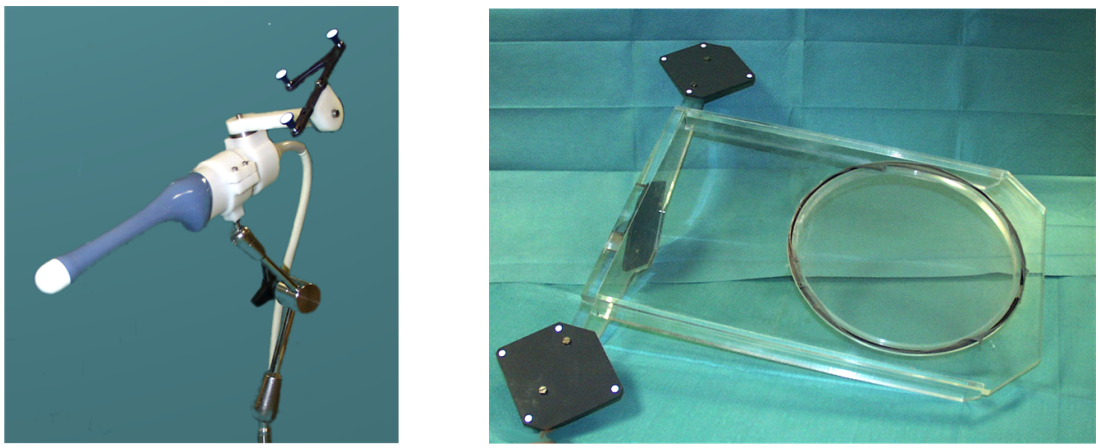

(a) (b)

Figure A.1: Calibration hardware. Figure (a) shows an endorectal US probe mounted on an articulated arm. Figure (b) shows the membrane phantom. Both the probe and the membrane are equipped with infra-red reflecting passive rigid bodies for tracking.

phantom presented in [97], which is based on imaging the bottom wall of a water tank. The geometric form of the wall, which is a line in 2-D and a plane in 3-D, can be very robustly extracted from the US data using statistical algorithms like the Hough transform. This makes it possible to fully automate the feature extraction process without significant loss in precision and accuracy. This represents a big advantage over semi-automatic point-detection based phantoms in terms of calibration speed and ease of use. To overcome the plane thickness and the reverberation problems observable in US images of rigid surfaces [97] a filigrane nylon mesh membrane, tightly spanned on a planar rigid support with a circular and about 20cm wide hole, is used as target (see Fig. A.1b). Reverberation is further reduced by inclining the membrane plane with respect to the water tank bottom by 45°. A tracking reference (rigid body) is mounted on the membrane frame for phantom localization. The phantom is filled with water and equipped with a thermometer to measure water temperature.

1.2.3 3-D Calibration Mathematics

As illustrated in Fig. A.2, four references are relevant for calibration: first of all, the membrane space **M** is defined as a reference in which the membrane lies in the origin and is parallel to e_x and e_y base vectors of **M**. In this space, every point with a zero z-ordinate is a membrane plane point. The phantom space **Ph** and the probe space **Pr** are defined by the rigid bodies that are attached on the phantom

and on the probe. Finally, the US volume space **U** corresponds to the voxel space of the 3-D images acquired by the ultrasound device. $\mathbf{T_{Ph2M}}$, $\mathbf{T_{Pr2Ph}}$ and $\mathbf{T_{U2Pr}}$ are homogenous 4x4 transformation matrices.

Suppose that we identified a point $\mathbf{p} = (x, y, z)$ in a US volume **U** as a point belonging to the membrane. With $\mathbf{s} = (s_x, s_y, s_z)$ denoting the voxel scale factors, it is verified that

$$\begin{pmatrix} m_1 \\ m_2 \\ 0 \\ 1 \end{pmatrix} = \mathbf{T_{Ph2M}} \cdot \mathbf{T_{Pr2Ph}} \cdot \mathbf{T_{U2Pr}} \cdot \begin{pmatrix} s_x x \\ s_y y \\ s_z z \\ 1 \end{pmatrix}. \tag{A.1}$$

where $\mathbf{T_{Ph2M}}$ is known from membrane pre-calibration (see Chap. 1.2.4) and $\mathbf{T_{Pr2Ph}}$ is given by the tracking system. Further, the scaling vector **s** is communicated by the US hardware. The remaining unknown element is the homogenous rigid transformation $\mathbf{T_{U2P}}$. For convenience we define the elements of $\mathbf{T_{Pr2Ph}} \cdot \mathbf{T_{Ph2M}}$ as a_{ij} and the elements of $\mathbf{T_{U2P}}$ as b_{ij} $(i, j \in 1..4)$. The zero component of Eqn. A.1 yields then

$$\begin{aligned} 0 \;=\; & a_{31}\,(s_x x b_{11} + s_y y b_{12} + s_z z b_{13} + b_{14}) + \\ & a_{32}\,(s_x x b_{21} + s_y y b_{22} + s_z z b_{23} + b_{24}) + \\ & a_{33}\,(s_x x b_{31} + s_y y b_{32} + s_z z b_{33} + b_{34}) + \\ & a_{34}\ . \end{aligned} \tag{A.2}$$

Using Euler angles and a three-dimensional vector we can represent $\mathbf{T_{U2P}}$ with six variables, which leaves us in total with 6 unknowns to solve for. A detected plane can be added to the equation system by adding at least three plane points (Using of course the $\mathbf{T_{Pr2Ph}}$ measured while acquiring the US volume in which the plane was detected).

1.2.4 Membrane Pre-calibration

To reduce the number of degrees of freedom of the calibration process, the membrane space **M** is determined using a pointer equipped with a tracking reference. A large number of surface points of the membrane-supporting structure is acquired in order to compute the plane equation using a least square approximation combined with a simple M-estimator to increase robustness. Since the phantom rigid body is permanently fixed on the phantom, pre-calibration has to be carried out only once.

1.2.5 Acquisition Protocol

The major drawback of single-plane phantoms resides in their barely discriminative geometry. A plane can be described with only three variables, from which

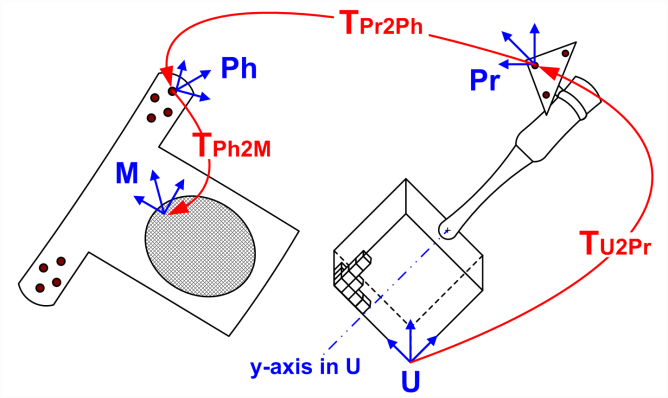

Figure A.2: Illustration of the transformations involved in the calibration process. Note that scaling is omitted from the scheme for simplification.

it follows that even with an optimal acquisition protocol, a minimum number of two acquisitions is necessary to cover all degrees of freedom. To obtain robust results the twelve-step acquisition protocol presented in [97] is used. The protocol improvement presented in [120] mainly addresses the z-axis imprecision problem inherent of most 2-D calibration systems. Since 3-D probes give as much information in z-direction as in x- or y- direction this modification yields no particular advantage in the 3-D domain but requires at least 18 steps. For that reason we stick to the original version.

Sweeping and volume reconstruction being a continuous process of 1 to 4 seconds, significant distortions can be observed in the US volume when the probe is moved rapidly. Also, no direct access to the digital data is available which prevented synchronization of probe position measurement with US image acquisition. Therefore an articulated arm for complete probe immobilization during acquisition is used, eliminating all motion-induced artifacts and time lags. Furthermore, immobilizing the probe makes it possible to perform high precision position measuring based on a large number of measures and outlier elimination.

1.2.6 Feature Extraction

The first step of the feature extraction process consists in correcting the distortion caused by the difference between US speed in water at room temperature and in human tissue at 37°. To determine US speed in water in function of temperature the polynomial formula established by Bilaniuk and Wong was used [16, 17]. A distortion geometry overview for all common probe types is given in [52]. The distortion geometry of a sectorial probe is given in Fig. A.3a. With v_W^t being the US velocity in water for a given temperature t and v_T being the velocity in tissue,

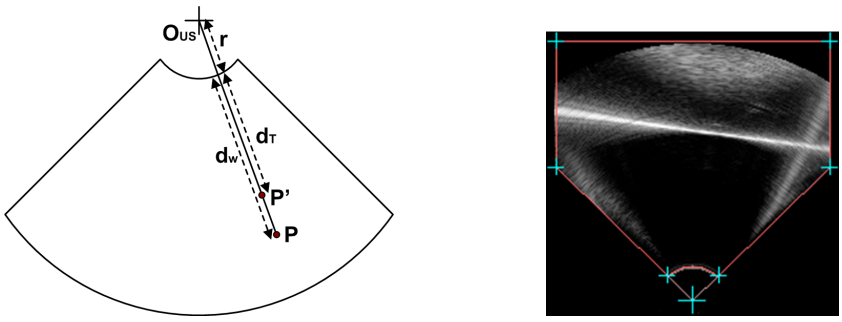

(a) (b)

Figure A.3: US speed correction. Fig. (a) illustrates the correction geometry of a sectorial probe. $\mathbf{O_{US}}$ is the probe origin, r the probe surface radius, d_W the point's $\mathbf{P}$ distorted distance from the probe surface in water, d_T the corrected distance and $\mathbf{P}'$ the corrected point. Fig. (b) shows the probe mask used to determine US origin and probe surface radius.

d_T is determined using the following formula:

$$d_T = \frac{v_T}{v_W} \cdot d_W. \tag{A.3}$$

Sectorial probe speed correction requires manual definition of the US origin and the scan head surface radius. A graphical user interface was developed for this purpose (see Fig. A.3b). Origin and surface radius have to be defined only once during calibration. Plane detection can be carried out using the 3-D Hough transformation, but it would take several minutes to compute the result. Fortunately it is possible to determine the plane with good precision by simply extracting its intersection with two arbitrary volume slices, using the 2-D Hough transform. To facilitate and to accelerate US speed correction the xy and zy planes passing through the scan head origin were used. The Hough transform implementation uses intensity accumulation and the following threshold s_H for an image $\mathbf{I}$:

$$s_H = \max\{i \in Hist(\mathbf{I})\} + (\max\{i \in \mathbf{I}\} - \min\{i \in \mathbf{I}\})/3\ . \tag{A.4}$$

The purpose of s_H is to ignore the low-intensity water background, which represents the largest part of the image.

1.2.7 Optimization

Optimization of Eqn. A.2 is carried out with the non-linear Levenberg-Marquardt implementation given in [98]. A random restart scheme within a range of reasonable initialization values robustifies this process.

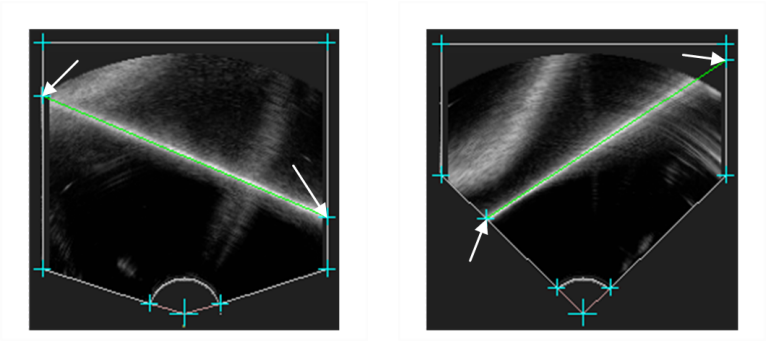

(a) (b)

Figure A.4: Screenshot of a successful automatic plane extraction. Note that lines are correctly detected in spite of a degraded membrane image caused by a steep scan angle. The arrows point at the line intersections with the mask.

1.2.8 Visual Back-tests

The plane coordinates as resulting from the optimization process are visualized as a line in the slices used for feature extraction, and distance plus rotational errors between the segmented and the calculated line are evaluated. This allows for manual replacement of evident outliers with new acquisitions and for recomputation of the calibration without requiring a complete restart.

1.3 Experimental Results

1.3.1 Test Configuration

Precision and accuracy assessments were carried out using a membrane plane precalibration with an RMS surface distance error of 0.43mm for the measured surface points. A total of ten calibrations were performed using the twelve-step protocol. The probe rigid body was not moved between calibrations. The probe was mounted on an articulated arm and immobilized during position and image acquisition. The water temperature was 23°. The acquired US volumes had a size of (199, 199, 199) isotropic voxels with 0.477mm side lengths.

1.3.2 Feature Extraction Quality

The Hough transform extracted lines correctly for 238 out of 240 acquired images. Line detection failed for images on which only a very small part of the membrane was visible. In these cases, lines had to be manually determined. Note that

	Distance Error [mm]	Distance Error [vox]	Angular Error [deg]
RMS Error	0.37	0.77	0.26
Max Error	1.30	2.73	1.09

Table A.1: Aggregated Feature Extraction Precision.

no manual outlier elimination was carried out. As detection failures were rare, calibration was in average carried out in about twenty minutes. To get a better idea about the quality of feature extraction and about the presence of distortions in the membrane images we measured the detection precision: using Eqn. A.2 we can calculate the distance error between a measured plane point and the computed plane as follows:

$$\begin{aligned} \epsilon(x,y,z) \;=\; & a_{31}\,(s_x x b_{11} + s_y y b_{12} + s_z z b_{13} + b_{14}) + \\ & a_{32}\,(s_x x b_{21} + s_y y b_{22} + s_z z b_{23} + b_{24}) + \\ & a_{33}\,(s_x x b_{31} + s_y y b_{32} + s_z z b_{33} + b_{34}) + \\ & a_{34}\ . \end{aligned} \tag{A.5}$$

For each calibration, the average and the root mean square (RMS) distance of a set of points to the pre-calibration plane was computed using Eqn. A.5. For each line we computed ten equidistant points between the extreme points on the line segment inside the US volume. The angular feature extraction error is defined as the angle between the computed plane normal and the cross product of the directional vectors of the two extracted lines. Based on this definition the maximum and the RMS angular errors were computed for each acquired volume of the calibration. The aggregated errors for all calibrations can be found in Tab. A.1.

1.3.3 Calibration Precision

The calibration precision measures the reproducibility of calibration results. Again, both the translational and the angular errors were assessed. The translational error is defined as the standard deviation of the volume center after scaling and right-hand multiplication to the different calibration transformations $\mathbf{T^i_{U2Pr}}$. The angular error is measured as the standard deviation of angular differences between the $(0,0,1)$ vector after scaling and right-hand multiplication to the different calibration transformations $\mathbf{T^i_{U2Pr}}$ (see Tab. A.2).

	Distance Error [mm]	Distance Error [vox]	Angular Error [deg]
Standard Deviation	1.15	2.41	0.61
Max Error	1.99	4.03	1.12

Table A.2: Calibration Precision.

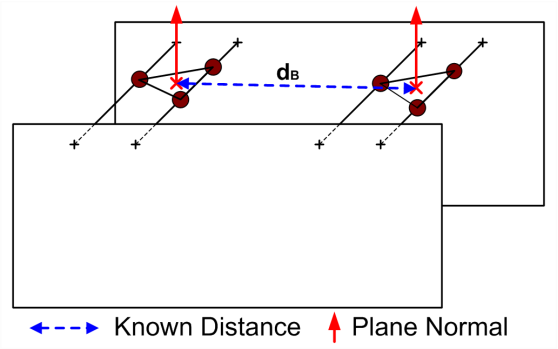

Figure A.5: Reconstruction Accuracy Measurement Phantom.

1.3.4 Reconstruction Accuracy

Reconstruction accuracy was assessed using the bead phantom illustrated in Fig. A.5. Note that the beads are co-planar within a precision of 0.25mm (RMS). The left-hand three beads form the left triangle while the right-hand beads form the right triangle. The distance d_B of the triangle barycenters was evaluated with an estimated accuracy of about 0.5mm.

Twenty US images of the phantom were acquired, ten imaging the left triangle and ten the right triangle. Note that images of one triangle did not intersect with the second one. The bead centers were manually extracted from the images. The reconstructed triangle barycenters and normals were then projected into probe space for each calibration, which yielded 100 point and vector pairs. The distance error for a point pair is defined as the difference between their Euclidean distance and d_B. The angular error for a vector pair is defined as the angle between both vectors. The results are given in Tab. A.3.

1.4 Discussion

Probe calibration for currently available swept-volume 3-D probes makes only sense for static applications. Real-time access to acquired 3D images is currently

	Distance Error [mm]	Distance Error [vox]	Angular Error [deg]
RMS Error	0.90	1.90	1.79
Max Error	2.44	5.11	3.03

Table A.3: Reconstruction Accuracy Results.

not provided by the Voluson hardware. Also, depending on scan parameters, the duration of the volume sweeping process ranges from 1s to 4s. Due to probe or tissue motion, the physical location of voxels can therefore be way off the position indicated by the calibration. The latter problem could be reduced by scan-heads equipped with high-frequency sweeping-devices, but it will disappear only with non-sweeping 2-D piezo array US probes. Until then, it is better to calibrate probes using an articulated arm for immobilization.

Passing from 2-D to 3-D calibration improves calibration results (for information on 2-D precision of the membrane phantom see [58]) because the z-axis-uncertainty inherent of 2-D calibration is eliminated: when giving a plane instead of a line as input to the optimizer the rotational degrees of freedom are significantly better covered. This allowed us to reduce the number of acquisitions for calibration while still achieving precise and accurate results.

Feature extraction from membrane phantom images showed both robust and precise results. On our set of test images the line blurring and intensity degradation effects occurring when scanning a wall phantom from an oblique angle were correctly handled by the Hough transform: lines were consistently placed in the center of the beam width. Note that the feature extraction precision RMS and maximum errors reported in Tab. A.1 are relatively small, which indicates that the physical plane location corresponds indeed with the beam width center line. Membrane reverberation was not observable and did therefore not disturb the detection process. Also, the membrane phantom was not exposed to line thickness problems. Due to these characteristics, feature extraction could be fully automated (up to the manual US origin and probe radius determination required for US speed correction).

The user-independency resulting from automated feature extraction is partially counter-balanced by the necessity to follow a protocol for data acquisition, which can re-introduce user bias. Nevertheless we believe that it is more convenient to follow a simple acquisition protocol instead of extracting features semi-automatically from US volumes. Further we preferred to correct for US speed errors instead of requiring 50° water or phantom fill materials that have the same US characteristics than human tissue, which makes it possible to use the phantom in a clinical

context.

The overall calibration time of about 20 minutes is mostly due to manipulation of the articulated arm, high precision probe position measurement and data transfer, which requires several manual interventions. Feature extraction and optimization is computed in several seconds. In cases where the feature extraction precision evaluation shows poor results, a visual verification and eventually a correction have to be carried out, which requires some additional minutes. Significant speed-up could be achieved by automating the communication between the US scanner and the calibration computer.

The presented calibration system assumes that the SoS in human tissue is uniform and that it corresponds to the SoS internally used by the US scanner, which is in general the mean SoS in tissue of 1540m/s. However, SoS varies with different types of tissue: The SoS in fat is approximately 1450m/s, in blood 1570m/s, in the brain 1541m/s and in water 1480m/s. As the US generally crosses tissue layers of different thickness and different types on its way through the body, and as the target tissue is often viewed from various positions, the in vivo accuracy of the calibration may show fluctuations of more than 5 per cents in extreme cases during an examination. Also, for some applications it would be appropriate to use a different mean SoS than the 1540m/s for calibration, but this is beyond the scope of this study.

Future work will address the twelve-step acquisition protocol which contains a lot of redundancy. Also, our system does currently not provide a foolproof indicator for missing coverage of degrees of freedom. We therefore started experiments with an Eigenvalue system similar to the one presented in [59].

1.5 Conclusion

A robust 3-D US probe calibration method designed for clinical usage was presented. Calibration can be carried out in about twenty minutes due to fully automatic 3-D Plane extraction based on robust and efficient 2-D line detection. The point reconstruction accuracy of our phantom can compete with previously presented 3-D phantoms: Lange and Eulenstein communicated RMS errors between 2.0mm and 2.2mm [72], Bouchet et al were confronted to 1.1mm RMS point accuracy [22] while Pohn and Rohling published errors between 1.52mm for their IXI-wire, 1.59mm for the cube and 1.85mm for their stylus approach. With 0.90mm RMS point accuracy (see Tab. A.3) we achieved slightly better results. Finally, the proposed method is temperature-independent and uses water as transmission matter which facilitates its usage.

Bibliography

[1] H. U. Ahmed, D. Pendse, R. Illing, C. Allen, J. H. Van der Meulen, and M. Emberton, "Will focal therapy become a standard of care for men with localized prostate cancer?" *Nature Clinical Practice*, vol. 4, no. 11, november 2007.

[2] R. Alterovitz, K. Goldberg, J. Pouliot, I.-C. J. Hsu, Y. Kim, S. M. Noworolski, and J. Kurhanewicz, "Registration of MR prostate images with biomechanical modeling and nonlinear parameter estimation," *Medical Physics*, vol. 33, no. 2, 2006.

[3] Y. Amit, U. Grenander, and M. Piccioni, "Matching 3D anatomical surfaces with non-rigid deformations using octree-splines," *Am. Statis. Ass.*, vol. 86, no. 414, pp. 376–387, 1991.

[4] N. Arad and D. Reisfeld, "Image warping using few anchor points and radial functions," *Comp. Graphics Forum*, vol. 14, no. 1, pp. 35–46, 1995.

[5] X. Artignan, M. Rastkhah, J. Balosso, P. Fourneret, O. Gilliot, and M. Bolla, "Quantification of prostate movements during radiotherapy treatment," *Cancer Radiotherapy*, vol. 10, no. 6–7, pp. 381–387, november 2006.

[6] B. B. Avants, M. Grossmann, and G. J. C., "Symmetric diffeomorphic image registration: Evaluating automated labeling of elderly and neurodegenerative cortex and frontal lobe," *In Proc. Biomed. Image Registration*, vol. 4057, pp. 50–57, 2006.

[7] R. Bajcsy and S. Kovacic, "Toward an individualized brain atlas elastic matching," Tech. Rep., Dept. of Computer and Information Science, Moore School, University of Philadelphia, Tech. Rep., 1986.

[8] J. M. Balter, M. H. Sandler, K. Klam, L. R. Bree, A. S. Lichter, and A. K. Tenhaken, "Measurement of prostate movement over the course of routine," *Int. J. Radiation Oncology Biol. Phys.*, vol. 31, no. 1, pp. 113–118, 1995.

[9] H. G. Barrow, J. M. Tenenbaum, R. C. Bolles, and H. C. Wolf, "Parametric correspondence and chamfer matching: two new techniques for image matching," *Proc. of the Fifth International Joint Conference on Artificial Intelligence*, pp. 659–663, 1977.

[10] M. Barva, J. Kybic, H. Liebgott, C. Cachard, and V. Hlaváč, "Comparison of methods for tool localization in biological tissue from 3D ultrasound data," *IEEE Ultrasonics Symposium*, pp. 1983–1986, 2006.

[11] C. J. Beard, P. Kijewski, M. Bussière, R. Gelman, D. Gladstone, K. Shaffer, M. Plunkett, P. Costello, and N. Coleman, "Analysis of prostate and seminal vesicle motion: Implications for treatment planning," *Int. J. Radiation Oncology Biol. Phys.*, vol. 34, no. 2, pp. 451–458, 1996.

[12] S. Benhimane and E. Malis, "Real-time image-based tracking of planes using second-order minimization techniques," in *Intelligent Robots and Systems 2004 (IROS 2004). IEEE/RSJ International Conference on*, vol. 1, 2004, pp. 943–948.

[13] P. J. Besl and N. D. McKay, "A method for registration of 3D shapes," *IEEE Trans. on Pattern Anal. and Mach. Intell.*, vol. 14, pp. 239–254, 1992.

[14] D. Beyersdorff, A. Winkel, B. Hamm, S. Lenk, S. A. Loening, and M. Taupitz, "MR imaging-guided prostate biopsy with a closed MR unit at 1.5 T: initial results," *Radiology*, vol. 234, no. 2, pp. 576–581, 2005.

[15] A. Bharatha, M. Hirose, N. Hata, S. K. Warfield, M. Ferrant, K. H. Zou, E. Suarez-Santana, J. Ruiz-Azola, A. D'Amico, C. M. C. Cormack, R. Kikinis, F. A. Jolesz, and C. M. C. Tempany, "Evaluation of three-dimensional finite element-based deformable registration of pre- and intraoperative prostate imaging," *Medical Physics*, vol. 28, no. 12, December 2001.

[16] N. Bilaniuk and G. S. K. Wong, "Speed of sound in pure water as a function of temperature," *J. Acoust. Soc. Am.*, vol. 93, no. 3, pp. 1609–1612, 1993.

[17] ——, "Erratum: Speed of sound in pure water as a function of temperature [j. acoust. soc. am. 93, 1609-1612 (1993)]," *J. Acoust. Soc. Am*, vol. 99, no. 5, p. 3257, 1996.

[18] J. M. Blackall, G. P. Penney, A. P. King, and D. J. Hawkes, "Alignment of sparse freehand 3-D ultrasound with preoperative images of the liver using models of respiratory motion and deformation," *IEEE Trans Med Imag*, vol. 24, no. 11, 2005.

[19] J. M. Blackall, D. Rueckert, C. R. Maurer, P. P. Graeme, D. L. G. Hill, and D. J. Hawkes, "An image registration approach to automated calibration for freehand 3-D ultrasound," *MICCAI Proc*, vol. 1935, pp. 462–471, 2000.

[20] F. L. Bookstein, "Principal warps: thin-plate splines and the decomposition of deformations," *IEEE Trans. on Pattern Anal. and Machine Intell.*, vol. 11, no. 6, pp. 567–585, 1989.

[21] G. Borgefors, "Hierarchical chamfer matching: a parametric edge matching algorithm," *IEEE Trans. on Pattern Anal. and Mach. Intell.*, vol. 10, pp. 849–865, 1998.

[22] L. G. Bouchet, S. L. Meeks, G. Goodchield, F. J. Bova, J. M. Buatti, and W. A. Friedmann, "Calibration of three-dimensional ultrasound images for image-guided radiation therapy," *Physics Med Biol*, vol. 46, pp. 559–577, 2001.

[23] D. Boukerroui, J. A. Noble, and M. Brady, "Velocity estimation in ultrasound images: a block matching approach," *Information Processing in Medical Imaging*, vol. 2732, 2003.

[24] R. P. Brent, *Algorithms for Minimization without derivatives.* Englewood Cliffs, N. J., 1973.

[25] W. Briggs, v. E. Henson, and S. F. McCormick, *A multigrid tutorial*, 2nd ed. Soc for Industrial & Applied Math, 2008.

[26] M. Bro-Nielsen, "Medical image registration and surgery simulation," Ph.D. thesis, Institut for Matematisk Modellering, Danmarks Tekniske Universitet, Lyngby, Denmark, Tech. Rep., 1996.

[27] C. Broit, "Optimal registration of deformed images," Ph.D. thesis, Computer and Information Science, University of Pennsylviania, Tech. Rep., 1981.

[28] M. Bucki, B. Dauliac, V. Daanen, R. Moalic, J. L. Descotes, and J. Troccaz, "ProNav: a navigation software for prostate biopsies," in *Surgetica'05*, 2005.

[29] P. Cachier, E. Bardinet, D. Dormont, X. Pennec, and N. Ayache, "Iconic feature based nonrigid registration: the PASHA algorithm," *Computer Vision and Image Understanding*, vol. 89, pp. 272–298, 2003.

[30] P. Cachier and D. Rey, "Symmetrization of the non-rigid registration problem using inversion-invariant energies: Application to multiple sclerosis." *MICCAI'00, LNCS*, vol. 1935, pp. 472–481, 2000.

[31] H. B. Carter, "Assessing risk: Does this patient have prostate cancer?" *Journal of the National Cancer Institute*, vol. 98, no. 8, pp. 506–507, 2006.

[32] G. E. Christensen, "Consistent linear-elastic transformations for image matching," *Proc. Inf. Process. Med. Imag.*, pp. 224–237, 1999.

[33] G. E. Christensen and H. J. Johnson, "Consistent image registration," *IEEE Trans. on Med. Imag.*, vol. 20, no. 7, pp. 568–582, 2001.

[34] G. E. Christensen, S. C. Joshi, and M. I. Miller, "Volumetric transformation of brain anatomy," *IEEE Trans. on Med. Imag.*, vol. 06, no. 6, pp. 864–877, 1997.

[35] G. E. Christensen, M. I. Miller, and M. Vannier, "A 3D deformable magnetic resonance textbook based on elasticity," *AAAI Spring Symposium Series: Applications of Computer Vision in Medical Image Processing*, pp. 153–156, 1994.

[36] G. E. Christensen, R. D. Rabbitt, and M. I. Miller, "Deformable templates using large deformation kinematics," *IEEE Trans. on Image Processing*, vol. 5, no. 10, pp. 1435–1447, 1996.

[37] B. Cohen and I. Dinstein, "New maximum likelihood motion estimation schemes for noisy ultrasound images," *Pattern Recognition*, no. 35, pp. 455–463, 2002.

[38] A. Collignon, F. Maes, D. Delaere, D. Vandermeulen, P. Suetens, and G. Marchal, "Automated multimodality image registration using information theory," *Proc. Information Processing in Medical Imaging (Brest)*, pp. 263–274, 1995.

[39] R. M. Comeau, A. Fenster, and T. M. Peters, "Integrated MR and ultrasound imaging for improved image guidance in neurosurgery," *SPIE Proc*, vol. 3338, pp. 747–754, 1998.

[40] M. Ding and A. Fenster, "A real-time biopsy needle segmentation technique using the hough transform," *Medical Physics*, vol. 30, no. 8, pp. 2222–2233, 2003.

[41] T. M. Downs, G. D. Grossfeld, K. Shinohara, and P. R. Carroll, "Transrectal ultrasound-guided prostate biopsy," *Image-Guided Diagnosis and Treatment of Cancer*, pp. 3–34, 2003.

[42] J. Duchon, "Interpolation des fonctions de deux variables suivant le principle de la flexion des plaques minces," *RAIRO Analyse Numérique*, pp. 5–12, 1976.

[43] L. H. Eadie, D. d. Cunha, R. B. Davidson, and A. M. Seifalian, "A real time pointer to a pre-operative surgical planning index block of ultrasound images for image guided surgery," in *Visualization and Data Analysis 2004, SPIE Vol. 5295*, San Jose, California, USA, Jan 2004, pp. 14–23.

[44] J. I. Epstein, P. C. Walsh, M. Carmichael, and C. B. Brendler, "Pathologic and clinical findings to predict tumor extent of nonpalpable (stage T1c) prostate cancer." *JAMA: the journal of the American Medical Association*, vol. 271, no. 5, pp. 368–374, 1994.

[45] A. C. Evans, W. Dai, L. Collins, P. Neelin, and S. Marrett, "Warping of a computerized 3-D atlas to match brain image volumes for quantitative neuroanatomical and functional analysis," *SPIE 1445, Medical Imaging V*, pp. 236–246, 1991.

[46] J. Ferlay, P. Autier, M. Boniol, M. Heanue, M. Colombet, and P. Boyle, "Estimates of the cancer incidence and mortality in europe in 2006," *Annals of Oncology*, February 2007.

[47] B. Fischer and J. Modersitzki, "Fast inversion of matrices arising in image processing." *Numerical Algorithms*, vol. 22, pp. 1–11, 1999.

[48] M. Fornefett, K. Rohr, and H. Stiehl, "Radial basis functions with compact support for elastic registration of medical images," in *Image and Vision Computing*, vol. 19, 2001, pp. 97–96.

[49] P. Foroughi and P. Abolmaesumi, "A modified HAMMER algorithm for deformable registration of ultrasound images," *In Proc. CARS'05*, pp. 236–241, June 2005.

[50] P. Foroughi, P. Abolmaesumi, and K. Hashtrudi-Zaad, "Intra-subject elastic registration of 3D ultrasound images," *MedIA*, vol. 10, pp. 713–725, 2006.

[51] A. Gillers, D. Mouton, X. Rébillard, D. Chautard, A. Ruffion, F. Staerman, and F. Cornud, "Conditions de réalisation et schéma de ponctions lors d'une première série de biopsies prostatiques," in *Progrès en Urologie*, vol. 14, 2004.

[52] A. Goldstein, "The effect of acoustic velocity on phantom measurements," *Ultrasound Med Biol*, vol. 26, no. 7, pp. 1133–1143, 2000.

[53] C. Gramkow, "On averaging rotations," *Journal of Mathematical Imaging and Vision*, vol. 15, pp. 7–16, 2001.

[54] A. P. Gu'eziec, X. Pennec, and N. Ayache, "Medical image registration using geometric hashing," *IEEE Comput. Sci. Eng.*, vol. 4, no. 4, pp. 29–41, 1997.

[55] E. Haber and J. Modersitzki, "Intensity gradient based registration and fusion of multi-modal images," *MICCAI'06*, vol. 3216, pp. 591–598, 2006.

[56] N. Hata, M. Jinzaki, D. Kacher, R. Cormak, D. Gering, A. Nabavi, S. G. Silverman, A. V. D'Amico, R. Kikinis, F. A. Jolesz, and C. M. Tempany, "MR imaging-guided prostate biopsy with surgical navigation software: device validation and feasibility," *Radiology*, vol. 220, no. 1, pp. 263–268, 2001.

[57] M. Hirose, A. Bharatha, N. Hata, K. H. Zou, S. K. Warfield, R. A. Cormack, A. D'Amico, R. Kikinis, F. A. Jolesz, and C. M. C. Tempany, "Quantitative MR imaging assessment of prostate gland deformation before and during MR imaging-guided brachytherapy," *Acad Radiol*, vol. 9, pp. 906–912, 2002.

[58] I. Hook, *Probe calibration for 3-D ultrasound image localization. Internship report.* TIMC laboratory, GMCAO, 2003.

[59] P. W. Hsu, R. W. Prager, A. H. Gee, and G. M. Treece, "Rapid, easy and reliable calibration for freehand 3-D ultrasound," *MIUA Proc*, pp. 91–94, July 2005.

[60] X. Huang, N. A. Hill, J. Ren, G. Guiraudon, D. R. Boughner, and T. M. Peters, "Dynamic 3D ultrasound and MR image registration of the beating heart," *MICCAI 2005*, vol. 2, pp. 171–178, 2005.

[61] D. I. Hughes and F. A. Duck, "Automatic attenuation compensation for ultrasonic imaging," *Ultrasound in Medicine & Biology*, vol. 23, pp. 651–664, 1997.

[62] A. Jemal, R. Siegel, E. Ward, T. Murray, J. Xu, and M. J. Thu, "Cancer statistics," *CA Cancer J Clin*, vol. 57, pp. 43–66, 2007.

[63] S. C. Joshi, "Large deformation diffeomorphisms and gaussian random fields for statistical characterization of brain sub-manifolds," Ph.D. thesis, Dep. of Electrical Engineering, Sever Institute of Technology, Washington University, St. Louis, MO, Tech. Rep., 1997.

[64] S. C. Joshi and M. I. Miller, "Landmark matching via large deformation diffeomorphisms," *IEEE Trans. on Medical Imaging*, vol. 9, no. 8, 2000.

[65] S. C. Joshi, M. I. Miller, G. E. Christensen, A. Banerjee, T. A. Coogan, and U. Grenander, "Hierarchical brain mapping via a generalized dirichlet solution for mapping brain manifolds," *SPIE'95*, vol. 2573, pp. 278–289, 1995.

[66] S. Kaneko, I. Murase, and S. Igarashi, "Robust image registration by increment sign correlation," *Pattern Recognition*, vol. 35, no. 10, pp. 2223–2234(12), 2003.

[67] D. W. Keetch, J. M. McMurtry, D. S. Smith, G. L. Andriole, and W. J. Catalona, "Prostate specific antigen density versus prostate specific antigen slope as predictors of prostate cancer in men with initially negative prostatic biopsies," *J Urol*, vol. 156, pp. 428–431, 1996.

[68] J. F. Krücker, G. L. LeCarpentier, J. B. Fowlkes, and P. L. Carson, "Rapid elastic image registration for 3-D ultrasound," *IEEE Trans. on Med. Imag.*, vol. 21, no. 11, pp. 1384–1394, 2002.

[69] A. Krieger, C. Csoma, I. I. Iordachita, P. Guion, A. K. Singh, G. Fichtinger, and L. L. Whitcomb, "Design and preliminary accuracy studies of an MRI-guided transrectal prostate intervention system," *MICCAI'07*, vol. 4792, pp. 59–67, 2007.

[70] A. Krieger, R. C. Susil, C. Menard, J. A. Coleman, G. Fichtinger, E. Atalar, and L. L. Whitcomb, "Design of a novel MRI compatible manipulator for image guided prostate interventions," *IEEE Trans. on Biomed. Eng.*, vol. 52, no. 2, pp. 306–313, 2005.

[71] V. Lagerburg, M. A. Moerland, J. J. Lagendijk, and J. J. Battermann, "Measurement of prostate rotation during insertion of needles for brachytherapy," *Radiother Oncol*, vol. 77, no. 3, december 2005.

[72] T. Lange and S. Eulenstein, "Calibration of swept-volume 3-D ultrasound," *MIUA Proc*, vol. 99, no. 3, pp. 29–32, july 2002.

[73] D. LaRose, "Iterative X-ray/CT registration using accelerated volume rendering," Ph.D. Thesis, Carnegie Mellon University, Tech. Rep., 2001.

[74] D. F. Leotta, "An efficient calibration method for freehand 3-D ultrasound imaging systems." *Ultrasound Med Biol*, vol. 30, pp. 999–1008, 2004.

[75] A. Leow, S. C. Huang, A. Geng, J. Becker, S. Davis, A. Toga, and P. Thompson, "Inverse consistent mapping in 3-D deformable image registration: its construction and statistical properties." *Proc. Inf. Process. Med. Imag.*, vol. 3565, pp. 493–503, 2005.

[76] J. Lian, L. Xing, and S. Hunjan, "Mapping of the prostate in endorectal coil-based MRI/MRSI and CT: a deformable registration and validation study," *Medical Physics*, vol. 31, no. 11, 2004.

[77] F. Linseth, G. A. Tangen, and T. Langø, "Probe calibration for freehand 3-D ultrasound," *Ultrasound Med Biol*, vol. 29, pp. 49–69, 2003.

[78] J. A. Little, D. L. G. Hill, and D. J. Hawkes, "Deformation incorporating rigid structures," in *Computer Vision and Image Understanding*, vol. 66, 1997, pp. 223–232.

[79] F. Maes, D. Vandermeulen, and P. Suetens, "Comparative evaluation of multiresolution optimization strategies for multimodality image registration by maximization of mutual information," *Medical Image Analysis*, vol. 3, no. 4, 1999.

[80] S. Malone, J. M. Crook, W. S. Kendal, and J. Szanto, "Respiratory-induced prostate motion: Quantification and characterization," *Int. J. Radiat. Oncol. Biol. Phys.*, vol. 48, pp. 105–109, 2000.

[81] J. E. McNeal, "Regional morphology and pathology of the prostate," *Am J Clin Pathol*, vol. 49, no. 3, pp. 347–57, march 1968.

[82] ——, "Normal and pathologic anatomy of the prostate," *Urology*, vol. 17, pp. 11–16, 1981.

[83] ——, "The zonal anatomy of the prostate," *Prostate*, vol. 2, no. 1, pp. 25–49, 1981.

[84] L. Mercier, T. Langø, F. Lindseth, and D. L. Collins, "A review of calibration techniques for freehand 3-D ultrasound systems," *Ultrasound Med Biol*, vol. 31, no. 2, pp. 143–165, 2004.

[85] J. Modersitzki and S. Wirtz, "Combining homogenization and registration," *Biomedical Image Registration: Third International Workshop, WBIR*, pp. 257–263, 2006.

[86] J. Modersitzki, *Numerical Methods for Image Registration*, C. Golub, C. Schwab, and E. Süli, Eds. Oxford University Press, 2004.

[87] A. A. Morsy and O. T. Von Ramm, "3D ultrasound tissue motion tracking using correlation search," *Ultrason. Imaging*, vol. 20, no. 3, pp. 151–159, 1998.

[88] D. M. Muratore and R. J. Galloway, "Beam calibration without a phantom for creating a 3-D freehand ultrasound system," *Ultrasound Med Biol*, vol. 27, no. 11, pp. 1557–1566, November 2001.

[89] P. Novotny, J. Cannon, and R. Howe, "Tool localization in 3D ultrasound images," *MICCAI'03*, vol. 2879, pp. 969–977, 2003.

[90] P. M. Novotny, J. A. Stoll, N. V. Vasilyev, P. J. del Nido, P. E. Dupont, T. E. Zickler, and R. D. Howe, "GPU based real-time instrument tracking with three-dimensional ultrasound," *Medical Image Analysis*, vol. 11, no. 5, pp. 458–464, 2007.

[91] M. Ohori, J. A. Eastham, H. Koh, K. Kuroiwa, K. M. Slawin, and T. M. Wheeler, "Is focal therapy reasonable in patients with early stage prostate cancer (CAP) - an analysis of radical prostatectomy specimens," *presented at the AUA annual meeting, Atlanta, GA, USA*, May 2006.

[92] T. E. Oliphant, A. Manduca, R. L. Ehman, and J. F. Greenleaf, "Complex-valued stiffness reconstruction for magnetic resonance elastography by algebraic inversion of the differential equation." *Magn. Reson. Med.*, vol. 45, pp. 299–310, 2001.

[93] R. A. Padhani, V. S. Khoo, J. Suckling, J. E. Husband, M. O. Leach, and D. P. Dearnaley, "Evaluating the effect of rectal distension and rectal movement on prostate gland position using cine MRI," *Int J. Radiation Oncology Biol. Phys.*, vol. 44, no. 3, pp. 525–533, 1999.

[94] N. Pagoulatos, D. N. Haynor, and Y. Kim, "A fast calibration method for 3-D tracking of ultrasound images using a spatial localizer," *Ultrasound Med Biol*, vol. 27, pp. 1219–1229, 2001.

[95] G. P. Penney, J. M. Blackall, M. S. Hamady, T. Sabharawal, A. Adam, and D. J. Hawkes, "Registration of freehand 3D ultrasound and magnetic resonance liver images," *Medical Image Analysis*, vol. 8, pp. 81–91, 2004.

[96] T. C. Poon and R. N. Rohling, "Comparison of calibration methods for spatial tracking of a 3-D ultrasound probe," *Ultrasound Med Biol*, vol. 31, no. 8, pp. 1095–1108, August 2005.

[97] R. W. Prager, R. N. Rohling, A. H. Gee, and L. Berman, "Rapid calibration for 3-D freehand ultrasound," *Ultrasound Med Biol*, vol. 24, no. 6, pp. 855–869, 1998.

[98] W. H. Press, B. P. Flannery, S. A. Teukolsky, and V. W. T., *Numerical Recipies in C, 2nd edition.* Cambridge University Press, 1992.

[99] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical recipes in C++*, 2nd ed. Oxford Press, 2002.

[100] A. Roche, G. Malandain, and N. Ayache, "Unifying maximum likelihood approaches in medical image registration," *Int. J. of Imaging Syst. and Techn.*, vol. 11, no. 1, pp. 71–80, 2000.

[101] A. Roche, G. Malandain, X. Pennec, and N. Ayache, "The correlation ratio as a new similarity measure for multimodal image registration," *MICCAI'98*, pp. 1115–1124, 1998.

[102] P. Rogelj and S. Kovacic, "Symmetric image registration," *Medical Image Analysis*, vol. 10, no. 3, pp. 484–493, 2006.

[103] R. Rohling, A. Gee, and L. Berman, "A comparison of freehand three-dimensional ultrasound reconstruction techniques," *Medical Image Analysis*, vol. 3, no. 4, pp. 339–359, december 1999.

[104] D. Rueckert, L. Sonoda, C. Hayes, D. L. G. Hill, M. O. Leach, and D. Hawkes, "Non-rigid registration using free-form deformations: Application to breast MR images." *IEEE Trans. on Med. Imag.*, vol. 18, no. 8, pp. 712–721, 1999.

[105] A. Sawada, K. Yoda, M. Kokubo, T. Kunieda, Y. Nagata, and M. Hiraoka, "A technique for noninvasive respiratory gated radiation treatment based on a real time 3d ultrasound image correlation: A phantom study," *Medical Physics*, vol. 31, no. 2, pp. 245–250, 2004.

[106] W. Shao, R. Wu, C. H. Thng, K. V. Ling, H. S. S. Ho, C. W. S. Cheng, and W. S. Ng, "Deformable registration for integration of MRI/MRSI information in TRUS-guided prostate biopsy," *SPIE, Medical Imaging 2005: Image Processing*, vol. 5747, pp. 1263–1273, 2005.

[107] D. Shen and C. Davatzikos, "HAMMER: hierarchical attribute matching mechanism for elastic registration," *IEEE Trans. on Med. Imag.*, vol. 21, no. 11, 2002.

[108] D. Shen, Z. Lao, J. Zeng, W. Zhang, I. A. Sesterhenn, L. Sun, J. W. Moul, E. H. Herskovits, G. Fichtinger, and C. Davatzikos, "Optimization of biopsy strategy by a statistical atlas of prostate cancer distribution," *Medical Image Analysis*, vol. 8, no. 2, pp. 139–150, 2004.

[109] R. Souchon, "Prostate cancer detection and HIFU therapy monitoring using elastography," PhD thesis. Institut National des sciences appliquées de Lyon, Lyon., Tech. Rep., 2004.

[110] T. A. Stamey, F. S. Freiha, J. E. McNeal, E. A. Redwine, A. S. Whittemore, and H. P. Schmid, "Localized prostate cancer. relationship of tumor volume to clinical significance for treatment of prostate cancer," *Cancer*, vol. 71, no. 3, pp. 933–938, 1993.

[111] A. State, D. T. Chen, C. Tektor, A. Brandt, H. Chen, R. Ohbuchi, M. Bajura, and H. Fuchs, "Case study: Observing a volume rendered fetus within a pregnant patient," *IEEE Visualization Proc*, pp. 364–368, 1994.

[112] R. Stefanescu, "Grid powered nonlinear image registration with locally adaptive regularization," Ph.D. Thesis, Université de Nice Sophia-Antipolis, Tech. Rep., 2005.

[113] D. Stoianovici, D. Song, D. Petrisor, D. Ursu, D. Mazilu, M. Mutener, M. Schar, and A. Patriciu, ""MRI stealth" robot for prostate interventions," *Minimally Invasive Therapy*, vol. 16, no. 4, pp. 241–248, 2007.

[114] M. G. Strintzis and I. Kokkinidis, "Maximum likelihood motion estimation in ultrasound image sequences," *IEEE Signal Proc. Letters*, vol. 4, no. 6, 1997.

[115] C. Studholme, D. J. Hawkes, and D. L. G. Hill, "An overlap-invariant entropy measure of 3D medical image alignment," *Pattern Recognition*, vol. 32, pp. 71–86, 1999.

[116] R. C. Susil, K. Camphausen, P. Choyke, E. R. McVeigh, G. S. Gustafson, H. Ning, R. W. Miller, E. Atalar, C. N. Coleman, and C. Ménard, "System for prostate brachytherapy and biopsy in a standard 1.5 T MRI scanner," *Magn. Reson. Med.*, vol. 52, no. 3, pp. 683–687, 2004.

[117] L. T., *Ultrasound guided surgery: Image processing and navigation. Ph.D. thesis.* Trondheim: Norvegian University of Science and Technology, 2000.

[118] D. Terzopoulos, "Regularization of inverse visual problems involving discontinuities," *IEEE Trans. on Pattern Anal. Machine Intell.*, vol. 8, no. 4, pp. 413–424, 1986.

[119] J. P. Thirion, "Image matching as a diffusion process: Analogy with maxwell's demons," *Medical Image Analysis*, vol. 2, no. 3, pp. 243–260, 1995.

[120] G. M. Treece, A. H. Gee, R. W. Prager, C. C. J., and B. L. H., "High-definition freehand 3-D ultrasound," *Ultrasound Med Biol*, vol. 29, no. 4, pp. 529–46, April 2003.

[121] J. Tsao, "Interpolation artifacts in multimodality image registration based on maximization of mutual information," *IEEE Trans. on Med. Imaging*, vol. 22, no. 7, pp. 854–864, 2003.

[122] G. H. Valadez, "Variational methods for multimodal image matching," PhD thesis. INRIA Nice - Sophia Antipolis, Tech. Rep., mai 2002.

[123] M. van Herk, A. Bruce, G. Kroes, T. Shouman, A. Touw, and J. V. Lebesque, "Quantification of organ motion during conformal radiotherapy of the prostate by three dimensional image registration," *Int. J. Radiation Oncology Biol. Phys.*, vol. 33, no. 5, pp. 1311–1320, 1995.

[124] T. Vercauteren, X. Pennec, A. Perchant, and N. Ayache, "Non-parametric diffeomorphic image registration with the demons algorithm," *Proc. of Information Processing in Medical Imaging (IPMI'07)*, vol. 4792, July 2007.

[125] A. Villers, D. Mouton, X. Rébillard, D. Chautard, A. Ruffion, F. Staerman, and F. Cornud, "Conditions de réalisation et schéma de ponctions lors d'une première série de biopsies prostatiques," *Progrès en Urologie*, vol. 14, pp. 1–11, 2004.

[126] A. Villers, X. Rébillard, M. Soulié, J. L. Davin, P. Coloby, J. L. Moreau, A. Mejean, J. Irani, C. Coulange, and P. Mangin, "Dépistage du cancer de la prostate," *Progrès en Urologie*, vol. 13, pp. 209–214, 2003.

[127] P. A. Viola and W. M. Wells, "Alignment by maximization of mutual information," *ICCV'95*, pp. 16–23, 1995.

[128] C. Wachinger, W. Wein, and N. Navab, "Three-dimensional ultrasound mosaicing," *MICCAI'07*, vol. 4792, pp. 327–335, 2007.

[129] S. Wirtz, N. Papenberg, B. Fischer, and O. Schmitt, "Robust and staining-invariant elastic registration of a series of images from histologic slices," *Proc. of the SPIE, Medical Imaging*, 2005.

[130] S. Xu, J. Krucker, P. Guion, N. Glossop, Z. Neeman, P. Choyke, A. K. Singh, and B. J. Wood, "Closed-loop control in fused MR-TRUS image-guided prostate biopsy," *MICCAI'07*, vol. 4791, pp. 128–135, 2007.

[131] Y. Zhan, M. Feldman, J. Tomaszeweski, C. Davatzikos, and D. Shen, "Registering histological and MR images of the prostate for image-based cancer detection," *MICCAI'06*, vol. 4191, pp. 620–628, 2006.

[132] Z. Zhang, Y. Jiang, and H. Tsui, "Consistent multi-modal non-rigid registration based on a variational approach," *Pattern Recognition Letters*, pp. 715–725, 2006.

[133] D. Zikic, W. Wein, A. Khamene, D. A. Clevert, and N. Navab, "Fast deformable registration of 3D-ultrasound data using a variational approach," *MICCAI*, vol. 9, no. 1, pp. 915–923, January 2006.

Index