

Analysis of Ultrasound and Optical Coherence Tomography for Markerless Volumetric Image Guidance in Robotic Radiosurgery

Vom Promotionsausschuss der
Technischen Universität Hamburg
zur Erlangung des akademischen Grades
Doktor-Ingenieur (Dr.-Ing.)
genehmigte Dissertation

von
Matthias Schlüter

aus
Kiel

2021

1. Gutachter: Prof. Dr.-Ing. Alexander Schlaefer
 2. Gutachter: Prof. Dr.-Ing. Rolf-Rainer Grigat
- Tag der mündlichen Prüfung: 29.06.2021

Summary

This thesis analyzes different aspects of image guidance in robotic radiosurgery considering the modalities ultrasound imaging and optical coherence tomography (OCT). Robotic radiosurgery is a specific type of radiation therapy. It employs an X-ray source, which is mounted to a robot, to generate high-energy beams for cancer treatment. Due to the radiation's harmful effect on healthy tissue, a highly accurate delivery of the dose is necessary to fully cover the target while sparing healthy structures. For this purpose, beams are delivered from many different directions such that the dose accumulates only within the target area and falls off quickly outside of the target. The robot-mounted source allows to actively compensate for patient motion occurring during an irradiation session by adjusting the beam directions in real time. However, this requires imaging throughout a treatment to reliably detect any motion of the patient or any internal organ motion.

Common imaging modalities for guiding radiation therapy are X-ray-based techniques, optical tracking systems, electromagnetic tracking, and magnetic resonance imaging. However, each of these methods comes with some drawbacks. For example, some of them are harmful and do not allow for continuous imaging, require indirect marker tracking, only provide projective imaging, or their integration in existing treatment systems is challenging. For this reason, this thesis considers harmless modalities providing direct volumetric tracking of targets without requiring markers. On the one hand, we study aspects of ultrasound-based image guidance for abdominal treatment sites. This modality allows for markerless volumetric tracking of internal soft-tissue targets. However, manual operation of an ultrasound transducer is not feasible in a treatment room, requiring some kind of robotic system to hold the transducer steadily onto the patient's abdominal wall. Previous studies have shown that such approaches and systems are feasible. However, there are remaining issues in practical application, because a robot is typically radiopaque and thereby interferes with the beam delivery by blocking directions for treatment beams. Therefore, we propose and evaluate methods to optimize and analyze such a robot setup in order to mitigate deterioration of treatment quality. We demonstrate the effects of our methods on cases of prostate cancer. On the other hand, this thesis considers OCT for markerless volumetric tracking of, for example, a patient's head during cranial radiosurgery. The sub-surface information accessible for OCT might allow for direct tracking of poorly structured surfaces for which conventional superficially scanning optical systems fail. We develop and experimentally characterize a prototypical OCT-based tracking system.

Abstract

Für eine präzise Strahlenchirurgie müssen Bewegungen während der Behandlung zuverlässig detektiert und kompensiert werden. Diese Arbeit untersucht Ansätze zur markerlosen volumetrischen Bildführung. Der Einfluss robotischer transabdominaler Ultraschallbildgebung wird analysiert und optimiert. Für kraniale Ziele wird ein neuartiger Ansatz mittels optischer Kohärenztomographie beschrieben.

An accurate dose delivery in radiosurgery requires to reliably detect and compensate any motion of the target during the treatment. In this thesis, we study approaches for markerless volumetric image guidance. For abdominal targets, we analyze and optimize the impact of robotic transabdominal ultrasound imaging. For cranial targets, we describe a novel setup using optical coherence tomography.

Acknowledgements

Diese Arbeit entstand während meiner Zeit als wissenschaftlicher Mitarbeiter am Institut für Medizintechnische und Intelligente Systeme der TUHH. Für diese Möglichkeit und die Betreuung möchte ich mich herzlich bei Herrn Prof. Schlaefer bedanken. Außerdem danke ich Herrn Prof. Grigat für die Begutachtung dieser Arbeit und sein ausführliches Feedback.

Ich möchte mich bei allen ehemaligen Kollegen für die gemeinsame Zeit bedanken, auch über die Arbeit hinaus. Insbesondere danke ich Christoph und Thore, von denen ich sehr viel lernen konnte. Michael, Marco und Katrin danke ich für ihre organisatorische und administrative Unterstützung. Außerdem Omer, Sven, Sarah, Nils, Johanna, Martin, Max und Stefan für die gemeinsamen Projekte sowie Lukas für seine Unterstützung als HiWi. Darüber hinaus möchte ich mich für die Zusammenarbeit bei den Kollegen des Instituts für Biomedizinische Bildgebung am UKE und der Max Planck Gruppe für Dynamik in Atomarer Auflösung bedanken. Ganz besonders danke ich Herrn Christoph Fürweger und seinen Kollegen am Cyberknife Zentrum in München.

Für das Gegenlesen dieser Arbeit möchte ich Lukas, Johanna und Sven danken. Außerdem danke ich meiner Familie, Freunden sowie ganz besonders Johanna für die ständige Unterstützung. Meinen neuen Kollegen bei Ibeo möchte ich für die sehr angenehme Arbeitsatmosphäre danken, die auch zur Fertigstellung dieser Arbeit beitrug.

Contents

Summary	iii
Abstract	v
Acknowledgements	vii
List of Acronyms	xiii
1 Introduction	1
1.1 Motivation	1
1.2 Research Questions	4
1.3 Organization	4
2 Radiation Therapy and Robotic Radiosurgery	7
2.1 Medical and Physical Background	7
2.1.1 Types of Radiation Therapy	7
2.1.2 Radiation Therapy with Photon Beams	10
2.2 The CyberKnife System	12
2.3 Inverse Planning	14
2.3.1 Data for Treatment Planning	14
2.3.2 Clinical Goals	17
2.3.3 Dose Model	18
2.3.4 Sampling of Cylindrical Beams	20
2.3.5 Linear-Programming Approach	21
3 Image Guidance in Radiation Therapy	25
3.1 Immobilization of the Patient	26
3.2 X-ray Imaging	27
3.2.1 Basic Principles	27
3.2.2 Stereoscopic X-ray Imaging	28
3.2.3 Cone Beam CT	29
3.2.4 Portal Imaging	30
3.3 Optical Tracking	30
3.3.1 Basic Concepts	31
3.3.2 Surface Imaging for Patient Setup	31
3.3.3 Surrogate Tracking for Internal Respiratory Motion	32
3.3.4 Head Tracking	33
3.4 Electromagnetic Tracking	34

Contents

3.5	Magnetic Resonance Imaging	35
3.5.1	Basic Principles	35
3.5.2	Integration in the Treatment Room	35
3.6	Ultrasound Imaging	36
3.6.1	Basic Principles	37
3.6.2	Inter-Fractional Positioning of the Patient	38
3.6.3	Intra-Fractional Systems and Setups	40
4	Basic Treatment-Planning Setup for Ultrasound-Guided Radiosurgery	45
4.1	Patient Data and Planning Parameters	45
4.2	Robot-Arm Model for Image Guidance	46
4.3	Finding Blocked Beams	48
4.4	Experimental Setup and Evaluation Metrics	50
4.5	Results	51
4.6	Discussion	55
5	Optimizing a Robotic Setup for Ultrasound-Guided Radiosurgery	57
5.1	Optimization of the Robot Positioning	58
5.1.1	Search Space	58
5.1.2	Objective Functions	59
5.1.3	Optimization Strategies	60
5.1.4	Experimental Setup	64
5.2	Selection of an Efficient Set of Robot Configurations	65
5.2.1	Discretization	65
5.2.2	Post-Processing as a Set Cover Problem	66
5.2.3	Direct Integration into the Planning LP	67
5.2.4	Experimental Setup	69
5.3	Results	70
5.3.1	Optimized Positioning	70
5.3.2	Efficient Configurations	75
5.4	Discussion	77
6	Synchronizing Robot Motion and Beam Delivery in Ultrasound-Guided Radiosurgery	81
6.1	Model for Efficient Beam Delivery	82
6.1.1	The Traveling Salesman Problem	82
6.1.2	Modeling the Beam Delivery	84
6.1.3	Solving TSPs	84
6.2	Model for Synchronized Robot Motion	85
6.2.1	The Generalized Traveling Salesman Problem	86
6.2.2	Modeling the Robot Synchronization	87
6.2.3	Solving GTSPs	88
6.3	Experimental Setup	89
6.3.1	Robot Models and Parameters	89

6.3.2	Methods for Synchronization	91
6.4	Results	93
6.5	Discussion	94
7	Towards Optical Coherence Tomography for Image Guidance in Cranial Radiosurgery	101
7.1	Basic Principles	102
7.1.1	Low-Coherence Interferometry	103
7.1.2	Fourier-Domain OCT	104
7.1.3	Lateral Scanning	105
7.1.4	Practical SS-OCT Image Reconstruction	106
7.2	OCT-Based Navigation and Tracking	107
7.3	Design of an OCT-Based Tracking System	108
7.3.1	Hardware Setup	110
7.3.2	Data Processing Pipeline	113
7.3.3	Calibration to Cartesian Coordinate System	115
7.3.4	Multi-Template Tracking	116
7.4	Tracking Algorithms	117
7.4.1	Phase Correlation	118
7.4.2	MOSSE Filter	119
7.4.3	Deep Learning Methods	121
7.5	Experimental Setup for System Characterization	121
7.5.1	Phantoms	121
7.5.2	Simulating Motion	123
7.5.3	Experiments	124
7.5.4	Evaluation Methods	125
7.6	Results	127
7.6.1	Phase Correlation	127
7.6.2	Cartesian Calibration	127
7.6.3	Single-Template Tracking	128
7.6.4	Multi-Template Tracking	132
7.7	Discussion	136
8	Conclusion and Outlook	141
	Bibliography	145

List of Acronyms

AP	anterior-posterior
ATSP	asymmetric traveling salesman problem
CBCT	cone beam computed tomography
CT	computed tomography
CTSP	clustered traveling salesman problem
CTV	clinical target volume
DOF	degree of freedom
DRR	digitally reconstructed radiograph
DVH	dose-volume histogram
EBRT	external beam radiation therapy
EM	electromagnetic
FDML	Fourier-domain mode locking
FD-OCT	Fourier-domain optical coherence tomography
FFT	fast Fourier transform
FOV	field of view
GPU	graphics processing unit
GTSP	generalized traveling salesman problem
GTV	gross tumor volume
HPP	shortest Hamiltonian path problem
iff	if and only if
IGRT	image-guided radiation therapy
ILP	integer linear program
IMRT	intensity-modulated radiation therapy
LINAC	linear accelerator
LP	linear program
MIP	mixed-integer linear program
MLC	multi-leaf collimator
MRI	magnetic resonance imaging
MU	monitor unit
NMR	nuclear magnetic resonance
OAR	organ at risk
OCO	optimization of the coverage
OCO-OV	OCO objective value
OCT	optical coherence tomography
PRV	planning organ-at-risk volume
PTV	planning target volume
RF	radio frequency

List of Acronyms

RL	right-left
RMSE	root-mean-square error
SAD	source-axis distance
SBRT	stereotactic body radiation therapy
SCP	set cover problem
SD-OCT	spectral-domain optical coherence tomography
SI	superior-inferior
SimAn	simulated annealing
SIRC	segmental inverse robot construction
SRS	stereotactic radiosurgery
SS-OCT	swept-source optical coherence tomography
STSP	symmetric traveling salesman problem
SVD	singular value decomposition
TAUS	transabdominal ultrasound
TD-OCT	time-domain optical coherence tomography
TPUS	transperineal ultrasound
TRUS	transrectal ultrasound
TSP	traveling salesman problem
VMAT	volumetric arc therapy
VOI	volume of interest

1 Introduction

1.1 Motivation

Cancer is a frequent cause of death world-wide. The WHO's International Agency for Research on Cancer reported for 2018 a total of 9 555 027 deaths over all ages and sexes [107]. Lung cancer has caused 18.4% of these deaths and is the most frequent type with twice as many deaths as the following colorectal cancer. Liver and prostate cancer caused 8.2% and 3.8%, respectively. The most-frequently diagnosed types are lung and breast cancer, which have each been diagnosed in 11.6% of 2018's 18 078 957 new cases. Prostate cancer has been diagnosed in 7.1% of them. For Germany, the Robert Koch Institut reported in 2013 a total of 223 093 deaths and 482 470 new cases, which is close to the average rate within the European Union [213]. From the 252 550 new cases for men, 59 620 were due to prostate cancer. Siegel et al. [251] estimate for the USA that one in three persons will develop some type of invasive cancer in their life.

Different treatment approaches exist for cancer and applying combinations of them is common. They include a surgical resection, chemotherapy, and radiation therapy. Chemotherapy and radiation therapy use drugs and ionizing radiation, respectively, to kill cancer cells. All approaches carry risks, can cause secondary diseases, and cannot guarantee a cure. The decision for the most suitable approach depends, for example, on site, size, and spread of the tumor, on the impact on quality of life, and on the age of the patient. Miller et al. [179] report extensive statistics on the cancer treatments of tumors at different stages in the USA.

In the field of radiation therapy, a variety of approaches and systems exist (see Section 2.1). One of them is the external delivery of photon beams with energies of some megaelectronvolts. Delivering a sequence of beams from multiple directions, as illustrated in Figure 1.1, allows to deposit a high amount of energy in the tumor area for cell killing while other healthy structures can be spared to minimize their damage and the probability of secondary diseases. One specific treatment system is the CyberKnife (Accuray, USA), which uses a robot arm to deliver a high number of beams from directions all around the patient (see Section 2.2). It has been designed in particular for hypofractionated radiosurgery, which means that only a few irradiation sessions, called fractions, are necessary but a high dose is delivered in each of these. In contrast, the conventional radiation-therapy approach considers more fractions with lower doses each. Employing higher doses requires a higher spatial accuracy to ensure that the actual target is irradiated and the dose to healthy structures is minimized. This requires a sophisticated planning of the treatment based on volumetric image data (see Section 2.3). However, differ-

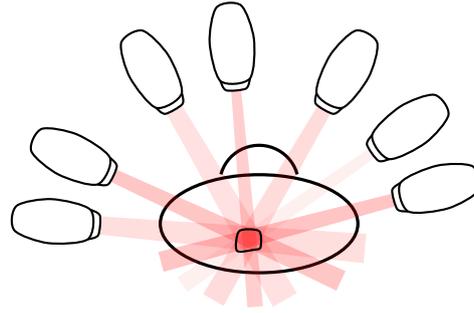


Figure 1.1: Beam delivery from multiple directions with different beam shapes and activation times allows to realize a high dose (red) in the target and low doses to surrounding healthy structures due to superposition.

ent kinds of motion will occur during the treatment which change the patient's anatomic geometry relative to the treatment machine. Thereby, the pre-treatment image data used for planning is no longer reflecting the actual treatment geometry leading to irradiation of wrong areas if the plan is not corrected. Common intra-fractional motion events include involuntary patient movements, respiration-induced motion, and spontaneous organ motion. Furthermore, it is necessary to verify before each fraction that there is no inter-fractional motion, i.e., the patient is positioned exactly as assumed by the treatment plan and no relevant anatomical changes, for example due to different bladder or rectum filling, appeared. Therefore, a fast, accurate, and reliable image guidance is needed to detect and, ideally, compensate any motion w.r.t. the treatment plan. Several imaging modalities are commonly used for image-guided radiation therapy (IGRT). Their application depends on treatment site and system, but they often involve X-ray-based imaging supported by optical systems, electromagnetic (EM) tracking, ultrasound imaging, or magnetic resonance imaging (MRI) (see Chapter 3).

In the course of this thesis, we consider the following properties as interesting for effective intra-fractional image guidance in radiosurgery.

Continuous Imaging Only continuous acquisition and real-time processing of images can ensure to detect all relevant motion events. The first aspect mainly requires that the duration of imaging is not safety critical for the patient. The latter requires sufficiently fast imaging procedures and algorithms.

Markerless Imaging The direct imaging and tracking of the actual target instead of tracking surrogates or markers allows, in general, for lower uncertainty and less complex setup. In contrast, we would need an accurate and reliable model to correctly derive the motion of the target from the observed motion of a surrogate. This requires a sophisticated design and placement of fiducial markers, which additionally have to be compatible to the treatment environment.

Volumetric Imaging Reconstructing 3D information from projective images is feasible but introduces additional errors. Furthermore, these errors can be non-isotropically distributed resulting in a varying accuracy depending on the direction of motion. In contrast, volumetric imaging allows to directly access 3D shapes and their motion or deformation.

Additionally, an image guidance system should have as few requirements on the environment, i.e., treatment room and patient setup, as possible. This aspect allows for an easy integration and flexible use in existing setups.

If we consider these properties and radiosurgery targets in the abdomen, like prostate or liver, ultrasound imaging is an interesting candidate modality. It is not harmful and provides direct real-time volumetric imaging of organs which are not shadowed by bones or air (see Section 3.6.1). In the context of radiation therapy, ultrasound imaging has been mainly used for inter-fractional imaging and patient positioning in the past (see Section 3.6.2), but it recently received increasing interest for intra-fractional imaging and guidance. The most flexible approach for abdominal targets is to place an ultrasound transducer onto the abdominal wall close to the target. However, the ionizing radiation in the treatment room forbids a manual operation of the transducer as in conventional diagnostic ultrasound imaging. For this reason, different robotic systems have been proposed (see Section 3.6.3). However, both transducers and robots contain radiopaque materials which interfere with treatment beams. The beam directions shadowed by the transducer or robot are not available for treatment which can result in unacceptable deterioration of the plan quality [79]. Therefore, transabdominal robotic ultrasound requires a sophisticated setup w.r.t. the beam delivery for practical application.

For cranial targets, like for example brain lesions, ultrasound imaging is not applicable due to the skull. We have for these sites, however, in general a more rigid relation between external motion of the head and the internal motion of the target than for abdominal sites. Hence, tracking external motion can be sufficient. Considering our defined properties of interest, we find that optical coherence tomography (OCT) might be an interesting approach. Employing infrared light and interferometric measurement, it provides volumetric imaging at tens-of-micrometers spatial resolution and allows to visualize also structures a few millimeters below a surface (see Section 7.1). The basic principle is comparable to ultrasound imaging, because OCT measures path length differences of received reflections to reconstruct depth profiles. The additional information about sub-surface structures acquired by OCT compared to conventional optical imaging techniques could allow for tracking targets without requiring markers even if their surface is poorly structured. Thereby, OCT might outperform the common superficially scanning modalities.

1.2 Research Questions

The purpose of this thesis is to investigate aspects of markerless volumetric image guidance in the context of robotic radiosurgery. As motivated in the previous section, we consider ultrasound imaging for abdominal targets and OCT-based tracking for cranial targets.

Robotic ultrasound guidance for abdominal treatments has already been studied. Especially, there is previous work on employing a kinematically redundant robot arm to efficiently avoid blocking of beam directions by the ultrasound transducer or the robot arm [78, 79]. The kinematic redundancy further allows to move the robot arm while maintaining a steady pose of the transducer as required for continuous imaging. This, in turn, allows to actively elude beams during treatment to some degree and achieve even less impact on the treatment plans. However, these studies only showed that there are positionings and configurations of the robot arm leading to acceptable treatment plans. They did not show how to determine them and how to implement an actual synchronization between the motion of the ultrasound robot and the beam delivery.

In this thesis, our first research question is how to automatically find suitable positionings and configurations of a kinematically redundant robot arm for ultrasound guidance (Chapter 5). This question implies an analysis and optimization of the setup and its relation to the quality of treatment plans.

The second research question is how to achieve an efficient synchronization between movements of the robot arm and the beam delivery (Chapter 6). Both the motion of the ultrasound robot and the motion of the robot for beam delivery take time prolonging the treatment and should be minimized.

Regarding OCT-based tracking, there are no direct previous studies on which we could build because it has not been used for such kind of application yet. Therefore, the third research question of this thesis is whether it is possible to develop a markerless volumetric tracking system based on OCT (Chapter 7). This involves both the design of a prototypical system and an experimental characterization on phantoms.

1.3 Organization

The remaining chapters of this thesis are organized as follows. The second chapter introduces the physical and mathematical background of radiation therapy and treatment planning. Subsequently, we review the state-of-the-art imaging modalities for IGRT in Chapter 3 and especially discuss the application of ultrasound imaging. In Chapter 4, we introduce our basic experimental setup and parameters for treatment plan optimization during robotic-ultrasound guidance, which we consider in the two following chapters as well. The first of these two chapters proposes and evaluates methods to optimize the robotic setup, addressing the first research question. The second one describes and evaluates methods for the syn-

chronization between the robot for ultrasound guidance and the robot for beam delivery as outlined in the second research question. In Chapter 7, we move from abdominal to cranial targets. Following the third research question, we introduce OCT and describe and evaluate a prototypical tracking system based on this modality. Finally, we summarize the results of this thesis and provide an outlook for future research.

In the following, we outline the contents of each chapter in more detail. Parts of this thesis have been published in journals [230, 232, 233] and have been presented at conferences [228, 231, 234, 235].

Chapter 2: Radiation Therapy and Robotic Radiosurgery After starting with a short general overview of the different kinds of radiation therapy, basic physical properties, and available clinical systems, this chapter focuses on introducing treatment planning for robotic radiosurgery as implemented by the CyberKnife system. It describes the required input data, dose calculations, and the mathematical foundations of the inverse planning problem.

Chapter 3: Image Guidance in Radiation Therapy In this chapter, we review the state-of-the-art modalities for image guidance and motion tracking in radiation therapy. While the main part of this thesis only considers the task of intra-fractional tracking, i.e., tracking motion occurring during an irradiation session, we also discuss systems for patient positioning and compensation of inter-fractional changes in this chapter. The chapter covers X-ray-based imaging, optical systems, EM tracking, and MRI. Furthermore, we put a major emphasis on ultrasound imaging. In particular, we will discuss an intra-fractional setup for markerless volumetric image guidance employing a kinematically redundant robot arm, which we analyze and optimize in the following three chapters.

Chapter 4: Basic Treatment-Planning Setup for Ultrasound-Guided Radiosurgery This chapter introduces the basic treatment planning setup, which we consider for image guidance by robotic transabdominal ultrasound imaging during radiosurgery. A description of the robot model under consideration is given and the patient data of real cases of prostate cancer, which we use for evaluation, are presented. Furthermore, we provide and discuss basic planning aspects and results when employing a kinematically redundant robot arm for ultrasound guidance.

Chapter 5: Optimizing a Robotic Setup for Ultrasound-Guided Radiosurgery In this chapter, we propose and evaluate methods for optimal positioning of the ultrasound-guiding robot arm in terms of plan quality. We especially analyze the selection of suitable configurations provided by the arm's redundant kinematics. Parts of this chapter have been published in a journal [232] and have been presented at conferences [228, 231].

Chapter 6: Synchronizing Robot Motion and Beam Delivery in Ultrasound-Guided Radiosurgery When exploiting the kinematic redundancy of our robot arm under consideration for ultrasound guidance, a synchronization to the beam delivery is necessary. This chapter describes and evaluates methods to model and optimize this synchronization problem. Parts of this chapter have been published in a journal [230].

Chapter 7: Towards Optical Coherence Tomography for Image Guidance in Cranial Radiosurgery Ultrasound imaging is not suitable for cranial targets. Therefore, we propose a prototypical tracking system based on OCT for markerless volumetric image guidance during cranial treatments in this chapter. We describe the hardware setup and data processing and provide an experimental characterization of the prototypical system on phantoms. Parts of this chapter have been published in a journal [233] and have been presented at conferences [234, 235].

Chapter 8: Conclusion and Outlook This final chapter summarizes the findings of this thesis and provides an outlook on aspects for future research.

2 Radiation Therapy and Robotic Radiosurgery

In this chapter, we discuss the basic aspects of radiation therapy and especially the sub-fields of stereotactic body radiation therapy (SBRT) and robotic radiosurgery. After a general introduction to the medical and physical background of radiation therapy, we add more details on treatment with a CyberKnife system. Besides the technical aspects, we describe algorithms for dose calculation and inverse treatment planning.

2.1 Medical and Physical Background

Today, various methods for treating patients with cancer exist. Besides for example surgical resection or chemotherapy, there are different approaches which employ ionizing radiation to kill the cancerous cells and they are commonly summarized as radiation therapy or radiotherapy. Also combinations of the mentioned techniques are common [179]. For example, radiation therapy can follow a surgical resection of the macroscopic parts of a tumor.

A basic measure in radiation therapy is the delivered dose. It describes the absorbed energy per unit mass and has the SI-unit gray (symbol Gy), which is joule per kilogram in base units. The more energy is absorbed, the more ionization occurs in the material, i.e., the more pairs of electrons and positive ions are produced. In the human body, ionization causes damaging or death of cells. An important property of tissue and tumors for radiation therapy is their radiosensitivity. The higher the radiosensitivity of a tissue, the less dose is required to damage or kill the cells. In particular, a high radiosensitivity of tumor cells compared to surrounding healthy tissue can facilitate safe and efficient treatment by irradiation. Furthermore, the applicability of radiation therapy typically requires that the tumor is rather local. If the tumor is spreading throughout the whole body, a reasonable irradiation is often not possible.

2.1.1 Types of Radiation Therapy

Radiation therapy is divided into different sub-classes which differ in the type of radiation and the way of delivery. For the latter, we can differentiate between external beam radiation therapy (EBRT), which we will consider throughout this thesis, and brachytherapy.

In brachytherapy, the source of radiation is placed very close to the target area. For example, sealed or unsealed seeds are placed directly into or next to the tumor tissue. These seeds contain a radiation source with only a very short range such that a precise treatment of the cancer is possible while the surrounding healthy tissue is spared. Depending on the application, either a superficial placement is feasible or a brachytherapy needle is used for insertion of the seeds. However, this technique is only suitable for localized and rather small tumors in sufficiently accessible areas. Brachytherapy involves many further details which have resulted in many sub-types of it. For example, we have to distinguish how long the source remains within the patient, the dose rate of the delivery, and the type of the employed radionuclide.

The name EBRT already implies that in this technique the beams of ionizing radiation are created outside of the patient. Therefore, a beam has to, at least partially, pass through the patient's body in order to reach an internal target, like for example the prostate, the liver, or the lungs. In contrast to brachytherapy, it is thus unavoidable to also irradiate and potentially damage healthy structures, which also leads to longer treatment durations.

In EBRT, we can use different types of radiation. The most important difference between them is the process of energy absorption when penetrating through tissue. In the following we will briefly compare the behavior of electron, proton, and photon beams. The latter is the type of radiation employed by the CyberKnife. For a more detailed physical description, see for example Pawlicki et al. [199].

Electron beams A beam of electrons continuously interacts with other electrons while passing through a material. At each interaction, an electron of the beam is scattered, i.e., changes its direction of travel, and transmits a minor part of its kinetic energy to an electron of the material. Eventually, all energy of an electron has been transmitted by these interactions and the electron stops. This is an interesting aspect, because no energy will be deposited beyond a certain depth at which all electrons of the beams have stopped with a high probability (Figure 2.1a). Thereby, no deeper healthy structures will be damaged in a treatment. The falloff of the dose, however, happens rather shortly below the skin. This approach is therefore mainly used for superficial targets or intra-operatively.

If an electron of the material is ejected from its atom by an interaction with the electron beam, then this atom is ionized. The ejected electron, in turn, can interact with electrons of other atoms and ionize these as well. If further electrons are ejected this way, they are referred to as secondary electrons.

Proton beams Proton-beam and electron-beam radiation therapy can be summarized as charged-particle therapy as they have similar properties. However, protons are much heavier making them less subject to scattering. Furthermore, they penetrate much farther into tissue and deliver a very high dose shortly before they stop. This so-called Bragg peak is shown in Figure 2.1a. It allows to con-

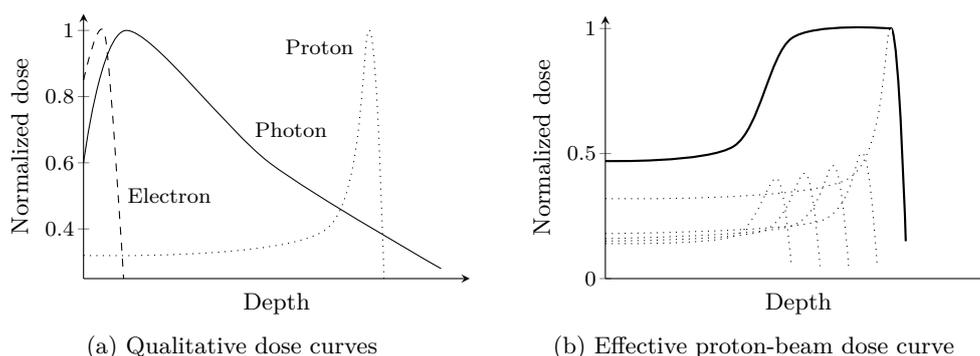


Figure 2.1: Qualitative sketch of the dose curves for electron (a, dashed), proton (a, dotted), and photon beams (a, solid). The effective dose curve (b, solid) for proton beams can be flattened by using beams with different energies (b, dotted).

concentrate the dose delivery in an internal target, if its depth matches the depth of the Bragg peak. Note that this does not relate to the geometrical depth, but the effective depth considering the varying tissue densities along the path. Typically, this is expressed as the equivalent depth in water or another reference material. While a sharp peak is typically too extreme for a homogeneous treatment of a target, the dose curve can be flattened by using beams of different energy because the energy relates directly to the depth of the Bragg peak (Figure 2.1b).

Photon beams The radiation in photon-beam-based radiation therapy are X-rays or γ -rays, though their distinction is not fully consistent in the literature. In earlier days, their distinction was often based on the energy, or equivalently the wavelength, of the photons. Photons with energies below arbitrary limits in the order of 100 keV were considered as X-rays and photons with higher energies as γ -rays. Today, γ -rays are rather defined as the product of nuclear decay and X-rays as the product of electron interactions outside of the nucleus. For the radiation therapy itself, the distinction does not matter as both X-rays and γ -rays show equivalent behavior within tissue for the same energy.

The main type of interaction between tissue and photon beams with energies of some megaelectronvolt, as employed in the CyberKnife, is Compton scattering [199]. An incident photon collides with an outer electron of an atom. If the photon's energy is above the electron's binding energy, it ejects the electron. The photon and the electron then move at different angles and the sum of their new kinetic energies equals the previous energy of the incident photon reduced by the rather small binding energy. The higher the electron density of a material, i.e., the number of electrons per unit mass, the more Compton scattering will occur. However, the mass per volume has to be considered as well when comparing the attenuation of photons in different materials. For example, the high density of bones compensates their rather low electron density compared to fat tissue leading

in consequence to a much higher attenuation. Furthermore, the ejected electrons can ionize further atoms as a secondary effect, similar to electron beams.

In contrast to proton beams, the dose curve of photon beams has its peak close to the skin (Figure 2.1). For the CyberKnife system, it is about 15 mm below the skin. Nevertheless, the exposure to the skin is considered as lower than for proton beams [252]. The subsequent falloff of the dose curve is rather flat allowing to deliver dose to deep targets. However, dose delivery does not end at a specific depth but continues throughout the whole body. Besides these dosimetric differences, photon-therapy devices are considerably cheaper than proton-therapy devices.

2.1.2 Radiation Therapy with Photon Beams

Photon beams are typically generated by a linear accelerator (LINAC) today. A prominent exception is the Gamma Knife (Elekta, Sweden), which is in clinical use since 1968 [157]. It uses Cobalt-60 sources to produce γ -rays by radioactive decay. The Gamma Knife is designed for treatment of brain tumors and for this reason it consists of about 200 sources arranged as a helmet-like device around the patient's head.

In contrast, a LINAC accelerates electrons with a voltage of some megavolts and directs them onto a target material like tungsten, similar to conventional X-ray tubes. When colliding with the target, X-rays are released. Alternatively, accelerated electrons could be used directly for electron-beam therapy. The activation time of a treatment beam produced by a LINAC is typically quantified in terms of so-called monitor units (MUs) rather than a unit like seconds. The MUs describe the ionization by the beam occurring within an ionization chamber and they are proportional to the beam's intensity. Therefore, delivering a defined number of MUs results in a defined delivered dose, even if the beam's intensity produced by the LINAC varies over time [199]. Typically, the chambers are calibrated such that 1 MU corresponds to the delivery of 1 cGy at a reference distance in a reference material with a reference beam shape.

For radiation therapy, the profile of the photon beams is an important issue. Many therapy systems employ a flattening filter to achieve a plateau-like beam profile, i.e., the delivered dose remains approximately constant with increasing distance to the beam's central axis within the maximum field size. Naturally, there would be a decay of the dose with increasing distance. The actual shape of a treatment beam is defined by a subsequent collimator. In the simplest cases, these collimators have a circular hole. Due to the opening angle and thereby cone-like shape of the resulting beams, the diameter of a collimator is typically not expressed as the actual physical diameter of its hole, but as the effective diameter at a reference source-axis distance (SAD). The SAD is the distance between the LINAC and a reference point, in gantry-based systems typically the isocenter. A more complex type of collimators are multi-leaf collimators (MLCs) as sketched in Figure 2.2. They consist of a number of movable leaf pairs and allow to form

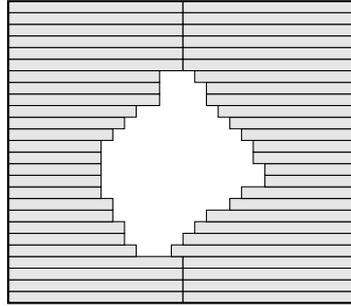


Figure 2.2: Collimation with an MLC allows to define beam shapes which closely follow the shape of a target. Each row of the collimator consists of a motorized pair of leaves which block radiation.

beam shapes which closely follow the cross-sectional shape of the target.

Due to the dose curve of photon beams, irradiation of a target from one single direction does not provide a meaningful dose distribution for treatment. Irradiation from different directions with appropriate beam shapes and different activation times of the beams, in contrast, allows to shape dose distributions which show a high dose within the target and low doses to other tissue. This technique is called conformal radiation therapy. A closely related concept is intensity-modulated radiation therapy (IMRT). If we consider conformal radiation therapy with an MLC, the shape of a beam matches the shape of the target w.r.t. this beam's direction and the beam is active for some time to deliver a certain amount of dose to the target. In IMRT, we can additionally vary the activation times of different parts of the beam. This is realized by closing some leaves of the collimator for some time in order to reduce the dose delivered to that part of the target. In other words, we do not deliver a single beam from every direction but a weighted sum of multiple beams with varying shape and activation time.

Many treatment machines are realized as some type of gantry, i.e., the LINAC rotates around the patient similar to a computed tomography (CT) acquisition. In the simplest case, this leads to coplanar and isocentric treatment beams which means that the central axes of all beams are within a plane and they intersect in one point, the isocenter. This requires to position the patient such that the isocenter is in the target. Gantry-based systems allow for an extension of IMRT which is called volumetric arc therapy (VMAT) [194] and is commercially implemented for example as RapidArc (Varian Medical Systems, USA). Instead of delivering a discrete sequence of beams, the LINAC is continuously moving and delivering dose in a VMAT treatment. To realize a specified dose distribution, we need to optimize its trajectory which includes variation over time of the velocity, of the shape of the MLC, and of the beam intensity. The main advantage of this technique are shorter treatment times due to more time-efficient dose delivery [256, 276].

The limitation to coplanar and isocentric beams in gantry-based systems can be

removed by employing a robotic treatment couch. Moving the couch allows to shift the isocenter w.r.t. the target and, by rotations, to irradiate from non-coplanar directions. Non-coplanar and non-isocentric beams typically allow for a steeper dose falloff around the target and better sparing of healthy structures [103, 249]. In helical tomotherapy, especially implemented in the TomoTherapy (Accuray, USA) system, the patient couch is axially moved through the rotating gantry to realize a continuous helical beam delivery [170]. The combination of couch rotation and VMAT has been proposed, sometimes under the name 4π therapy, and allows for continuous spherical trajectories around the patient providing time-efficient irradiation from more directions [48, 253, 282].

A treatment is separated into multiple sessions, which is called fractionation [260]. In each fraction, only some fraction of the total dose is delivered. The time between the fractions allows healthy cells to recover from collateral damage, while tumor cells are often not able to repair themselves sufficiently fast. Thereby, side effects can be reduced. Fractionation can result in daily treatment over several weeks. Therefore, the accurate and reproducible positioning and aligning of the patient w.r.t. the treatment system is an essential issue. Any variation occurring between two fractions is referred to as inter-fractional while a variation in the course of the delivery of a fraction is referred to as intra-fractional.

Another approach is hypofractionation, which only uses a few fractions with high doses each, or even only one single fraction [89, 127]. These approaches are commonly called stereotactic, which refers to a fixation of the patient to allow for a very accurate delivery of the dose to the target. This accuracy is especially needed when delivering high doses. Stereotactic treatments are further separated into stereotactic radiosurgery (SRS) in case of cranial targets and SBRT or stereotactic ablative body radiotherapy (SABR) in case of other sites, including lungs or prostate. Due to the high dose per fraction, the dose has to be delivered from many different directions to efficiently spare surrounding healthy tissue. While proper fixation is feasible for cranial targets, accurate dose delivery to other sites always requires high-quality image guidance, i.e., IGRT, to ensure that the target has not moved and is still correctly irradiated [16]. The Gamma Knife has been pioneering in SRS. Today, the CyberKnife is an important system for SRS and SBRT.

2.2 The CyberKnife System

The CyberKnife has been originally proposed for SRS of intra-cranial targets. The aim was to develop a system which does not require a fixating stereotactic frame [2, 3]. Instead, it uses image guidance and a robot-mounted LINAC (Figure 2.3) to allow for delivery of non-coplanar and non-isocentric beams and to adapt the irradiation to motion [240]. It is now used for treatment of targets throughout the body [125], including SBRT of lungs, liver, and prostate. Its current version is the CyberKnife M6, which is technically summarized in Kilby et al. [126].

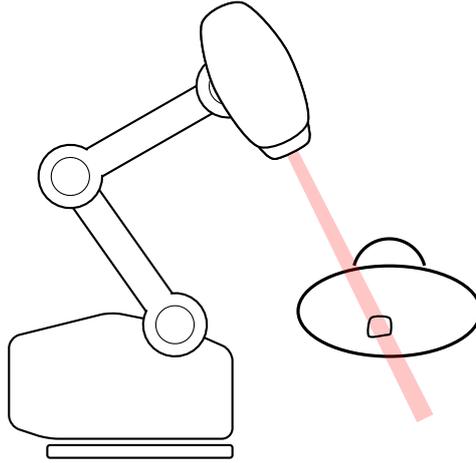


Figure 2.3: Illustration of treatment with a robot-mounted LINAC as implemented in the CyberKnife.

The CyberKnife employs a LINAC without flattening filter, which produces 6 MeV photon beams at a dose rate of 1000 MU min^{-1} . Under reference conditions, 1 MU corresponds to 1 cGy. The LINAC is mounted to an industrial robot arm with six degrees of freedom (DOFs). This allows to deliver dose from arbitrary points around the patient with beams having arbitrary orientations. Thus, CyberKnife treatments can be highly non-coplanar and non-isocentric without moving the patient. Note that mimicking isocentric treatments is also possible by defining one or more virtual isocenters within the target. In practice, there are some restrictions for the beam delivery, however. Beams only start at a discrete set of around 100 to 200 points, called beam nodes, which are located within a half spherical shell around the patient. Their average distance to the target is about 800 mm, which serves as the reference SAD value for several specifications due to the lack of an isocenter. However, there has to be a safe path for the robot arm to reach all beam nodes with the equipped LINAC, including sufficient margins allowing to compensate potential motion. The orientations of the beams are arbitrary, though intersection with the target is obviously mandatory.

Initially, the CyberKnife only provided 12 fixed-size circular collimators to realize beams with diameters from 5 mm to 60 mm at an SAD of 800 mm. The collimators can be exchanged automatically by the robot arm at a dedicated tool table. Nevertheless, the path of beam nodes has to be driven by the robot arm for each collimator which substantially prolongs the treatment, because the robot arm moves its end-effector for safety reasons only at up to 3% or 4% of its maximum speed [175, 230]. The maximum speed is about 2 m s^{-1} in Cartesian space. To reduce the overhead of changing the collimator, the IRIS collimator has been introduced which changes its diameter by opening and closing a mechatronic iris [53]. However, it is only used to realize the same discrete set of diameters as provided by the fixed-size collimators. Recently, MLCs became available for the

Cyberknife [8, 75]. The InCise 2 MLC provides 26 pairs of tungsten leaves which have a width of 3.85 mm and provide a total treatment field of size $115 \times 100 \text{ mm}^2$ at an SAD of 800 mm. As for MLCs in general, the idea is to provide more efficient treatments of similar quality as with circular collimators [175]. Nevertheless, the fixed-size collimators still provide the beam shapes with lowest uncertainty [126]. The CyberKnife only provides step-and-shoot treatments where a sequence of individual beams is delivered. However, there have been experimental approaches described recently for continuous irradiation with the MLC along arcs, similar to non-coplanar VMAT [13, 123, 124].

For intra-fractional image guidance, the CyberKnife is equipped with stereoscopic X-ray imaging and an optical tracking system. This basic setup has already been introduced in 2000 [241]. Detected target motion can be compensated by adjusting the pose of the LINAC by moving the robot arm. If the motion exceeds some centimeters or some degrees depending on the axis, an emergency stop is triggered in order to re-position the patient. Furthermore, the CyberKnife provides a robotic treatment couch. However, this couch is only intended to support positioning of the patient and does not provide active motion compensation during an irradiation, because there is no rigid relation between the couch and the patient [126]. Positioning is further supported by, for example, a laser pointing from the LINAC along the beams' central axes.

2.3 Inverse Planning

In forward planning for radiation therapy, the planner manually selects a set of beam directions and shapes. The computer then computes the activation times of these beams to match a given dose and visualizes the resulting dose distribution. Based on this feedback, the planner modifies the input parameters until they result in an acceptable distribution. This approach is only suitable for rather easy cases and few beams. Not only for complex scenarios like CyberKnife treatments with a few hundred non-coplanar and non-isocentric beams, forward planning has been mostly replaced by inverse planning. In inverse planning, the physician defines volumes of interest (VOIs) in image data and assigns goals and constraints for the dose distribution. Based on this input, the planning software automatically sets up and solves an optimization problem. It outputs a set of beam directions, shapes, and activation times which fulfills all constraints and optimizes one or more selected goals.

2.3.1 Data for Treatment Planning

In order to compute the dose delivered by a treatment beam to a region within the patient, a high-resolution planning CT is required. From this volumetric scan, the electron densities can be derived. The first planning step, however, is the delineation and contouring of relevant VOIs by a physician [30]. Due to the rather poor soft-tissue contrast of CT, delineation can be supported by other imaging

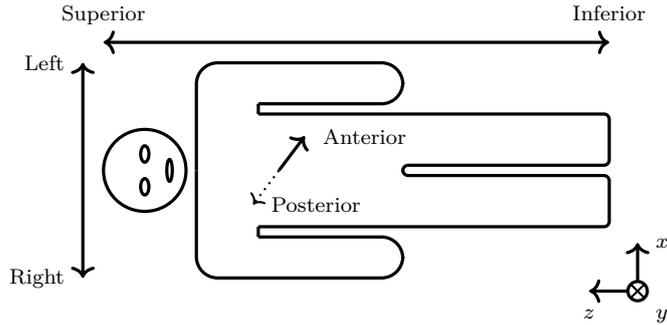


Figure 2.4: Relation between anatomical coordinate axes and the axes of our world coordinate system.

modalities like MRI or a positron-emission tomography (PET) scan with a suitable tracer [117]. Via multi-modal image registration, the defined contours can be transferred to the planning CT for all subsequent planning steps. Delineation is typically done slice-wise in volumetric images, either manually or supported by an image-segmentation algorithm.

Different anatomical coordinate systems are used in medical contexts. The following axes can be defined with respect to the patient (Figure 2.4). The superior-inferior (SI) axis points from the head (superior) towards the feet (inferior), the anterior-posterior (AP) axis points from the front (anterior) to the back (posterior), and the right-left (RL) axis points from the right shoulder to the left shoulder. The planning CT typically consists of cross-sectional slices acquired along the SI-axis. If a beam intersects with the patient outside of the planning CT scan, that beam has to be rejected because the dose delivered to the patient could not be completely evaluated. Therefore, the planning CT should not only be centered around the target and all relevant VOIs but also cover sufficient additional tissue to allow for a wide range of beam directions. Exemplary slices from a contoured planning CT are shown in Figure 2.5. Furthermore, one point within the planning CT is defined as the reference point for all alignments relative to the treatment room, including the beam-node positions and the field of view (FOV) of the stereoscopic X-ray imaging. In this thesis, we use a so-called LPS coordinate system for treatment planning, i.e., the x -axis is the RL-direction, the y -axis is the AP-direction, and the z -axis is the negative SI-direction. In terms of the treatment room, in which the patient is lying on a couch, the y -axis points towards the floor.

The most essential VOI to be delineated is the tumor area, which is the target for the irradiation. The resulting VOI is called gross tumor volume (GTV) and describes the macroscopically identifiable part of the tumor. An extension of the GTV by microscopic parts of the tumor leads to the clinical target volume (CTV). This accounts for the typical spreading of cancerous cells and ensures that the whole actual tumor will be treated. If radiation therapy follows a resection, there

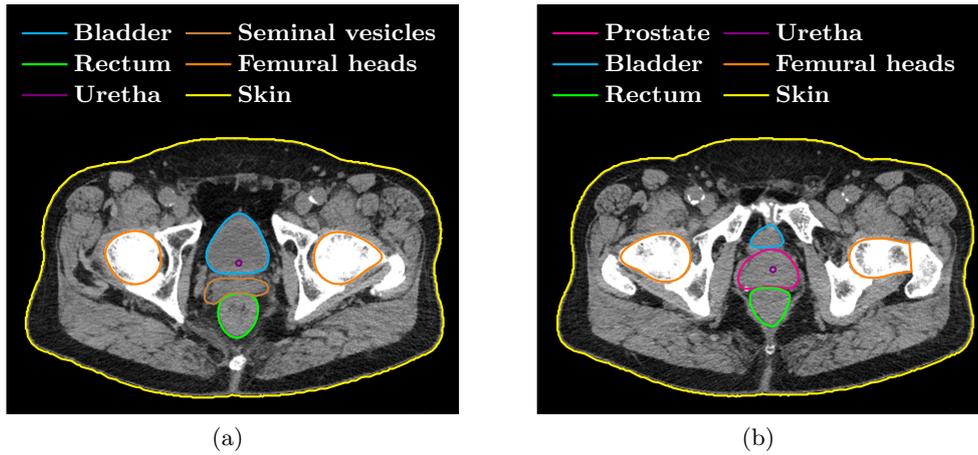


Figure 2.5: Two contoured CT slices from a planning CT for prostate cancer. The horizontal and vertical image axes correspond to the RL-direction and AP-direction, respectively. The offset of the two slices is 10 mm along the SI-axis.

might not even be a GTV anymore because all macroscopic parts have already been removed during surgery. Lastly, a further extension by safety margins is made and eventually leads to the planning target volume (PTV). The additional margins account for a range of inaccuracies which might occur. These include motion due to breathing, shifts and deformations due to for example bladder filling, patient movement, and variations and uncertainties in patient setup. On the one hand, a large margin increases the likelihood that the CTV actually receives the intended dose and the tumor cells are killed. On the other hand, larger margins result in higher doses to healthy tissue and thereby increase the likelihood of side effects. Therefore, a fundamental issue in radiation therapy is to achieve as small as possible margins by reducing the uncertainties, especially by high-quality image guidance during patient setup and treatment. See for example Yartsev et al. [280] for a survey on image-guidance approaches and their impact on margins for treatment of prostate cancer.

In addition, critical healthy structures are delineated. These so-called organs at risk (OARs) are typically radio-sensitive tissues which lead to significant medical consequences when receiving too much radiation. In case of prostate cancer, OARs can be for example bladder, rectum, or urethra. If an OAR is extended by a safety margin, it is referred to as a planning organ-at-risk volume (PRV). In this thesis, however, we only consider OARs without additional margins, which is typical for CyberKnife treatments. Figure 2.6 illustrates the overall nomenclature for delineated VOIs.

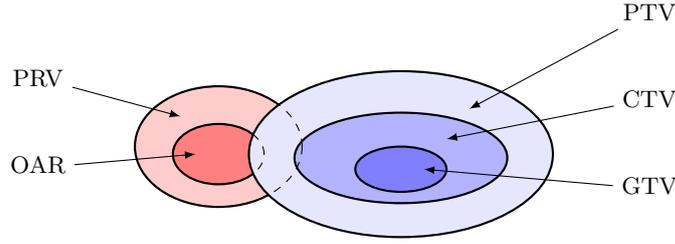


Figure 2.6: Illustration of GTV (macroscopic tumor), CTV (microscopic spread), and PTV (safety margins). In this sketch, we have an overlap with an OAR, which can also be extended by safety margins to a PRV.

2.3.2 Clinical Goals

To describe and compare the quality of treatment plans quantitatively, several figures of merit are in use. These are sometimes called clinical goals, because they are typically not directly equivalent to the objective functions of the actual mathematical optimization problems which are solved to create the plans. They are rather related to the shape of the dose distribution w.r.t. to the VOIs. In the following, we will introduce some of the most prominent clinical goals.

Dose Volume The dose volume V_d^{VOI} describes which percentage of a certain VOI receives at least the dose d . A common goal described by dose volumes is to restrict higher doses to a small fraction of an OAR. Furthermore, restrictions on under-dosage within the PTV can be expressed by dose volumes.

Volume Dose Complementary to the dose volume, the volume dose d_p^{VOI} is the highest dose which is received by at least p percent of a VOI.

Coverage If there is a prescribed dose d_{ps} for the PTV, then the coverage COV is the percentage of the PTV which receives at least d_{ps} . Therefore, it is simply the special dose volume

$$\text{COV} = V_{d_{\text{ps}}}^{\text{PTV}}. \quad (2.1)$$

and should be as high as possible, ideally 100 %.

Homogeneity Index There are different definitions for the homogeneity index HI in the literature [119]. They all relate in some way the maximum and minimum dose in the PTV to describe how homogeneous the dose is delivered. Sometimes the extreme values are replaced, for example, by the 5 % and 95 % quantiles of the dose distribution, respectively. This leads to

$$\text{HI} = \frac{d_5^{\text{PTV}}}{d_{95}^{\text{PTV}}}. \quad (2.2)$$

For a homogeneous irradiation, HI should be close to 1.

Conformity Index The conformity index CI measures how much tissue has been irradiated with at least d_{ps} relative to the size of the PTV V^{PTV} , i.e.,

$$CI = \frac{V_{d_{ps}}^{patient}}{V^{PTV}}. \quad (2.3)$$

The basic idea is that ideally only the PTV receives a high dose, there is a steep dose gradient around it, and there are no high-dose areas farther away from the PTV. There exist alternative formulations for (2.3) [59]. One of them is called the new conformity index

$$nCI = \frac{V_{d_{ps}}^{patient}}{V_{d_{ps}}^{PTV}} \cdot \frac{V^{PTV}}{V_{d_{ps}}^{PTV}} \quad (2.4)$$

and is the inverse of the Paddick conformity index [195]. Ideally, conformity indices reach a value of 1.

Note that none of these goals is individually meaningful, because each of them only describes one specific aspect of the dose distribution. Beside these quantitative values, the qualitative goal is to achieve a high dose throughout the target and spare the OARs as much as possible. Also the dose to other healthy tissue should be minimized, which requires a sharp falloff of the dose outside the PTV. Furthermore, there are more general goals as, for example, short treatment time which improves aspects like patient comfort, additional dose due to X-ray-based image guidance, and patient throughput. All these goals, however, are conflicting in general. Higher dose in the target requires higher dose in surrounding tissue due to the nature of dose delivery by photon beams. A high and homogeneous coverage conflicts with a sharp dose falloff. Employing more beam directions can improve OAR sparing but increases the treatment time. Therefore, treatment planning for radiation therapy is a multi-criteria optimization problem with conflicting objectives and both hard and soft constraints.

A standard method to visualize and evaluate treatment plans with respect to the dose distribution in a VOI are dose-volume histograms (DVHs) [50]. This type of diagram has the received dose on the horizontal axis and the percentage of a VOI, which receives at least that dose, on the vertical axis. However, both DVHs and the values of the clinical goals do not provide any spatial information. For spatial visualization, isodose curves can be plotted within slices of the planning CT or the deviation from a prescribed dose can be encoded by overlaid colors. Strategies have been proposed to provide the planner a more visual and interactive feedback during in the planning process [222, 268].

2.3.3 Dose Model

To evaluate and optimize the beams for a treatment, we need to have a physical model of the dose delivery and energy deposition of each beam within the body

depending on its activation time. For photon-beam therapy, the dose delivered by a beam depends on the local electron densities in the tissue. These densities can be derived from the intensities in the planning CT. For very precise calculation of the delivered dose, Monte-Carlo simulations are employed. However, these typically have a high computational effort and require detailed information about the treatment machine and setup. As an alternative, we can use simpler models yielding algorithms which are fast but still sufficiently accurate for basic planning tasks. In practice, however, final validation of the plan with a Monte-Carlo algorithm might be necessary [192, 274]. The CyberKnife provides a ray-tracing algorithm as an alternative to a Monte-Carlo-based algorithm [125]. Note that this algorithm is only suitable for treatments with the fixed-size or iris collimators. In case of MLC-based treatments, the CyberKnife provides a finite-size pencil beam (FSPB) algorithm [112, 113] as a fast alternative to Monte Carlo, which accounts for the arbitrary beam shapes realizable by an MLC.

In this thesis, we only consider the circular collimators and dose calculations by ray tracing. We follow Accuray’s physics guide. Note that some terms are adapted from classical gantry-based systems, although the CyberKnife clearly has more degrees of freedom. This leads to some discrepancies in the meaning of the terms compared to other literature. The ray-tracing algorithm considers the central axis of each beam and computes the dose for discrete points on discrete planes orthogonal to that axis. As sketched in Figure 2.7, we determine for each dose plane its SAD, i.e., distance to the LINAC, and its depth within the patient, both measured w.r.t. the central axis. For the depth, however, we do not consider the geometrical depth but correct it for the passed electron densities relative to water to obtain the effective depth. Several parameters of the treatment system’s beams are commonly calibrated within a homogeneous water phantom and would therefore be incorrect if we apply them considering the geometrical depth of the dose plane. One of these parameters is the tissue-phantom ratio (TPR), which describes the dose curve along the central axis. It is measured at the reference SAD of 800 mm in different depths for each collimator diameter C_{800} . Therefore, it is a function of the effective depth and the effective diameter

$$C_{\text{SAD}} = C_{800} \cdot \frac{\text{SAD}}{800 \text{ mm}}. \quad (2.5)$$

Another parameter is the output factor (OF). It describes scattering and the effect that photons can be reflected outside of the collimator’s open area. These distort the MU measurement in the ionization chamber.

To calculate the dose at a specific point \mathbf{p} within the dose plane, we need to consider further parameters which depend on its distance R_{SAD} to the central axis. The off-center ratio (OCR) describes the lateral falloff of the beam intensity and obviously depends on the diameter of the employed collimator. This parameter is only measured for the reference SAD and therefore we need to convert R_{SAD} to the corresponding

$$R_{800} = R_{\text{SAD}} \cdot \frac{800 \text{ mm}}{\text{SAD}}. \quad (2.6)$$

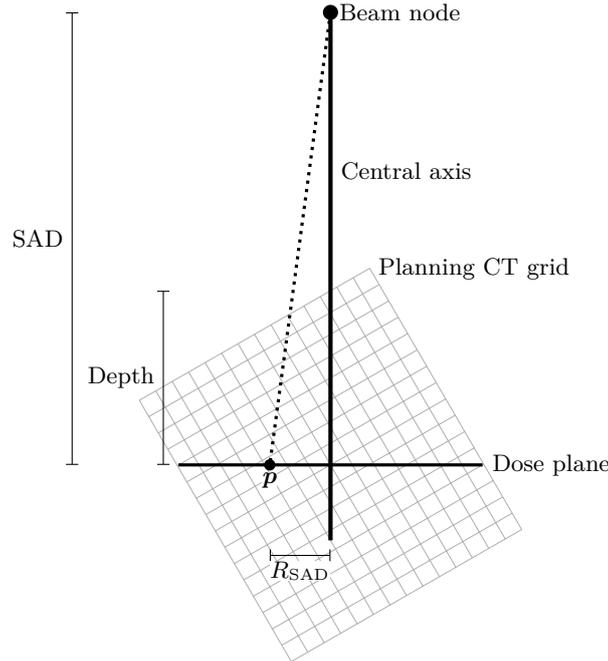


Figure 2.7: Two-dimensional sketch of a dose calculation point \mathbf{p} . The SAD and radius R_{SAD} are geometrically measured, while we have to distinguish between geometrical depth and effective depth w.r.t. electron densities along the path. Note that both SAD and depth are measured w.r.t. central axis and dose plane, independent of the actual position of \mathbf{p} within the plane.

With these parameters, which all provide only relative values, we can calculate the dose per MU delivered to point \mathbf{p} as

$$D(C_{800}, \text{SAD}, \text{depth}_{\text{eff}}, R_{\text{SAD}}) = 1 \text{ cGy MU}^{-1} \cdot \text{OCR}(C_{800}, R_{800}, \text{depth}_{\text{eff}}) \cdot \text{TPR}(C_{\text{SAD}}, \text{depth}_{\text{eff}}) \cdot \text{OF}(C_{800}, \text{SAD}) \cdot \left(\frac{800 \text{ mm}}{\text{SAD}}\right)^2, \quad (2.7)$$

which additionally includes the inverse-square law. The system is calibrated to provide a dose rate of 1 cGy MU^{-1} in (2.7) for the 60 mm collimator at the reference SAD of 800 mm in an effective depth of 15 mm on the beam's central axis.

2.3.4 Sampling of Cylindrical Beams

The theoretical freedom of the CyberKnife's beam delivery is already practically limited by its concept of beam nodes. Furthermore, some beams have to be rejected due to the limited extend of the planning CT. In practice, there might be further undesired beams, like beams which intersect with an OAR before intersecting the PTV. Nevertheless, it is still intractable to consider all possible combinations of beam node, collimator diameter, and orientation for planning.

The common strategy for CyberKnife treatment planning is to randomly sample a set of N beams. Thereby, only the activation times of these so-called candidate beams are variables for planning. We refer to the subset of the candidate beams which is actually active after planning as the set of treatment beams. With increasing number of candidate beams, we can expect the difference to the solution of the complete problem to vanish. Thus, we have to decide for a proper trade-off between the computational complexity and the quality of the solution, which is the typical trade-off decision in heuristic optimization. Different approaches have been proposed which aim to find more efficient candidate beams than random sampling. They consider, for example, anatomical projections from beam nodes to derive promising orientations from previous plans by techniques like matching with a database [223] or deep learning [77, 80]. Alternatively, the planning problem can be solved multiple times and all inactive beams of the previous run are replaced by newly sampled beams [242].

Given a set of candidate beams, we can setup a matrix D , the dose-coefficient matrix, which maps the activation times \mathbf{x} of the candidate beams to the dose $\mathbf{d} = D\mathbf{x}$ delivered to the voxels of the planning CT. Therefore, each row of D corresponds to one voxel and each column to one candidate beam. We obtain the entries by calculating point doses for every beam by (2.7) and resampling the results to the grid of the planning CT. If we are only interested in the dose to a certain VOI, we simply consider the related subset of rows of D .

2.3.5 Linear-Programming Approach

The CyberKnife provides a stepwise multi-criteria optimization approach based on solving a sequence of linear programs (LPs) [159, 221]. An LP can be solved with guaranteed global optimality by, for example, the simplex algorithm or interior point methods. In the stepwise approach, we do not combine multiple objectives by weighting factors to a single objective function. Instead, we only optimize one objective in each iteration. The result of the current iteration is then transformed into a new constraint. Therefore, it cannot be accidentally deteriorated in any following iteration. Nevertheless, the constraints can be manually relaxed if they do not allow for an acceptable result in any later iteration. Furthermore, there are post-processing routines which allow to maintain a similar plan quality but with fewer beams and nodes in order to reduce the treatment time. Recently, the VOLO algorithm has been added to the CyberKnife’s planning software, especially because the stepwise implementation is not designed for optimization of MLC apertures [126, 237]. VOLO optimizes a weighted sum of multiple quadratic objective functions and additionally provides methods to include optimization of MLC apertures.

Several objective functions are available for the stepwise optimization. They include maximizing the minimum dose in a target VOI, minimizing the maximum dose in an OAR, minimizing the mean dose in an OAR, or minimizing the total number of MU. In this thesis, we focus on optimization of the coverage (OCO)

and employ hard constraints on the other objectives. Thereby, we can more easily compare results for different scenarios, because only one objective varies and the hard constraints prevent unacceptable results for the other aspects.

In order to not only control the dose in the VOIs but also the dose gradient around the PTV, additional virtual VOIs are commonly introduced and called SHELLs [236]. Note that the term shell is sometimes also used in the radiation-therapy context for masks which immobilize the patient during a treatment. We consider two SHELLs around the PTV. As illustrated in Figure 2.8, the SHELLs arise from isotropic dilation of the PTV. By setting upper dose bounds for the SHELLs, we can enforce a certain dose falloff around the PTV. Therefore, the dose bound should become lower from inner SHELLs to outer SHELLs. The basic LP structure which we consider for OCO can be described as

$$\min_{\mathbf{x}, \mathbf{s}} \mathbf{1}^\top \mathbf{s} \tag{2.8}$$

$$\text{s.t. } \mathbf{d}_{\text{PTV}} + \mathbf{s} \geq d_{\text{ps}} \tag{2.9}$$

$$\mathbf{d}_{\text{PTV}} \geq d_{\text{min}} \tag{2.10}$$

$$\mathbf{d}_{\text{PTV}} \leq d_{\text{max}} \tag{2.11}$$

$$\mathbf{d}_{\text{OAR}_i} \leq d_{\text{OAR}_i} \tag{2.12}$$

$$\mathbf{d}_{\text{SHELL}_j} \leq d_{\text{SHELL}_j} \tag{2.13}$$

$$\mathbf{x} \leq x_{\text{beam}} \tag{2.14}$$

$$\mathcal{N}\mathbf{x} \leq x_{\text{node}} \tag{2.15}$$

$$\mathbf{1}^\top \mathbf{x} \leq x_{\text{total}} \tag{2.16}$$

$$\mathbf{x} \geq 0 \tag{2.17}$$

$$\mathbf{s} \geq 0 \tag{2.18}$$

with the vector of non-negative activation times \mathbf{x} . Note that we imply element-wise comparison if we have a vector on the left-hand side but a scalar value on the right-hand side of a constraint. A dose vector \mathbf{d} containing the dose of each voxel can be calculated using the dose-coefficient matrix D as $\mathbf{d} = D\mathbf{x}$. If only the voxels of a certain VOI, e.g., \mathbf{d}_{PTV} for the PTV, are needed, this can be achieved by only using the corresponding rows of D . The coverage is modeled by (2.9) in combination with the objective function (2.8). If the dose in a PTV voxel is below d_{ps} , then the corresponding entry of the non-negative auxiliary vector \mathbf{s} has to have a sufficiently high positive value to not violate the constraint (2.9). By minimizing the sum of the entries in \mathbf{s} , we enforce that the doses are pushed towards d_{ps} and in the ideal case we would achieve a dose of at least d_{ps} in all PTV voxels. Further, all voxels of the PTV get a lower dose limit d_{min} and an upper dose limit d_{max} in (2.10) and (2.11), respectively. For the voxels of the i -th OAR, there is an upper limit of d_{OAR_i} set in (2.12). Analogously, upper limits for the SHELLs are set in (2.13). Also the activation times themselves are restricted. An individual beam may not exceed the activation time x_{beam} in (2.14), all beams of a node, which are summed by a matrix \mathcal{N} having one row per node, may not

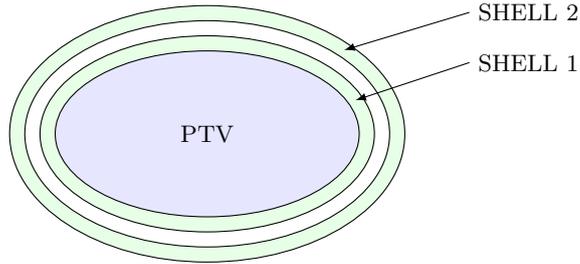


Figure 2.8: Illustration of two SHELLs around a PTV. These virtual VOIs arise from dilation of the PTV.

exceed x_{node} in (2.15), and the overall activation time may not exceed x_{total} in (2.16).

Note that the LP with objective function (2.8) does not directly optimize the coverage. Therefore, an improvement of the mathematical objective values, which we call OCO objective value (OCO-OV), does not necessarily lead to an improvement of the clinical goal. The coverage could be directly maximized by introducing a binary variable for each PTV voxel, which is one if and only if (iff) that voxel receives at least d_{ps} , and adding some modified constraints to relate the binary variables to the corresponding entries of \mathbf{d}_{PTV} . However, such an approach would lead to a combinatorial mixed-integer linear program (MIP) and would be much harder to solve than an LP. Our modified formulation with (2.8) and (2.9) follows the common approximation of the l_0 -norm (number of entries with a zero) by the l_1 -norm (sum of the absolute values), which does not guarantee but often provides a sparse solution for \mathbf{s} which we desire for a high coverage. Furthermore, LPs for radiosurgery planning are typically degenerated in the sense that there are many very different solutions which all have similar objective values. Thus, the optimizer often finds a very good solution fast but needs a lot of time to eventually find the actual global optimum.

3 Image Guidance in Radiation Therapy

The dose delivery in radiation therapy is planned on a CT scan of the patient, which has been acquired beforehand. This allows to design highly conformal dose distributions, which accurately cover the target but spare the OARs. However, the patient has to be positioned carefully before the delivery of each fraction. The planning process assumes a certain spatial relation between the coordinate frame of the treatment machine, the planning CT, and the patient's actual anatomy. Therefore, the patient has to be positioned accurately and remaining deviations from the setup assumed during planning have to be detected and accounted for in the treatment plan. Deviations can also arise from different conditions of the patient, for example a different filling of the bladder which moves and deforms also neighboring structures [36]. We refer to daily anatomical changes w.r.t. the planning CT as inter-fractional motion. The anatomical changes observed in daily pre-treatment imaging can be used to generate a new plan [142, 201]. Remaining positioning errors can partly be compensated by margins which are added to the CTV (compare Section 2.3.1). The margins result in a larger area which receives a high treatment dose. Thereby, the tumor is adequately irradiated as long as the positioning errors actually remain smaller than the margins. The severe drawback of larger margins is, however, that also more surrounding healthy tissue is irradiated. For this reason, it is a very important issue in radiation therapy and radiosurgery to minimize remaining uncertainty and thereby minimize the required margins.

Furthermore, there will be some inevitable motion by the patient during the actual treatment, i.e., intra-fractional motion, which changes the anatomical geometry compared to the planning CT. This motion includes not only predictable movements caused by, for example, respiration [46] but also irregular spontaneous movements. For example, prostate motion shows properties of a random walk [10] and the likelihood of larger motion increases with increasing treatment duration [197, 263]. During longer treatments, patients might even fall asleep resulting in involuntary and uncontrolled motion [34]. All kinds of movements, or more generally unexpected events, during the treatment need to be detected reliably in real time using appropriate imaging modalities [19]. Ensuring a precise dose delivery is especially important during hypofractionated treatments. The CyberKnife system is able to compensate for small motion of the target by adjusting the pose of the LINAC. Larger motion of the target, however, results in an emergency stop and a repositioning of the patient. Note that also relative motion between the target and an OAR might occur and could be additionally considered during planning [220]. As for inter-fractional motion management, the remaining uncer-

tainty can be accounted for by introducing further margins. Systems without a robot-mounted LINAC or a gimballed LINAC, as implemented in the Vero system (Brainlab, Germany, and Mitsubishi Heavy Industries, Japan) [44, 254], can use a motorized treatment couch to correct for motion [41, 143]. Another approach is an online adaption of MLC leaves to compensate for shifts or shape variations of the target [122, 131]. Nevertheless, remaining uncertainties still require adding further margins to the VOIs to ensure sufficient dose delivery to the target.

Another approach to reduce margins is to directly include expected motion into planning. For example, instead of a single plan, we can create a library of plans considering different realistic states of critical organs. Before each fraction, we select the plan which fits best to the current anatomical state of the patient [4, 14, 38]. In robust optimization, we consider probability distributions for the entries of the dose-coefficient matrix [37], which effectively models uncertainty in the dose being actually delivered by a beam to each voxel of the planning CT. Thereby, the objective function of the treatment plan becomes a function of random variables and optimization is employed, for example, in terms of the average objective value or the worst-case objective value [266]. Such approaches can also be extended to explicitly reflect the multi-criteria nature of planning [35]. A recent study considers online re-planning within a fraction to deliver the remaining dose correctly after a detected motion event without interruption [130].

In this chapter, we first discuss approaches which reduce motion of the patient by fixation. Subsequently, we review the most common imaging modalities for IGRT, which are X-ray-based imaging, optical systems, EM tracking, and MRI. Finally, we provide an extended review and discussion of ultrasound imaging for inter-fractional and intra-fractional guidance.

3.1 Immobilization of the Patient

As motion needs to be reliably detected and compensated in radiation therapy, an intuitive step is to firstly minimize the likelihood and amplitude of potential motion by immobilization of the patient. In case of cranial radiosurgery, the classical approach is to attach a stereotactic frame to the patient's head [158, 167, 169]. The frame is also fixed to parts of the treatment room or device, for example the couch, and does not only limit the motion of the patient but also defines a rigidly attached reference coordinate frame. Already the planning CT is acquired with the frame attached. Therefore, the spatial relation between the internal target and the frame can be measured in the CT scan. Later on, the pose of the target can be determined from this mapping in combination with external localization of the frame. Stereotactic frames are invasively attached to the skull and may be supported by other fixations like a bite block, ear plugs, or a head rest. Such fixations also allow to develop relocatable and less invasive frames [84]. An obvious disadvantage of frames is the resulting patient discomfort and the treatment might require at least local anesthesia. The main advantage is that

only very small remaining motion w.r.t. the frame is possible. This implies that the spatial relation between the coordinate frame and the internal target is very static.

Another approach for immobilization during cranial treatments are face masks. These are also employed for treatments with the CyberKnife. They can cover only the face or even cover the upper torso as well. To further improve patient comfort, there are also open masks which do not cover eyes, nose, and mouth [161]. The masks are often thermoplastic, i.e., made of a plastic mesh which becomes soft by heating to precisely fit the head geometry and becomes hard again after cooling. Similar to frames, masks can be combined with a special head rest or a bite block which further fixate the head. Furthermore, masks allow to attach markers for localization. In general, the remaining motion within masks is larger compared to invasive frames and requires correction by image guidance to achieve very small errors [74, 173, 207]. However, they provide higher patient comfort than frames.

For non-cranial targets, there is no direct equivalent for patient immobilization. Nevertheless, there are tools to minimize patient motion on the treatment couch and to ease repositioning. Also vests similar to face masks exist. However, targets in the abdomen move also due to respiration, for example, and external fixation does not provide a reference coordinate frame which is rigid w.r.t. the internal target. Furthermore, aspects like the filling of the bladder can cause deformations and motion and need to be considered [259]. Implantation of fiducials in, for example, the prostate can be used to establish a coordinate frame directly at the target and is common for CyberKnife treatments [102]. However, these fiducials are not externally visible as markers on a frame or mask are.

3.2 X-ray Imaging

X-ray imaging is one of the two standard imaging modalities during CyberKnife treatments. It is used to localize the target structure during patient positioning and treatment. During treatment, the information about the current pose of the target can be used to adjust the directions of the treatment beams. However, the ionizing nature of X-rays further increases the exposure of the patient [183]. Therefore, the applicability of this modality for continuous imaging is limited.

3.2.1 Basic Principles

X-ray imaging is based on the same physical principles as photon-beam radiation therapy. For medical imaging, however, the employed energy is only in the order of 25 keV to 120 keV. Low energies are used especially in mammography while high energies are necessary, for example, for imaging of the lungs. The energy has to be sufficiently high that a reasonable amount of the photons are able to pass through the whole body. These photons can then be measured on a detector, for example with an analog photo-sensitive film or with a two-dimensional sensor grid. The latter digital approach is typically employed today. However, analog

films are still frequently used for, for example, calibration and quality assurance of the radiation-therapy treatment machines [177]. X-ray-based images are typically output as negative images, i.e., darker image intensities correspond to more measured photons. Therefore, strongly absorbing structures like bones appear bright in an X-ray image while air appears dark. Depending on the target, X-ray imaging might provide an insufficient soft-tissue contrast. For targets like, for example, the prostate, gold fiducials can be implanted [102]. Due to their strong attenuation of X-rays, these fiducials can be easily detected in the images and registered to the planning CT in order to determine the change of the target's pose or, to a certain degree, its deformation.

While classical X-ray imaging only provides two-dimensional projections of the body, CT allows to reconstruct volumetric images of a patient. For this purpose, X-ray images are acquired sequentially from many different directions by rotating the source, commonly around the SI-axis. The individual images are one-dimensional projections but their combination allows to reconstruct a cross-sectional slice through the patient. Sequential acquisition of parallel slices at different positions provides volumetric image data. CT data provides quantitative values about the local attenuation properties of the scanned materials, from which the electron-density distribution in the tissue structures can be derived. The electron density, in turn, is the property which is required to calculate the dose deposition in tissue for photon beams. For this reason, CT is the most important imaging modality for radiation therapy treatment planning even though its soft-tissue contrast is rather limited. The high dose exposure, however, makes CT imaging unsuitable for continuous online monitoring during a treatment.

3.2.2 Stereoscopic X-ray Imaging

Similar to stereovision with conventional camera systems, stereoscopic X-ray imaging allows to reconstruct three-dimensional scenes from two projective 2D images of the same object from different directions. Especially if fiducials have been implanted, the pose of the marker formed by the fiducials can be reconstructed. See Figure 3.1 for an illustration of an orthogonal setup. In order to compare stereoscopic images with a planning CT, a common approach is to use digitally reconstructed radiographs (DRRs). These are virtual 2D X-ray images derived from a 3D CT image by ray tracing, i.e., they simulate projective X-ray imaging from some direction. During patient positioning, DRRs can be compared to the actual stereoscopic images to align the patient. For gantry-based treatment systems, it is common to attach an X-ray imaging source to the gantry, referred to as on-board imaging. Recently, real-time 6D tracking of abdominal targets by projective imaging of implanted gold fiducials has been proposed and demonstrated [186, 258]. By repeated image acquisition during VMAT rotations of the gantry, accurate reconstruction of the marker pose is feasible without explicit stereoscopy.

The CyberKnife system employs two orthogonal X-ray sources mounted to the ceiling and a detector, which is located on the floor below the treatment couch.

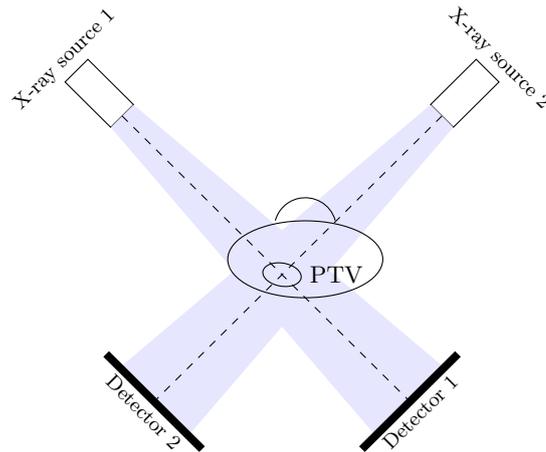


Figure 3.1: Cross-sectional sketch of orthogonal 2D X-ray imaging with the isocenter inside the PTV.

The setup is aligned to the reference point defined in the planning CT by the planner. The X-ray imaging can be used both for positioning and image guidance during treatment. Due to the additional exposure of the patient to ionizing radiation, the CyberKnife acquires images during treatment only every 30 s to 60 s [126]. This interval-based approach might therefore detect movements only with a certain delay. Furthermore, it might completely miss motion events which exceed the tolerable limits only in the time frame between two sequent acquisitions. In addition, a temporal synchronization between X-ray imaging and motion of the robot-mounted LINAC is necessary, because otherwise the robot arm can block the line-of-sight and thereby the imaging intervals become too large.

Conventionally, fiducials are implanted for tracking. However, the CyberKnife offers the fiducial-free systems Xsight lung [70, 278] and Xsight spine [100, 182]. The Xsight tracking system for lungs uses directly the soft-tissue contrast and registers the stereoscopic images to DRRs to derive 6D motion information. This is only feasible if the tumor has a certain extent. For spine tracking, static bony structures are localized in the images and compared to data from the planning CT. Both systems are also used for patient positioning. For SRS, the CyberKnife provides 6D tracking of the skull based on skeletal features [71, 118].

3.2.3 Cone Beam CT

Classical CT is based in its simplest implementation on the sequential acquisition of one-dimensional measurements with line detectors and is also called fan beam computed tomography (FBCT). In the advanced multi slice computed tomography (MSCT), multiple adjacent lines are acquired simultaneously. Another alternative is cone beam computed tomography (CBCT), in which the X-ray beams have a three-dimensional conical shape to directly cover a whole volume instead

of only a single or a few slices. For detection, two-dimensional flat panel detectors are used. On the one hand, CBCT's general image quality is lower, estimated electron densities are less accurate and precise, and the acquisition time is substantially longer making it more susceptible for motion artifacts. Additionally, the complexity of the image reconstruction is higher. On the other hand, it provides a higher spatial sensitivity and a smaller dose exposure. Furthermore, its form factor is substantially smaller making it suitable for C-arm-based imaging or integration in other systems.

In IGRT, using CBCT is interesting because it can be directly integrated into the treatment room and machines. Thereby, no transfer of the patient is necessary for volumetric imaging, but the setup is equal to the treatment setup. Furthermore, the image data can be acquired immediately before and even during treatment. Therefore, CBCT is very interesting for patient positioning and verification. CBCT could be used to directly acquire the planning CT in the treatment setup [281]. Due to the limited image quality and inaccurate measurement of electron densities, however, this is not common practice and applying correction algorithms is already necessary for adaptive planning [261, 277].

3.2.4 Portal Imaging

The fact that photon-beam radiation therapy and X-ray imaging employ the same physical effects has been exploited by using the treatment machine also for imaging. This results in imaging with several megaelectronvolts, which is one to two magnitudes higher than in conventional diagnostic X-ray imaging. Therefore, the main interaction process of the X-rays in tissue is Compton scattering instead of photoelectric interactions as in conventional kiloelectronvolt imaging. The resulting images have generally a substantially worse contrast. Nevertheless, portal imaging is even used for CBCT acquisition [181, 215]. While the CyberKnife does not use portal imaging, it is used in other systems for in-vivo dosimetry, verification of patient positioning, or intra-fractional motion estimation [9, 178, 262].

3.3 Optical Tracking

A non-ionizing and rather low-cost approach for motion estimation are optical tracking systems. However, they only provide surface information and thereby are not able to directly track internal targets which are not rigidly related to the surface of the patient. Nevertheless, they are an important tool in radiation-therapy systems for inter- and intra-fractional pose estimation and also the CyberKnife uses optical tracking, especially for tracking of respiration. In contrast to gating techniques [85], active compensation of respiratory motion allows for continuous irradiation, even without requiring breath-holding by the patient.

3.3.1 Basic Concepts

Optical tracking typically refers to systems which employ visible or infrared light. Many commercial tracking systems use dedicated infrared markers which are localized in stereoscopic images to estimate the position and orientation of a target. For this purpose, markers have to be rigidly attached closely to the actual structure of interest. Marker geometry and placement can have a substantial influence on the tracking errors [272]. A further distinction is made between active and passive marker-based systems. Passive systems measure the reflection of light. In contrast, active systems employ light emitting markers, for example a set of LEDs. An advantage of active markers is that they can be turned on sequentially to directly provide a unique identification in image sequences. Furthermore, they can create a strong and easy-to-detect signal. Disadvantages are the required power supply and electronics for control, often including cable connections.

Common devices for markerless pose estimation are depth-sensing cameras. The two major variants are time-of-flight (ToF) cameras and cameras employing structured light. The former emits light, typically infrared, and measures the time until the reflection is received. From this time difference, the spatial distance to the reflecting object can be derived. The whole three-dimensional scene is illuminated simultaneously and a two-dimensional sensor matrix allows to reconstruct a point cloud representing the surfaces in the scene. Therefore, tracking algorithms can directly exploit height variations on a target surface to match the measured point cloud with previous scans or a model and to estimate the pose. In structured light, light is emitted which forms a known geometric pattern, for example a grid. If the pattern is reflected on an object with a non-flat surface, we receive on the detector a deformed version of the pattern. From the deformation, we can recover the shape of the reflecting surface. Structured light can employ infrared light in order to minimize interference with other visible-light optical devices like RGB cameras.

A common issue of optical tracking systems is that they require a steady line-of-sight. Redundant markers and multiple cameras viewing from different directions can reduce this effect. However, this is still difficult in scenarios like surgical tracking with limited space, moving obstacles, and frequently changing environments. The advantages of optical tracking compared to other imaging modalities are the rather cheap and light-weight systems, no harmful radiation, no electromagnetic fields interfering with other devices or materials, fast acquisition and processing, and the common integration into small medical instruments.

3.3.2 Surface Imaging for Patient Setup

A common application for optical systems in radiation therapy is patient setup. The most basic approach are visible-light laser beams or patterns in the treatment room, which guide drawing marks on the patient's skin or on an immobilization device. Due to the markings and the fixed laser setup, it is possible to reproduce the

patient setup in subsequent fractions. However, also advanced systems exist which acquire a surface scan of the patient and exploit this 3D map for more automated positioning [101]. The wide use of such systems even led to the term of surface-guided radiation therapy (SGRT) for surface-scanning-supported treatments. The AlignRT system (Vision RT, UK) projects visible-light patterns on the patient and derives the surface from the deformed pattern using three cameras. It outputs the 6D deviation from a reference surface to enable correction of the patient setup. A similar approach implements the IDENTIFY system (Varian Medical Systems, USA). The Catalyst system (C-RAD, Sweden) employs near-invisible-violet light to project its patterns. While all these systems are mostly used to guide patient setup, they also provide intra-fractional monitoring, for example to guide gating. For this application, however, also specialized marker-based systems exist, as outlined in the next sub-section.

3.3.3 Surrogate Tracking for Internal Respiratory Motion

Respiration results in clear superficial motion of the patient. Therefore, optical tracking systems are an intuitive choice to observe respiration-induced motion in IGRT, because they are fast, harmless, cheap, and easy to integrate. When attaching markers to the patient, the external respiratory motion is measurable. However, the internal target and OARs are not rigidly connected to the patient surface. Hence, optical approaches only provide indirect measurements for treatment sites severely affected by respiratory motion like lung, liver, or kidneys. For this reason, they are typically combined with a cross-sectional imaging modality, like X-ray imaging, to estimate and verify a correct mapping from externally observed motion to the internal motion of a VOI.

The CyberKnife provides the Synchrony respiration tracking system, which employs three ceiling-mounted infrared tracking cameras. These continuously estimate the pose defined by three active markers attached to the patient [126]. Thereby, no breath-holding or gating is necessary, but the robot-mounted LINAC synchronously adapts the beam direction to the respiration phase, as illustrated in Figure 3.2. Synchrony establishes a correlation-based model to infer internal motion from the externally observed motion and verifies and updates this model periodically with the information obtained from stereoscopic X-ray [55, 241]. X-ray imaging involves either implanted fiducials or the fiducial-free Xsight technology described in the previous section. Due to the substantial payload on the robot arm, inertia is an issue when dynamic motion is required for compensation and causes a latency of about 115 ms [126]. For this reason, the implemented motion estimation includes prediction [54]. Including further tiny sensors in the setup has been proposed to further improve the predictions [52]. However, there might occur unexpected motion which completely invalidates the correlation model and requires to interrupt treatment. These events also need to be detected reliably [5].

Other commercial implementations include the Real-time Position Management (Varian Medical Systems, USA) which also employs infrared markers to track res-

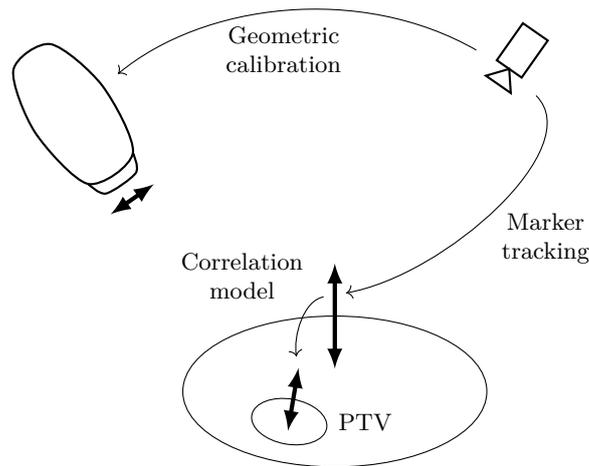


Figure 3.2: A correlation model relates the tracked motion of external markers to the internal motion of the target. The geometrical calibration between tracking system and robotic LINAC allows to adjust the beam delivery accordingly.

piratory motion, in this case for gating. A markerless alternative is ExacTrac (Brainlab, Germany) which combines optical surface imaging since lately with thermal imaging to enhance accuracy [28]. Similar to Synchrony, ExacTrac complements the optical tracking with stereoscopic X-ray imaging to verify and update its patient model. It is designed for use in both patient setup and intra-fractional tracking.

3.3.4 Head Tracking

The task of tracking a patient’s head motion during an intervention does not only appear in cranial radiation therapy, but also in other fields like transcranial magnetic stimulation (TMS) [88, 212] or artifact-free brain imaging [69, 140]. In radiation therapy with stereotactic frames, markers can be attached to the frame for tracking by an optical tracking system. Attachment of markers to masks or bite blocks [185] are alternative approaches in case of frame-less treatments, which improve patient comfort as discussed in Section 3.1. In any case, marker-based approaches have the drawback of the critical placement procedure and limited rigidity w.r.t. the actual target. An alternative is markerless surface imaging [34, 39, 160]. However, extracting motion in surface imaging requires to perform image registration and the accuracy of rigid registration suffers from any occurring non-rigidity [90]. Non-rigid registration, in turn, has higher requirements on exploitable features in the data. For this reason, different techniques have been proposed to enhance surface data with additional information, like thermal imaging [28] or local tissue thickness estimation with infrared light [275].

3.4 Electromagnetic Tracking

EM tracking is a technique which effectively localizes small markers by sensing variations of electromagnetic fields. It is widely employed in medical interventions [68]. In contrast to optical tracking, it does not require a line-of-sight and allows to track internal targets. For the latter, however, it is necessary to place the markers invasively in the target.

There are mainly three different approaches for EM tracking systems [20]. The first approach utilizes alternating magnetic fields which are externally generated and measured by small coils attached to the target. The received signals are read and evaluated for localization of the coils w.r.t. to the field. This typically requires a wired connection between the coils and the computer system which processes the signals. The second approach employs fluxgate sensors which are able to measure static magnetic fields. The remaining setup is essentially equal to that of systems using alternating fields.

The third approach uses passive transponders which are attached to the target. If these transponders are excited by a magnetic field, they emit a radio signal. This signal is received by antennas and the position of the transponder is derived from signal properties like the delay or the intensity. If each employed transponder requires excitation with a different frequency and they form a known geometry, they can be localized sequentially and the pose of their geometry can be estimated w.r.t. the magnetic field.

A drawback of EM tracking is its sensitivity to distortions of the fields caused by other devices or materials [279], which compromise the localization procedure. Furthermore, there has to be sufficient space for appropriate integration of the required hardware, which might include cable connections. Lastly, the accuracy of EM tracking systems is lower than the accuracy of optical tracking systems. However, EM tracking systems allow for integration into small instruments, like catheters or endoscopes, and tracking of those instruments inside the patient without requiring a line-of-sight.

The transponder-based approach allows to implant wireless transponders in organs and to track their motion. This procedure is similar to the implantation of fiducials to enable X-ray-based localization but suitable for continuous monitoring. In radiation therapy, this is implemented by the Calypso system (Varian Medical Systems, USA) for prostate [139] and lung [110, 248] tracking. The system employs three cylindrical transponders with diameters of 1.4 mm or 2.1 mm which are suitable for permanent implantation. The EM tracking system itself is localized within the treatment room by an optical tracking system. For prostate tracking, accuracies similar to those achieved by orthogonal X-ray imaging have been reported [267]. It has also been used to guide MLC tracking [25, 122]. Another clinical system for prostate tracking is RayPilot (Micropos Medical, Sweden) which uses wired transponders which are retracted after the treatment. The CyberKnife does not natively provide EM tracking.

3.5 Magnetic Resonance Imaging

3.5.1 Basic Principles

In MRI, we use a physical effect called nuclear magnetic resonance (NMR) to generate high-resolution volumetric images of a patient. The images especially provide a very good contrast between different soft tissues. For this reason, MRI is often used during treatment planning for delineation of the target and the OARs and registered to the planning CT. It only uses magnetic fields which are considered to be safe for humans as long as they are operated reasonably. Nevertheless, there are remaining risks [43]. The basic component of an MRI scanner is a strong static magnetic field having a flux density in the order of 0.5 T to 7 T. Therefore, MRI cannot be used for patients with artificial cardiac pacemakers or metallic pieces or implants. Furthermore, it requires a shielded room which only contains NMR-compatible materials and devices. The strong magnetic field is typically realized with a magnet cooled by liquid helium to establish superconductivity. This cooling procedure is expensive and time-consuming. For this reason, MRI systems are kept active even during standby phases.

The imaging physics and practical instrumentation in MRI are complex and beyond the scope of this section. Details can be found in Brown et al. [29], for example. However, the basic idea is that the strong static magnetic field aligns hydrogen nuclei inside the human body. This results in a small but measurable magnetization. The orientation of this magnetization can be manipulated by applying radio frequency (RF) pulses with a suitable frequency, the Larmor frequency, which depends on the magnetic flux density of the static field. After the RF pulse is turned off, relaxation processes start, which align the magnetization with the static field again and can be measured with external coils. The time required for relaxation is tissue-dependent and provides the contrast in the images when spatially resolved. Many different sequences of RF pulses and measurement protocols exist, which result in different contrasts in the images and allow to suppress signals from specific structures. Furthermore, various contrast agents and boluses exist.

The static magnetic field is typically aligned with the patient's SI-axis and, comparable to CT, we either acquire a sequence of axial slices or a whole volume at once. Spatial encoding of the signals received in the coils can be realized by three small gradient fields which slightly modify the frequency and the phase in the received signal depending on the 3D location of the source. This encoding allows to reconstruct three-dimensional images from the signals recorded by the receive coils.

3.5.2 Integration in the Treatment Room

MRI provides harmless volumetric imaging with a large FOV and very good soft-tissue contrast. However, image acquisition is rather slow for conventional diagnostic scanning. Therefore, interventional MRI employs advanced strategies like

sparser sampling, restriction to 2D imaging, or parallel acquisition techniques. The term cine MRI arose in the context of cardiovascular imaging for fast imaging protocols allowing to acquire multiple images during a single cardiac cycle.

In contrast to the imaging modalities discussed in the previous sections, MRI cannot be simply added to an existing radiation-therapy setup. Due to the strong magnetic field, MRI can only be employed in specially shielded rooms which, on the one hand, prevent any interference of the fields with the outside environment and, on the other hand, only contain material which does not interact with or distort the fields. Therefore, integration of MRI in a treatment room is difficult and requires a dedicated and sophisticated design of all involved components [205]. In addition to the implications on hardware design, X-ray beams are affected by the magnetic fields and this has to be considered in dose calculations and planning [176, 190]. For these reasons, it seems rather unlikely that already existing systems, like the CyberKnife, will receive an MRI-guidance upgrade in near future.

Only a few clinical MRI-integrated systems exist up to now. The MRIdian system (Viewray, USA) has been in clinical use for some years now, but only provides 0.35 T imaging [62]. The Unity system (Elekta, Sweden) is the first system providing imaging with 1.5 T [204]. It has been shown that MLC tracking is feasible with this system [87] and also novel approaches for online re-planning based on the detected motion have been proposed [129, 130]. In addition to the two mentioned systems, there are further systems under development [57, 121]. Beside others, Hunt et al. [106] and Paganelli et al. [196] provide reviews on the state-of-the-art MRI-based image guidance in radiation therapy.

3.6 Ultrasound Imaging

Every imaging modality discussed for IGRT in the previous sections has some drawbacks. X-ray imaging, and especially CT, increases the dose exposure and has a low soft-tissue contrast which often requires invasive placement of fiducials in the target. Optical tracking systems only allow for indirect tracking of internal targets based on externally visible motion, typically induced by respiration. Furthermore, it requires periodic verification by a cross-sectional imaging modality. EM tracking is sensitive to distortions caused by other devices and requires invasive markers. Lastly, the very good soft-tissue contrast of MRI comes with high costs and a difficult integration. Therefore, another interesting modality is ultrasound imaging. While it clearly also has several drawbacks, it provides some very interesting aspects for image guidance in radiation therapy. Several recent review articles on this topic are available [32, 64, 191, 273].

Ultrasound imaging is established in prostate radiation therapy for patient setup and verification [64]. It does not use ionizing radiation and is considered to be safe as long as operated reasonably, similar to MRI. Instead of radiation, it transmits acoustic waves into the body and measures the delay of echoes to estimate the location of reflecting layers. Thereby, it allows for cross-sectional imaging. Wave

transmission requires direct contact of a so-called transducer or probe, which is conventionally hand-held, with the patient surface. The penetration and imaging depth is severely limited in contrast to X-ray imaging or MRI. Furthermore, structures like bones or air enclosures further limit the effective imaging depth and can produce artifacts. Nevertheless, ultrasound imaging allows to visualize internal structures in real-time as one-, two-, or three-dimensional images, referred to as A-scans, B-scans, and C-scans, respectively. The systems are rather cheap compared to CT or MRI and do not require a specific environment or shielding. Recently, there has been an increasing interest in employing ultrasound for intra-fractional guidance [191, 273]. This has been encouraged by recent advances in the real-time capability of volumetric imaging with reasonable resolution and FOV size [108]. However, intra-fractional application requires to find practical solutions for several issues. The conventional manual operation of the ultrasound transducer is not feasible in the treatment room during a treatment due to the radiation. Further issues include the blocking of beam directions by the transducer and its mounting and the limited reproducibility of a transducer's position and its pressure onto the abdomen.

In this section, we first describe the basic principles of ultrasound imaging. Subsequently, we review the state-of-the-art application of ultrasound imaging for patient setup and inter-fractional guidance. Lastly, we discuss recent approaches to intra-fractional guidance and the remaining difficulties and open questions. In particular, we review a recent approach employing a general-purpose robot arm with redundant kinematics for holding the transducer. For this specific setup, we will analyze and optimize different aspects in the following chapters.

3.6.1 Basic Principles

In ultrasound imaging, we send acoustic waves into the patient's body and measure the time until we receive echoes due to reflections. The relevant tissue property is the tissue-dependent acoustic impedance and reflections mainly occur where the impedance abruptly changes. The frequency of the waves is typically in the order of 1 MHz to 20 MHz, but depends on the specific application. The higher the frequency, the higher the spatial resolution. However, waves with higher frequencies experience a stronger attenuation in tissue and therefore ultrasound imaging with lower frequencies provides signals from deeper structures. The imaging depth can reach from a few centimeters to beyond 20 cm and the spatial resolution is typically in the order of 0.1 mm to 3 mm depending on the parameters of the system's components [257]. The axial resolution, i.e., in depth direction, is often better than the lateral resolution.

Generation of the ultrasound waves is based on the piezo-electric effect. An ultrasound transducer, also called probe, contains piezo crystals which mechanically oscillate in resonance when supplied with a proper alternating voltage. If we have a direct contact between the transducer and the surface of a patient, the oscillation creates an acoustic wave which propagates through the tissue. The

oscillation is only excited for a small time interval and thereby we obtain a short pulsed wave. After excitation, the same piezo crystals can be used to sense echoes by employing the reciprocal effect until we excite them again to generate the next pulse. This repetition time determines from how deep we can receive echoes, because the echo delay is directly related to the traveled distance. In the simplest case, the temporal sequence of echoes from each pulse corresponds to one one-dimensional depth scan. This so-called A-scan contains the amplitudes of the received echoes depending on the time which has passed since excitation. In practice, the echo time is converted to a distance by assuming a homogeneous medium with a speed of sound of about 1540 m s^{-1} in case of soft-tissue imaging. This simplifying assumption is clearly never true in reality and this fact is important when interpreting the images because they do not provide correct geometric distances [21, 216]. Therefore, correction algorithms considering the actual local speed of sound should be applied for quantitative measurements [65, 66]. Furthermore, the attenuation of a wave increases with increasing depth which requires to apply an echo-time-dependent intensity compensation in the images to obtain reasonable and comparable contrast over the whole depth range. Contrast agents can help to image structures not having a good natural contrast in ultrasound images.

Several transducer designs exist which allow for two- and three-dimensional imaging. The most intuitive approach are flat one- and two-dimensional arrays of piezo crystals. The crystals are sequentially excited and thereby record depth scans at different locations. Other transducers use so-called beam forming. By exciting multiple crystals simultaneously but with small delays, we can define an arbitrary orientation for the resulting superimposed wave. Given only a 2D imaging transducer, mechanical sweeping or wobbling can alter the direction of the 2D plane to assemble 3D images. Alternatively, manual sweeping can be done to create 3D images with a 2D transducer. More detailed information about transducer designs and the principles of ultrasound imaging in general can be found, for example, in Szabo [257].

Other diagnostic modes of ultrasound include Doppler imaging to estimate flows and elastography to derive mechanical tissue properties. In photoacoustic imaging, we record ultrasonic emissions in tissue caused by excitation with a laser pulse. Furthermore, integration of ultrasound imaging into catheters for intra-vascular imaging is an established technique. Similar to MRI, ultrasound imaging is considered as safe when operated reasonable. Risks arise in longer procedures especially from thermal effects caused by the attenuated sound waves. Ultrasound is also employed as a therapeutic tool in different variants. They include high-intensity focused ultrasound (HIFU), which allows for insertion-free thermal ablation, and low-intensity stimulation, which can accelerate healing of fractures.

3.6.2 Inter-Fractional Positioning of the Patient

Ultrasound imaging is used in the context of radiation therapy mostly for inter-fractional imaging of the prostate. Treatment sites like the brain or lungs are

hardly accessible with ultrasound due to their bony and air-filled environment, respectively. Nevertheless, lung surface tracking has been demonstrated [214]. Although studies have investigated ultrasound for various other sites, like liver, pancreas, kidneys, gallbladder, bladder, and most notably breast cancer [64, 191], the application is not clinically established and the few commercial systems focus on prostate treatments. For prostate imaging, there are three different approaches to place a transducer such that the prostate will be visible in the images [32].

In transrectal ultrasound (TRUS), a specially designed ultrasound probe is inserted in the rectum for imaging close to the prostate with good image quality. However, this approach is uncomfortable for the patient and not used for positioning or tracking in radiation therapy [32]. Nevertheless, TRUS is suitable for diagnostic imaging and can also be used to guide placement of fiducials in the prostate [114, 145]. In brachytherapy, TRUS can guide needle insertions [11].

The second approach is transperineal ultrasound (TPUS), in which we position the transducer non-invasively onto the perineum, i.e., the area between scrotum and rectum. This requires a specific positioning of the patient in a rather uncomfortable pose. The approach is commercially implemented in the Clarity Autoscan system (Elekta, Sweden) [1, 93, 141]. Its mechanically-swept 2D transducer allows for automatic 3D imaging. The system allows for both inter- and intra-fractional application, where we discuss the latter in the next sub-section.

Transabdominal ultrasound (TAUS) is the third and most flexible approach in the sense that other abdominal treatment sites, like the liver, could be imaged the same way. A transducer is simply placed onto the abdominal wall close to the prostate. Typically, the bladder provides a suitable imaging window if filled appropriately. However, the farther distance to the prostate compared to TRUS or TPUS requires to image at lower frequencies to achieve the necessary imaging depth resulting in lower image quality. The TAUS approach has been commercially implemented for inter-fractional motion detection. Two systems are SonArray (Varian Medical Systems, USA) [200, 219] and BAT, later called BATCAM, (Best Nomos, USA) [144, 148, 217]. Another system is Clarity (Elekta, Sweden) which is an earlier and simpler version of Clarity Autoscan and requires manual sweeping of the 2D transducer to acquire and reconstruct 3D images [162]. Hence, the system requires some experience of the operator to achieve good quality and reproducible imaging [60]. Severe operator and training dependence, however, is a general issue in ultrasound guidance [180].

Localizing the transducer within in the treatment room is essential for registration and patient alignment. For this purpose, optical tracking can be employed. An alternative are passive probe holders, which are fixed in the treatment room and sense the pose of the transducer. This has been used in the BAT system [144]. Also EM tracking of the transducer is an option. Furthermore, a spatial calibration between the estimated transducer pose and the actual ultrasound image is necessary to directly relate points in the image to room coordinates [202, 226]. Acquiring a reference ultrasound image for daily alignment along with the planning CT is commonly considered as more accurate and reliable than employ-

ing inter-modality alignment with contours extracted from the planning CT [40, 180]. The former approach avoids, for example, registration issues related to the geometric distortion in ultrasound images or related to displacements due to the transducer's pressure onto the patient, which are not present in the planning CT.

3.6.3 Intra-Fractional Systems and Setups

It has been demonstrated that real-time volumetric ultrasound imaging is suitable for intra-fractional motion monitoring [239] and even tasks like guiding dynamic MLC tracking [58, 109]. However, it requires to solve some practical problems depending on the specific system setup. The most important issue in intra-fractional ultrasound imaging arises from the fact that we need to ensure continuous contact between the transducer and the patient. Normally, a physician holds the transducer and as a matter of routine adapts its pose to compensate, for example, breathing-induced motion of the patient. During a radiation-therapy treatment, however, manual operation is not feasible due to the ionizing radiation in the treatment room. Therefore, some kind of automatic system is required which ensures a steady contact to the patient's surface. Only an initial setup immediately before the treatment could be done manually. However, even this can be difficult due to the generally strong operator dependence. Robotic ultrasound systems are also relevant in other medical applications, especially in needle guidance and some kinds of surgeries [203].

While the Clarity system requiring manual sweeping clearly disqualifies from intra-fractional application, its TPUS-based motorized version Clarity Autoscan has been evaluated for intra-fractional imaging in various studies [92, 141, 210, 211]. The transperineal transducer placement rarely interferes with treatment beams and only leads to minor pressure-induced artifacts. This system is the only established commercial solution for intra-fractional ultrasound guidance.

TAUS has the advantage over TPUS that it is more generic as several treatment sites could be imaged with the same system. In contrast, the TPUS approach is limited to prostate imaging. Furthermore, TAUS does not require a special pose of the patient for prostate imaging. However, especially when considering the beam directions used in CyberKnife treatments which are mainly from anterior, blocking of treatment beams by the transducer is a bigger issue than for TPUS. The feasibility of radiolucent transducers for 3D imaging has been demonstrated [225], but all commercial transducers contain metal and other materials which interfere with X-ray beams. Bazalova-Carter et al. [12] have proposed Monte-Carlo-based dose models which can account for such interfering transducers during treatment planning. However, the authors recommend to rather avoid intersections of the treatment beams with the transducer, if possible, because these beams become very inefficient, the model requires precise reproduction of the transducer setup during planning and all fractions, and the long-term effect of the radiation on the transducers is not yet clear. It has been demonstrated for liver [224, 285] and prostate treatments [79, 227] with different treatment devices and systems for

ultrasound guidance that the sparing beam directions which would interfere with the transducer can generally still lead to acceptable plans.

Şen et al. [245] developed a collaborative robotic system for intra-fractional ultrasound imaging. The in-house-constructed system integrates force sensing and provides five active DOFs and additionally six passive DOFs. Note that a conventional robot arm which can realize any end-effector pose within its workspace requires six DOFs. Adding more DOFs allows to realize the same pose in different configurations which is interesting when, for example, reaching around obstacles is necessary. The research group has put a major emphasis on the workflow for reproducible transducer placement [246]. After acquiring a reference ultrasound image, they replace the transducer for acquisition of the planning CT with a model which, in contrast to the actual transducer, does not produce artifacts in the CT image. The model is automatically positioned based on the previous pose and reaction forces of the actual transducer [15]. For transducer placement before the daily alignment, a similar protocol is used in which first the operator performs a rough manual alignment and afterwards the robot guides the precise alignment with force feedback [244, 247]. To account for anatomical changes, however, the operator decides the eventual placement. In Şen et al. [246], the authors conclude to replace their specialized robot design with a commercial general-purpose robot arm to ease handling.

Another research group started with tele-robotic 2D imaging [227] and developed an autonomous robotic system for 3D image guidance during treatments of the liver [224]. For the latter, they have put emphasis on investigating not only the blocking of beams by the transducer but also by the robot arm. The in-house-constructed robot arm provides nine DOFs allowing to flexibly handle a wide range of patient geometries.

The approach which we will consider in the subsequent chapters of this thesis has been developed at University of Lübeck [133, 134, 135]. Similar to the previously described approaches, it considers employing a robotic arm and force sensing to realize a continuous contact between the ultrasound transducer and the patient. However, instead of an in-house-built specialized robotic system, it considers a commercially-available general-purpose robot arm with seven DOFs and integrated force sensing. Thereby, the mechatronic part is already mature and does not require major additional engineering. Furthermore, this system is flexible and does not make any assumption on the environment or treatment site. Currently, two commercial light-weight robot arms with seven DOFs are available, namely the LBR iiwa (KUKA Roboter, Germany) and the Panda (Franka Emika, Germany) which we show in Figure 3.3a and Figure 3.3b, respectively. Note that Figure 3.3a shows an LBR iiwa version allowing for up to 14 kg payload. There is a slimmer version limited to 7 kg which is more similar to the shown Panda having a maximum payload of 3 kg. The two robot arms have very similar kinematics and both provide integrated force sensing. Seitz et al. [243] have demonstrated how the LBR iiwa’s force sensing allows to compensate for breathing and other motion and thereby enables continuous contact and ultrasound imaging. While

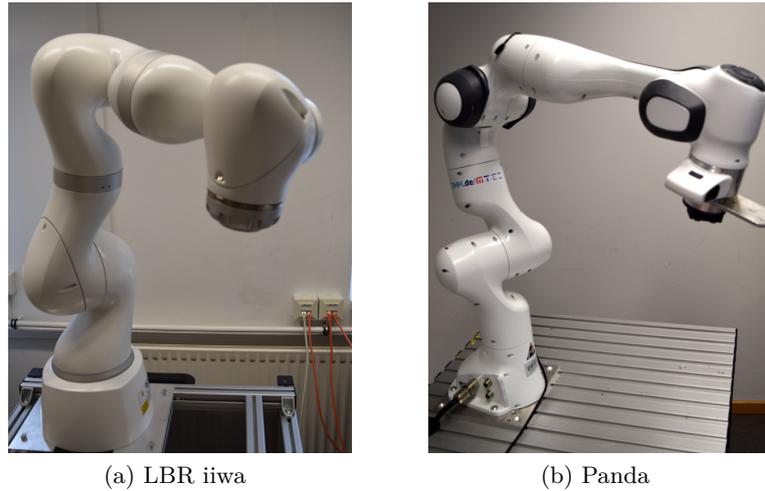


Figure 3.3: A high-payload version of the LBR iiwa (a) and a Panda (b) robot arm with seven DOFs and integrated force sensing.

the Panda is rather intended for application in research, the LBR iiwa is specifically designed for safe human-robot collaboration and also available in specialized versions for operation in clean rooms and for medical applications. Note, that there are further commercial 7-DOF robot arms like, for example, the PA 10-7C (Mitsubishi Heavy Industries, Japan) and the Single-arm YuMi (ABB, Switzerland). However, they are either rather heavy systems, intended for industry and logistics, lack force sensing, or are not designed for collaboration with humans.

Similar to the other robotic-ultrasound systems, the 7th DOF leading to a redundant kinematic increases the dexterity of the robot arm [136] and thereby provides additional flexibility for the ultrasound setup. This is interesting for avoiding the blocking of important beam directions. However, because all DOFs are actively actuated, it is even possible to vary the robot arm's configuration during a treatment without manual intervention. The kinematic design of both the LBR iiwa and the Panda allows to continuously rotate the arm's elbow while maintaining the same end-effector pose, as illustrated in Figure 3.4 for two angles. This allows to further reduce blocking of beam directions, because the robot arm could literally elude treatment beams. Kuhlemann et al. [137] derived for the LBR iiwa an analytic inverse kinematic, i.e., an explicit and deterministic algorithm to compute all joint angles given the desired end-effector pose, which allows to control the angle of the elbow, which we call LIFT, directly with a single additional parameter. This severely simplifies the control of elbow rotations with steady end-effector pose.

Using a kinematically redundant robot arm, however, cannot fully circumvent beam blocking. Both the pose of the ultrasound transducer and the robot's base are fixed and mainly the intermediate links not directly connected to the end-

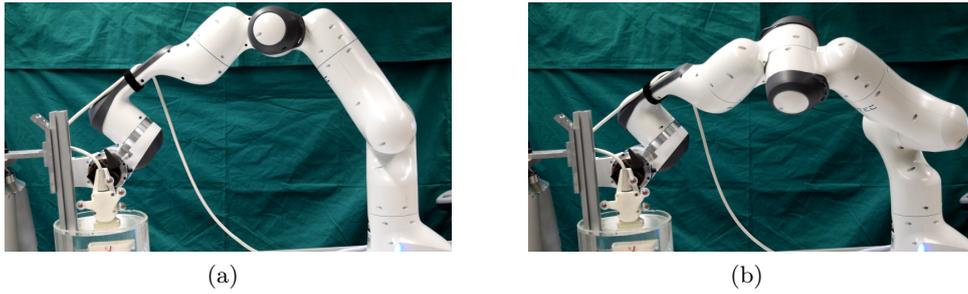


Figure 3.4: A 7-DOF robot arm maintains the same pose of a 2D ultrasound transducer in different configurations by continuous rotation of its elbow.

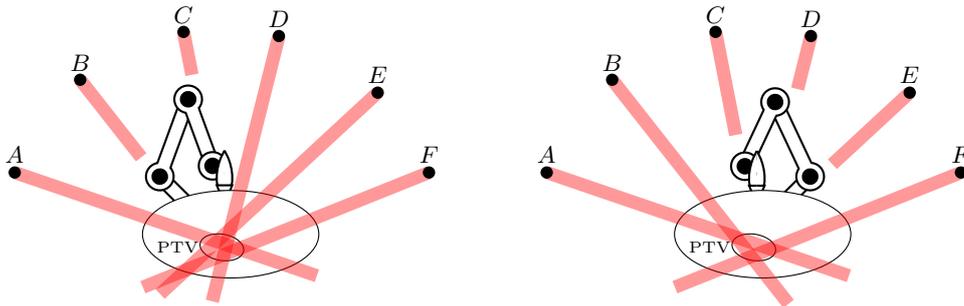


Figure 3.5: Two-dimensional sketch of the advantage when employing a kinematically redundant robot arm. The same pose of the ultrasound transducer can be realized with different configurations (left and right). In this example, beams A and F are feasible in both configurations. Beams D and E are only feasible in the configuration on the left, while beam B is only feasible on the right. Beam C is infeasible in both configurations.

effector or the base move during an elbow rotation. Therefore, some beam direction will remain blocked, as illustrated in Figure 3.5, even if we employ elbow rotations during the treatment. At the same time, other beam directions will not be blocked for any LIFT angle. Gerlach et al. [79] showed that it is possible to position an LBR iiwa with attached 3D transducer such that acceptable prostate plans can be obtained. However, it requires to find suitable positions for the robot base around the patient and to select a suitable LIFT angle. Furthermore, it has been demonstrated that considering multiple LIFT angles during treatment can further improve the achievable plan quality [78].

4 Basic Treatment-Planning Setup for Ultrasound-Guided Radiosurgery

If we consider a kinematically redundant robot arm for image guidance by trans-abdominal ultrasound imaging, as introduced in Section 3.6.3, this has an impact on treatment planning for radiosurgery. Ultrasound imaging requires a steady contact between the ultrasound transducer and the patient. Therefore, the robot arm which holds the transducer has to be positioned close to the patient and it has to place and maintain the transducer directly onto the abdominal wall at a suitable viewport w.r.t. the imaging target. However, conventional robot arms are radiopaque, i.e., we cannot deliver dose through them. Therefore, we can only consider candidate beams for planning, which do not intersect with any part of the robot arm or the ultrasound transducer, which conventionally is radiopaque as well. This can deteriorate the achievable treatment plan quality if the blocked directions cannot be adequately compensated by beams from the remaining feasible directions. The redundant kinematics of our 7-DOF robot arm under consideration provides additional flexibility in comparison to a setup with a conventional 6-DOF robot arm, because an end-effector pose can be realized by infinitely many configurations.

In this chapter, we review and discuss the basic treatment planning aspects which are required for an ultrasound-guidance setup with a kinematically redundant robot arm. We study planning on real cases of prostate cancer. For this purpose, we first introduce the patient data and planning parameters, which we employ in the following chapters as well. Afterwards, we introduce the robot model which we consider for image guidance. We then review an algorithm to determine whether a beam is feasible or blocked given a positioning and configuration of the guiding robot arm [79]. Finally, we present some basic results for our patient data and three manually selected robot-arm setups, which will serve as a reference for our methods and experiments in the following chapters.

4.1 Patient Data and Planning Parameters

In this thesis, we will consider ten cases of prostate cancer, which have been treated with the CyberKnife at Europäisches Cyberknife Zentrum München-Großhadern. The data has been provided by Christoph Fürweger. Besides the planning CTs and contours of the VOIs, the data also contains the parameters for the ray-tracing-based dose model. Our planning data and parameters essentially follow previous feasibility studies [78, 79]. We consider treatment with cylindrical collimators with

Table 4.1: Common treatment planning parameters for all patients. They contain the prescribed dose, lower and upper dose bound for the PTV, upper dose bounds for the OARs, i.e., bladder and rectum, and the upper MU limits for each individual beam, all beams of a node, and all beams (compare Section 2.3.5 for the notation).

d_{ps}	d_{min}	d_{max}	$d_{bladder}$	d_{rectum}
3625 cGy	0 cGy	4050 cGy	3600 cGy	3600 cGy
x_{beam}	x_{node}	x_{total}		
300 MU	10 000 MU	40 000 MU		

diameters of 10 mm, 15 mm, 20 mm, 30 mm, and 40 mm at an SAD of 800 mm. For all patients, we consider 99 beam nodes, which are the union of the actually employed nodes of all patients. However, some of the beam nodes cannot produce feasible beams for some patients due to the specific geometries of the setups.

We focus on the OCO approach introduced in Section 2.3.5, which maximizes the coverage of, in our case, the PTV. Basic parameters for the LP which we will solve for planning are listed in Table 4.1. They have been adapted from a five-fraction protocol and the prescribed dose for the PTV is set to 3625 cGy. We do not consider an explicit lower bound for the dose in the PTV but an upper bound of 4050 cGy. We consider two OARs, the bladder and the rectum. The dose to both structures is limited to 3600 cGy. Furthermore, we have upper bounds on the MUs. Any individual beam has at most 300 MU and the sum of all beams of a node is limited to 10 000 MU. This bound, however, has rather low impact due to the bound for the individual beams in combination with a total MU bound of 40 000 MU.

To obtain similar planning results for the different patients, we consider different upper dose limits for two SHELLs around the PTV. Table 4.2 shows the bounds, which range from 3500 cGy to 3660 cGy and from 2700 cGy to 2975 cGy for the inner and outer SHELL, respectively. They have been manually tuned to achieve a mean coverage of about 95 % for each patient. The table also lists the VOI sizes for planning of each patient.

4.2 Robot-Arm Model for Image Guidance

We consider an LBR iiwa robot arm for ultrasound guidance to follow the setup described in Section 3.6.3. In more detail, this light-weight model is specifically designed for human-robot collaboration tasks, especially in medical environments. It has seven rotational joints, which provide seven DOFs. While six DOFs are sufficient to realize arbitrary poses, i.e., positions and orientations of the end-effector, within a robot’s workspace, the additional DOF provides further flexibility by allowing to realize the same end-effector pose in, in this case, infinitely many configurations which are continuously connected. Continuously connected means

4.2 Robot-Arm Model for Image Guidance

Table 4.2: Upper dose bounds on the inner and outer SHELL structures and the PTV and OAR sizes of each patient. Volumes are given w.r.t. their number of voxels for planning after overlap removal and a voxel size of $3 \times 3 \times 3 \text{ mm}^3$.

Patient	PTV size	Bladder size	Rectum size	d_{SHELL_1}	d_{SHELL_2}
0	71.0 cm ³	80.6 cm ³	53.2 cm ³	3600 cGy	2700 cGy
1	98.9 cm ³	98.3 cm ³	90.5 cm ³	3660 cGy	2900 cGy
2	87.2 cm ³	98.2 cm ³	47.9 cm ³	3660 cGy	2900 cGy
3	86.4 cm ³	125.8 cm ³	40.4 cm ³	3525 cGy	2750 cGy
4	116.3 cm ³	80.2 cm ³	71.5 cm ³	3650 cGy	2975 cGy
5	74.4 cm ³	196.2 cm ³	94.6 cm ³	3575 cGy	2875 cGy
6	122.7 cm ³	58.8 cm ³	61.3 cm ³	3625 cGy	2925 cGy
7	90.5 cm ³	40.4 cm ³	56.3 cm ³	3600 cGy	2900 cGy
8	92.9 cm ³	116.6 cm ³	43.7 cm ³	3600 cGy	2875 cGy
9	56.4 cm ³	269.6 cm ³	125.0 cm ³	3500 cGy	2700 cGy

that we can keep the end-effector steady while changing through the different configurations. This is important for safe operation close to a patient like in our ultrasound-guidance scenario. The analytical solution for the inverse kinematics of this robot arm derived by Kuhlemann et al. [137] allows us to parametrize the configurations which result in the same end-effector pose by a single parameter which we refer to as LIFT. As already mentioned in Section 3.6.3, this parameter is an angle which describes the rotation of the robot’s elbow. A detailed geometric description of this angle and the elbow rotation can be found in Kuhlemann et al. [137]. See Figure 4.1 for a CAD model of the robot arm and an illustration of different LIFT angles. Note, however, that the elbow rotation does not affect every link of the robot arm. The inverse kinematics yield the required joint angles of the robot arm for each input consisting of the desired end-effector pose, LIFT angle, and three binary parameters ARM, ELBOW, and FLIP (also called WRIST). The three binary parameters are common to describe the remaining ambiguity of 6-DOF robot arms [111, 155]. However, the eight configurations which they describe represent fundamentally different arrangements of a robot and changing from one to another typically requires large motion.

Furthermore, our robot arm under consideration provides force sensing, i.e., it can estimate the forces applied to its end-effector. Although we do not consider this aspect explicitly in this thesis, it is very helpful for ultrasound guidance because thereby the robot can sense motion of the patient, like breathing, and compensate for it automatically. Thereby, a steady contact between the ultrasound transducer and the patient can be ensured. In our experimental setup, we equip the robot arm with a 3D ultrasound transducer. Specifically, we consider a GE 3V (General Electric, USA), which is a 2D matrix array probe. We will refer to the introduced 7-DOF robot arm equipped with the transducer as the GUIDE-robot in the following. Remember that there is also a LINAC-robot in CyberKnife which carries the LINAC for beam delivery from many directions around the patient.

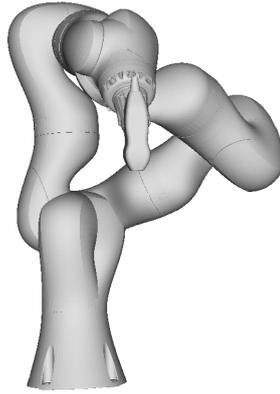


Figure 4.1: Overlay of the robot arm equipped with an ultrasound transducer when realizing the same end-effector pose with the LIFT angle 0° and the LIFT angle -90° .

We fix the orientation of the ultrasound transducer on the patient’s abdominal wall and consider three reference positions of the GUIDE-robot’s base around the patients, which we call A, B, and C. They are illustrated in Figure 4.2 when the LIFT angle is 0° and they all use the same ARM, ELBOW, and FLIP parameters. The positions shall reflect different choices for placing the robot, namely on the patient’s right next to the PTV (A), between the patient’s legs (B), and next to the patient’s right leg (C). Due to our defined orientation of the transducer, placing the robot at the left side of the patient would be less intuitive. Note that the fixed pose of the ultrasound transducer also fixes the GUIDE-robot’s last link and decreasingly limits the freedom of following links. Therefore, some beam directions will remain blocked independent of the robot positioning. Furthermore, note that the robot base is placed at different heights in the three reference setups in order to obtain a broad range of LIFT angles without collisions with the patient for each setup.

4.3 Finding Blocked Beams

We use the algorithm proposed by Gerlach et al. [79] in this thesis to determine whether a beam is blocked by the GUIDE-robot and will review its concept in this section. The algorithm is employed beam-node-wise. As sketched in Figure 4.3, we setup a projection plane which intersects the centroid of the PTV and is perpendicular to the line from the beam node through the centroid of the PTV. We compute the central projection of the GUIDE-robot, i.e., robot arm and ultrasound transducer, w.r.t. the beam node onto that projection plane. Only points of the GUIDE-robot are considered which lie between the projection plane and a parallel plane through the beam node. During the treatment, however, there will be some motion of the PTV. This motion, in turn, requires to adapt the

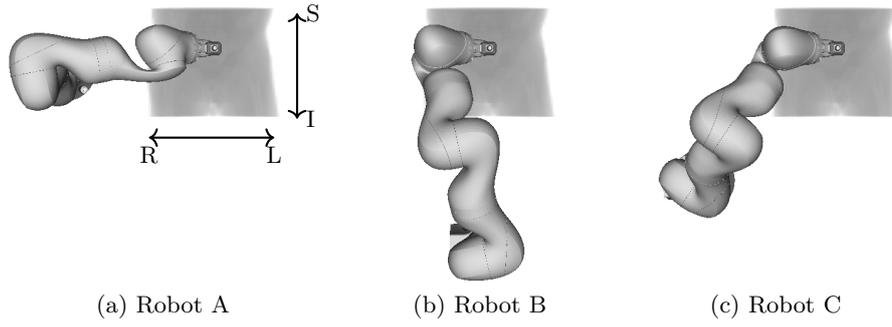


Figure 4.2: Anterior views on the three reference positions, which we call A, B, and C, of the GUIDE-robot holding the ultrasound transducer when the LIFT angle is 0° . The planning CT scan of patient 0 is shown.

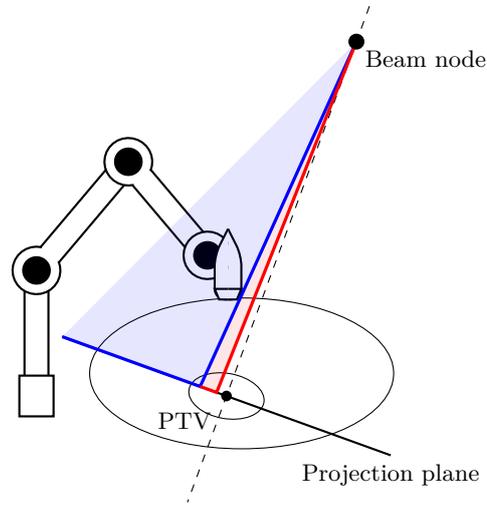


Figure 4.3: Central projection of the GUIDE robot from a beam node onto a plane orthogonal to the line from beam node to PTV centroid (blue). The projection is extended in the projection plane by a safety margin to account for potential motion during treatment (red).

positioning of the ultrasound transducer and of the delivered beam. For this reason, we extend the GUIDE-robot's projection by a safety margin of 20 mm. This margin size has been shown to be a reasonable choice [79]. The information of the projection is stored as a discrete grid with pixel size of $3 \times 3 \text{ mm}^2$.

Having the projection of the GUIDE-robot extended by the safety margin, we check for every candidate beam starting at that beam node whether it intersects with the projection. For this purpose, we need the extension of the intersection of the beam with the projection plane. We estimate that as the distance-transformed diameter of its collimator w.r.t. the distance from beam node to PTV centroid, extended by a diameter-dependent margin factor. This value describes a diameter

around the intersection of the beam’s central axis and the projection plane, which may not intersect with the GUIDE-robot’s extended projection. Otherwise, we define the beam as blocked and thereby infeasible for treatment.

When setting up a set of candidate beams, we have to check for all randomly sampled beams (see Section 2.3.4) whether they are actually feasible or intersect with the GUIDE-robot. Only feasible beams are added to the set of candidate beams and random sampling is repeated until the set contains our desired number of beams. If we employ multiple configurations of the GUIDE-robot during treatment in this thesis, we only discard beams which are infeasible in all of the employed configurations.

4.4 Experimental Setup and Evaluation Metrics

For all treatment planning tasks in this thesis, we use and extend an in-house planning software. The software is written in Java and uses the tool CPLEX (IBM, USA) for solving LPs. Our planning approach relies on random sampling of candidate beams. Depending on the sampled set of beams, the quantitative values of the resulting treatment plans vary. Therefore, we repeat all experiments several times with the same parameters but different random seeds for sampling in order to report statistical results. In this chapter’s experiments, we employ 50 random seeds for each parameter combination and consider sets of 6000 candidate beams. If we consider using multiple LIFT angles of the GUIDE-robot, we resample beams which are not feasible in any of the configurations.

As we employ OCO, the natural quantitative value which we obtain is the OCO-OV, i.e., the value of the mathematical objective function after optimization. Furthermore, we obtain the coverage of the PTV defined in (2.1) which is the related clinical goal. In our statistical evaluations, we will frequently consider correlation coefficients. The most famous correlation coefficient is Pearson’s correlation coefficient, which measures linear correlation with values ranging from -1 to 1 . A second one is Spearman’s rank correlation coefficient whose values also range from -1 to 1 . However, it computes the linear correlation not w.r.t. the data values themselves but w.r.t. to their rank in the datasets if sorted. Thereby, it is a measure for the monotony of their relation.

In this chapter’s experiments, we evaluate the impact of our three reference positionings of the GUIDE-robot on the treatment plan quality. For this purpose, we compare the results when using different numbers of LIFT angles during the treatment. If we consider only one LIFT angle, we arbitrarily select 0° as shown in Figure 4.2. If we consider multiple angles, we select them equidistantly from -90° to 90° . For example, 5 LIFT angles will mean in the following that we consider the LIFT angles -90° , -45° , 0° , 45° , and 90° . We will only consider odd numbers, such that each set of angles contains 0° . In particular, we consider 3, 5, 7, 19, and 37 angles, corresponding to step sizes of 90° , 45° , 30° , 10° , and 5° , respectively. While larger LIFT angles are feasible in principle, we have to consider that the

Table 4.3: Treatment planning results for each patient without a GUIDE-robot. The means and standard deviations are shown for the OCO-OV, the coverage, and the number of active beams. Additionally, Spearman’s correlation coefficient c is shown for the pairs of OCO-OV and coverage of the resulting treatment plans.

Patient	OCO-OV	COV in %	c	# treat. beams
0	2720 ± 1110	95.1 ± 1.0	-0.54	409 ± 10
1	3967 ± 654	95.3 ± 0.6	-0.32	435 ± 11
2	1747 ± 260	96.3 ± 0.8	-0.00	454 ± 10
3	2344 ± 682	95.8 ± 0.8	-0.43	478 ± 12
4	4825 ± 672	95.6 ± 0.5	-0.21	448 ± 10
5	2318 ± 476	95.5 ± 0.9	-0.10	415 ± 9
6	4532 ± 1066	95.6 ± 0.7	-0.71	464 ± 12
7	4190 ± 727	95.2 ± 0.7	-0.46	418 ± 11
8	2580 ± 708	95.7 ± 0.8	-0.63	452 ± 9
9	1790 ± 436	95.0 ± 1.0	-0.30	401 ± 10

GUIDE-robot would collide with the patient when rotating the elbow too far.

4.5 Results

As a reference, we report in Table 4.3 the treatment planning results which we obtain without considering a GUIDE-robot. The mean coverage ranges from 95.0% to 96.3% among the patients. However, the second highest coverage, which we obtain for patient 3, is already only 95.8%. The mean OCO-OVs of the patients range from 1747 for patient 2 to 4825 for patient 4, but should not be compared between different patients because they are not normalized and are based on completely different data. If we consider different plans for the same patient, however, the OCO-OV is a more comparable measure than the coverage, because it is the actual objective function of the optimization. The OCO-OV is related to the coverage but the relation is in general complex and not monotonous. Table 4.3 provides for each patient Spearman’s correlation coefficient for the two quantities. Note that the coefficients are only based on the final optimization results and do not describe the situation during a specific optimization run. For some patients’ results, we observe a good negative correlation. For other patients, like patient 2, there is no correlation, instead. This is illustrated in Figure 4.4 and further motivates to compare the OCO-OV rather than the coverage, especially if the coverage values are high and close to each other. The number of treatment beams, i.e., candidate beams with a non-zero number of MUs after optimization, is typically between 400 and 500.

When considering the three reference robot positionings A, B, and C with only a LIFT angle of 0° , we observe both a higher OCO-OV and a lower coverage in Figure 4.5a and Figure 4.5b, respectively. We observe that robot B provides the

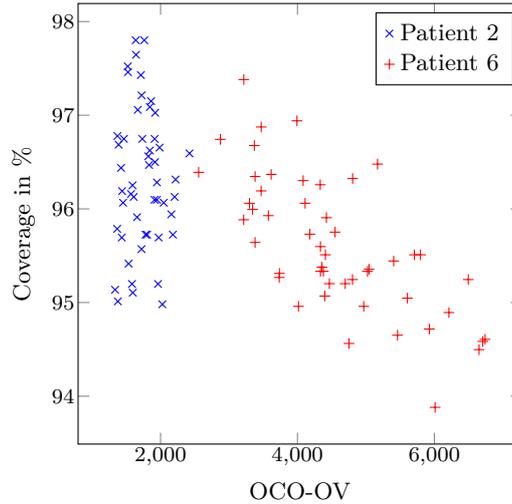
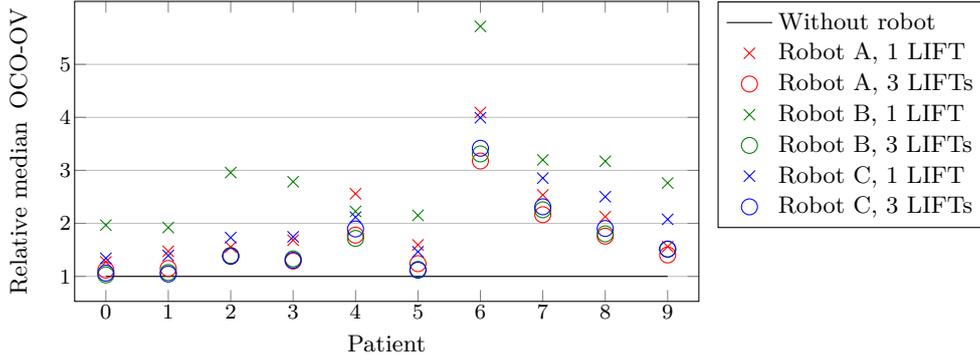


Figure 4.4: Scatter plot of the OCO-OV and coverage of the final treatment plans from all random seeds for patients 2 and 6.

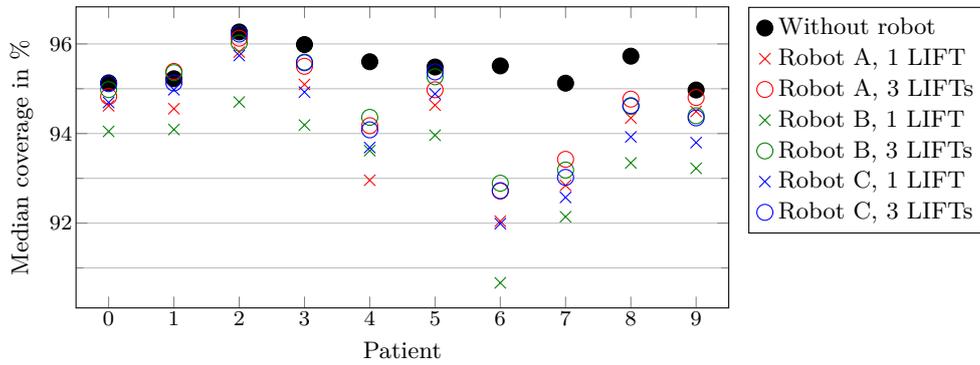
worst results, except for patient 4. If we employ three LIFT angles, i.e., -90° , 0° , and 90° , then we observe better results. For some patients, results are already close to the results without a GUIDE-robot. For patient 1, the planning results with robots A and B even have higher median coverages than without the GUIDE-robot. However, this is not the case for the OCO-OV, i.e., the actual objective value from planning. We further observe that the influence of the robot position is much smaller when employing three LIFT angles compared to the results for one LIFT angle. This observation highlights the impact and importance of the kinematic redundancy of the GUIDE-robot.

In more detail, Table 4.4 provides the mean OCO-OVs for robot A and different numbers of employed LIFT angles. It shows that employing three instead of only one LIFT angle provides a large improvement compared to the further improvements by employing at least five LIFT angles. Especially the results for patient 4 show no clear monotonicity. This is also the only patient for which five LIFT angles provide the lowest mean OCO-OV. For all others, we obtain the lowest result for 37 LIFT angles. However, the differences for the higher numbers of LIFT angles are about one order of magnitude smaller than the observed standard deviations.

Table 4.5 underlines that the impact of the robot position reduces when employing many LIFT angles during the treatment. For employing a single LIFT angle, the OCO-OVs are typically notably higher for robot B than for robot A and C. An exception is again patient 4, for which we observe the largest value for robot A. In summary, robot A provides the best mean values for six of the patients and robot C for the other four of them. This situation is different when considering 37 LIFT angles. The results of all three robots are rather close and robot B provides the lowest mean for five patients. Robot A is now only the best choice for three



(a) Relative OCO-OVs. Note that the result of robot A with 3 LIFT angles is hidden by robot B's result with 3 LIFT angles for patient 2 and robot B's result with 3 LIFT angles for patient 9 is hidden by robot C with 3 LIFT angles.



(b) Coverage

Figure 4.5: Comparison of the median OCO-OV relative to the setup without a GUIDE-robot (a) and the median coverage with and without GUIDE-robot (b). For the single LIFT angle cases, we select 0° . We additionally select -90° and 90° in the case of three LIFT angles.

Table 4.4: Means and standard deviations of the OCO-OVs when employing robot A with different numbers of LIFT angles. Angles are selected equidistantly from -90° to 90° . The first column is the patient index and the lowest mean achieved is printed bold for each patient.

	1 LIFT	3 LIFTs	5 LIFTs	7 LIFTs	19 LIFTs	37 LIFTs
0	3630 \pm 1222	3085 \pm 1098	3061 \pm 1067	3073 \pm 1073	3067 \pm 1044	2960 \pm 1041
1	6023 \pm 1049	4616 \pm 779	4545 \pm 771	4586 \pm 763	4497 \pm 748	4321 \pm 729
2	2784 \pm 479	2450 \pm 412	2398 \pm 396	2404 \pm 405	2402 \pm 393	2334 \pm 374
3	3950 \pm 762	2918 \pm 625	2938 \pm 635	2933 \pm 632	2935 \pm 631	2809 \pm 582
4	12352 \pm 1439	8550 \pm 1045	8399 \pm 1081	8584 \pm 1113	8455 \pm 1109	8445 \pm 1111
5	3725 \pm 615	2988 \pm 574	2800 \pm 545	2818 \pm 550	2793 \pm 560	2730 \pm 543
6	18114 \pm 3118	14575 \pm 2531	14182 \pm 2450	14261 \pm 2425	14150 \pm 2492	14030 \pm 2491
7	10250 \pm 1279	8889 \pm 1154	8632 \pm 1153	8719 \pm 1175	8621 \pm 1162	8586 \pm 1139
8	5363 \pm 1095	4620 \pm 875	4537 \pm 841	4538 \pm 840	4535 \pm 831	4527 \pm 829
9	2789 \pm 581	2521 \pm 554	2510 \pm 531	2528 \pm 540	2521 \pm 534	2509 \pm 548

4 Basic Treatment-Planning Setup for Ultrasound-Guided Radiosurgery

Table 4.5: Mean and standard deviation of the OCO-OVs for reference robot A, B, and C when employing only 1 LIFT angle or 37 LIFT angles. For each patient and number of LIFT angles, the lowest mean is printed bold.

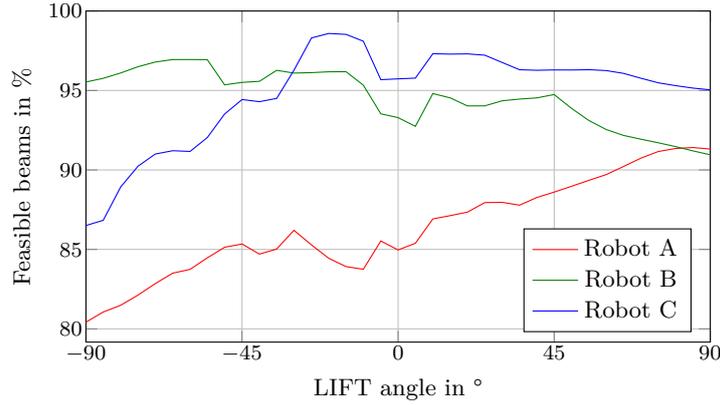
	1 LIFT			37 LIFTs		
	Robot A	Robot B	Robot C	Robot A	Robot B	Robot C
0	3630 ± 1222	5554 ± 1206	3659 ± 988	2960 ± 1041	2845 ± 969	2850 ± 966
1	6023 ± 1049	7728 ± 1092	5642 ± 840	4321 ± 729	4203 ± 747	4194 ± 741
2	2784 ± 479	5222 ± 769	2993 ± 439	2334 ± 374	2330 ± 349	2333 ± 358
3	3950 ± 762	6221 ± 1101	3973 ± 738	2809 ± 582	2973 ± 604	2972 ± 608
4	12352 ± 1439	10663 ± 1447	10284 ± 1356	8445 ± 1111	8106 ± 1036	8156 ± 1027
5	3725 ± 615	5086 ± 694	3454 ± 616	2730 ± 543	2641 ± 496	2640 ± 496
6	18114 ± 3118	25056 ± 4153	17766 ± 2893	14030 ± 2491	13733 ± 2410	13845 ± 2407
7	10250 ± 1279	13155 ± 1723	11762 ± 1535	8586 ± 1139	8872 ± 1265	8892 ± 1219
8	5363 ± 1095	8023 ± 1284	6316 ± 942	4527 ± 829	4490 ± 775	4496 ± 768
9	2789 ± 581	5070 ± 818	3772 ± 688	2509 ± 548	2584 ± 584	2607 ± 568

Table 4.6: Average percentages of the candidate beams and of the treatment beams which are not feasible in all of the employed configurations, but blocked in at least one of them. The percentages are presented for employing 3, 7, and 37 equidistant configurations and are shown separately for the robot positions A, B, and C.

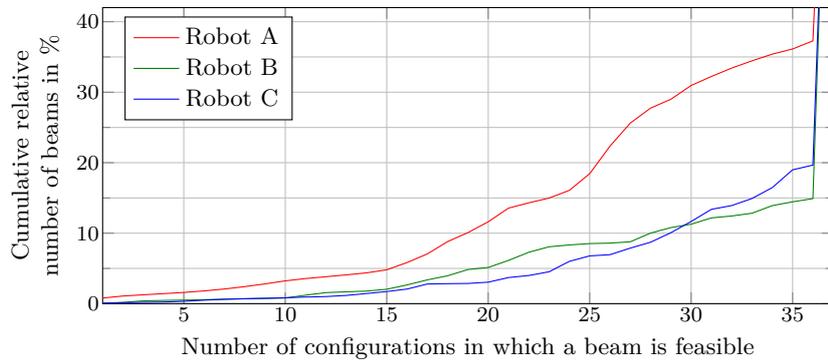
# LIFTs	3			7			37			
	Robot position	A	B	C	A	B	C	A	B	C
Candidate beams in %		31.5	13.1	17.1	36.1	13.5	19.0	37.3	14.9	19.7
Treatment beams in %		33.1	32.9	30.6	37.2	33.5	33.7	38.4	36.0	34.5

patients and robot C for two patients.

Another interesting aspect is in how many of the employed robot configurations a beam is feasible. Figure 4.6a shows that, in the 37 LIFT angles scenario, in every configuration of every robot setup more than 80% of the candidate beams are feasible. For robot A, we generally observe lower percentages than for the other two. If we consider the beams instead of the configurations, we see from Figure 4.6b that on average less than 0.8%, i.e., less than 48, of the candidate beams are only feasible in a single configuration. Note that due to the resampling every beam is feasible in at least one configuration. Nevertheless, this shows that 5° sampling of the LIFT angles does not yet result in complete redundancy of the configurations. If we consider the 10th percentiles, we obtain feasibility in at most 19, 28, and 29 configurations for robot A, B, and C, respectively. An interesting difference appears when comparing the statistics for candidate beams and for the subset of actual treatment beams. Table 4.6 shows that the mean percentage of beams which are infeasible in at least one configuration is higher for the treatment beams than for the candidate beams. Furthermore, the percentages largely differ between robot A and robots B and C for the candidate beams, but they are much closer for the treatment beams. Around a third of all treatment beams are infeasible in at least one of the employed configurations.



(a) Feasible candidate beams in each configuration



(b) Cumulative histogram of feasibility

Figure 4.6: Statistics on the feasibility of the candidate beams when employing 37 equidistant LIFT angles from -90° to 90° . Shown are the average percentages of candidate beams which are feasible in each individual configuration (a) and a cumulative histogram of in how many configurations candidate beams are feasible (b). Each candidate beam is feasible at least in one configuration, due to the resampling process, and at most in all 37 employed configurations.

4.6 Discussion

In this chapter, we presented the basic treatment planning aspects required for planning in presence of a GUIDE-robot. Furthermore, we introduced the patient data, planning parameters, and three robot settings which serve as a reference in the subsequent chapters. Our results show that employing the reference GUIDE-robot setups with only one configuration has an impact on the planning results (Figure 4.5). The impact decreases when employing more configurations, i.e., changing the GUIDE-robot's LIFT angle during the treatment to allow for more beam directions. However, there is a saturation and simply increasing the number of configurations rather converges to a constant gap w.r.t. to the plans without any GUIDE-robot (Table 4.4). Interestingly, the remaining gap is on a similar

level for all three reference robots while their results in case of employing only one LIFT angle differ notably (Table 4.5). This indicates that the remaining gap should be mainly caused by the ultrasound transducer and the robot parts close to the end-effector, which are positioned similarly for the three robots due to the defined pose of the ultrasound transducer. Lastly, we showed that beams which are not feasible in all employed configurations are over-proportionately selected for treatment (Table 4.6). This emphasizes again the impact and importance of employing multiple configurations.

Although our results indicate that employing multiple configurations during treatment seems to be a good choice for planning, it has some practical difficulties. First, changing the configuration takes time and, second, it has to be synchronized with the beam delivery to ensure that we can always actually deliver the current beam. Third, the robot motion is quite close to the patient. With respect to both safety and patient comfort, we should aim to minimize that motion. We will further discuss these issues in Chapter 6, where we will present methods to synchronize the motion of the GUIDE-robot and the beam delivery. Prior to this, however, we will study in Chapter 5 how we can optimize the GUIDE-robot setup in order to obtain better planning results with low numbers of configurations in order to keep the synchronization effort already as low as possible. We will both investigate positioning of the robot, which we have shown to be especially important when employing only a few configurations, and the selection process of suitable LIFT angles for treatment.

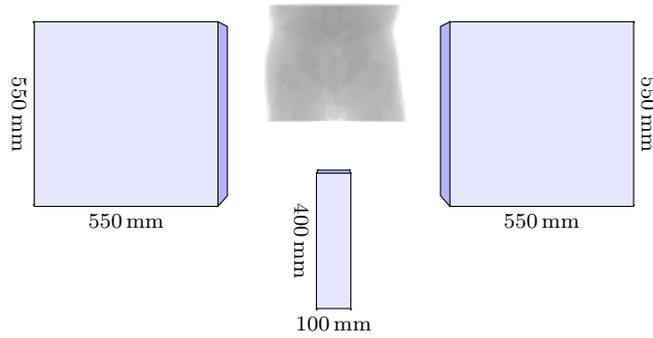
5 Optimizing a Robotic Setup for Ultrasound-Guided Radiosurgery

In the previous chapter, we confirmed and extended results from previous studies [78, 79] which have shown that ultrasound guidance using a kinematically redundant robot arm is feasible for prostate treatments from a treatment plan quality's point of view. However, these studies also indicated that the achievable plan quality can severely vary depending on the positioning of the robot around the patient. The studies came to this conclusion by evaluating a variety of manually defined positionings. For this reason, the first aspect of this chapter, is how the robot positioning problem can be formalized for systematic evaluation and automatic optimization, which we will analyze in Section 5.1.

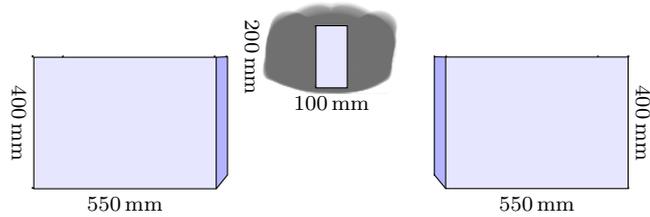
The mentioned studies further showed that exploiting the kinematic redundancy of the robot arm during the treatment can further mitigate the negative impact of beam blocking on treatment plan quality. This was tested by evaluating three configurations of the robot's elbow which correspond to a rotation of the central configuration by -90° and 90° . We confirmed this observation in the previous chapter but additionally observed that adding a very high number of configurations arbitrarily is not meaningful. This leads us to the second aspect of this chapter. We will analyze in Section 5.2 the problem to decide which configurations should actually be used and provide two optimization approaches. While using more configuration would obviously never decrease the expectable plan quality, it still has some other disadvantages. On the one hand, it requires robot motion close to the patient which might in general be undesired due to safety aspects and also patient comfort. On the other hand, it requires a synchronization with the beam delivery, because not every treatment beam can be delivered in every configuration of the GUIDE-robot. This potentially increases the overall treatment time due to additional idle times of the LINAC during robot movement if no efficient synchronization between the GUIDE-robot and the LINAC-robot exists.

In summary, this chapter addresses the first of the research questions formulated in Section 1.2. We separately consider the analysis and optimization of the GUIDE-robot's positioning and the selection of suitable configurations given a positioning. Furthermore, we evaluate the methods and their impact on plan quality in comparison to our reference setups introduced in the previous chapter. Parts of this chapter have been published in a journal [232] and have been presented at conferences [228, 231].

5 Optimizing a Robotic Setup for Ultrasound-Guided Radiosurgery



(a) View onto xz -plane (anterior view)



(b) View onto xy -plane (inferior view)

Figure 5.1: Search space for positioning of the GUIDE-robot's base. It consists of three disjoint boxes which are aligned around the planning CT of, in this example, patient 0.

5.1 Optimization of the Robot Positioning

5.1.1 Search Space

As a general setup for the optimization, we consider three areas in which the base of the GUIDE-robot can be placed (Figure 5.1). To the left and to the right of the patient, we define the areas as boxes of size $550 \times 400 \times 550 \text{ mm}^3$. Additionally, we define a box of size $100 \times 200 \times 400 \text{ mm}^3$ between the patient's legs. The exact positions of the boxes are defined relative to the CT center of the treatment plan. Despite the position of the robot base, we have the LIFT angle as another variable of our search. However, the actual freedom in choosing the LIFT angle is limited due to potential collisions with the patient. Therefore, its feasible range of values depends on the position of the base.

Considering our three reference positions of the GUIDE-robot's base from Section 4.2, positions A and C are within the box to the right of the patient. Position C is within the box between the legs of the patient. We did not consider positions to the left of the patient for the reference positions due to our defined pose of the ultrasound transducer. Nevertheless, we now consider this rather unintuitive choice for manual positioning in our automated analysis and optimization.

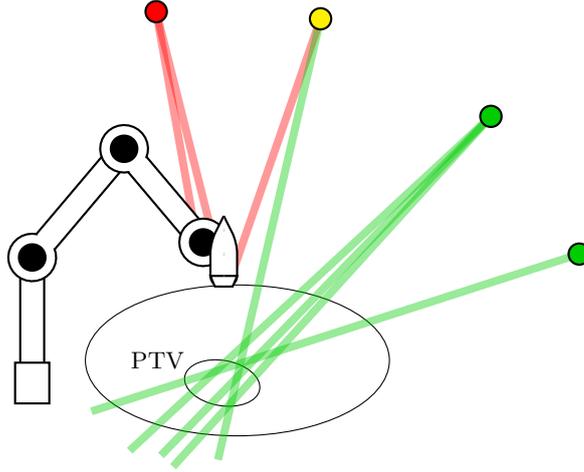


Figure 5.2: Sketch of a setup with three blocked beams (red lines), one blocked node (red circle), and two free nodes (green circles). Following our definition, the fourth node (yellow circle) is neither blocked nor free because both blocked and non-blocked beams (green lines) start at this node.

5.1.2 Objective Functions

In order to formulate an optimization problem for the positioning of the GUIDE-robot, we need to define meaningful objective functions. All objective functions which we propose in the following assume that we already have a fixed set of candidate beams, which we can use to evaluate them.

Number of Blocked Beams (BB) For this objective, we count how many of the candidate beams are blocked by the current robot positioning. The less beams are blocked, the less impact the GUIDE-robot has on the treatment planning. Figure 5.2 illustrates a small example, in which the BB objective value is 3.

Overlap of Projected Robot and PTV (OP) The beam-blocking algorithm projects the GUIDE-robot onto a plane which is orthogonal to the line from a beam node to the centroid of the PTV (see Section 4.3). For this objective function, we additionally project the PTV onto the same plane and count how many pixels of the projected GUIDE-robot, including the safety margin, and the projected PTV overlap. For the objective value, we simply sum up the overlapping pixels w.r.t. all beam nodes. The more overlap we observe, the more shadowed is the PTV by the GUIDE-robot. In comparison to the BB objective, this objective is less dependent on the candidate beam set. However, we only consider projections from beam nodes at which at least one candidate beam starts. If the candidate beam set is sufficiently large, we can assume that the other nodes are not usable for treatment at all.

Number of Blocked Nodes (BN) We count the number of beam nodes for which all candidate beams are blocked and define them as blocked nodes. These nodes effectively have to be removed from treatment planning. Figure 5.2 illustrates an example, in which the BN objective value is 1.

Number of Free Nodes (FN) In contrast to the previous objective, we count the number of beam nodes for which no candidate beam is blocked and define them as free nodes. Therefore, these nodes are fully available for planning and not influenced by the image-guidance setup, at least w.r.t. the candidate beams at hand. Figure 5.2 illustrates an example, in which the FN objective value is 2.

All proposed objective functions require to apply the beam-blocking algorithm from Section 4.3 for evaluation. Therefore, it is not possible to give analytic expressions for them. Furthermore, they are highly discontinuous due to counting and the discrete nature of the beams, for which we assume a finite diameter to calculate blocking. For this reason, we will have to employ heuristic optimization algorithms which do not give any guarantees for the quality of their solutions.

5.1.3 Optimization Strategies

We propose two optimization strategies for optimizing the positioning of the GUIDE-robot. The first approach, which we call segmental inverse robot construction (SIRC), is a problem-specific heuristic search strategy. The second one is an implementation of the generic strategy simulated annealing (SimAn).

Segmental Inverse Robot Construction (SIRC)

A robot arm is a sequence of links and sequential links are connected by a joint. The typical task in robotics is to determine suitable joint angles given a fixed pose of the robot base and a desired end-effector pose. In our case, however, we do not have a fixed pose of the base but aim to optimize it w.r.t. the objective functions proposed in the previous section. Nevertheless, we consider a fixed pose of the ultrasound transducer fixing the GUIDE-robot's end-effector pose, i.e., the angles of the last joints. In the SIRC approach which we propose, we iteratively fix the remaining joints.

For this purpose, we describe the GUIDE-robot by five points P_1, \dots, P_5 where P_1 is at the base and P_5 at the end-effector. This concept is sketched in Figure 5.3. The points are chosen in such a way that they need to maintain a defined physical distance in every configuration of the robot due to the geometry of the links. This is the case for the centers of joints not forming a line with their two neighboring joints in every configuration. P_5 and P_4 are already fixed due to the fixed position and orientation of the ultrasound transducer. The remaining three points are subject to our optimization. Due to their physically required distances, each remaining point P_i , where $i \in \{1, 2, 3\}$, has to be located on a sphere around P_{i+1} .

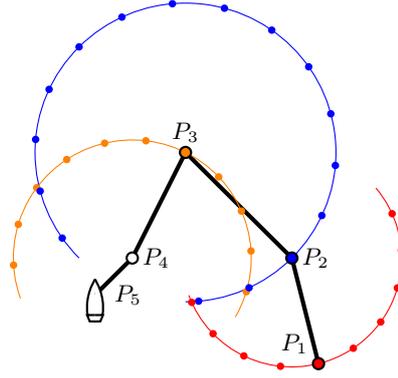


Figure 5.3: 2D sketch of the SIRC approach. The robot configuration is described by five points. P_5 and P_4 are fixed due to the pose of the ultrasound transducer. The remaining points are selected iteratively on spheres around the previous points.

From this observation, we derive the following optimization strategy. Given P_5 and P_4 , we position P_3 on the sphere around P_4 such that it yields the best objective value when considering beam blocking by only the parts of the GUIDE-robot which are fixed so far. Subsequently, we select P_2 and P_1 in the same manner. Effectively, selecting a point corresponds to selecting the angles of two joints. The angle of the very first joint at the base is arbitrary. The resulting robot configuration, however, will not be optimal because the selection of a point limits the freedom to select the subsequent points. Therefore, we employ a depth-first search for better configurations and truncate branches of the search tree that cannot lead to better objective values. Furthermore, we have to discard solutions which would collide with the patient or position the robot base outside of the search space defined in Section 5.1.1. Note that this approach assumes that the base can be mounted in the treatment room with an arbitrary orientation.

In our implementation, we aim to sample at least 36 feasible points on each sphere. However, we limit the maximum number of resampling steps for finding feasible points. The points are generated by sampling angles for the two related joints, because this is easier from an implementation's point of view than deriving the angles from the points. The pose of the robot base can be derived from the final joint angles employing the forward kinematics backwards. If explicit knowledge of the parameters LIFT, ARM, ELBOW, and FLIP is required, they are harder to derive from the solution than the base pose. One possibility is to consider the four parameters as variables and setup an optimization which varies them until the inverse kinematics yield the same joint angles as we obtained from SIRC.

Simulated Annealing (SimAn)

SimAn is a well-known generic concept for probabilistic search to approximate complex optimization problems. It does not require to calculate derivatives but

considers sampling from a neighborhood of the current solution for iteration. Strictly local search algorithms, like hill climbing in which we always move towards the best neighboring solution, often get stuck in low-quality local optima. To reduce this issue, SimAn allows to move towards worse neighboring solution with some probability hoping to leave a local optimum and to find better solutions elsewhere in the search space. This probability is higher the closer the worse solution's objective value is to the current objective value. Furthermore, this probability decreases over time, i.e., with increasing number of iterations, in order to eventually force convergence to a local optimum. Informally, we can say that the algorithm is initially allowed to freely jump around in the search space, but eventually has to settle down. The best solution found during the whole process is the final result of SimAn.

In particular, SimAn considers minimization of a scalar function $f(\mathbf{y})$ with a vector of decision variables \mathbf{y} and an initial solution \mathbf{y}_0 . In every iteration $i > 0$, we first randomly select a solution $\mathbf{y}^n \in \mathcal{N}(\mathbf{y}_{i-1})$, where $\mathcal{N}(\mathbf{y})$ defines the neighborhood of a solution and \mathbf{y}_{i-1} is the solution from the previous iteration. We then calculate the objective value $f(\mathbf{y}^n)$ and compare it to $f(\mathbf{y}_{i-1})$. We set the solution of the current iteration as

$$\mathbf{y}_i = \begin{cases} \mathbf{y}^n & \text{if } f(\mathbf{y}^n) \leq f(\mathbf{y}_{i-1}) \\ \mathbf{y}^n & \text{if } f(\mathbf{y}^n) > f(\mathbf{y}_{i-1}) \wedge p_i \leq E_i(\mathbf{y}^n, \mathbf{y}_{i-1}) \\ \mathbf{y}_{i-1} & \text{else} \end{cases} \quad (5.1)$$

where $p_i \in [0, 1]$ is a real-valued random number drawn from a uniform distribution and determines whether we accept \mathbf{y}^n although it provides a worse objective value. For this purpose, we calculate the so-called energy

$$E_i(\mathbf{a}, \mathbf{b}) = \exp\left(-\frac{f(\mathbf{a}) - f(\mathbf{b})}{\Theta_i}\right), \quad f(\mathbf{a}) > f(\mathbf{b}), \quad \Theta_i > 0 \quad (5.2)$$

where Θ_i is called temperature and monotonously decreases with increasing i . Thereby, the likelihood to accept a worse solution decreases with increasing difference between the objective values and additionally decreases over time. Note that SimAn approaches a hill-climbing algorithm for $\Theta_i \rightarrow 0$. After we have determined \mathbf{y}_i , we start the next iteration unless some termination criterion is met. In the simplest case, this criterion is a maximum number of iterations. Other criteria are, for example, no improvement of the objective for a defined number of iterations or finding an objective value below a predefined threshold. As mentioned before, the final solution is the solution with the lowest objective value we found in any iteration, which is not necessarily the solution of the last iteration. The overall process is sketched in Figure 5.4.

In our implementation, we consider a problem with four decision variables, i.e., $\mathbf{y} \in \mathbb{R}^4$, namely the three-dimensional position of the GUIDE-robot's base and its LIFT angle. The orientation of the base is fixed to be orthogonal to the floor. We define preferred values for the three binary parameters ARM, ELBOW, and

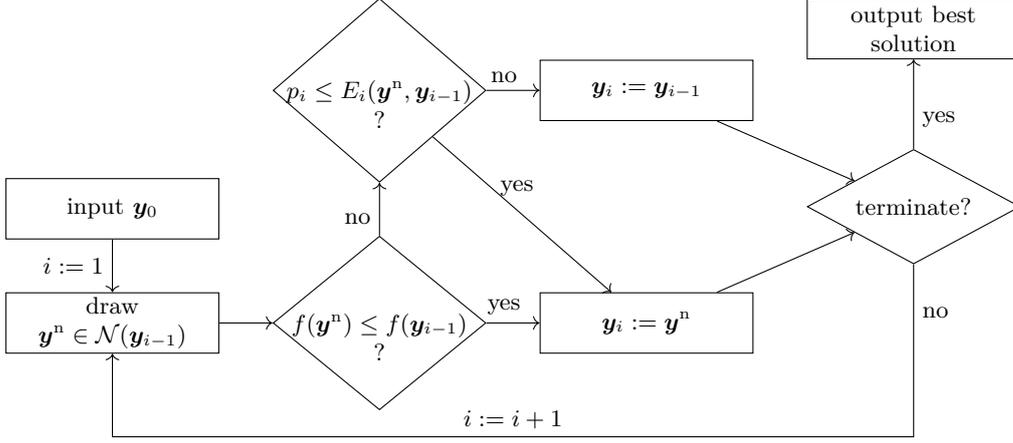


Figure 5.4: Basic steps of SimAn.

FLIP (see Section 4.2) but try all combinations if the preferred combination is not feasible for a given base position and LIFT angle. Feasibility of a solution also includes that the GUIDE-robot does not collide with the patient. If no feasible configuration exists for a solution, then this solution is discarded. Mostly, SimAn is applied to discrete variables. However, we employ a continuous approach allowing arbitrary real values for \mathbf{y} . For this purpose, we define drawing a neighboring solution \mathbf{y}^n for \mathbf{y}_{i-1} as

$$\mathbf{y}^n = \mathbf{y}_{i-1} + \frac{\Theta_i}{50} \mathbf{r}_i \odot \mathbf{l}, \quad (5.3)$$

where $\mathbf{r}_i \in \mathbb{R}^4$ is a vector of random numbers drawn from a standard normal distribution, which is element-wisely multiplied with a vector $\mathbf{l} \in \mathbb{R}^4$ containing the range between lower and upper bound for each of the four decision variables. Therefore, the neighboring solution is obtained from the last iteration's solution by making a step of random size in each dimension. The step size decreases over time because we include the temperature Θ_i . Another uncommon aspect of our problem-at-hand is that the search space consists of three unconnected boxes. Therefore, we implement the possibility to move from one box to another when the neighboring solution reaches a border in direction to the other box.

We define a maximum number of 500 iterations as the termination criterion. The temperature is initially $\Theta_0 = 10$ and is halved after every 100 iterations. SimAn requires to provide an initial solution \mathbf{y}_0 which, in general, influences the result. We consider five initial solutions for each parameter combination. They all have a LIFT angle of 0° and are located far left, close left, inferior, close right, and far right as shown in Figure 5.5.

On the one hand, the SimAn approach is less problem-specific than SIRC. Furthermore, our implementation only considers a fixed orientation of the base to have a small number of decision variables. On the other hand, it provides direct access to the variables-of-interest. In particular, it allows to explicitly control the LIFT

5 Optimizing a Robotic Setup for Ultrasound-Guided Radiosurgery

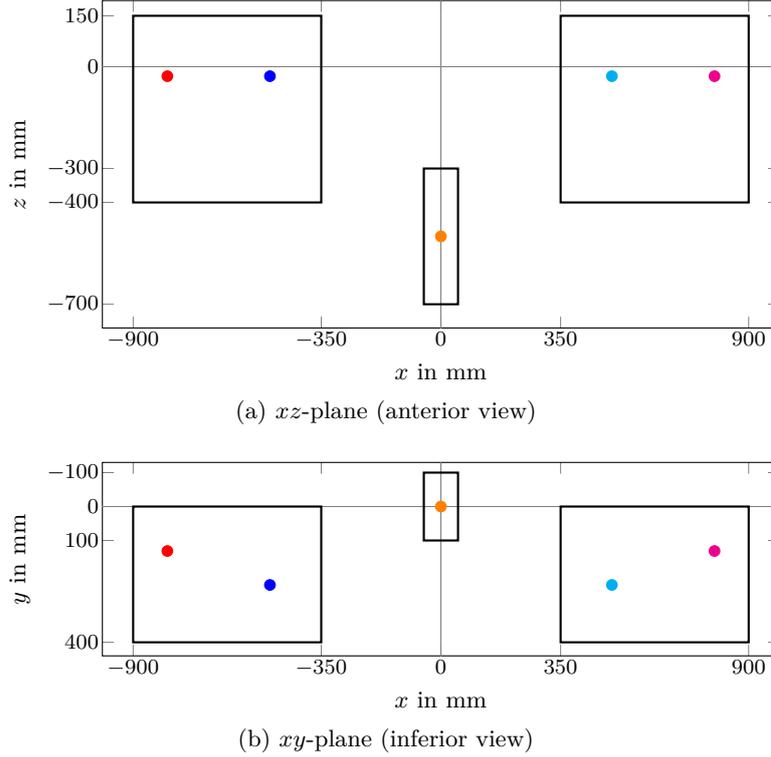


Figure 5.5: The five initial solutions for the position of the robot base projected onto the xz -plane (a) and the xy -plane (b). Note that the LIFT angle is another decision variable, but set to 0° for all five initial solutions.

angle. Additionally, its generic structure allows to easily add further variables, for example the orientation of the robot base or additional LIFT angles. However, increasing the dimension of the search space also increases the complexity and more iterations will be required to find a solution with a similar quality as for the lower-dimensional problem.

5.1.4 Experimental Setup

We evaluate our two optimization strategies and the four proposed objective functions by comparing their treatment planning results to the three reference setups, which are described in Section 4.4 and evaluated in Section 4.5. Because the SimAn-based strategy allows to directly control the LIFT angle, we also provide a study considering three angle values. However, in order to keep the same dimensionality of the search space, we only consider one of them as a decision variable and define the other two as rotated by -90° and 90° . To evaluate the objective functions, we sample a random set of candidate beams. After optimization of the robot, we resample beams which are blocked before solving the planning LP.

5.2 Selection of an Efficient Set of Robot Configurations

The SimAn approach which we introduced in the previous section is very generic and can easily be extended to optimize further variables. However, if we add variables, i.e., increase the dimensionality of the search space, it becomes harder to obtain high-quality solutions. We can expect to need more iterations to adequately explore the larger search space. If we specifically consider the selection of multiple configurations for the GUIDE-robot, then we will also find some problem-specific issues. First, we have to define beforehand how many configuration we want to consider. Second, we have to re-compute the beam blocking in each iteration for the current positioning of the robot and all current LIFT angles, which might be many. Third, if the optimization finishes, then it is still not clear whether our number of configurations is actually a good choice or whether it is lower or higher than necessary. This issue could be fixed by extending the objective function to optimize further criteria related to the configurations, but this would increase the complexity and obscure the underlying model even more.

For these reasons, we introduce in this section two different approaches towards the configuration selection, which split the problem from the positioning problem. For a fixed robot positioning, we first describe a post-processing approach which reduces the number of configurations given a set of treatment beams resulting from inverse planning. Second, we show that the selection problem can be directly integrated into the treatment planning LP, which allows for more control over the optimization process and direct inclusion of the treatment planning objectives. In contrast, the generic SimAn approach has to use a surrogate objective function, because solving the planning LP in each iteration is computationally not tractable. With the direct integration in the planning problem, we can hope to obtain better insights into the problem and the influence of the number and distribution of configurations on treatment plan quality.

5.2.1 Discretization

If we allow for arbitrary configurations, i.e., have continuous variables describing the angle of a configuration, we will always end up with a highly non-linear problem, because it implies computing the blocking or non-blocking by the current configurations for each beam throughout the optimization steps. Therefore, it seems more promising to discretize the infinite set of configurations to a finite set. This is especially meaningful as we only consider a finite set of beams as well. From a certain level of granularity on, a finer discretization cannot create new beam blocking patterns, but they will be equal to those of their neighboring configurations. Configurations blocking the exactly same beams are for our application equivalent, as long as we do not assign further objectives or figures of merit. In theory, there are 2^N different blocking patterns if we consider N beams. In practice, however, we expect only a much smaller subset to be actually realizable due to the underlying geometry of the GUIDE-robot and the spatial ar-

rangement of the beams. We will only consider finite sets of configurations in the following, namely the set of 37 equidistant LIFT angles introduced in Section 4.4 which span the range from -90° to 90° in steps of 5° . In consequence, our task is to select a subset of these configurations, which is optimal in some sense. This task leads to combinatorial optimization problems.

5.2.2 Post-Processing as a Set Cover Problem

If we solve the inverse planning problem considering a finite set of candidate configurations, we obtain a set of treatment beams. We can now check in a post-processing step how many of the configurations are actually required to allow for delivering these treatment beams. The problem of finding the minimum number of configurations to deliver all treatment beams can be identified with the set cover problem (SCP).

The SCP considers a set of elements U , called universe, and a set S which consists of subsets of U such that their union equals U , i.e.,

$$S \subseteq \mathcal{P}(U) \text{ with } \bigcup_{\hat{S} \in S} \hat{S} = U \quad (5.4)$$

where $\mathcal{P}(U)$ is the power set of U , which is the set of all subsets of U . Consider a subset $T \subseteq S$. An element $u \in U$ is called *covered* by T , iff it is an element of the union of the sets in T . The task in the SCP is to solve

$$\min_{T \subseteq S} |T| \quad \text{s.t.} \quad \bigcup_{\hat{T} \in T} \hat{T} = U, \quad (5.5)$$

i.e., to find the smallest set T which covers all elements of U . Figure 5.6 show a small example for an SCP. In our application, we can identify the elements of U as our treatment beams, which we want to deliver. S contains one subset of U per GUIDE-robot configuration consisting of these treatment beams which are feasible in that configuration. Note that the condition in (5.4) is ensured by our requirement for beam sampling that a beam has to be feasible in at least one configuration. After optimization, $|T|$ will be the smallest number of configurations which we have to employ in order to deliver all treatment beams.

The SCP is NP-hard and can be described by an integer linear program (ILP). If we enumerate the elements of S , we can introduce a binary variable $x_i \in \{0, 1\}$ for each $S_i \in S$, which is one iff the set is an element of T , and minimize their total sum. Furthermore, we have for each $u \in U$ a constraint

$$\sum_{i \in \{j: u \in S_j \subseteq S\}} x_i \geq 1, \quad (5.6)$$

ensuring that T covers u .

There is a simple greedy algorithm to approximate this problem. It starts with an empty set $T = \emptyset$. In each iteration, it counts for each $\hat{S} \in S$ how many

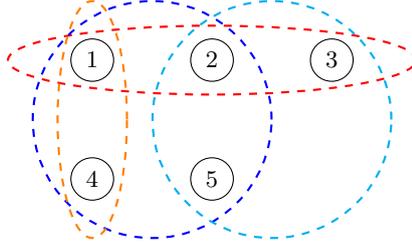


Figure 5.6: Illustration of an SCP with the universe $U = \{1, 2, 3, 4, 5\}$ and the set $S = \{\{1, 2, 3\}, \{1, 2, 4, 5\}, \{2, 3, 5\}, \{1, 4\}\}$, whose elements are illustrated in red, blue, cyan, and orange, respectively. Covering all elements of U requires to select at least two sets from S , for example $\{1, 2, 3\}$ (red) and $\{1, 2, 4, 5\}$ (blue).

elements of U which are not yet covered by T are in \hat{S} . The set \hat{S} containing the most yet-uncovered elements is then added to T . The algorithm terminates when T covers all elements of U , which is the case after at most $|S|$ iterations. As a simple example in which the greedy algorithm will not provide an optimal solution, consider $U = \{1, 2, 3, 4, 5, 6\}$ and $S = \{\{1, 2, 3, 4\}, \{1, 2, 5\}, \{3, 4, 6\}\}$. In the first iteration, the largest set in S will always be selected. In this example, this selection requires to subsequently select the other two sets as well in order to cover the elements 5 and 6. However, selecting only the two sets having three elements would already be sufficient to cover all elements of U .

5.2.3 Direct Integration into the Planning LP

A disadvantage of the described post-processing approach is that it leads to two decoupled optimization problems, the inverse planning and the selection of configurations. However, these problems are actually strongly coupled. Consider the case that we can cover all treatment beams except for one with K configurations. In the previous approach, we are then forced to use another configuration as well to cover the last remaining beam. However, a single beam has typically no major influence on the inverse planning but can be replaced by one or more other candidate beams. Therefore, it might be possible to identify other beams which are already covered by the K configurations without any change in the treatment plan quality. For this reason it is interesting to look into a coupling of the two problems.

One possibility to combine the two optimization problems is to combine their objective functions, i.e., considering a weighted sum of the planning objective and a term penalizing the number of required configurations. However, this kind of combination always requires to fine-tune the weighting factor to find a proper trade-off between them. A different approach is a sequential optimization, as employed in the clinically implemented step-wise inverse planning, in which we first optimize one objective, turn the result into a constraint, relax some constraint if necessary, and then optimize the second objective.

We decide for the following sequential and iterative approach which allows to

analyze the actually required number of configurations. First, we sample a finite set of candidate configurations and sample random candidate beams which are feasible in at least one of the configurations. Using these candidate beams, we solve the inverse planning LP for OCO. This results in an objective value s^* describing the remaining gap to full coverage. In the second step, we minimize the number of configurations but convert (2.8) to a constraint

$$\mathbf{1}^\top \mathbf{s} \leq r \cdot s^* \quad (5.7)$$

with a relaxation factor $r > 1$ close to one. In other words, we seek for the smallest subset of configurations which provides the same or a better OCO-OV. Note that the activation times of all candidate beams are still decision variables and the set of treatment beams might therefore change during this second optimization step. However, we have to provide a mathematical model to describe the mentioned objective and the fact that a beam might become infeasible if we deselect some configurations.

The probably most intuitive approach introduces M binary variables $c_i \in \{0, 1\}$, $i = 1, \dots, M$, one for each of the M candidate configurations. The variable c_i is zero iff configuration i is not used. The objective to be optimized then simply is

$$\min \sum_{i=1}^M c_i. \quad (5.8)$$

We can further introduce binary constants $f_{ij} \in \{0, 1\}$ which are one iff beam j is feasible in configuration i . The constraint that a beam j cannot be active, i.e., we require $x_j = 0$, if it is not feasible in any of the configurations in use then becomes

$$x_j \leq H \cdot \sum_{i=1}^M f_{ij} c_i \quad (5.9)$$

where H is a sufficiently large constant, like for example the upper MU bound x_{beam} for each beam from (2.14). If there is at least one i with $f_{ij} c_i = 1$, then the constraint does not actively limit the activation time x_j . Otherwise, this beam is no longer feasible and x_j is forced by the constraint to be zero.

This approach turns the inverse planning problem into a MIP, which is much harder to solve. In the worst case, for each of the 2^M theoretical combinations of configurations a modified planning LP has to be solved to find the optimal MIP solution. Therefore, we relax the problem above by defining c_i as non-negative real-valued variables, instead, to obtain an LP. Nevertheless, our previous definition that c_i is zero iff the configuration is not used still holds. We can also still use (5.8) and (5.9), but their numerical values are now more difficult to interpret. Instead of directly looking at the value of (5.8), we have to count the variables c_i which have a non-zero value to get the actual number of configurations in use. In (5.9), a beam still has to be inactive if none of the configurations in which it is feasible is used. However, we could now even use any constant $H > 0$, and especially

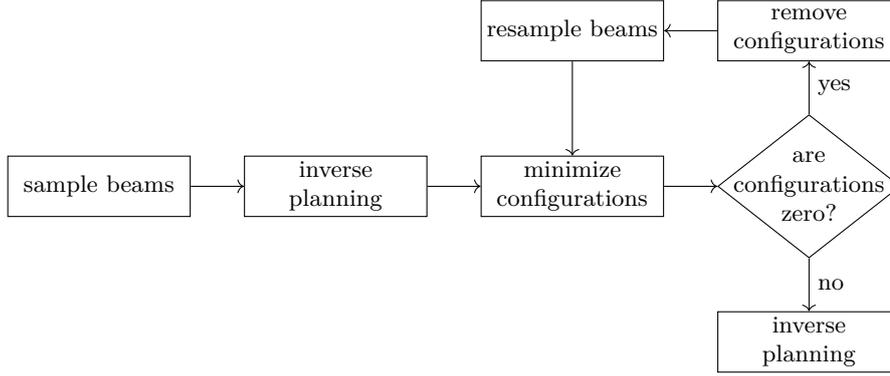


Figure 5.7: Workflow for the integrated minimization of the number of configurations.

$H = 1$, because it only scales the range of the values of the variables c_i which we will observe. The constraint (5.7) is not affected by the definition of c_i . Of course, we have to include all other constraints of the original inverse planning problem as well.

We propose to realize the integration of the relaxed minimization of configurations into the planning process as illustrated in Figure 5.7. In the first step, we randomly sample beams with respect to all candidate configurations. Afterwards, we solve the inverse planning problem. In the third step, we solve the relaxed problem which minimizes the number of configurations. If any c_i is zero afterwards, we fix these variables and resample all beams which have become infeasible. This is repeated until no further configuration can be removed. Finally, the inverse planning is solved again to establish optimality with respect to the OCO-OV. Note that due to the resampling, this last step can lead to a better OCO-OV than the initial treatment plan which considers all candidate configurations.

5.2.4 Experimental Setup

Similar to Section 4.4, we consider the three reference robot positions A, B, and C and initially employ 37 LIFT angles as the candidate configurations. We consider the greedy algorithm for the SCP model to approximate in a post-processing step the minimum number of configurations, which is actually required to deliver all treatment beams.

Furthermore, we study minimizing the number of configurations directly by extending planning with another LP as illustrated in Figure 5.7. We set the relaxation factor r for the original OCO-OV to 1.001. Note that due to the resampling of blocked beams due to removed configurations, the final OCO-OV is typically even smaller than the original one. To potentially avoid the effort of patient-specific optimization, we also study whether the most-frequently selected configurations for a set of training patients can be applied to other patients without further optimization.

5.3 Results

5.3.1 Optimized Positioning

In Section 5.1.2, we proposed four objectives to optimize the positioning of the GUIDE-robot, namely minimizing the number of blocked beams (BB), minimizing the overlap between GUIDE-robot projection and PTV projection (OP), minimizing the number of blocked beam nodes (BN), and maximizing the number of free beam nodes (FN). To solve the resulting optimization problems, we proposed the two strategies simulated annealing (SimAn) and segmental inverse robot construction (SIRC) in Section 5.1.3. We compare the results of these approaches in Table 5.1. For the SimAn strategy, we observe that the BB objective outperforms the other objectives. Only for patient 7, OP provides a higher improvement. While all approaches achieve better median OCO-OVs than reference robot B, improving robot A turns out to be challenging. No optimization approach is able to improve robot A’s result for patient 9. BB is at least able to improve the results for all other patients. OP leads to slightly worse results for two further patients, while BN and FN lead to worse results for five and six patients, respectively. For patient 6, employing FN even leads to a 30.4% higher median OCO-OV.

When considering the SIRC strategy in Table 5.1, the differences between the four proposed objectives are less clear. In fact, BB leads to the best result for only a single patient. OP provides the best results for four patients, BN for one, and FN for four. Nevertheless, FN leads to bad results for patients 7, 8, and 9 and BN fails for patient 6 compared to the other objectives. Therefore, only BB and OP seem to be rather consistent. Major differences between these two only appear for patients 2 and 9. For patient 9, we are again not able to achieve an improvement of robot A’s result, though the gap when employing OP or BN is clearly smaller than for SimAn.

If we compare the results of SimAn and SIRC obtained with BB, for patients 6 and 9 we observe rather clear outperformance by SimAn and for patients 2, 7, and 8 by SIRC. Note, however, that both are outperformed for patient 9 by SIRC with OP or BN. If we consider for each patient the best results obtained with any objective, SimAn is only advantageous for patient 6.

The observation that there is a substantial dependency on the patient under consideration is confirmed by Figure 5.8. It shows the OCO-OVs which we obtain by grid-based sampling of the search space and compares them to the OCO-OVs when only considering the grid points whose BB, OP, BN, or FN objective value belongs to the best 3-percentile. In Figure 5.8a, we mainly observe that the very high OCO-OVs do not appear when considering only the points with good objective value. However, neither the median nor the box limits are substantially better for this patient. In contrast, considering any objective improves the OCO-OV distribution in Figure 5.8b. Nevertheless, relying on BB would clearly be the best choice in this case. Figure 5.8c is an example for the case that one objective fails while the other three objectives perform well. Lastly, Figure 5.8d shows how the

Table 5.1: For each patient, the improvement of the median OCO-OV due to the proposed optimization approaches is shown. Results are shown relative to the reference robots A, B, and C. Remember that we solve a minimization problem for treatment planning and therefore better plans have a lower objective value. The objective functions yielding the best result for SimAn and SIRC are indicated bold for each patient.

	SimAn				SIRC				
	BB	OP	BN	FN	BB	OP	BN	FN	
Reference robot A	0	-8.0 %	0.4 %	8.5 %	1.5 %	-7.8 %	-8.5 %	-6.9 %	-10.8 %
	1	-14.8 %	-11.8 %	-4.7 %	-5.4 %	-15.5 %	-14.6 %	-11.5 %	-17.2 %
	2	-8.2 %	-0.5 %	8.4 %	4.1 %	-21.2 %	-6.1 %	-5.6 %	-19.2 %
	3	-11.8 %	-8.1 %	3.0 %	0.2 %	-12.7 %	-12.1 %	-9.9 %	-13.7 %
	4	-23.2 %	-21.9 %	-18.1 %	-13.4 %	-23.4 %	-24.1 %	-9.4 %	-1.6 %
	5	-16.7 %	-9.7 %	-3.3 %	-2.5 %	-18.6 %	-19.1 %	-11.4 %	-23.5 %
	6	-16.8 %	-11.2 %	-4.4 %	30.4 %	-10.3 %	-10.7 %	13.3 %	6.0 %
	7	-0.7 %	-4.0 %	7.2 %	19.6 %	-10.7 %	-11.0 %	-7.1 %	16.0 %
	8	-1.6 %	1.4 %	10.0 %	19.6 %	-6.3 %	-7.9 %	-6.9 %	22.0 %
	9	15.7 %	21.0 %	29.3 %	30.9 %	19.0 %	9.2 %	8.9 %	43.1 %
Reference robot B	0	-40.5 %	-35.1 %	-29.9 %	-34.4 %	-40.4 %	-40.9 %	-39.8 %	-42.3 %
	1	-34.7 %	-32.4 %	-27.0 %	-27.5 %	-35.3 %	-34.5 %	-32.1 %	-36.5 %
	2	-51.8 %	-47.7 %	-43.1 %	-45.3 %	-58.6 %	-50.7 %	-50.4 %	-57.5 %
	3	-46.7 %	-44.5 %	-37.8 %	-39.5 %	-47.2 %	-46.9 %	-45.6 %	-47.9 %
	4	-11.7 %	-10.2 %	-5.8 %	-0.4 %	-11.9 %	-12.7 %	4.2 %	13.2 %
	5	-38.3 %	-33.1 %	-28.3 %	-27.7 %	-39.7 %	-40.0 %	-34.3 %	-43.3 %
	6	-40.5 %	-36.4 %	-31.6 %	-6.6 %	-35.8 %	-36.1 %	-18.9 %	-24.2 %
	7	-21.2 %	-23.9 %	-15.0 %	-5.1 %	-29.2 %	-29.5 %	-26.4 %	-8.0 %
	8	-34.0 %	-32.0 %	-26.3 %	-19.8 %	-37.2 %	-38.3 %	-37.6 %	-18.2 %
	9	-34.6 %	-31.6 %	-26.9 %	-26.0 %	-32.7 %	-38.3 %	-38.4 %	-19.1 %
Reference robot C	0	-12.8 %	-4.8 %	2.8 %	-3.8 %	-12.6 %	-13.3 %	-11.8 %	-15.4 %
	1	-10.0 %	-6.8 %	0.7 %	-0.0 %	-10.8 %	-9.7 %	-6.5 %	-12.5 %
	2	-17.6 %	-10.7 %	-2.7 %	-6.6 %	-29.3 %	-15.7 %	-15.3 %	-27.5 %
	3	-15.0 %	-11.5 %	-0.7 %	-3.4 %	-15.8 %	-15.3 %	-13.1 %	-16.8 %
	4	-7.2 %	-5.6 %	-0.9 %	4.7 %	-7.4 %	-8.2 %	9.5 %	19.0 %
	5	-9.3 %	-1.7 %	5.3 %	6.2 %	-11.4 %	-11.9 %	-3.5 %	-16.7 %
	6	-14.8 %	-9.0 %	-2.1 %	33.7 %	-8.1 %	-8.5 %	16.1 %	8.6 %
	7	-11.7 %	-14.7 %	-4.7 %	6.3 %	-20.7 %	-20.9 %	-17.5 %	3.1 %
	8	-16.4 %	-13.8 %	-6.6 %	1.6 %	-20.4 %	-21.8 %	-20.9 %	3.6 %
	9	-13.0 %	-9.0 %	-2.7 %	-1.5 %	-10.5 %	-17.9 %	-18.1 %	7.6 %

best OCO-OV of the grid might not be achievable with any objective. In general, however, Figure 5.8 shows that considering BB seems to be a reasonable choice to obtain setups with appropriate OCO-OV.

When analyzing the solutions we obtained with SimAn and BB over all patients, we obtain the spatial distribution shown in Figure 5.9. Interestingly, the solutions are not close to reference robot A's position, although this robot setup performed well and it is quite close to one of the initial guesses for SimAn (compare the blue point in Figure 5.5). Instead, the solutions accumulate around a diagonal position, at which also robot C is located, and in the inferior box of the search space, in which we also find robot B. In the box to the patient's left, we only find a few

5 Optimizing a Robotic Setup for Ultrasound-Guided Radiosurgery

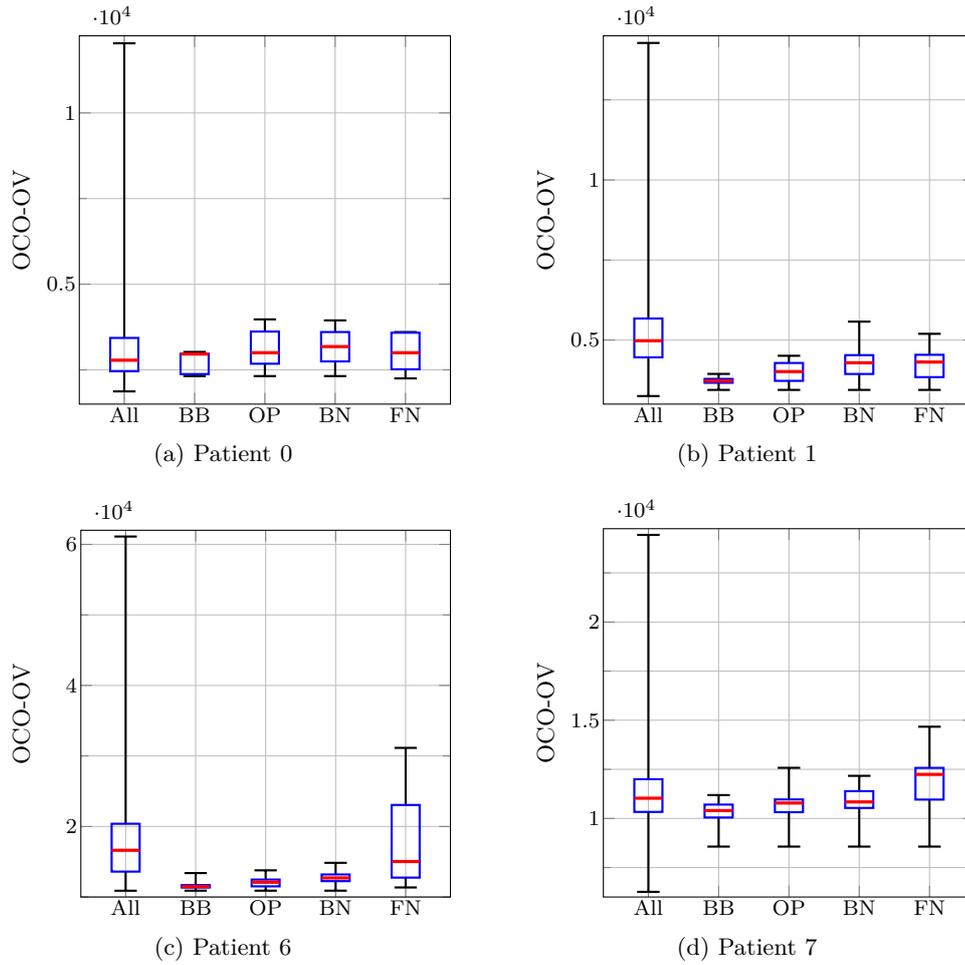


Figure 5.8: Results for a patient-wise grid-based sampling of the search space. The distribution of the OCO-OVs for all sampling points is compared to the OCO-OVs of only the sampling points belonging to best 3-percentile objective values w.r.t. BB, OP, BN, and FN. The grid-based sampling led to 1171, 1184, 1202, and 1121 feasible configurations for patients 0, 1, 6, and 7, respectively. Note that the vertical axes have different scales.

solutions. As mentioned before, having the robot in this box is rather cumbersome for realizing our defined transducer orientation. Also the optimization does not find many suitable setups in this area. In terms of the LIFT angle, the results show some structure in Figure 5.9c. This distribution is shown in more detail in Figure 5.10 for all solutions. A LIFT angle close to 0° is rarely chosen. In contrast, there is an interval from about -30° to -5° containing the most solutions. It is accompanied by peaks around -55° and 10° .

The preference to position the GUIDE-robot to the patient’s right rather than to the patient’s left is also reflected in the quality of the five initial guesses. Table 5.2

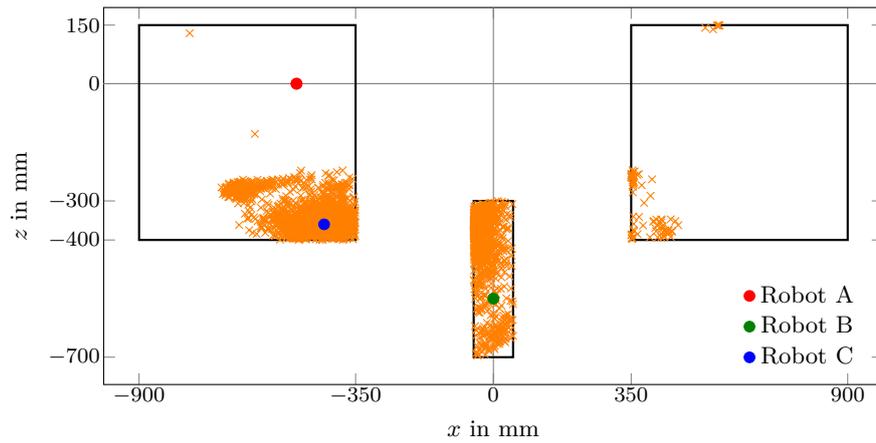
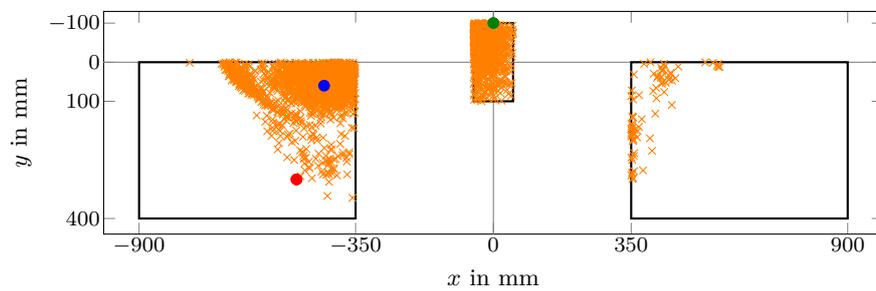
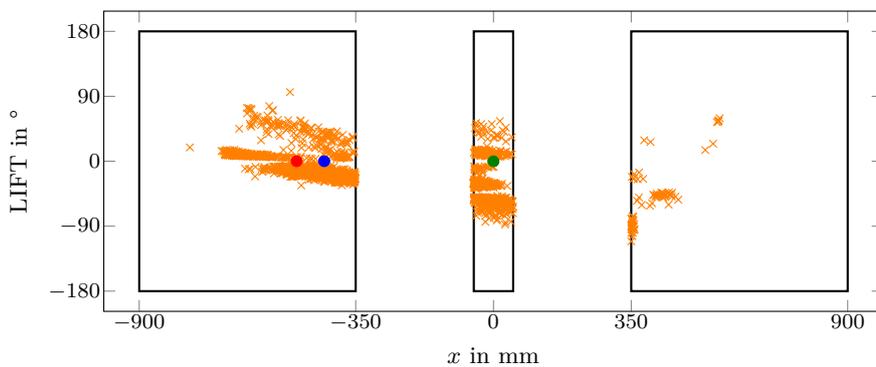
(a) xz -plane (anterior view)(b) xy -plane (inferior view)(c) x -LIFT-plane

Figure 5.9: Locations of all solutions for all patients obtained from SimAn with objective BB. The reference robot positions A, B, and C are indicated by a red, green, and blue point, respectively.

shows that the initial guess located far right leads to the best results for seven patients. For the remaining three patients, we obtain the best results for the initial guess close to the patient's right.

We already discussed in Chapter 4 that the differences between the three ref-

5 Optimizing a Robotic Setup for Ultrasound-Guided Radiosurgery

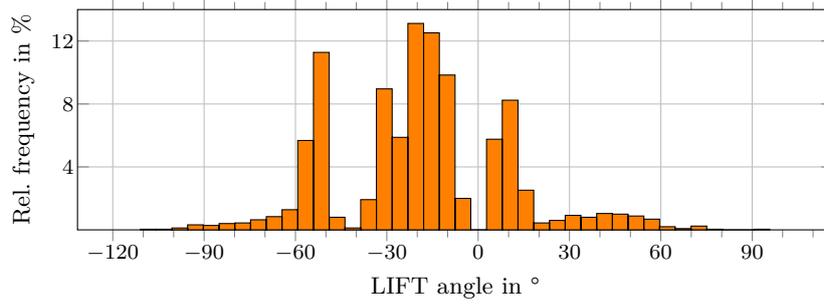


Figure 5.10: Binned histogram of the LIFT angles for all patients' solutions obtained from SimAn with objective BB.

Table 5.2: Mean and standard deviation for each patient's OCO-OV obtained from SimAn with objective BB. Results are shown depending on the position of the initial guess and the lowest mean for each patient is marked bold. Note that the terms left and right refer to the patient's RL-axis.

	Far right	Close right	Inferior	Close left	Far left
0	3271 ± 1087	3209 ± 1077	3329 ± 1147	3307 ± 1141	3353 ± 1168
1	4870 ± 881	4911 ± 858	5051 ± 877	4948 ± 902	5065 ± 926
2	2366 ± 336	2406 ± 293	2509 ± 320	2488 ± 447	2450 ± 342
3	3297 ± 744	3342 ± 763	3418 ± 776	3395 ± 771	3455 ± 827
4	9154 ± 1368	9225 ± 1256	9488 ± 1304	9547 ± 1379	9457 ± 1338
5	3007 ± 550	3035 ± 503	3157 ± 561	3198 ± 586	3065 ± 572
6	14840 ± 2443	14750 ± 2226	15033 ± 2119	15250 ± 2868	15632 ± 2498
7	10122 ± 1465	10185 ± 1416	10518 ± 1460	10604 ± 1515	10400 ± 1560
8	5076 ± 899	5178 ± 996	5362 ± 826	5520 ± 1252	5422 ± 1150
9	3120 ± 615	3080 ± 537	3326 ± 611	3216 ± 531	3319 ± 622

erence robot setup vanish when increasing the number of employed LIFT angles. In particular, employing three angles provides a major improvement of the plan quality compared to employing only one angle. While SIRC does not allow for a direct control of the LIFT angle, it is an explicit variable in SimAn and we could include multiple LIFT angles directly in the optimization process. However, to maintain the same dimension of the search space, we restrict ourselves in Table 5.3 to one LIFT angle variable and only additionally evaluate this angle rotated by -90° and 90° . This additional evaluation happens either already during optimization, labeled as BB_3 , or only afterwards on the final solution, labeled as BB_1 , i.e., the subscript i in BB_i indicates the number of LIFT angles evaluated during optimization for calculating the value of the BB objective. Only for patient 2, the optimization result BB_1 with evaluation of a single LIFT angle provides a slightly better mean OCO-OV than reference robot A with evaluation of three LIFT angles. Interestingly, patient 9 is still difficult for the optimizer and even using three LIFT angles cannot achieve the single-LIFT result of robot A. For the remaining patients, using three LIFT angles is always better than using only a single angle.

Table 5.3: Comparison of the OCO-OVs when using 1 or 3 LIFT angles for each patient. For 1 LIFT angle, we consider reference robot A and optimization of BB with SimAn. For 3 LIFT angles, we consider reference robot A and two optimization schemes with SimAn. BB_1 evaluates the optimal solution obtained for 1 LIFT angle with two additional angles rotated by $\pm 90^\circ$, while BB_3 considers the additional angles already during optimization. The lowest mean for each patient is marked bold.

	1 LIFT		3 LIFTs		
	Robot A	BB_1	Robot A	BB_1	BB_3
0	3630 \pm 1222	3294 \pm 1117	3085 \pm 1098	3005 \pm 1061	2981 \pm 1058
1	6023 \pm 1049	4969 \pm 885	4616 \pm 779	4573 \pm 866	4520 \pm 861
2	2784 \pm 479	2444 \pm 353	2450 \pm 412	2249 \pm 290	2222 \pm 286
3	3950 \pm 762	3381 \pm 772	2918 \pm 625	3068 \pm 692	3018 \pm 630
4	12352 \pm 1439	9374 \pm 1328	8550 \pm 1045	8907 \pm 1210	8587 \pm 1167
5	3725 \pm 615	3093 \pm 556	2988 \pm 574	2808 \pm 474	2784 \pm 466
6	18114 \pm 3118	15101 \pm 2446	14575 \pm 2531	14862 \pm 2294	14404 \pm 2270
7	10250 \pm 1279	10366 \pm 1484	8889 \pm 1154	9345 \pm 1226	8853 \pm 984
8	5363 \pm 1095	5312 \pm 1041	4620 \pm 875	4807 \pm 926	4635 \pm 849
9	2789 \pm 581	3212 \pm 589	2521 \pm 554	2881 \pm 513	2800 \pm 479

Table 5.4: Means and standard deviations for the number of configurations, which are actually required to deliver all candidate or treatment beams, obtained from the greedy algorithm for the SCP.

	Robot A	Robot B	Robot C
Candidate beams	11.1 \pm 2.1	9.0 \pm 2.4	7.8 \pm 2.4
Treatment beams	5.2 \pm 1.2	4.0 \pm 1.1	3.3 \pm 1.3

Further, for four patients neither BB_1 nor BB_3 with three LIFT angles can achieve the results of robot A with three LIFT angles. Considering three angles already during optimization via BB_3 is always better than considering only the central angle during optimization via BB_1 .

5.3.2 Efficient Configurations

To analyze the efficient selection of GUIDE-robot configurations, we consider 37 equidistant LIFT angles from -90° to 90° defining the candidate configurations for each of the three reference robot setups A, B, and C. Considering the candidate configurations and candidate beams as an SCP and solving it with the greedy algorithm described in Section 5.2.2, however, we see from Table 5.4 that on average already 11.1, 9.0, and 7.8 configurations are sufficient to realize all the candidate beams for robots A, B, and C, respectively. For the much smaller set of actual treatment beams, even selecting 5.2, 4.0, and 3.3 of the 37 candidate configurations is sufficient on average.

When integrating the iterative minimization of the number configuration in

5 Optimizing a Robotic Setup for Ultrasound-Guided Radiosurgery

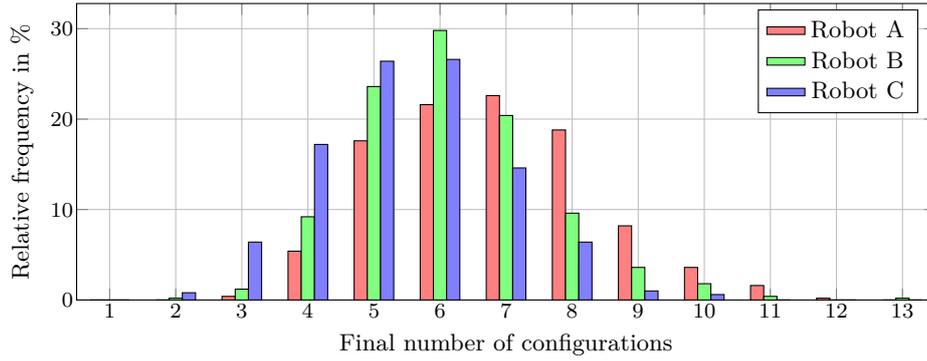


Figure 5.11: Histogram of the number of remaining configurations for robots A, B, and C after optimization.

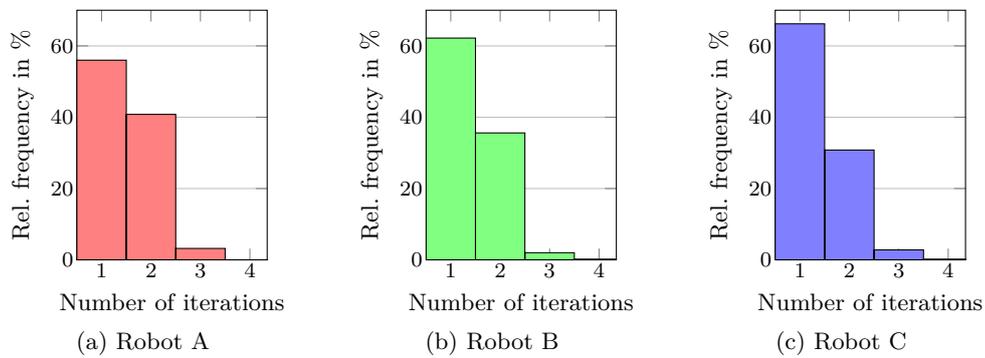


Figure 5.12: Number of iterations after which the final number of configurations is reached. A number of two means that the second iteration successfully removed at least one configuration, but the third iteration was not able to remove any further configurations.

the plan optimization workflow, as introduced in Section 5.2.3, we obtain the results shown in Figure 5.11. The modes of the distributions for the number of remaining configurations are 7, 6, and 6 for robots A, B, and C, respectively. The minimum observed numbers are 3, 2, and 2 and the maximum numbers 12, 13, and 10. As Figure 5.12 shows, the iterative process mostly requires only one or two iterations to obtain its final solution, i.e., the subsequent iteration is not able to remove any further configuration. Only for robots B and C, it happens rarely that the fourth iteration still provides an improvement. If we compare the results of the iterative minimization to the previous results by simply solving the SCP, as presented in Table 5.5, we see that about two configurations more remain on average. However, if we additionally post-process the iterative result as an SCP, the differences vanish.

Figure 5.13 shows which LIFT angles remain how frequently after iterative minimization. For robot A, there are three angles which appear in more than

Table 5.5: Comparison of the means and standard deviations for the number of configurations obtained by heuristically solving the SCP in a post-processing step, by iterative minimization, and by iterative minimization and subsequent post-processing.

	SCP	Iter	Iter+SCP
Robot A	5.2 ± 1.2	6.8 ± 1.6	4.8 ± 1.0
Robot B	4.0 ± 1.1	6.1 ± 1.5	3.8 ± 1.1
Robot C	3.3 ± 1.3	5.5 ± 1.4	3.4 ± 1.1

50 % of the solutions. In contrast, only one angle appears in more than half of the solutions for robot B. Interestingly, the angle -90° appears as a rather high peak for all three robots. An angle of 0° , which served as a reference for optimizing the robot positioning in the previous sub-section, only seems to be a suitable choice for robot A. Altogether, the three distributions do not have completely different shapes but show some similarities in the number and location of more-frequently selected angles.

If we apply the 1, 3, and 5 most-frequently selected LIFT angles for patient 0 to 5 to all patients, we obtain the results presented in Table 5.6. Considering patients 0 to 5 as a training set and patients 6 to 9 as a validation set, we see that a single LIFT angle provides clearly worse results compared to the patient-specific optimization of the angles. An interesting exception is patient 0 in case of robot A. Already the most-frequently selected angle is sufficient to provide the same treatment plan quality. Also for patient 9, which belongs to the validation set, we observe with 9.7 % a comparably small deviation from the patient-specific optimization result. For the other combinations of patients and robots, however, the increase in the OCO-OV is 16.8 % to 48.3 % when only applying one LIFT angle. For robots A and B, selecting five angles is necessary to obtain an increase smaller than 6 % for all patients. For robot C, this is already achieved by selecting three angles.

5.4 Discussion

In this chapter, we analyzed and optimized the setup of the GUIDE-robot. We focused on two aspects separately, the general positioning of the GUIDE-robot and the selection of an efficient set of configurations given a fixed positioning. For optimizing the positioning of the GUIDE-robot, we proposed and evaluated four objective functions and two strategies to solve the resulting optimization problems approximately. Generally, we were able to improve our three manually selected reference setups by these approaches (Table 5.1). An exception is patient 9, for which robot A’s plan quality could not be reached by any optimization approach. For the SimAn strategy, we found that the BB objective counting the number of beams blocked by the GUIDE-robot should be used. In contrast, the results for

5 Optimizing a Robotic Setup for Ultrasound-Guided Radiosurgery

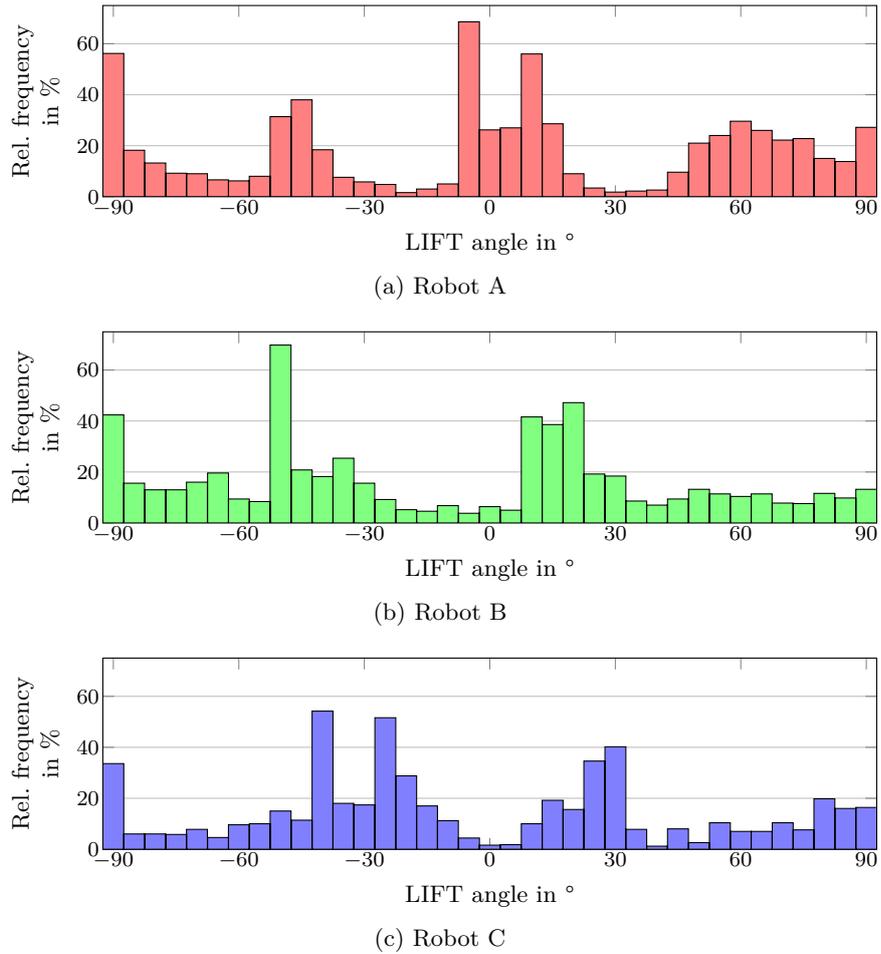


Figure 5.13: Histogram of the LIFT angles remaining after optimization. Note that the sum of the bars is greater than 100 % because usually multiple configurations remain.

the SIRC strategy showed no clear trend towards a single objective function. The BB objective and OP, which evaluates the overlap of projected GUIDE-robot and PTV, are most consistent and perform similarly for most patients.

A grid-based sampling and evaluation confirmed that generally BB is a suitable choice (Figure 5.8). However, an ideal objective function would perfectly correlate with the treatment plan quality. This is clearly neither achieved by our proposed methods nor a realistic goal, because of the complexity of the planning problem which is not accessible by only considering geometric relations between candidate beams and the GUIDE-robot. However, we showed that it is feasible to fully automatically find suitable positionings for the GUIDE-robot in terms of the achievable treatment plan quality. If for this task solving an optimization problem is not desired, the structure of the solutions which we obtained might be

Table 5.6: Mean change of the median OCO-OV when applying the 1, 3, or 5 most-frequently selected LIFT angles for patients 0 to 5 directly to all patients without patient-specific optimization. The most-frequent angles are specific to the robots A, B, and C. Note that smaller OCO-OVs are better.

		Robot A			Robot B			Robot C		
		1	3	5	1	3	5	1	3	5
Training	0	-0.8 %	-1.2 %	0.1 %	32.6 %	1.5 %	0.7 %	27.0 %	0.1 %	0.0 %
	1	34.2 %	4.2 %	2.3 %	33.0 %	3.5 %	-0.8 %	19.5 %	0.8 %	0.4 %
	2	16.8 %	3.6 %	2.1 %	32.5 %	8.6 %	3.5 %	37.7 %	5.7 %	4.9 %
	3	28.3 %	2.6 %	1.6 %	34.3 %	3.1 %	0.1 %	25.8 %	0.1 %	0.0 %
	4	48.3 %	8.0 %	2.0 %	24.9 %	9.2 %	0.5 %	17.4 %	0.4 %	-0.5 %
	5	37.8 %	6.9 %	0.4 %	33.3 %	4.2 %	1.1 %	27.2 %	0.9 %	0.9 %
Validation	6	35.7 %	13.8 %	2.7 %	34.1 %	8.9 %	3.8 %	27.4 %	1.9 %	1.6 %
	7	30.6 %	6.7 %	4.7 %	30.4 %	11.2 %	1.5 %	31.2 %	1.3 %	1.1 %
	8	47.9 %	12.6 %	5.3 %	41.5 %	8.8 %	1.1 %	34.6 %	0.7 %	1.4 %
	9	9.7 %	5.6 %	0.4 %	39.3 %	9.2 %	1.7 %	34.0 %	0.9 %	0.9 %

a helpful guide for manual positioning, because they accumulated in some areas around the patients (Figure 5.9). Note, however, that these results are partly biased by our defined orientation of the ultrasound transducer. If the transducer is rotated around the AP-axis, the preferable locations for the robot base should be expected to rotate around the patient as well. This should also be considered for the initial guess when employing the SimAn approach (Table 5.2). Our results confirmed furthermore that the choice of the LIFT angle is important. Interestingly, the optimizer frequently selected them from a rather small range of angles (Figure 5.10).

Eventually, we confirmed the previous observation that considering more LIFT angles during the treatment is in general advantageous (Table 5.3). In this point, there is an essential difference between the SimAn and the SIRC strategy. SIRC appeared to explore the search space better because it provided satisfying results with multiple objectives functions. By design, however, it does not allow to control the GUIDE-robot explicitly, while evaluation and optimization of arbitrarily many LIFT angles can easily be realized in SimAn. However, this increases both the computational effort for evaluation of an objective function and the dimensionality of the search space. For this reason, we also discussed in this chapter separately how to select efficient configurations of the GUIDE-robot given a fixed positioning for the robot base.

We proposed two approaches to determine an efficient subset of a set of candidate configurations for the GUIDE-robot. The SCP formulation of this problem, which we proposed, is very inexpensive when solved by a greedy algorithm. We applied it after treatment plan optimization to obtain a small subset of configurations which is actually required to deliver all treatment beams. From 37 candidate configurations, only about three to five were actually required on average, depending on the reference robot setup under consideration (Table 5.4). This shows that considering many configurations for planning is also practically meaningful, be-

cause eventually only a small number of them is actually required. Therefore, the synchronization problem between the GUIDE-robot and the beam delivery, which we will discuss in the next chapter in detail, remains tractable and realistic while we can fully exploit the kinematic redundancy for minimizing beam blocking and plan deterioration.

Additionally, we proposed an extension of the treatment planning LP which iteratively tries to minimize the number of configurations. We relaxed the exact MIP model to an easier to handle LP and obtained numbers of configurations which were slightly higher than for the SCP approach (Figure 5.11). This shows that our simple relaxation leads to a relevant gap to the optimal MIP solution. In particular, because we can further reduce the number of configurations when applying the SCP approach subsequently (Table 5.5). Nevertheless, the iterative LP approach yields some further insights in the nature of the problem. In contrast to the SCP approach which is a post-processing step, the iterative LP approach directly interacts with the planning problem. Hence, its solutions indicate which LIFT angles are generally preferable for a GUIDE-robot setup (Figure 5.13). These angles were even applicable to new patients to achieve reasonable plan qualities without repeating the whole optimization for these patients (Table 5.6).

In summary, we saw in this chapter that it is feasible to automatically determine suitable GUIDE-robot setups. Finding multiple suitable configurations can either be directly included in this procedure or realized separately. We saw that even when we initially considered many candidate configurations in order to block as few beam directions as possible for plan optimization, we can eventually reduce them easily to a reasonable number without any deterioration of the plan quality.

6 Synchronizing Robot Motion and Beam Delivery in Ultrasound-Guided Radiosurgery

In the two preceding chapters, we showed that considering multiple configurations of the GUIDE-robot during a treatment can severely decrease the impact of blocked beam directions on the treatment plan quality. However, employing multiple configurations requires a synchronization between the GUIDE-robot and the beam delivery, as outlined in Section 1.2 as the second research question of this thesis. A beam can only be delivered if the GUIDE-robot is in one of its non-blocking configurations. A simple approach would be to start in one configuration of the GUIDE-robot, deliver all feasible beams, move to the next configuration, deliver all feasible and not yet delivered beams, and so on. However, this approach will lead to rather inefficient trajectories for the LINAC-robot, which in turn increase the total treatment time. Furthermore, this approach results in some remaining questions. Two of them are, in which order we should use the configurations and whether we really should deliver a feasible beam directly or whether it would be more time-efficient to wait for another configuration.

In this chapter, we develop models to describe and optimize the synchronization problem between the GUIDE-robot's configuration changes and the beam delivery of the LINAC-robot. For this purpose, we start in Section 6.1 by describing a model to determine the most efficient trajectory for a beam delivery without the GUIDE-robot. This problem is of course not new but relevant for any CyberKnife treatment. We describe how the problem can be modeled as a traveling salesman problem (TSP), which conceptually is also employed in the current generation of the CyberKnife system [126]. However, in Section 6.2, we show how our TSP-based approach can be directly extended to a generalized traveling salesman problem (GTSP) in order to model the synchronization problem between the configuration changes by the GUIDE-robot and the beam delivery under the assumption that the GUIDE-robot does not move during active beam delivery. This assumption reduces the effort for ensuring safety, because beam blocking has only to be verified for a small finite set of static robot configurations. We discuss and evaluate different approaches to solve the GTSP model, which is known to be practically hard to solve. Parts of this chapter have been published in a journal [230].

Throughout this chapter, we assume that the elements of a set A have an arbitrary enumeration, i.e., we write $A = \{a_1, \dots, a_{|A|}\}$ without implying an ordering

relation. Depending on the context and convenience, we refer to an element of the set either as $a \in A$ or as $a_i \in A$ and implicitly assume $i \in \{1, \dots, |A|\}$ for the latter. In contrast, we refer to ordered tuples either as (a, b, \dots) where $a, b, \dots \in A$ or as (a_i, a_j, \dots) where $a_i, a_j, \dots \in A$.

6.1 Model for Efficient Beam Delivery

In order to deliver a treatment beam, the LINAC-robot has to position its end-effector, and thereby the attached LINAC, accordingly. This has to be done for each treatment beam in some sequence. While the order of the beams does not affect the outcome of the treatment, it affects the idle times of the LINAC due to the required robot motion in between.

In this section, we describe a mathematical model, the TSP, and describe how we apply it to model the beam sequencing problem of the LINAC-robot. We also discuss different approaches for solving TSPs. While the first versions of the CyberKnife only used predefined paths through the beam nodes [45], the current system also solves a TSP-like problem to obtain efficient paths [126]. In contrast to our simplified model, the actual implementation considers further obstacles and the requirement that the LINAC-robot must not block the stereoscopic X-ray imaging for too long time intervals.

6.1.1 The Traveling Salesman Problem

Informally, the TSP can be illustrated by a salesman who is going to visit customers in different cities. He would like to know the shortest route which starts at his home town, takes him to each of the other cities exactly once, and finally takes him back home. Finding an answer to this question requires to define the term *shortest* in more detail. We could select, for example, the distance to travel, the time traveling takes, or the resulting financial costs. This decision, in turn, leads to new details which have to be considered. For example, the bee-line distance from city A to city B is the same as from B to A. However, this symmetry might not hold if we consider the time for traveling instead, for example if the speed limits on a highway are different for each direction due to construction work. This property is distinguished in the literature by the terms symmetric traveling salesman problem (STSP) and asymmetric traveling salesman problem (ATSP). Furthermore, the bee-line distance fulfills the triangle inequality as going from A to B via C is never shorter than going from A to B directly. This, again, does not have to be true in general, for example if we consider the time or costs of different train connections. Fulfillment of the triangle inequality results in another distinction of different subtypes of the TSP and there are several more subtypes and variations in the literature [96]. They have been applied for modeling and optimization in various fields.

Mathematically, the TSP is defined on a graph (V, E) . V is a set of vertices, corresponding to the cities to visit in our illustrative example. The vertices are

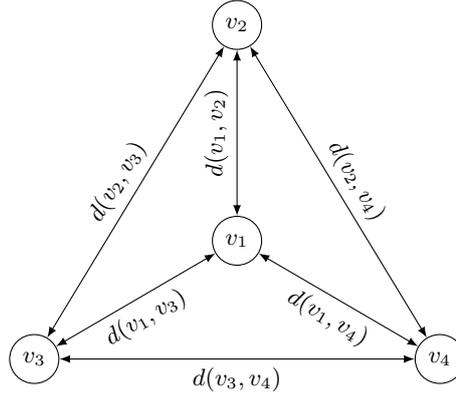


Figure 6.1: Illustration of an STSP with four vertices.

connected by a set of edges E . If we consider a fully-connected graph, E contains an edge for every pair of vertices $v_i, v_j \in V$, i.e.,

$$E = \{(v_i, v_j) : v_i, v_j \in V\}. \quad (6.1)$$

More generally, however, there can also be pairs of vertices without direct connections. Furthermore, edges can be directed or undirected, i.e., (v_i, v_j) and (v_j, v_i) are equivalent for the latter case. For a TSP, each edge is associated with a cost $d(v_i, v_j)$, as illustrated in Figure 6.1, which is also called weight or distance. Often, only non-negative and integer-valued costs are used, although this is not a formal requirement.

On such a graph, we have typically multiple possibilities to get from one vertex to another vertex. These so-called *walks* between two vertices can be uniquely described by providing the sequence of vertices, which are visited. In the simplest case, a walk only consists of the two vertices themselves. Note that a walk, in contrast to a directed walk, does not have a defined direction in which the sequence is interpreted. If a walk contains only distinct vertices, then we call it a *path*. A walk for which only the first and last vertex are equal is a *cycle*. In the TSP, however, we ask for the shortest *tour*. A tour is a cycle containing each vertex of the graph. The costs of a tour are the summed costs of all edges, which connect the sequentially visited vertices of the tour. When solving a TSP, the task is to find a tour with minimum costs.

If a fully-connected graph has N vertices, there are actually not $N!$ different tours in terms of solving the related TSP, but only $(N - 1)!$ in case of asymmetric costs because a tour does not have a defined start. In case of symmetric costs, the number of non-equivalent tours further reduces to $(N - 1)!/2$ because the direction is not relevant. Therefore, for an exemplary fully-connected graph with $V = \{v_1, v_2, v_3\}$ and symmetric costs, the tours (v_1, v_2, v_3, v_1) , (v_2, v_3, v_1, v_2) , and (v_1, v_3, v_2, v_1) are all equivalent in terms of solving the related TSP. In our illustrative example, this means that for the optimal tour it is neither relevant which of

the cities is the salesman's home town nor in which direction the salesman travels the tour.

6.1.2 Modeling the Beam Delivery

In the beam sequencing problem for the LINAC-robot, we identify the set of treatment beams B as the set of vertices of our graph for the TSP model, i.e., $V := B$. In this thesis, we assume that each vertex is connected to each other vertex. In practice, however, there might be collisions with the patient or other obstacles in the treatment room for some pairs of beams. We could account for these by assigning a very high cost to the related edges in the graph in order to enforce that they cannot be part of an optimal tour. We define the costs $d(b_i, b_j)$, $b_i, b_j \in B$, of an edge as the time required for moving the LINAC-robot and assume that this motion results in symmetric costs, i.e., an STSP.

A major discrepancy between the described TSP model and the real scenario is the fact that a tour, by definition, is a cycle. This implies that the LINAC-robot eventually moves back to the pose for delivery of the first beam again. This is unnecessary because this beam has already been delivered. A more appropriate model would be a so-called shortest Hamiltonian path problem (HPP). In this problem, we search for a Hamiltonian path, which is a path containing each vertex of the graph, i.e., a walk containing each vertex exactly once. Note that this problem's solution provides a defined first and last vertex, i.e., treatment beam.

An HPP can easily be transformed into a TSP and after solving this TSP its solution can easily be transformed into the solution of the original HPP [96]. The transformation into a TSP requires to add another vertex to the graph, which is connected to every other vertex by zero-cost edges, in case of only non-negative costs. It can be interpreted as a kind of joker which allows to skip the costs for moving between one pair of vertices. In the resulting tour, the neighboring vertices of the additional vertex are the first and last vertex of the corresponding optimal Hamiltonian path. As an example, consider a fully connected graph $G = (V, E)$ with symmetric non-negative costs and vertices $V := \{v_1, v_2, v_3\}$. To solve the HPP on this graph, we solve the TSP on the fully connected graph $\hat{G} = (\hat{V}, \hat{E})$ with $\hat{V} = V \cup \{z\}$ and the costs $d(v, z) = 0$ for edges between the additional vertex z and any $v \in V$. If the optimal tour is (v_3, v_1, z, v_2, v_3) , then we obtain the solution for the HPP through opening the cycle by removing z , i.e., the solution is (v_2, v_3, v_1) .

6.1.3 Solving TSPs

Due to the described transformation of an HPP into a TSP, we can employ any TSP solver for HPPs as well. Therefore, we only discuss the more common task of solving TSPs in the following. However, we will only briefly cover two selected approaches in the following.

Integer Linear Programming

The TSP is an NP-hard combinatorial optimization problem. One approach for solving is to formulate it as an ILP. One prominent formulation for the STSP is the *double index* or *subtour elimination* formulation [96]. It introduces binary variables x_e , $e \in E$, which are one iff the related edge is part of the solution. Therefore, the objective function to be minimized sums the costs of an edge e multiplied with x_e over all $e \in E$. Only two main types of constraints are required to enforce that solutions are valid tours. First, for each vertex, there have to be exactly two edges in the solution which are connected to that vertex. Second, a solution has to actually be a single cycle and not a set of multiple disjoint cycles covering all vertices. The latter aspect led to the name *subtour elimination* of this formulation.

In practice, the binary constraints on the variables x_e have to be relaxed in order to apply an LP solver with efficient branch-and-bound or cutting-plane methods. Therefore, the properties of the relaxation determine how suitable a formulation is for practical solving. More details including different preprocessing, branch-and-bound, and cutting-plane techniques can be found, for example, in Gutin et al. [96].

Lin-Kernighan Heuristic

One specific heuristic approach for solving TSPs is based on so-called k -opt improvements [163]. Given a valid tour for the problem, we select k edges and replace them by k shorter edges. In 2-opt, the selection of two edges uniquely defines the two replacing edges, because we have to maintain a valid tour. We repeatedly check for all pairs of edges whether replacement would provide a better tour. If no improvement is possible anymore, we found a 2-optimal solution. The general approach k -opt uses the same concept, but the number of possibilities to select the k edges to be replaced follows a binomial coefficient and the number of combinations for selecting k replacing edges increases exponentially. Therefore, k is typically only chosen as 2 to 4 to maintain a reasonable computational effort. A strategy building upon the basic k -opt strategy is the Lin-Kernighan heuristic [164], which adaptively chooses k and introduces additional rules for the selection of edges to find efficient improvements in short time.

Several modifications and extensions have been proposed for the Lin-Kernighan heuristic [6, 270]. One particular variant of its steps and rules is the Lin-Kernighan-Helsgaun heuristic [97, 98].

6.2 Model for Synchronized Robot Motion

We now have to extend the model for the beam delivery to also account for the required configuration changes by the GUIDE-robot. For this purpose, we first introduce the GTSP, which is a generalization of the TSP. Afterwards, we show

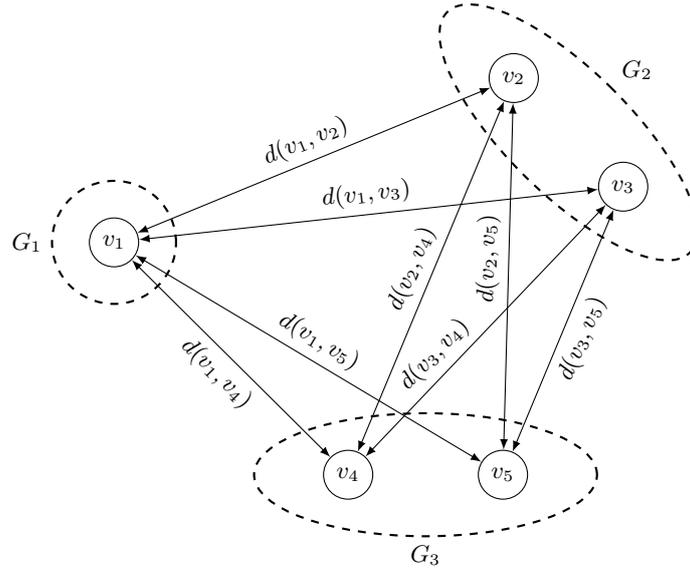


Figure 6.2: Illustration of a GTSP with five vertices v_1, \dots, v_5 and three groups G_1, G_2, G_3 . Vertices belonging to the same group are never connected.

how this model can be used to describe and optimize the synchronization problem between the two robots for beam delivery and ultrasound guidance.

6.2.1 The Generalized Traveling Salesman Problem

In the GTSP, we do no longer want to visit every vertex of our graph. Instead, we define K groups $\{G_1, \dots, G_K\}$, $K \geq 3$, assign each vertex $v \in V$ to exactly one group, and search for the shortest cycle which contains exactly one vertex of each group. Because a valid solution only visits exactly one vertex of each group, there are no edges in the graph between vertices of the same group. Figure 6.2 shows an example for this problem. In fact, a TSP is a special instance of a GTSP in which each group only contains one vertex. Thereby, groups and vertices become equivalent.

Another view on the GTSP is that it consists of two coupled problems. Firstly, we have to select a tour w.r.t. to the groups. For K groups, there are $(K-1)!/2$ tours as for STSPs. Secondly, we have to select one vertex from each group. The number of combinations is the product of the sizes of the groups, i.e., there are $\prod_{i=1}^K |G_i|$ combinations. If the N vertices are equally distributed among the groups, i.e., each group contains N/K vertices, there are $(K-1)!/2 \cdot (N/K)^K$ feasible solutions for the GTSP.

Note that there exist different names for this problem in the literature. They include international TSP, group TSP, and set TSP. Further, our definition of the GTSP is sometimes referred to as the equality GTSP. In this case, GTSP refers to the related problem that at least one vertex of each group has to be visited,

instead.

6.2.2 Modeling the Robot Synchronization

Given a set of GUIDE-robot configurations C which we employ for the treatment, we have to consider the sequencing of combinations (b, c) of beams $b \in B$ and configurations $c \in C$. However, not all beams are feasible in all configurations. Therefore, we define the set of configurations in which a beam $b \in B$ is feasible as

$$F(b) = \{c : b \text{ feasible in } c \in C\}. \quad (6.2)$$

The set of vertices in our model can then be defined as

$$V := \{(b, c) : b \in B, c \in F(b)\}. \quad (6.3)$$

We only have to deliver each beam once. Therefore, all vertices representing the same beam b_i form the group G_i of the GTSP, i.e.,

$$G_i := \{(b_i, c) : b_i \in B, c \in F(b_i)\}, \quad i = 1, \dots, |B|. \quad (6.4)$$

By definition, vertices of the same group are not connected by an edge. This implies, on application level, that we cannot change the GUIDE-robot's configuration without also moving the LINAC-robot for delivery of the next beam. However, this is not limiting due to our initially formulated assumption that the GUIDE-robot never moves during beam delivery. Therefore, it would not be meaningful to move the GUIDE-robot without also moving the LINAC-robot.

Both robots can move simultaneously, because we assume that there is no active beam delivery. Thus only the robot which needs longer for its motion contributes to the costs of an edge. This leads us to the costs

$$\begin{aligned} d((b_i, c_k), (b_j, c_l)) &= \max \{d(b_i, b_j), d(c_k, c_l)\}, \\ &\text{for } b_i, b_j \in B, b_i \neq b_j, c_k \in F(b_i), c_l \in F(b_j) \end{aligned} \quad (6.5)$$

for edges between vertices of different groups. As in the TSP model for the beam delivery, $d(b_i, b_j)$ is the symmetric time needed for moving the LINAC-robot between delivery of two different beams b_i and b_j . In addition, we consider the symmetric time $d(c_k, c_l)$ to change between the GUIDE-robot's configurations c_k and c_l . In case of $c_k = c_l$, we have $d(c_k, c_k) = 0$ and the LINAC-robot's motion defines the costs of the edge.

In the model for beam delivery, we discussed that the HPP is closer to our actual problem than the TSP. We have the same issue when considering the actual task of the GTSP as we are rather interested in having a defined first and last group, i.e., treatment beam to deliver, and not a closed cycle. We fix this, as before, by adding another group with a single vertex, which is connected to all other vertices with zero costs, solving the resulting GTSP, and deriving our actual solution from this GTSP's solution.

6.2.3 Solving GTSPs

In the following, we briefly discuss some selected approaches to practically solve GTSPs.

Integer Linear Programming

Formulating a GTSP as an ILP is similar to the formulation of the TSP in terms of introducing binary variables x_e representing the edges $e \in E$ and formulating the objective function. However, we need to introduce one additional binary variable y_v for each vertex $v \in V$, which is one iff v is the vertex of its group which is actually visited. The sum over all y_v of the vertices of a group has to be one, i.e., exactly one of the group's vertices is visited. As for the TSP, we need also constraints which ensure that the solution is a single cycle. There are branch-and-bound and cutting-plane strategies exploiting the ILP formulation of the GTSP [63], but due to the computational effort they are still limited to rather small instances of the problem.

Transformation into a TSP

In principle, any GTSP can be converted to an STSP in a straight-forward manner and then be solved by any standard solver for TSPs. Noon et al. [189] describe a method to first convert the GTSP to a clustered traveling salesman problem (CTSP), then to an ATSP, and eventually to an STSP. The CTSP has some similarities to the GTSP extension of the TSP. Each vertex is assigned to exactly one so-called cluster. However, we search for an optimal tour visiting actually each vertex of each cluster exactly once. The special constraint is that when visiting a vertex of a cluster, we have to visit all other vertices of that cluster as well before we are allowed to visit a vertex of any other cluster. In terms of our initial illustrative example, we can consider the clusters as cities and the vertices as the customers, which are located within a city. When entering a city, the salesman has to visit all the customers before moving on to the next city. He is not allowed to partition them into multiple visits. Although it seems intuitive for many applications, it is formally not required that the costs of edges within a cluster are small compared to the costs of edges connecting different clusters.

Transforming GTSPs to TSPs and applying standard solvers seems to be a satisfying strategy. However, the transformation has been shown to produce TSPs with structures which are very unsuitable for the usually very efficient heuristics of exact TSP solvers [116, 146]. The transformation involves adding many edges with very high costs to enforce that the TSP solution actually fulfills the constraints of the GTSP after back transformation. This requires to compute exact solutions for the TSP because near-optimal solutions from heuristic TSP solvers can correspond to infeasible GTSP solutions. Therefore, several strategies and heuristics have been developed specifically for GTSPs [24, 95, 116, 250]. However, Helsgaun [99] showed recently that the Lin-Kernighan-Helsgaun heuristic for TSPs can perform

very well on GTSPs transformed to ATSPs. No modification of the heuristics is necessary, but some additional post-processing steps are performed after back transformation of the ATSP solution.

Decoupling Approximation

Because the GTSP can be interpreted as two coupled problems, a trivial simplification is to fix the solution to one of them and only optimize the other one. Depending on the structure of an instance, this simplification provides arbitrarily bad results, of course. We can only expect reasonable results, if one of the two problems substantially dominates the structure of the coupled problem. Only considering the optimal tour w.r.t. the groups and fixing the selection of vertices, we remain with a conventional TSP. In contrast, when fixing the tour and only considering the problem of selecting one vertex for each group, we obtain a shortest cycle problem on a graph with special structure. The structure is illustrated in Figure 6.3. Due to the fixed tour of groups, the vertices of a group are only connected to the vertices of exactly two other groups. We need unidirectional edges to reflect visiting exactly one node of each group, although the shortest cycle is an equivalent solution in both directions. In case we are rather interested in having a defined first and last group to visit instead of a closed cycle, the structure becomes even simpler. We can add an artificial vertex connected to the first group’s vertices with zero costs and another vertex to which the last group’s vertices are connected with zero costs. On the resulting graph, which is exemplarily illustrated in Figure 6.4, we only have to search for the shortest path between the two artificial vertices to obtain the optimal selection of vertices w.r.t. the fixed sequence of groups.

6.3 Experimental Setup

6.3.1 Robot Models and Parameters

Following the previous chapters, we consider for the GUIDE-robot an LBR iiwa with seven joints and the three reference setups introduced in Section 4.2. We consider employing three, five, or seven equidistant LIFT angles during treatment, where the smallest angle is always -90° and the largest 90° . Due to the results of Section 5.3.2, we find that these numbers of configurations represent reasonable choices. For the LINAC-robot, we consider a KR 300 R2500 ultra (KUKA Roboter, Germany), which is a 6-DOF robot arm. We consider the time for robot motion as the costs $d(b_i, b_j)$ and $d(c_k, c_l)$ in our models. In a safety-critical environment, like radiosurgery with the robots operating close to a patient, we cannot realistically assume robot motion at maximum speed. The actual CyberKnife operates its robot arm with at most 4% of its maximum speed w.r.t. the end-effector motion in Cartesian space [175, 230]. However, our GUIDE-robot does not move its end-effector. Therefore, we consistently consider limiting the joint

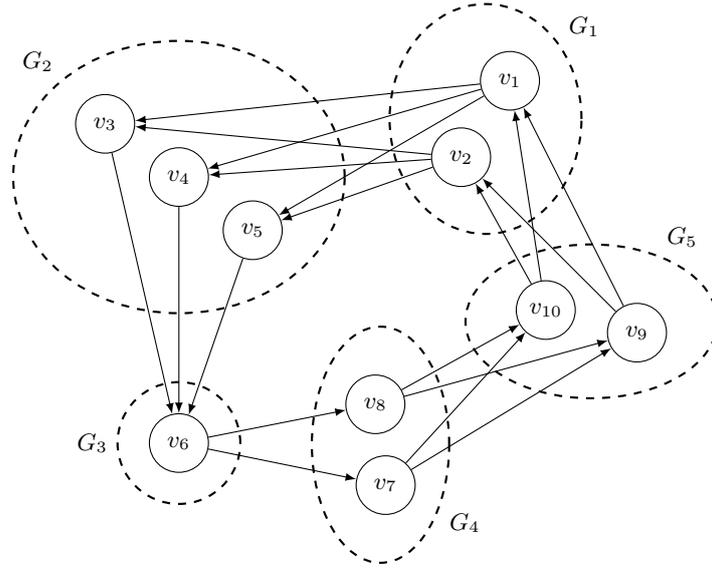


Figure 6.3: Illustration of the graph structure for the shortest cycle problem, which remains when fixing the tour for visiting the groups of a GTSP. The costs of the edges are omitted in this plot.

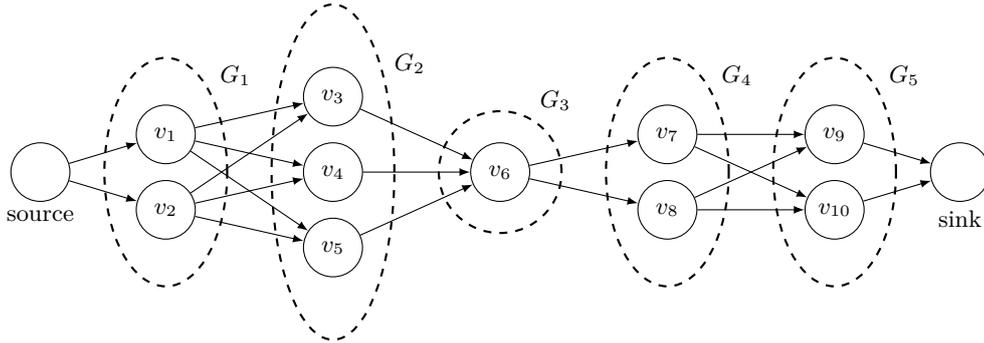


Figure 6.4: Illustration of the shortest path problem, which remains if we consider a fixed Hamiltonian path for visiting the groups of a GTSP. The artificially added vertices are labeled as source and sink. The costs of the edges are omitted in this plot.

speeds for both robot arms. Table 6.1 presents for each joint θ_g of each robot arm the maximum rotational speeds s_g^{LINAC} and s_g^{GUIDE} , respectively, converted from radians to degree. We limit the speed of each joint of the LINAC-robot and the GUIDE-robot to a fraction λ_{LINAC} and λ_{GUIDE} , respectively, of its maximum speed. Therefore, we estimate the time required for robot motion as

$$d(b_i, b_j) = \max_{g \in \{1, \dots, 6\}} \frac{|\theta_g^i - \theta_g^j|}{\lambda_{\text{LINAC}} \cdot s_g^{\text{LINAC}}} \quad (6.6)$$

Table 6.1: Maximum rotational speed of each joint of the KR 300 R2500 ultra and the LBR iiwa, which we consider as the LINAC-robot and the GUIDE-robot, respectively.

Joint	Maximum speed in $^{\circ}\text{s}^{-1}$	
	LINAC-robot	GUIDE-robot
θ_1	105	98
θ_2	101	98
θ_3	107	100
θ_4	122	130
θ_5	113	140
θ_6	175	180
θ_7	–	180

for moving the LINAC-robot between delivery of beams b_i and b_j and as

$$d(c_k, c_l) = \max_{g \in \{1, \dots, 7\}} \frac{|\theta_g^k - \theta_g^l|}{\lambda_{\text{GUIDE}} \cdot s_g^{\text{GUIDE}}} \quad (6.7)$$

for changing between the GUIDE-robot’s configurations c_k and c_l . The superscripts of θ_g in the equations denote that this is the angle obtained from inverse kinematics to realize the respective beam delivery or configuration. As mentioned before, this assumes that moving between each pair of beams is feasible for the LINAC-robot. We calculate the differences between the joint angles in radians under consideration of $\theta \equiv \theta + n \cdot 2\pi$, $n \in \mathbb{Z}$, and wrap them to the half-open interval $[-\pi, \pi)$. This, from a practical point of view, neglects the existence of joint limits. Furthermore, we neglect acceleration and deceleration. We consider the combinations of the fractions 1%, 4%, and 7% for λ_{LINAC} and λ_{GUIDE} to study different scenarios.

The overall treatment time does not only involve the time required for robot motion. We consider dose delivery by the LINAC at a rate of 1000 MU min^{-1} . We neglect the time required to change the diameter of the collimator between the delivery of sequent beams. Due to the random beam sampling and plan optimization, it is rather unlikely that two treatment beams starting at the same node only differ in their collimator’s diameter and we assume that opening or closing of the collimator’s iris takes less time than the LINAC-motion for realizing the different orientations.

6.3.2 Methods for Synchronization

We evaluate four methods in this work to approach the synchronization problem. We label them, with increasing complexity, as methods A, B, C, and D. They all consider a set of treatment beams obtained from inverse planning. We generate 30 different sets of candidate beams for each combination of patient, reference robot, and number of LIFT angles.

Method A: No GUIDE-robot In this method, we only consider optimizing the beam sequence for the LINAC-robot as an HPP and ignore the GUIDE-robot. Note that the treatment beams are based on the setup including the GUIDE-robot, nevertheless, and hence we should rather interpret this method as considering an infinitely fast GUIDE-robot. We use the software tool Concorde (version 03.12.19, written by David Applegate, Robert Bixby, Vašek Chvátal, and William Cook) which, besides providing heuristic algorithms, allows to solve TSPs exactly and very efficiently. The solution of method A will serve as the baseline for evaluating the following methods, which actually consider the GUIDE-robot and therefore have to lead to solutions with at least the same costs.

Method B: Non-optimized GUIDE-robot In this method, we fix the beam sequence obtained from method A and naively switch the GUIDE-robot’s configurations as necessary. For this purpose, we start in an arbitrary configuration in which the first beam is feasible. If the subsequent beam is not feasible in the current configuration, we switch to the closest configuration in which the beam is actually feasible. Therefore, there is no real optimization of the sequence of configurations involved.

Method C: Decoupled robot optimization We employ the decoupling described in Section 6.2.3. As in method B, we fix the beam sequence obtained from method A. However, we use the Dijkstra algorithm [47] to solve the remaining shortest path problem. Thereby, we obtain the optimal sequence of GUIDE-robot configurations given the fixed beam sequence. The solution of method B is part of method C’s search space. As method C is exact, we cannot obtain a worse solution. However, we can only expect reasonable results from this decoupling if the costs for moving the LINAC-robot dominate the quality of solutions. In contrast, results will be of low quality if we need to severely adapt the sequence of treatment beams for efficient synchronous motion.

Method D: GTSP-based coupled optimization We use the software tool GLKH (version 1.0, written by Keld Helsgaun) to approximately solve the full GTSP model. It implements the approach mentioned in Section 6.2.3 which converts the GTSP to an ATSP and employs the Lin-Kernighan-Helsgaun heuristic with GTSP-specific post-processing routines. We extract a path visiting all groups from the solution. Both the solutions from method B and method C are part of the search space of method D. However, method D is not exact and therefore we have no guarantee to find a solution of equal or better quality. This could be fixed by providing method B’s or method C’s solution as an initial guess. However, we only evaluate random initial guesses in this thesis to obtain information about the complexity of solving the full GTSP model without bias.

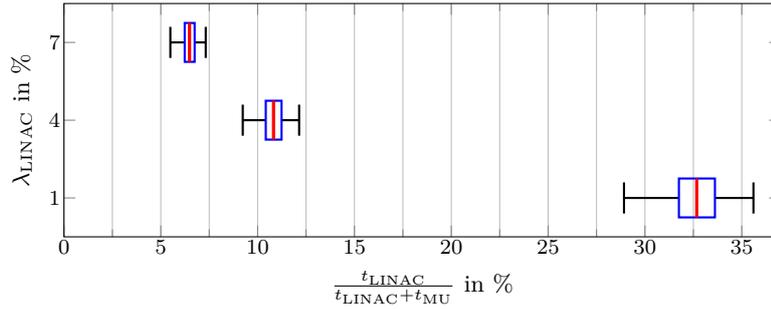


Figure 6.5: Time for LINAC-robot motion with method A relative to the overall treatment time, i.e., LINAC-robot motion and beam delivery.

6.4 Results

Robot motion by the LINAC-robot contributes already in conventional CyberKnife treatments to the overall treatment time. We define the overall treatment time as the combined time for robot motion and beam delivery. Due to the hard upper limit of 40 000 MU for the total activation time of the treatment beams in (2.16), which is typically exactly met in our resulting treatment plans, the time for beam delivery is essentially a constant. The time for LINAC-robot motion depends on the number of treatment beams, their spatial distribution, and the speed of the LINAC-robot, which we represent by a fraction λ_{LINAC} of its maximum speed. If we consider optimizing the LINAC-robot motion with method A, i.e., solving the TSP, we obtain the results presented in Figure 6.5. The median proportion of the overall treatment time contributed by robot motion is 32.7 %, 10.8 %, and 6.5 % when considering 1 %, 4 %, and 7 % for λ_{LINAC} , respectively.

We propose three methods to synchronize the motion of the LINAC-robot with the motion of the GUIDE-robot. We present in Figure 6.6 the resulting relative treatment times for these methods. A relative time of 100 % corresponds to the solution by method A, i.e., no overhead caused by the GUIDE-robot. A general observation is the trivial fact, that the impact of the additional GUIDE-robot becomes smaller when the GUIDE-robot become faster or the LINAC-robot becomes slower. Furthermore, the impact is slightly lower when considering seven instead of three LIFT angles. We observe that method B, i.e., no optimization of the GUIDE-robot motion but only changing its configuration if necessary for delivering the next beam from the sequence obtained by method A, that it performs worst. It both shows the highest medians as well as the largest variation. The more complex method C, which also considers the beam sequence from method A but optimizes the sequence of GUIDE-robot configurations, provides better results. In fact, it cannot give worse solutions because the solution by method B is part of its search space and the approach is exact. In contrast, the most complex approach by method D, which considers the full GTSP model, is not exact and not always able to find better solutions than method C. However, method D

especially outperforms the other approaches if we consider the GUIDE-robot to be slow relative to the LINAC-robot. For a rather fast GUIDE-robot, method C typically provides better results. In general, method D shows the lowest variation in the quality of its solutions.

In Table 6.2, we compare the results separately for the three positionings of the GUIDE-robot. We consider the best result from methods B, C, and D for each parameter combination. When employing three LIFT angles, robot A seems to be the best choice in terms of maintaining a short overall treatment time. In contrast, robot B provides the smallest increases for seven LIFT angles. For this setting, robot A performs worse and its difference to the other robots is bigger than observed for the worst robot in the other settings. In case of five LIFT angles, the results vary but are mostly best for robot C.

Another interesting aspect besides the overall treatment time is the number of configuration changes performed by the GUIDE-robot. In Figure 6.7, we observe that method D provides in some settings solutions with fewer configuration changes compared to method C. Major differences are visible if $\lambda_{\text{GUIDE}} \leq \lambda_{\text{LINAC}}$, i.e., the GUIDE-robot is relatively slow. Note that these are the cases in which moving the GUIDE-robot is expensive. In the other cases, the results of the two methods are rather similar. A robot-wise analysis of method D provides Table 6.3. Robot B provides the lowest numbers of configuration changes when employing three or five LIFT angles. For seven LIFT angles, robot C requires the lowest numbers of changes.

6.5 Discussion

In this chapter, we proposed and evaluated methods to efficiently synchronize the configuration changes by the GUIDE-robot and the motion of the LINAC-robot for beam delivery. For this purpose, we described a TSP-based model for optimized LINAC-motion w.r.t. time and used this as a reference. We extended this model to a GTSP in order to account for the GUIDE-robot as well. Because the resulting problem is hard to solve, we studied three approaches to solve it approximately with increasing complexity. We refer to the reference as method A and to the three other methods as B, C, and D where method D has the highest complexity.

In our setup, the overall treatment time consists of the activation times of the treatment beams and the robot motion. Considering method A, we showed that the movement speed of the LINAC-robot severely influences the general impact of time for robot motion on the overall treatment time (Figure 6.5). However, motion of a robot close to a patient is a highly safety-critical action. Therefore, we only considered moving at 1% to 7% of the LINAC-robot's maximum speed. Note that an increased treatment time due to slower robot motion has several effects. Besides reduced patient comfort, the likelihood of larger target motion increases, exposure to imaging dose increases in case of additional X-ray imaging at fixed time intervals, and costs of a treatment increase due to, for example,

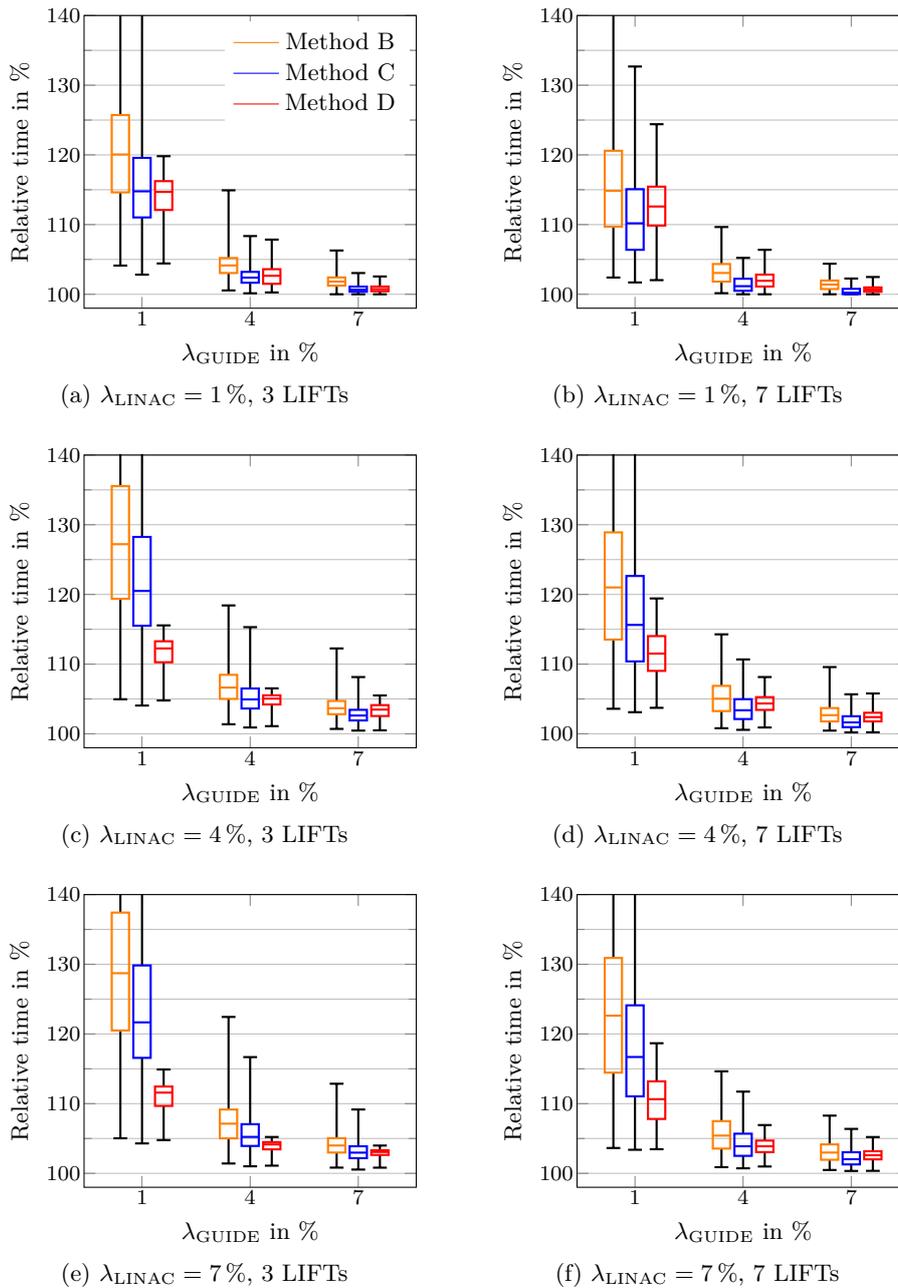


Figure 6.6: Relative overall treatment time (robot motion and beam delivery) when employing method B (orange), method C (blue), or method D (red). 100 % corresponds to the solution without GUIDE-robot from method A and therefore cannot be undercut by the other methods additionally considering the GUIDE-robot. Note that the plots are cut at 140 %.

Table 6.2: Mean and standard deviation of the relative increase of the overall treatment time (robot motion and treatment time) in percent for the three GUIDE-robot setups. The increase is relative to method A which does not consider the GUIDE-robot. For each experiment, the best of the three results obtained from method B, method C, and method D is considered. For each parameter combination, the best result is marked bold.

# LIFTS	λ_{LINAC}	λ_{GUIDE}	robot A	robot B	robot C
3	1%	1%	11.4 ± 3.3	12.6 ± 4.1	11.7 ± 3.7
		4%	1.6 ± 0.8	2.6 ± 1.2	2.0 ± 1.1
		7%	0.6 ± 0.3	0.7 ± 0.5	0.6 ± 0.4
	4%	1%	9.8 ± 2.4	12.1 ± 3.4	11.4 ± 2.0
		4%	3.9 ± 1.2	4.3 ± 1.4	4.0 ± 1.2
		7%	2.2 ± 0.8	2.7 ± 1.0	2.4 ± 0.9
	7%	1%	9.1 ± 2.3	11.8 ± 3.2	11.1 ± 1.6
		4%	3.3 ± 0.9	3.8 ± 1.1	3.6 ± 0.9
		7%	2.4 ± 0.7	2.5 ± 0.8	2.4 ± 0.7
5	1%	1%	13.6 ± 3.5	12.3 ± 4.5	11.5 ± 4.1
		4%	1.9 ± 0.9	1.7 ± 1.0	1.6 ± 0.8
		7%	0.8 ± 0.3	0.4 ± 0.4	0.4 ± 0.3
	4%	1%	12.3 ± 2.4	13.3 ± 3.7	13.2 ± 3.2
		4%	4.6 ± 1.1	4.1 ± 1.5	3.8 ± 1.4
		7%	2.6 ± 0.8	2.2 ± 0.9	2.0 ± 0.9
	7%	1%	11.6 ± 2.1	13.0 ± 3.4	12.3 ± 2.7
		4%	4.1 ± 1.0	4.2 ± 1.3	3.9 ± 1.2
		7%	2.8 ± 0.7	2.5 ± 0.9	2.3 ± 0.9
7	1%	1%	12.8 ± 3.4	7.8 ± 3.8	8.2 ± 4.1
		4%	1.8 ± 0.8	0.8 ± 0.7	0.8 ± 0.6
		7%	0.7 ± 0.3	0.2 ± 0.3	0.2 ± 0.2
	4%	1%	12.5 ± 2.6	9.3 ± 3.9	10.2 ± 3.9
		4%	4.3 ± 1.1	2.6 ± 1.3	2.8 ± 1.4
		7%	2.3 ± 0.7	1.2 ± 0.7	1.3 ± 0.7
	7%	1%	11.8 ± 2.3	8.6 ± 3.5	9.8 ± 3.8
		4%	4.0 ± 1.0	2.7 ± 1.2	3.0 ± 1.3
		7%	2.6 ± 0.7	1.5 ± 0.8	1.7 ± 0.8

higher personal expenses for each treatment and lower patient throughput on the machine.

To minimize the increase in treatment time due to additional time required for moving the GUIDE-robot, we employed the three methods B, C, and D (Figure 6.6). Clearly, the speed of the GUIDE-robot relative to the LINAC-robot dictates the general increase in treatment time. We observed that the simplicity of method B, which does not optimize the sequence of configuration changes at all, leads to worse solutions than the other two methods. Interestingly, method C which fixes the beam sequence from method A and only optimizes the sequence of configuration changes outperformed in some scenarios method D which simultaneously optimizes both sequences. This especially happened if the GUIDE-robot was

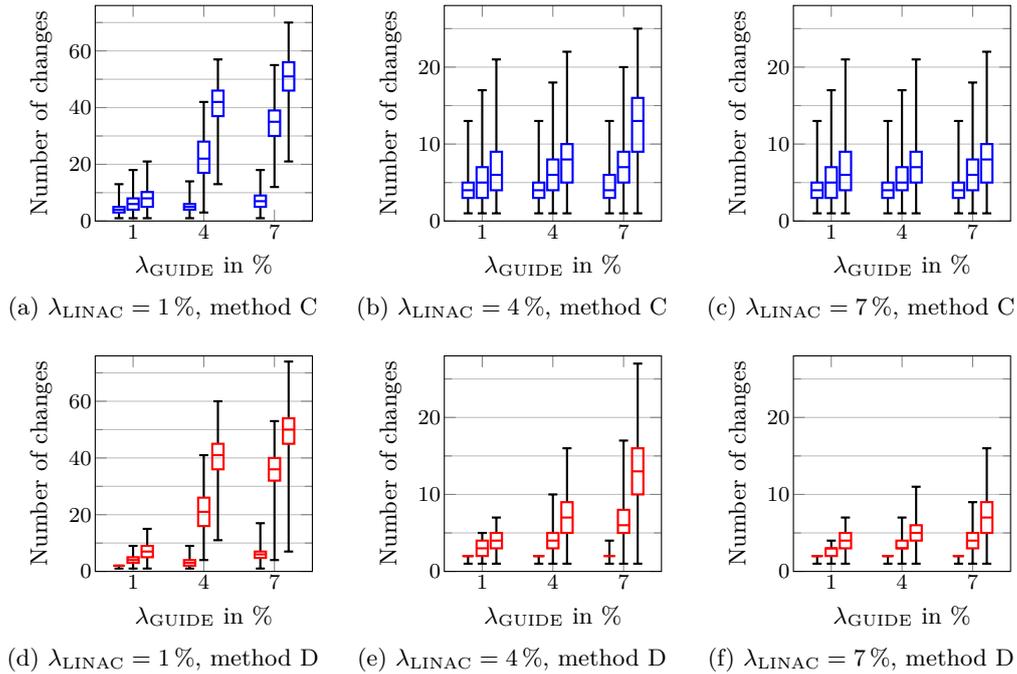


Figure 6.7: Resulting number of configuration changes by the GUIDE-robot when employing method C (a-c) or method D (d-f). Note that the y -axis scaling in (a) and (d) is different to that in the other sub-figures. Further, note that some boxes in (d), (e), and (f) collapse due to the small and discrete variation.

moving faster than the LINAC-robot. In this case, the influence of the GUIDE-robot is rather small and the most-efficient beam sequence from method A is still a good choice. The results showed that solving the full problem is practically difficult. However, a simple improvement would be to consider the result from method C as the initial guess for the underlying heuristic solver of method D. Thereby, the result of method D cannot be worse than the result of method C. Because we can employ method C without relevant computational effort, this option should definitely be considered in practical application of this synchronization procedure, although it might also add an unfavorable bias in some cases.

Furthermore, we showed that the increase of the overall treatment time also differs for the three reference positionings of the GUIDE-robot (Table 6.2). Depending on the number of employed LIFT angles, a different setup was advantageous. We already observed in previous chapters that the influence of the positioning vanishes w.r.t. the plan quality with increasing number of LIFT angles under consideration. Hence, this chapter's analysis might provide a suitable metric for deciding between different setups in case of employing multiple angles as it considers a completely different figure of merit.

An aspect which our proposed methods do not explicitly take into account is the

Table 6.3: Mean and standard deviation for the number of configuration changes by the GUIDE-robot in the solutions obtained with Method D. The best result for each parameter combination is marked bold.

# LIFTs	λ_{LINAC}	λ_{GUIDE}	Robot A	Robot B	Robot C
3	1%	1%	1.8 ± 0.4	1.7 ± 0.5	1.9 ± 0.3
		4%	3.3 ± 1.6	2.9 ± 0.9	3.4 ± 1.3
		7%	7.5 ± 2.6	4.6 ± 1.5	6.5 ± 1.9
	4%	1%	1.8 ± 0.4	1.7 ± 0.5	1.9 ± 0.3
		4%	1.8 ± 0.4	1.7 ± 0.4	1.9 ± 0.2
		7%	1.9 ± 0.6	1.7 ± 0.4	2.0 ± 0.3
	7%	1%	1.8 ± 0.4	1.7 ± 0.5	1.9 ± 0.3
		4%	1.8 ± 0.4	1.7 ± 0.5	1.9 ± 0.3
		7%	1.8 ± 0.4	1.7 ± 0.5	2.0 ± 0.2
5	1%	1%	4.5 ± 1.5	3.5 ± 1.3	3.9 ± 1.3
		4%	25.8 ± 6.2	16.8 ± 5.3	21.2 ± 6.0
		7%	36.9 ± 5.8	33.7 ± 6.5	35.8 ± 6.4
	4%	1%	3.1 ± 0.7	2.8 ± 0.8	3.0 ± 0.8
		4%	4.8 ± 1.5	3.5 ± 1.3	3.9 ± 1.2
		7%	7.8 ± 2.3	5.3 ± 1.9	5.9 ± 2.0
	7%	1%	3.2 ± 0.7	2.7 ± 0.8	2.9 ± 0.8
		4%	3.6 ± 0.9	3.1 ± 1.1	3.3 ± 1.0
		7%	4.7 ± 1.4	3.6 ± 1.3	4.1 ± 1.3
7	1%	1%	9.2 ± 2.3	6.6 ± 2.3	5.7 ± 2.6
		4%	40.0 ± 8.0	42.4 ± 6.5	37.7 ± 7.4
		7%	49.8 ± 8.0	50.6 ± 7.5	48.1 ± 6.9
	4%	1%	4.4 ± 1.1	3.8 ± 1.3	3.6 ± 1.3
		4%	9.5 ± 2.4	6.2 ± 2.3	6.0 ± 2.3
		7%	16.7 ± 3.6	11.9 ± 3.5	10.8 ± 3.8
	7%	1%	4.3 ± 1.0	3.7 ± 1.1	3.5 ± 1.4
		4%	5.8 ± 1.7	4.7 ± 1.7	4.2 ± 1.8
		7%	9.2 ± 2.3	6.5 ± 2.3	5.9 ± 2.5

number of configuration changes, which we need to realize the sequence of configurations (Figure 6.7 and Table 6.3). This number severely depends on the number of LIFT angles we consider and the speeds of the two robots. Although fewer configuration changes obviously correspond to less GUIDE-robot motion, the actual relationship between the number of configuration changes and the increase of the overall treatment time is more complicated. The faster the GUIDE-robot is compared to the LINAC-robot, the more motion it can perform without affecting the overall time. For this reason, we observed even beyond 50 configuration changes in some of our solutions although we only considered seven different configurations. It might be interesting to limit this number in order to decrease the frequency of GUIDE-robot motion. One motivation for this action is again the patient comfort. The GUIDE-robot is very close to the patient and frequent motion might cause discomfort.

In the actual CyberKnife system, the LINAC-robot moves at about 4% of its

maximum speed. Considering the same limitation for the GUIDE-robot, we obtained in our evaluations moderate increases in the overall treatment time in the order of 3% to 5%, depending on the exact setup (Table 6.2). Note that we considered an optimal trajectory of the LINAC-robot as the reference for these values. In practice, there are additional factors already leading to less efficient trajectories. First, there might be obstacles or other limitations for safe robot motion forbidding to travel directly between some pairs of beam nodes. This aspect would directly affect our ultrasound-guidance setup as well. Second, the real system requires X-ray imaging in defined intervals for reliable tracking of target motion. The LINAC-robot can block these imaging X-ray beams similar to our GUIDE-robot blocking treatment beams. Thus, more time-consuming trajectories might be necessary to comply with the required imaging intervals throughout the whole course of the treatment. This aspect has an interesting effect when considering the impact on treatment time of our ultrasound-guidance setup, because a fundamental motivation of the ultrasound guidance is to reduce the necessity for X-ray imaging. Thus, future studies should compare the impact of more efficient LINAC-robot trajectories due to less X-ray imaging to the impact of less efficient trajectories due to motion of the GUIDE-robot.

7 Towards Optical Coherence Tomography for Image Guidance in Cranial Radiosurgery

In the previous chapters, we discussed different treatment-planning aspects appearing in robotic ultrasound guidance. Ultrasound imaging is a promising modality for abdominal targets because it is non-ionizing, allows for real-time volumetric imaging, and provides direct and markerless imaging of internal targets. The integration of intra-fractional ultrasound imaging in the treatment room is not as expensive as MRI, but requires to replace the conventional manual operator with a robotic system and to consider the blocking of treatment beams by the ultrasound transducer and the robot.

However, ultrasound is not applicable for guiding cranial treatments due to the skull. We discussed in Section 3.3.4 that optical tracking systems can be employed to track head motion during cranial radiation therapy. These are also non-ionizing, low-cost, allow for real-time imaging, and integration in the treatment room mainly requires to establish a steady line-of-sight. A drawback compared to other imaging modalities like ultrasound, CT, MRI, or EM tracking is the superficial scanning of conventional optical tracking systems. While external tracking of head motion itself is rather uncritical for cranial treatment sites because the internal targets do not move much relative to the skull, it requires either to attach markers which increase the effort for integration or, in case of surface scanning, to have clear and rigid features for precise image registration.

A relatively new and less common optical imaging modality is OCT. Similar to the measurement of acoustic reflections from different layers within tissue in ultrasound imaging, OCT measures reflections of light to reconstruct depth profiles. Employing near-infrared light allows, in addition to surface reflections, to receive reflections from structures which are a few millimeters within the tissue. In consequence, one of OCT's main characteristics is its small FOV. However, state-of-art systems provide volumetric imaging with a spatial resolution in the order of $10\ \mu\text{m}$. Nevertheless, the temporal resolution has been a limiting factor for a long time. This has changed recently due to major improvements in laser technology and high-resolution volumetric OCT imaging is feasible today at video rate and with few motion artifacts. For this reason, we propose to employ OCT imaging to realize motion tracking. Due to its high spatial resolution, it can acquire tiny superficial structures which cannot be exploited in larger-scale optical systems. Furthermore, OCT can exploit structures below the surface that are

not accessible at all by superficial scanning. Therefore, we claim that OCT is a promising choice to implement precise and reliable markerless tracking, even in scenarios with poorly structured surfaces.

In this chapter, we propose an OCT-based tracking system following the third research question formulated in Section 1.2. We describe and evaluate a prototypical system, which we design for general-purpose tracking though a major motivation is to develop a novel approach for head tracking and image guidance for cranial radiation therapy. We first introduce the fundamentals of OCT. Subsequently, we discuss recent applications of OCT for guidance of interventions and tracking. In the main part of this chapter, we design and describe an OCT-based general-purpose tracking system. It provides fast volumetric imaging and markerless tracking of moving structures, which do not need to be superficially visible. We finally show the feasibility of 3D- and 6D-motion tracking with our system design and present quantitative tracking results. Parts of this chapter have been published in a journal [233] and presented at conferences [234, 235].

7.1 Basic Principles

This section introduces the basic principles of OCT. It is a non-invasive imaging modality employing reflections of light and has been proposed in 1991 [104] for the first time as a dedicated cross-sectional imaging concept. It does not only allow for distance scans, but reflections are also received from depths up to about 1 mm to 2 mm in tissue. This is possible because the typically employed wavelengths like 800 nm, 900 nm, 1050 nm, 1300 nm, and 1550 nm experience a relatively low absorption and reduced scattering in water or tissue. The small size of the FOV, which is typically a few millimeters in air, is compensated by a very high spatial resolution in the order of 10 μm . While the basic principle of OCT only allows for acquisition of one-dimensional depth profiles along the beam direction, using additional scanning mechanisms which laterally deflect the light beam provide 2D and 3D imaging. Figure 7.1 exemplarily shows a cross-sectional scan of a human finger tip, which was acquired with a Telesto system (Thorlabs, Germany/USA). There are also advanced approaches which acquire axial sequences of lateral image slices [51, 156]. Similar to ultrasound imaging, one-, two-, and three-dimensional scans are commonly termed A-scans, B-scans, and C-scans, respectively.

Today, OCT is an established technique in ophthalmology [73]. Further, it is used for intravascular imaging [27] by integration of an optical fiber into catheters and combination with a larger-scale modality like angiography or recently magnetic particle imaging [91, 149]. It also experiences increasing interest in other fields, for example dentistry [168] and dermatology [193]. Other applications include elastography [147, 150], force sensing [82], and flow estimation [255]. Recent advantages in imaging speed allow for new applications in the context of 4D imaging [128] and imaging with substantially decreased motion artifacts [269].

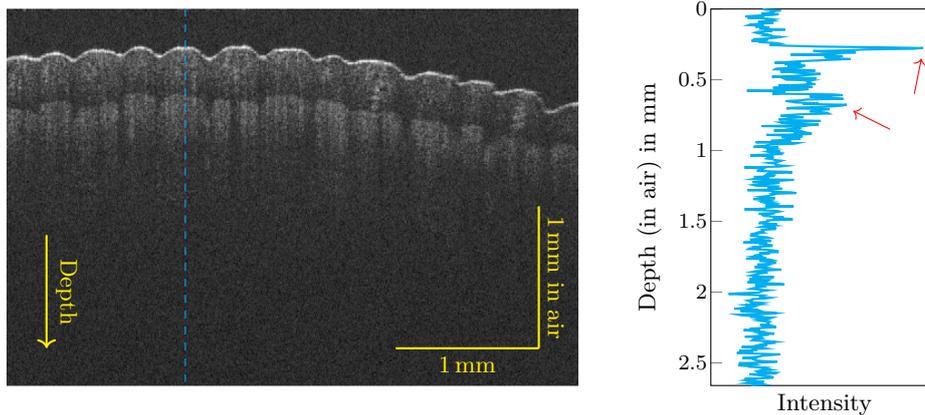


Figure 7.1: Cross-sectional B-scan of a finger tip. It consists of 512 A-scans with 512 pixels each and covers $4 \times 2.66 \text{ mm}^2$ in air. The normalized intensity of the exemplary A-scan shows, as indicated by the red arrows, one major peak at the surface depth and a minor peak at the transition to the second tissue layer.

7.1.1 Low-Coherence Interferometry

The basic concept of OCT relies on an interferometer, typically in the Michelson configuration. As illustrated in Figure 7.2, a light beam is, often equally, split into two beams. One of the beams is directed onto a mirror and reflected. This part between beam splitter and mirror is referred to as the reference arm. The other beam is directed onto a target which we want to measure. This path starting at the beam splitter is referred to as sample arm. If we have an at least partially reflecting or backscattering target in the sample arm, a fractional intensity of the beam reaches back to the beam splitter. There, it superimposes with the reflected beam from the reference arm. The intensity of the superposition is then measured by an infrared-sensitive detector.

In classical Michelson interferometry, we typically employ a laser as the light source, which is monochromatic. It produces coherent light, i.e., the superimposed reflected beams give an interference signal on the detector which depends on the path length difference between the reference arm and the reflector in the sample arm. If the difference is an integer multiple of the wavelength, we observe constructive interference. If we have an odd integer multiple of the half wavelength, then we observe destructive interference. In practice, however, we never have perfectly coherent light. Instead, the coherence is limited to the so-called coherence time τ_c . If the temporal delay between the beam from the sample arm and the beam from the reference arm is longer than τ_c , we do not observe interference anymore. At least for OCT, it is more common to describe the coherence by the coherence length l_c which is the distance the light travels during τ_c in some medium, i.e., the maximum path length difference for interference.

Having limited coherence, i.e., a rather short l_c , is a fundamental concept of

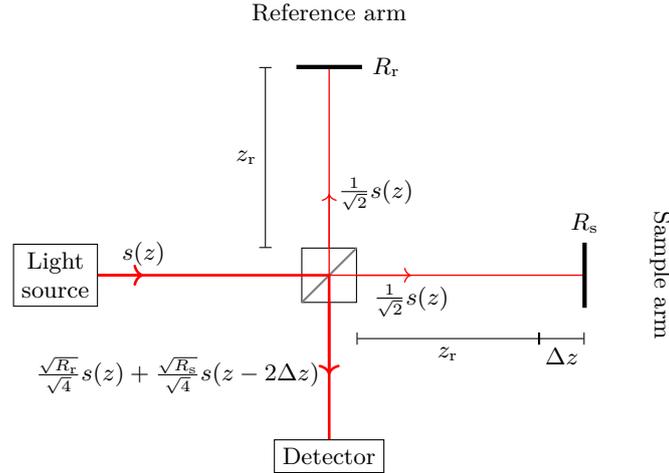


Figure 7.2: Sketch of a Michelson interferometer. A light beam is splitted into two beams with half intensity. The beam in the reference arm is reflected at a mirror with power reflectivity R_r at distance z_r . The beam in the second arm is reflected at a sample with power reflectivity R_s at distance $z_r + \Delta z$. Eventually, the intensity of their superposition is recorded at a detector and contains an interference signal depending on Δz and the wavelength.

OCT. We only obtain an interference signal from reflecting structures whose path difference to the reference mirror is covered by the coherence length. This provides a spatial encoding. If we have multiple reflecting layers in the sample arm, we can iteratively increase the length of the reference arm and only observe constructive interference at the detector, if the current length matches the distance to one of the layers. Thereby, we can record a depth profile and this concept results in time-domain optical coherence tomography (TD-OCT). It also illustrates that the axial resolution in OCT is directly related to the coherence length of the light source. The lateral resolution is mainly determined by the employed optics.

7.1.2 Fourier-Domain OCT

TD-OCT allows to record arbitrarily long depth profiles. However, an iterative mechanical movement of the reference mirror is very slow for reasonable numbers of measurements and therefore does not allow for satisfying A-scan rates. The approach has been almost completely replaced by Fourier-domain optical coherence tomography (FD-OCT). As the name suggests, it records the spectral information of the interference signal. There are two main versions of FD-OCT, spectral-domain optical coherence tomography (SD-OCT) and swept-source optical coherence tomography (SS-OCT). Both have in common that the reference mirror, i.e., the reference arm length, is fixed during A-scan acquisition. Their difference is the way we record the spectral information.

In SD-OCT, we use a broad-band light source and a spectrometer. Each element

of a spectrometer's line sensor measures the incident intensity of a small range of frequencies. If we consider N discrete layers in the sample arm with reflectivities R_{s_i} , $i = 1, \dots, N$, and path differences Δz_i compared to the reference mirror, then the signal $I_D(k)$ at the detector is

$$\begin{aligned}
 I_D(k) \propto & \frac{S(k)}{4} \left(R_r + \sum_{i=1}^N R_{s_i} \right) \\
 & + \frac{S(k)}{2} \sum_{i=1}^N \sqrt{R_r R_{s_i}} \cos(2k\Delta z_i) \\
 & + \frac{S(k)}{2} \sum_{i=1}^N \sum_{j=i+1}^N \sqrt{R_{s_i} R_{s_j}} \cos(2k(\Delta z_i - \Delta z_j))
 \end{aligned} \tag{7.1}$$

with the power spectrum $S(k)$ of the light source. More details can be found, for example, in Drexler et al. [49]. The first term in (7.1) describes the constant received intensities. The second term is the interference signal we are actually interested in. The third term describes pairwise interference between the layers in the sample arm. If we only consider the second term, we observe a set of cosines at the detector and their frequencies correspond to twice the path differences Δz_i . If we apply a Fourier transform, we obtain peaks at exactly these frequencies and this is the A-scan signal. Typically, one re-scales the depth axis to have peaks at Δz_i rather than at $2\Delta z_i$. The power spectrum is related to the coherence function via Fourier transform and smears the peaks in the A-scan depending on the coherence length. The fact that we cannot retrieve the sign of the path difference from (7.1) results in ambiguities. The signals of layers with difference Δz and with difference $-\Delta z$ appear at the same position in the A-scan. This effect can be described as a mirroring artifact.

The idea of SS-OCT is to avoid using speed-limiting line cameras. Instead, a narrow-band light source varies its frequency over time and we can thereby record the interference spectrum with a single detector element as a function of time. Typically, the frequency variation, called sweep, is assumed to have a linear relationship between frequency and time and each sweep yields the spectrum of one A-scan. Especially Fourier-domain mode locking (FDML) lasers [105] can realize sweeps with high repetition rates.

7.1.3 Lateral Scanning

In order to acquire two- or three-dimensional OCT images, the light beam has to be deflected such that we acquire a sequence of A-scans at different positions. The classical approach for volumetric scanning is a scan head which contains two galvanometer mirrors (galvos) as illustrated in Figure 7.3. Each mirror is responsible for imaging along one axis and changes the angle of the outgoing beam by rotation. These two mirrors alone would position the A-scans radially. Ideally, a fully telecentric f-theta lens as shown in Figure 7.4 is used for scanning.

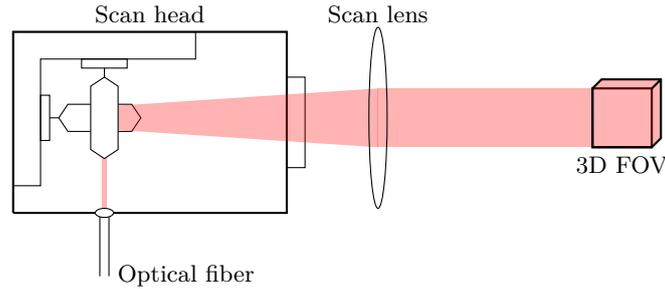


Figure 7.3: A scan head uses two galvos to deflect the light beam laterally over time. Thereby, they realize a 2D grid of A-scans resulting in a 3D FOV and volumetric images.

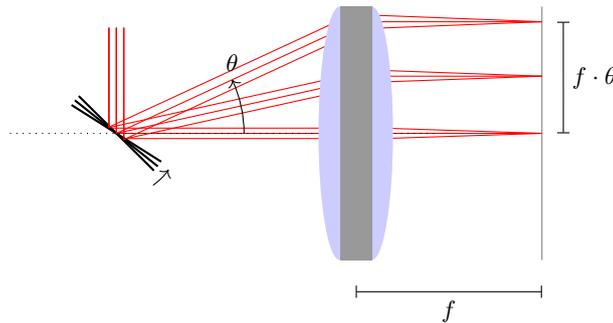


Figure 7.4: Illustration of galvo-based scanning with an ideal telecentric f-theta lens.

Its f-theta design allows to scan within a plane and with a linear relation between ingoing beam angle and outgoing lateral shift. Furthermore, telecentricity results in parallel A-scans, although there is always a remaining angle in real lenses.

7.1.4 Practical SS-OCT Image Reconstruction

The following steps are typical for reconstructing an A-scan for an SS-OCT device based on an FDML laser. We adopt the terminology used by the company Optores (Germany) for their systems. Before we can reconstruct a digital raw A-scan signal $r = (r_1, \dots, r_N)$ consisting of N samples recorded by the detector, basic calibration signals have to be acquired and processed. This is necessary, because real systems do not produce a perfectly linear frequency sweep. Due to any non-linearity, linear sampling in time results in non-linear sampling in frequency space, resulting in chirp-like effects in the spectrum. If a Fourier transform were directly applied to these samples, the A-scans would suffer from severe artifacts and distortions. Therefore, the following steps are made initially.

- Acquire a background signal b of length N by closing the sampling arm with a shutter. Thereby, the signal does not contain any actual information but only the constant signal from the reference arm (compare (7.1)) and possibly

noise due to internal reflections and interferences.

- Acquire a recalibration signal c of length N by closing the sampling arm and opening a so-called recalibration arm. This arm is split from the reference arm and contains a mirror which is aligned to the mirror in the reference arm.
- Estimate the non-linearity of the sweep from the background-corrected recalibration signal $c - s$ and determine points for linear re-sampling.

The last step requires to solve an inverse problem. It is based on the assumption that the mirror in the recalibration arm leads to a single interference signal, i.e., $c - b$ is a single cosine with a certain frequency. As one approach, one can identify the zero crossings of this signal which would be equidistantly separated in case of a linear sweep. From the actual distances one can estimate the non-linearity and determine sampling points which represent measuring a linear sweep.

After the initial calibration, we can start to reconstruct A-scans. This involves the following steps which are executed on each A-scan separately.

- Acquire a raw A-scan r .
- Subtract the background signal b from r and resample it, for example by linear interpolation, to compensate for the non-linear sweep.
- Multiply the corrected signal with a windowing function, e.g., a Hann window, and apply a discrete Fourier transform (DFT), commonly implemented as a computationally more efficient fast Fourier transform (FFT).
- Take the absolute values of the result's complex values which correspond to positive frequencies and apply logarithmic compression.

All these steps have rather low computational effort. Due to the fact that all A-scans of a C-scan can be reconstructed independently of each other, reconstruction can easily be parallelized. If we acquire B-scans in both scanning directions of the resonant galvo, we eventually have to reverse the order of the A-scans in every other B-scan.

7.2 OCT-Based Navigation and Tracking

OCT imaging is established in ophthalmology and therefore compensating retinal motion is an important issue for artifact-free imaging and reliable diagnosis [31, 151, 218]. In this thesis, however, we do not consider imaging of the eye but are interested in rather generic tracking targets. Therefore, we focus in this section on discussing previous application of OCT to navigation of non-ophthalmologic interventions by tissue tracking.

Many studies investigated OCT for guidance of different types of medical lasers [26, 72, 120, 264]. One particular setting is guidance during cochleostomy. The

ability to identify borders below the surface has also been exploited to guide ablation [283]. Also rotary OCT probes have been studied for imaging and pre-operative planning [94]. Furthermore, Zhang et al. [284] proposed to ablate tiny dots superficially and use them as artificial markers for 6-DOF tracking. A similar approach has been used for calibration between a robotic ablation system and an OCT scan head to enable targeted soft-tissue ablation and spatial mapping of mass-spectrometric characterizations [229]. Gessert et al. [83] showed that deep learning approaches can estimate the pose of a known marker in volumetric OCT data even in occluded settings. This task is related to OCT-based tracking of instruments like needles [271], which can rely on information about the artificial tools in contrast to direct tissue tracking. A markerless concept to realize visual servoing for precise repetitive positioning of an OCT probe has been described by Dahroug et al. [42].

The limited size of the FOV motivated studies considering image stitching. One strategy is to use a robotic system which moves the OCT scanner. Stitching is then employed purely relying on a calibration [206] or with additional post-processing with some registration method [61]. Other studies focused on the algorithmic stitching part instead of the hardware setups to realize acquisition [76, 153, 166]. Several studies suggested the prominent feature descriptor SIFT [165] to register OCT volumes [76, 166, 187]. Other approaches rather considered maximum intensity projections or depth maps [152, 154]. Also an explicit separation between estimation of the lateral and axial motion has been proposed [76, 188].

7.3 Design of an OCT-Based Tracking System

The basis of our tracking system is a commercial OCT device including a scan head for C-scan acquisition. The design of the actual tracking system is based on the following specifications and considerations.

Volumetric imaging In contrast to depth cameras which rather provide 2.5D information, we want to exploit the full 3D imaging capability of OCT. This potentially allows to track targets with structure-less surfaces as long as there is structure below the surface.

Markerless tracking We do not want to require prior knowledge about a target for tracking. Instead, the C-scans should be evaluated for generic information about motion. Typically, this results in feature-based or intensity-based approaches which determine motion with respect to a defined template C-scan. This requires to maintain a certain amount of overlap between subsequent C-scans.

High imaging and tracking rate The slower the acquisition process of the imaging system, the more motion artifacts we observe in the images. For OCT, the

acquisition rate of commercial systems is mainly advertised in terms of the A-scan rate. The more A-scans we acquire per C-scan, the lower is our C-scan rate and the more motion effects are present within a C-scan. Therefore, using volumes consisting of as few A-scans as reasonably possible automatically reduces the sensitivity to motion artifacts. Additionally, these small C-scans reduce the computational effort for image processing and thereby allow for higher tracking rates.

Motion beyond the FOV size OCT only provides a small FOV in the order of some millimeters and an arbitrarily moving target will quickly leave it. Therefore, we desire to have a tracking system which is able to adapt the position of the FOV in order to follow larger-scale motion as well. This also allows to maintain a proper overlap of the C-scans as required for tracking algorithms.

6D motion In general, natural motion of rigid objects consists of translations and rotations. Therefore, we aim to track these types of motion and to quantitatively estimate the trajectory.

We design the tracking system mainly with respect to these five considerations and the details will be described in the following subsections. In terms of image processing, we need a computationally efficient volumetric approach which does not require predefined structures in the image data. For this purpose, we will consider correlation filters which work intensity-based and can be efficiently calculated in Fourier domain. In terms of hardware, we need a setup which allows to move the FOV in space to follow a moving target. As a side effect, this allows to keep the same part of the target in the FOV which reduces the risk for drifting effects in template-based tracking algorithms compared to using different locations as templates over time [154, 174]. One approach would be mounting an OCT scan head to a robot [206]. While such an approach allows to realize 6D repositioning of the FOV, it has limited dynamics due to the mechanical movements and additionally vibrations might distort the scanning. Instead, we directly exploit the fact that axial shifting can be realized by adjusting the length of the reference arm. Further, we adapt the galvo-based scanning process which realizes the transition from A-scans to C-scans. By setting up a second scanning stage, we can laterally shift a whole C-scan to different scanning positions. A disadvantage of this approach is that the range of motion is much more limited than for a robot-mounted scan head and that we cannot follow rotational movements of the target directly. However, we will show how we can overcome this limitation, at least for rotations within a certain range, by tracking multiple points on a target periodically.

7.3.1 Hardware Setup

OCT System

We employ a commercially available SS-OCT system (OMES, Optores, Germany) with a center wavelength of 1315 nm and an A-scan rate of about 1.59 MHz. We crop the A-scans to 480 pixels corresponding to an imaging range of about 3.5 mm in air. For C-scan imaging, we have two scan heads available. They employ a resonant galvo with a frequency of about 15 kHz and 5.67 kHz, respectively, for the fast scanning direction to allow for high acquisition rates. To further increase the rate, one B-scan is acquired during forward movement of the galvo and another one during backward movement, i.e., the A-scans of every second B-scan are spatially acquired in reverse order, resulting in scan rates of about 30 kHz and 11.3 kHz, respectively. Not all A-scans during a half-period of the galvo are meaningful, however. Close to its turning points, the galvo has to accelerate and decelerate which leads to non-linear sampling effects. For an equidistantly sampled volume, as typically used in imaging, only about 85% of the A-scans can be used and the remaining are discarded. The general setup is illustrated in Figure 7.5. Furthermore, a synchronization is necessary in order to have the forward and the backward acquired B-scans properly aligned. For a resonant galvo with frequency f_{RG} and a system with A-scan rate $f_{\text{A-scan}}$, we can acquire at most

$$N_{\text{max}} = \frac{f_{\text{A-scan}}}{f_{\text{B-scan}}} \cdot 0.85 \quad (7.2)$$

linearly sampled A-scans, where $f_{\text{B-scan}} = 2f_{\text{RG}}$ if we employ both forward and backward acquisition of B-scans. Interestingly, the C-scan rate of an OCT system with a resonant galvo does not directly depend on the system's A-scan rate. Instead, we have to consider the fixed B-scan rate $f_{\text{B-scan}}$ and the number of B-scans per C-scan N , which is essentially unlimited for non-resonant galvos but assumed to be equal to the number of A-scans per B-scan throughout this thesis. Therefore, we have $N \leq N_{\text{max}}$ and the C-scan rate is

$$f_{\text{C-scan}} = \frac{f_{\text{B-scan}}}{N + 4}. \quad (7.3)$$

The additive constant of four is due to the fact that we have to move back the non-resonant galvo to its origin after acquisition of each C-scan which takes some time, in our case equal to the acquisition time of four B-scans. This time is rather chosen empirically and conservatively and in principle would not be necessary if C-scans were, similar to the B-scans, acquired during both forward and backward motion of the galvo.

The resonant galvos have two major drawbacks. Firstly, their amplitude is regulated with an analog rotary knob. This limits the reproducibility of a specific FOV size compared to the digitally controlled slow scanning direction. Secondly, the frequency of the galvo is in reality not fixed but slightly varies and drifts, for example due to varying temperature. This complicates the synchronization

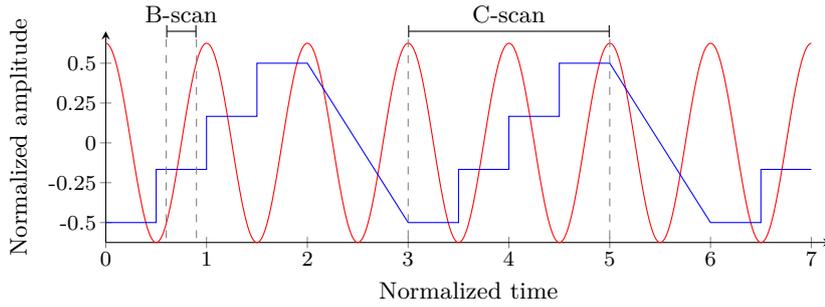


Figure 7.5: Illustration of the amplitudes of a resonant (red) and a stepper galvo (blue) during C-scan acquisition with an FOV of size 1×1 in normalized units. In this example, a C-scan consists of only four B-scans. A-scans acquired around the resonant galvo's turning points have to be discarded due to the geometric distortions. Therefore, the amplitude of the resonant galvo has to be chosen larger than the desired FOV of the B-scans. Furthermore, B-scans have to be discarded while the stepper galvo moves back after a C-scan.

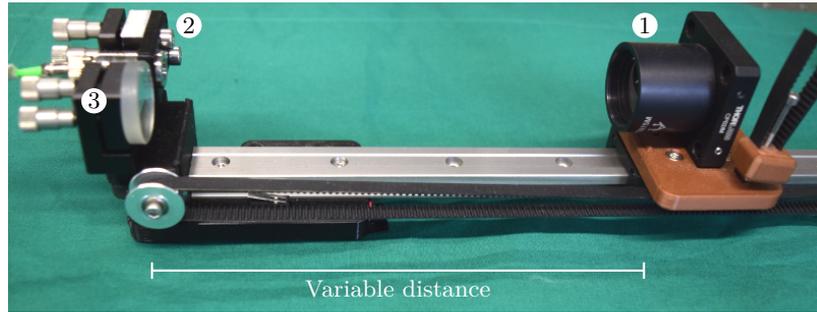
between the forward and the backward acquired B-scans, which might have to be adapted if drifting effects occur over time.

Computing Device

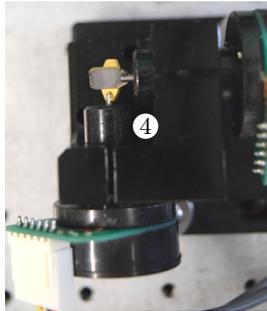
A computer with a digitizer board (ATS9373, Alazartech, Canada) records 2432 raw samples for each A-scan. The computer is equipped with two graphics processing units (GPUs), one Geforce GTX Titan X and one Geforce GTX 980 Ti (Nvidia, USA). The first GPU is intended for image reconstruction while the second GPU is intended for running the tracking algorithms.

Axial and Lateral Repositioning

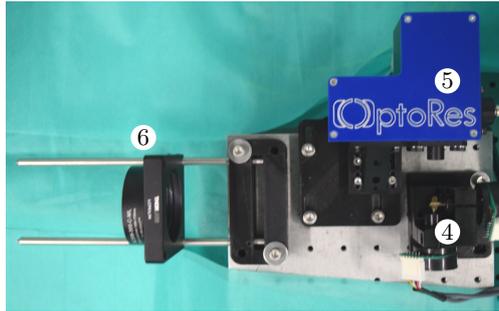
If we change the length of an OCT system's reference arm, then the length of the sample arm, i.e., the axial position of the FOV, changes equally. Therefore, it is intuitive to equip the reference arm with a motorized mirror to automatically adapt the axial FOV position. While our OCT system already provides such a motorized reference arm internally, we use a self-built external reference arm shown in Figure 7.6a, instead, which allows for faster motion. Note that, equal to the internal reference arm, not the actual mirror is mounted to a stepper motor but a retro-reflector. Therefore, the light travels from the collimator to the retro-reflector and then the same distance again to the mirror. Therefore, moving the motor along a distance Δz shifts the FOV by $2\Delta z$. The dynamics of stepper motors, however, are not designed for applications like ours. They are especially slow if they change their direction of travel, which can be expected to happen rather frequently during tracking and is further increased by noise in the results of the tracking algorithm or oscillations caused by the closed control loop.



(a) Moving the retro-reflector (1) with the stepper motor changes the distance to the collimator (2) and the mirror (3).



(b) Second pair of galvos (4) for lateral positioning of the FOV.



(c) The galvo pair (4) shifts the C-scan created by the scan head (5) and the beams are then focussed by an achromatic lens (6).

Figure 7.6: Motorized reference arm (a), second pair of galvos (b), and the setup with scan head, second pair of galvos, and lens (c).

For lateral repositioning of the FOV, we use the same concept as employed in scan heads. A pair of stepper galvos (Figure 7.6b) shifts the whole C-scan in space. The concept is illustrated in Figure 7.7. To focus the beam, we use an achromatic lens with a diameter of 50.8 mm (2 inch) and a focal distance of 300 mm. The first version of the second scanning stage was setup by Schünemann [238]. In comparison to more sophisticated scan lenses, our employed lens provides a rather large diameter, focal distance, and depth of view while having much lower costs. The diameter is a critical parameter because it limits how far we can shift the FOV laterally. The focal length of 300 mm provides us a reasonable distance between the target and the tracking system. Finally, a large depth of view, i.e., range in which the spot size of the beam is reasonably small, is important for axial tracking.

Both galvos and the stepper motor of the reference arm are controlled by an STM32F407 microcontroller (STMicroelectronic, Switzerland) which communicates with the measurement computer via a USB connection. This connection introduces some delay. On average, a ping, i.e., sending a command to the microcontroller and receiving an answer, takes 6 ms and a minimum of 4 ms. In the worst case, the delay can be arbitrarily large because the computer is not running

7.3 Design of an OCT-Based Tracking System

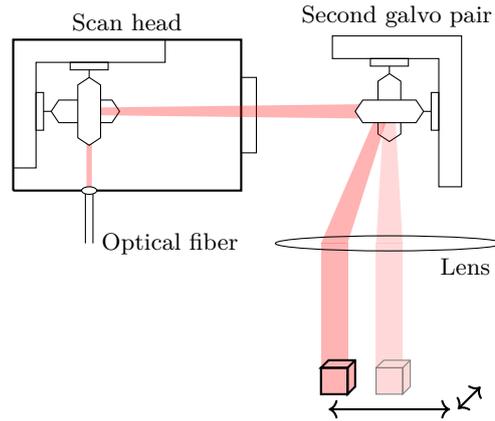


Figure 7.7: Illustration of the second galvo pair in front of the scan head to shift the whole C-scans laterally.

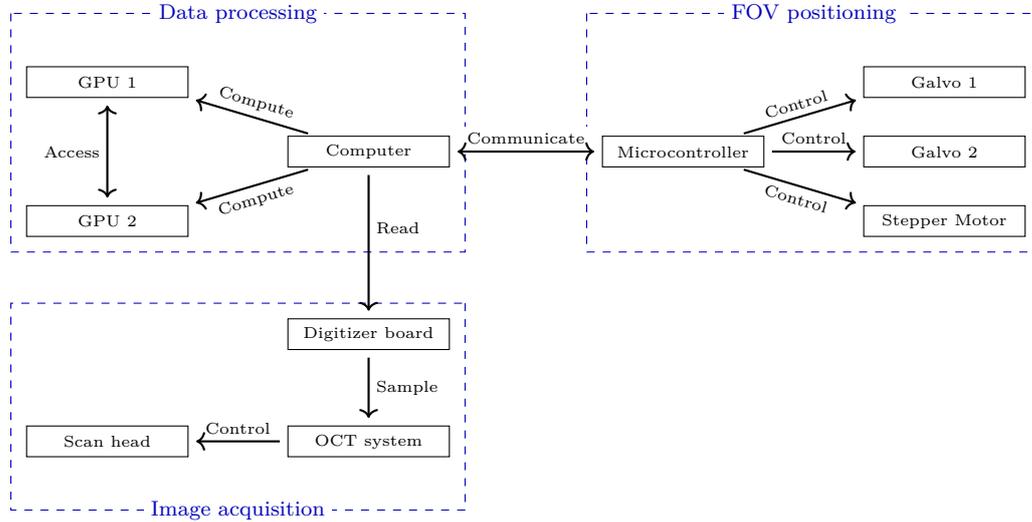


Figure 7.8: Interaction of the tracking system's main components.

a real-time operating system.

7.3.2 Data Processing Pipeline

The data processing pipeline consists of several parallel units shown in Figure 7.9 which asynchronously provide their results to other units. This allows for overlapping computations and furthermore allows for dynamic processing rates of the individual units. The first unit is responsible to organize the transfer of the raw OCT signals from the digitizer board into the computer's main memory.

The raw OCT signals then have to be reconstructed and assembled to C-scans.

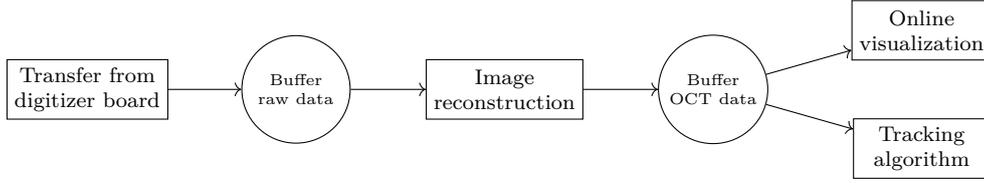


Figure 7.9: Data processing units (rectangles) of the tracking system asynchronously connected by thread-safe buffers (circles).

Image reconstruction is running on one of the GPUs and is implemented in CUDA. While there is a proprietary reconstruction library called OGOP (Optores, Germany) provided, we also implemented our own reconstruction in CUDA. First of all, however, the OCT data has to be transferred from main memory to GPU memory. To avoid costly memory allocations and frees, a class has been implemented which wraps the allocation and free calls and internally maintains lists of CUDA memory areas of different sizes, which have already been allocated but are currently not used. Thereby, actual allocation is only done if the list of memory with the requested size is empty.

The image reconstruction unit gets the raw signals of one C-scan and employs the reconstruction steps. It puts the resulting C-scans, which are still in GPU memory, into another buffer. A small fraction of these C-scans are used by a visualization unit, but the remaining C-scans are available for the tracking algorithm. The output of the tracking algorithm, i.e., the translation $\mathbf{t} = (\Delta x, \Delta y, \Delta z)^\top$ compared to the template, is then used to adapt the position of the FOV. For this purpose, the translation, which is in voxels, has to be converted to steps of the galvos and the stepper motor. At this point, we assume that each axis of the C-scan is proportional to one motor axis, i.e., either one of the galvos or the stepper motor. Furthermore, a proportional gain is applied and we therefore obtain the relative motorsteps \mathbf{m} to compensate the translation \mathbf{t} as

$$\mathbf{m}_{\text{rel}} = \begin{pmatrix} p_{xy} & 0 & 0 \\ 0 & p_{xy} & 0 \\ 0 & 0 & p_z \end{pmatrix} \begin{pmatrix} c_{xy} & 0 & 0 \\ 0 & c_{xy} & 0 \\ 0 & 0 & c_z \end{pmatrix} \mathbf{t}. \quad (7.4)$$

We apply the same gain p_{xy} and conversion factor c_{xy} between voxels and galvo steps to both lateral axes of the C-scan. For the axial axis, we use different values p_z and c_z . Recently, Gessert et al. [81] proposed a deep learning approach for our system which estimates the translation directly in motor steps. While this approach would still require to tune the proportional gains, it could account for the non-ideal mapping between the image axes and the motor axes and thereby improve the feedback loop. However, the online performance in the actual closed-loop system has not been shown yet.

The microcontroller moves the galvos and the stepper motor to a new absolute position \mathbf{m} according to the determined \mathbf{m}_{rel} . Note that we log these absolute

positions to determine the current position of our target. We do not use the calculated translations directly for this purpose. This allows that the microcontroller could discard remaining motion of the stepper motor as soon as it received new motorsteps and this information is not contained in the translations.

7.3.3 Calibration to Cartesian Coordinate System

The quantitative values recorded by our tracking system are the positions of the two galvos and of the stepper motor over time. Firstly, this results in values with no meaningful units. Secondly, the galvos and the stepper motor do not form an orthogonal coordinate system, but we can even expect some spherical effects due to the rotating galvos. Therefore, we introduce a calibration procedure requiring a 3D motion stage to transform the motor positions to a Cartesian coordinate frame with physical units.

The general idea is to use single-template tracking of a phantom which is moved by a motion stage in order to acquire calibration data, i.e., pairs of motor positions and corresponding positions of the motion stage. In principle, this data could be acquired continuously. However, we use a more discrete approach by moving the motion stage to a position, wait there for some time, record the current motor position several times, and then move on to the next position. The waiting allows the tracking system to reach a steady position eliminating any delay occurring during dynamic tracking. By recording the motor position for the same motion-stage position several times, we account for potential residual noise, i.e., the motor position slightly alters although the target is steady. After obtaining M pairs of corresponding positions \mathbf{p}^i of the motion stage and motor positions \mathbf{m}^i , we can establish a spatial transformation between the two coordinate frames. Note that we actually do not need the absolute positions \mathbf{p}^i , but it is sufficient to consider them relative to the initial position to obtain a Cartesian coordinate frame for our tracking system. Thereby, we can freely define which motor steps correspond to the origin.

Typically, spatial calibrations involve estimating a rigid transformation, which consists of a translation and rotation. In our case, we need more degrees of freedom. Firstly, each axis can have an individual scaling from motor steps to millimeter. Secondly, the motor axes are not necessarily orthogonal. Thirdly, the motor steps might not even define a proper right-handed coordinate frame and we might need to allow for reflections as well. These aspects can be represented by an affine transformation, represented as a matrix $A \in \mathbb{R}^{3 \times 4}$. In contrast to rigid transformations, it has full twelve degrees of freedom. The only formal constraint is that the 3×3 sub-matrix formed by the first three columns of A is regular. To obtain the coefficients of A , we can solve

$$\min_A \sum_{i=0}^{M-1} \left\| \mathbf{p}^i - A \begin{bmatrix} \mathbf{m}^i \\ 1 \end{bmatrix} \right\|^2. \quad (7.5)$$

However, as mentioned before, we have to expect spherical distortions due to

the nature of the galvos and the employed lens. These cannot be corrected by an affine transform. Therefore, we propose to additionally fit a function $\mathbf{f}_B : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ having a set of parameters B in a second step. In a least-squares sense, this leads to solving

$$\min_B \sum_{i=0}^{M-1} \left\| \mathbf{p}^i - \mathbf{f}_B \left(A \begin{bmatrix} \mathbf{m}^i \\ 1 \end{bmatrix} \right) \right\|^2. \quad (7.6)$$

We especially propose to fit a quadratic function, which in total has 30 parameters, or a cubic function, which has 57 parameters. Afterwards, we can transform any motor position \mathbf{m} into the corresponding Cartesian motor position

$$\widehat{\mathbf{m}} = \mathbf{f}_B \left(A \begin{bmatrix} \mathbf{m} \\ 1 \end{bmatrix} \right). \quad (7.7)$$

7.3.4 Multi-Template Tracking

If we consider a rigid body, we can track its orientation by observing the motion of at least three non-collinear points of it. Therefore, we propose to employ a multi-template tracking approach to estimate orientations. It essentially consists of applying single-template tracking to at least three spatially separated parts of the same target simultaneously. For this purpose, we define N initial motor positions \mathbf{m}_n^0 , $n = 1, \dots, N$, on a circle around $\mathbf{c} = (x^0, y^0, z^0)^\top$ with radius r as

$$\mathbf{m}_n^0 = \mathbf{c} + \begin{pmatrix} r \cdot \cos \left(2\pi \frac{n-1}{N} \right) \\ r \cdot \sin \left(2\pi \frac{n-1}{N} \right) \\ 0 \end{pmatrix} \quad (7.8)$$

when tracking $N \geq 3$ templates simultaneously. This simple approach is illustrated in Figure 7.10, but it is only applicable if the target is sufficiently flat within r , i.e., the target is within the FOV at every motor position. Otherwise we should replace z^0 with adequate individual values z_n^0 . Note that non-circular arrangements are also possible.

At each initial motor position \mathbf{m}_n^0 , we acquire a template volume which we denote with V_n^0 . Afterwards, we continuously iterate through the positions. If we are in iteration $i \geq 1$ and consider template n , then we first move to \mathbf{m}_n^{i-1} and acquire V_n^i . Second, we calculate the translation \mathbf{t}_n^i between V_n^0 and V_n^i and use it to determine \mathbf{m}_n^i . However, we do not yet move to this new motor position and start iteration $i + 1$. Instead, we repeat the procedure for template $n + 1$ or start iteration $i + 1$ with template 0 in case of $n = N$. Thereby, we implement a multi-template tracking. In fact, we can increase efficiency by pipelining some steps. After we acquired V_n^i , we can immediately move from \mathbf{m}_n^{i-1} to \mathbf{m}_{n+1}^{i-1} and do the calculation of \mathbf{m}_n^i in parallel.

To actually estimate the current orientation in iteration i , we need all N Cartesian motor positions $\widehat{\mathbf{m}}_n^i$. However, these do not represent the same point in time

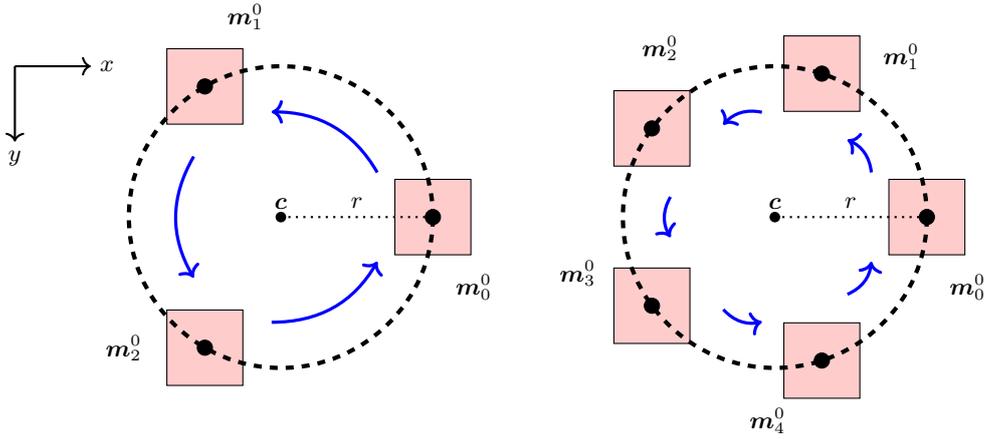


Figure 7.10: Illustration of the arrangement of the initial motor positions for multi-template tracking with $N = 3$ (left) and $N = 5$ (right) templates. They form a circle in the xy -plane.

due to our sequential tracking. Therefore, we also use $\widehat{\mathbf{m}}_n^{i-1}$ and interpolate the motor positions to one, in principle arbitrary, common point in time. We also assign this point in time to our estimated orientation. To actually estimate the orientations, we employ the Kabsch algorithm [115]. It is a closed-form solution to the problem of finding the optimal rotation matrix R , in least-squares sense, between two sets of three-dimensional points, whose centers-of-mass are in the origin. To estimate the orientation represented by our Cartesian motor positions in iteration i w.r.t. the initial motor positions, we therefore first calculate and subtract their centers-of-mass. Afterwards, the algorithm essentially requires to compute the singular value decomposition (SVD) of a matrix, which arises from forming $3 \times N$ matrices from both sets of positions and multiplying them. The product of the two orthogonal matrices from the SVD result is either the rotation R or a reflection. The latter can be converted to the former by correcting the sign of its negative singular value.

7.4 Tracking Algorithms

The cross-correlation between two signals $f(t)$ and $g(t)$ is a measure of their similarity depending on a shift between them. If we consider continuous signals, i.e., $t \in \mathbb{R}$, their cross-correlation is

$$(f \star g)(t) = \int f^*(\tau) \cdot g(\tau + t) d\tau \quad (7.9)$$

where $f^*(t)$ is the complex-conjugate of $f(t)$. The extension for two or more dimensions is straightforward by introducing a shift vector and integrating over all dimensions. In case of discrete signals, the integration reduces to a summation.

If $g(t)$ is a shifted version of $f(t)$, i.e., $g(t) = f(t - \Delta t)$, we can determine the shift Δt by finding the shift which maximizes their cross-correlation. However, this requires to evaluate the cross-correlation for every possible shift leading to extensive computational effort. This effort can be reduced by employing Fourier-domain methods. Because cross-correlating $f(t)$ and $g(t)$ is equal to convolving $f^*(-t)$ and $g(t)$, we can employ a fast convolution, i.e., a multiplication in Fourier domain, instead of direct evaluation of (7.9) in spatial domain. Therefore, methods exploiting this relationship are very interesting for applications which require high processing rates.

One approach to estimate translations between images is the phase correlation method which is based on evaluating cross-correlation in Fourier domain [132]. In its basic form, it assumes purely translational motion and can be implemented very efficiently. We can define one image as the template and then determine the translations of all following images with respect to that template. Phase correlation, however, can be expected to fail as soon as the appearance of the target severely changes in the image data because it will no longer be sufficiently similar to the template. Such large changes have to be expected if the target's motion cannot be approximated by translations but it substantially rotates, which cannot be compensated by shifting the position of the FOV. For this reason, Glandorf [86] proposed to employ the MOSSE filter [22] in this tracking setup instead. This algorithm is a special kind of correlation filter with a concept very similar to phase correlation, but it allows for a more general definition of the template and furthermore allows to update the template during tracking. Thereby, this filter can adapt to appearance changes over time. However, template updates always have the risk to cause drifting effects, i.e., after some time the current template does not correspond to the original target anymore but represents a different structure [174].

7.4.1 Phase Correlation

Consider a signal $f(\mathbf{x})$, $\mathbf{x} \in \mathbb{R}^D$, and its shifted version $g(\mathbf{x}) = f(\mathbf{x} - \Delta\mathbf{x})$. Let $F(\mathbf{k})$ be the D -dimensional Fourier transform of $f(\mathbf{x})$. From the Fourier shift theorem, we obtain

$$G(\mathbf{k}) = F(\mathbf{k}) \cdot e^{-i2\pi\mathbf{k}^\top \Delta\mathbf{x}} \quad (7.10)$$

as the Fourier transform of $g(\mathbf{x})$, where i is the imaginary unit. The so-called normalized cross-power spectrum

$$R_{f,g}(\mathbf{k}) = \frac{F(\mathbf{k}) \cdot G^*(\mathbf{k})}{|F(\mathbf{k}) \cdot G^*(\mathbf{k})|} \quad (7.11)$$

is the normalized Fourier transform of the signals' cross correlation. Due to (7.10), in our case we can write

$$\begin{aligned}
 R_{f,g}(\mathbf{k}) &= \frac{F(\mathbf{k}) \cdot F^*(\mathbf{k}) \cdot e^{i2\pi\mathbf{k}^\top \Delta\mathbf{x}}}{|F(\mathbf{k}) \cdot F^*(\mathbf{k}) \cdot e^{i2\pi\mathbf{k}^\top \Delta\mathbf{x}}|} \\
 &= \frac{F(\mathbf{k}) \cdot F^*(\mathbf{k}) \cdot e^{i2\pi\mathbf{k}^\top \Delta\mathbf{x}}}{|F(\mathbf{k}) \cdot F^*(\mathbf{k})|} \\
 &= e^{i2\pi\mathbf{k}^\top \Delta\mathbf{x}}
 \end{aligned} \tag{7.12}$$

and see that the result is a pure phase term which depends on the shift $\Delta\mathbf{x}$. By applying the inverse Fourier transform, we obtain

$$r_{f,g}(\mathbf{x}) = \delta(\mathbf{x} + \Delta\mathbf{x}), \tag{7.13}$$

i.e., a Dirac peak at the position which corresponds to the shift of the signals.

This algorithm is known as the phase-correlation method [132]. In practical image processing, it has several issues. We typically apply it to discrete image data and therefore the derivation would require that all shifts are cyclic. This is not the case in reality, however, and introduces distortions. Additionally, a shift is typically not aligned with the image grid and extensions of the algorithm are necessary to have sub-pixel accuracy [67]. Furthermore, we typically do not have images which are related by a pure translation and therefore we obtain further distortions in the result due to the additional transformations missing in the model.

The basic form of phase correlation is computationally very efficient. If we assume using a fixed template image, we need only two (multi-dimensional) FFTs, simple element-wise operations, and a search for the location of the maximum intensity to determine the shift between the template and another image. For the FFT, efficient implementations exist and element-wise operation can easily be parallelized on, for example, a GPU. The search for the location of the maximum intensity can be parallelized by employing the parallel-reduction design pattern. In our implementation, we additionally apply a Hann-shaped low pass to $R_{f,g}(\mathbf{k})$ to increase robustness.

Phase correlation has not only been extended for sub-pixel accuracy. In case of two-dimensional images, we can also use it to estimate rotations [33] and scaling [208]. Reddy et al. [208] showed that both transformations can be obtained via phase correlation with deterministic computational effort after applying suitable coordinate transformations. There is also an extension which allows to determine rotations of 3D images [184]. However, this involves an iterative estimation of the parameters and furthermore assumes 3D objects which can be represented by polygon meshes.

7.4.2 MOSSE Filter

The general setup for correlation filters is a large image which contains a small target to be tracked. Classical correlation filters construct an optimal filter which

we then apply to smaller parts, i.e., sub-images, of the large image. The filter should only output a high value for the sub-image which corresponds to the target. Thus, the natural output of the filter is a single correlation-like value for each sub-image. The optimal filter is constructed from a training set consisting of matching and non-matching small input images and corresponding training output values. There are many different variants of correlation filters, which mainly differ in the way the filter is constructed and the constraints it has to fulfill [138, 171, 172, 209].

This general approach has been substantially extended by the so-called ASEF filters which allow to specify training output not only as a single correlation value but as a full cross-correlation-like map [23]. This allows to use large images directly as the training input and to define a peak in the output map at the position of the target within this image. The actual application of the constructed filter $h(\mathbf{x})$ is done in frequency domain. If we apply the filter to an image $f(\mathbf{x})$, we obtain the frequency-domain output as

$$O(\mathbf{k}) = F(\mathbf{k}) \cdot H^*(\mathbf{k}) \quad (7.14)$$

and in spatial domain it should show a peak at the position of the target. This operation is very similar to the operation (7.11) and the result (7.13) in the phase-correlation method. However, correlation filters allow for much more flexibility due to the training process which can consider arbitrarily many training images $f_j(\mathbf{x})$. These can, for example, be images of the same object but with different noise or even under different conditions like illumination. Each training output $o_j(\mathbf{x})$ would still be a peak at the same location in this case. However, we can also train the filter with images in which the target is located at different positions. This is reflected by the training outputs in which the peak is shifted accordingly. After training, applying the filter to a new image has the same low computational effort as phase correlation.

One particular variant specifically designed for tracking tasks is the MOSSE (minimum output sum of squared error) filter [22]. We will employ this variant in this thesis. The training outputs are typically not defined as ideal peaks, but as Gaussians

$$o_j(\mathbf{x}) = \frac{1}{(\sqrt{2\pi}\sigma)^D} \cdot \exp\left(-\frac{(\mathbf{s}_j - \mathbf{x})^\top(\mathbf{s}_j - \mathbf{x})}{2\sigma^2}\right) \quad (7.15)$$

with standard deviation σ and shift \mathbf{s}_j . The MOSSE filter computes the optimal filter, in the discrete case, by minimizing the sum over all training images of the summed squared differences between the actual and the target output. The closed-form solution to this optimization problem presented by Bolme et al. [22] is

$$H^*(\mathbf{k}) = \frac{\sum_j O_j(\mathbf{k}) \cdot F_j^*(\mathbf{k})}{\sum_j F_j(\mathbf{k}) \cdot F_j^*(\mathbf{k})}, \quad (7.16)$$

but typically a regularization constant λ is added to the denominator.

7.5 Experimental Setup for System Characterization

Training the filter with many images helps to increase its robustness. However, there is also a simple extension which allows to adapt the filter online during tracking. Let upper indices denote images obtained during tracking in contrast to lower indices for the training images. We obtain our initial filter $H^0(\mathbf{k})$ from the training process (7.16). If we separate in iteration i the current complex-conjugated filter $H^{i*}(\mathbf{k})$ into its numerator, $A^i(\mathbf{k})$, and its denominator, $B^i(\mathbf{k})$, we can update the filter with a learning rate η via

$$A^i(\mathbf{k}) = \eta \cdot \left(O^i(\mathbf{k}) \cdot F^{i*}(\mathbf{k}) \right) + (1 - \eta) \cdot A^{i-1}(\mathbf{k}) \quad (7.17)$$

and

$$B^i(\mathbf{k}) = \eta \cdot \left(F^i(\mathbf{k}) \cdot F^{i*}(\mathbf{k}) + \lambda \right) + (1 - \eta) \cdot B^{i-1}(\mathbf{k}), \quad (7.18)$$

where $O^i(\mathbf{k})$ is a Gaussian with a shift corresponding to the result when applying $H^{i-1}(\mathbf{k})$ to $F^i(\mathbf{k})$. This highlights the advantage of allowing shifted training images, because during online tracking we will typically not find the target centered.

7.4.3 Deep Learning Methods

Recently, different methods employing deep learning have been proposed in the context of our proposed tracking setup [17, 18, 81]. The studies investigate both 3D and 4D image data as input for tracking, where 4D refers to processing of small temporal sequences of C-scans. However, they have only been studied for offline estimation of 3D motion up to now. While the results show promising accuracy and the methods also allow to directly estimate the motion in steps of the galvos and the stepper motor, their generalization is limited so far and requires to train the underlying networks on quite specific structures. Therefore, it is still an open question how markerless these methods actually are in practice. Furthermore, the complexity of the networks does not necessarily allow for processing rates as achievable with methods like phase correlation, yet. We do not consider these approaches in this thesis, but only employ the described classical correlation-based methods.

7.5 Experimental Setup for System Characterization

To characterize the proposed tracking system, we employ the following methodology. We will consider different phantoms for tracking. Motion is simulated using a robot arm, which allows to drive predefined trajectories reproducibly and accurately. Therefore, we are able to derive quantitative values by comparing the trajectories estimated by our tracking system with the actual trajectories driven by the robot arm.

7.5.1 Phantoms

As phantoms for our tracking experiments, we employ thin plates which have been 3D printed and are made of photopolymer resin. They have a thickness of about

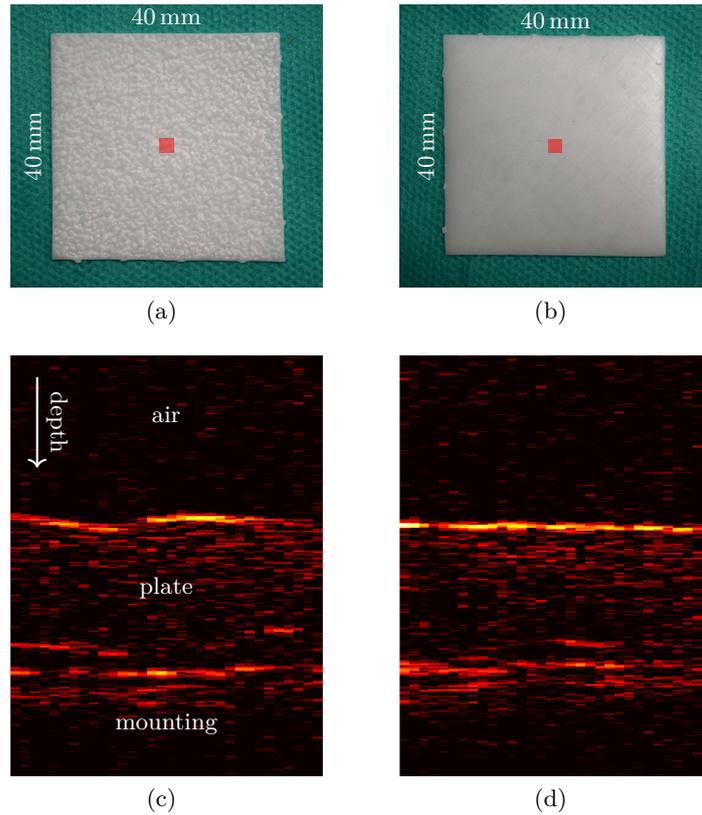


Figure 7.11: Photographies (top row) and B-scans (bottom row) of the structured-surface phantom (left column) and the flat-surface phantom (right column). Red squares illustrate the size of the OCT's lateral FOV which is $2.5 \times 2.5 \times 3.5 \text{ mm}^3$.

1 mm and a lateral size of $40 \times 40 \text{ mm}^2$. As Figure 7.11 shows, both the front- and the backside of a plate are visible in the OCT images. We consider two different plates. One of them has a fine random structure, i.e., height variations of few hundreds of micrometers, on both sides. The other one only has one structured side while the other side is flat. We use the flat side as the front side and the structured side as the backside during tracking, i.e., there is visually no information on the surface from the tracking system's point of view.

Furthermore, we use samples of chicken breast to simulate tracking of tissue. In Figure 7.12, an exemplary chicken-breast sample is attached to a sample holder. We arbitrarily select spots on the chicken breast for tracking and only adjust the axial position in order to have the surface properly within the OCT's FOV.

7.5 Experimental Setup for System Characterization



Figure 7.12: View on chicken breast attached to a sample holder.

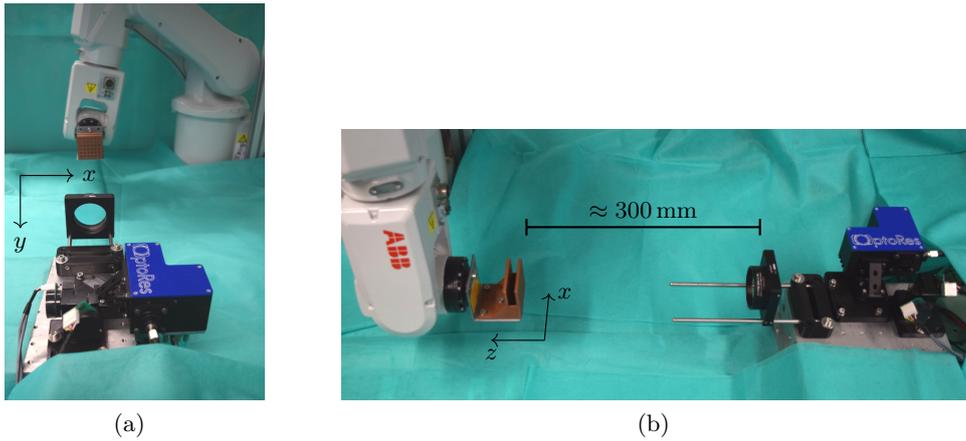


Figure 7.13: Front view (a) and side view (b) on the robot arm with attached sample holder in front of the tracking system. The two galvos realize moving the FOV along x and y , respectively, and the motorized reference arm (not shown) moves the FOV along z .

7.5.2 Simulating Motion

We simulate motion with a 6-axis robot arm (IRB 120, ABB, Switzerland), which has a repeatability of 0.01 mm. For simplicity, we also use this robot arm for the Cartesian calibration of the tracking system, although it is not necessary to use the same system for both and a 3D motion stage would be sufficient to perform the calibration. We mount a sample holder for the phantoms to the end-effector of the robot arm as shown in Figure 7.13. The end-effector is positioned in front of the tracking system's lens such that the surface of the sample is close to the lens' focal spot. Its orientation is approximately aligned to the axes of the C-scans and of the coordinate frame defined by the two galvos and the stepper motor. This allows to evaluate the different axes separately by only moving the corresponding robot axes.

If we simulate translational motion, we use the current position of the end-effector, define a destination position, and a define target velocity. This is the

maximum velocity in Cartesian space, i.e., the constant velocity we obtain after the necessary acceleration phase and before deceleration. Afterwards, we directly move to the next position until we reach the experiment's defined total duration. We will on the one hand consider motion back and forth along a line. This line is either along one robot axis, the diagonal of a square spanned by two axes, or the diagonal of the cube spanned by all three axes. In particular, we will distinguish lateral motion, i.e., motion at constant depth from the OCT's point of view, and axial motion, i.e., motion along the OCT's depth axis. On the other hand, we will consider random motion in the sense that we define a random sequence of positions and drive from one to the next with the pre-defined maximum velocity.

For rotational motion, we define the maximum angular velocity. Analogously to the translations along a line, we consider rotations around two fixed axes to evaluate different components. The first axis of rotation is the axial OCT axis. Rotating around this axis can be described as in-plane rotations with respect to the OCT's lateral plane. The second axis equals one of the lateral axes and therefore results in out-of-plane rotations. The first should only require moving the galvos, while the second should involve one galvo and the stepper motor.

Furthermore, we consider random rotations. For this purpose, we define a sequence of random axes of rotation and random angles. They are always relative to the initial orientation of the end-effector. Thereby, we can define a maximum angle α_{\max} and draw the actual angles uniformly from the interval $[0, \alpha_{\max}]$. To create a random axis, we can define a vector $\mathbf{a} = (x, y, z)^T$, draw each of its elements from a Gaussian distribution, and use $\mathbf{a}/\|\mathbf{a}\|$ as the axis of rotation. Because we can rotate around any axis in space, drawing non-negative angles is sufficient.

If we consider 6D motion, we define the translational and rotational part independently. Therefore, we obtain a sequence of positions and a sequence of orientations and realize them in parallel.

7.5.3 Experiments

For our tracking experiments, we employ the fast scan head. It provides imaging with 32×32 A-scans per C-scans at a volume rate of 831 Hz. The FOV is about $2.5 \times 2.5 \times 3.5 \text{ mm}^3$ when considering 480 pixels along the depth direction. We employ the slower scan head only for evaluating the performance of the phase correlation implementation.

Before the actual tracking experiments, we have to perform a calibration to Cartesian coordinates to allow for a quantitative evaluation of the tracking performance. For this purpose, we move the robot arm at a speed of 2 mm s^{-1} to the calibration positions while employing single-template tracking with phase correlation.

Each experiment consists of 45s with active tracking. During this time, the robot arm moves continuously and we log its poses and the related timestamps. For each repetition of an experiment with the same parameters, we reposition the target in front of the tracking system. Thereby, we track different parts of the

target in each repetition.

For single-template tracking, we log the motor positions at about 80 Hz asynchronously to tracking and employ phase correlation. The template C-scan is acquired before the robot arm starts to move, i.e., while the target is static.

For multi-template tracking, we log the motor positions for each of the N tracked templates directly. At each template-defining motor position \mathbf{m}_n^0 , $n = 1, \dots, N$, we initially train one MOSSE filter on ten C-scans of the static target. Therefore, the training outputs are centered smoothed peaks following (7.15) and the training mainly accounts for the noise during static OCT acquisition. Subsequently, we will update the filters each time we apply them. When we acquire the new C-scan in iteration $i \geq 1$ at motor position \mathbf{m}_n^i , we apply the corresponding filter. This yields us the translation which we use to determine \mathbf{m}_n^{i+1} and which we also use for updating the filter by converting it to the required shifted filter output in (7.17).

7.5.4 Evaluation Methods

To quantitatively evaluate tracking, the trajectories of the tracking system and of the reference system, in our case the robot arm, have to be transformed into the same coordinate frame. For purely translational motion, we can simply use the result of the calibration to Cartesian coordinates and evaluate the relative motion performed by the robot arm and the relative motion estimated by the tracking system. If we want to evaluate full 6D motion, we also have to describe the rotations in a common coordinate frame. The problem of finding the required transformations is commonly referred to as a hand-eye calibration. The term *hand* refers to the dynamic system, typically some kind of robot, realizing motion and providing ground-truth information. The term *eye* refers to some kind of static camera system. The transformations are assumed to be rigid, i.e., we can describe them with homogeneous matrices

$$T = \begin{bmatrix} R & \mathbf{t} \\ \mathbf{0}^\top & 1 \end{bmatrix} \quad (7.19)$$

with a rotation matrix $R \in \mathbb{R}^{3 \times 3}$, a translation vector $\mathbf{t} \in \mathbb{R}^3$, and a vector $\mathbf{0} \in \mathbb{R}^3$ only containing zeros. If we realize several poses ${}^B T_E^{(i)}$ of the end-effector w.r.t to the base coordinate frame of the robot and measure the corresponding poses ${}^C T_M^{(i)}$ of some suitable marker w.r.t. the camera coordinate frame, we obtain equations

$${}^B T_E^{(i)} {}^E T_M = {}^B T_C {}^C T_M^{(i)}. \quad (7.20)$$

when considering the chain of transformations illustrated in Figure 7.14. The constant spatial relations between the marker and the end-effector ${}^E T_M$ and between the camera and the robot base ${}^B T_C$ are the unknown parameters we need to solve for. Common algorithms for this problem exploit the structure of the transformations and separate estimation of the translational and the rotational parts [198, 265]. There are other approaches which, for example, ignore the structure of the

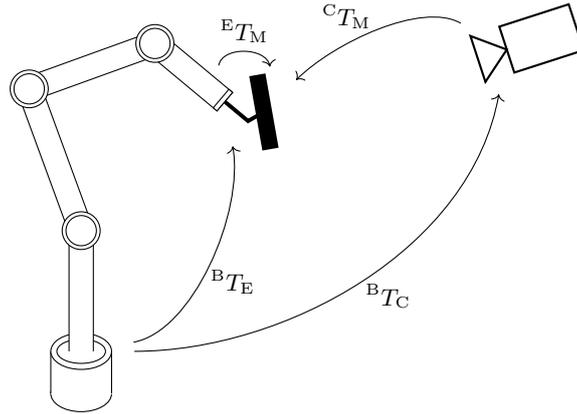


Figure 7.14: Sketch of the general hand-eye calibration setup. It involves the pose of the end-effector frame w.r.t. the robot base frame ${}^B T_E$, the marker frame w.r.t. the end-effector frame ${}^E T_M$, the marker frame w.r.t. the camera frame ${}^C T_M$, and the camera frame w.r.t. the robot base frame ${}^B T_C$.

transformations in order to simplify the problem and to account for real-world deviations from the ideal rigidity [56].

There are studies describing hand-eye calibrations involving OCT imaging. For example, Rajput et al. [206] proposed a setup in which an OCT scan head can be calibrated to both a robot and a Kinect RGBd camera (Microsoft, USA). Formally, this has been an eye-in-hand setup as the scan head was mounted onto the robot to move the OCT's FOV in space. Zhou et al. [286] considered needle-based ophthalmologic interventions. They used localization of such a needle to develop a calibration of a microscope-integrated OCT system to an ophthalmologic robot, which can be done in the beginning of an intervention. One major difference of such studies to our setup is the employed C-scan grid. With denser A-scan sampling, localization becomes precise which is crucial for the calibration results. Our lateral scanning grid is rather coarse, in contrast. A second major difference is that we do not consider calibration to an image coordinate frame but to the frame defined by the motor positions.

We decide to not employ a full hand-eye calibration at this point. The calibration data would have to cover the whole volume in which we track motion later on and should, furthermore, cover substantial rotations. In principle, we could acquire such calibration data similar to the calibration to Cartesian coordinates, but using multi-template tracking. This approach, however, has the severe issue that the residual calibration error will already include the tracking error. This is critical when we actually want to use the calibration result to evaluate the tracking accuracy. Typically, the poses of the marker in camera coordinates should be estimated by a different method which has a substantially smaller error than the method to be evaluated. To directly obtain quantitative results without designing a new calibration procedure, we evaluate tracking of rotations only by using the

axis-angle representation of orientations instead, i.e., we describe an orientation relative to another orientation by one angle around one axis of rotation. In this representation, the absolute value of the angle is independent of the coordinate systems.

As a quantitative measure, we will mostly consider the root-mean-square error (RMSE), which is

$$\text{RMSE}(e) = \sqrt{\frac{1}{N} \sum_{i=1}^N e_i^2} \quad (7.21)$$

for a vector $e = (e_1, \dots, e_N)^\top$ containing N error values. Its unit is equal to the unit of the underlying errors. We apply the RMSE to either the norm of the difference between the relative positions recorded by the robot arm and the tracking system or to the difference of their axis-angle angles relative to the initial orientation. We use timestamps to resample the poses smoothly logged by the robot arm to the poses logged by the tracking system.

7.6 Results

Before presenting the results of the actual tracking experiments, we firstly evaluate the performance of our implementation of the phase-correlation algorithm. We then show results for the Cartesian calibration of the system. Subsequently, we evaluate single-template tracking for purely translational motion and multi-template tracking for 6D motion.

7.6.1 Phase Correlation

We see from Table 7.1 that our implementation of the phase correlation is sufficiently fast to handle all relevant volume sizes of our two scan heads under consideration. The average computational times are 712 μs , 1863 μs , and 6568 μs for 32×32 , 64×64 , and 116×116 A-scans per C-scan, respectively, but were measured on an idle system. If we consider actual tracking with the fast scan head, i.e., having overhead especially due to copying the raw data to the GPU and reconstructing the A-scans, we only achieve a full tracking rate of 821 Hz and visualization rate of 10 Hz when employing our own implementation of the image reconstruction. Otherwise, we achieve a tracking rate slightly above 600 Hz.

7.6.2 Cartesian Calibration

Our calibration to a Cartesian coordinate system has two major parameters, the type of the fitted function and the number of positions which we record. Figure 7.15 compares the residual calibration errors and the errors on a test set for different settings. The calibration positions are within a cylinder with diameter and depth of 40 mm. The test set's positions are only within a diameter and depth

Table 7.1: Achievable volume rates for our phase correlation implementation. Measurements are based on evaluation of 100 000 volumes without further computations in parallel. For reference, the volume rates of the 11.3 kHz and the 30 kHz scan head are given for the same lateral A-scan grid sizes. Note that the fast scanner is limited to about 45 A-scans per B-scan due to (7.2).

Voxel grid	$32 \times 32 \times 480$	$64 \times 64 \times 480$	$116 \times 116 \times 480$
Phase correlation	1404 Hz	537 Hz	152 Hz
30 kHz scanner	831 Hz	–	–
11.3 kHz scanner	315 Hz	167 Hz	94 Hz

of 30 mm. Note that there might be test positions which are close to calibration positions.

When fitting only the affine transformation, we observe larger residual errors than errors on the test set. This indicates that especially the outer positions, which are not covered by the test set, cannot be adequately described by the affine transformation. The additionally fitted quadratic decreases the residual errors to less than 0.1 mm. However, we observe clear overfitting when using 36 positions. For all other numbers, the testing errors do not exceed 0.1 mm. If we replace the quadratic by a cubic function, we need 154 positions to avoid overfitting and to actually obtain testing errors which are lower than for the quadratic function. The resulting errors, however, are only slightly lower. Therefore, we decide to use the calibration strategy with 80 positions and fitting of an affine transformation and a quadratic function for evaluation of all following experiments.

7.6.3 Single-Template Tracking

We manually tuned the proportional gains in (7.4) to achieve small tracking errors on all axes over a wide range of speeds and present the results in Figure 7.16. We employ identical settings to tracking of the structured-surface plate and to tracking of chicken breast samples. Considering linear motion along x and y (Figures 7.16a and 7.16b), we observe in general similar RMSEs for the structured-surface plate and samples of chicken breast. For the plate, however, the results are very consistent and we do not observe any failed tracking, which we define as an RMSE above 1 mm, for speeds up to 40 mm s^{-1} . In contrast, we observe for the chicken samples some outliers at 10 mm s^{-1} and 25 mm s^{-1} and tracking failures start to occur at 17.5 mm s^{-1} . Nevertheless, the errors of the majority of the successful trials remain below about 0.2 mm as for the plate.

For motion along z (Figures 7.16c and 7.16d), we obtain RMSEs in the order of 0.2 mm only for speeds up to 25 mm s^{-1} . Beyond that speed, the errors starts to almost linearly increase for the plate. For the chicken samples, tracking already always fails for 40 mm s^{-1} . Furthermore, the errors for the lower speeds are higher than for motion along xy . Note, however, that the effective speed per axis is lower for xy motion than for z motion. An exemplary trajectory for motion along z is

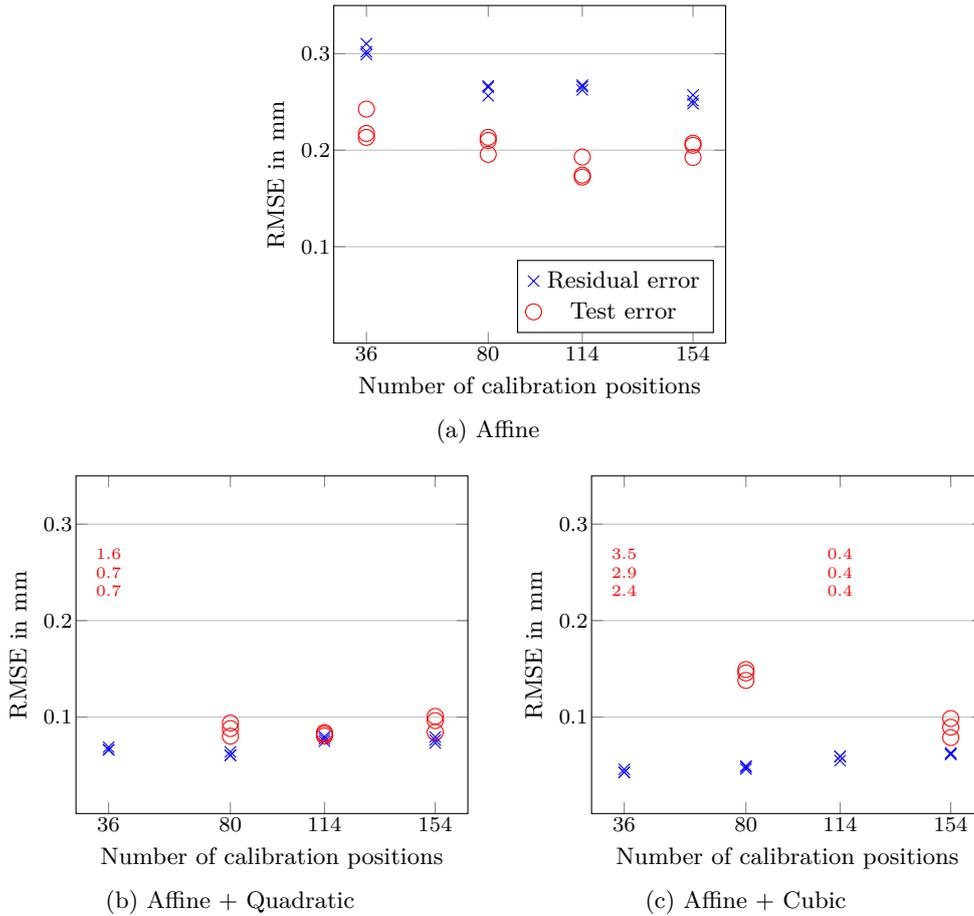


Figure 7.15: Residual RMSEs (blue crosses) of the Cartesian calibration and RMSEs on a test set consisting of 252 positions (red circles). The positions are within a cylinder with equal diameter and depth, which are 40 mm and 30 mm for calibration and testing, respectively. Three calibrations for each number of positions were performed and calibration involved fitting an affine transform (a) and fitting additionally a quadratic (b) or cubic (c) function. RMSE values above 0.35 mm are printed as text.

shown in Figure 7.17.

When considering linear motion along all three axes (Figures 7.16e and 7.16f), we mainly observe RMSEs up to about 0.2 mm for up to 40 mm s^{-1} . For the chicken samples, however, the results are again less consistent and we have some higher errors and few tracking failures. Beyond 40 mm s^{-1} , the errors are clearly increasing. For random motion along all axes as illustrated in Figure 7.18, the error is worse than for the linear motion (Figures 7.16g and 7.16h). We observe tracking failures for lower speeds and the errors start to increase earlier. Nevertheless, the plate can still be tracked without failures for up to 40 mm s^{-1} and successful trials at 40 mm s^{-1} have errors about 0.3 mm.

7 Towards Optical Coherence Tomography for Image Guidance in Cranial Radiosurgery

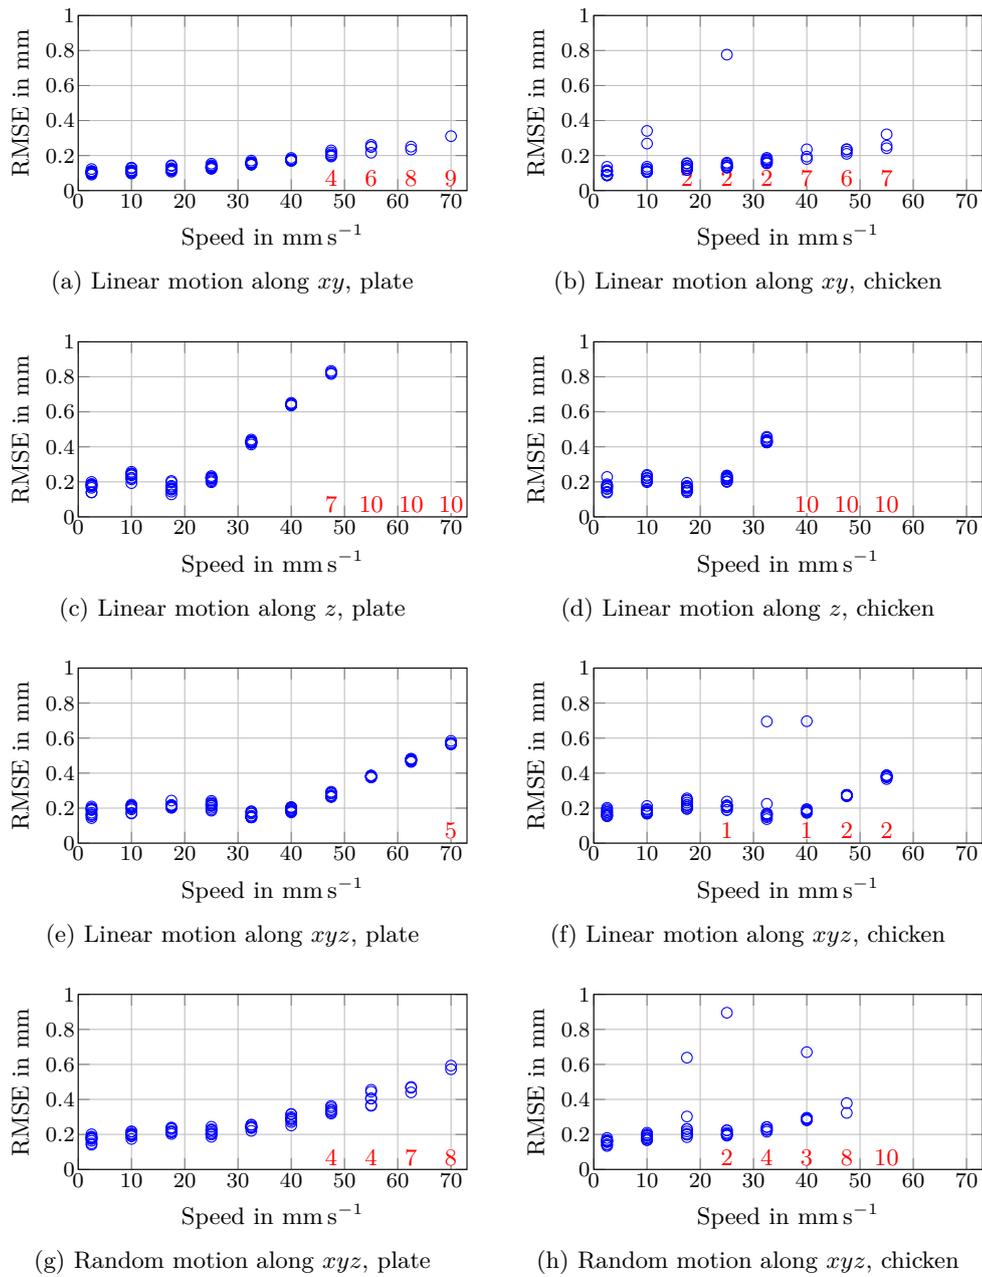
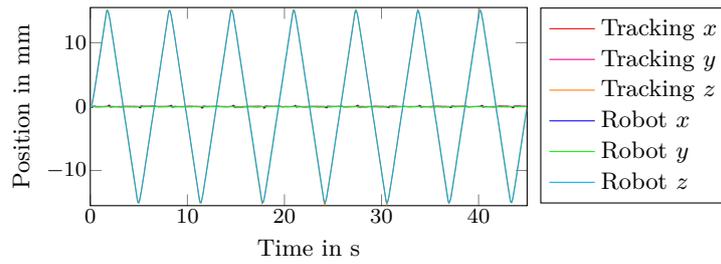
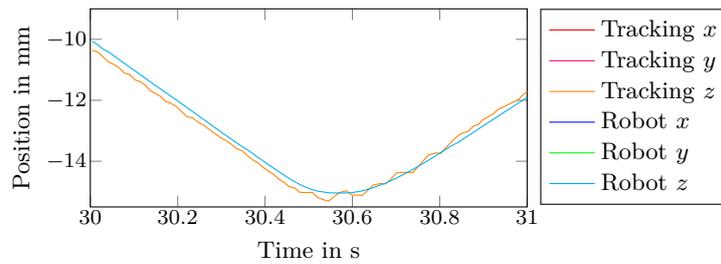
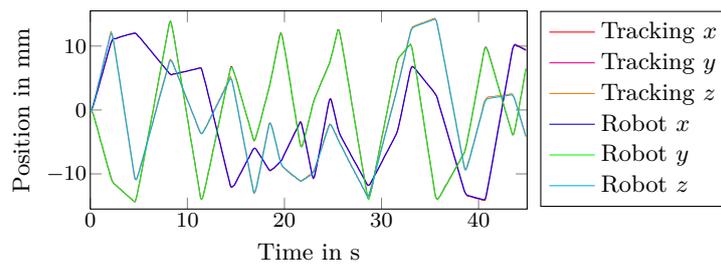
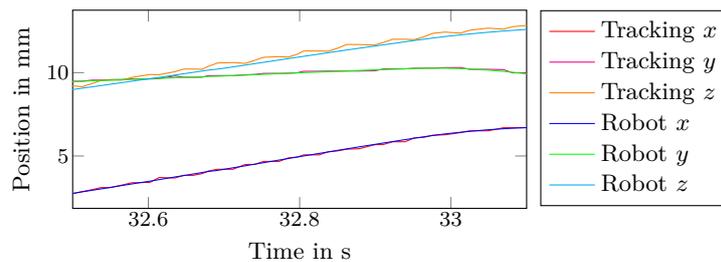


Figure 7.16: Results for tracking of the structured-surface plate phantom (left column) and samples of chicken breast (right column) at different speeds. Linear motion was along a line of length 30 mm and random motion was within a cube with side length of 30 mm. We recorded ten repetitions for each velocity. Experiments with an RMSE above 1 mm are considered as failures and not shown. Instead, their number is printed in red.

(a) Linear motion along z 

(b) Zoom into (a)

Figure 7.17: Linear motion trajectory of a chicken-breast sample at 10 mm s^{-1} along the axial z -axis. Therefore, the x and y components are almost zero. We show the full trajectory (a) and a zoom into a turning point (b).

(a) Random motion along xyz 

(b) Zoom into (c)

Figure 7.18: Random motion trajectory of a chicken-breast sample at 10 mm s^{-1} involving all axes. We show the full trajectory (a) and a zoom (b).

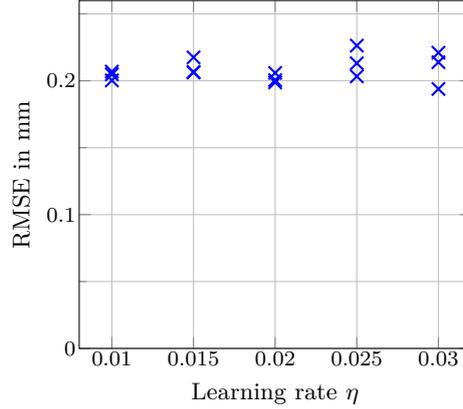


Figure 7.19: RMSEs depending on the learning rate η for translational motion with 12 mm s^{-1} . The trajectory is a line covering all three axes.

7.6.4 Multi-Template Tracking

While the single-template tracking with phase correlation is limited to 3D motion, our proposed multi-template tracking with the adaptive MOSSE filter allows to track 6D motion. Before the experiments, we again manually tune the proportional gains in (7.4) and keep these parameters for all experiments. We mainly consider tracking the structured-surface phantom (Figure 7.11a) and employ as the baseline parameters for multi-template tracking (compare Figure 7.10) $N = 3$ templates arranged at a radius of $r = 200$ in motor steps, which corresponds to about 5.4 mm. Figure 7.19 shows that the results obtained with multi-template tracking and the adaptive MOSSE filter for purely translational motion are comparable to the results we obtained in the previous sub-section for single-template tracking. The evaluated learning rates η do not seem to impact the results. Note that $\eta = 0$, i.e., a conventional non-adaptive MOSSE filter, is not able to track any rotations with substantial angle and therefore not considered in this section. In the following, we first analyze tracking of purely rotational motion. We separate this analysis into in-plane and out-of-plane rotations. During in-plane rotations, the phantom rotates approximately around the OCT's z -axis, i.e., within the xy -plane. During out-of-plane rotations, the rotation is around one of the OCT's lateral axes. Second, we analyze true 6D motion.

In-Plane Rotations

In-plane rotations keep the phantom orientated almost parallel to the lateral plane of the C-scans. Therefore, no major motion of the motorized reference arm is necessary for compensation. For the baseline parameters and different learning rates η , we obtain the results shown in Figure 7.20. When rotating at 4° s^{-1} , we obtain for $\eta = 0.01$ RMSEs in the order of 0.2° for both maximum rotation angles of 8° and 13° . Higher learning rates slightly increase the errors to about

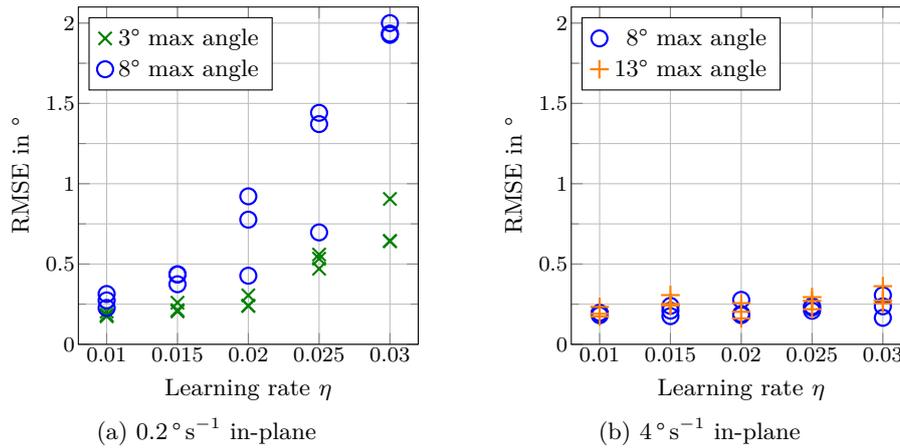


Figure 7.20: Results for in-plane rotations with different speeds and maximum angles depending on the learning rate η .

0.3° for a maximum angle of 13° . In contrast, we observe a much higher influence of the learning rate for slow rotations at only 0.2° s^{-1} . While the errors for $\eta = 0.01$ are around 0.25° , a rate of $\eta = 0.025$ increases them to about 0.5° for a maximum angle of 3° and even up to almost 1.5° for a maximum angle of 8° . Note that we do not consider slow rotations of 13° because the duration of the experiments is too short to reach the turning point. The increased error arises from a systematic underestimation of the actual rotation, as exemplified in Figure 7.21. The deviation of the estimation from the actual angle realized by the robot arm increases with increasing learning rate. The general progress, however, is typically reflected correctly by the estimated trajectory.

In Figure 7.22, we present the results for in-plane rotations of the flat-surface phantom (Figure 7.11a). Although this phantom does not provide superficially visible variations, the results for $\eta = 0.01$ are similar to those for the structured-surface phantom. However, setting $\eta = 0.02$ provides better results for the faster rotation at 4° s^{-1} .

Out-of-Plane Rotations

Out-of-plane rotation of the structured-surface phantom not only requires moving the galvos for compensation but also substantial motion of the motorized reference arm. For single-template tracking, we saw that this is a limiting component in our setup. The out-of-plane results in Figure 7.23 are clearly worse than the in-plane results. The RMSEs are rather in the order of 1° and the trackable rotation speed is lower. Furthermore, we need a rather high learning rate to successfully track at 1.5° s^{-1} . However, the learning rate has no clear impact on the results for slow rotations. The increased errors are due to a systematic overestimation of the rotation angle. Figure 7.24 shows that there is a steady offset for angles

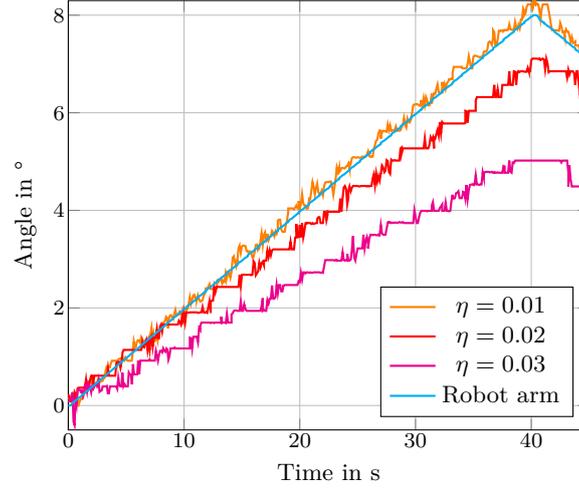


Figure 7.21: Exemplary illustration of the underestimation of the actual angle caused by higher learning rates η . Data corresponds to in-plane rotation with 0.2°s^{-1} and a maximum angle of 8° .

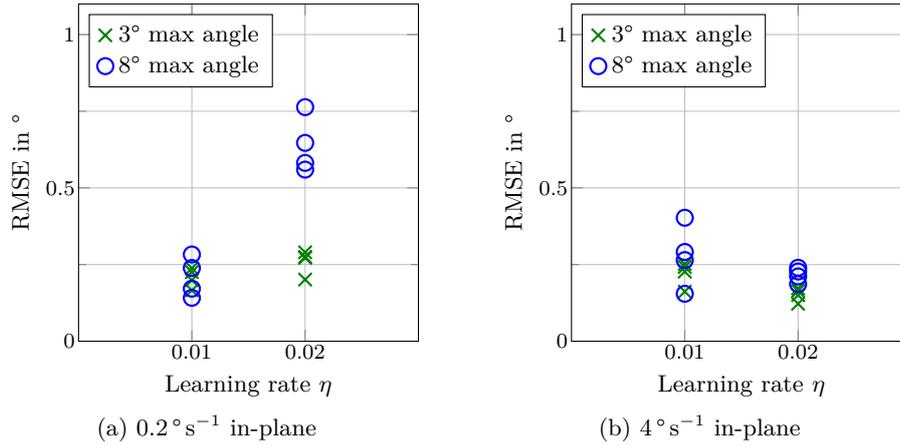


Figure 7.22: Results for in-plane rotations with the flat-surface phantom for selected learning rates η .

above 2° resulting in similar RMSEs for different speeds and maximum angles. The qualitative shapes of the angle trajectories are estimated correctly, however.

Up to now, we only considered our baseline number of templates $N = 3$ and radius $r = 200$. We see in Figure 7.25 that increasing r decreases the RMSE for the out-of-plane rotations. However, the diameter of our scan lens limits the range in which we can move the C-scan with the second pair of galvos. Therefore, increasing the radius reduces the amplitude of motion which we can follow. If we set $r = 300$ and additionally set the number of templates to $N = 7$, we obtain the results in Table 7.2. We can track rotations with a maximum angle of 3° at 0.2°s^{-1}

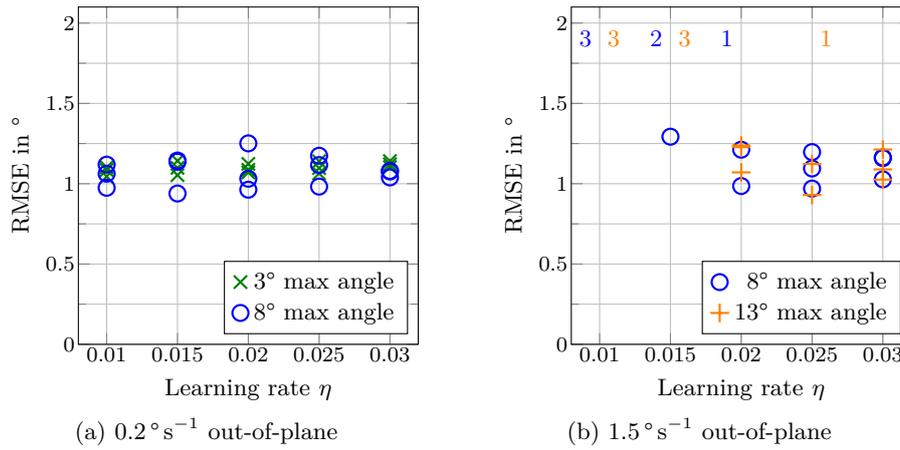


Figure 7.23: Results for out-of-plane rotations depending on the learning rate η for different speeds and maximum angles. The numbers at the top indicate the number of trials resulting in an RMSE above 2° .

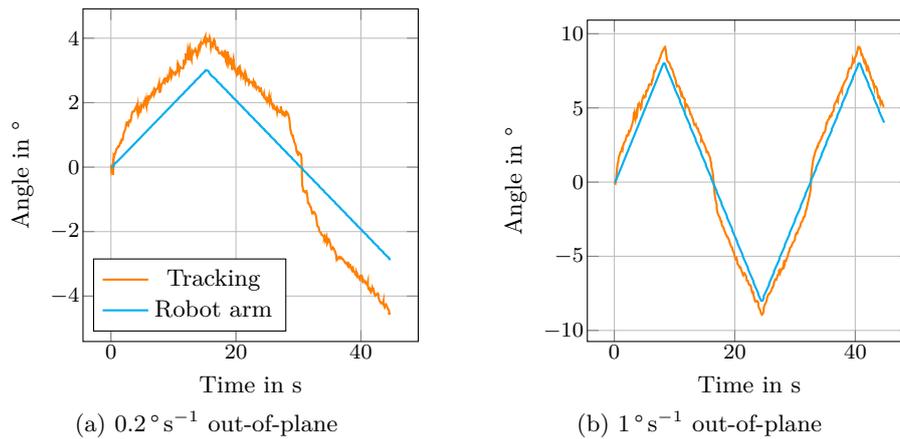


Figure 7.24: Exemplary illustration of the systematic overestimation of the angle occurring for out-of-plane rotations and $\eta = 0.02$ when the maximum angle is 3° (a) or 8° (b).

and 1.5°s^{-1} with RMSEs around 0.6° . As a drawback of the higher number of templates, however, tracking a 8° rotation already failed two out of three times for a speed of 1°s^{-1} . Furthermore, the effective tracking rate is decreased due to the evaluation of more templates per cycle. For slower rotations, the results are comparable to the results for 3° maximum angle.

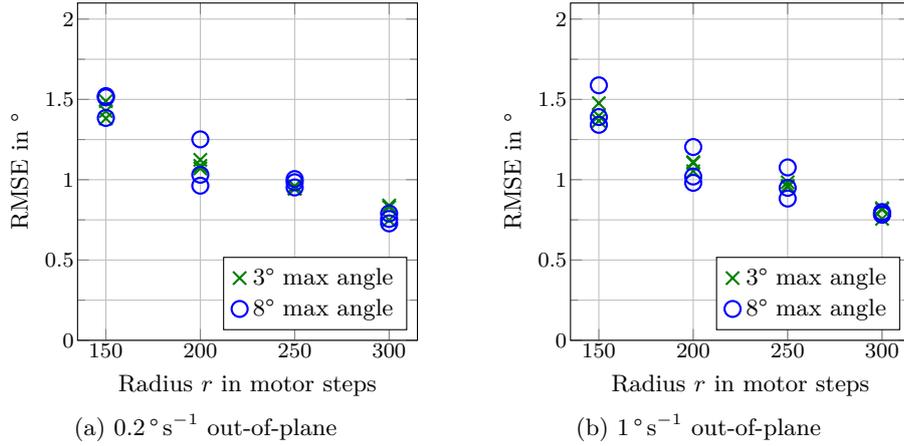


Figure 7.25: Results for out-of-plane rotations when combining $\eta = 0.02$ with different radii r . Note that previous results are based on $r = 200$. A radius of 150, 200, 250, and 300 motor steps corresponds to about 4.1 mm, 5.4 mm, 6.8 mm, and 8.1 mm, respectively.

Table 7.2: RMSEs obtained for out-of-plane rotations when employing $N = 7$ templates, radius $r = 300$, and $\eta = 0.02$.

RMSE	Maximum angle 3°			Maximum angle 8°					
	0.2 ° s ⁻¹		1.5 ° s ⁻¹	0.2 ° s ⁻¹					
	0.62°	0.63°	0.64°	0.60°	0.62°	0.63°	0.55°	0.63°	0.63°

Arbitrary 6D Motion

Finally, we consider combined translational and rotational motion. Both the sequence of the end-effector’s position and its orientation are generated randomly. Figure 7.26 shows exemplary resulting trajectories. The positions are within a cube and the orientations limited to a maximum angle w.r.t. the initial orientation. For a rotational speeds of 1 ° s^{-1} and the baseline settings $N = 3$ and $r = 200$, Figure 7.27 shows that the results are comparable to previous results on purely rotational motion. The lower learning rate shows higher errors for slower motion, but also fails more often during faster translations. For a maximum rotation of 3 ° and simultaneous translation with 8 mm s^{-1} , rotational errors around 0.75 ° are feasible. Furthermore, Figure 7.26b illustrates that there is still an overestimation of the actual angle while the general progress is correctly followed.

7.7 Discussion

We described and evaluated an OCT-based hardware and software setup which allows to track 3D and 6D motion without requiring markers or any knowledge about the target. Our hardware setup consisting of a second pair of galvos and

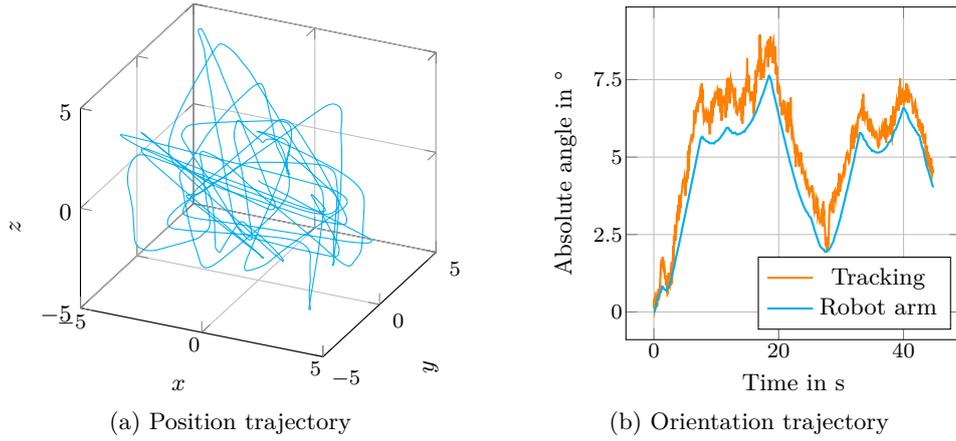


Figure 7.26: Exemplary trajectory of positions (a) and orientation (b) to simulate arbitrary 6D motion. In this case, the translation is random within a cube of size $10 \times 10 \times 10 \text{ mm}^3$ and has a speed of 8 mm s^{-1} . The random rotation has a speed of 1° s^{-1} and the maximum angle w.r.t. the initial orientation is 8° .

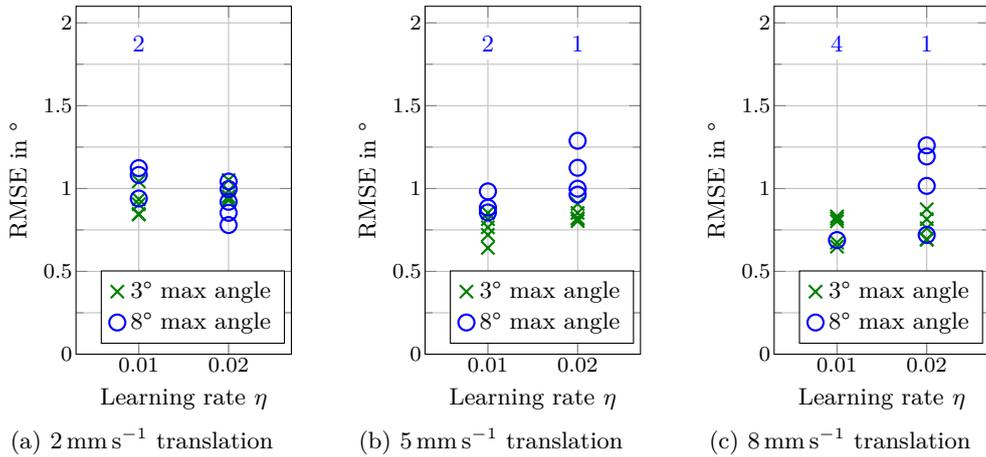


Figure 7.27: Results for arbitrary 6D motion. The speed of rotation is fixed to 1° s^{-1} . We report five repetitions for each combination of parameters and the numbers at the top indicate the number of trials resulting in an RMSE above 2° .

a motorized reference arm allows to shift the FOV in space to follow the motion of a target. The range for lateral shifting by the galvos is more limited than axial shifting due to the diameter of the lens. Nevertheless, we demonstrated 3D tracking of motion over 30 mm, which substantially exceeds the OCT's FOV of $2.5 \times 2.5 \times 3.5 \text{ mm}^3$ in our setup. For this purpose, single-template tracking with the phase correlation algorithm showed promising tracking results, both w.r.t. achievable tracking rates (Table 7.1), which equaled our C-scan acquisition rate of 831 Hz, and the achievable tracking errors (Figure 7.16). Due to our calibration

procedure between the tracking system and the reference robot arm with residual errors below 0.1 mm on a test set (Figure 7.15), we are able to quantify the tracking accuracy. We obtain RMSEs in the order of 0.2 mm for arbitrary 3D motion with speeds up to about 25 mm s^{-1} for both our phantom and chicken breast samples. Tracking of up to 40 mm s^{-1} is feasible with errors around 0.3 mm. For chicken breast samples, however, there are outliers indicating that reliability might be improvable by adding a method to estimate the quality of the template and, if necessary, slightly move the FOV to another location.

Furthermore, our results show that the motorized reference arm is a limiting factor. While we have been able to track lateral motion at 70 mm s^{-1} , the fastest axial motion has been 47.5 mm s^{-1} . However, successful trials with errors below 0.4 mm are even limited to 25 mm s^{-1} . Note that the proportional gains allow for some tuning of the tracking performance for specific scenarios, like an emphasis on speed or on tracking errors. In the future, we should consider a different technical realization for adjusting the reference arm length, if tracking of faster motion is required. Stepper motors are especially limited if the direction of travel changes frequently and if acceleration and deceleration dominate compared to traveling at full speed.

In order to track 6D motion, i.e., not only translations but also rotations, we evaluated a multi-template tracking approach. While we are still not able to rotate the FOV in space, tracking at least three spatially separated templates allows us to derive the full pose of a target. This requires us to apply methods which are robust to appearance changes. Therefore, we proposed to use an adaptive MOSSE filter. Tracking pure translations is still feasible with RMSEs in the order of 0.2 mm and did not show sensitivity to the MOSSE filter's learning rate (Figure 7.19).

When tracking pure rotations, however, we observed multiple important dependencies and properties. For in-plane rotations around the OCT's z -axis, we are able to track rotations by 13° with 4° s^{-1} with low RMSEs around 0.25° (Figure 7.20b). However, if the learning rate is set too high, severely increased errors occur for slow rotations (Figure 7.20a). Because there is no major influence of the learning rate on the results for faster rotations, drifting effects seem to occur in which the template update suppresses the detection of motion resulting in underestimation of the angle (Figure 7.21). Nevertheless, tracking in-plane rotations by a flat-surface phantom, i.e., no superficially visible changes, has also been demonstrated (Figure 7.22) with similar errors when selecting an appropriate learning rate. This emphasizes the potential benefit of exploiting the sub-surface information provided by OCT in contrast to other optical imaging modalities.

While in-plane rotations do not require shifting the FOV by moving the motorized reference arm, this is necessary to follow out-of-plane rotations. Due to the fact that we need to move the FOV to the next template's current position before we can acquire the next C-scan for evaluation, the motion of the motorized reference arm can severely decrease the tracking rate, especially during large rotation angles. Our results for pure out-of-plane rotations show, firstly, RMSEs in the order of 1° which is a factor of four higher than for in-plane rotations (Figure 7.23)

and, secondly, a lower speed which can be tracked. Furthermore, tracking faster rotations requires to increase the learning rate to be successful. The increased errors are due to an overestimation of small rotations which remains as a constant offset for rotations beyond 2° (Figure 7.24). Nevertheless, the progress of the rotation trajectory, especially turning points and zeros crossings, is estimated correctly. Additionally, we showed that the RMSE for out-of-plane rotations can be decreased by tracking more templates arranged on a larger circle (Figure 7.25). Our most extreme settings provided errors around 0.6° (Table 7.2). However, the drawback of such a parameter setting is a reduced trackable range of motion due to the larger area already initially covered by the templates and the decreased tracking rate due to moving to and evaluating more templates.

When considering true 6D motion with simultaneous arbitrary translations and rotations, we obtain errors comparable to the previous experiments (Figure 7.27). However, selection of the learning rate can again be a critical factor. A limitation in this evaluation, however, is the lack of a full hand-eye calibration as described in Section 7.5.4. For this reason, we cannot evaluate the error of the full estimated pose, but only of the orientation as for pure rotations. Therefore, a protocol for an accurate hand-eye calibration should be developed for our proposed tracking system in the future.

In summary, our proposed scheme for 6D tracking showed, on the one hand, promising results but, on the other hand, also some limitations of the prototypical setup. One approach to improve the performance on hardware level could be replacing the second galvo pair and the motorized reference arm with a different setup which allows to move the FOV in 6D space. However, this might not be achievable without mechanically moving parts which show the same limited dynamics as we already observed for the motorized reference arm. A less invasive approach could be employing multiple motorized reference arms to reduce waiting times caused by traveling to the next template's position. Switching sequentially through these reference arms would allow the currently unused arms to already move to the positions of subsequent templates. Thereby, a substantial increase of the tracking rate allowing for tracking of more templates and more robust arrangements could be possible. On software level, improvements could be made by adapting the learning rate during tracking. Recently, it has been shown that adapting the learning rate to the estimated current speed of the target can further improve tracking results [7]. In any case, the learning rate should, if possible, be selected application-specific depending on the expected motion patterns.

To conclude this chapter, we can summarize that tracking of 3D and 6D motion with our proposed markerless tracking system based on OCT is feasible. The range of motion we can follow is currently limited to a few centimeters, but this already substantially exceeds the FOV size of OCT. It is sufficient for applications in which the maximum range of motion needed to be continuously tracked is limited and known beforehand, as for example in radiosurgery where large motion triggers an emergency stop anyhow. The sub-surface information obtained from volumetric OCT imaging provides information which is not accessible for super-

ficially scanning optical systems. While we found efficient algorithms to extract motion without specific knowledge about the target, the prototypical hardware setup, especially the motorized reference arm for axially shifting the FOV, shows some limitations which should be tackled. Nevertheless, the presented results show that OCT is promising for general-purpose motion tracking with high accuracy in scenarios like, for example, head tracking and image guidance during cranial radiosurgery.

8 Conclusion and Outlook

In this thesis, we analyzed different aspects of ultrasound imaging and OCT for markerless and volumetric intra-fractional image guidance during robotic radio-surgery. Our analysis has been guided by the three research questions which we formulated in Section 1.2. The first and the second question are related to trans-abdominal ultrasound guidance with the kinematically redundant robot-arm setup reviewed in Section 3.6.3 and further evaluated in Chapter 4. The third question relates to the feasibility of OCT-based tracking as a novel alternative for cranial image guidance. In this chapter, we summarize the findings of this thesis w.r.t. these three research questions and provide outlooks for aspects requiring further research or providing starting points for further improvement.

In the first research question, we asked how to automatically determine a suitable robot setup for ultrasound guidance in order to minimize the negative impact on treatment plan quality due to beam blocking by the robot arm and the ultrasound transducer. This question has been addressed in Chapter 5, where we proposed automatic strategies and demonstrated their effectiveness. We proposed different optimization schemes for the GUIDE-robot holding the ultrasound transducer. First, we proposed and compared strategies to optimize the position of the GUIDE-robot's base and simultaneously select a suitable LIFT angle. We evaluated the methods based on prostate cases and compared the results to manually determined reference setups. The results showed that the optimization strategies actually allow to determine good setups in most cases. However, the quality depended on the patient and in one case we were not able to outperform one of the reference setups with our approach. Second, we provided methods to select efficient sets of LIFT angles for treatment given a base position for the GUIDE-robot. The main result was that even if we initially consider a large number of configurations for the GUIDE-robot in order to have many beam directions available for planning, we can eventually reduce this number to a reasonably small number while maintaining the same plan quality.

The second research question considered the synchronization problem arising if we consider multiple configurations for the GUIDE-robot. We showed in Chapter 6 how this problem can be modeled and demonstrated that the resulting robot trajectories require only a reasonable amount of temporal overhead. Starting from a model to optimize the LINAC-robot's trajectory for beam delivery, we extended the model to account for the GUIDE-robot changing its configuration. We showed for robot speeds which were in the same order as the actual speed of the LINAC-robot during real CyberKnife treatments that it is feasible to generate efficient trajectories for both robot arms. These trajectories result in only small overhead

compared to the ideal trajectory of the LINAC-robot. However, our model did not provide a direct control of the number of configuration changes occurring in the solution. This might be an interesting future aspect if it is desired to limit that number.

Note that we evaluated our proposed methods addressing the first and second research question in case studies on prostate cancer. However, both the GUIDE-robot setup and our methods are flexible and can be applied to other abdominal targets like the liver as well. By defining a new target pose of the ultrasound transducer, we can immediately evaluate beam blocking for guiding treatment of another site. Our optimization of the robot setup only requires to additionally define the new search space for positioning the robot's base. Our selection of efficient LIFT angles and our synchronization between GUIDE-robot and LINAC-robot do not require any changes. However, the importance of some aspects and the room for improvements might vary with the treatment site due to a different general treatment geometry. Therefore, the methods should be evaluated for other treatment sites in the future to analyze the differences and prove the practical applicability of such generic ultrasound guidance.

In general, the robot setup could be further extended to provide even more flexibility. Instead of a fixed mounting of the robot's base, we could consider moving the base, for example on a rail, during treatment. This could allow for more motion of the links during treatment effectively leading to less blocked beam directions. Some flexibility in mounting the robot base might be required in practical installations anyway in order to allow for a flexible adaption to different patients, couch poses, treatment sites, or potential obstacles in the room. Note that we always considered the robot as being positioned relative to the patient rather than at absolute coordinates in the treatment room.

Regarding the third research question, we provided in Chapter 7 a design and feasibility study for markerless volumetric tracking based on OCT. The motivation for this study was to develop an alternative approach for image guidance during cranial radiosurgery. However, we did not yet consider specific tracking scenarios but demonstrated that OCT can be a valuable basis for general tracking tasks in cases where conventional superficially-scanning systems fail. Nevertheless, future work should consider more specifically the illustrated application and prove that the proposed system design is appropriate for this specific task. Furthermore, our prototypical system can be improved in several aspects. In particular, the axial repositioning of the FOV to follow a moving target currently suffers from the limited dynamics of the motorized reference arm. Further improvements could consider a predictive tracking approach and an automatic local search for optimal template locations.

A practical implementation of OCT guidance in a treatment room would lead to similar issues as our transabdominal ultrasound guidance. The OCT-based tracking system should be mounted to some kind of holder which allows to align it with the patient's head. Even a passive probe holder could, in addition to the tracking system's components, interfere with the delivery of treatment beams.

Therefore, mounting the tracking system to a kinematically redundant robot arm could again be an interesting option. In this case, we could directly apply our methods developed for transabdominal ultrasound to optimize the setup. However, our OCT tracking, in contrast to the ultrasound guidance, does not require direct contact to the patient. The distance to the surface should have a positive effect on the beam blocking issue. Furthermore, we have more freedom in choosing the pose of the tracking system. For ultrasound guidance, we need to have a good view on the rather small target which is limited by bones, air, or other surrounding structures. For head tracking, however, we can consider different regions of the head. Due to the small FOV of the OCT, it might even be possible to adapt an immobilizing face mask with minor effort in order to allow for tracking of certain spots identified as very suitable. Furthermore, we could switch between different spots during treatment to allow for even more beam directions. Having the system mounted to a robot arm, this switching could happen automatically and fast.

To conclude this thesis, we can summarize that we provided and evaluated methods which solve some important issues arising in practical application of ultrasound-based image guidance with a kinematically redundant robot arm in abdominal radiosurgery. Furthermore, we provided a prototypical design and experimental characterization for a novel OCT-based tracking system. This approach can potentially overcome some limitations of conventional tracking systems currently used for guiding cranial radiosurgery.

Bibliography

- [1] M. C. Abramowitz, E. Bossart, R. Flook, X. Wu, R. Brooks, M. Lachaine, F. Lathuiliere, and A. Pollack. “Noninvasive real-time prostate tracking using a transperineal ultrasound approach”. In: *International Journal of Radiation Oncology*Biophysics* 84.3 (2012), S133. DOI: 10.1016/j.ijrobp.2012.07.145.
- [2] J. R. Adler, M. J. Murphy, S. D. Chang, and S. L. Hancock. “Image-guided Robotic Radiosurgery”. In: *Neurosurgery* 44.6 (1999), pp. 1299–1306. DOI: 10.1097/00006123-199906000-00079.
- [3] J. R. Adler Jr, S. D. Chang, M. J. Murphy, J. Doty, P. Geis, and S. L. Hancock. “The Cyberknife: A Frameless Robotic System for Radiosurgery”. In: *Stereotactic and Functional Neurosurgery* 69.1-4 (1997), pp. 124–128. DOI: 10.1159/000099863.
- [4] M. Antico, P. Prinsen, F. Cellini, A. Fracassi, A. A. Isola, D. Cobben, and D. Fontanarosa. “Real-time adaptive planning method for radiotherapy treatment delivery for prostate cancer patients, based on a library of plans accounting for possible anatomy configuration changes”. In: *PLOS ONE* 14.2 (2019), e0213002. DOI: 10.1371/journal.pone.0213002.
- [5] S.-T. Antoni, J. Rinast, X. Ma, S. Schupp, and A. Schlaefer. “Online model checking for monitoring surrogate-based respiratory motion tracking in radiation therapy”. In: *International Journal of Computer Assisted Radiology and Surgery* 11.11 (2016), pp. 2085–2096. DOI: 10.1007/s11548-016-1423-2.
- [6] D. Applegate, W. Cook, and A. Rohe. “Chained Lin-Kernighan for Large Traveling Salesman Problems”. In: *INFORMS Journal on Computing* 15.1 (2003), pp. 82–92. DOI: 10.1287/ijoc.15.1.82.15157.
- [7] C. S. Asha and A. V. Narasimhadhan. “Adaptive Learning Rate for Visual Tracking Using Correlation Filters”. In: *Procedia Computer Science* 89 (2016), pp. 614–622. DOI: 10.1016/j.procs.2016.06.023.
- [8] G. Asmerom, D. Bourne, J. Chappelow, L. M. Goggin, R. Heitz, P. Jordan, W. Kilby, T. Laing, C. R. Maurer Jr, J. M. Noll, S. Sayeh, and A. Weber. “The design and physical characterization of a multileaf collimator for robotic radiosurgery”. In: *Biomedical Physics & Engineering Express* 2.1 (2016), p. 017003. DOI: 10.1088/2057-1976/2/1/017003.
- [9] J. D. Azcona, R. Li, E. Mok, S. Hancock, and L. Xing. “Automatic Prostate Tracking and Motion Assessment in Volumetric Modulated Arc Therapy With an Electronic Portal Imaging Device”. In: *International Journal of Radiation Oncology*Biophysics* 86.4 (2013), pp. 762–768. DOI: 10.1016/j.ijrobp.2013.03.007.
- [10] H. Ballhausen, M. Li, N.-S. Hegemann, U. Ganswindt, and C. Belka. “Intra-fraction motion of the prostate is a random walk”. In: *Physics in Medicine & Biology* 60.2 (2014), pp. 549–563. DOI: 10.1088/0031-9155/60/2/549.

Bibliography

- [11] D. Batchelar, M. Gaztañaga, M. Schmid, C. Araujo, F. Bachand, and J. Crook. “Validation study of ultrasound-based high-dose-rate prostate brachytherapy planning compared with CT-based planning”. In: *Brachytherapy* 13.1 (2014), pp. 75–79. DOI: 10.1016/j.brachy.2013.08.004.
- [12] M. Bazalova-Carter, J. Schlosser, J. Chen, and D. Hristov. “Monte Carlo modeling of ultrasound probes for image guided radiotherapy”. In: *Medical Physics* 42.10 (2015), pp. 5745–5756. DOI: 10.1118/1.4929978.
- [13] J. L. Bedford, H. S. Tsang, S. Nill, and U. Oelfke. “Treatment planning optimization with beam motion modelling for dynamic arc delivery of SBRT using Cyberknife with multileaf collimation”. In: *Medical Physics* 46.12 (2019), pp. 5421–5433. DOI: 10.1002/mp.13848.
- [14] C. Beekman, B. van Triest, S. van Beek, J.-J. Sonke, and P. Remeijer. “Margin and PTV volume reduction using a population based library of plans strategy for rectal cancer radiotherapy”. In: *Medical Physics* 45.10 (2018), pp. 4345–4354. DOI: 10.1002/mp.13137.
- [15] M. A. L. Bell, H. T. Sen, I. I. Iordachita, P. Kazanzides, and J. Wong. “In vivo reproducibility of robotic probe placement for a novel ultrasound-guided radiation therapy system”. In: *Journal of Medical Imaging* 1.2 (2014), p. 025001. DOI: 10.1117/1.JMI.1.2.025001.
- [16] S. H. Benedict, K. M. Yenice, D. Followill, J. M. Galvin, W. Hinson, B. Kavanagh, P. Keall, M. Lovelock, S. Meeks, L. Papiez, T. Purdie, R. Sadagopan, M. C. Schell, B. Salter, D. J. Schlesinger, A. S. Shiu, T. Solberg, D. Y. Song, V. Stieber, R. Timmerman, W. A. Tomé, D. Verellen, L. Wang, and F.-F. Yin. “Stereotactic body radiation therapy: The report of AAPM Task Group 101”. In: *Medical Physics* 37.8 (2010), pp. 4078–4101. DOI: 10.1118/1.3438081.
- [17] M. Bengs, N. Gessert, and A. Schlaefer. “A Deep Learning Approach for Motion Forecasting Using 4D OCT Data”. In: *International Conference on Medical Imaging with Deep Learning*. 2020.
- [18] M. Bengs, N. Gessert, M. Schlüter, and A. Schlaefer. “Spatio-temporal deep learning methods for motion estimation using 4D OCT image data”. In: *International Journal of Computer Assisted Radiology and Surgery* 15.6 (2020), pp. 943–952. DOI: 10.1007/s11548-020-02178-z.
- [19] J. Bertholet, A. Knopf, B. Eiben, J. McClelland, A. Grimwood, E. Harris, M. Menten, P. Poulsen, D. T. Nguyen, P. Keall, and U. Oelfke. “Real-time intrafraction motion monitoring in external beam radiotherapy”. In: *Physics in Medicine & Biology* 64.15 (2019), 15TR01. DOI: 10.1088/1361-6560/ab2ba8.
- [20] W. Birkfellner, J. Hummel, E. Wilson, and K. Cleary. “Tracking Devices”. In: *Image-Guided Interventions. Technology and Applications*. Ed. by T. Peters and K. Cleary. Springer US, 2008, pp. 23–44. DOI: 10.1007/978-0-387-73858-1_2.
- [21] T. Bland, J. Tong, B. Ward, and N. G. Parker. “Geometric distortion of area in medical ultrasound images”. In: *Journal of Physics: Conference Series* 797 (2017), p. 012002. DOI: 10.1088/1742-6596/797/1/012002.
- [22] D. S. Bolme, J. R. Beveridge, B. A. Draper, and Y. M. Lui. “Visual object tracking using adaptive correlation filters”. In: *IEEE Conference on Computer Vision and Pattern Recognition*. 2010, pp. 2544–2550. DOI: 10.1109/cvpr.2010.5539960.

- [23] D. S. Bolme, B. A. Draper, and J. R. Beveridge. “Average of Synthetic Exact Filters”. In: *IEEE Conference on Computer Vision and Pattern Recognition*. 2009, pp. 2105–2112. DOI: 10.1109/cvpr.2009.5206701.
- [24] B. Bontoux, C. Artigues, and D. Feillet. “A Memetic Algorithm with a large neighborhood crossover operator for the Generalized Traveling Salesman Problem”. In: *Computers & Operations Research* 37.11 (2010), pp. 1844–1852. DOI: 10.1016/j.cor.2009.05.004.
- [25] J. T. Booth, V. Caillet, N. Hardcastle, R. O’Brien, K. Szymura, C. Crasta, B. Harris, C. Haddad, T. Eade, and P. J. Keall. “The first patient treatment of electromagnetic-guided real time adaptive radiotherapy using MLC tracking for lung SABR”. In: *Radiotherapy and Oncology* 121.1 (2016), pp. 19–25. DOI: 10.1016/j.radonc.2016.08.025.
- [26] S. A. Boppart, J. Herrmann, C. Pitris, D. L. Stamper, M. E. Brezinski, and J. G. Fujimoto. “High-Resolution Optical Coherence Tomography-Guided Laser Ablation of Surgical Tissue”. In: *Journal of Surgical Research* 82.2 (1999), pp. 275–284. DOI: 10.1006/jsre.1998.5555.
- [27] B. E. Bouma, M. Villiger, K. Otsuka, and W.-Y. Oh. “Intravascular optical coherence tomography [Invited]”. In: *Biomedical Optics Express* 8.5 (2017), p. 2660. DOI: 10.1364/boe.8.002660.
- [28] Brainlab AG. *Potential and Challenges of Surface Guidance in Radiation Therapy. White Paper*. Version Rev1. 2019. URL: https://www.brainlab.com/de/potential-and-challenges-of-sgrt_brainlab/ (visited on 05/26/2020).
- [29] R. W. Brown, Y.-C. N. Cheng, E. M. Haacke, M. R. Thompson, and R. Venkatesan. *Magnetic Resonance Imaging. Physical Principles and Sequence Design*. 2nd ed. John Wiley & Sons, 2014. DOI: 10.1002/9781118633953.
- [30] N. G. Burnet. “Defining the tumour and target volumes for radiotherapy”. In: *Cancer Imaging* 4.2 (2004), pp. 153–161. DOI: 10.1102/1470-7330.2004.0054.
- [31] A. Camino, M. Zhang, S. S. Gao, T. S. Hwang, U. Sharma, D. J. Wilson, D. Huang, and Y. Jia. “Evaluation of artifact reduction in optical coherence tomography angiography with real-time tracking and motion correction technology”. In: *Biomedical Optics Express* 7.10 (2016), p. 3905. DOI: 10.1364/boe.7.003905.
- [32] S. M. Camps, D. Fontanarosa, P. H. N. de With, F. Verhaegen, and B. G. L. Vanneste. “The use of ultrasound imaging in the external beam radiotherapy workflow of prostate cancer patients”. In: *BioMed Research International* 2018 (2018), p. 7569590. DOI: 10.1155/2018/7569590.
- [33] E. de Castro and C. Morandi. “Registration of Translated and Rotated Images Using Finite Fourier Transforms”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* PAMI-9.5 (1987), pp. 700–703. DOI: 10.1109/tpami.1987.4767966.
- [34] L. I. Cerviño, N. Detorie, M. Taylor, J. D. Lawson, T. Harry, K. T. Murphy, A. J. Mundt, S. B. Jiang, and T. A. Pawlicki. “Initial clinical experience with a frameless and maskless stereotactic radiosurgery treatment”. In: *Practical Radiation Oncology* 2.1 (2012), pp. 54–62. DOI: 10.1016/j.prro.2011.04.005.

Bibliography

- [35] W. Chen, J. Unkelbach, A. Trofimov, T. Madden, H. Kooy, T. Bortfeld, and D. Craft. “Including robustness in multi-criteria optimization for intensity-modulated proton therapy”. In: *Physics in Medicine & Biology* 57.3 (2012), pp. 591–608. DOI: 10.1088/0031-9155/57/3/591.
- [36] Z. Chen, Z. Yang, J. Wang, and W. Hu. “Dosimetric impact of different bladder and rectum filling during prostate cancer radiotherapy”. In: *Radiation Oncology* 11.1 (2016), p. 103. DOI: 10.1186/s13014-016-0681-z.
- [37] M. Chu, Y. Zinchenko, S. G. Henderson, and M. B. Sharpe. “Robust optimization for intensity modulated radiation therapy treatment planning under uncertainty”. In: *Physics in Medicine & Biology* 50.23 (2005), pp. 5463–5477. DOI: 10.1088/0031-9155/50/23/003.
- [38] S. D. Collins and M. M. Leech. “A review of plan library approaches in adaptive radiotherapy of bladder cancer”. In: *Acta Oncologica* 57.5 (2018), pp. 566–573. DOI: 10.1080/0284186x.2017.1420908.
- [39] E. L. Covington, J. B. Fiveash, X. Wu, I. Brezovich, C. D. Willey, K. Riley, and R. A. Popple. “Optical surface guidance for submillimeter monitoring of patient position during frameless stereotactic radiotherapy”. In: *Journal of Applied Clinical Medical Physics* 20.6 (2019), pp. 91–98. DOI: 10.1002/acm2.12611.
- [40] F. L. B. Cury, G. Shenouda, L. Souhami, M. Duclos, S. L. Faria, M. David, F. Verhaegen, R. Corns, and T. Falco. “Ultrasound-based image guided radiotherapy for prostate cancer – comparison of cross-modality and intramodality methods for daily localization during external beam radiotherapy”. In: *International Journal of Radiation Oncology*Biophysics* 66.5 (2006), pp. 1562–1567. DOI: 10.1016/j.ijrobp.2006.07.1375.
- [41] W. D. D’Souza, S. A. Naqvi, and C. X. Yu. “Real-time intra-fraction-motion tracking using the treatment couch: a feasibility study”. In: *Physics in Medicine & Biology* 50.17 (2005), pp. 4021–4033. DOI: 10.1088/0031-9155/50/17/007.
- [42] B. Dahroug, B. Tamadazte, and N. Andreff. “PCA-Based Visual Servoing Using Optical Coherence Tomography”. In: *IEEE Robotics and Automation Letters* 5.2 (2020), pp. 3430–3437. DOI: 10.1109/lra.2020.2977259.
- [43] J. G. Delfino, D. M. Krainak, S. A. Flesher, and D. L. Miller. “MRI-related FDA adverse event reports: A 10-yr review”. In: *Medical Physics* 46.12 (2019), pp. 5562–5571. DOI: 10.1002/mp.13768.
- [44] T. Depuydt, K. Poels, D. Verellen, B. Engels, C. Collen, C. Haverbeke, T. Gevaert, N. Buls, G. van Gompel, T. Reynders, M. Duchateau, K. Tournel, M. Boussaer, F. Steenbeke, F. Vandenbroucke, and M. de Ridder. “Initial assessment of tumor tracking with a gimbaled linac system in clinical circumstances: A patient simulation study”. In: *Radiation Oncology* 106.2 (2013), pp. 236–240. DOI: 10.1016/j.radonc.2012.12.015.
- [45] S. Dieterich and I. C. Gibbs. “The CyberKnife in clinical use: current roles, future expectations”. In: *IMRT, IGRT, SBRT – Advances in the Treatment Planning and Delivery of Radiotherapy*. Vol. 43. Karger, 2011, pp. 181–194. DOI: 10.1159/000322423.
- [46] S. Dieterich, O. Green, and J. Booth. “SBRT targets that move with respiration”. In: *Physica Medica* 56 (2018), pp. 19–24. DOI: 10.1016/j.ejmp.2018.10.021.

- [47] E. W. Dijkstra. “A note on two problems in connexion with graphs”. In: *Numerische Mathematik* 1.1 (1959), pp. 269–271. DOI: 10.1007/bf01386390.
- [48] P. Dong, P. Lee, D. Ruan, T. Long, E. Romeijn, Y. Yang, D. Low, P. Kupelian, and K. Sheng. “ 4π Non-Coplanar Liver SBRT: A Novel Delivery Technique”. In: *International Journal of Radiation Oncology*Biophysics* 85.5 (2013), pp. 1360–1366. DOI: 10.1016/j.ijrobp.2012.09.028.
- [49] W. Drexler and J. G. Fujimoto, eds. *Optical Coherence Tomography. Technology and Applications*. Springer Berlin Heidelberg, 2008. DOI: 10.1007/978-3-540-77550-8.
- [50] R. E. Drzymala, R. Mohan, L. Brewster, J. Chu, M. Goitein, W. Harms, and M. Urie. “Dose-volume histograms”. In: *International Journal of Radiation Oncology*Biophysics* 21.1 (1991), pp. 71–78. DOI: 10.1016/0360-3016(91)90168-4.
- [51] A. Dubois, K. Grieve, G. Moneron, R. Lecaque, L. Vabre, and C. Boccara. “Ultra-high-resolution full-field optical coherence tomography”. In: *Applied Optics* 43.14 (2004), p. 2874. DOI: 10.1364/ao.43.002874.
- [52] R. Dürichen, T. Wissel, F. Ernst, A. Schlaefer, and A. Schweikard. “Multivariate respiratory motion prediction”. In: *Physics in Medicine & Biology* 59.20 (2014), pp. 6043–6060. DOI: 10.1088/0031-9155/59/20/6043.
- [53] G. G. Echner, W. Kilby, M. Lee, E. Earnst, S. Sayeh, A. Schlaefer, B. Rhein, J. R. Dooley, C. Lang, O. Blanck, E. Lessard, C. R. Maurer Jr, and W. Schlegel. “The design, physical properties and clinical utility of an iris collimator for robotic radiosurgery”. In: *Physics in Medicine & Biology* 54.18 (2009), p. 5359. DOI: 10.1088/0031-9155/54/18/001.
- [54] F. Ernst, R. Dürichen, A. Schlaefer, and A. Schweikard. “Evaluating and comparing algorithms for respiratory motion prediction”. In: *Physics in Medicine & Biology* 58.11 (2013), pp. 3911–3929. DOI: 10.1088/0031-9155/58/11/3911.
- [55] F. Ernst, R. Bruder, A. Schlaefer, and A. Schweikard. “Correlation between external and internal respiratory motion: a validation study”. In: *International Journal of Computer Assisted Radiology and Surgery* 7.3 (2012), pp. 483–492. DOI: 10.1007/s11548-011-0653-6.
- [56] F. Ernst, L. Richter, L. Matthäus, V. Martens, R. Bruder, A. Schlaefer, and A. Schweikard. “Non-orthogonal tool/flange and robot/world calibration”. In: *The International Journal of Medical Robotics and Computer Assisted Surgery* 8.4 (2012), pp. 407–420. DOI: 10.1002/rcs.1427.
- [57] B. G. Fallone. “The Rotating Biplanar Linac-Magnetic Resonance Imaging System”. In: *Seminars in Radiation Oncology* 24.3 (2014), pp. 200–202. DOI: 10.1016/j.semradonc.2014.02.011.
- [58] M. F. Fast, T. P. O’Shea, S. Nill, U. Oelfke, and E. J. Harris. “First evaluation of the feasibility of MLC tracking using ultrasound motion estimation”. In: *Medical Physics* 43.8Part1 (2016), pp. 4628–4633. DOI: 10.1118/1.4955440.
- [59] L. Feuvret, G. Noël, J.-J. Mazon, and P. Bey. “Conformity index: A review”. In: *International Journal of Radiation Oncology*Biophysics* 64.2 (2006), pp. 333–342. DOI: 10.1016/j.ijrobp.2005.09.028.

Bibliography

- [60] C. Fiandra, A. Guarneri, F. Muñoz, F. Moretto, A. R. Filippi, M. Levis, R. Ragona, and U. Ricardi. “Impact of the observers’ experience on daily prostate localization accuracy in ultrasound-based IGRT with the Clarity platform”. In: *Journal of Applied Clinical Medical Physics* 15.4 (2014), pp. 168–173. DOI: 10.1120/jacmp.v15i4.4795.
- [61] M. Finke, S. Kantelhardt, A. Schlaefel, R. Bruder, E. Lankenau, A. Giese, and A. Schweikard. “Automatic scanning of large tissue areas in neurosurgery using optical coherence tomography”. In: *The International Journal of Medical Robotics and Computer Assisted Surgery* 8.3 (2012), pp. 327–336. DOI: 10.1002/racs.1425.
- [62] B. W. Fischer-Valuck, L. Henke, O. Green, R. Kashani, S. Acharya, J. D. Bradley, C. G. Robinson, M. Thomas, I. Zoberi, W. Thorstad, H. Gay, J. Huang, M. Roach, V. Rodriguez, L. Santanam, H. Li, H. Li, J. Contreras, T. Mazur, D. Hallahan, J. R. Olsen, P. Parikh, S. Mutic, and J. Michalski. “Two-and-a-half-year clinical experience with the world’s first magnetic resonance image guided radiation therapy system”. In: *Advances in Radiation Oncology* 2.3 (2017), pp. 485–493. DOI: 10.1016/j.adro.2017.05.006.
- [63] M. Fischetti, J. J. Salazar González, and P. Toth. “A branch-and-cut algorithm for the symmetric generalized traveling salesman problem”. In: *Operations Research* 45.3 (1997), pp. 378–394. DOI: 10.1287/opre.45.3.378.
- [64] D. Fontanarosa, S. van der Meer, J. Bamber, E. Harris, T. O’Shea, and F. Verhaegen. “Review of ultrasound image guidance in external beam radiotherapy: I. Treatment planning and inter-fraction motion management”. In: *Physics in Medicine & Biology* 60.3 (2015), R77–R114. DOI: 10.1088/0031-9155/60/3/r77.
- [65] D. Fontanarosa, S. van der Meer, E. Harris, and F. Verhaegen. “A CT based correction method for speed of sound aberration for ultrasound based image guided radiotherapy”. In: *Medical Physics* 38.5 (2011), pp. 2665–2673. DOI: 10.1118/1.3583475.
- [66] D. Fontanarosa, S. Pesente, F. Pascoli, D. Ermacora, I. A. Rumeileh, and F. Verhaegen. “A speed of sound aberration correction algorithm for curvilinear ultrasound transducers in ultrasound-based image-guided radiotherapy”. In: *Physics in Medicine & Biology* 58.5 (2013), pp. 1341–1360. DOI: 10.1088/0031-9155/58/5/1341.
- [67] H. Foroosh, J. B. Zerubia, and M. Berthod. “Extension of phase correlation to subpixel registration”. In: *IEEE Transactions on Image Processing* 11.3 (2002), pp. 188–200. DOI: 10.1109/83.988953.
- [68] A. M. Franz, T. Haidegger, W. Birkfellner, K. Cleary, T. M. Peters, and L. Maier-Hein. “Electromagnetic Tracking in Medicine—A Review of Technology, Validation, and Applications”. In: *IEEE Transactions on Medical Imaging* 33.8 (2014), pp. 1702–1725. DOI: 10.1109/tmi.2014.2321777.
- [69] R. Frost, P. Wighton, F. I. Karahanoğlu, R. L. Robertson, P. E. Grant, B. Fischl, M. D. Tisdall, and A. van der Kouwe. “Markerless high-frequency prospective motion correction for neuroanatomical MRI”. In: *Magnetic Resonance in Medicine* 82.1 (2019), pp. 126–144. DOI: 10.1002/mrm.27705.

- [70] D. Fu, R. Kahn, B. Wang, H. Wang, Z. Mu, J. Park, G. Kuduvalli, and C. R. Maurer Jr. “Xsight Lung Tracking System: A Fiducial-Less Method for Respiratory Motion Tracking”. In: *Treating Tumors that Move with Respiration*. Ed. by H. C. Urschel, J. J. Kresl, J. D. Luketich, L. Papiez, and R. D. Timmerman. Springer Berlin Heidelberg, 2007, pp. 265–282. DOI: 10.1007/978-3-540-69886-9_26.
- [71] D. Fu, G. Kuduvalli, V. Mitrovic, W. Main, and L. Thomson. “Automated skull tracking for the CyberKnife image-guided radiosurgery system”. In: *SPIE Medical Imaging*. 2005, pp. 366–377. DOI: 10.1117/12.594984.
- [72] A. Fuchs, S. Pengel, J. Bergmeier, L. A. Kahrs, and T. Ortmaier. “Fast and automatic depth control of iterative bone ablation based on optical coherence tomography data”. In: *SPIE European Conferences on Biomedical Optics*. 2015, 95420P. DOI: 10.1117/12.2183695.
- [73] J. Fujimoto and E. Swanson. “The Development, Commercialization, and Impact of Optical Coherence Tomography”. In: *Investigative Ophthalmology & Visual Science* 57.9 (2016), OCT1. DOI: 10.1167/iovs.16-19963.
- [74] J. Furtak and M. Harat. “Large In-mask Motion during Frameless Radiosurgery of a Brain Metastasis”. In: *Journal of Neurological Surgery Part A: Central European Neurosurgery* 79.04 (2018), pp. 341–343. DOI: 10.1055/s-0037-1604285.
- [75] C. Fürweger, P. Prins, H. Coskan, and B. J. M. Heijmen. “Characteristics and performance of the first commercial multileaf collimator for a robotic radiosurgery system”. In: *Medical Physics* 43.5 (2016), pp. 2063–2071. DOI: 10.1118/1.4944740.
- [76] Y. Gan, W. Yao, K. M. Myers, and C. P. Hendon. “An automated 3D registration method for optical coherence tomography volumes”. In: *Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. 2014, pp. 3873–3876. DOI: 10.1109/embc.2014.6944469.
- [77] S. Gerlach, C. Fürweger, T. Hofmann, and A. Schlaefer. “Feasibility and analysis of CNN based candidate beam generation for robotic radiosurgery”. In: *Medical Physics* 47.9 (2020), pp. 3806–3815. DOI: 10.1002/mp.14331.
- [78] S. Gerlach, I. Kuhlemann, F. Ernst, C. Fürweger, and A. Schlaefer. “Impact of robotic ultrasound image guidance on plan quality in SBRT of the prostate”. In: *The British Journal of Radiology* 90.1078 (2017), p. 20160926. DOI: 10.1259/bjr.20160926.
- [79] S. Gerlach, I. Kuhlemann, P. Jauer, R. Bruder, F. Ernst, C. Fürweger, and A. Schlaefer. “Robotic ultrasound-guided SBRT of the prostate: feasibility with respect to plan quality”. In: *International Journal of Computer Assisted Radiology and Surgery* 12.1 (2017), pp. 149–159. DOI: 10.1007/s11548-016-1455-7.
- [80] S. Gerlach, M. Schlüter, C. Fürweger, and A. Schlaefer. “Machbarkeit CNN-basierter Erzeugung von Kandidatenstrahlen für Radiochirurgie der Prostata”. In: *Jahrestagung der Deutschen Gesellschaft für Computer- und Roboter Assistierte Chirurgie*. 2019, pp. 128–129.
- [81] N. Gessert, M. Gromniak, M. Schlüter, and A. Schlaefer. “Two-path 3D CNNs for calibration of system parameters for OCT-based motion compensation”. In: *SPIE Medical Imaging*. 2019, p. 1095108. DOI: 10.1117/12.2512823.

Bibliography

- [82] N. Gessert, T. Priegnitz, T. Saathoff, S.-T. Antoni, D. Meyer, M. F. Hamann, K.-P. Jünemann, C. Otte, and A. Schlaefer. “Needle Tip Force Estimation Using an OCT Fiber and a Fused convGRU-CNN Architecture”. In: *International Conference on Medical Image Computing and Computer Assisted Intervention*. 2018, pp. 222–229. DOI: 10.1007/978-3-030-00937-3_26.
- [83] N. Gessert, M. Schlüter, and A. Schlaefer. “A deep learning approach for pose estimation from volumetric OCT data”. In: *Medical Image Analysis* 46 (2018), pp. 162–179. DOI: 10.1016/j.media.2018.03.002.
- [84] S. S. Gill, D. G. T. Thomas, A. P. Warrington, and M. Brada. “Relocatable frame for stereotactic external beam radiotherapy”. In: *International Journal of Radiation Oncology*Biophysics* 20.3 (1991), pp. 599–603. DOI: 10.1016/0360-3016(91)90076-g.
- [85] P. Giraud and A. Houle. “Respiratory Gating for Radiotherapy: Main Technical Aspects and Clinical Benefits”. In: *ISRN Pulmonology* 2013 (2013), pp. 1–13. DOI: 10.1155/2013/519602.
- [86] L. Glandorf. “Multi-Volume OCT-Tracking for Robust 6-DOF Motion Estimation”. Master’s thesis. Hamburg University of Technology, Dec. 17, 2019.
- [87] M. Glitzner, P. L. Woodhead, P. T. S. Borman, J. J. W. Lagendijk, and B. W. Raaymakers. “Technical note: MLC-tracking performance on the Elekta Unity MRI-linac”. In: *Physics in Medicine & Biology* 64.15 (2019), 15NT02. DOI: 10.1088/1361-6560/ab2667.
- [88] S. M. Goetz, I. C. Kozyrkov, B. Lubner, S. H. Lisanby, D. L. K. Murphy, W. M. Grill, and A. V. Peterchev. “Accuracy of robotic coil positioning during transcranial magnetic stimulation”. In: *Journal of Neural Engineering* 16.5 (2019), p. 054003. DOI: 10.1088/1741-2552/ab2953.
- [89] K. A. Goodman, E. A. Wiegner, K. E. Maturen, Z. Zhang, Q. Mo, G. Yang, I. C. Gibbs, G. A. Fisher, and A. C. Koong. “Dose-Escalation Study of Single-Fraction Stereotactic Body Radiotherapy for Liver Malignancies”. In: *International Journal of Radiation Oncology*Biophysics* 78.2 (2010), pp. 486–493. DOI: 10.1016/j.ijrobp.2009.08.020.
- [90] O. Gopan and Q. Wu. “Evaluation of the Accuracy of a 3D Surface Imaging System for Patient Setup in Head and Neck Cancer Radiotherapy”. In: *International Journal of Radiation Oncology*Biophysics* 84.2 (2012), pp. 547–552. DOI: 10.1016/j.ijrobp.2011.12.004.
- [91] F. Griese, S. Latus, M. Schlüter, M. Graeser, M. Lutz, A. Schlaefer, and T. Knopp. “In-Vitro MPI-guided IVOCT catheter tracking in real time for motion artifact compensation”. In: *PLOS ONE* 15.3 (2020), e0230821. DOI: 10.1371/journal.pone.0230821.
- [92] A. Grimwood, H. A. McNair, T. P. O’Shea, S. Gilroy, K. Thomas, J. C. Bamber, A. C. Tree, and E. J. Harris. “In Vivo Validation of Elekta’s Clarity Autoscan for Ultrasound-based Intrafraction Motion Estimation of the Prostate During Radiation Therapy”. In: *International Journal of Radiation Oncology*Biophysics* 102.4 (2018), pp. 912–921. DOI: 10.1016/j.ijrobp.2018.04.008.

- [93] A. Grimwood, H. Rivaz, H. Zhou, H. A. McNair, K. Yakubowski, J. C. Bamber, A. C. Tree, and E. J. Harris. “Improving 3D ultrasound prostate localisation in radiotherapy through increased automation of interfraction matching”. In: *Radiotherapy and Oncology* 149 (2020), pp. 134–141. DOI: 10.1016/j.radonc.2020.04.044.
- [94] S. S. Gurbani, P. Wilkening, M. Zhao, B. Gonenc, G. W. Cheon, I. I. Iordachita, W. Chien, R. H. Taylor, J. K. Niparko, and J. U. Kang. “Robot-assisted three-dimensional registration for cochlear implant surgery using a common-path swept-source optical coherence tomography probe”. In: *Journal of Biomedical Optics* 19.5 (2014), p. 057004. DOI: 10.1117/1.jbo.19.5.057004.
- [95] G. Gutin and D. Karapetyan. “A memetic algorithm for the generalized traveling salesman problem”. In: *Natural Computing* 9.1 (2010), pp. 47–60. DOI: 10.1007/s11047-009-9111-6.
- [96] G. Gutin and A. P. Punnen, eds. *The Traveling Salesman Problem and Its Variations*. Springer US, 2007. DOI: 10.1007/b101971.
- [97] K. Helsgaun. “An effective implementation of the Lin–Kernighan traveling salesman heuristic”. In: *European Journal of Operational Research* 126.1 (2000), pp. 106–130. DOI: 10.1016/s0377-2217(99)00284-2.
- [98] K. Helsgaun. “General k-opt submoves for the Lin–Kernighan TSP heuristic”. In: *Mathematical Programming Computation* 1.2 (2009), pp. 119–163. DOI: 10.1007/s12532-009-0004-6.
- [99] K. Helsgaun. “Solving the equality generalized traveling salesman problem using the Lin–Kernighan–Helsgaun Algorithm”. In: *Mathematical Programming Computation* 7.3 (2015), pp. 269–287. DOI: 10.1007/s12532-015-0080-8.
- [100] A. K. Ho, D. Fu, C. Cotrutz, S. L. Hancock, S. D. Chang, I. C. Gibbs, C. R. Maurer Jr, and J. R. Adler. “A Study of the Accuracy of CyberKnife Spinal Radiosurgery Using Skeletal Structure Tracking”. In: *Operative Neurosurgery* 60.suppl_2 (2007), ONS-147–ONS-156. DOI: 10.1227/01.neu.0000249248.55923.ec.
- [101] J. D. P. Hoisak and T. Pawlicki. “The Role of Optical Surface Imaging Systems in Radiation Therapy”. In: *Seminars in Radiation Oncology* 28.3 (2018), pp. 185–193. DOI: 10.1016/j.semradonc.2018.02.003.
- [102] O. E. Holmes, J. Gratton, J. Szanto, E. Vandervoort, J. Doody, E. Henderson, S. C. Morgan, J. O’Sullivan, and S. Malone. “Reducing errors in prostate tracking with an improved fiducial implantation protocol for CyberKnife based stereotactic body radiotherapy (SBRT)”. In: *Journal of Radiosurgery and SBRT* 5.3 (2018), pp. 217–227.
- [103] S. Hossain, P. Xia, K. Huang, M. Descovich, C. Chuang, A. R. Gottschalk, M. Roach, and L. Ma. “Dose Gradient Near Target-Normal Structure Interface for Nonisocentric CyberKnife and Isocentric Intensity-Modulated Body Radiotherapy for Prostate Cancer”. In: *International Journal of Radiation Oncology*Biophysics*Physics* 78.1 (2010), pp. 58–63. DOI: 10.1016/j.ijrobp.2009.07.1752.
- [104] D. Huang, E. A. Swanson, C. P. Lin, J. S. Schuman, W. G. Stinson, W. Chang, M. R. Hee, T. Flotte, K. Gregory, C. A. Puliafito, and J. G. Fujimoto. “Optical coherence tomography”. In: *Science* 254.5035 (1991), pp. 1178–1181. DOI: 10.1126/science.1957169.

Bibliography

- [105] R. Huber, M. Wojtkowski, and J. G. Fujimoto. “Fourier Domain Mode Locking (FDML): A new laser operating regime and applications for optical coherence tomography”. In: *Optics Express* 14.8 (2006), p. 3225. DOI: 10.1364/oe.14.003225.
- [106] A. Hunt, V. N. Hansen, U. Oelfke, S. Nill, and S. Hafeez. “Adaptive Radiotherapy Enabled by MRI Guidance”. In: *Clinical Oncology* 30.11 (2018), pp. 711–719. DOI: 10.1016/j.clon.2018.08.001.
- [107] International Agency for Research on Cancer. *Cancer Fact Sheets. All cancers 2018*. 2019. URL: <https://gco.iarc.fr/today/fact-sheets-cancers>.
- [108] S. Ipsen, R. Bruder, V. García-Vázquez, A. Schweikard, and F. Ernst. “Assessment of 4D Ultrasound Systems for Image-guided Radiation Therapy – Image Quality, Framerates and CT Artifacts”. In: *Current Directions in Biomedical Engineering* 5.1 (2019), pp. 245–248. DOI: 10.1515/cdbme-2019-0062.
- [109] S. Ipsen, R. Bruder, R. O’Brien, P. J. Keall, A. Schweikard, and P. R. Poulsen. “Online 4D ultrasound guidance for real-time motion compensation by MLC tracking”. In: *Medical Physics* 43.10 (2016), pp. 5695–5704. DOI: 10.1118/1.4962932.
- [110] M. Jaccard, A. Champion, A. Dubouloz, C. Picardi, J. Plojoux, P. Soccal, R. Miralbell, G. Dipasquale, and F. Caparrotti. “Clinical experience with lung-specific electromagnetic transponders for real-time tumor tracking in lung stereotactic body radiotherapy”. In: *Physics and Imaging in Radiation Oncology* 12 (2019), pp. 30–37. DOI: 10.1016/j.phro.2019.11.002.
- [111] R. N. Jazar. *Theory of Applied Robotics*. 2nd ed. Springer US, 2010. DOI: 10.1007/978-1-4419-1750-8.
- [112] U. Jeleń and M. Alber. “A finite size pencil beam algorithm for IMRT dose optimization: density corrections.” In: *Physics in Medicine & Biology* 52.3 (2007), pp. 617–633. DOI: 10.1088/0031-9155/52/3/006.
- [113] U. Jeleń, M. Söhn, and M. Alber. “A finite size pencil beam for IMRT dose optimization.” In: *Physics in Medicine & Biology* 50.8 (2005), pp. 1747–1766. DOI: 10.1088/0031-9155/50/8/009.
- [114] I. Kably, M. Bordegaray, K. Shah, J. Salsamendi, and G. Narayanan. “Single-Center Experience in Prostate Fiducial Marker Placement: Technique and Midterm Follow-up”. In: *Journal of Vascular and Interventional Radiology* 25.7 (2014), 1125–1132.e1. DOI: 10.1016/j.jvir.2014.03.017.
- [115] W. Kabsch. “A solution for the best rotation to relate two sets of vectors”. In: *Acta Crystallographica Section A* 32.5 (1976), pp. 922–923. DOI: 10.1107/s0567739476001873.
- [116] D. Karapetyan and G. Gutin. “Lin-Kernighan heuristic adaptations for the generalized traveling salesman problem”. In: *European Journal of Operational Research* 208.3 (2011), pp. 221–232. DOI: 10.1016/j.ejor.2010.08.011.
- [117] K. Karki, S. Saraiya, G. D. Hugo, N. Mukhopadhyay, N. Jan, J. Schuster, M. Schutzer, L. Fahrner, R. Groves, K. M. Olsen, J. C. Ford, and E. Weiss. “Variabilities of Magnetic Resonance Imaging-, Computed Tomography-, and Positron Emission Tomography-Computed Tomography-Based Tumor and Lymph Node Delineations for Lung Cancer Radiation Therapy Planning”. In: *International Journal of Radiation Oncology*Biophysics*Physics* 99.1 (2017), pp. 80–89. DOI: 10.1016/j.ijrobp.2017.05.002.

- [118] T. Kataria, K. Narang, D. Gupta, S. S. Bisht, A. Abhishek, S. Goyal, T. Basu, and K. P. Karrthick. “Analysis of intrafraction motion in CyberKnife-based stereotaxy using mask based immobilization and 6D-skull tracking”. In: *Journal of Radio-surgery and SBRT* 4.3 (2016), pp. 203–212.
- [119] T. Kataria, K. Sharma, V. Subramani, K. P. Karrthick, and S. S. Bisht. “Homogeneity Index: An objective tool for assessment of conformal radiation treatments”. In: *Journal of Medical Physics* 37.4 (2012), p. 207. DOI: 10.4103/0971-6203.103606.
- [120] N. Katta, A. B. McElroy, A. D. Estrada, and T. E. Milner. “Optical coherence tomography image-guided smart laser knife for surgery”. In: *Lasers in Surgery and Medicine* 50.3 (2017), pp. 202–212. DOI: 10.1002/lsm.22705.
- [121] P. J. Keall, M. Barton, and S. Crozier. “The Australian Magnetic Resonance Imaging-Linac Program”. In: *Seminars in Radiation Oncology* 24.3 (2014), pp. 203–206. DOI: 10.1016/j.semradonc.2014.02.015.
- [122] P. J. Keall, E. Colvill, R. O’Brien, V. Caillet, T. Eade, A. Kneebone, G. Hruba, P. R. Poulsen, B. Zwan, P. B. Greer, and J. Booth. “Electromagnetic-Guided MLC Tracking Radiation Therapy for Prostate Cancer Patients: Prospective Clinical Trial Results”. In: *International Journal of Radiation Oncology*Biophysics*Physics* 101.2 (2018), pp. 387–395. DOI: 10.1016/j.ijrobp.2018.01.098.
- [123] V. Kearney, J. P. Cheung, C. McGuinness, and T. D. Solberg. “CyberArc: a non-coplanar-arc optimization algorithm for CyberKnife”. In: *Physics in Medicine & Biology* 62.14 (2017), p. 5777. DOI: 10.1088/1361-6560/aa6f92.
- [124] V. Kearney, M. Descovich, A. Sudhyadhom, J. P. Cheung, C. McGuinness, and T. D. Solberg. “A continuous arc delivery optimization algorithm for CyberKnife M6”. In: *Medical Physics* 45.8 (2018), pp. 3861–3870. DOI: 10.1002/mp.13022.
- [125] W. Kilby, J. R. Dooley, G. Kuduvalli, S. Sayeh, and C. R. Maurer Jr. “The CyberKnife® Robotic Radiosurgery System in 2010”. In: *Technology in Cancer Research & Treatment* 9.5 (2010), pp. 433–452. DOI: 10.1177/153303461000900502.
- [126] W. Kilby, M. Naylor, J. R. Dooley, C. R. Maurer Jr, and S. Sayeh. “A Technical Overview of the CyberKnife System”. In: *Handbook of Robotic and Image-Guided Surgery*. Ed. by M. H. Abedin-Nasab. Elsevier, 2020, pp. 15–38. DOI: 10.1016/b978-0-12-814245-5.00002-5.
- [127] J. P. Kirkpatrick, S. G. Soltys, S. S. Lo, K. Beal, D. C. Shrieve, and P. D. Brown. “The radiosurgery fractionation quandary: single fraction or hypofractionation?” In: *Neuro-Oncology* 19 (2017), pp. ii38–ii49. DOI: 10.1093/neuonc/now301.
- [128] J. P. Kolb, W. Draxinger, J. Klee, T. Pfeiffer, M. Eibl, T. Klein, W. Wieser, and R. Huber. “Live video rate volumetric OCT imaging of the retina with multi-MHz A-scan rates”. In: *PLOS ONE* 14.3 (2019), e0213144. DOI: 10.1371/journal.pone.0213144.
- [129] C. Kontaxis, G. H. Bol, J. J. W. Lagendijk, and B. W. Raaymakers. “A new methodology for inter- and intrafraction plan adaptation for the MR-linac”. In: *Physics in Medicine & Biology* 60.19 (2015), pp. 7485–7497. DOI: 10.1088/0031-9155/60/19/7485.

Bibliography

- [130] C. Kontaxis, G. H. Bol, B. Stemkens, M. Glitzner, F. M. Prins, L. G. W. Kerkmeijer, J. J. W. Lagendijk, and B. W. Raaymakers. “Towards fast online intrafraction replanning for free-breathing stereotactic body radiation therapy with the MR-linac”. In: *Physics in Medicine & Biology* 62.18 (2017), pp. 7233–7248. DOI: 10.1088/1361-6560/aa82ae.
- [131] A. Krauss, M. F. Fast, S. Nill, and U. Oelfke. “Multileaf collimator tracking integrated with a novel X-ray imaging system and external surrogate monitoring”. In: *Physics in Medicine & Biology* 57.8 (2012), p. 2425. DOI: 10.1088/0031-9155/57/8/2425.
- [132] C. D. Kuglin and D. C. Hines. “The phase correlation image alignment method”. In: *IEEE Conference on Cybernetics and Society*. 1975, pp. 163–165.
- [133] I. Kuhlemann, R. Bruder, F. Ernst, and A. Schweikard. “WE-G-BRF-09: Force- and Image-Adaptive Strategies for Robotised Placement of 4D Ultrasound Probes”. In: *Medical Physics* 41.6Part30 (2014), p. 523. DOI: 10.1118/1.4889502.
- [134] I. Kuhlemann, P. Jauer, A. Schweikard, and F. Ernst. “SU-G-JeP3-08: Robotic System for Ultrasound Tracking in Radiation Therapy”. In: *Medical Physics* 43.6 (2016), p. 3672. DOI: 10.1118/1.4957073.
- [135] I. Kuhlemann. “Methods for Quasi-Static Tasks with Redundant Manipulators. Advances in Kinematics, Dexterity and Sensitivity”. PhD thesis. University of Lübeck, 2019.
- [136] I. Kuhlemann, P. Jauer, F. Ernst, and A. Schweikard. “Robots with seven degrees of freedom: Is the additional DoF worth it?” In: *IEEE International Conference on Control, Automation and Robotics*. 2016, pp. 80–84. DOI: 10.1109/iccar.2016.7486703.
- [137] I. Kuhlemann, A. Schweikard, P. Jauer, and F. Ernst. “Robust inverse kinematics by configuration control for redundant manipulators with seven DoF”. In: *IEEE International Conference on Control, Automation and Robotics*. 2016, pp. 49–55. DOI: 10.1109/iccar.2016.7486697.
- [138] B. V. K. V. Kumar. “Minimum-variance synthetic discriminant functions”. In: *Journal of the Optical Society of America A* 3.10 (1986), pp. 1579–1584. DOI: 10.1364/josaa.3.001579.
- [139] P. Kupelian, T. Willoughby, A. Mahadevan, T. Djemil, G. Weinstein, S. Jani, C. Enke, T. Solberg, N. Flores, D. Liu, D. Beyer, and L. Levine. “Multi-institutional clinical experience with the Calypso System in localization and continuous, real-time monitoring of the prostate gland during external radiotherapy”. In: *International Journal of Radiation Oncology*Biophysics*Physics* 67.4 (2007), pp. 1088–1098. DOI: 10.1016/j.ijrobp.2006.10.026.
- [140] A. Z. Kyme, S. Se, S. R. Meikle, and R. R. Fulton. “Markerless motion estimation for motion-compensated clinical brain imaging”. In: *Physics in Medicine & Biology* 63.10 (2018), p. 105018. DOI: 10.1088/1361-6560/aabd48.
- [141] M. Lachaine and T. Falco. “Intrafractional prostate motion management with the Clarity Autoscan system”. In: *Medical Physics International Journal* 1.1 (2013), pp. 72–80.

- [142] J. Lamb, M. Cao, A. Kishan, N. Agazaryan, D. H. Thomas, N. Shaverdian, Y. Yang, S. Ray, D. A. Low, A. Raldow, M. L. Steinberg, and P. Lee. “Online Adaptive Radiation Therapy: Implementation of a New Process of Care”. In: *Cureus* 9.8 (2017), e1618. DOI: 10.7759/cureus.1618.
- [143] S. Lang, J. Zeimetz, G. Ochsner, M. S. Daners, O. Riesterer, and S. Klöck. “Development and evaluation of a prototype tracking system using the treatment couch”. In: *Medical Physics* 41.2 (2014), p. 021720. DOI: 10.1118/1.4862077.
- [144] K. M. Langen, J. Pouliot, C. Anezinos, M. Aubin, A. R. Gottschalk, I.-C. Hsu, D. Lowther, Y.-M. Liu, K. Shinohara, L. J. Verhey, V. Weinberg, and M. Roach. “Evaluation of ultrasound-based prostate localization for image-guided radiotherapy”. In: *International Journal of Radiation Oncology*Biography*Physics* 57.3 (2003), pp. 635–644. DOI: 10.1016/s0360-3016(03)00633-3.
- [145] J. F. Langenhuijsen, E. N. J. T. van Lin, L. A. Kiemeney, L. P. van der Vight, G. M. McColl, A. G. Visser, and J. A. Witjes. “Ultrasound-Guided Transrectal Implantation of Gold Markers for Prostate Localization During External Beam Radiotherapy: Complication Rate and Risk Factors”. In: *International Journal of Radiation Oncology*Biography*Physics* 69.3 (2007), pp. 671–676. DOI: 10.1016/j.ijrobp.2007.04.009.
- [146] G. Laporte and F. Semet. “Computational Evaluation Of A Transformation Procedure For The Symmetric Generalized Traveling Salesman Problem”. In: *INFOR: Information Systems and Operational Research* 37.2 (1999), pp. 114–120. DOI: 10.1080/03155986.1999.11732374.
- [147] K. V. Larin and D. D. Sampson. “Optical coherence elastography – OCT at work in tissue biomechanics [Invited]”. In: *Biomedical Optics Express* 8.2 (2017), p. 1172. DOI: 10.1364/boe.8.001172.
- [148] J. Lattanzi, S. McNeeley, W. Pinover, E. Horwitz, I. Das, T. E. Schultheiss, and G. E. Hanks. “A comparison of daily CT localization to a daily ultrasound-based system in prostate cancer”. In: *International Journal of Radiation Oncology*Biography*Physics* 43.4 (1999), pp. 719–725. DOI: 10.1016/s0360-3016(98)00496-9.
- [149] S. Latus, F. Griese, M. Schlüter, C. Otte, M. Möddel, M. Graeser, T. Saathoff, T. Knopp, and A. Schlaefer. “Bimodal intravascular volumetric imaging combining OCT and MPI”. In: *Medical Physics* 46.3 (2019), pp. 1371–1383. DOI: 10.1002/mp.13388.
- [150] S. Latus, C. Otte, M. Schlüter, J. Rehra, K. Bizon, H. Schulz-Hildebrandt, T. Saathoff, G. Hüttmann, and A. Schlaefer. “An Approach for Needle Based Optical Coherence Elastography Measurements”. In: *International Conference on Medical Image Computing and Computer Assisted Intervention*. 2017, pp. 655–663. DOI: 10.1007/978-3-319-66185-8_74.
- [151] J. L. Lauermann, M. Treder, P. Heiduschka, C. R. Clemens, N. Eter, and F. Alten. “Impact of eye-tracking technology on OCT-angiography imaging quality in age-related macular degeneration”. In: *Graefe’s Archive for Clinical and Experimental Ophthalmology* 255.8 (2017), pp. 1535–1542. DOI: 10.1007/s00417-017-3684-z.
- [152] M.-H. Laves, S. Ihler, L. A. Kahrs, and T. Ortmaier. “Deep-learning-based 2.5D flow field estimation for maximum intensity projections of 4D optical coherence tomography”. In: *SPIE Medical Imaging*. 2019, 109510R. DOI: 10.1117/12.2512952.

Bibliography

- [153] M.-H. Laves, L. A. Kahrs, and T. Ortmaier. “Volumetric 3D stitching of optical coherence tomography volumes”. In: *Current Directions in Biomedical Engineering* 4.1 (2018), pp. 327–330. DOI: 10.1515/cdbme-2018-0079.
- [154] M.-H. Laves, A. Schoob, L. A. Kahrs, T. Pfeiffer, R. Huber, and T. Ortmaier. “Feature tracking for automated volume of interest stabilization on 4D-OCT images”. In: *SPIE Medical Imaging*. 2017, 101350W. DOI: 10.1117/12.2255090.
- [155] C. S. G. Lee and M. Ziegler. “Geometric Approach in Solving Inverse Kinematics of PUMA Robots”. In: *IEEE Transactions on Aerospace and Electronic Systems* AES-20.6 (1984), pp. 695–706. DOI: 10.1109/TAES.1984.310452.
- [156] R. A. Leitgeb. “En face optical coherence tomography: a technology review [Invited]”. In: *Biomedical Optics Express* 10.5 (2019), p. 2177. DOI: 10.1364/boe.10.002177.
- [157] L. Leksell. “Cerebral Radiosurgery. I. Gammathalamotomy in Two Cases of Intractable Pain.” In: *Acta Chirurgica Scandinavica* 134 (1968), pp. 585–589.
- [158] L. Leksell. “Stereotactic radiosurgery.” In: *Journal of Neurology, Neurosurgery & Psychiatry* 46.9 (1983), pp. 797–803. DOI: 10.1136/jnnp.46.9.797.
- [159] E. Lessard, W. Kilby, J. Dooley, C. Sims, A. Schlaefel, O. Blanck, and C. R. Maurer Jr. “Sequential Optimization Scripts to Facilitate Treatment Planning for Robotic Radiosurgery Clinical Studies for Prostate and Lung Cancers”. In: *IFMBE Proceedings*. 2009, pp. 1031–1034. DOI: 10.1007/978-3-642-03474-9_290.
- [160] G. Li, Å. Ballangrud, L. C. Kuo, H. Kang, A. Kirov, M. Lovelock, Y. Yamada, J. Mechalakos, and H. Amols. “Motion monitoring for cranial frameless stereotactic radiosurgery using video-based three-dimensional optical surface imaging”. In: *Medical Physics* 38.7 (2011), pp. 3981–3994. DOI: 10.1118/1.3596526.
- [161] G. Li, D. M. Lovelock, J. Mechalakos, S. Rao, C. Della-Biancia, H. Amols, and N. Lee. “Migration from full-head mask to "open-face" mask for immobilization of patients with head and neck cancer”. In: *Journal of Applied Clinical Medical Physics* 14.5 (2013), pp. 243–254. DOI: 10.1120/jacmp.v14i5.4400.
- [162] M. Li, H. Ballhausen, N.-S. Hegemann, U. Ganswindt, F. Manapov, S. Tritschler, A. Roosen, C. Gratzke, M. Reiner, and C. Belka. “A comparative assessment of prostate positioning guided by three-dimensional ultrasound and cone beam CT”. In: *Radiation Oncology* 10.1 (2015), p. 82. DOI: 10.1186/s13014-015-0380-1.
- [163] S. Lin. “Computer Solutions of the Traveling Salesman Problem”. In: *Bell System Technical Journal* 44.10 (1965), pp. 2245–2269. DOI: 10.1002/j.1538-7305.1965.tb04146.x.
- [164] S. Lin and B. W. Kernighan. “An effective heuristic algorithm for the traveling-salesman problem”. In: *Operations Research* 21.2 (1973), pp. 498–516. DOI: 10.1287/opre.21.2.498.
- [165] D. G. Lowe. “Distinctive Image Features from Scale-Invariant Keypoints”. In: *International Journal of Computer Vision* 60.2 (2004), pp. 91–110. DOI: 10.1023/b:visi.0000029664.99615.94.
- [166] K. L. Lurie and A. K. Ellerbee. “Volumetric mosaicing for optical coherence tomography for large area bladder wall visualization”. In: *SPIE BiOS*. 2014, 89261P. DOI: 10.1117/12.2036325.

- [167] W. Lutz, K. R. Winston, and N. Maleki. “A system for stereotactic radiosurgery with a linear accelerator”. In: *International Journal of Radiation Oncology*Biophysics* 14.2 (1988), pp. 373–381. DOI: 10.1016/0360-3016(88)90446-4.
- [168] M. Machoy, J. Seeliger, L. Szyszka-Sommerfeld, R. Koprowski, T. Gedrange, and K. Woźniak. “The Use of Optical Coherence Tomography in Dental Diagnostics: A State-of-the-Art Review”. In: *Journal of Healthcare Engineering* 2017 (2017), p. 7560645. DOI: 10.1155/2017/7560645.
- [169] R. J. Maciunas, R. L. Galloway Jr, J. Latimer, C. Cobb, E. Zaccharias, A. Moore, and V. R. Mandava. “An Independent Application Accuracy Evaluation of Stereotactic Frame Systems”. In: *Stereotactic and Functional Neurosurgery* 58.1-4 (1992), pp. 103–107. DOI: 10.1159/000098981.
- [170] T. R. Mackie, J. Balog, K. Ruchala, D. Shepard, S. Aldridge, E. Fitchard, P. Reckwerdt, G. Olivera, T. McNutt, and M. Mehta. “Tomotherapy”. In: *Seminars in Radiation Oncology* 9.1 (1999), pp. 108–117. DOI: 10.1016/s1053-4296(99)80058-7.
- [171] A. Mahalanobis, B. V. K. V. Kumar, and D. Casasent. “Minimum average correlation energy filters”. In: *Applied Optics* 26.17 (1987), pp. 3633–3640. DOI: 10.1364/ao.26.003633.
- [172] A. Mahalanobis, B. V. K. V. Kumar, S. Song, S. R. F. Sims, and J. F. Epperson. “Unconstrained correlation filters”. In: *Applied Optics* 33.17 (1994), pp. 3751–3759. DOI: 10.1364/ao.33.003751.
- [173] L. Masi, F. Casamassima, C. Polli, C. Menichelli, I. Bonucci, and C. Cavedon. “Cone Beam CT Image Guidance for Intracranial Stereotactic Treatments: Comparison With a Frame Guided Set-Up”. In: *International Journal of Radiation Oncology*Biophysics* 71.3 (2008), pp. 926–933. DOI: 10.1016/j.ijrobp.2008.03.006.
- [174] L. Matthews, T. Ishikawa, and S. Baker. “The template update problem”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 26.6 (2004), pp. 810–815. DOI: 10.1109/tpami.2004.16.
- [175] C. M. McGuinness, A. R. Gottschalk, E. Lessard, J. L. Nakamura, D. Pinnaduwege, J. Pouliot, C. Sims, and M. Descovich. “Investigating the clinical advantages of a robotic linac equipped with a multileaf collimator in the treatment of brain and prostate cancer patients”. In: *Journal of Applied Clinical Medical Physics* 16.5 (2015), pp. 284–295. DOI: 10.1120/jacmp.v16i5.5502.
- [176] M. J. Menten, M. F. Fast, S. Nill, C. P. Kamerling, F. McDonald, and U. Oelfke. “Lung stereotactic body radiotherapy with an MR-linac – Quantifying the impact of the magnetic field and real-time tumor tracking”. In: *Radiotherapy and Oncology* 119.3 (2016), pp. 461–466. DOI: 10.1016/j.radonc.2016.04.019.
- [177] M. Miften, A. Olch, D. Mihailidis, J. Moran, T. Pawlicki, A. Molineu, H. Li, K. Wijesooriya, J. Shi, P. Xia, N. Papanikolaou, and D. A. Low. “Tolerance limits and methodologies for IMRT measurement-based verification QA: Recommendations of AAPM Task Group No. 218”. In: *Medical Physics* 45.4 (2018), e53–e83. DOI: 10.1002/mp.12810.

Bibliography

- [178] B. J. Mijnheer, P. González, I. Olaciregui-Ruiz, R. A. Rozendaal, M. van Herk, and A. Mans. “Overview of 3-year experience with large-scale electronic portal imaging device-based 3-dimensional transit dosimetry”. In: *Practical Radiation Oncology* 5.6 (2015), e679–e687. DOI: 10.1016/j.prro.2015.07.001.
- [179] K. D. Miller, L. Nogueira, A. B. Mariotto, J. H. Rowland, K. R. Yabroff, C. M. Alfano, A. Jemal, J. L. Kramer, and R. L. Siegel. “Cancer treatment and survivorship statistics, 2019”. In: *CA: A Cancer Journal for Clinicians* 69.5 (2019), pp. 363–385. DOI: 10.3322/caac.21565.
- [180] J. A. Molloy, G. Chan, A. Markovic, S. McNeeley, D. Pfeiffer, B. Salter, and W. A. Tome. “Quality assurance of U.S.-guided external beam radiotherapy for prostate cancer: Report of AAPM Task Group 154”. In: *Medical Physics* 38.2 (2011), pp. 857–871. DOI: 10.1118/1.3531674.
- [181] O. Morin, A. Gillis, J. Chen, M. Aubin, M. K. Bucci, M. Roach, and J. Pouliot. “Megavoltage cone-beam CT: System description and clinical applications”. In: *Medical Dosimetry* 31.1 (2006), pp. 51–61. DOI: 10.1016/j.meddos.2005.12.009.
- [182] A. Muacevic, M. Staehler, C. Drexler, B. Wowra, M. Reiser, and J.-C. Tonn. “Technical description, phantom accuracy, and clinical feasibility for fiducial-free frameless real-time image-guided spinal radiosurgery”. In: *Journal of Neurosurgery: Spine* 5.4 (2006), pp. 303–312. DOI: 10.3171/spi.2006.5.4.303.
- [183] M. J. Murphy, J. Balter, S. Balter, J. A. BenComo Jr, I. J. Das, S. B. Jiang, C.-M. Ma, G. H. Olivera, R. F. Rodebaugh, K. J. Ruchala, H. Shirato, and F.-F. Yin. “The management of imaging dose during image-guided radiotherapy: Report of the AAPM Task Group 75”. In: *Medical Physics* 34.10 (2007), pp. 4041–4063. DOI: 10.1118/1.2775667.
- [184] T. Nakamori, H. Okabayashi, and A. Kawanaka. “3-D Object matching using phase correlation method”. In: *IEEE International Conference on Signal Processing*. 2010, pp. 1275–1278. DOI: 10.1109/icosp.2010.5656976.
- [185] S. K. Nath, J. D. Lawson, J.-Z. Wang, D. R. Simpson, C. B. Newman, J. F. Alksne, A. J. Mundt, and K. T. Murphy. “Optically-guided frameless linac-based radiosurgery for brain metastases: clinical experience”. In: *Journal of Neuro-Oncology* 97.1 (2009), pp. 67–72. DOI: 10.1007/s11060-009-9989-y.
- [186] D. T. Nguyen, R. O’Brien, J.-H. Kim, C.-Y. Huang, L. Wilton, P. Greer, K. Legge, J. T. Booth, P. R. Poulsen, J. Martin, and P. J. Keall. “The first clinical implementation of a real-time six degree of freedom target tracking system during radiation therapy based on Kilovoltage Intrafraction Monitoring (KIM)”. In: *Radiotherapy and Oncology* 123.1 (2017), pp. 37–42. DOI: 10.1016/j.radonc.2017.02.013.
- [187] M. Niemeijer, M. K. Garvin, K. Lee, B. van Ginneken, M. D. Abràmoff, and M. Sonka. “Registration of 3D spectral OCT volumes using 3D SIFT feature point matching”. In: *SPIE Medical Imaging*. 2009, p. 72591I. DOI: 10.1117/12.811906.
- [188] M. Niemeijer, K. Lee, M. K. Garvin, M. D. Abràmoff, and M. Sonka. “Registration of 3D spectral OCT volumes combining ICP with a graph-based approach”. In: *SPIE Medical Imaging*. 2012, 83141A. DOI: 10.1117/12.911104.
- [189] C. Noon and J. Bean. “An Efficient Transformation Of The Generalized Traveling Salesman Problem”. In: *INFOR: Information Systems and Operational Research* 31.1 (1993), pp. 39–44. DOI: 10.1080/03155986.1993.11732212.

- [190] D. J. O'Brien, J. Dolan, S. Pencea, N. Schupp, and G. O. Sawakuchi. "Relative dosimetry with an MR-linac: Response of ion chambers, diamond, and diode detectors for off-axis, depth dose, and output factor measurements". In: *Medical Physics* 45.2 (2017), pp. 884–897. DOI: 10.1002/mp.12699.
- [191] T. O'Shea, J. Bamber, D. Fontanarosa, S. van der Meer, F. Verhaegen, and E. Harris. "Review of ultrasound image guidance in external beam radiotherapy part II: intra-fraction motion management and novel applications". In: *Physics in Medicine & Biology* 61.8 (2016), R90–R137. DOI: 10.1088/0031-9155/61/8/r90.
- [192] C. C. Okoye, R. B. Patel, S. Hasan, T. Podder, A. Khouri, J. Fabien, Y. Zhang, D. Dobbins, J. W. Sohn, J. Yuan, M. Yao, M. Machtay, A. E. Sloan, J. Miller, and S. S. Lo. "Comparison of Ray Tracing and Monte Carlo Calculation Algorithms for Thoracic Spine Lesions Treated With CyberKnife-Based Stereotactic Body Radiation Therapy". In: *Technology in Cancer Research & Treatment* 15.1 (2015), pp. 196–202. DOI: 10.1177/1533034614568026.
- [193] J. Olsen, J. Holmes, and G. B. E. Jemec. "Advances in optical coherence tomography in dermatology—a review". In: *Journal of Biomedical Optics* 23.04 (2018), p. 040901. DOI: 10.1117/1.jbo.23.4.040901.
- [194] K. Otto. "Volumetric modulated arc therapy: IMRT in a single gantry arc". In: *Medical Physics* 35.1 (2007), pp. 310–317. DOI: 10.1118/1.2818738.
- [195] I. Paddick. "A simple scoring ratio to index the conformity of radiosurgical treatment plans". In: *Journal of Neurosurgery* 93.supplement_3 (2000), pp. 219–222. DOI: 10.3171/jns.2000.93.supplement_3.0219.
- [196] C. Paganelli, B. Whelan, M. Peroni, P. Summers, M. Fast, T. van de Lindt, J. McClelland, B. Eiben, P. Keall, T. Lomax, M. Riboldi, and G. Baroni. "MRI-guidance for motion management in external beam radiotherapy: current status and future challenges". In: *Physics in Medicine & Biology* 63.22 (2018), 22TR03. DOI: 10.1088/1361-6560/aaebcf.
- [197] E. P. P. Pang, K. Knight, S. Y. Park, W. Lian, Z. Master, M. Baird, J. W. X. Chan, M. L. C. Wang, T. W. K. Tan, M. L. K. Chua, E. T. Chua, W. S. Looi, W. L. Nei, and J. K. L. Tuan. "Duration-dependent margins for prostate radiotherapy—a practical motion mitigation strategy". In: *Strahlentherapie und Onkologie* 196.7 (2020), pp. 657–663. DOI: 10.1007/s00066-019-01558-y.
- [198] F. C. Park and B. J. Martin. "Robot sensor calibration: solving $AX=XB$ on the Euclidean group". In: *IEEE Transactions on Robotics and Automation* 10.5 (1994), pp. 717–721. DOI: 10.1109/70.326576.
- [199] T. Pawlicki, D. J. Scanderbeg, and G. Starkschall. *Hendee's Radiation Therapy Physics*. 4th ed. John Wiley & Sons Inc, 2016. DOI: 10.1002/9781118575338.
- [200] K. Peignaux, G. Truc, I. Barillot, A. Ammor, S. Naudy, G. Créhange, and P. Maingon. "Clinical assessment of the use of the Sonarray system for daily prostate localization". In: *Radiotherapy and Oncology* 81.2 (2006), pp. 176–178. DOI: 10.1016/j.radonc.2006.08.027.
- [201] C. Peng, G. Chen, E. Ahunbay, D. Wang, C. Lawton, and X. A. Li. "Validation of an online replanning technique for prostate adaptive radiotherapy". In: *Physics in Medicine & Biology* 56.12 (2011), pp. 3659–3668. DOI: 10.1088/0031-9155/56/12/013.

Bibliography

- [202] T. C. Poon and R. N. Rohling. “Comparison of calibration methods for spatial tracking of a 3-D ultrasound probe”. In: *Ultrasound in Medicine & Biology* 31.8 (2005), pp. 1095–1108. DOI: 10.1016/j.ultrasmedbio.2005.04.003.
- [203] A. M. Priester, S. Natarajan, and M. O. Culjat. “Robotic ultrasound systems in medicine”. In: *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* 60.3 (2013), pp. 507–523. DOI: 10.1109/tuffc.2013.2593.
- [204] B. W. Raaymakers, I. M. Jürgenliemk-Schulz, G. H. Bol, M. Glitzner, A. N. T. J. Kotte, B. van Asselen, J. C. J. de Boer, J. J. Bluemink, S. L. Hackett, M. A. Moerland, S. J. Woodings, J. W. H. Wolthaus, H. M. van Zijp, M. E. P. Philippens, R. Tijssen, J. G. M. Kok, E. N. de Groot-van Breugel, I. Kiekebosch, L. T. C. Meijers, C. N. Nomden, G. G. Sikkes, P. A. H. Doornaert, W. S. C. Eppinga, N. Kasperts, L. G. W. Kerkmeijer, J. H. A. Tersteeg, K. J. Brown, B. Pais, P. Woodhead, and J. J. W. Lagendijk. “First patients treated with a 1.5 T MRI-Linac: clinical proof of concept of a high-precision, high-field MRI guided radiotherapy treatment”. In: *Physics in Medicine & Biology* 62.23 (2017), pp. L41–L50. DOI: 10.1088/1361-6560/aa9517.
- [205] B. W. Raaymakers, J. J. W. Lagendijk, J. Overweg, J. G. M. Kok, A. J. E. Raaijmakers, E. M. Kerkhof, R. W. van der Put, I. Meijsing, S. P. M. Crijs, F. Benedosso, M. van Vulpen, C. H. W. de Graaff, J. Allen, and K. J. Brown. “Integrating a 1.5 T MRI scanner with a 6 MV accelerator: proof of concept”. In: *Physics in Medicine & Biology* 54.12 (2009), N229–N237. DOI: 10.1088/0031-9155/54/12/n01.
- [206] O. Rajput, S.-T. Antoni, C. Otte, T. Saathoff, L. Matthäus, and A. Schlaefer. “High accuracy 3D data acquisition using co-registered OCT and kinect”. In: *IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems*. 2016, pp. 32–37. DOI: 10.1109/mfi.2016.7849463.
- [207] N. Ramakrishna, F. Rosca, S. Friesen, E. Tezcanli, P. Zygmanszki, and F. Hacker. “A clinical comparison of patient setup and intra-fraction motion using frame-based radiosurgery versus a frameless image-guided radiosurgery system for intracranial lesions”. In: *Radiotherapy and Oncology* 95.1 (2010), pp. 109–115. DOI: 10.1016/j.radonc.2009.12.030.
- [208] B. S. Reddy and B. N. Chatterji. “An FFT-based technique for translation, rotation, and scale-invariant image registration”. In: *IEEE Transactions on Image Processing* 5.8 (1996), pp. 1266–1271. DOI: 10.1109/83.506761.
- [209] P. Refregier. “Optimal trade-off filters for noise robustness, sharpness of the correlation peak, and Horner efficiency”. In: *Optics Letters* 16.11 (1991), pp. 829–831. DOI: 10.1364/ol.16.000829.
- [210] A. K. Richardson and P. Jacobs. “Intrafraction monitoring of prostate motion during radiotherapy using the Clarity® Autoscan Transperineal Ultrasound (TPUS) system”. In: *Radiography* 23.4 (2017), pp. 310–313. DOI: 10.1016/j.radi.2017.07.003.
- [211] A. Richter, F. Exner, S. Weick, I. Lawrenz, B. Polat, M. Flentje, and F. Mantel. “Evaluation of intrafraction prostate motion tracking using the Clarity Autoscan system for safety margin validation”. In: *Zeitschrift für Medizinische Physik* 30.2 (2020), pp. 135–141. DOI: 10.1016/j.zemedi.2019.12.004.

- [212] L. Richter, P. Trillenber, A. Schweikard, and A. Schlaefer. “Stimulus Intensity for Hand Held and Robotic Transcranial Magnetic Stimulation”. In: *Brain Stimulation* 6.3 (2013), pp. 315–321. DOI: 10.1016/j.brs.2012.06.002.
- [213] Robert Koch Institut. *Bericht zum Krebsgeschehen in Deutschland 2016*. 2016. URL: https://www.krebsdaten.de/Krebs/DE/Content/Publikationen/Krebs_in_Deutschland/krebs_in_deutschland_node.html (visited on 06/28/2020).
- [214] J. M. Rubin, M. Feng, S. W. Hadley, J. B. Fowlkes, and J. D. Hamilton. “Potential Use of Ultrasound Speckle Tracking for Motion Management During Radiotherapy”. In: *Journal of Ultrasound in Medicine* 31.3 (2012), pp. 469–481. DOI: 10.7863/jum.2012.31.3.469.
- [215] K. J. Ruchala, G. H. Olivera, E. A. Schloesser, and T. R. Mackie. “Megavoltage CT on a tomotherapy system”. In: *Physics in Medicine & Biology* 44.10 (1999), pp. 2597–2621. DOI: 10.1088/0031-9155/44/10/316.
- [216] B. J. Salter, B. Wang, M. W. Szegedi, P. Rassiah-Szegedi, D. C. Shrieve, R. Cheng, and M. Fuss. “Evaluation of alignment error due to a speed artifact in stereotactic ultrasound image guidance”. In: *Physics in Medicine & Biology* 53.23 (2008), N437–N445. DOI: 10.1088/0031-9155/53/23/n03.
- [217] B. J. Salter, M. Szegedi, C. Boehm, V. Sarkar, P. Rassiah-Szegedi, B. Wang, H. Zhao, J. Huang, L. Huang, K. Kokeny, and J. D. Tward. “Comparison of 2 transabdominal ultrasound image guidance techniques for prostate and prostatic fossa radiation therapy”. In: *Practical Radiation Oncology* 7.2 (2017), e99–e107. DOI: 10.1016/j.prro.2016.07.001.
- [218] E. A. T. Say, S. Ferenczy, G. N. Magrath, W. A. Samara, C. T. L. Khoo, and C. L. Shields. “Image Quality and Artifacts on Optical Coherence Tomography Angiography”. In: *Retina* 37.9 (2017), pp. 1660–1673. DOI: 10.1097/iae.0000000000001414.
- [219] T. J. Scarbrough, N. M. Golden, J. Y. Ting, C. D. Fuller, A. Wong, P. A. Kupelian, and C. R. Thomas. “Comparison of ultrasound and implanted seed marker prostate localization methods: Implications for image-guided radiotherapy”. In: *International Journal of Radiation Oncology*Biophysics*Physics* 65.2 (2006), pp. 378–387. DOI: 10.1016/j.ijrobp.2006.01.008.
- [220] A. Schlaefer, J. Fisseler, S. Dieterich, H. Shiomi, K. Cleary, and A. Schweikard. “Feasibility of four-dimensional conformal planning for robotic radiosurgery”. In: *Medical Physics* 32.12 (2005), pp. 3786–3792. DOI: 10.1118/1.2122607.
- [221] A. Schlaefer and A. Schweikard. “Stepwise multi-criteria optimization for robotic radiosurgery”. In: *Medical Physics* 35.5 (2008), pp. 2094–2103. DOI: 10.1118/1.2900716.
- [222] A. Schlaefer, T. Viulet, A. Muacevic, and C. Fürweger. “Multicriteria optimization of the spatial dose distribution”. In: *Medical Physics* 40.12 (2013), p. 121720. DOI: 10.1118/1.4828840.
- [223] A. Schlaefer and S. Dieterich. “Feasibility of case-based beam generation for robotic radiosurgery”. In: *Artificial Intelligence in Medicine* 52.2 (2011), pp. 67–75. DOI: 10.1016/j.artmed.2011.04.008.

Bibliography

- [224] J. Schlosser, R. H. Gong, R. Bruder, A. Schweikard, S. Jang, J. Henrie, A. Kamaya, A. C. Koong, D. T. Chang, and D. Hristov. “Robotic intrafractional US guidance for liver SABR: System design, beam avoidance, and clinical imaging”. In: *Medical Physics* 43.11 (2016), pp. 5951–5963. DOI: 10.1118/1.4964454.
- [225] J. Schlosser and D. Hristov. “Radiolucent 4D Ultrasound Imaging: System Design and Application to Radiotherapy Guidance”. In: *IEEE Transactions on Medical Imaging* 35.10 (2016), pp. 2292–2300. DOI: 10.1109/tmi.2016.2559499.
- [226] J. Schlosser, C. Kirmizibayrak, V. Shamdasani, S. Metz, and D. Hristov. “Automatic 3D ultrasound calibration for image guided therapy using intramodality image registration”. In: *Physics in Medicine & Biology* 58.21 (2013), pp. 7481–7496. DOI: 10.1088/0031-9155/58/21/7481.
- [227] J. Schlosser, K. Salisbury, and D. Hristov. “Telerobotic system concept for real-time soft-tissue imaging during radiotherapy beam delivery”. In: *Medical Physics* 37.12 (2010), pp. 6357–6367. DOI: 10.1118/1.3515457.
- [228] M. Schlüter, C. Fürweger, and A. Schlaefer. “Optimizing Configurations for 7-DoF Robotic Ultrasound Guidance in Radiotherapy of the Prostate”. In: *Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. 2019, pp. 6983–6986. DOI: 10.1109/EMBC.2019.8857245.
- [229] M. Schlüter, M. M. Fuh, S. Maier, C. Otte, P. Kiani, N.-O. Hansen, R. J. D. Miller, H. Schlüter, and A. Schlaefer. “Towards OCT-Navigated Tissue Ablation with a Picosecond Infrared Laser (PIRL) and Mass-Spectrometric Analysis”. In: *Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. 2019, pp. 158–161. DOI: 10.1109/embc.2019.8856808.
- [230] M. Schlüter, C. Fürweger, and A. Schlaefer. “Optimizing robot motion for robotic ultrasound-guided radiation therapy”. In: *Physics in Medicine & Biology* (2019), p. 195012. DOI: 10.1088/1361-6560/ab3bfb.
- [231] M. Schlüter, S. Gerlach, C. Fürweger, and A. Schlaefer. “Analysis and optimization of the robot setup for robotic-ultrasound-guided radiation therapy”. Presented at International Congress and Exhibition on Computer Assisted Radiology and Surgery. 2019.
- [232] M. Schlüter, S. Gerlach, C. Fürweger, and A. Schlaefer. “Analysis and optimization of the robot setup for robotic-ultrasound-guided radiation therapy”. In: *International Journal of Computer Assisted Radiology and Surgery* 14.8 (2019), pp. 1379–1387. DOI: 10.1007/s11548-019-02009-w.
- [233] M. Schlüter, L. Glandorf, M. Gromniak, T. Saathoff, and A. Schlaefer. “Concept for Markerless 6D Tracking Employing Volumetric Optical Coherence Tomography”. In: *Sensors* 20.9 (2020), p. 2678. DOI: 10.3390/s20092678.
- [234] M. Schlüter, L. Glandorf, J. Sprenger, M. Gromniak, M. Neidhardt, T. Saathoff, and A. Schlaefer. “High-Speed Markerless Tissue Motion Tracking Using Volumetric Optical Coherence Tomography Images”. In: *IEEE International Symposium on Biomedical Imaging*. 2020, pp. 1979–1982. DOI: 10.1109/ISBI45749.2020.9098448.
- [235] M. Schlüter, C. Otte, T. Saathoff, N. Gessert, and A. Schlaefer. “Feasibility of a markerless tracking system based on optical coherence tomography”. In: *SPIE Medical Imaging*. 2019, p. 1095107. DOI: 10.1117/12.2512178.

- [236] M. Schlüter, D. Schmitt, C. Fürweger, A. Schweikard, and A. Schlaefer. “Treatment Planning”. In: *CyberKnife NeuroRadiosurgery. A practical Guide*. Ed. by A. Conti, P. Romanelli, E. Pantelis, S. G. Soltys, Y. H. Cho, and M. Lim. Springer, 2020, pp. 59–73. DOI: 10.1007/978-3-030-50668-1_5.
- [237] E. Schüler, A. Lo, C. F. Chuang, S. G. Soltys, E. L. Pollom, and L. Wang. “Clinical impact of the VOLO optimizer on treatment plan quality and clinical treatment efficiency for CyberKnife”. In: *Journal of Applied Clinical Medical Physics* 21.5 (2020), pp. 38–47. DOI: 10.1002/acm2.12851.
- [238] P. Schünemann. “Design and Implementation of a 3D-OCT Tracking System”. Master’s thesis. Hamburg University of Technology, Apr. 19, 2016.
- [239] J. Schwaab, M. Prall, C. Sarti, R. Kaderka, C. Bert, C. Kurz, K. Parodi, M. Guenther, and J. E. Jenne. “Ultrasound tracking for intra-fractional motion compensation in radiation therapy”. In: *Physica Medica* 30.5 (2014), pp. 578–582. DOI: 10.1016/j.ejmp.2014.03.003.
- [240] A. Schweikard, M. Bodduluri, and J. R. Adler. “Planning for camera-guided robotic radiosurgery”. In: *IEEE Transactions on Robotics and Automation* 14.6 (1998), pp. 951–962. DOI: 10.1109/70.736778.
- [241] A. Schweikard, G. Glosser, M. Bodduluri, M. J. Murphy, and J. R. Adler. “Robotic Motion Compensation for Respiratory Movement during Radiosurgery”. In: *Computer Aided Surgery* 5.4 (2000), pp. 263–277. DOI: 10.3109/10929080009148894.
- [242] A. Schweikard, A. Schlaefer, and J. R. Adler. “Resampling: an optimization method for inverse planning in robotic radiosurgery”. In: *Medical Physics* 33.11 (2006), pp. 4005–4011. DOI: 10.1118/1.2357020.
- [243] P. K. Seitz, B. Baumann, W. Johnen, C. Lissek, J. Seidel, and R. Bendl. “Development of a robot-assisted ultrasound-guided radiation therapy (USgRT)”. In: *International Journal of Computer Assisted Radiology and Surgery* 15.3 (2019), pp. 491–501. DOI: 10.1007/s11548-019-02104-y.
- [244] H. T. Şen, M. A. L. Bell, I. Iordachita, J. Wong, and P. Kazanzides. “A cooperatively controlled robot for ultrasound monitoring of radiation therapy”. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*. 2013, pp. 3071–3076. DOI: 10.1109/iros.2013.6696791.
- [245] H. T. Şen, M. A. L. Bell, Y. Zhang, K. Ding, J. Wong, I. Iordachita, and P. Kazanzides. “System integration and preliminary in-vivo experiments of a robot for ultrasound guidance and monitoring during radiotherapy”. In: *International Conference on Advanced Robotics*. 2015, pp. 53–59. DOI: 10.1109/icar.2015.7251433.
- [246] H. T. Şen, M. A. L. Bell, Y. Zhang, K. Ding, E. Boctor, J. Wong, I. Iordachita, and P. Kazanzides. “System Integration and In Vivo Testing of a Robot for Ultrasound Guidance and Monitoring During Radiotherapy”. In: *IEEE Transactions on Biomedical Engineering* 64.7 (2017), pp. 1608–1618. DOI: 10.1109/tbme.2016.2612229.
- [247] H. T. Şen, A. Cheng, K. Ding, E. Boctor, J. Wong, I. Iordachita, and P. Kazanzides. “Cooperative Control with Ultrasound Guidance for Radiation Therapy”. In: *Frontiers in Robotics and AI* 3 (2016), p. 49. DOI: 10.3389/frobt.2016.00049.

Bibliography

- [248] A. P. Shah, P. A. Kupelian, B. J. Waghorn, T. R. Willoughby, J. M. Rineer, R. R. Mañon, M. A. Vollenweider, and S. L. Meeks. “Real-Time Tumor Tracking in the Lung Using an Electromagnetic Tracking System”. In: *International Journal of Radiation Oncology*Biophysics* 86.3 (2013), pp. 477–483. DOI: 10.1016/j.ijrobp.2012.12.030.
- [249] A. W. M. Sharfo, M. L. P. Dirkx, S. Breedveld, A. M. Romero, and B. J. M. Heijmen. “VMAT plus a few computer-optimized non-coplanar IMRT beams (VMAT+) tested for liver SBRT”. In: *Radiotherapy and Oncology* 123.1 (2017), pp. 49–56. DOI: 10.1016/j.radonc.2017.02.018.
- [250] X. H. Shi, Y. C. Liang, H. P. Lee, C. Lu, and Q. X. Wang. “Particle swarm optimization-based algorithms for TSP and generalized TSP”. In: *Information Processing Letters* 103.5 (2007), pp. 169–176. DOI: 10.1016/j.ip1.2007.03.010.
- [251] R. L. Siegel, K. D. Miller, and A. Jemal. “Cancer statistics, 2019”. In: *CA: A Cancer Journal for Clinicians* 69.1 (2019), pp. 7–34. DOI: 10.3322/caac.21551.
- [252] A. R. Smith. “Vision 20/20: Proton therapy”. In: *Medical Physics* 36.2 (2009), pp. 556–568. DOI: 10.1118/1.3058485.
- [253] G. Smyth, J. C. Bamber, P. M. Evans, and J. L. Bedford. “Trajectory optimization for dynamic couch rotation during volumetric modulated arc radiotherapy”. In: *Physics in Medicine & Biology* 58.22 (2013), pp. 8163–8177. DOI: 10.1088/0031-9155/58/22/8163.
- [254] T. Sothmann, O. Blanck, K. Poels, R. Werner, and T. Gauer. “Real time tracking in liver SBRT: comparison of CyberKnife and Vero by planning structure-based γ -evaluation and dose-area-histograms”. In: *Physics in Medicine & Biology* 61.4 (2016), pp. 1677–1691. DOI: 10.1088/0031-9155/61/4/1677.
- [255] V. J. Srinivasan, S. Sakadžić, I. Gorczynska, S. Ruvinskaya, W. Wu, J. G. Fujimoto, and D. A. Boas. “Quantitative cerebral blood flow with Optical Coherence Tomography”. In: *Optics Express* 18.3 (2010), pp. 2477–2494. DOI: 10.1364/oe.18.002477.
- [256] M. T. Studenski, V. Bar-Ad, J. Siglin, D. Cignetti, J. Curry, M. Tuluc, and A. S. Harrison. “Clinical experience transitioning from IMRT to VMAT for head and neck cancer”. In: *Medical Dosimetry* 38.2 (2013), pp. 171–175. DOI: 10.1016/j.meddos.2012.10.009.
- [257] T. L. Szabo. *Diagnostic Ultrasound Imaging: Inside Out*. 2nd ed. Elsevier, 2014. DOI: 10.1016/c2011-0-07261-7.
- [258] J. N. Tehrani, R. T. O’Brien, P. R. Poulsen, and P. Keall. “Real-time estimation of prostate tumor rotation and translation with a kV imaging system based on an iterative closest point algorithm”. In: *Physics in Medicine & Biology* 58.23 (2013), pp. 8517–8533. DOI: 10.1088/0031-9155/58/23/8517.
- [259] R. K. ten Haken, J. D. Forman, D. K. Heimbürger, A. Gerhardsson, D. L. McShan, C. Perez-Tamayo, S. L. Schoepfel, and A. S. Lichter. “Treatment planning issues related to prostate movement in response to differential filling of the rectum and bladder”. In: *International Journal of Radiation Oncology*Biophysics* 20.6 (1991), pp. 1317–1324. DOI: 10.1016/0360-3016(91)90244-x.
- [260] The Royal College of Radiologists. *Radiotherapy dose fractionation. BFCO(19)3*. 2019. URL: <https://www.rcr.ac.uk/publication/radiotherapy-dose-fractionation-third-edition> (visited on 06/30/2020).

- [261] R. S. Thing, U. Bernchou, E. Mainegra-Hing, O. Hansen, and C. Brink. “Hounsfield unit recovery in clinical cone beam CT images of the thorax acquired for image guided radiation therapy”. In: *Physics in Medicine & Biology* 61.15 (2016), pp. 5781–5802. DOI: 10.1088/0031-9155/61/15/5781.
- [262] M. S. Thomsen, U. Harrov, W. Fledelius, and P. R. Poulsen. “Inter- and intra-fraction geometric errors in daily image-guided radiotherapy of free-breathing breast cancer patients measured with continuous portal imaging”. In: *Acta Oncologica* 53.6 (2014), pp. 802–808. DOI: 10.3109/0284186x.2014.905700.
- [263] X. Tong, X. Chen, J. Li, Q. Xu, M.-h. Lin, L. Chen, R. A. Price, and C.-M. Ma. “Intrafractional prostate motion during external beam radiotherapy monitored by a real-time target localization system”. In: *Journal of Applied Clinical Medical Physics* 16.2 (2015), pp. 51–61. DOI: 10.1120/jacmp.v16i2.5013.
- [264] B. A. Torkian, S. Guo, A. W. Jahng, L.-H. L. Liaw, Z. Chen, and B. J. F. Wong. “Noninvasive Measurement of Ablation Crater Size and Thermal Injury after CO₂ Laser in the Vocal Cord with Optical Coherence Tomography”. In: *Otolaryngology–Head and Neck Surgery* 134.1 (2006), pp. 86–91. DOI: 10.1016/j.otohns.2005.09.016.
- [265] R. Y. Tsai and R. K. Lenz. “A new technique for fully autonomous and efficient 3D robotics hand/eye calibration”. In: *IEEE Transactions on Robotics and Automation* 5.3 (1989), pp. 345–358. DOI: 10.1109/70.34770.
- [266] J. Unkelbach, M. Alber, M. Bangert, R. Bokrantz, T. C. Y. Chan, J. O. Deasy, A. Fredriksson, B. L. Gorissen, M. van Herk, W. Liu, H. Mahmoudzadeh, O. Nohadani, J. V. Siebers, M. Witte, and H. Xu. “Robust radiotherapy planning”. In: *Physics in Medicine & Biology* 63.22 (2018), 22TR02. DOI: 10.1088/1361-6560/aae659.
- [267] A. Vanhanen, H. Syrén, and M. Kapanen. “Localization accuracy of two electromagnetic tracking systems in prostate cancer radiotherapy: A comparison with fiducial marker based kilovoltage imaging”. In: *Physica Medica* 56 (2018), pp. 10–18. DOI: 10.1016/j.ejmp.2018.11.007.
- [268] T. Violet. “An interactive approach to localised dose planning for radiotherapy”. PhD thesis. University of Lübeck, 2016.
- [269] T. Wang, T. Pfeiffer, E. Regar, W. Wieser, H. van Beusekom, C. T. Lancee, G. Springeling, I. Krabbendam-Peters, A. F. van der Steen, R. Huber, and G. van Soest. “Heartbeat OCT and Motion-Free 3D In Vivo Coronary Artery Microscopy”. In: *JACC: Cardiovascular Imaging* 9.5 (2016), pp. 622–623. DOI: 10.1016/j.jcmg.2015.08.010.
- [270] Y. Wang, J. Li, K. Gao, and Q. Pan. “Memetic Algorithm based on Improved Inver-over operator and Lin-Kernighan local search for the Euclidean traveling salesman problem”. In: *Computers & Mathematics with Applications* 62.7 (2011), pp. 2743–2754. DOI: 10.1016/j.camwa.2011.06.063.
- [271] J. Weiss, N. Rieke, M. A. Nasser, M. Maier, A. Eslami, and N. Navab. “Fast 5DOF needle tracking in iOCT”. In: *International Journal of Computer Assisted Radiology and Surgery* 13.6 (2018), pp. 787–796. DOI: 10.1007/s11548-018-1751-5.

Bibliography

- [272] J. B. West, J. M. Fitzpatrick, S. A. Toms, C. R. Maurer Jr, and R. J. Maciunas. “Fiducial Point Placement and the Accuracy of Point-based, Rigid Body Registration”. In: *Neurosurgery* 48.4 (2001), pp. 810–817. DOI: 10.1097/00006123-200104000-00023.
- [273] C. Western, D. Hristov, and J. Schlosser. “Ultrasound imaging in radiation therapy: From interfractional to intrafractional guidance”. In: *Cureus* 7.6 (2015), e280. DOI: 10.7759/cureus.280.
- [274] E. E. Wilcox, G. M. Daskalov, H. Lincoln, R. C. Shumway, B. M. Kaplan, and J. M. Colasanto. “Comparison of Planned Dose Distributions Calculated by Monte Carlo and Ray-Trace Algorithms for the Treatment of Lung Tumors With CyberKnife: A Preliminary Study in 33 Patients”. In: *International Journal of Radiation Oncology*Biography*Physics* 77.1 (2010), pp. 277–284. DOI: 10.1016/j.ijrobp.2009.08.001.
- [275] T. Wissel, P. Stüber, B. Wagner, R. Bruder, C. Erdmann, C.-S. Deutz, B. Sack, J. Manit, A. Schweikard, and F. Ernst. “Enhanced Optical Head Tracking for Cranial Radiation Therapy: Supporting Surface Registration by Cutaneous Structures”. In: *International Journal of Radiation Oncology*Biography*Physics* 95.2 (2016), pp. 810–817. DOI: 10.1016/j.ijrobp.2016.01.041.
- [276] D. Wolff, F. Stieler, G. Welzel, F. Lorenz, Y. Abo-Madyan, S. Mai, C. Herskind, M. Polednik, V. Steil, F. Wenz, and F. Lohr. “Volumetric modulated arc therapy (VMAT) vs. serial tomotherapy, step-and-shoot IMRT and 3D-conformal RT for treatment of prostate cancer”. In: *Radiotherapy and Oncology* 93.2 (2009), pp. 226–233. DOI: 10.1016/j.radonc.2009.08.011.
- [277] Y. Xu, T. Bai, H. Yan, L. Ouyang, A. Pompos, J. Wang, L. Zhou, S. B. Jiang, and X. Jia. “A practical cone-beam CT scatter correction method with optimized Monte Carlo simulations for image-guided radiation therapy”. In: *Physics in Medicine & Biology* 60.9 (2015), pp. 3567–3587. DOI: 10.1088/0031-9155/60/9/3567.
- [278] Z.-Y. Yang, Y. Chang, H.-Y. Liu, G. Liu, and Q. Li. “Target margin design for real-time lung tumor tracking stereotactic body radiation therapy using CyberKnife Xsight Lung Tracking System”. In: *Scientific Reports* 7.1 (2017), p. 10826. DOI: 10.1038/s41598-017-11128-w.
- [279] Z. Yaniv, E. Wilson, D. Lindisch, and K. Cleary. “Electromagnetic tracking in the clinical environment”. In: *Medical Physics* 36.3 (2009), pp. 876–892. DOI: 10.1118/1.3075829.
- [280] S. Yartsev and G. Bauman. “Target margins in radiotherapy of prostate cancer”. In: *The British Journal of Radiology* 89.1067 (2016), p. 20160312. DOI: 10.1259/bjr.20160312.
- [281] S. Yoo and F.-F. Yin. “Dosimetric feasibility of cone-beam CT-based treatment planning compared to CT-based treatment planning”. In: *International Journal of Radiation Oncology*Biography*Physics* 66.5 (2006), pp. 1553–1561. DOI: 10.1016/j.ijrobp.2006.08.031.

- [282] V. Y. Yu, A. Landers, K. Woods, D. Nguyen, M. Cao, D. Du, R. K. Chin, K. Sheng, and T. B. Kaprealian. “A Prospective 4π Radiation Therapy Clinical Study in Recurrent High-Grade Glioma Patients”. In: *International Journal of Radiation Oncology*Biography*Physics* 101.1 (2018), pp. 144–151. DOI: 10.1016/j.ijrobp.2018.01.048.
- [283] Y. Zhang, T. Pfeiffer, M. Weller, W. Wieser, R. Huber, J. Raczkowsky, J. Schipper, H. Wörn, and T. Klenzner. “Optical Coherence Tomography Guided Laser Cochleostomy: Towards the Accuracy on Tens of Micrometer Scale”. In: *BioMed Research International* 2014 (2014), p. 251814. DOI: 10.1155/2014/251814.
- [284] Y. Zhang and H. Wörn. “Optical coherence tomography as highly accurate optical tracking system”. In: *IEEE/ASME International Conference on Advanced Intelligent Mechatronics*. 2014, pp. 1145–1150. DOI: 10.1109/aim.2014.6878235.
- [285] Y. Zhong, K. Stephans, P. Qi, N. Yu, J. Wong, and P. Xia. “Assessing feasibility of real-time ultrasound monitoring in stereotactic body radiotherapy of liver tumors”. In: *Technology in Cancer Research & Treatment* 12.3 (2013), pp. 243–250. DOI: 10.7785/tcrt.2012.500323.
- [286] M. Zhou, M. Hamad, J. Weiss, A. Eslami, K. Huang, M. Maier, C. P. Lohmann, N. Navab, A. Knoll, and M. A. Nasser. “Towards Robotic Eye Surgery: Marker-Free, Online Hand-Eye Calibration Using Optical Coherence Tomography Images”. In: *IEEE Robotics and Automation Letters* 3.4 (2018), pp. 3944–3951. DOI: 10.1109/lra.2018.2858744.