

Using deep learning generated CBCT contours for on-line dose assessment of
prostate SABR treatments

by

Conor Sinclair Smith

B.Sc. (Hons.), University of British Columbia, 2020

A Thesis Submitted in Partial Fulfillment of the
Requirements for the Degree of

MASTER OF SCIENCE

in the Department of Physics and Astronomy

© Conor Sinclair Smith, 2023
University of Victoria

All rights reserved. This Thesis may not be reproduced in whole or in part, by
photocopy or other means, without the permission of the author.

Using deep learning generated CBCT contours for on-line dose assessment of
prostate SABR treatments

by

Conor Sinclair Smith

B.Sc. (Hons.), University of British Columbia, 2020

Supervisory Committee

Dr. E. Chin, Co-Supervisor
(Department of Physics and Astronomy)

Dr. M. Bazalova-Carter, Co-Supervisor
(Department of Physics and Astronomy)

Dr. I. Gagne, Departmental Member
(British Columbia Cancer Agency & Department of Physics and Astronomy)

ABSTRACT

Prostate Stereotactic Ablative Body Radiotherapy (SABR) is a hypofractionated treatment regime where small errors in patient setup have the potential to deliver much higher doses to Organs at Risk (OARs). Inter-fraction variability is reduced by having patients follow a bowel and bladder preparation protocol, however patient adherence is inconsistent and OAR variability is still present. At our centre, radiation therapists use non-dosimetric decision trees (DT) to determine if the patient is sufficiently prepared and positioned, however the application of these decision trees varies between therapists. We used deep neural-network-generated cone-beam computed tomography (CBCT) contours to estimate the daily dose delivered to the rectum and bladder. We compared these estimates to the planned dose-volume metrics to facilitate future development of personalized and quantitative decision trees for prostate SABR treatments.

Two-hundred pre-treatment CBCT scans from 40 prostate SABR patients, each prescribed 40 Gy in five fractions, were contoured using Limbus Contour AI based segmentation software (v1.5.0-D2). Dose-volume histogram (DVH) curves for each fraction's OARs were calculated by combining the planned dose distribution and AI-generated contours using on-line rigid registration data. To evaluate the performance of patient bowel & bladder preparation and SABR decision trees, we calculated the same dose-volume metrics that were used as goals for treatment planning using the DVH curves generated from the daily CBCT contours. CBCT dose-volume metrics that met the preferred goal were considered as having 'no violations', and metrics that did not meet the mandatory goals were classified as 'major violations', with values in between categorized as 'minor violations.'

Twenty-seven percent of fractions had at least one major violation in one of the bladder planning goals, and 34% had minor violations. Fourteen percent of fractions had a rectum metric major violation, with 10% categorized as minor violations. Across their five fraction treatment, five patients had recurring bladder V37 Gy major violations, and 2 patients had rectum V36 Gy major violations. Bowel and bladder preparation played a significant role as OAR differences in position and volume led to many instances of mandatory goals not being met. These results show a clear need for the development of rigorous and personalized decision trees. Online quantitative OAR volume data incorporated into a future DT would likely make it successful.

Contents

Supervisory Committee	ii
Abstract	iii
Table of Contents	iv
List of Tables	vii
List of Figures	viii
Acknowledgements	x
Dedication	xi
1 Introduction	1
2 Background	3
2.1 Physics of Radiation Therapy	3
2.1.1 Photon Interactions	3
2.1.2 Electron Interactions	7
2.2 Computed Tomography	8
2.2.1 Computed Tomography - X-Ray Production	8
2.2.2 Computed Tomography - X-ray Detection	10
2.2.3 Computed Tomography - Image Reconstruction	11
2.2.4 Computed Tomography - Beam Geometry	15
2.3 Medical Linear Accelerators	17
2.4 Treatment Planning in Radiation Therapy	18
2.4.1 Irradiation Target Definitions	19
2.4.2 Dose-Volume constraints	20
2.4.3 External Beam Planning	21

2.4.4	Volumetric Modulated Arc Therapy	22
2.4.5	Dose Calculation	22
2.5	Image-Guided Radiation Therapy	23
2.5.1	Prostate IGRT	24
2.6	Adaptive Radiation Therapy	24
2.7	Auto-contouring	26
2.7.1	Previous Approaches	26
2.7.2	Machine Learning	26
2.7.3	Neural Networks	27
2.7.4	Convolutional Neural Networks	28
2.7.5	U-Net	28
2.7.6	Limbus AI Autocontouring	31
2.8	Prostate Stereotactic Ablative Body Radiotherapy	33
2.8.1	Fractionation	33
2.8.2	SABR dose-volume constraints	34
2.8.3	Preparation and positioning challenges	35
3	Methods	37
3.1	Prostate SABR	37
3.1.1	Patient Preparation	37
3.1.2	Procedure	38
3.1.3	Decision Trees	38
3.2	Calculation of CBCT OAR Dose	42
3.2.1	CBCT Selection	43
3.2.2	CBCT OAR Dose Calculation Method	43
3.2.3	Workflow Automation	46
3.3	Dose Shift Study	46
4	Results	49
4.1	Dose Shift Study	49
4.2	Verification of Method Implementation	53
4.3	DVH Curves Calculated Using CBCT Contours	55
4.4	Bladder	57
4.5	Rectum	63
5	Discussion & Conclusion	68

5.1 Discussion	68
5.2 Conclusion	73
Bibliography	74
A Original Decision Trees	83
B Dose Shift Study	85

List of Tables

Table 2.1 BC Cancer Victoria Prostate SABR planning goals for 40 Gy/5 fraction prescription.	35
Table 3.1 Steps taken to obtain reference and evaluated dose distributions	48
Table 4.1 Average gamma passing rate (2%/2 mm criteria) for five dose approximation scenarios	50
Table 4.2 Dose distributions compared using planning goals for five dose approximation scenarios.	52
Table 4.3 Example Patient CBCT Dose-Volume Metric Dataset	55

List of Figures

Figure 2.1	Photoelectric Effect Diagram [4].	4
Figure 2.2	Compton Scattering Diagram [4].	5
Figure 2.3	Pair Production Diagram [4].	6
Figure 2.4	Bremsstrahlung process [4].	8
Figure 2.5	Basic x-ray tube schematic	9
Figure 2.6	Distribution of x-rays for different x-ray tube voltages [4].	10
Figure 2.7	CT Detector Components	11
Figure 2.8	Measured projections in CT imaging [15].	12
Figure 2.9	Backprojection in CT Imaging	13
Figure 2.10	Fourier reconstruction [12]	14
Figure 2.11	Iterative reconstruction [11]	14
Figure 2.12	Computed Tomography beam geometries	15
Figure 2.13	Cone Beam CT using a flat panel detector	16
Figure 2.14	Core components of a linear accelerator treatment head [22]	18
Figure 2.15	ICRU Report No. 50 Planning Volumes [24].	20
Figure 2.16	Dose-volume constraints	21
Figure 2.17	A simple neural network	27
Figure 2.18	Example of a max pooling operation.	28
Figure 2.19	Example of U-Net image segmentation trained on a lung X-Ray dataset.	29
Figure 2.20	U-Net Architecture [46]	31
Figure 2.21	Comparison of Limbus deep-learning-generated contours to expert inter-observer-variability	33
Figure 2.22	Example of inter-fraction variation of the rectum and bladder.	36
Figure 3.1	Bladder Decision Tree for Prostate SABR Patients.	40
Figure 3.2	Rectum Decision Tree for Prostate SABR Patients.	41
Figure 3.3	Interpolation of dose onto CBCT slices	45

Figure 3.4	CBCT OAR Dose Calculation Workflow	45
Figure 4.1	Verification of Method Implementation	54
Figure 4.2	Fraction classification frequency.	56
Figure 4.3	Classification of dose-volume metrics using daily CBCT bladder contours	57
Figure 4.4	Bladder - Patients by number of major dose violations over the course of their treatment.	58
Figure 4.5	pCT and CBCT bladder volume distributions.	60
Figure 4.6	Bladder - Dose-volume metrics compared to absolute change in volume from pCT to CBCT.	61
Figure 4.7	Bladder - Dose-volume metrics compared to relative change in volume from pCT to CBCT.	62
Figure 4.8	Classification of dose-volume metrics using daily CBCT rectum contours.	63
Figure 4.9	Rectum - Patients by number of major dose violations.	64
Figure 4.10	pCT and CBCT rectum volume distributions.	65
Figure 4.11	Rectum - Dose-volume metrics compared to absolute change in volume from pCT to CBCT.	66
Figure 4.12	Rectum - Dose-volume metrics compared to relative change in volume from pCT to CBCT.	67
Figure A.1	ASSERT Trial Rectum Decision Tree.	83
Figure A.2	ASSERT Trial Bladder Decision Tree.	84
Figure B.1	Dose difference and gamma index plots for Case Ia	85
Figure B.2	Dose difference and gamma index plots for Case Ib	85
Figure B.3	Dose difference and gamma index plots for Case II	86
Figure B.4	Dose difference and gamma index plots for Case IIa	86
Figure B.5	Dose difference and gamma index plots for Case IIb	86

Acknowledgements

I would like to thank:

Dr. Erika Chin for her continued guidance, for making time to meet regularly, for useful and actionable feedback, and for supporting and advocating for me, the list goes on. I tell anyone willing to listen that I had a great supervisor.

Dr. Karl Otto for also taking time to attend meetings, for providing some great code ideas, and for his help preparing me and sending me to COMP.

Dr. Isabelle Gagne for answering all kinds of clinical questions, for also helping me prepare my COMP abstract and presentation, and for providing feedback on some sections of my thesis.

My RO Army Dr. Jordan Wong, Dr. Soha Atallah, Dr. Saibishkumar Elantholi Parameswaran, Dr. Howard Pai, Dr. Abraham Alexander who took the divide and conquer approach to reviewing 200 CBCT contour sets. A special thanks to Dr. Abraham Alexander who also helped with many of my clinical questions and provided insight on clinical significance of our results.

“We’ll make it”

Master Chief

Dedication

I would like to dedicate this work to my parents, Ron Smith and Lynn Sinclair, without whom I would not have had the means and support I needed to pursue my post-secondary education. I am deeply grateful for their encouragement, their unwavering love, and for the sacrifices they have made on my behalf. I will always carry their lessons with me and strive to make them proud. I love you both and I am forever thankful.

Chapter 1

Introduction

Cancer is the cause of death for one in four Canadians, making it the leading cause of death in Canada [1]. It is estimated that 43% of all Canadians will receive a cancer diagnosis in their lifetime [2]. With a growing and aging population, the number of new cancer cases is predicted to increase by 84% in males, and by 74% in females between 2007 and 2032 [3]. According to the Canadian Cancer Society, 24,600 Canadian men will be diagnosed with prostate cancer in 2022, representing 20% of all cancer cases in men for the year. 4,600 of those will sadly succumb to the disease, which is 10% of all cancer-related deaths in males for 2022. In their lifetime, 1 in 8 Canadian men will develop prostate cancer, and 1 in 29 will die from it. With such significant statistics, it is imperative that we are constantly striving for improvements in patient outcomes. Approximately half of all patients diagnosed with cancer receive some form of radiation therapy (RT) as part of their treatment. Considering the forecasted increase in cancer incidence, we also must continue to make improvements in radiation therapy workflows such that centers can maximize their throughput without sacrificing patient outcomes.

Conventional hypofractionation schedules may deliver 60 Gy in 20 fractions for 3 Gy per fraction. Prostate Stereotactic Ablative Body Radiotherapy (SABR) is an ultra-hypofractionated radiotherapy treatment where dose is delivered in 2-5 fractions. For comparison, a SABR treatment of 40 Gy given in 5 fractions delivers 8 Gy per fraction. The higher ablative doses per fraction and small errors in patient setup can potentially result in delivery of higher doses to Organs-at-risk (OARs) and increase the probability of toxicity. The bladder and rectum are adjacent to the prostate and are considered OARs in prostate radiotherapy treatments. Naturally, these OARs are constantly changing shape and volume throughout the day, therefore, the patient's anatomy from their planning computed tomography (pCT) scan is un-

likely to be identical to their anatomy on the day of treatment 2-3 weeks later. The change of rectum and bladder volume from plan imaging to subsequent treatment fractions is known as inter-fraction variation and is a significant source of uncertainty in prostate RT. The general approach to reduce OAR inter-fraction variation is to use bowel and bladder preparation protocols. These protocols are designed to maintain a consistently full bladder and empty rectum both when the patient presents for their initial pCT and throughout their subsequent treatment appointments. Despite this strict protocol, it is impossible to completely eliminate all OAR volume variations due to variability in patient compliance.

To account for this reduced but remaining daily anatomical variation, BC Cancer Victoria's radiation therapists use specially designed bowel and bladder decision trees to visually assess patient pre-treatment Cone-Beam Computed Tomography (CBCT) images and determine if the patient setup is acceptable. However, these decision trees are qualitative in nature, complicated, time consuming, and not personalized to the patient's specific anatomy.

In this work, our goal was to develop the foundation of an efficient tool that could be used at treatment to quantify the expected OAR dose of the day and guide online clinical decision-making that is personalized to the patient's specific anatomy. We combined Limbus AI's auto-contouring software with patients' planned dose data to estimate the dose-volume histogram (DVH) curves for the bladder and rectum in prostate SABR patients for each treatment fraction. We retrospectively reviewed 40 prostate SABR patients who were treated with 40 Gy in 5 fractions and had a SpaceOAR hydrogel implant.

Chapter 2

Background

2.1 Physics of Radiation Therapy

2.1.1 Photon Interactions

Ionizing Electromagnetic Radiation

Photons consist of packets of electromagnetic radiation that oscillate their electric and magnetic field vectors perpendicular to the direction of the photon's motion. Each photon carries energy $h\nu$, where h is Planck's constant, and ν is the photon's frequency. In radiation therapy treatments, we are interested in photons with sufficient energy to ionize matter, as this is what ultimately leads to the death of cancer cells after a chain of events. In Computed Tomography (CT) imaging, the beam of photons must be energetic enough for some of the photons to pass through the patient and for some to be attenuated, but not so energetic that unnecessary dose is delivered. After all, we are often CT scanning patients in order to figure out precisely where to put the therapeutic dose. CT scanners and linear accelerators both use ionizing radiation at different energy ranges for different intended outcomes. CT scanner photon energies range from 40 kV to 140 kV whereas linear accelerator beams can range from 1 MV up to 20 MV. It is therefore important to understand the ways that ionizing photons interact with matter. As medical physicists, we are particularly interested in photon beams at clinical energies and how they interact with human tissues. While there are more interaction modes than the three discussed in the following sections, the photoelectric effect, Compton scattering, and pair production are the dominant modes of interactions between diagnostic and therapeutic photon energies and human tissues.

Photoelectric Effect

The photoelectric effect is an interaction where a photon is absorbed by an atom and one of the electrons from the atom's K, L, M, or N shells is subsequently ejected (Figure 2.1). The atom absorbs the entirety of the photon's energy $h\nu$ and the ejected photoelectron's energy is therefore $h\nu - E_B$, where E_B is the binding energy of the shell from which the electron was ejected. With a vacancy in the atom's electron shells, an electron from a higher shell may de-excite to fill the vacancy emitting a 'characteristic x-ray' photon in the process. The energy of the characteristic x-ray is equal to the difference in binding energies of the initial and final shell. Alternatively, the energy from the relaxing electron may be transferred to another electron in an outer shell, which is then ejected from the atom as an 'Meitner-Auger electron' [4].

The K-shell binding energy depends on the atomic number of the atom in the interaction; as atomic number increases, so does the K-shell binding energy. In clinical scenarios, patient tissue mostly consists of low atomic number ($Z_{eff} \approx 6.46 - 7.64$ [5]) elements with low K-shell binding energies (K-edge < 1 keV [6]). The result is that the characteristic x-rays produced in clinical photoelectric interactions are low energy and are considered to be locally absorbed [7]. The mass attenuation coefficient for the photoelectric effect is proportional to Z^3/ν^3 , where Z is the atomic number of the target atom. With the low effective atomic numbers of human tissues, the photoelectric effect is only dominant in low-energy diagnostic imaging applications, and is responsible for the contrast between bone and soft tissue. Due to the inverse cubic dependence on frequency, this interaction mode is relatively unimportant in higher energy therapeutic beams [8].




Figure 2.1: Photoelectric Effect Diagram [4].

Compton Scattering

In the Compton interaction, a high energy photon collides with a loosely bound ($E_B \ll h\nu_0$) electron, after which the photon scatters with reduced energy ($h\nu'$) at an angle ϕ , and the electron is ejected from the atom at angle θ leaving behind an ion. This process is depicted in Figure 2.2. The physics of the Compton interaction is conventionally explained through the lens of conservation of energy and momentum. Using these conservation laws, the following relationship for the energy of the scattered photon can be derived [8]:

$$h\nu' = h\nu_0 \frac{1}{1 + \alpha(1 - \cos\phi)}. \quad (2.1)$$

Where $\alpha = h\nu_0/m_0c^2$, and m_0c^2 is the rest energy of the electron (0.511 MeV). Furthermore, the kinetic energy of the Compton electron is [8]:

$$E_{\bar{e}} = h\nu_0 \frac{\alpha(1 - \cos\phi)}{1 + \alpha(1 - \cos\phi)}. \quad (2.2)$$

The probability of a Compton interaction is considered to be independent of atomic number Z , as it primarily depends on the electron density of the medium. In water, the Compton interaction probability increases with photon energy from 10 to 150 keV, after which it decreases with photon energy. Despite this, it is still the predominant mode of interaction between water and photons ranging from 30 keV to 24 MeV, which encompasses photon beams used in radiation therapy [9].




Figure 2.2: Compton Scattering Diagram [4].

Pair Production

Pair production interactions occur when a photon with at least 1.022 MeV interacts with the electromagnetic field of an atomic nucleus and creates an electron-positron pair (Figure 2.3). Photons with less than 1.022 MeV are unable to undergo pair production since the energy is less than the sum of the electron and positron rest mass energies. This would violate conservation of energy if it were to occur. The total energy available to the electron-positron pair is therefore $h\nu - 1.022 \text{ MeV}$. While the energy of the electron and positron can take any distribution, from the electron receiving all the excess energy to the positron receiving it, the most probable energy distribution is that the electron and positron will have equal energies and scatter at the same angle relative to the original photon's direction [5].

The positron from pair production interactions travels a short distance from the site it was born, losing energy through ionization, excitation, and bremsstrahlung. As the energy of the positron decreases, the probability of it interacting with an electron in the medium increases. Eventually the positron will annihilate with an electron creating two 0.511 MeV photons that travel opposite directions in order to conserve momentum. Occasionally, the positron annihilates before losing its kinetic energy, in these *in-flight-annihilations* the two photons will have energy greater than 0.511 MeV.

Since pair production occurs when a photon interacts with a nucleus, the attenuation coefficient for pair production varies approximately with Z^2 per atom, thus it is often not the dominant mode of interaction in human tissues. However, pair productions do occur more often as photon energy increases above the 1.022 MeV threshold, increasing from 6% at 4 MeV to 20% at 7 MeV. These correspond to clinical beams of 12 MV and 21 MV respectively. Photon energy needs to be increased to at least 24 MeV to become the dominant mode of interaction, which approximately corresponds to a 70 MV linac beam [9].




Figure 2.3: Pair Production Diagram [4].

2.1.2 Electron Interactions

The aforementioned modes of interactions between photons and matter create a number of free high energy electrons. The dominant mode of electron energy loss in low atomic number media is through inelastic collisions with other atomic electrons, this leads to either excitation or ionization of the atomic electron. Ionizations may occur if the free electron has more energy than the binding energy of the orbital electron. The atom, missing an electron, becomes a negatively charged ion and the free electron binds with another atom thus creating a positively charged ion. The creation of ions is how the chemical reactions leading to cell death occur, however most of the absorbed energy is converted into heat, which produces no biological effect at the energy levels relevant in radiation therapy [9]. Occasionally, the ionized electron will have enough energy to ionize another electron and creates a secondary electron or δ -ray.

If the free electron does not have sufficient energy to ionize the orbital electron, the interaction will result in the excitation of the orbital electron. As the excited electron returns to its ground state it will emit an x-ray photon, which would continue on to interact with other matter in the ways discussed in Section 2.1.1. At each interaction, the original free electron loses a certain amount of kinetic energy, and the amount of energy lost depends on the electron's energy prior to the interaction and the composition of the medium; the electron can undergo many inelastic collisions before losing all of its kinetic energy.

Alternatively, the free electron can undergo an inelastic collision with the atomic nucleus via the Coulomb force with the nucleus' electromagnetic field. When this occurs, the electron is decelerated rapidly and the change in energy gives rise to 'Bremsstrahlung' photons [10] (Figure 2.4). The interaction can cause the electron to lose any amount of energy up to its original kinetic energy, therefore the Bremsstrahlung photons have a continuous distribution of energies. The energy of the Bremsstrahlung photon depends on the interaction distance between the electron and the atomic nucleus; a particle in a close interaction experiences a large Coulomb force and loses a greater amount of energy, creating a high energy photon. Conversely, electrons in distant interactions create lower energy photons. The probability, however, of close interactions is low and therefore fewer high energy photons are created [11]. This mode of energy loss is more dominant in high atomic number media as the electron's energy loss through Bremsstrahlung is proportional to Z^2 [4]. High

Z materials such as tungsten are used for the production of x-rays in both imaging x-ray tubes and linear accelerator targets.




Figure 2.4: Bremsstrahlung process [4].

2.2 Computed Tomography

Computed Tomography (CT) is an x-ray imaging technique that is used to painlessly peer inside the human body. The invention won Sir Godfrey Hounsfield and Alan Cormack the 1979 Nobel Prize in Medicine. In CT imaging, a patient is positioned in between an x-ray source and detector array in a combination that is known as *the gantry*. The gantry rotates around the patient and acquires multiple measurements of the changes in x-ray intensity as the gantry angle changes. The raw data from a CT scanner forms a sinogram from which the anatomical information can be teased out using advanced reconstruction algorithms. The final result of a CT scan is a three dimensional array of data where the array elements represent the mean electron density distribution throughout the imaged volume. This data is used in treatment planning to optimize the delivery of radiation according to the patient's specific anatomy.

2.2.1 Computed Tomography - X-Ray Production

CT imaging beams are generated in x-ray tubes wherein electrons, thermionically emitted from a cathode, are accelerated towards a high Z anode by a potential difference (Figure 2.5). This potential difference is often in the 40 to 150 kV range for diagnostic imaging applications. Since the definition of an electronvolt is the amount of kinetic energy gained by an electron accelerated by a 1 V potential difference, the electrons in x-ray tubes will have anywhere from 40 keV to 150 keV of kinetic energy

prior to striking the anode. The tube current is proportional to the number of electrons per second accelerated from the cathode to the anode ($1 \text{ mA} = 6.24 \times 10^{15} e^-/\text{s}$), and with greater tube current comes greater x-ray output and anode heat.




Figure 2.5: Basic x-ray tube schematic [4].

The multitude of charged particles bombarding the anode create a continuous energy distribution of Bremsstrahlung rays, along with peaks in the distribution caused by characteristic x-rays (Figure 2.6). If there is no filtration, the Bremsstrahlung portion of the spectrum follows a relationship known as Kramer’s equation [4]:

$$I_E = KZ(qV - E), \quad (2.3)$$

where I_E is the intensity of photons with energy E , Z is the anode atomic number, qV is the electron kinetic energy, and K is a constant. The intensity of photons drops to zero when $E = qV$ because the Bremsstrahlung photon energy cannot exceed the electron kinetic energy, and the probability of creating high energy Bremsstrahlung photons is smaller than creating low energy photons. The reality, however, is that there is inherent filtration in real-world x-ray tubes as there are several mechanisms for absorption of photons in the tube. While this reduces the efficiency of the x-ray tube, in clinical applications this is beneficial as the low energy x-rays do not have the penetrating power to be useful for CT imaging and only contribute to absorbed tissue dose. Depending on the region being imaged, oftentimes additional filtering is added to *harden* the beam and attenuate more low-energy photons, thus increasing the mean energy of the spectrum.




Figure 2.6: Distribution of x-rays for different x-ray tube voltages [4].

2.2.2 Computed Tomography - X-ray Detection

X-ray detection in CT is done with either a gas-filled detector or a solid state detector (SSD); however gas filled detectors are antiquated and are rarely used clinically. The SSDs in modern CT scanners consist of a scintillator attached to a photodiode. Scintillators are materials such as CdWO_4 , $\text{Gd}_2\text{O}_2\text{S}$, HiLightTM, or GemstoneTM [12], that produce visible light from interactions with ionizing radiation. The attached photodiode semiconductor converts the visible light into an electric current proportional to the light intensity. SSDs have a high detection efficiency and are physically smaller in size. As a result, their use can reduce imaging dose and increase spatial resolution [13].

SSDs are made more efficient by *scintering* the scintillator crystals. This process consists of heating the scintillation crystals to just below their melting point for a long period of time (hours). The result is a high density ceramic with greater light output. The ceramic is then carved with a saw or laser to produce a large number of detector elements (e.g. 64×64). The gaps in the detector elements are packed with an opaque filler to reduce cross talk between elements (Figure 2.7a). The combined scintillator-photodiode is connected to a larger stack of electronics which provides power and receives the photodiode's electronic signal. The electronics stack contains gain channels for each detector and an analog-to-digital converter, having these components close to the SSD has the added benefit of a reduction in electronic noise [14]. Finally, several SSD-electronics units are combined together to form the full array of the CT detector (Figure 2.7b).

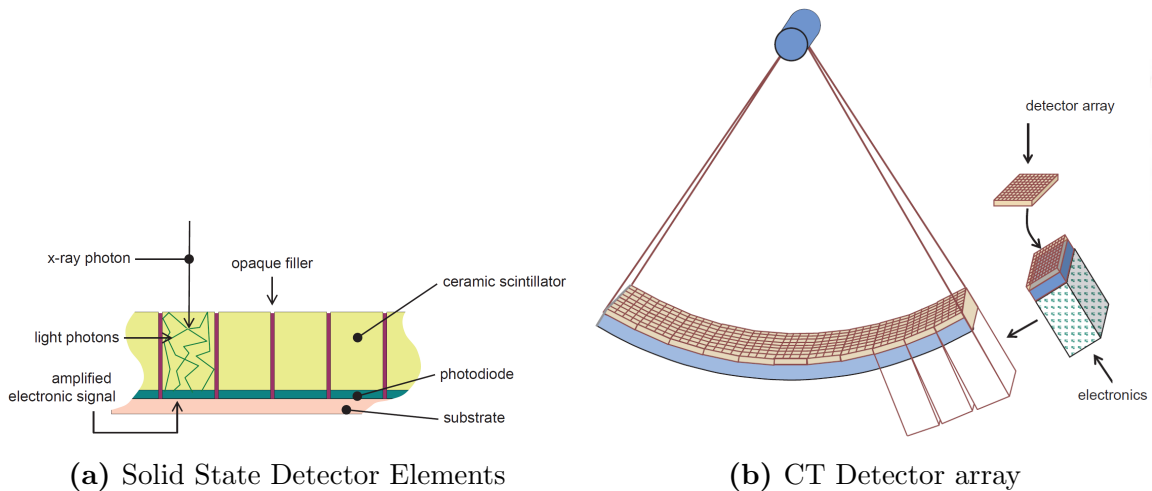


Figure 2.7: CT Detector Components [11].

2.2.3 Computed Tomography - Image Reconstruction

This section examines the basics of reconstructing the anatomical information from raw projection data. To begin, it is helpful to linearize the problem at hand. In order to correct for the exponential attenuation of x-ray interactions, the projection data undergoes a logarithmic transformation and normalization. At each projection angle, the intensity I reaching a single detector element through a non-uniform object is:

$$I = I_0 e^{-(\mu_1 \Delta x_1 + \mu_2 \Delta x_2 + \mu_3 \Delta x_3 + \dots + \mu_N \Delta x_N)}. \quad (2.4)$$

Where I_0 is the entrance x-ray intensity, and Δx_i is the thickness of material with linear attenuation coefficient μ_i . Dividing both sides of Equation 2.4 by I_0 , and taking the negative logarithm we obtain the linear equation:

$$P = -\ln \left(\frac{I}{I_0} \right) = \sum_{n=1}^N \mu_n \Delta x. \quad (2.5)$$

Where P is the projection measurement. Allowing Δx to approach zero gives a line integral of attenuation coefficients along the x-ray path to the detector.

$$P = -\ln \left(\frac{I}{I_0} \right) = \int_L \mu(x) dx. \quad (2.6)$$

Thus, the ratio of input and output intensity is related to the attenuation coefficients encountered by the x-ray along its path [12]. The fundamental problem of CT reconstruction is to determine the attenuation coefficient distribution given the measured line integrals from many different angles (Figure 2.8).




Figure 2.8: Measured projections in CT imaging [15].

There are four primary methods for calculating image slices given a set of projections. The first method is based on solving a set of many simultaneous linear equations. If there are $N \times N$ pixels in the image, the pixel values can be calculated by taking at least N^2 projections. Often CT scanners acquire at least 50% more projections making the problem overdetermined. This reduces image noise and artifacts and also accounts for the possibility of errors in measurement. With so many simultaneous linear equations (almost 400,000 for just one 512×512 image), this method is often impractical as it requires vast amounts of computational resources [12].

The second method is *simple backprojection* (Figure 2.9a) where each measured projection is ‘smeared’ back across the image. More specifically, each pixel along the projection direction is set to the corresponding projection value, then the pixel values are divided by the number of pixels the projection was smeared across. This gives each pixel’s average attenuation coefficient value. An image can be constructed by summing a large number of smears from many different angles [16]. This method is associated with a characteristic $1/r$ blurring because, in the fourier domain of these images, low frequencies (associated with smooth surfaces) are sampled much more densely than high frequencies (associated with details and sharp edges).

Filtered backprojection corrects for the blurring in simple backprojection by applying a filter kernel to each projection prior to the backprojection process. This is commonly done by convolving the projections with a ramp filter in the spatial domain, which suppresses the characteristic blurring and enhances edges between tissues

of different attenuation coefficients. Different filters can be chosen depending on the imaging site, and some can be used to increase spatial resolution at the expense of image noise [15].

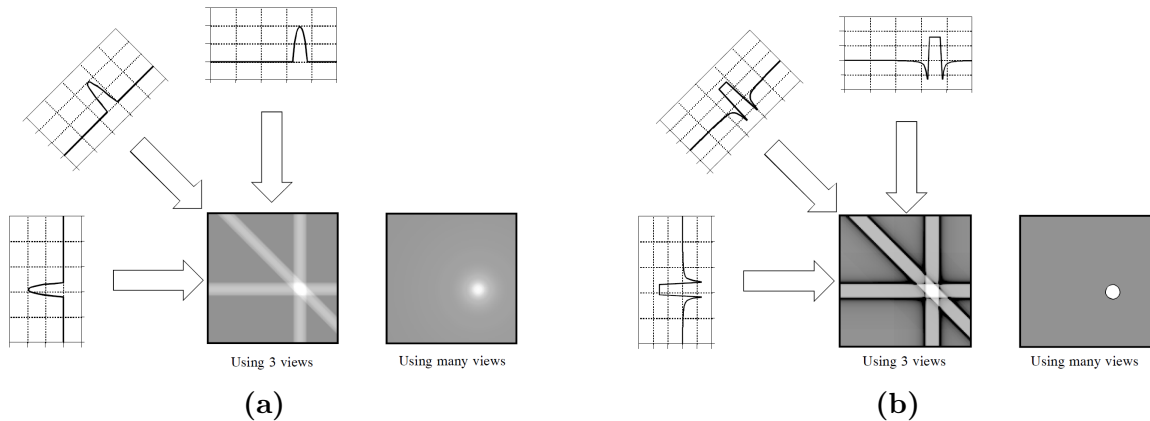


Figure 2.9: Backprojection in CT Imaging. **(a)** Simple backprojection with characteristic $1/r$ blurring [15]. **(b)** Filtered backprojection. Filters are applied to projections prior to backprojection [15].

Fourier reconstruction methods exploit the Fourier slice theorem. The theorem states that the Fourier transform of a parallel projection of an object $f(x, y)$ obtained at angle θ equals a line in a 2D Fourier transform of $f(x, y)$ taken at the same angle. That is, if we take the Fourier transform of a projection at angle θ , this is equal to a slice along the same angle of the object in the frequency domain (Figure 2.10). In an imaging pipeline, as each projection is acquired, the 1D fast Fourier transform of each projection is computed and placed as a slice in the Fourier domain. Since the sampling density along each slice is uniform, the sample distribution of the image in the Fourier domain is arranged radially leading to overall denser sampling closer to the origin. This is the reason for simple backprojection's $1/r$ blurring. Next, the radial sampling distribution is interpolated to obtain a rectangular grid distribution. Finally, the spatial domain image is reconstructed by taking the inverse Fourier transform of the frequency domain image [15].

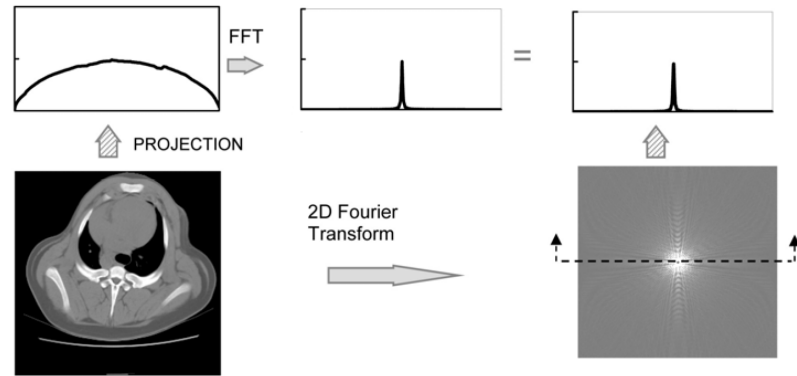


Figure 2.10: Fourier reconstruction [12]

The final reconstruction method, *iterative reconstruction*, begins with a seed image based on a generic model or the filtered backprojection image. The forward projections of the seed image are directly compared to the measured projections and an error matrix is computed using the differences in projections. The error matrix is used to update the seed image so that the error matrix of the next iteration is reduced. This computationally intensive process is repeated over multiple iterations to obtain a final, high quality image (Figure 2.11). Until recently, filtered backprojection was the primary method of image reconstruction in CT imaging, however advances in computation power have allowed for wider use of iterative reconstruction. Iterative algorithms will likely become the preferred mode of reconstruction as they can improve signal-to-noise ratio (SNR) for the same imaging dose, or reduce dose for the same SNR compared to filtered backprojection methods [11].

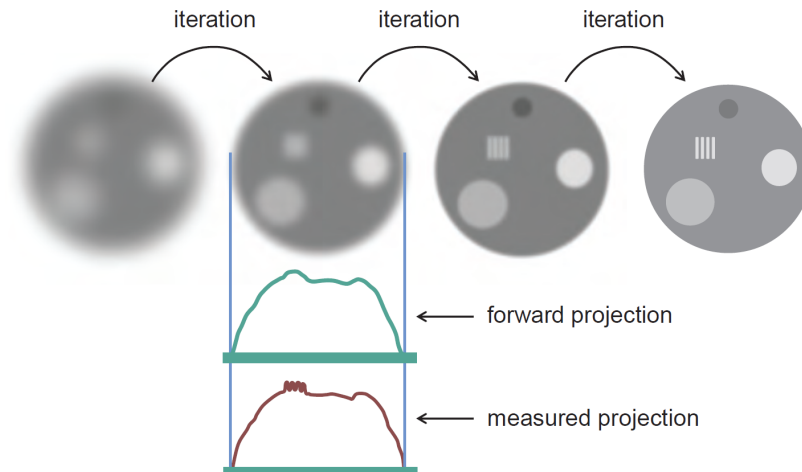


Figure 2.11: Iterative reconstruction [11]

A post-processing step common to all reconstruction methods is the conversion from attenuation coefficient pixel values to Hounsfield Units (HUs).

$$HU_{(x,y,z)} = 1000 \frac{\mu_{(x,y,z)} - \mu_{water}}{\mu_{water}}. \quad (2.7)$$

Where $\mu_{(x,y,z)}$ is the average attenuation coefficient for a given voxel. By Equation 2.7, when a voxel contains primarily water, its HU value will be near zero; when the voxel contains air the attenuation coefficient is zero and therefore the HU value is -1000 . X-ray tube outputs are often calibrated according to these values. The conversion to HU values gives a single channel image which gives CT images their characteristic greyscale appearance [11].

2.2.4 Computed Tomography - Beam Geometry

The reconstruction methods outlined in the previous section are valid for parallel beam geometries (Figure 2.12a). Other beam geometries such as fan-beam and cone-beam are used more frequently in clinical settings, which both come with distinct advantages and disadvantages.




Figure 2.12: Computed Tomography beam geometries: (a) parallel beam geometry, (b) fan beam geometry. (c) cone beam geometry with a curved detector [12]

Fan beam geometries (Figure 2.12b) are most commonly used in dedicated CT scanners. In this geometry the x-rays originate from a single source and diverge on a plane as they radiate outward. Reconstruction of fan beam geometries is done by either accounting for the beam divergence in the smearing of the backprojection, or by binning sinogram projections so that the fan beam geometry can be treated as a parallel beam geometry. Detector arrays are commonly arranged in an arc so that the source-to-detector distance is a constant radius, thus minimizing the difference in photon fluence between each detector. Furthermore, the primary x-rays strike the

detector at normal incidents as this reduces lateral positioning errors due to x-ray beam parallax. In this geometry the patient is moved through the gantry as multiple image slices are acquired so that a volume of the patient can be imaged [11].

The *cone beam* geometry stacks multiple fan beams in the z direction to allow for the acquisition of multiple slices simultaneously (Figure 2.12, 2.13). Cone beam geometries are often used in conjunction with flat panel detectors in image-guided radiotherapy (IGRT) as it can image a large volume in a single arc, thus allowing for the verification of patient set up prior to irradiation [11]. The Feldkamp-Davis-Kress algorithm is used for cone beam image reconstruction as it accounts for beam divergence in the z direction as well as using the source-to-isocenter and source-to-detector distances to fix image distortions [17].

While cone-beam CTs (CBCTs) are able to achieve a higher spatial resolution (likely due to higher density of detector elements) [18], the spatial resolution is highly dependant on imaging protocol and is therefore not a consistent advantage for CBCT [19]. On average, CBCTs also deliver higher imaging doses than fan beam CTs [20]. This is of particular concern in IGRT where the patient is imaged before every treatment and potentially accrues more dose than prescribed if not accounted for. Efforts to reduce imaging dose often involves a reduction in x-ray tube mAs, which, combined with the fact that CBCTs are intrinsically more prone to scattering, beam hardening, and artifacts, the resulting images are often poor quality compared to fan beam CT [18]. However, the ability to acquire a 3D image in a single arc makes CBCT a strong imaging modality for IGRT as gantry rotation speed is typically limited for safety.




Figure 2.13: Cone Beam CT using a flat panel detector [11], this beam geometry is used in linac-mounted CBCT imagers.

2.3 Medical Linear Accelerators

A medical linear accelerator (linac) is a device that can accelerate electrons to high energies. It uses either travelling or standing microwave-frequency electromagnetic fields to accelerate the electrons. The high energy electrons are directed towards a high Z target material, such as tungsten, where a beam of photons is created through bremsstrahlung interactions. The resulting photon beam has a spectrum of energies, with the maximum energy equal to the incident electron kinetic energy, and an average energy of approximately one third of the electron kinetic energy. In some cases the target is removed and replaced with a scattering foil allowing the electron beam to be used to treat superficial tumors since electrons do not penetrate as deeply into tissue as photons.

The treatment head of a Varian TrueBeam linear accelerator (Figure 2.14) consists of several parts that monitor, shape, and modulate the photon beam. After the electron beam strikes the target, a portion of the subsequent photon beam is attenuated by fixed position primary collimators for principal beam shaping. The conic flattening filter ensures that the photon fluence is spatially uniform at a certain depth, conventionally 10 cm. Alternatively, a thin filter made of either brass, stainless steel, or aluminium can be used in ‘flattening filter free’ (FFF) treatments that allows the dose rate to be 2-4 times greater [21]. This is useful in stereotactic radiosurgery treatments since the dose can be delivered in a shorter period of time, thus preventing treatments from becoming overly long and decreasing the likelihood of patient movement due to discomfort. After the flattening filter comes the monitor chamber: a set of ionization chambers that are calibrated to monitor dose-rate, integrated dose, and field symmetry. The beam is then further collimated by a set of jaws that shape the radiation into a rectangular field up to $40 \times 40 \text{ cm}^2$ in size at the position of the isocenter. High-resolution shaping of the beam is provided by the multi-leaf collimator (MLC) system. The MLC system consists of a large number of individually motor-driven tungsten leaves so that near-arbitrary field shapes can be achieved. The field shapes that can be achieved are limited by the width of the MLC leaves which can be as thin as 2.5 mm in modern high definition MLC systems.

The treatment head is mounted to a gantry which rotates around a point in space known as the isocenter. Typically, the isocenter is positioned within the patient's target volume so that the radiation beam is always directed towards the target. The combination of MLC high-resolution field shaping and gantry rotation allows for the delivery of highly conformal treatment doses.




Figure 2.14: Core components of a linear accelerator treatment head [22]

2.4 Treatment Planning in Radiation Therapy

The process of treatment planning involves determining the the most appropriate way to irradiate the patient given their unique pathology and anatomy. In general, treatment planning follows five crucial steps:

1. Selection of patient positioning and immobilization methods for reproducible treatments.
2. Imaging of treatment site with CT or MRI.
3. Delineation of the tumor site and nearby organs-at-risk (OARs).
4. Treatment plan optimization.
5. Evaluation of the resulting dose distribution.

Data of the patient's electron density distribution is required for dose calculations so patients are imaged using CT prior to treatment planning. The CT planning images are used to create target volume structures, as well as to contour healthy tissue avoidance structures. However, sometimes MRIs or positron emission tomography (PET) CT scans are also taken to aid in contouring procedures.

Optimization and plan evaluation often involve the use of DVH curves which depict the relationship between target/OAR volume and dose. They summarize 3D dose distributions into a 2D graph, as such they lack spatial information on the position of high or low dose regions. In differential DVH curves, dose values are divided into bins and the number of dose voxels in each bin range is counted and then multiplied by the voxel size. Cumulative DVH curves represent the volume receiving a certain dose or higher plotted as a function of dose. Cumulative DVH curves allow one to assess target coverage with metrics such as 95% of volume receiving at least 40 Gy, or OAR metrics such as restricting the volume that gets more than 36 Gy to less than 1 cm³.

2.4.1 Irradiation Target Definitions

The International Commission on Radiation Units and Measurements (ICRU) Report No. 50 [23] standardizes the volume definitions used in radiation therapy (Figure 2.15). The ICRU volume definitions are as follows:

Gross Tumor Volume (GTV) is the gross demonstrable location and extent of tumour. This is the tumor volume that can be seen or imaged.

Clinical Target Volume (CTV) encompasses the GTV, but is extended to include any microscopic extension of the primary tumor. While the CTV cannot be directly imaged, it is important that the entire volume be treated to prevent cancer recurrence.

Planning Target Volume (PTV) contains both the CTV and GTV with a margin added to account for organ motion and patient setup uncertainties.

Treated Volume (TV) is the volume enclosed by a selected isodose surface (e.g. 90%)

Irradiated Volume (IV) is the volume receiving a radiation dose that is considered significant relative to its normal tissue tolerance.

Organs at Risk (OARs) are organs near the PTV that do not contain malignant cells and are not the intended target for radiation, yet they often lie within the irradiated volume.

Previously, structures such as the CTV, PTV, and OARs were manually contoured on each slice of the imaging volume. This is a time consuming and repetitive task for highly trained staff considering scanners can often acquire over 100 image slices. As a time saving measure, many centers have incorporated AI-based auto-contouring software into their treatment planning workflow where contours simply need to be reviewed and, if necessary, edited.

The end goal of treatment planning in radiation therapy is to devise a plan that minimizes dose to OARs whilst maintaining sufficient dose coverage to the PTV in order to achieve tumor control. This poses a significant challenge in treatment sites where the size, shape, and position of OARs can vary significantly between fractions.




Figure 2.15: ICRU Report No. 50 Planning Volumes [24].

2.4.2 Dose-Volume constraints

Goals in treatment planning are often outlined by a set of dose-volume metrics for the CTV, PTV and OARs. For OARs, dose-volume metrics indicate the maximum volume V_{max} that can receive a certain dose D in either absolute (cm^3 for volume, or Gy for dose) or relative terms (% of total volume, or % of prescribed/maximum dose) (Figure 2.16a). OARs often have a set of several dose-volume constraints that are related to toxicities, or to allow for finer control of optimization during treatment planning. For target volumes, a combination of maximum and minimum dose-volume metrics are used so that coverage is near uniform throughout the volume. The full target volume should receive at least dose D_{min} , but not exceed D_{max} (Figure 2.16b).

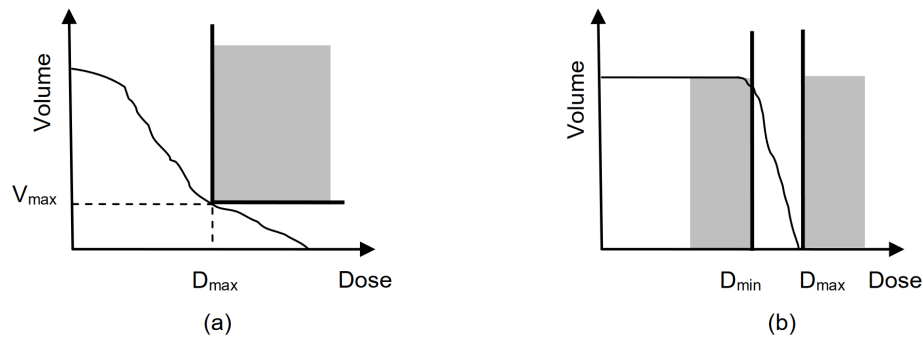


Figure 2.16: Dose-volume constraints for (a) OARs, and (b) target volumes [25].

2.4.3 External Beam Planning

Historically, external beam radiation therapy (EBRT) was forward planned, where the treatment planner manually chose parameters in a trial and error process. Modern treatment planning workflows incorporate inverse planning, where a number of dosimetric requirements are stipulated and an algorithm in the treatment planning system (TPS) searches for a set of optimal beam parameters to best meet those requirements. Beam parameters may include fluence, energy, number and direction of beams, collimator settings, and MLC leaf positions. A set of dose-volume constraints for target structures and OARs constitute what an ideal dose distribution should adhere to. The dose-volume metrics are associated with weights and are incorporated into a cost function which quantifies the difference between the desired and actual dose distributions. The weights indicate the relative importance of each planning objective. They can be changed by the treatment planner as in some cases it may be more important to maintain PTV dose coverage than to meet a minor OAR dose-volume metric [26]. While there are multiple cost functions that may be used in treatment planning, one of the most common involves minimizing the mean square difference between the actual dose distribution and the treatment planning goals given by the dose-volume constraints. In its most general form, this cost function is written as:

$$f = \sum_i f_{OAR_i} + \sum_j f_{target_j}$$

where f is the total objective function given by the sum of the OAR objective functions, f_{OAR} , and the sum of the target objective functions, f_{target} .

2.4.4 Volumetric Modulated Arc Therapy

While there are a great number of ways to deliver external beam radiation, volumetric modulated arc therapy (VMAT) [27] is the primary method of delivery for prostate SABR treatments at BC Cancer Victoria. In VMAT treatments, the linac gantry rotates continuously in an arc of up to 360° while the MLCs shape the radiation beam. VMAT treatments are able to deliver more conformal treatments as the radiation is delivered from a continuum of gantry angles.

At the start of VMAT optimization, the continuous linac rotation and MLC positions are approximated with a small number of static beams. As optimization progresses, new intermediary beam angles are added. The optimizer determines the MLC leaf positions and dose rate as a function of gantry angle, aiming to adhere to the required target volume coverage and OAR sparing. The optimization process also incorporates objective functions to balance the trade-offs between PTV coverage, dose homogeneity, and OAR sparing. Commonly used optimization algorithms in treatment planning systems include gradient descent, simulated annealing, and genetic algorithms [28]. To reduce the number of degrees of freedom, the couch position and the collimator angle are often fixed. Constraints are also placed on gantry rotation speed, leaf motion speed, and dose rate are considered to ensure the plan is deliverable, safe, and adheres to the hardware limitations of the linear accelerator. The gantry rotation speed is typically 6 degrees/s (1 RPM), with MLC leaf motion limited to approximately 2-3 cm/s. Dose rates for 6FFF and 10FFF are capped at 1400 MU/min and 2400 MU/min respectively.

The progressive increase in beam angle sampling density converges on an accurate representation of the continuous gantry rotation and MLC motion. This iterative optimization results in a VMAT plan that delivers the prescribed dose to the tumor while minimizing the dose to surrounding healthy tissues and OARs. These treatments are fast to deliver as the entire radiation dose is given in one to two arcs.

2.4.5 Dose Calculation

The gold standard for modelling dose distributions are Monte-Carlo (MC) transport simulations. In MC simulations, a large number of photons are modelled according to probability distributions which is very computationally intensive. MC simulations can be made arbitrarily accurate by increasing the number of photons simulated at the cost of even greater computation times. Many other photon dose-

deposition algorithms using different methods of computation exist, such as the Analytical Anisotropic Algorithm (AAA), which is used in the Varian Eclipse TPS, or more recently, Acuros XB (AXB).

AAA is a 3D pencil beam convolution-superposition algorithm that models primary photons, scattered extra-focal photons, and electron scatter from beam-attenuating devices separately. Prior to dose calculation, the AAA algorithm is configured using Monte-Carlo principles to pre-compute the physical parameters that characterize the fluence, energy spectra, and scattering properties of photons and electrons in water equivalent media. These parameters are fitted to treatment-unit-specific beams to ensure subsequent dose calculations are fast and accurate [29].

In dose calculation, AAA uses individual convolution kernels for the modelling of primary photons, scattered photons, and scattered electrons. The clinical beam is divided up into small beamlets with which the kernels are convolved to obtain the individual dose contributions from each model. The final dose distribution is simply the sum of the contributions from primary photons, scattered photons, and scattered electrons [29].

Acuros XB (AXB) is another dose calculation algorithm that uses a grid-based Boltzmann transport equation solver [30]. AXB has been shown to be more accurate when the PTV or OAR involves bone or air [31]. A study comparing AXB to AAA in 50 Gy/25 fraction nasopharyngeal carcinoma treatments [32] showed that AAA tends to overestimate PTV median dose by 0.9 Gy and GTV mean dose by 1.0 Gy. Plans calculated with AXB gave mean doses to lung and heart were lower by 1.7% and 2.4% respectively. We deemed AAA sufficient for this project as prostate EBRT treatments do not generally involve any air-soft tissue interfaces.

2.5 Image-Guided Radiation Therapy

In IGRT, imaging devices are mounted to the gantry of linear accelerators. These images can be used to verify patient setup through comparisons of treatment images to planning images and corrections can then be made to patient setup prior to irradiation. Reproducibility of patient setup can also be evaluated by observing images of the patient's setup over multiple treatments. If necessary, images from previously delivered fractions can also be used to aid replan decision making [26]. Varian TrueBeam linear accelerators are equipped with both kilovoltage (kV) x-ray imagers as well as megavoltage (MV) electronic portal imaging devices for either 2D radiogra-

phy or CBCT. Use of these options depends on the application at hand. kVCBCTs are able to achieve better soft tissue contrast and spatial resolution at lower doses compared to MVCBCT. However, MVCBCTs are less susceptible to metallic artifacts and do not need to extrapolate measured attenuation coefficient values from kV to MV for dosimetric corrections [4].

2.5.1 Prostate IGRT

A method for improving prostate treatment accuracy involves implanting three fiducial markers directly into a patient's prostate prior to treatment planning. At least three non-aligned fiducial markers are required for 3D registration to the planned position. At treatment, the position of the fiducial markers are compared to their position on the pCT using two orthogonal 2D kV images. This comparison is made using a digitally reconstructed radiograph (DRR) of the pCT, which is a 2D image created from projecting 3D imaging data onto a 2D plane. The use of implanted fiducial markers for patient setup has been shown to improve treatment accuracy and allow for reduced PTV margins by 2–4 mm [33, 34]. Comparisons between non-IGRT and IGRT with fiducials have shown that IGRT can reduce the 3-year likelihood of grade 2 or higher urinary toxicity from 20.0% to 10.4% ($p = 0.02$) [35]. In addition, prostate-specific antigen relapse-free survival outcome can improve from 77.7% to 97% after 3 years in high-risk patients with fiducial-based IGRT [35].

While marker matching with kV images reliably ensures good positioning of the PTV, further refinement of patient setup is possible by taking a kVCBCT afterwards. OARs such as rectum and bladder pose a significant inter-fraction variability problem and do not appear on the kV marker match. By taking a kVCBCT of the patient between a marker match and the delivery of radiation, small corrections can be made according to the position of the OARs on the day of the treatment [36].

2.6 Adaptive Radiation Therapy

The success of radiation therapy treatments are highly dependent on the accurate delivery of radiation dose. In treatment sites such as prostate, optimal treatment is limited by variations in the delivered dose caused by inter-fraction or intra-fraction variation in organ position, size, and shape, as well as patient positioning. To account for these, one can increase the margin around the target volume to maintain target coverage, but the consequence is an increase of normal tissue toxicity. Adaptive

radiation therapy (ART) is a concept where treatment gets further optimized when given new information over the course of a patient's radiation therapy treatment. In its most basic form, ART may involve re-planning the patient part way through treatment to account for situations such as target volume shrinkage or patient weight loss. This method of ART does not require any extra technology, however the increase in staff workload scales with the frequency of re-planning. As ART treatments move from replanning the remainder of a patient's course to dynamically replanning each fraction, ART solutions will rely more heavily on automation [37].

As technology develops, ART treatments can become increasingly sophisticated. In-room full field volumetric imaging and auto-contouring allow for calculation of daily dose thus facilitating decisions based on dose rather than geometry. This uses calibration curves to perform dose calculations directly on the CBCT images [38]. Advances in deformable image registration allow for the calculation of accumulated dose over the course of treatment on the original pCT scan [39]. This has the benefit of being able to link accumulated dose to outcomes data for the refinement of tumor control probability models and normal tissue complication, further aiding in clinical decision making in ART [40].

The most advanced ART treatments include online replanning, where the treatment plan is tuned to the anatomy of the patient on the day of treatment. This involves imaging the patient, contouring relevant structures, and generating a new treatment plan, all while the patient is waiting on the treatment couch [41]. Ahunbay et al. [42] found that this could increase minimum PTV dose by 13% and reduce equivalent uniform dose to the rectum by 13% in prostate treatments using a Varian Ethos system. While the outcomes for online ART are promising, the contouring and replanning systems rely heavily on automation. Clinical staff are required to be present for each treatment to review contours and treatment plans. Auto-segmentation is an active field where deep learning is used to create structure contours from image data. Auto-contouring has the potential to offer considerable time savings while generating quality contours, contributing to the automation of workflows required by adaptive radiotherapy.

2.7 Auto-contouring

2.7.1 Previous Approaches

There have been many historical approaches to CT biomedical image segmentation. Most, like modern deep-learning-based methods, intended to alleviate clinical workload. Image thresholding involves setting intensity-based thresholds in order to segment structures based on their pixel values. These methods are some of the simplest to implement however they struggle to segment structures that overlap, are touching, or have metal artifacts [43]. Edge detection methods use filters to detect boundaries between different structures. These methods are good at detecting boundaries between structures with differing intensity values, providing sharp and precise segmentation masks. Although, these methods often require additional post-processing in order to form a complete, continuous anatomical contour [44]. Atlas-based segmentation methods involve using a pre-defined ‘atlas’ of anatomical structures to guide the segmentation process. These methods allow one to incorporate prior anatomical knowledge into the contouring process and are useful when segmenting structures that have consistent shapes and sizes across images. However, the opposite is also true, they often give poor results when the structure size and shape varies [45]. These are just a handful of the plethora of methods that have been used for auto-contouring of anatomical images. Many of them have notable downsides and are therefore being eclipsed by modern deep learning-based methods.

2.7.2 Machine Learning

In general, a machine learning algorithm is a computer program that uses input data to generate a desired output, without explicitly coding for the particular outcome. All machine learning models require some amount of training and validation. During model training, a large amount of input data is combined with expected model outputs, and the model learns how to reproduce those desired outcomes. Ideally, the trained model is generalized, that is, it is neither over nor under trained and can make reasonable predictions from novel input data. The impact that machine learning has had on technology cannot be overstated. Its use has led to advances in all areas where data can be collected, from maximizing social media engagement to spacecraft engineering.

2.7.3 Neural Networks

Artificial neural networks (Figure 2.17) are inspired by the biological neural networks found in animal brains. They consist of a collection of nodes that act much like neurons, and the nodes are connected together similar to synapses. Data is transmitted from one node to another via a connection wherein some mathematical operation is performed. Each pair of two neurons and their connection has an associated weight. The value of the weights is iteratively increased or decreased to reflect the strength of the node connections as learning proceeds. Input nodes are connected to one or more layers of hidden nodes which are subsequently connected to the output layer.

Learning in neural networks involves supplying the model with an adequate amount of examples, each with an input and desired output. Input data is repeatedly processed through the network and the model output is compared to the ground truth; as learning progresses an error function is used to adjust the weights of a network's connections, strengthening node connections that result in desired outputs, and penalizing connections that produce unwanted results. After many iterations, the network is ideally able to produce acceptable outputs using input data it has never encountered before.




Figure 2.17: A simple neural network. Line thickness indicates strength of weights between nodes.

2.7.4 Convolutional Neural Networks

Convolutional neural networks (CNNs) are a type of neural network that have been successful in many computer vision tasks. A core feature of CNNs is the inclusion of convolutional layers, where dot products of a convolution kernel are systematically performed with many subsections layer's input. The kernel values at each layer are one of the learnable parameters in CNNs. Pooling layers (Figure 2.18) often follow convolutional layers to non-linearly downsample the data and reduce its dimensionality while preserving the relative location of image features. Pooling layers are a cost saving measure to reduce computation time with minimal loss of information. CNNs may also use fully connected layers similar to those shown in Figure 2.17.




Figure 2.18: Example of a max pooling operation.

2.7.5 U-Net

U-Net is a CNN architecture developed by Ronnenberger et. al. [46] that excels at image segmentation. It takes an image as an input and creates a segmentation mask, classifying each pixel according to the training data given (Figure 2.19). The advantage of this is that the network output has the same number and position of pixels as the input image. As a CNN architecture with kernel values as a learnable parameter, it is able to precisely localize and distinguish borders of features in the input image. Additionally, data-augmentation allows the network to be trained off of relatively few training images, which is an advantage in a biomedical imaging scenario where the creation of training datasets can be a time consuming task for highly trained individuals.

In many machine learning problems, it is helpful to perform several data augmentation operations prior to model training. This helps improve model performance and its ability to generalize, making greater use of limited datasets. In biomed-




Figure 2.19: Example of U-Net image segmentation trained on a lung X-Ray dataset. The left image is the original mask created by an expert using the center image. The right image is the output of a trained U-Net model that has never seen the center image before.

cal image segmentation, there are several types of data augmentation that can help yield better results when using a U-Net CNN. Rotations of images by random angles help models recognize anatomical structures from different orientations. Scaling or elastic deformations of images help models recognize structures that can change in size and shape. Translations can improve the model’s ability to handle differences in patient positioning. Flipping images vertically or horizontally shows the model how to recognize features that have left-right or up-down symmetry. Random cropping can help models focus on local regions of interest. Modifying pixel values with noise (gaussian, salt-and-pepper) or intensity-based transformations (contrast stretching, gamma correction, or histogram equalization) can make models more robust in the face of image noise and artifacts, or brightness and contrast variations in the input data. All of these methods artificially increase the size of the dataset, maximizing the models ability to generalize from size-limited datasets.

The original U-Net architecture consists of several pooling layers followed by an equivalent number of up-convolution layers. Portions of the network prior to pooling layers are cropped and concatenated to layers after the up-convolutions. These features are what gives the architecture its characteristic ‘U’ shape (Figure 2.20). Broadly, the architecture can be thought of as an encoder followed by a decoder. The contracting path (encoder) follows the pattern of two 3×3 un-padded convolutions, each followed by a rectified linear unit (ReLU) activation function, then a stride 2 max pooling operation for downsampling. The max pooling operation sweeps over the input feature map and takes the maximum value of many 2×2 sections of the

input. At each downsampling, the number of feature channels is also doubled. For example, a single channel CT image would get two channels, or a three channel RGB image would get six channels. In the initial layers of the encoding path, the network extracts low level features such as edges, corners, and textures. These represent the basic spatial and structural information present in an anatomical image. The following intermediate layers identifies more abstract features such as shapes and patterns which are typically larger structures such as organs or bones. Deeper layers in the network’s encoding path capture higher-level semantic information, such as the relative spatial relationships between larger anatomical structures.

The expanding (decoder) path consists of an up-sampling of the feature map followed by a 2×2 convolution which halves the number of feature channels, this is concatenated with a cropped feature map from the contracting path which is then subjected to two 3×3 convolutions, each followed by a ReLU activation function. This pattern is also repeated four times to match the number of steps in the contracting path. The final layer is a 1×1 convolution which maps the component feature vectors to the desired number of pixel classes. In the expanding path, the high-level semantic information is combined with the lower-level spatial information from the encoding path via the skip connections. Doing so helps the network restore spatial resolution into the final segmented output. In auto-contouring applications, the final layer of the network generates a segmentation map that classifies each pixel of the input image into different anatomical structures

U-Net has superseded many previous approaches to biomedical image segmentation. Methods based on U-Net are robust to variations in data, such as noise, intensity differences, and deformations, making it suitable for processing biomedical images with diverse characteristics. It can also be easily adapted to various image segmentation tasks by training it on site-specific data.

Common metrics to assess the performance of AI-generated contours are the Dice Similarity Coefficient (DSC), and Hausdorff Distance (HD). Dice similarity coefficient is a measure of similarity between two datasets, X and Y . It is defined as:

$$DSC = \frac{2|X \cap Y|}{|X| + |Y|}.$$

The DSC is two times the number of elements that are common to both dataset divided by the sum of the number of elements in each dataset. A DSC value of 1 implies a perfect overlap, and a DSC value of zero means there is no overlap. Hausdorff

distance measures how far two subsets of a metric space are from each other, its value represents the largest distance of a point in one set to the closest point in the other set. Mathematically, it is defined as:

$$d(\mathcal{X}, \mathcal{Y}) = \sup \left\{ \sup_{x \in \mathcal{X}} \inf_{y \in \mathcal{Y}} d(x, y), \sup_{y \in \mathcal{Y}} \inf_{x \in \mathcal{X}} d(x, y) \right\}. \quad (2.8)$$

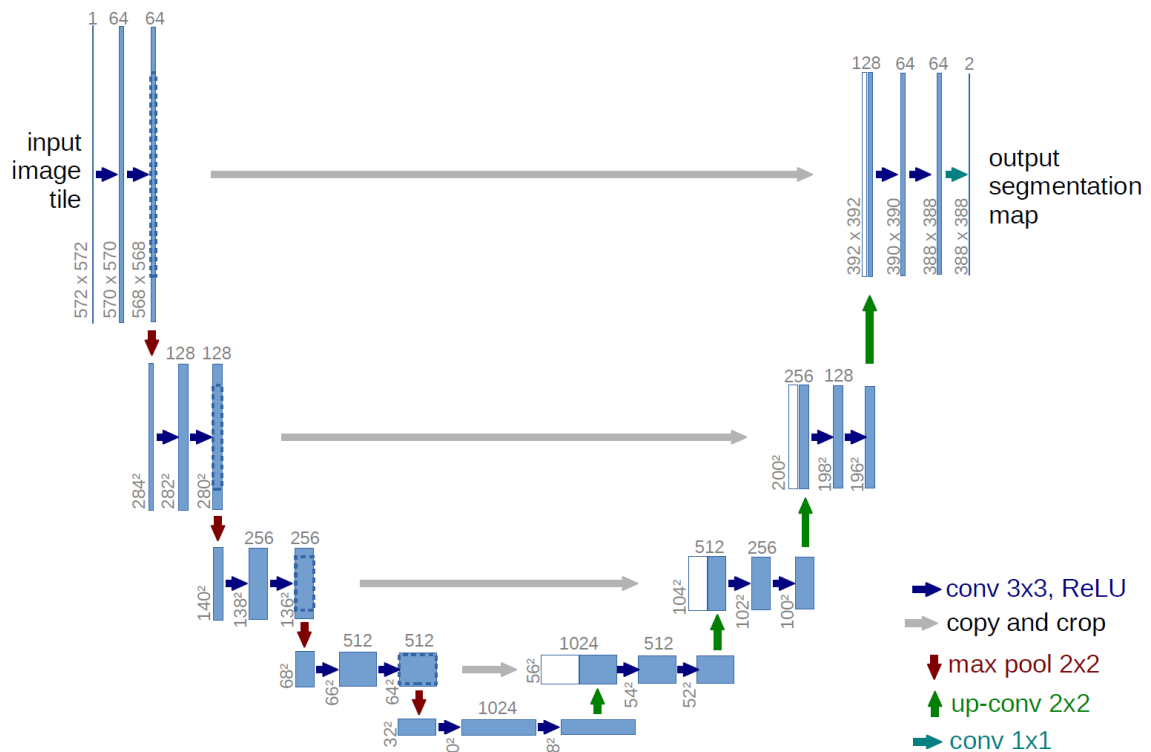


Figure 2.20: U-Net Architecture [46]

2.7.6 Limbus AI Autocontouring

Limbus AI is a company specializing in machine-learning-based software solutions for radiation therapy. Their auto-contouring software currently supports a number of imaged treatment sites such as pelvis on CT or central nervous system (CNS) with MRI. Their prostate CBCT models were developed using an adapted U-Net architecture and trained off of approximately 500 CBCT scans from 250 patients, and 1000 CT scans each from a different patient. The training and validation datasets included

full and narrow field of view CBCTs from both Varian and Elekta and all contour data was corrected by an in-house RO to ensure consistency across the dataset. To balance the dataset between CBCT and CT scans, half of CT scans/contours were randomly selected to use in training in order to avoid training a bias towards CT input data. A distinct model was trained for each OAR over several epochs and approximately 30 scans/contours were used as validation. Models were chosen such that the validation set had the highest average dice similarity coefficient (DSC). Some post-processing is also performed to improve contour accuracy. For example, it is unlikely that disconnected pixels classified as OAR are actually part of the OAR and are therefore removed. Model performance was evaluated by computing DSC and HD between deep-learning-generated contours and expert contours, in addition to expert contours to other expert contours (expert inter-observer variability). Deep-learning based contouring is becoming increasingly widespread as it offers quality contouring in shorter time frames.

Figure 2.21 shows the performance of Limbus prostate CBCT contours to expert inter-observer variability (IOV). To obtain these values, Limbus AI internally conducted a study using CBCT scans from 20 different patients. Expert IOV was obtained by computing DSC and HD values between two ROs, RO_1 and RO_2 , for each structure in the 20 CBCTs and then averaged. RO_1 and RO_2 contours were separately compared to Limbus-generated contours and again averaged to yield a single DSC and HD value for each structure.

In terms of dice similarity coefficient, Limbus scores were 0.94, 0.82, 0.81, and 0.55 for bladder, rectum, prostate, and seminal vesicles respectively, while expert IOV values were similar at 0.96, 0.81, 0.80, and 0.59. Bladder, rectum, prostate, and seminal vesicles Hausdorff distance values were also close at 77.1, 66.6, 48.9, and 43.7 for Limbus, with expert IOV scoring 79.6, 68.5, 49.7, and 50.2. Limbus contouring for each of these structures performed as well as expert IOV, which is considered the gold standard that AI-generated contours are evaluated against.

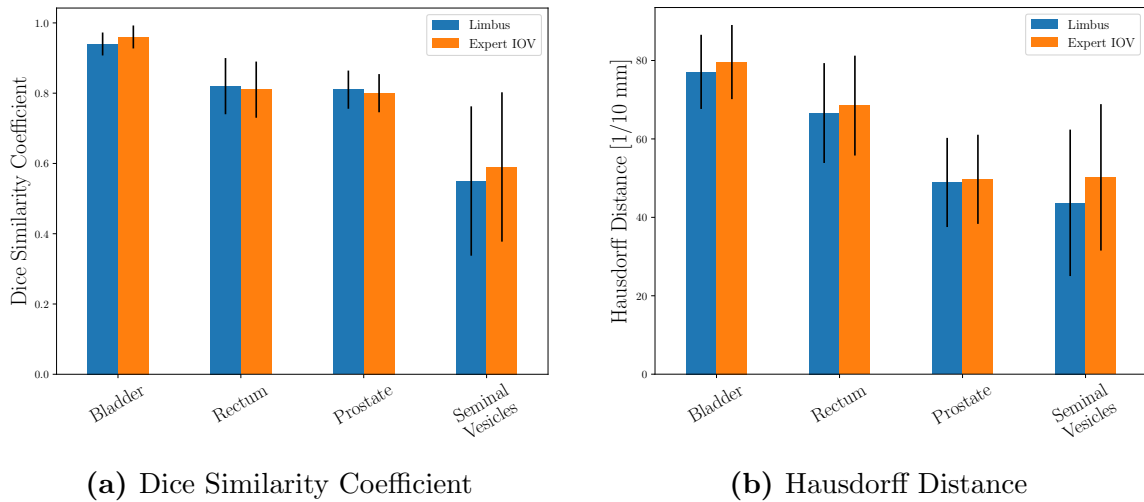


Figure 2.21: Comparison of Limbus deep-learning-generated contours to expert inter-observer-variability. Error bars represent one standard deviation.

2.8 Prostate Stereotactic Ablative Body Radiotherapy

2.8.1 Fractionation

Delivery of radiation dose is commonly divided up into several fractions and given over a period of time. This is done in order to maximize the therapeutic effect on cancerous cells while minimizing damage to the surrounding healthy tissue. Malignant and normal cells typically have differing sensitivity to radiation, healthy cells also have an improved ability to repair themselves between fractions. The linear-quadratic model is a widely used mathematical tool for evaluating the radiobiological effects of various fractionation schedules, helping to optimize treatment plans. It incorporates the α/β ratio, which represents the relative contributions of linear (α) and quadratic (β) components of cell killing, providing insight into the inherent radiosensitivity of different tissue types. The optimal number of fractions varies depending on factors such as tumor size, location, and the radiosensitivity of surrounding tissues, as characterized by their α/β ratios. In some cases, fewer fractions (hypofractionation) can be beneficial, offering a shorter treatment duration and greater patient convenience. Target cells in treatments that benefit from hypofractionation typically have a low α/β ratio. When there is a higher probability of late normal tissue complications, a greater number of fractions (hyperfractionation) may be necessary to reduce side effects and protect healthy tissue. Conversely, these treatments that benefit from hyperfractionation have a high α/β ratio. Ultimately, the fractionation schedule is

designed to strike a balance between maximizing tumor control and maintaining the function of normal tissues [47, 48].

Conventionally fractionated prostate EBRT is delivered in fractions of 1.8–2 Gy, five days per week for over 30 fractions. However, studies have shown that prostate cancer has a low α/β ratio of approximately 1.5–3.0 Gy [49, 50], therefore equivalent tumor control with comparable late side-effects should be achievable by treating with larger doses per fraction over a shorter time period (hypofractionation) [51, 52]. Indeed, ultra-hypofractionated (UHF) prostate EBRT treatments, where fractionated doses exceed 5 Gy, have been shown to be non-inferior to conventionally fractionated EBRT for intermediate-to-high risk prostate cancer regarding failure-free survival [53].

When ultra-hypofractionated EBRT is delivered in five or less fractions it is termed stereotactic ablative body radiotherapy (SABR), or stereotactic body radiation therapy (SBRT). This has the benefit of being both cost effective for clinics, as well as less logistically challenging for patients, especially for those from remote areas. Studies comparing a moderately hypofractionated prostate EBRT schedule to UHF regimes delivered using VMAT have shown no statistically significant differences in toxicity effects and patient quality of life 6 months post-treatment [54]. Furthermore, prostate SABR is also associated with comparable rates of tumor control probability (TCP) to another standard of care, high dose-rate (HDR) brachytherapy [55], where HDR achieves TCP of 96%–98% [56] and SABR TCP is in the range of 90%–95% [57].

2.8.2 SABR dose-volume constraints

The dose-volume metric protocol for 40 Gy/5 fraction prostate SABR at BC Cancer Victoria is outlined in Table 2.1. OARs such as the rectum, sigmoid and bladder have more than one primary goal to be incorporated into treatment planning. Several structures have a secondary goal if the primary goal cannot be met during optimization. The protocol stipulates that no secondary goal should be exceeded. If the secondary goal cannot be met at the lowest possible target coverage, then a radiation oncologist (RO) is brought in to decide whether target coverage or OAR sparing is more important.

Structure	Primary Goal	Secondary Goal
CTV	$V_{40\text{Gy}} \geq 95\%$	$V_{39.2\text{Gy}} \geq 95\%$
PTV	$V_{36.25\text{Gy}} \geq 95\%$	-
Rectum + 5mm	$D_{\text{max}} \leq 105\%$	$D_{\text{max}} \leq 107\%$
Rectum	$V_{36\text{Gy}} \leq 1 \text{ cm}^3$	$V_{36\text{Gy}} \leq 2 \text{ cm}^3$
	$V_{33\text{Gy}} \leq 10\%$	$V_{33\text{Gy}} \leq 12.5\%$
	$V_{29\text{Gy}} \leq 20\%$	$V_{29\text{Gy}} \leq 22.5\%$
	$V_{18\text{Gy}} \leq 40\%$	$V_{18\text{Gy}} \leq 50\%$
Sigmoid	$V_{30\text{Gy}} \leq 1 \text{ cm}^3$	-
	$V_{18\text{Gy}} \leq 5 \text{ cm}^3$	-
Bladder	$V_{37\text{Gy}} \leq 5 \text{ cm}^3$	$V_{37\text{Gy}} \leq 10 \text{ cm}^3$
	$V_{36\text{Gy}} \leq 5\%$	$V_{36\text{Gy}} \leq 10\%$
	$V_{33\text{Gy}} \leq 10\%$	$V_{33\text{Gy}} \leq 20\%$
	$V_{18\text{Gy}} \leq 30\%$	$V_{18\text{Gy}} \leq 45\%$
Penile Bulb	$V_{30\text{Gy}} \leq 50\%$	-
Femoral Heads	$V_{14\text{Gy}} \leq 5\%$	-

Table 2.1: BC Cancer Victoria Prostate SABR planning goals for 40 Gy/5 fraction prescription.

2.8.3 Preparation and positioning challenges

With such high doses per fraction, the plan optimization creates very steep dose gradients to adhere to both target coverage and OAR sparing. Isodose line fall-off from the target can go from 100% of prescription dose to 50% in just 2 cm. As such, patient setup and position in prostate SABR treatments has a significant impact on whether or not PTV coverage is achieved and OAR dose is spared. For example, an anterior-posterior mispositioning of a patient by just 8 mm, while unlikely to occur with modern IGRT, could potentially bring the patient's rectum from touching the 80% isodose surface to being within the 100% isodose volume.

An additional challenge in prostate treatments in general is the degree of inter-fraction OAR variation (Figure 2.22). The bladder and rectum naturally change shape and volume all the time. If unaccounted for, the PTV may have poor dose

coverage thus reducing tumor control probability. Furthermore, portions of OARs could move into a high dose region, increasing toxicity. Currently, this challenge is in part addressed by the use of bowel and bladder preparation protocols to ensure that the shape and volume of the rectum and bladder are as similar as possible to the pCT. Details of BC Cancer Victoria’s OAR preparation protocols will be discussed in Section 3.1.1. While this does significantly reduce inter-fraction OAR motion, and by extension improve setup reproducibility, it does not completely eliminate it, and variation can cause OAR dose to be higher than planned [58]. For this reason, adaptive radiotherapy solutions have the potential for significant OAR dose sparing.




Figure 2.22: Example of inter-fraction variation of the rectum and bladder in prostate radiotherapy. Images are CBCTs from a single patient’s third and fourth fractions.

Chapter 3

Methods

3.1 Prostate SABR

3.1.1 Patient Preparation

Patients in this study were SABR technique prostate cancer patients treated with 40 Gy in five fractions. Each patient had a SpaceOAR Hydrogel injection in addition to three fiducial markers implanted directly into the prostate. The SpaceOAR gel is a polyethylene glycol-based hydrogel, it is injected between the rectum and the prostate, increasing the distance between the two structures. This helps reduce rectal dose, and by extension, late gastrointestinal and genitourinary toxicities. The gel stays in place for approximately 3 months and is absorbed into the body and removed through urine after about 6 months time [59]. Prior to both the pCT and each treatment fraction, patients were instructed to follow a bowel and bladder preparation protocol to reduce the geometric inter-fraction variability of the rectum and bladder. Bowel preparation included two Gas-X tablets the night before, the morning of, and again 1 hour prior to imaging/treatment. Additionally, they had a Fleet saline enema 1-2 hours before imaging/treatment in order to empty the patient's bowels. The purpose of this was to reduce the amount of faeces and gas in the rectum thus reducing its volume and moves the rectum wall away from the PTV which minimizes dose. Preparation of the patient's bladder involved voiding 1 hour prior to imaging/RT then consuming 750 mL of water within 15 minutes and holding until after the procedure. If the patient was unable to consume and withhold this amount then either gentle (750 mL in 30 minutes) or ultra-gentle (500 mL in 30 minutes) protocols were used. The motivation for filling the bladder is to increase its volume and stretch the bladder wall out so that less bladder wall is exposed to higher doses of radiation. This also has the added effect of pushing the small bowel, a highly radio-sensitive tissue, farther away from the PTV.

3.1.2 Procedure

After the patient is set up on the treatment couch and positioned using external markers and lasers, two orthogonal kV images are taken. The position of the fiducial markers on the kV images is compared to the marker position on the DRR from the pCT and then the couch is shifted to minimize the difference between the two image sets. Next, a CBCT is taken to confirm that the fiducial markers lie within the pCT fiducial + 2 mm contour as well as to evaluate OAR position and volume using the rectum and bladder decision trees. If the CBCT anatomy is deemed acceptable, the first arc is delivered. A second set of kV images is taken to assess intra-fraction motion. If a couch shift of >4 mm in any direction is required to match the fiducial markers on the kV images to the DRR from the pCT, then another CBCT is taken. The decision trees are then navigated for a second time with this CBCT. Once any potential intra-fraction motion is corrected the second treatment arc is delivered. Finally, a post-treatment CBCT is taken for offline review by an RO.

3.1.3 Decision Trees

While the previously described preparation protocols reduce the amount of inter-fraction variability of the bladder and rectum, the bladder and rectum are ultimately non-rigid structures and it is impossible to completely eliminate variation in practical clinical scenarios. It is inevitable that some patients will come in for treatment poorly prepared and there must be systems in place to prevent mistreatment. The ASSERT clinical trial [60] compared SABR fractionation (36.25 Gy/5 fx) to mild hypofractionation (73.68 Gy/28 fx) in prostate patients. For the trial, our center developed decision trees to help determine if the patient is adequately set up and prepared. At each trial treatment, a RO and a medical physicist (MP) were present and they used the decision trees to evaluate the patient preparation and setup. After the trial, a registry was opened to provide more prostate patients with the convenience of SABR fractionation. Due to the higher expected number of SABR patients, it was no longer feasible for a MP/RO to be present at every treatment and the decision trees were adapted so that they could be applied by radiation therapists.

To develop the original rectum decision tree, the dosimetric impact of potential rectum distension scenarios was studied by comparing combinations of linear and angular shifts and stretches of the rectum. Ultimately, it was decided that the rectum should not distend anteriorly past the planning rectum + 10 mm contour for more than one fraction in order to minimize the percentage of patients receiving grade III toxicities. The original decision tree in Figure A.1 was developed to reflect this. Later, alterations were made to this decision tree to facilitate prostate SABR patients who had SpaceOAR hydrogel implants (Figure 3.2). Using this decision tree, patients are pulled from treatment less frequently as it allows for delivery to an imperfect setup in some select cases.

During the ASSERT trial, the RO and MP evaluated bladder preparation by measuring the distance between the femoral heads and bladder dome. The bladder dome had to be between 2–5 cm superior of the femoral heads to minimize bladder wall and small bowel dose. This distance corresponded to an ideal bladder volume between 250 and 500 cm³. The original process (See Figure A.2) was deemed too time intensive for therapists to conduct multiple times per day as it often involved estimating the position of two mis-aligned femoral heads. The revised bladder decision tree for patients with SpaceOAR gel (Figure 3.1) simply asks the therapist to compare the bladder volume at treatment to the pCT volume. The ideal case is that the bladder volume is 2/3 to full as compared to the volume on pCT. Overfull is generally safe to treat, but it can also be uncomfortable for the patient, and for some of these patients muscle clenching can lead to unexpected prostate motion. On the other end, if the bladder volume is less than 2/3, then the decision tree prompts the therapists to look at the small bowel. If more small bowel is inside the 80% isodose than approved by the RO at the time of planning, the patient is still treated. However, the RO is then tasked to review the patient's CBCTs prior to the next fraction. As per this decision tree, there are not any cases except for RO intervention where a patient might be pulled from the couch prior to treatment due to inadequate bladder preparation.

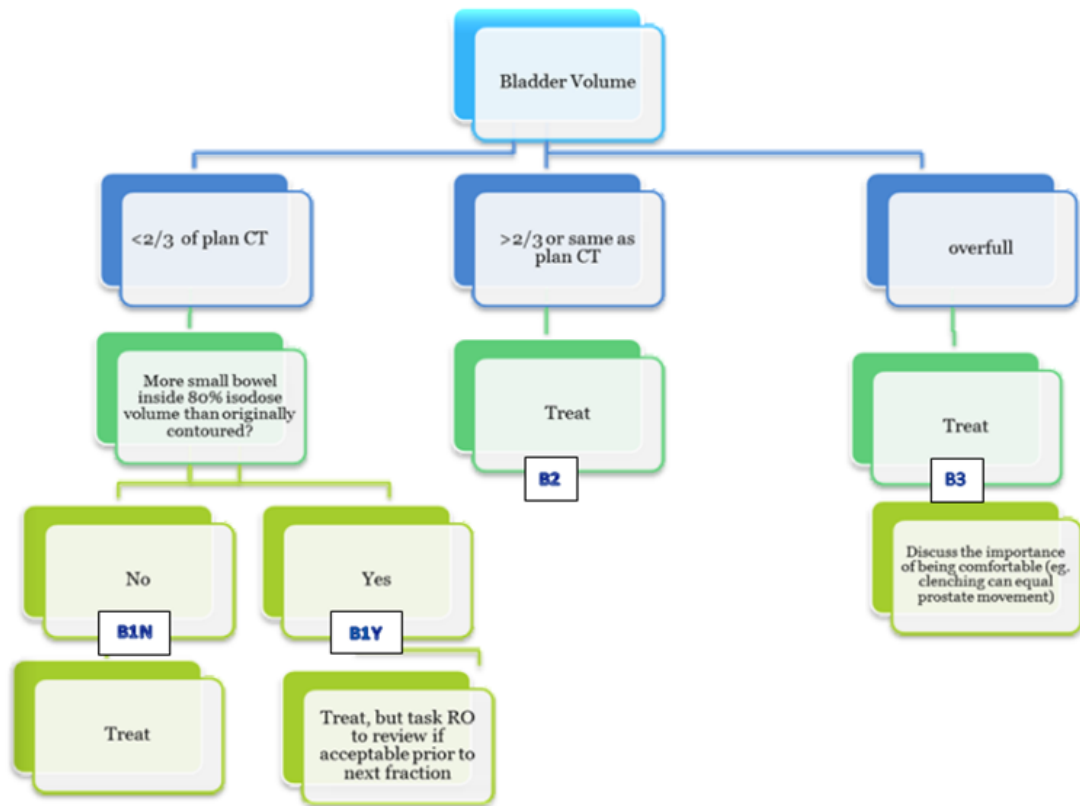


Figure 3.1: Bladder Decision Tree for Prostate SABR Patients.

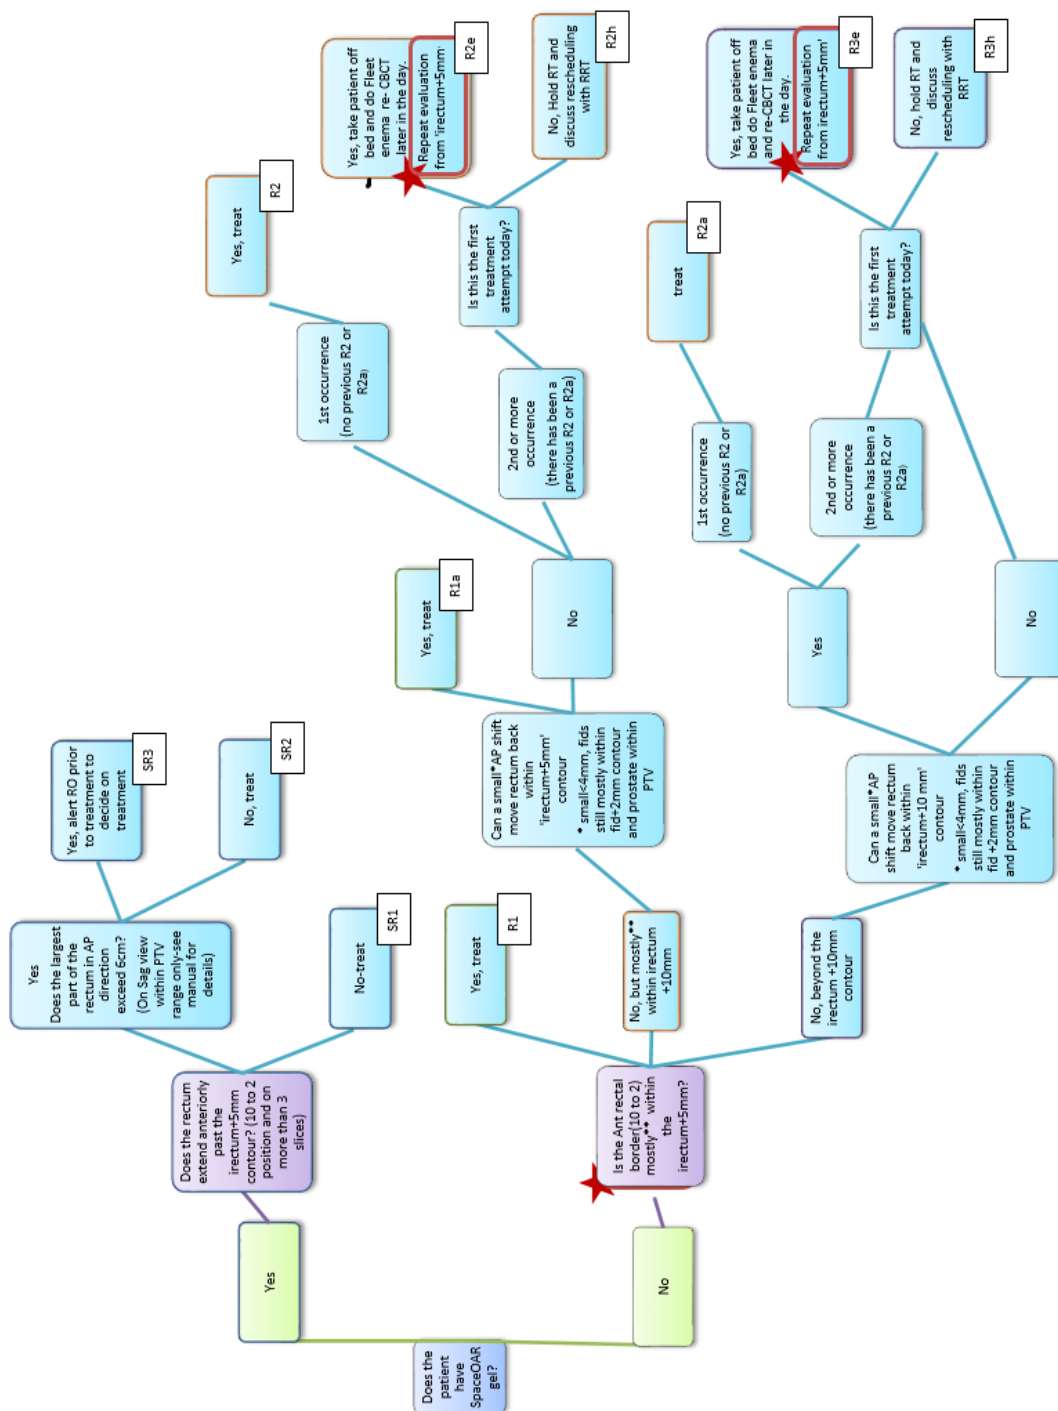


Figure 3.2: Rectum Decision Tree for Prostate SABR Patients.

3.2 Calculation of CBCT OAR Dose

Our goal was to use CBCT contours to calculate DVH curves for the bladder and rectum on each day of treatment. Our approach to this was to use Limbus AI's auto-contouring software to generate the CBCT contours, then superimpose the planned dose distribution onto the CBCT images and compute CBCT OAR DVH curves. Corrections to Limbus-generated contours were made by a radiation oncologist if any adjustments were required. Our ROs found that corrections to the rectum contours tended to be required at the mid-prostate level in approximately 80% of CBCTs. This was likely due to Limbus prostate CBCT models being trained on patient data without SpaceOAR gel. Furthermore, SpaceOAR gel is known to have poor contrast on CT images. At BC Cancer Victoria, MRI images are also taken prior to treatment planning to aid in the contouring of the SpaceOAR gel. The advent of SpaceOAR Vue [61] promises increased contrast on CT-based imaging devices and will likely result in improved rectum contours once approved for use in Canada. Bladder corrections were typically made around the prostate median lobe or a bladder diverticula, however bladder contour corrections were required less often than rectum corrections.

Superimposing the planned dose was deemed necessary as dose re-calculation would simply take too long to compute, all while the patient is waiting for treatment with a full bladder. Furthermore, our prostate SABR procedure involves CBCTs in Spotlight mode, which have been shown to be insufficient for accurate dose re-calculation due to both CBCT HU values not matching pCT HU values and the smaller field of view not capturing the entire patient body contour [62]. We hypothesized that the superimposing of dose for prostate SABR would be sufficiently accurate as the tissues in the immediate tumor region are fairly homogeneous, unlike lung where air causes dramatic density changes. Additionally, the VMAT treatment technique spreads the radiation delivery over multiple 360 degree arcs such that dose inaccuracy from any one angle is likely averaged out. We examined the feasibility of superimposing planned dose vs. full dose recalculation using the methods outlined in Section 3.3.

DVH curves for CTV or PTV structures were not examined in this work. Primarily, our focus was on the DVH curves of OARs, as the use of fiducial marker implants provides adequate prostate coverage [33, 34, 35]. Additionally, the intended application of Limbus Contour is to contour healthy tissue rather than tumor volumes. During the data collection period, the version of Limbus we used contoured

the prostate and seminal vesicles as separate structures. In order to generate a CTV structure, these contours would need to be combined in order for a RO to make corrections. However, current versions of Limbus Contour now combine the prostate and seminal vesicle structures, presenting the potential for future work to explore the calculation of DVH curves for CTV or PTV structures. We also did not examine the dose to the small bowel at treatment as this structure is not currently supported by Limbus Contour.

3.2.1 CBCT Selection

We preferentially used pre-treatment CBCTs to estimate OAR dose as our goal was to evaluate patient setup and preparation. Occasionally, the pre-treatment CBCT was not registered to the pCT and the post-treatment CBCT was used instead. In these cases we compared the pre-treatment and post-treatment CBCTs and only used a patient's data if the differences in contours were less than 1 mm.

3.2.2 CBCT OAR Dose Calculation Method

To calculate CBCT OAR DVH curves, we began by performing an inverse affine transformation (Equation 3.1) on the planned dose distribution voxels. The purpose of this was to obtain the position of each dose voxel in the CBCT coordinate system.

$$\begin{bmatrix} CBCT_x \\ CBCT_y \\ CBCT_z \\ 1 \end{bmatrix} = \begin{bmatrix} M_{11} & M_{12} & M_{13} & T_x \\ M_{21} & M_{22} & M_{23} & T_y \\ M_{31} & M_{32} & M_{33} & T_z \\ 0 & 0 & 0 & 1 \end{bmatrix}^{-1} \begin{bmatrix} pCT_x \\ pCT_y \\ pCT_z \\ 1 \end{bmatrix} \quad (3.1)$$

The first vector on the left-hand side is the x , y , and z positions of the dose distribution in CBCT-space. The matrix term is the registration matrix pulled directly from the associated registration DICOM file, and the final vector term is the x , y , and z positions of the dose distribution in pCT-space. The matrix term is inverted because the equation typically has the matrix on the left-hand side, allowing one to calculate the position of CBCTs in the pCT coordinate system. However, for our purpose it was simpler to shift the dose distribution into CBCT-space instead of the contour points into pCT space. By applying the inverted registration matrix to the position of every dose voxel, the location of the dose voxels in CBCT-space could be obtained.

Since the pCT slice thickness/position of the dose distribution does not always match the CBCT slice thickness/position, it was necessary to interpolate the dose distribution. The reason for this was to have the contour and dose data on the same image slice positions. We used SciPy’s `RegularGridInterpolator` function [63] to create a linear interpolator object that could calculate the dose at an arbitrary point in space, allowing for the calculation of dose at the position of the CBCT slices. This function takes advantage of the consistent spacing of dose voxels, thus avoiding the computationally expensive triangulation methods used in other more generalized interpolation methods.

Compute time was further reduced by only considering voxels within a CBCT contour, necessitating the creation of a boolean mask where array elements are true if the voxel lies within a contour. The mask was created using Matplotlib’s `path` function [64], which not only facilitates the creation of a closed polygon using CBCT contour points, but also has a built-in points-in-polygon function that can quickly determine which points on a grid are located within the contour. Using the interpolator object and the boolean mask, we computed the dose within a contour on the CBCT image slices, reducing the number of calculations to the minimum amount necessary for DVH calculation. As an additional time saving measure, we did not interpolate the dose onto a grid that matched the CBCT resolution (typically 512×512). Instead we used the same pixel spacing as the original dose distribution (approximately 136×167). See Figure 3.3 for a graphical representation of the interpolation process. The final result is a three dimensional array where the only non-zero values are the dose voxels that are within an OAR contour.

DVH curves were obtained by systematically summing the number of voxels within a contour that were less than a certain dose value, then multiplying by the voxel size. The DVH curves were evaluated at the same plan constraint values as those in Table 2.1. If the CBCT OAR dose-volume metric was less than the primary goal then it was considered optimal, if it was greater than the secondary goal, it was classified as a major violation, while values between the primary and secondary goals were categorized as minor violations. The entire CBCT OAR dose calculation process is summarized in Figure 3.4.




Figure 3.3: Interpolation of dose (blue) onto CBCT slices (orange), circles are points that lie within the CBCT slice's contour and are the only points where dose is interpolated. Vertices represent the center of each dose/CBCT voxel.

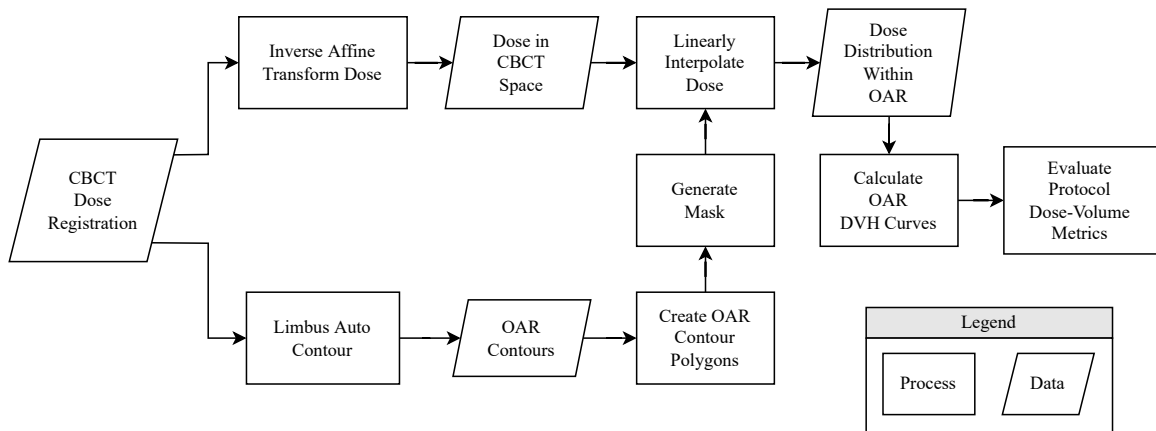


Figure 3.4: CBCT OAR Dose Calculation Workflow

3.2.3 Workflow Automation

While developing our method for CBCT OAR DVH calculation, we also implemented tools to help automate our workflow. We developed a simple GUI to prompt the user to select the OARs of interest. The list of available options is pulled directly from the contour DICOM, and the user can select as many structures as desired and the program will calculate CBCT DVH curves for each of those structures. In addition, our system identifies the prescription dose from the dose DICOM file, automatically finds the matching protocol. It then parses the protocol's XML file to dynamically obtain the treatment planning goals for the OARs selected by the user. The CBCT DVH curves are then evaluated against the planning goals and each metric is classified into optimal, minor, or major violation categories. The program was designed to be independent of treatment site as no clinical parameters are hard-coded in the program. All variables are read in from the relevant DICOM headers, and our program could theoretically be used for any treatment site that can be contoured by Limbus. Alternatively, flexibility could be traded off for automation and computation speed by pre-defining the OARs of interest for a treatment site, and calculation of CBCT OAR DVHs could begin as soon as the CBCT DICOM files are accessible by Limbus Auto-contour. In this scenario, it would be possible to have a set of templates for different OARs and treatment sites.

3.3 Dose Shift Study

We examined the accuracy of superimposing dose onto CBCT images instead of recalculating dose by running several worst case patient setup scenarios. These scenarios are shown in Table 3.1 where Reference Distribution refers to the “ground truth” and dose was recalculated as the final step, Evaluated Distributions were created to simulate superimposing of dose onto CBCT contours with the corresponding shift direction. These scenarios were performed on patient CT images due to the difficulties related to CBCT dose calculations as explained in Section 3.2.

The five scenarios listed in Table 3.1 are divided into two categories. Category I simulates the worst case of shifting the patient 4 mm away from the planned treatment isocenter. This shift distance of 4 mm was chosen as it is the maximum allowed couch shift in our BC Cancer Victoria prostate SABR IGRT protocols. This shift is applied when the tumor is accurately aligned to the radiation beam but the rectum is visually judged to have moved too close to the target and estimated to likely receive higher

dose than is clinically acceptable. Typically the treatment isocenter will then be shifted anteriorly to pull dose off the rectum (Case Ia) but for completeness of the study we simulated a 4 mm shift posteriorly (Case Ib).

Category II repeats the same shifts as category I, but includes the complicating factor of air in the rectum. While the strict bowel and bladder prep that patients follow is meant to prevent gas bubbles in the rectum, this is not always avoidable. Here we have simulated the highly unlikely, but worst case scenario of the entire rectum at the level of the prostate being filled with air. Case II assumes no isocenter shift while Case IIa is shifted anteriorly and Case IIb is shifted posteriorly.

The 5 scenarios were tested on 10 patient CT scans. All dose recalculations were done using Varian Eclipse AAA v15.6. For each case we determined the average local gamma passing rate using the open source PyMedPhys library [65]. Gamma calculation parameters consisted of a 2% dose difference tolerance, a 2 mm dose distance-to-agreement threshold, and an interpolation fraction of 10, which is the fraction that the distance threshold is divided into for interpolation. This analysis was repeated using lower dose thresholds of 10% and 50% of the maximum reference distribution dose.

While the gamma passing rate is helpful in evaluating global dose differences between two distributions, it is not a sensitive metric to local dose differences near the target that could have significant clinical impact. Furthermore, gamma analysis metrics typically would not be used to assess patient setup at the treatment, ROs are normally interested in whether or not target volumes or OARs are meeting certain dose-volume metrics. For these reasons, we also compared the dose-volume metrics that were used as planning goals (Table 2.1) for the Evaluated Distribution vs. the Reference Distribution. For each planning metric, we computed the average percent difference and average absolute difference between the two distributions for the 10 patients studied.

	Reference Distribution	Evaluated Distribution
Case Ia	1. Isocenter shifted 4 mm anteriorly 2. Recalculate dose	1. Original planned dose 2. Shift dose 4 mm anteriorly
Case Ib	1. Isocenter shifted 4 mm posteriorly 2. Recalculate dose	1. Original planned dose 2. Shift dose 4 mm posteriorly
Case II	1. Set rectum contour to air 2. Recalculate dose	1. Original planned dose
Case IIa	1. Isocenter shifted 4 mm anteriorly 2. Rectum contour set to air 3. Recalculate dose	1. Original planned dose 2. Shift dose 4 mm anteriorly
Case IIb	1. Isocenter shifted 4 mm posteriorly 2. Rectum contour set to air 3. Recalculate dose	1. Original planned dose 2. Shift dose 4 mm posteriorly

Table 3.1: Steps taken to obtain reference and evaluated dose distributions

Chapter 4

Results

4.1 Dose Shift Study

Average gamma passing rates for Cases Ia, Ib, II, IIa, and IIb from 10 different patients (Table 4.1) were 96.87%, 96.84%, 99.30%, 96.25%, and 96.17% respectively. These were very high average gamma passing rates given the strict 2%/2mm local gamma calculations these dose distributions were subjected to. Case II had no anterior/posterior shifts and scored the highest average gamma passing rate. Cases Ia and Ib were approximately a 2.5% decrease over Case II, but they involved a 4 mm shift after the evaluated dose distribution was calculated. Cases IIa and IIb had the lowest average gamma passing rates; however, these distributions had both a 4 mm shift and the reference distribution calculated with air in the rectum. The American Association of Physicists in Medicine (AAPM) Task Group 218 [66] recommends gamma passing rates greater than 90% using 3%/2mm criteria for universal action limits in IMRT QA, and 3%/2mm \geq 95% for universal tolerance limits. In all cases we saw average gamma passing rates of greater than 95% with our stricter criteria of 2%/2mm. Sample 2D plots of dose difference and gamma index distributions for each case can be seen in Appendix B. Voxels with gamma values greater than 1 tended to be at the skin in cases that involved shifts, or in the rectum in cases where it was calculated with air.

We repeated the same gamma analysis using a 50% dose threshold as this volume is typically of more interest to clinicians. We observed an increase in average passing rates compared to the values in Table 4.1. Because the largest dose differences were at the skin level, performing the gamma analysis at a 50% threshold meant that those voxels that would have failed the criteria were not included in the calculation

of the gamma passing rate. As a result, the average gamma passing rate increased to 100.00%, 100.00%, 99.47%, 99.76%, and 98.53% for Cases Ia, Ib, II, IIa, and IIb respectively.

	Reference Distribution	Evaluated Distribution	Average Gamma Passing Rate (10% threshold)	Average Gamma Passing Rate (50% threshold)
Case Ia	1. Isocenter shifted 4 mm anteriorly 2. Recalculate dose	1. Original planned dose 2. Shift dose 4 mm anteriorly	96.87%	100.00%
Case Ib	1. Isocenter shifted 4 mm posteriorly 2. Recalculate dose	1. Original planned dose 2. Shift dose 4 mm posteriorly	96.84%	100.00%
Case II	1. Original planned dose	1. Set rectum contour to air 2. Recalculate dose	99.30%	99.47%
Case IIa	1. Isocenter shifted 4 mm anteriorly 2. Rectum contour set to air 3. Recalculate dose	1. Original planned dose 2. Shift dose 4 mm anteriorly	96.25%	99.76%
Case IIb	1. Isocenter shifted 4 mm posteriorly 2. Rectum contour set to air 3. Recalculate dose	1. Original planned dose 2. Shift dose 4 mm posteriorly	96.17%	98.53%

Table 4.1: Average gamma passing rate (2%/2 mm criteria) for five dose approximation scenarios

When comparing cases using planning dose-volume metrics (Table 4.2), there were several instances where the average percent difference was large, however these were followed by a small absolute difference. This occurred in OAR dose-volume metrics that were numerically small. Comparing two small numbers often yields a large percent difference value, yet a small absolute difference. There were also occurrences of large absolute differences with small relative differences in cases where dose-volume metrics had large numerical values such as the target volumes. There were no metrics that had both a large average percent difference and a large average absolute difference.

Considering the rectum V36 Gy metric, Case IIb had the worst average absolute difference at $0.84 \pm 1.94 \text{ cm}^3$ while Case Ib had the overall largest singular difference at 7.51 cm^3 . However we considered these cases unrealistic as we would never shift the dose closer to the rectum in the clinic. The worst of the realistic scenarios was Case II, which had an absolute average difference in V36 Gy of $0.29 \pm 0.69 \text{ cm}^3$. However, our ROs considered this to not be clinically significant. For the rectum V18, V29, and

V33 Gy relative metrics, we saw absolute differences of less than 1% for the majority of metrics, with a difference of up to 1.6% for Case IIb's rectum V29 Gy metric, which we considered to be acceptable and within tolerance limits.

For the bladder V37 Gy metric, the largest average absolute difference occurred in Case IIa at $1.03 \pm 1.27 \text{ cm}^3$ and the largest singular absolute difference was given by Case IIb at 3.06 cm^3 . This was also not considered clinically significant by our ROs. The remaining cases were all less than 0.5 cm^3 different for the bladder V37 Gy metric. The bladder V18, V33, and V36 Gy relative metrics all had an average absolute difference of 0.22% or less. Finally, both left and right femur heads had absolute differences of less than 0.1%.


The left and right femur heads are 'large' structures and none of the average absolute differences for all cases were greater than 0.07%. As for target volumes, the average percent difference for the PTV D95% metric in each case was less than 0.5% and was considered acceptable. All cases, except for Case II, had an average percent difference and average absolute difference of around 6–7% for the D0.003cc metric. The D0.003cc metric is the highest dose delivered to 0.003 cm^3 of the PTV, its value at treatment planning should be less than 105% of the maximum dose for the primary goal, or less than 107% for the secondary goal. The large difference in the D0.003cc metric was likely caused by our dose calculation resolution of $0.125 \times 0.125 \times 0.125 \text{ cm}$ whereas a volume of 0.003 cm^3 was just under two voxels of our dose calculation voxel size, requiring interpolation to compute. For this reason, the D0.003cc metric is particularly sensitive and small differences in the reference and evaluated distributions may yield large differences in D0.003cc values. For the CTV D99%, Cases Ib and IIb had unacceptable average percent differences at 2.35% and 2.93% respectively, but as mentioned earlier we did not consider these cases to be clinically realistic. Cases Ia, II, and IIa had an average percent difference of less than 0.75% which was allowable according to our ROs.

ROI	Metric	Case Ia		Case Ib		Case II		Case IIa		Case IIb	
		Average %	Difference	Average %	Difference	Average %	Difference	Average %	Difference	Average %	Difference
Rectum	V36 Gy [cm ³]	0.00±0.00	0.00±0.00	1.38±2.25	0.01±0.02	65.92±82.02	0.29±0.69	20.00±60.00	0.00±0.01	30.79±33.96	0.84±1.94
Rectum	V33 Gy [% of volume]	0.75±1.82	0.00±0.01	0.76±1.18	0.02±0.03	27.99±36.89	0.08±0.07	22.04±31.28	0.10±0.25	7.44±8.61	0.19±0.17
Rectum	V29 Gy [% of volume]	6.82±19.95	0.00±0.00	0.51±0.64	0.03±0.04	9.51±10.52	0.13±0.17	23.4±59.23	0.03±0.05	4.34±4.41	0.23±0.27
Rectum	V18 Gy [% of volume]	0.89±1.21	0.04±0.06	0.23±0.32	0.06±0.07	6.61±3.29	0.98±0.81	4.41±3.97	0.41±0.67	6.59±3.19	1.60±0.99
Bladder	V37 Gy [cm ³]	8.72±12.51	0.45±0.40	9.11±11.66	0.20±0.13	13.77±29.2	0.52±0.89	18.21±31.92	1.03±1.27	12.82±23.15	0.29±0.51
Bladder	V36 Gy [% of volume]	2.17±0.74	0.12±0.08	3.39±1.55	0.07±0.04	3.28±3.76	0.10±0.11	4.62±3.42	0.22±0.15	3.82±3.86	0.07±0.07
Bladder	V33 Gy [% of volume]	1.17±0.41	0.09±0.07	1.63±0.70	0.06±0.03	1.40±1.65	0.06±0.07	2.40±1.38	0.17±0.10	1.82±1.67	0.05±0.04
Bladder	V18 Gy [% of volume]	0.33±0.13	0.08±0.07	0.51±0.23	0.09±0.07	0.64±0.65	0.10±0.08	0.83±0.50	0.17±0.11	0.70±0.64	0.09±0.06
Femur Head L	V14 Gy [% of volume]	7.77±8.39	0.04±0.04	4.59±4.33	0.04±0.05	2.24±5.40	0.01±0.01	6.43±7.24	0.04±0.04	3.93±3.82	0.03±0.04
Femur Head R	V14 Gy [% of volume]	4.85±7.03	0.04±0.05	4.13±5.90	0.04±0.05	9.24±17.21	0.04±0.04	12.09±25.47	0.07±0.06	7.65±12.65	0.02±0.02
PTV	D0.03cc [% of dose]	6.58±0.62	6.75±0.64	6.02±0.58	6.16±0.61	0.50±0.45	0.53±0.48	7.03±0.74	7.23±0.78	6.37±0.83	6.53±0.87
PTV	D95% [cGy]	0.08±0.08	2.86±2.56	0.47±0.13	16.51±4.36	0.46±0.87	16.79±31.23	0.41±0.74	13.94±25.59	0.46±0.44	16.03±15.42
CTV	D99% [cGy]	0.14±0.08	5.34±3.14	2.35±5.84	84.2±206.43	0.74±1.62	26.81±56.76	0.61±0.87	22.36±29.44	2.93±6.10	104.6±216.1

Table 4.2: Dose distributions compared using planning goals for five dose approximation scenarios.

4.2 Verification of Method Implementation

Calculating DVH curves for different organs and targets often involves interpolating which dose voxels are considered to belong inside or outside the organ contours of interest. To verify that our python implemented DVH calculator (as described in Section 3.2.2) was accurate, we compared it to the results of the commercial Varian Eclipse system used in the clinic. Since there were no ground truth DVH curves for CBCT OARs due to the previously described limitations of calculating dose on CBCTs, we used our method to re-calculate DVH curves on the pCT contours. We made this comparison for all 40 patients in our study. In the case of calculating DVH curves for pCT contours, the registration matrix was simply an identity matrix as there should be no spatial shift in the dose distribution. We compared planning and calculated bladder and rectum DVH curves for each patient; a sample set of these are shown in Figure 4.1. Differences were minimal and could only be seen with significant magnification, which gave us confidence in the execution of our method.



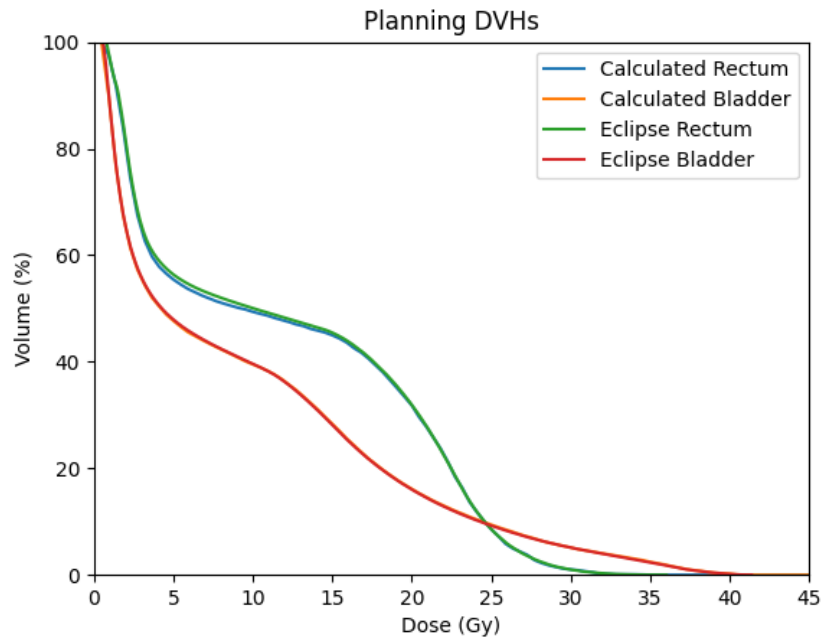



Figure 4.1: Example DVH curves using data directly from Eclipse compared to calculated DVH curves using pCT contours and the method described in Section [3.2.2](#)

4.3 DVH Curves Calculated Using CBCT Contours

OAR	Protocol				Plan	Fraction				
	Dose (Gy)	Minor	Major	Unit		1	2	3	4	5
Rectum	36	1	2	cm ³	0.55	3.22	5.33	3.66	3.35	2.92
Rectum	33	10	12.5	%	3.76	9.58	14.45	9.59	10.27	10.20
Rectum	29	20	22.5	%	8.86	15.02	21.45	14.92	16.41	16.55
Rectum	18	40	50	%	38.35	44.59	51.90	47.45	44.81	52.49
Bladder	37	5	20	cm ³	5.56	1.42	0.14	0.41	0.68	1.17
Bladder	36	5	10	%	1.57	0.71	0.12	0.27	0.23	0.69
Bladder	33	10	20	%	2.49	1.52	0.54	0.76	0.52	1.34
Bladder	18	30	45	%	8.14	8.08	6.11	5.45	3.25	6.82

Table 4.3: Example results from one patient whose CBCT scans were evaluated against the prostate SABR 40 Gy protocol (Table 2.1). Treatment planning values are given for comparison. Green indicates that the dose-volume metric calculated from CBCT data was optimal and less than the minor goal, yellow is when the dose-volume metric is between the minor and major goal, and metrics that exceeded the major goal are red. For this patient, all fractions were optimal in terms of bladder metrics, however all fractions had a rectum V36 Gy major violation and were therefore classified as fractions with a major violation.

We processed 40 patients, each with 5 treatment fractions, through our program to obtain a set of results similar to Table 4.3 for each patient. DVH curves for bladder and rectum at each fraction were evaluated against the planning goals from Table 2.1 and classified into optimal (green), minor (yellow) or major (red) categories. Comparing calculated CBCT OAR DVH curves to the planning goals gave us a picture of patient preparation and setup performance. To obtain the results shown in Figure 4.2 we classified the 200 fractions we processed according to the presence of a minor or major violation. Fractions where no planning metric was violated were classified as optimal. If a fraction contained a minor or a major violation in any metric it was categorized accordingly. Fractions with minor and major violations in different metrics were simply recorded as major violations so as to not double-count fractions.

No minor violations in any rectum metrics were in the initial CT treatment plans for our set of 40 patients. Despite rigorous patient bowel and bladder preparation, and our continued effort to create high quality treatment plans, Figure 4.2 shows an increased occurrence of both minor and major dose violations at treatment delivery.

Twenty-seven percent of fractions had at least one major violation in one of the bladder planning goals, while 34% had minor violations. Just 39% of treated fractions had neither minor nor major violations in bladder metrics. Fourteen percent of fractions had a rectum metric major violation, with 10% categorized as minor violations, while the remaining 76% of fractions were classed as optimal. Bladder metrics had more minor and major violations than rectum metrics at treatment delivery, perhaps due to the greater geometric variability of the bladder. The following sections will provide more detailed results for the rectum and bladder.

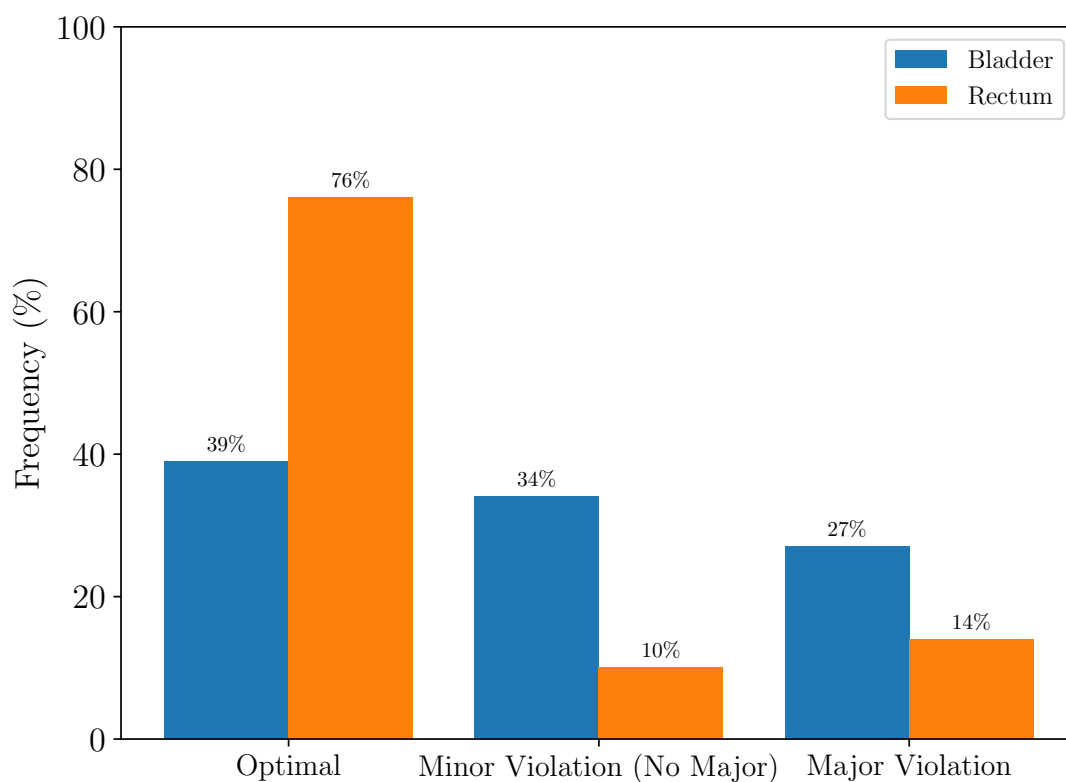


Figure 4.2: Fraction classification frequency. Fractions were classified according to the presence of any minor or major bladder and rectum dose violations. Fractions that had minor and major dose violations in different metrics were classified as major violations.

4.4 Bladder

Figure 4.3 shows the classification frequency for the 200 processed fractions broken down by bladder dose-volume constraint. If a fraction contained a planning goal violation for any metric it was classified as either minor or major according to its value. The percentage of fractions with major violations were less than 2.5% for V18, V33, and V36 Gy metric but rose to 25% for the V37 Gy metric. Minor violations were also higher for V37Gy at 33.0% compared to 7.0%, 8.5%, and 20.0% for V18, V33, and V36 Gy respectively.

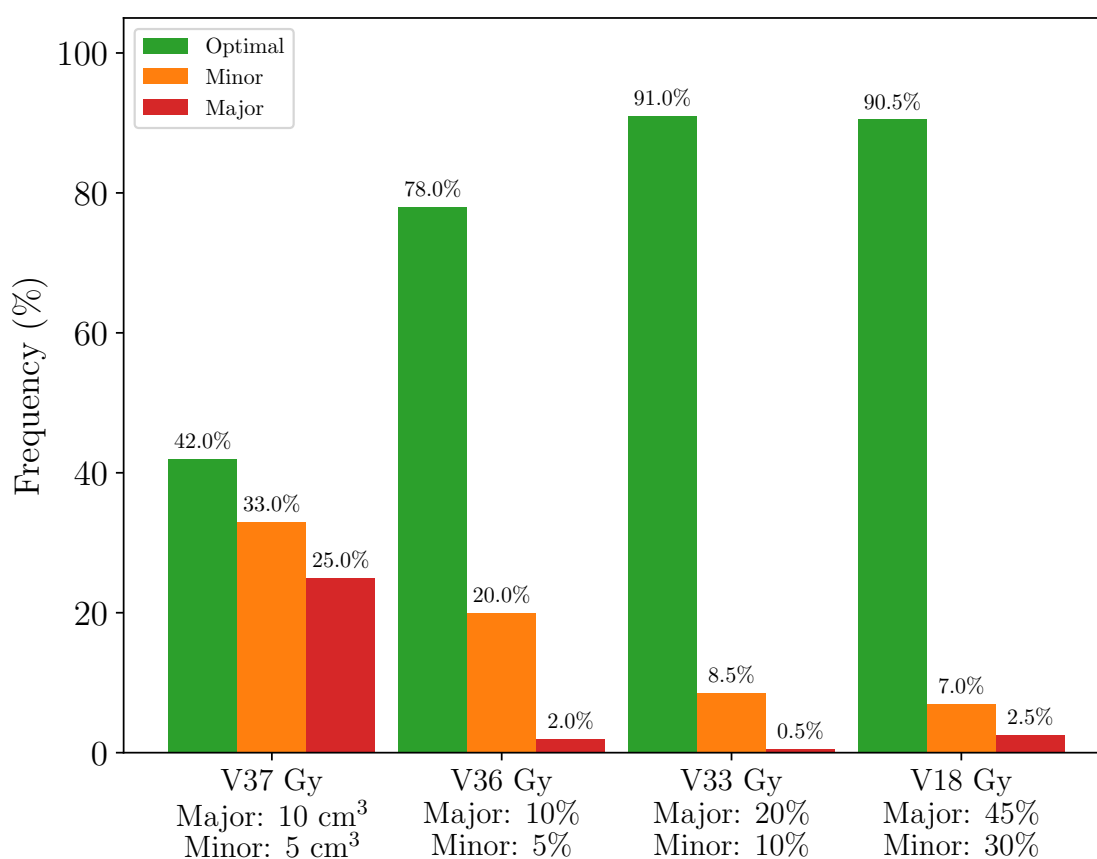


Figure 4.3: Classification of dose-volume metrics using daily CBCT bladder contours

At treatment planning, no patients had major violations of bladder metrics, however some minor violations were accepted. These occurred primarily in the V37 Gy metric where 8/40 patients had a planned minor violation. A single patient had a planned minor violation in the V36 Gy and V33 Gy metrics. For the V18 Gy metric, there were 2 patients with planned minor violations. Figure 4.4 shows how many pa-

tients received a certain number of major violations over the course of their 5-fraction treatment. Out of 40 patients, only 22 had no bladder major violations and of these, only 12 patients had all their bladder metrics classified as optimal over the course of their treatment. There were 5 patients who received a major dose violation of a bladder metric during each of their treated fractions. Of particular concern was that for 4 of these 5 patients, the 5 consecutive major violations occurred in the V37 Gy metric. This ‘high dose’ bladder metric is associated with worse Expanded Prostate Cancer Index Composite (EPIC) quality of life (QOL) score [67].

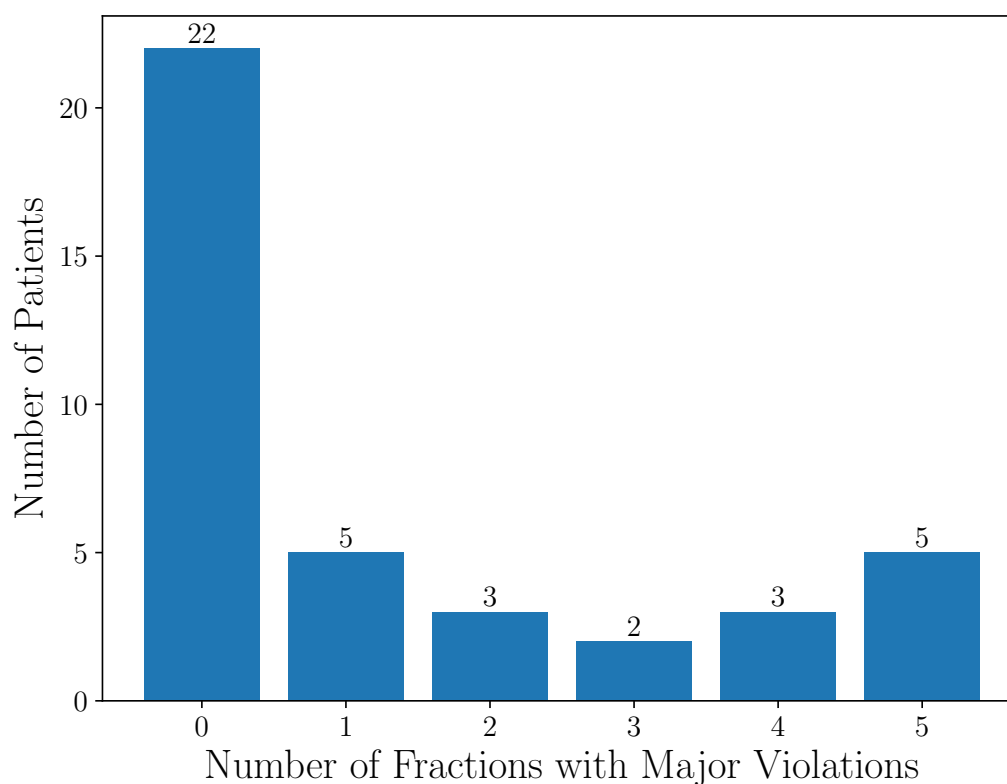


Figure 4.4: Bladder - Patients by number of major dose violations over the course of their treatment.

Bladder Volume

Figure 4.5 depicts the distribution of bladder volumes from pCT to CBCT. With only 40 data points for pCT bladder volumes, it is difficult to draw any conclusions about differences in pCT to CBCT bladder volume.

The scatter plots in Figure 4.6 show the four dose-volume metrics versus the change in bladder volume from CT to CBCT (ΔV). A positive value represents an increase in bladder volume at treatment. As expected, major violations for V18, V33, and V36 Gy occurred when the CBCT bladder volume was less than the pCT on account of the relative nature of these metrics. The V18, V33 and V36 Gy metrics are percent-volume metrics and are directly influenced by changes in total volume by definition. Even if the same absolute partial volume is irradiated to 36 Gy, a decrease in total bladder volume would increase the percent value of these metrics. This observation did not extend to the V37 Gy metric, which is based on absolute volume, where major violations occurred at a wide range of both positive and negative ΔV values. Inspecting these major violations showed portions of the bladder wall moved in the inferior and/or posterior direction, closer to the PTV. We saw up to 1 cm of posterior distension of the bladder within the 80% isodose. Therefore, for the bladder V37 Gy metric, the displacement direction of the bladder was more important than its actual change in volume from pCT.

We also found that the bladder volume at treatment was less than $2/3$ of the pCT volume in 59 of the 200 fractions (Figure 4.7). The bladder decision tree (Figure 3.1) recommends the bladder volume at treatment be at least $2/3$ of the pCT volume. Only 23/59 of the small bladders were correctly categorized as B1Y/N by the radiation therapists at time of treatment delivery. All of the major violations in the V18, V33, and V36 Gy metrics occurred when the bladder volume was at least $2/3$ less than the pCT volume. There were 50 fractions that had major violations of the important “high-dose” V37 Gy metric, 10 of which occurred when the bladder volume was $2/3$ less than at pCT, 4 of the 10 were correctly categorized as B1N by the therapists and proceeded to treatment and no follow-up by the was RO required. The remaining six were incorrectly categorized as B2. One patient had a V37 Gy major violation in every fraction treated. This patient had one of the largest bladder volumes (577 cm^3) on the pCT in this study, our data suggests that they had difficulty reproducing their pCT bladder volume at treatment. Their CBCT volumes ranged from 66.2% down to 23.6% of their pCT volume.




Figure 4.5: pCT and CBCT bladder volume distributions.

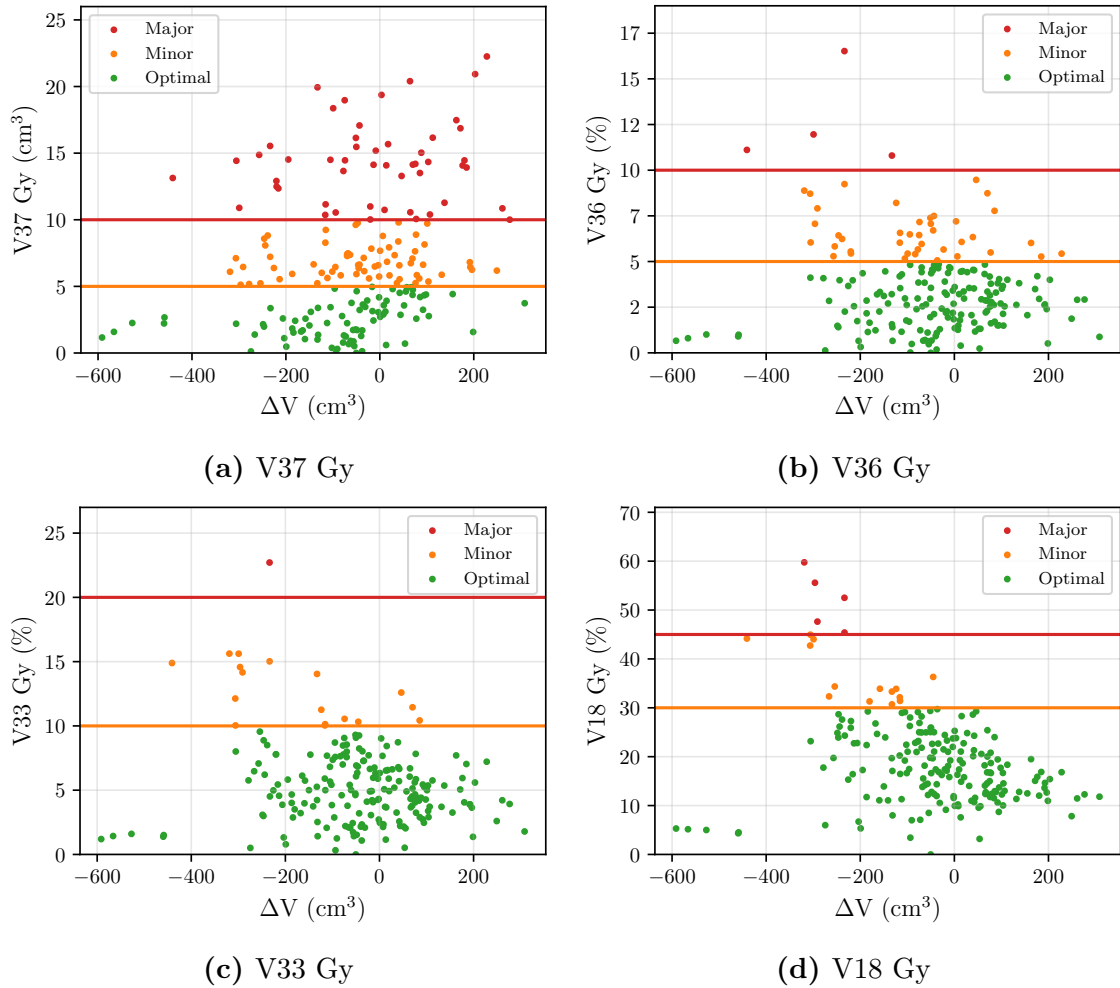


Figure 4.6: Bladder - Dose-volume metrics compared to absolute change in volume from pCT to CBCT ($V_{CBCT} - V_{pCT}$).

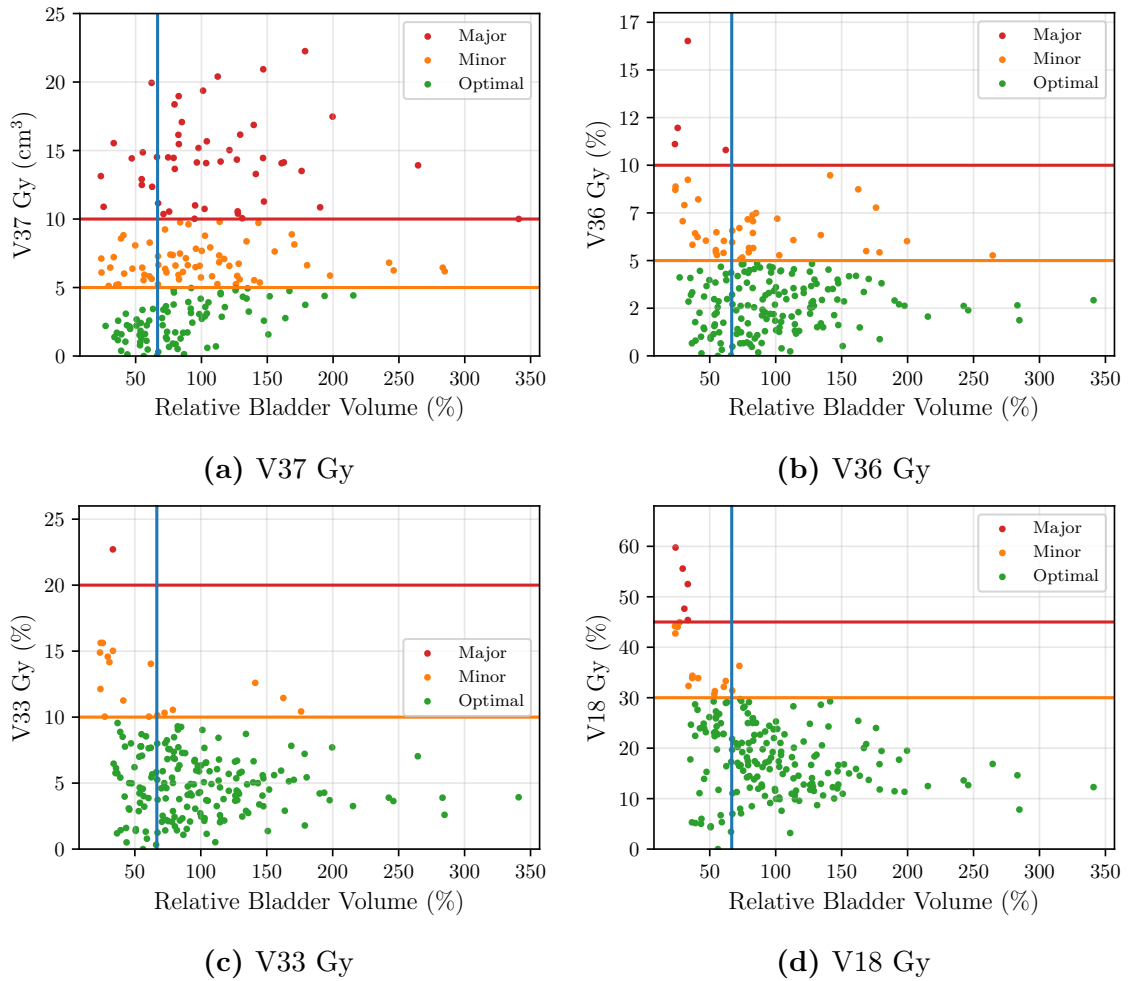


Figure 4.7: Bladder - Dose-volume metrics compared to relative change in volume from pCT to CBCT ($V_{CBCT}/V_{pCT} \times 100\%$). The blue vertical line represents the $2/3$ volume threshold from the bladder decision tree.

4.5 Rectum

Out of 200 analyzed fractions, major and minor violations for rectum were overall lower than bladder and the majority of fractions were fortunately optimal as shown in Figure 4.8. However, there were still an elevated number of “High-dose” rectum V36 Gy minor and major violations. Major violations for rectum V18, V29, and V33 Gy were all below 5%, but 12% of fractions had V36 Gy major violations. V29 and V33 Gy minor violations were both low at 0.5% and 3.0% respectively while V18 and V36 Gy minor violations were slightly more elevated at 6.0% and 7.0%.

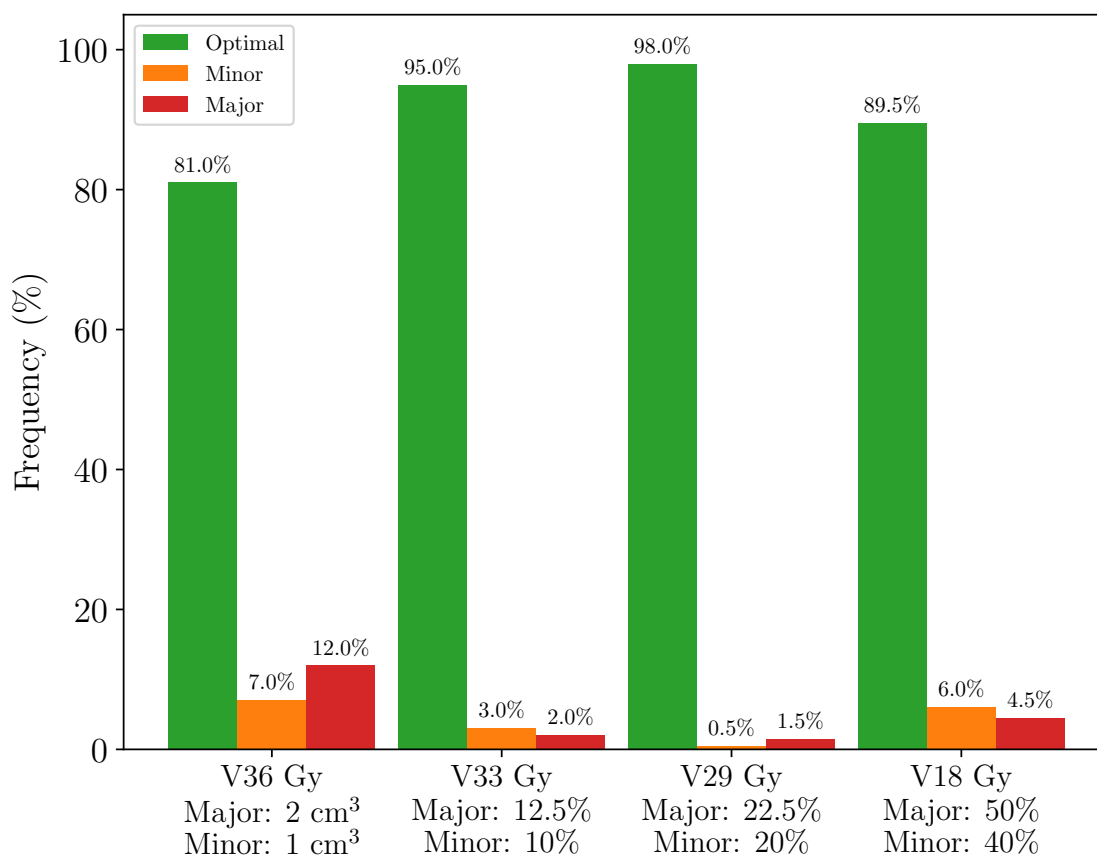


Figure 4.8: Classification of dose-volume metrics using daily CBCT rectum contours.

While the majority of patients had an ideal treatment in terms of rectum metrics, some patients had systematic major dose-volume violations (Figure 4.9). The 2 patients who had all 5 fractions with major violations, and 1 patient who had 4, all had their major violations occur in the V36 Gy metric. The raw data for one of these patients is shown in Table 4.3. These High dose-volume metrics are associated with worse EPIC bowel QOL and grade 2 or higher Common Terminology Criteria for Adverse Events (CTCAE) hematochezia [67].




Figure 4.9: Rectum - Patients by number of major dose violations.

Rectum Volume

The distribution of rectum volumes from pCT to CBCT (Figure 4.10) suffers from the same lack of pCT data points as the bladder volume distributions. More pCT data would permit the use of smaller bin sizes, giving an increased histogram resolution and allowing us to draw further conclusions.

We again examined the relationship between changes in rectum volume and the four rectal dose-volume metrics (Figure 4.11). Major violations for V18, V29, and

V33 Gy occurred at approximately $\Delta V = 0 \text{ cm}^3$. Larger rectum volumes were more often classified as optimal rather than the expected minor or major violations. Large total rectum volumes in these relative metrics may be reducing the metric's sensitivity to dose because a larger total volume would decrease the metric's value for a fixed absolute irradiated partial volume. Meanwhile, V36 Gy showed an increased range of volumes with major violations. Qualitative inspection again gave a picture that it was more important where the rectum was located on CBCT rather than its volume change. Fractions with major violations tended to have anterior rectum distension as compared to the pCT.

Rectum volume at treatment was measured to be at least 10% larger than the pCT volume in 60 of the 200 fractions investigated (Figure 4.12). Of the 24 V36 Gy major violations, 11 of them occurred when the rectum volume was at least 10% larger on CBCT than at pCT. None of the V18, V29, or V33 Gy major violations occurred when the rectum volume was large.




Figure 4.10: pCT and CBCT rectum volume distributions.

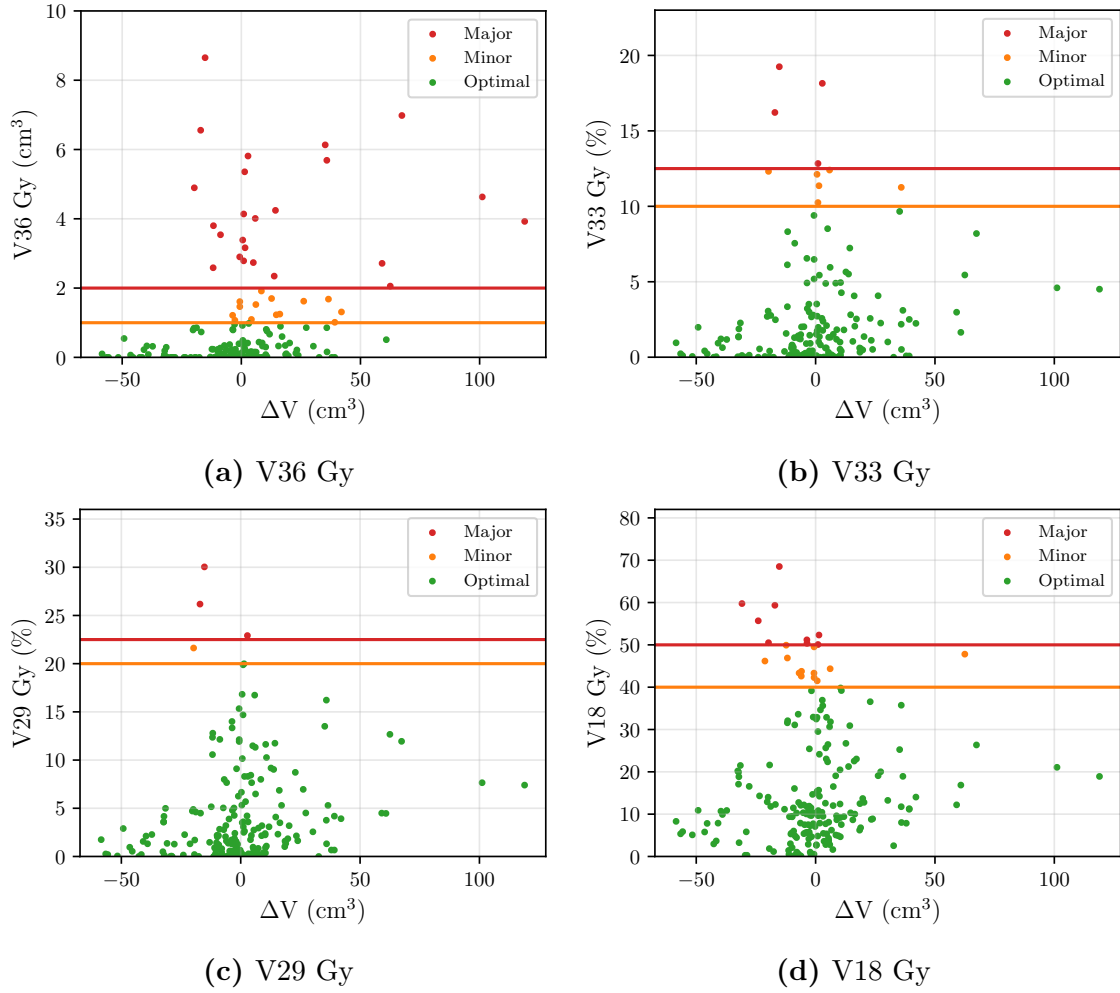


Figure 4.11: Rectum - Dose-volume metrics compared to absolute change in volume from pCT to CBCT ($V_{CBCT} - V_{pCT}$).

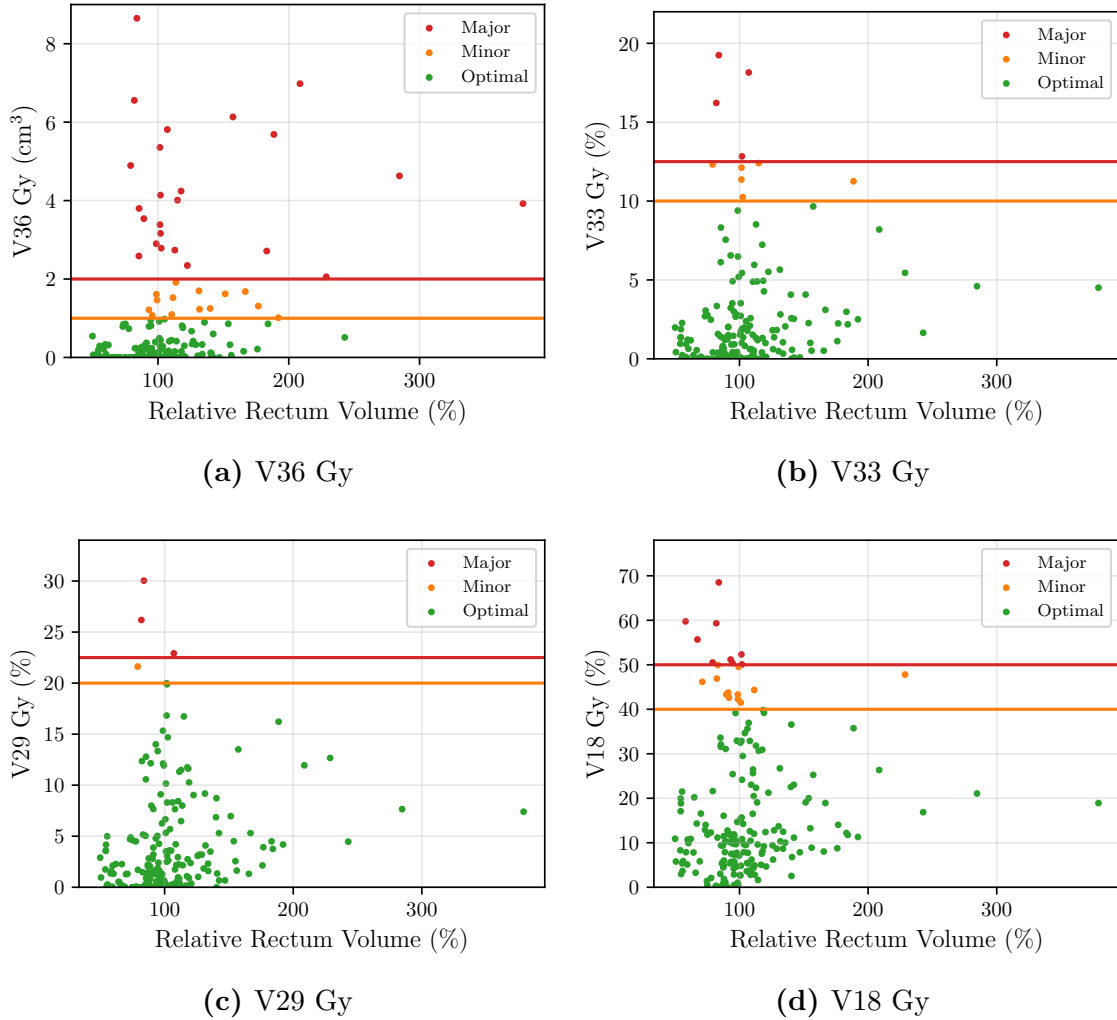


Figure 4.12: Rectum - Dose-volume metrics compared to relative change in volume from pCT to CBCT. ($V_{CBCT}/V_{pCT} \times 100\%$).

Chapter 5

Discussion & Conclusion

5.1 Discussion

In this set of 40 prostate SABR patients (40 Gy/5 fx) we used Limbus-generated CBCT contours to estimate per-fraction DVH curves for bladder and rectum. Until now we had never quantified dose-volume metrics for CBCT anatomy and our results showed that major violations in prostate SABR occurred despite careful planning and patient preparation protocols. Overall, we found that fractions with major violations occurred in bladder and rectum metrics at 27% and 14% respectively. Compared to rectum, major and minor violations occurred more frequently in bladder metrics. This was perhaps due to the greater variability in volume for bladder. The CBCT bladder volume coefficient of variation (standard deviation divided by the mean) was 47% whereas for rectum it was 35%. The high dose, absolute volume metrics had the worst rate of major violations, for bladder this was 25% of fractions, and for rectum this was 12% of fractions. These metrics did not appear to be dependent on the change in OAR volume from pCT to CBCT, but depended more on whether the OAR moved closer to the high dose region on CBCT. Some patients had recurring major violations in high dose metrics. There were four patients who had bladder V37 Gy major violations in all five of their treated fractions. Systematic violations also occurred in the rectum V36 Gy metric where two patients had five major violations, and one patient had four.

Patient bladder volume at treatment was less than $\frac{2}{3}$ of the pCT volume in 59 of the 200 fractions we observed. Only 23 of these instances were identified as category B1Y/N by the therapists. These results indicate that it is difficult to accurately estimate $\frac{2}{3}$ the volume of a three dimensional structure from observing only sets of two-dimensional images (sagittal, coronal, and axial planes) at the treatment console.

A workflow for calculating bladder volume using Limbus contouring is something that could be implemented in a short time frame and would yield more accurate bladder volume estimations.

Our study was similar to one conducted by Devlin et. al. [68] who also used the approach of overlaying a planned dose onto CBCT contours to evaluate dose-volume metrics at treatment. They studied prostate SABR patients treated with 35 Gy in five fractions using a VMAT technique. Our study differed in that we used SpaceOAR gel and auto-contouring for OARs on CBCT. Their study also saw higher than planned bladder and rectum DVHs due to OAR variation. Devlin et. al. suggested that a possible reason for high bladder DVHs was that they only had one planning goal ($V_{35\text{ Gy}} < 1\%$), however we observed higher than planned DVHs with four bladder planning goals. The group saw a negative correlation between their 1% dose metric and bladder volume, while we saw the same correlation in our V18, V33, and V36 Gy metrics. These are all metrics relative to the volume of the bladder. We did not observe the same volume correlation for our bladder V37 Gy absolute metric, where we saw no clear relationship (Figure 4.6). Our results also disagreed in the high dose rectum metrics where Devlin et. al. saw a positive correlation between rectum volume and their $V_{35\text{ Gy}} < 1\%$ metric, whereas we did not see a clear correlation with our $V_{36\text{ Gy}} < 2\text{ cm}^3$ metric (Figure 4.11). This is likely due to being a comparison between a relative and an absolute metric.

In a study examining inter-fraction variability of OARs in prostate SABR patients, who received 45 Gy in 25 fractions using helical tomotherapy and a subsequent 19 Gy in 2 fractions prostate SABR boost, Wahl et. al. [69] saw that rectal V75% had increased by a median of 0.93 cm^3 (range: -1.5 to 14.0 cm^3) compared to the initial treatment plan. They found that daily rectal V75% exceeded the plan constraint of $<2\text{ cm}^3$ in 47% of fractions. While not directly comparable to their V75% metric, we saw fewer violations in our rectum V36 Gy, V33 Gy, V29 Gy, and V18 Gy metrics. This was perhaps due to patients having SpaceOAR gel in our study. The group did not see statistically significant changes in bladder V75% dose at treatment compared to the treatment plans.

Fuchs et. al. [70] studied dose to variable OARs in prostate treatments using a Tomotherapy unit. In their study, they also saw a wide range of changes in rectum and bladder volume from plan to treatment despite their bowel and bladder preparation protocol. Rectal volume at treatment ranged from 62% up to 223% of the pCT volume, in our data this range was similar at 57.0% to 230.8%. As for bladder,

treatment volumes ranged from 22% up to 375% of the initial plan volume, we had a more narrow range of 40.1% to 185.7%. Similar to our results, they found that lower bladder volume was associated with increased mean dose to the bladder, the authors suggest that this is caused by a higher percentage of the bladder volume being closer to the target volume when not properly filled. They did not, however, see a strong correlation between rectum mean dose and rectum volume at treatment.

Another study, by Tanny et. al., [71] examined dose to OARs when using a CBCT-guided online ART platform. Patients were a mix of pelvis and prostate patients and were treated with 45 Gy in 25 fractions. Similar to our results, they found that bladder volumes showed the largest variation, while rectum volume varied to a lesser extent. They saw mean deviations ($(V_{CBCT} - V_{pCT}) / V_{pCT} \times 100\%$) in bladder volume range from -47% up to +12% whereas we saw deviations range from -59.9% up to 85.7% (average of -3.5%, standard deviation of 40%). For rectum, Tanny et. al. saw mean volume deviations from -31% to +16%, our data showed variations in the range of -42.9% up to 130.8% (average of +5.1%, standard deviation of 31.2%). The greater range could be due to patients in our study receiving 5 fractions compared to the 25 fractions in their study. Patients receiving 25 fractions of RT would be able develop a routine in the application of their preparation protocol, reducing the overall variability of OAR volumes. Despite the degree of OAR variability, the group saw favourable DVH distributions at treatment using online ART compared to conventionally scheduled plans. They found bladder V40 Gy, V30 Gy, and V20 Gy doses reduced by 3.9%, 29.8%, and 30.9%, as well as rectum dose decreased by 52.6%, 48.9%, and 24.6% for the same metrics. All the while, PTV coverage was maintained to within 1% of the reference value using the online ART platform, compared to reductions of up to 8% in the scheduled plans. They were able to maintain good target coverage while reducing OAR dose when moving to an online ART platform.

Our process involved overlaying the planned dose onto the CBCT contours, which, as our results in Section 4.1 showed, is a reasonably accurate approximation for anterior or posterior shifts. Each case had high gamma passing rates and no dose-volume metrics had both high percent difference and high absolute difference. Examining other scenarios that involve superior-inferior shifts or OAR rotations, translations, or deformations would further increase confidence in the practice of superimposing the planned dose onto CBCT structures. Additionally, it would be prudent to investigate large magnitude transformations to determine at what point the approximation breaks down and a full re-plan becomes necessary. While superimposing CBCT dose

is not as accurate as the full dose recalculation methods used in dedicated ART solutions, our method emphasizes computation speed over accuracy. It is intended to be used in clinics that do not currently have plans to implement cutting-edge units such as Elekta's Unity, Varian's Ethos, or MRIdian's ViewRay.

Another limitation of overlaying the planned dose is that we are treating each fraction as if the entire prescription dose was delivered during that fraction. This was done in order to classify metrics as optimal, minor, or major as per our planning goals. Furthermore, it was necessary to interpolate the dose distributions to find the dose on slices that matched the CBCT slices, we chose to perform the interpolation on a grid spacing that matched the original dose distribution as to minimize the degree of error resulting from linear interpolation. The approach by Devlin et. al. [68] was to create a verification plan calculated for a single fraction prescription of 700 cGy ($1/5$ of prescription dose), then overlay the CBCT contours. Both approaches lack dose accumulation for OARs as the individual voxels are not tracked across the course of treatment, a requirement for calculating dose accumulation. Deformable image registration (DIR) is a technique that could provide dose accumulation. However, accurate DIR requires thorough testing and evaluation with accurate deformable phantom measurements before it can be clinically implemented [72, 73] which we do not currently have available in our clinic.

Ideally, this work would have also included target volumes. Recent versions of Limbus Contour can create a prostate + seminal vesicle structure, which could be used to create CTV and PTV contours. With this update and our program's flexibility, DVH curves for these target volumes can now be generated. Although this is certainly an interesting application, the benefits from doing so may be limited as implanted fiducial markers are already standard practice for prostate SABR patients at our clinic.

The bladder decision tree was designed to ensure that the radiosensitive small bowel does not receive unnecessary radiation dose. However, Limbus Contour does not currently support contouring of the small bowel. If it could, our program would be able to calculate a DVH curve for the small bowel contour. However, the use of narrow field of view spotlight CBCTs could make auto-contouring of the small bowel challenging. There is a high probability that the small bowel may be excluded from the imaged volume, especially if the patient's bladder volume is large. Unless wide field of view CBCTs become a standard practice, this would be an inconsistent approach to calculating small bowel DVH curves from CBCT data.

Other centres in our BC Cancer network have utilized a theoretical minimum bladder volume contour during treatment planning. This contour represents the smallest bladder volume that the patient might have during treatment while still satisfying bladder constraints. This approach would likely aid in determining whether the bladder is adequately filled during treatment, as well as ensuring that the relative V36 Gy, V33 Gy, and V18 Gy bladder metrics are met. However, due to the steep dose gradients in prostate SABR treatments, this structure would need to account for posterior bladder distension in order to account for the position-dependant V37 Gy dose. Corrections to Limbus-generated bladder contours were infrequently required by our ROs, further studies using these contours could provide insights into the most probable amount of bladder distension. A minimum bladder volume structure that incorporates the most likely distension would help treatment planners better prepare for the inter-fraction variability of the bladder.

Our program takes approximately 90 seconds to calculate DVH curves for CBCT OARs, of which Limbus contouring occupies just 15 seconds. There is potential for further reduction in computation time. Since Limbus auto-contouring is based on U-Net, it first classifies each pixel as OAR or not OAR, creating a boolean mask. Then, it converts the mask into contour data points. Our program must convert the contour points back into a mask in order to perform the boolean operation with the dose data that is necessary for DVH calculation. The conversion from boolean mask to contour data in Limbus, then from contour to mask in our program is a redundant procedure. If our program were able to have direct access to the Limbus mask, we estimate that computation time could be halved. However, a contouring expert would still need to be present at treatment as Limbus contours may need corrections prior to beam delivery, especially when the patient has a SpaceOAR gel implant. We hope that the future transition to SpaceOAR Vue will reduce the frequency of rectum contour corrections.

In this work we developed the foundation for an on-line decision support tool. We hope to implement our program into our bladder and rectum decision trees so that patient setup decisions can be quantitatively informed by expected delivered dose rather than the current clinical practice of visual evaluation. At the moment, our center has limited resources and we are unable to accommodate a fully online ART system (Section 2.6). With our program, we hope to aid in clinical decision making without significantly increasing workload or the patient's time on the treatment table.

With some further improvements to our code's computation time, we hope to move

towards a semi-ART solution. Instead of fully re-planning the patient every fraction, we could auto-calculate small 3D couch shifts to optimize patient setup in a way that reduces the frequency of dose violations. Due to the observed major violations of the bladder V37 Gy and rectum V36 Gy, which occurred when the OAR was distended towards the target volume compared to the pCT, the ability to make a couch shift according to the estimated delivered dose would likely reduce OAR dose. While this may not be as dosimetrically effective as a full ART system, this is something our center could feasibly implement in the next year without extending a patient's time on the treatment couch or requiring extra hospital funds.

5.2 Conclusion

Bladder and rectum dose-volume metrics at treatment were higher than were anticipated during treatment planning. Relative metrics followed an inverse relationship with OAR volume at treatment, whereas absolute metrics did not. Bladder and rectum volume varied greatly during treatment despite the use of rigorous preparation protocols. The degree of volume variation was likely the cause of metric violation incidents. Our results motivate future work towards adaptive radiotherapy techniques. Our current population-based qualitative decision trees are not sufficient for prostate SABR patients where the high dose per fraction incurs greater risk of toxicity. Radiation therapy techniques that are able to adapt to the patient's anatomy as presented for treatment will likely result in lower DVH curves for OARs at treatment. We hope to use our system to calculate DVH curves for OARs in prostate treatment, but also other treatment sites provided that they can be contoured by Limbus and the majority of tissues are near water equivalent. This would also subject to site-specific investigations of superimposing the planned dose on to the CBCT to see if it is sufficiently accurate for assessing patient setup. This would allow for quantitative dose metrics to better inform clinicians and replace the current visual estimation used in on-line decision making.

Bibliography

- [1] Canadian Cancer Society. Canadian cancer statistics: a 2022 special report on cancer prevalence.
- [2] Darren R Brenner, Abbey Poirier, Ryan R Woods, Larry F Ellison, Jean-Michel Billette, Alain A Demers, Shary Xinyu Zhang, Chunhe Yao, Christian Finley, Natalie Fitzgerald, et al. Projected estimates of cancer in Canada in 2022. *CMAJ*, 194(17):E601–E607, 2022.
- [3] L Xie, R Semenciw, and L Mery. Cancer incidence in Canada: trends and projections (1983–2032). *Health promotion and chronic disease prevention in Canada: research, policy and practice*, 35(Suppl 1):1, 2015.
- [4] Faiz M Khan and John P Gibbons. *Khan's the physics of radiation therapy*. Lippincott Williams & Wilkins, 2014.
- [5] Eva Bezak, Alun H Beddoe, Loredana G Marcu, Martin Ebert, and Roger Price. *Johns and Cunningham's the Physics of Radiology*. Charles C Thomas Publisher, 2021.
- [6] Penelope J Allisy-Roberts and Jerry Williams. *Farr's physics for medical imaging*. Elsevier Health Sciences, 2007.
- [7] Albert C. Tompson, Doug Vaughan, and et al. *X-Ray Data Booklet*. Center for X-Ray Optics and Advanced Light Source, Sep 2009.
- [8] Lynn J. Verhey and Paula L. Petti. Principles of Radiation Physics. In Richard T. Hoppe, Theodore Locke Phillips, and Mack Roach, editors, *Leibel and Phillips Textbook of Radiation Oncology (Third Edition)*, chapter 7, pages 95–119. W.B. Saunders, Philadelphia, third edition edition, 2010.

- [9] Faiz M Khan, John Gibbons, Dimitris Mihailidis, and Hassaan Alkhatib. *Khan's lectures: handbook of the physics of radiation therapy*. Lippincott Williams & Wilkins, 2012.
- [10] Frank Herbert Attix. *Introduction to radiological physics and radiation dosimetry*. John Wiley & Sons, 2008.
- [11] Jerrold T Bushberg and John M Boone. *The essential physics of medical imaging*. Lippincott Williams & Wilkins, 2011.
- [12] Jiang Hsieh. *Computed tomography: principles, design, artifacts, and recent advances*. SPIE press, 2003.
- [13] Theobald Fuchs, Marc Kachelrieß, and Willi A Kalender. Direct comparison of a xenon and a solid-state CT detector system: measurements under working conditions. *IEEE transactions on medical imaging*, 19(9):941–948, 2000.
- [14] K. Iniewski. *Advanced X-ray Detector Technologies: Design and Applications*. Springer International Publishing, 2022.
- [15] Steven W Smith et al. *The scientist and engineer's guide to digital signal processing*. California Technical Pub. San Diego, 1997.
- [16] R Schofield, L King, U Tayal, I Castellano, J Stirrup, F Pontana, James Earls, and E Nicol. Image reconstruction: Part 1—understanding filtered back projection, noise and image acquisition. *Journal of cardiovascular computed tomography*, 14(3):219–225, 2020.
- [17] Lee A Feldkamp, Lloyd C Davis, and James W Kress. Practical cone-beam algorithm. *Josa a*, 1(6):612–619, 1984.
- [18] Julia Garayoa and Pablo Castro. A study on image quality provided by a kilovoltage cone-beam computed tomography. *Journal of Applied Clinical Medical Physics*, 14(1):239–257, 2013.
- [19] Lawrence Lechuga and Georg A Weidlich. Cone beam CT vs. fan beam CT: a comparison of image quality and dose delivered between two differing CT imaging modalities. *Cureus*, 8(9), 2016.

- [20] Parham Alaei and Emiliano Spezi. Imaging dose from cone beam computed tomography in radiation therapy. *Physica Medica*, 31(7):647–658, 2015.
- [21] Ying Xiao, Stephen F Kry, Richard Popple, Ellen Yorke, Niko Papanikolaou, Sotirios Stathakis, Ping Xia, Saiful Huq, John Bayouth, James Galvin, et al. Flattening filter-free accelerators: a report from the AAPM therapy emerging technology assessment work group. *Journal of applied clinical medical physics*, 16(3):12–29, 2015.
- [22] Chaoqiong Ma, Mingli Chen, Troy Long, David Parsons, Xuejun Gu, Steve Jiang, Qing Hou, and Weiguo Lu. Flattening filter free in intensity-modulated radiotherapy (IMRT)—Theoretical modeling with delivery efficiency analysis. *Medical physics*, 46(1):34–44, 2019.
- [23] Douglas Jones. ICRU report 50—prescribing, recording and reporting photon beam therapy, 1994.
- [24] Philip Mayles, Alan Nahum, and Jean-Claude Rosenwald. *Handbook of radiotherapy physics: theory and practice*. CRC Press, 2007.
- [25] Erika Ming Yee Chin. *A four dimensional volumetric modulated arc therapy planning system for stereotactic body radiation therapy in lung cancers*. PhD thesis, University of British Columbia, 2013.
- [26] Philip Mayles; Alan E. Nahum; J. C. Rosenwald. *Handbook of Radiotherapy Physics*. CRC Press, 2 edition, 2019.
- [27] Karl Otto. Volumetric modulated arc therapy: IMRT in a single gantry arc. *Medical physics*, 35(1):310–317, 2008.
- [28] Faiz M Khan, Paul W Sperduto, and John P Gibbons. *Khan’s Treatment Planning in Radiation Oncology*:. Lippincott Williams & Wilkins, 2021.
- [29] Janne Sievinen, Waldemar Ulmer, Wolfgang Kaissl, et al. AAA photon dose calculation model in Eclipse. *Palo Alto (CA): Varian Medical Systems*, 118:2894, 2005.
- [30] T Han, F Mourtada, K Kisling, J Mikell, D Followill, and R Howell. SU-E-T-720: Dosimetric verification of deterministic Acuros XB radiation transport

- algorithm for IMRT and VMAT plans with the RPC H&N phantom. *Medical Physics*, 38(6Part22):3656–3656, 2011.
- [31] Antonella Bufacchi, Orietta Caspiani, Giulia Rambaldi, Luca Marmiroli, Giuseppe Giovinazzo, and Mattia Polsoni. Clinical implication in the use of the AAA algorithm versus the AXB in nasopharyngeal carcinomas by comparison of TCP and NTCP values. *Radiation Oncology*, 15(1):1–8, 2020.
- [32] Sriram Padmanaban, Samantha Warren, Anthony Walsh, Mike Partridge, and Maria A Hawkins. Comparison of Acuros (AXB) and Anisotropic Analytical Algorithm (AAA) for dose calculation in treatment of oesophageal cancer: effects on modelling tumour control probability. *Radiation Oncology*, 9(1):1–6, 2014.
- [33] Angela GM O’Neill, Suneil Jain, Alan R Hounsell, and Joe M O’Sullivan. Fiducial marker guided prostate radiotherapy: a review. *The British journal of radiology*, 89(1068):20160296, 2016.
- [34] David Skarsgard, Pat Cadman, Ali El-Gayed, Robert Pearcey, Patricia Tai, Nadeem Pervez, and Jackson Wu. Planning target volume margins for prostate radiotherapy using daily electronic portal imaging and implanted fiducial markers. *Radiation oncology*, 5(1):1–11, 2010.
- [35] Michael J Zelefsky, Marisa Kollmeier, Brett Cox, Anthony Fidaleo, Dahlia Sperling, Xin Pei, Brett Carver, Jonathan Coleman, Michael Lovelock, and Margie Hunt. Improved clinical outcomes with high-dose image guided radiotherapy compared with non-IGRT for the treatment of clinically localized prostate cancer. *International Journal of Radiation Oncology, Biology, Physics*, 84(1):125–129, 2012.
- [36] Kavitha Srinivasan, Mohammad Mohammadi, and Justin Shepherd. Applications of linac-mounted kilovoltage cone-beam computed tomography in modern radiation therapy: A review. *Polish journal of radiology*, 79:181, 2014.
- [37] Kristy K Brock. Adaptive radiotherapy: moving into the future. In *Seminars in radiation oncology*, volume 29, page 181. NIH Public Access, 2019.
- [38] Valentina Giacometti, Raymond B King, Christina E Agnew, Denise M Irvine, Suneil Jain, Alan R Hounsell, and Conor K McGarry. An evaluation of techniques

- for dose calculation on cone beam computed tomography. *The British journal of radiology*, 92(1096):20180383, 2019.
- [39] Michael Velec, Joanne L Moseley, Tim Craig, Laura A Dawson, and Kristy K Brock. Accumulated dose in liver stereotactic body radiotherapy: positioning, breathing, and deformation effects. *International Journal of Radiation Oncology* Biology* Physics*, 83(4):1132–1140, 2012.
- [40] Molly M McCulloch, Daniel G Muenz, Matthew J Schipper, Michael Velec, Laura A Dawson, and Kristy K Brock. A simulation study to assess the potential impact of developing normal tissue complication probability models with accumulated dose. *Advances in Radiation Oncology*, 3(4):662–672, 2018.
- [41] Mikel Byrne, Ben Archibald-Heeren, Yunfei Hu, Amy Teh, Rhea Beserminji, Emma Cai, Guilin Liu, Angela Yates, James Rijken, Nick Collett, et al. Varian ethos online adaptive radiotherapy for prostate cancer: Early results of contouring accuracy, treatment plan quality, and treatment time. *Journal of Applied Clinical Medical Physics*, 23(1):e13479, 2022.
- [42] Ergun E Ahunbay, Cheng Peng, Shannon Holmes, Andrew Godley, Colleen Lawton, and X Allen Li. Online adaptive replanning method for prostate radiotherapy. *International Journal of Radiation Oncology* Biology* Physics*, 77(5):1561–1572, 2010.
- [43] Dzung L Pham, Chenyang Xu, and Jerry L Prince. Current methods in medical image segmentation. *Annual review of biomedical engineering*, 2(1):315–337, 2000.
- [44] Pei-Gee Ho. *Image segmentation*. BoD–Books on Demand, 2011.
- [45] W Jeffrey Zabel, Jessica L Conway, Adam Gladwish, Julia Skliarenko, Giulio Didiodato, Leah Goorts-Matthews, Adam Michalak, Sarah Reistetter, Jenna King, Keith Nakonechny, et al. Clinical evaluation of deep learning and atlas-based auto-contouring of bladder and rectum for prostate radiation therapy. *Practical Radiation Oncology*, 11(1):e80–e89, 2021.
- [46] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation, 2015.

- [47] Michael C Joiner and Albert J van der Kogel. *Basic clinical radiobiology*. CRC press, 2018.
- [48] Eric J Hall, Amato J Giaccia, et al. *Radiobiology for the Radiologist*, volume 6. Philadelphia, 2006.
- [49] David J Brenner and Eric J Hall. Hypofractionation in prostate cancer radiotherapy. *Transl Cancer Res*, 7(Suppl 6):S632–9, 2018.
- [50] JM Cosset, C Chargari, and G Créhange. Which alpha/beta ratio for prostate cancer in 2019? *Cancer Radiotherapie: Journal de la Societe Francaise de Radiotherapie Oncologique*, 23(4):342–345, 2019.
- [51] Raymond Miralbell, Stephen A Roberts, Eduardo Zubizarreta, and Jolyon H Hendry. Dose-fractionation sensitivity of prostate cancer deduced from radiotherapy outcomes of 5,969 patients in seven international institutional datasets: $\alpha/\beta= 1.4$ (0.9–2.2) Gy. *International Journal of Radiation Oncology* Biology* Physics*, 82(1):e17–e24, 2012.
- [52] David J Brenner and Eric J Hall. Fractionation and protraction for radiotherapy of prostate carcinoma. *International Journal of Radiation Oncology* Biology* Physics*, 43(5):1095–1101, 1999.
- [53] Anders Widmark, Adalsteinn Gunnlaugsson, Lars Beckman, Camilla Thellenberg-Karlsson, Morten Hoyer, Magnus Lagerlund, Jon Kindblom, Claes Ginman, Bengt Johansson, Kirsten Björnlinger, et al. Ultra-hypofractionated versus conventionally fractionated radiotherapy for prostate cancer: 5-year outcomes of the HYPO-RT-PC randomised, non-inferiority, phase 3 trial. *The Lancet*, 394(10196):385–395, 2019.
- [54] Winkle Kwan, Gaurav Bahl, David Kim, Allison Ye, Isabelle Gagne, and Abraham Alexander. Acute toxicity of ultrahypofractionation compared with moderate hypofractionation in prostate cancer treatment: A randomized trial. *International Journal of Radiation Oncology* Biology* Physics*, 113(5):1036–1043, 2022.
- [55] Ting Martin Ma, Oscar Lilleby, Wolfgang A Lilleby, and Amar U Kishan. Ablative radiotherapy in prostate cancer: stereotactic body radiotherapy and high dose rate brachytherapy. *Cancers*, 12(12):3606, 2020.

- [56] Christopher R King. LDR vs. HDR brachytherapy for localized prostate cancer: the view from radiobiological models. *Brachytherapy*, 1(4):219–226, 2002.
- [57] Trevor J Royce, Panayiotis Mavroidis, Kyle Wang, Aaron D Falchook, Nathan C Sheets, Donald B Fuller, Sean P Collins, Issam El Naqa, Daniel Y Song, George X Ding, et al. Tumor control probability modeling and systematic review of the literature of stereotactic body radiation therapy for prostate cancer. *International Journal of Radiation Oncology* Biology* Physics*, 110(1):227–236, 2021.
- [58] Lynsey Devlin, David Dodds, Azmat Sadozye, Philip McLoone, Nicholas MacLeod, Carolyn Lamb, Suzanne Currie, Stefanie Thomson, and Aileen Duffton. Dosimetric impact of organ at risk daily variation during prostate stereotactic ablative radiotherapy. *The British journal of radiology*, 93(1108):20190789, 2020.
- [59] Nigel Armstrong, Amit Bahl, Michael Pinkawa, Steve Ryder, Charlotte Ahmadu, Janine Ross, Samir Bhattacharyya, Emily Woodward, Suzanne Battaglia, Jean Binns, et al. SpaceOAR hydrogel spacer for reducing radiation toxicity during radiotherapy for prostate cancer. a systematic review. *Urology*, 156:e74–e85, 2021.
- [60] Abraham Alexander and Winkle Kwan. Androgen suppression with stereotactic body or external beam radiation therapy (ASSERT), Nov 2015.
- [61] Satvik R Hadigal and Atul K Gupta. Application of hydrogel spacer SpaceOAR Vue for prostate radiotherapy. *Tomography*, 8(6):2648–2661, 2022.
- [62] Sherry A Leeper, F Yin, and S Yoo. Evaluation of in-room cone-beam CT imaging techniques compared to CT for dose calculation. *International Journal of Radiation Oncology, Biology, Physics*, 75(3):S670–S671, 2009.
- [63] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa,

- Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020.
- [64] J. D. Hunter. Matplotlib: A 2D graphics environment. *Computing in Science & Engineering*, 9(3):90–95, 2007.
- [65] Simon Biggs, Matthew Jennings, Stuart Swerdloff, Phillip Chlap, Derek Lane, Jacob Rembish, Jacob McAloney, Paul King, Rafael Ayala, Fada Guan, et al. PyMedPhys: A community effort to develop an open, Python-based standard library for medical physics applications. *Journal of Open Source Software*, 7(78):4555, 2022.
- [66] Moyed Miften, Arthur Olch, Dimitris Mihailidis, Jean Moran, Todd Pawlicki, Andrea Molineu, Harold Li, Krishni Wijesooriya, Jie Shi, Ping Xia, et al. Tolerance limits and methodologies for imrt measurement-based verification qa: recommendations of aapm task group no. 218. *Medical physics*, 45(4):e53–e83, 2018.
- [67] Kyle Wang, Panayiotis Mavroidis, Trevor J Royce, Aaron D Falchook, Sean P Collins, Stephen Sapareto, Nathan C Sheets, Donald B Fuller, Issam El Naqa, Ellen Yorke, et al. Prostate stereotactic body radiation therapy: an overview of toxicity and dose response. *International Journal of Radiation Oncology* Biology* Physics*, 110(1):237–248, 2021.
- [68] Lynsey Devlin, David Dodds, Azmat Sadozye, Philip McLoone, Nicholas MacLeod, Carolyn Lamb, Suzanne Currie, Stefanie Thomson, and Aileen Duffton. Dosimetric impact of organ at risk daily variation during prostate stereotactic ablative radiotherapy. *The British journal of radiology*, 93(1108):20190789, 2020.
- [69] Michael Wahl, Martina Descovich, Erin Shugard, Dilini Pinnaduwege, Atchar Sudhyadhom, Albert Chang, Mack Roach, Alexander Gottschalk, and Josephine Chen. Interfraction anatomical variability can lead to significantly increased rectal dose for patients undergoing stereotactic body radiotherapy for prostate cancer. *Technology in Cancer Research & Treatment*, 16(2):178–187, 2017.
- [70] Frederik Fuchs, Gregor Habl, Michal Devečka, Severin Kampfer, Stephanie E Combs, and Kerstin A Kessel. Interfraction variation and dosimetric changes during image-guided radiation therapy in prostate cancer patients. *Radiation oncology journal*, 37(2):127, 2019.

- [71] S Tanny, NN Sperling, D Zheng, O Dona Lemus, N Joyce, HC Qiu Jr, and MA Cummings. Tracking OAR Volume and DVH variability in the initial cohort of pelvic patients treated with CBCT-guided, online adaptive therapy. *International Journal of Radiation Oncology, Biology, Physics*, 114(3):e598, 2022.
- [72] Seungjong Oh and Siyong Kim. Deformable image registration in radiation therapy. *Radiation oncology journal*, 35(2):101, 2017.
- [73] Haonan Xiao, Ge Ren, and Jing Cai. A review on 3D deformable image registration and its application in dose warping. *Radiation Medicine and Protection*, 1(04):171–178, 2020.

Appendix A

Original Decision Trees

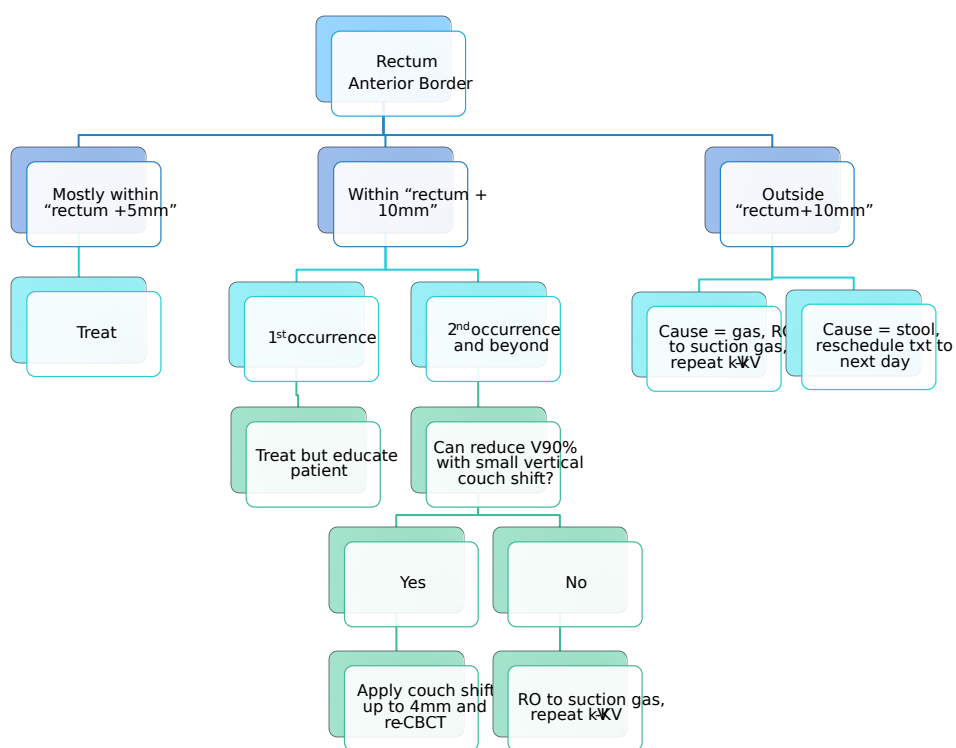


Figure A.1: ASSERT Trial Rectum Decision Tree.

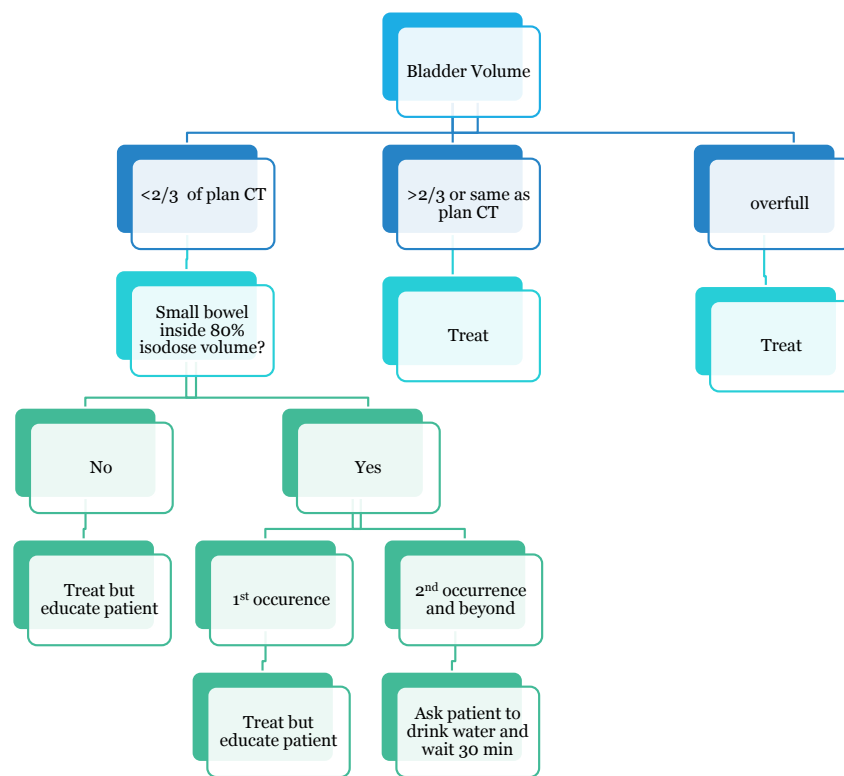


Figure A.2: ASSERT Trial Bladder Decision Tree.

Appendix B

Dose Shift Study




Figure B.1: Dose difference and gamma index plots for slice near the isocenter for Case Ia

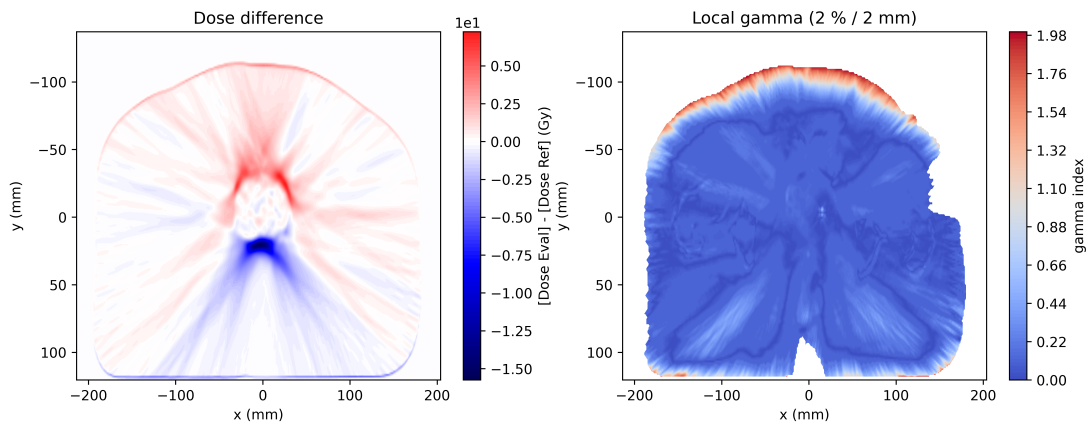


Figure B.2: Dose difference and gamma index plots for slice near the isocenter for Case Ib




Figure B.3: Dose difference and gamma index plots for slice near the isocenter for Case II

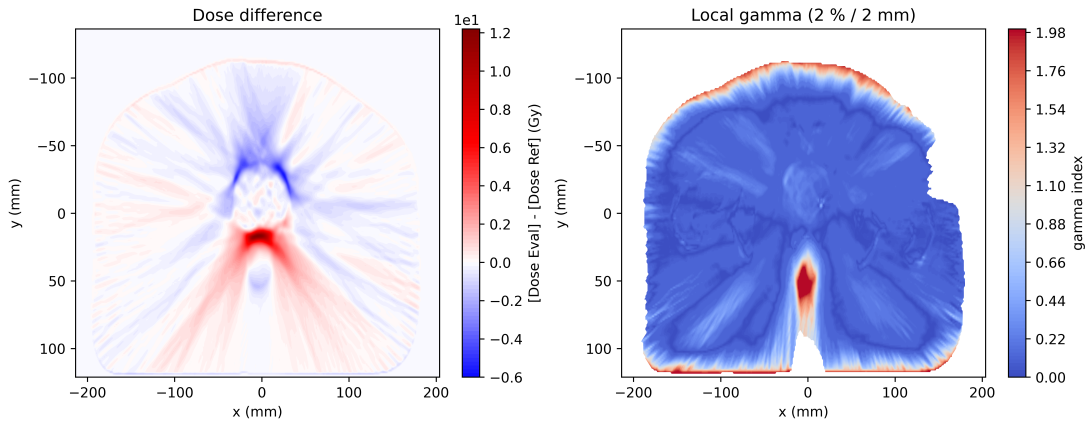


Figure B.4: Dose difference and gamma index plots for slice near the isocenter for Case IIa

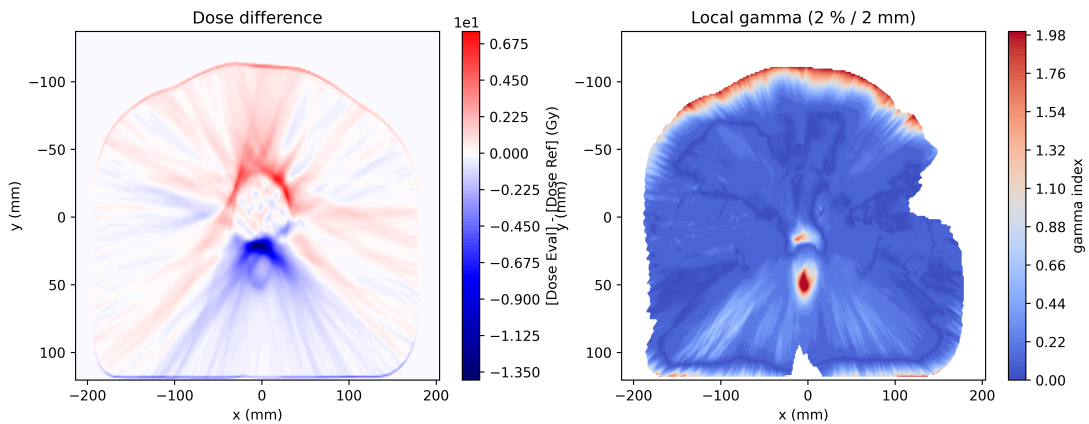


Figure B.5: Dose difference and gamma index plots for slice near the isocenter for Case IIb