Integer Programs for High Dose Rate Brachytherapy Needle and Dose Planning that Directly Optimize Clinical Objectives

by

Ko-Ay Timmy Siauw

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Engineering — Civil and Environmental Engineering

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Ken Y. Goldberg, Co-chair
Associate Professor Alexandre M. Bayen, Co-chair
Professor Filip C. Filippou
Professor Alper Atamtürk

Fall 2012

Integer Programs for High Dose Rate Brachytherapy Needle and Dose Planning that Directly Optimize Clinical Objectives

Abstract

Integer Programs for High Dose Rate Brachytherapy Needle and Dose Planning that Directly Optimize Clinical Objectives

by

Ko-Ay Timmy Siauw

Doctor of Philosophy in Engineering — Civil and Environmental Engineering

University of California, Berkeley

Professor Ken Y. Goldberg, Co-chair

Associate Professor Alexandre M. Bayen, Co-chair

High dose rate (HDR) brachytherapy is a radiation therapy for cancer in the prostate, cervix, breast, head, and neck, including other sites. In HDR brachytherapy, hollow needles are inserted or placed near the cancer site. Radiation is delivered to the patient by a radioactive source which is sequentially threaded through the needles. The dose distribution is controlled by altering the dwell times, the time spent at pre-defined positions on the needles.

HDR brachytherapy has a 90% cancer-free survival rate at 12 years when used for the treatment of prostate cancer, the focus of this dissertation. However, it can have serious negative side effects such as impotence and incontinence, which are caused by excess radiation exposure and needle puncture of healthy organs near the prostate, or organs at risk (OAR). A major goal of the field is to reduce side effects of HDR brachytherapy without compromising its therapeutic effectiveness. Towards this goal, this dissertation seeks to use mathematical optimization techniques to compute radiation dose distributions which meet clinical objectives and needle configurations which induce less trauma in the patient. We develop planning tools that directly optimize the dose distributions towards the RTOG-0321 standard dose objectives set by the Radiation Therapy Oncology Group and needle configurations which avoid puncturing OAR and use fewer needles than common practice. Specifically, this dissertation makes the following contributions.

Contributions:

1. We developed Inverse Planning by Integer Program (IPIP), the first integer program which directly optimizes dosimetric indices, the standard metrics used to evaluate HDR brachytherapy dose distributions. However, we showed that for anatomy data taken from patients previously treated at the UCSF clinic and the RTOG-0321 dose objectives, CPLEX could not solve IPIP within 30 minutes of computing time using its default parameters.

2. We developed a heuristic algorithm, IPIP-H, which uses two linear programs to compute feasible solutions for IPIP. Thus, it is a polynomial-time heuristic algorithm for IPIP. We used IPIP-H to compute dose plans for the same patients as IPIP. We showed that IPIP-H could compute a dose plan for each patient which met all the dose objectives specified by the RTOG-0321 protocol in less than 30 seconds of computing time (avg. 13 seconds). The solutions computed from IPIP-H were always feasible for IPIP and were within 5% of the optimal solution. We compared IPIP-H to Inverse Planning Simulated Annealing (IPSA), a dose planning model which is clinically deployed and has been used worldwide for over a decade. IPSA was not able to compute a dose plan which met all the dose objectives for any of the patients in our data set using its standard class solution. Therefore, IPSA would require iterations of manual fine tuning of its optimization parameters until a feasible dose plan was found. IPIP-H would not require iteration.

3. We formulated the problem of positioning HDR brachytherapy needles as a spatial coverage problem: given a large candidate set of needles for insertion, anatomy data, and a user parameter, δ, find the smallest candidate needle subset such that the minimum distance between any point in the prostate and a needle in the chosen set is less than δ. We showed that this problem could be represented as a set cover integer program.

4. We developed Needle Planning by Integer Program (NPIP), an algorithm which generates a set of candidate needles represented by skew-line segments, solves an integer program which chooses a candidate needle subset that covers the prostate according to the user-parameter, δ, and verifies that the final needle configuration meets dose objectives by computing a dose plan for it using IPIP. NPIP uses a candidate needle set which is approximately 10 times larger than considered with Hyrbid Inverse Planning Optimization (HIPO), the only other fully computerized needle planning system for HDR brachytherapy known to us. By construction, NPIP avoids choosing needles which penetrate OAR and needles which collide with each other. We used NPIP to compute needle configurations for patients previously treated at the UCSF clinic and compared the computed needle configurations to those implanted by the physician. NPIP could find needle configurations which met the RTOG-0321 dose objectives and used 10 or fewer needles; the physician used 16 needles. NPIP always computed a needle configuration that avoided puncturing the penile bulb; the average number of punctures made by the physician was 5. NPIP required an average of 5 minutes of computing time, but there was a wide range of run times, up to almost one hour. We also conducted a sensitivity analysis of NPIP-generated needle configurations to placement errors on the order expected from current needle insertion robots, which was about 2 mm. We showed that, although dose objectives could be met with 10 or fewer needles, 16 needles were required to meet dose objectives robustly.

5. We designed and implemented the first end-to-end robotic HDR brachytherapy experiment. Our experiment utilized Contributions 1 through 4, and Acubot-RND, a needle insertion robot specialized for needle insertion. We planned and executed NPIP-generated needle configurations in a fully equipped brachytherapy environment on two anatomically-

correct gelatin phantoms. There were non-trivial placement errors between the planned needle configuration and the implanted needle configuration. We separated the error into systematic error and random error. We computed the systematic error as the rigid least squares fit between points regularly sampled along the needles in the planned and actual needle configuration. The total RMS error between the planned and actual needle configuration was 3 mm for the first phantom and 5 mm for the second phantom. We computed the random error as the total RMS error between the planned and actual needle configuration after the systematic error was removed. The random error was 1.4 mm for the first phantom and 2.5 mm for the second phantom. Our random errors were close to the placement error of current needle insertion robots which have a more sophisticated calibration device. Although there were discrepancies between the planned and actual needle configuration, we showed that our end-to-end robotic experiment could execute the planned needle configurations with sufficient accuracy to meet the RTOG-0321 dose objectives and avoid puncturing OAR. We compared the needle configurations executed by our robotic workflow with a needle configuration executed by a world-class brachytherapist, who also used 16 needles, met dose objectives and avoided puncturing OAR. Therefore, the needle configurations executed in our experiment are comparable to an expert physician.

In summary, this dissertation has developed mathematical methods which improve the planning of HDR brachytherapy dose distributions and needle configurations. Dose distributions can be directly optimized towards the standard RTOG-0321 dosimetric protocol, or other dose objectives based on constraining dosimetric indices, and needle configurations can be computed which meet dose objectives, use fewer needles than standard practice, and avoid puncturing OAR. We have demonstrated the feasibility of using IPIP and NPIP in a clinical environment using a robotic clinical workflow. These planning methods are a significant step towards reducing side effect of brachytherapy. We leave a clinical translation of these tools to determine if, and the extent, side effects are actually reduced.

This dissertation is dedicated to my mom and dad,

the best parents anyone could hope for.

Acknowledgments

First, I thank my parents for their love, support, and encouragement throughout my life. All that I have accomplished I owe to them.

I thank my advisors Professor Ken Goldberg for his patience and diligence in training me as a scholar and for his constant sense of urgency that has kept me moving towards graduating, Professor Jean Pouliot for making me feel welcome at the Department of Radiation Oncology at the UCSF Mt. Zion campus and giving me the resources and first-hand experiences that have allowed me to succeed as a researcher in this field, Professor Alper Atamtürk for teaching me the fundamental mathematical tools that have been the foundation of my research, and Professor Alex Bayen for all our wonderful teaching experiences together and always "having my back".

I thank Professor Filip Filippou for mentoring me as an undergraduate and teaching me that excellence can be achieved if we improve ourselves a little bit every day. I also thank Professor Bozihidar Stoijanovic for introducing me to research and Professor Shaofan Lee for showing me the value of a patient and methodical mind. I am especially grateful to Professor Adam Cunha, a friend and mentor, for taking me under his wing at UCSF, for teaching me the in's and out's of brachytherapy, and for the whiteboard conversations we have had which have spawned and refined so many research ideas. I also thank Dr. I-Chow (Joe) Hsu for allowing me to be a guest at his brachytherapy clinic where I have, quite literally, learned from the best in the field.

I thank Professors Ron Alterovitz and Kris Hauser for supervising my early research while post-doctoral scholars at Berkeley. I am also especially thankful to Professor Dmitry Berenson for being both a friend and scholarly example while at Berkeley.

I thank Joan Chamberlain for her personal and professional support, especially during the many semesters I was a teaching assistant for E7. I also thank Shelley Okimoto for keeping me on track to graduate and for her patience helping me navigate the administrative maze of being a researcher at both UC Berkeley and UCSF.

I thank my fellow labmates in the Berkeley Laboratory for Automation Sciences and Engineering for their constant input, feedback, and encouragement throughout my graduate career. I especially thank Judy Hoffman, for giving me the kick in the head that led to my breakthrough in needle planning and Animesh Garg, for programming a robot to make my computational work come to life.

Finally, I thank my friends Mikey Gee, Andrew Tinka, Erik Keldrauk, Leah Anderson, Trucy Phan, Pranthik Samal, James Lew, Scott Payne, Claire Saint-Pierre, Anand Kulkarni, Ephrat Bitton, Stephanie Maroney, Rosalynn Chongchaikit, Sherry Chan, and Jenna Tower for making my years as a grad student enjoyable ones.

Contents

Contents **iii**

List of Figures **v**

List of Tables **vii**

1 Introduction **1**

2 Background **4**
- 2.1 Prostate High Dose Rate Brachytherapy 4
- 2.2 Inverse Planning 7
- 2.3 Computational Inverse Planning Tools 11
- 2.4 Mathematical Programming 12
  - 2.4.1 Linear Programming 13
  - 2.4.2 Integer Programming 14
  - 2.4.3 Mathematical Programming for Inverse Planning 16
- 2.5 Robot-Assisted Brachytherapy 17
- 2.6 Addressing Previous Work 17

3 Dose Planning with IPIP **19**
- 3.1 Overview 19
- 3.2 Background 19
- 3.3 Method and Materials 21
  - 3.3.1 Model Formulation 21
  - 3.3.2 IPIP Heuristic Formulation 24
  - 3.3.3 Patient Data Sets 27
  - 3.3.4 Dose Rate Calculations and Dose Objectives 28
  - 3.3.5 Method Evaluation 31
- 3.4 Results 32
- 3.5 Discussion 33
- 3.6 Conclusion 35

4 Skew-Line Needle Planning using NPIP 36
4.1 Overview 36
4.2 Background 37
4.3 Method and Materials 38
4.3.1 System Formulation 38
4.3.1.1 Candidate Needle Set Generation 38
4.3.1.2 Needle Selection 42
4.3.1.3 Dose Planning 44
4.3.2 Patient Data Sets 45
4.3.3 Method Evaluation 45
4.3.3.1 NPIP Evaluation 45
4.3.3.2 Sensitivity Analysis 46
4.4 Results 48
4.4.1 NPIP Results 48
4.4.2 Sensitivity Results 50
4.5 Discussion 55
4.6 Conclusion 56

5 Towards Robot-Assisted HDR Brachytherapy 57
5.1 Overview 57
5.2 Background 57
5.3 Method and Materials 58
5.3.1 Phantom 58
5.3.2 Planning 59
5.3.3 The Robot: Acubot-RND 60
5.3.4 Method Evaluation 63
5.3.4.1 Robotic Implant 63
5.3.4.2 Expert Implant 65
5.4 Results 65
5.4.1 Clinical Analysis 66
5.4.2 Placement Error Analysis 67
5.5 Discussion 70
5.6 Conclusion 71

6 Towards a Dose Planning Integer Program for Gamma Knife Perfexion 73
6.1 Overview 73
6.2 Background 73
6.3 Method and Materials 75
6.3.1 Model Formulation 75
6.3.2 Patient Data 77
6.3.3 Helmet Model 77
6.3.4 Beam Model 79
6.3.5 Method Evaluation 80

6.4 Results 81
6.5 Discussion 82
6.6 Conclusion 82

7 Conclusion 83
7.1 Contributions 83
7.2 Future Work 84
7.2.1 Dose Planning 84
7.2.2 Needle Planning 85
7.2.3 Robot-Assisted Brachytherapy 86
7.3 Closing Remarks 86

Bibliography 87

Appendices 99

A Dose Planning Appendix 100
A.1 Additional Numerical Results 100
A.2 Additional DVH Results 104

B Needle Planning Appendix 114
B.1 Additional Robustness Results 114

List of Figures

2.1 Regions of interest for prostate HDR brachytherapy 6
2.2 Example dose-volume histogram 9
2.3 Dose distribution heat map 10

3.1 Linear penalty function 21
3.2 IPIP heuristic intuition 26
3.3 Dose rate vs. distance approximation used for ^{192}Ir. 29

4.1 Method 1: Candidate needle set 40
4.2 Method 2: Random needle sampling 41
4.3 Method 2: False positive example 41
4.4 Method 2: Candidate needle set 42

4.5 NPIP failure case 44
4.6 NPIP needle configuration examples 49
4.7 Example actual vs. computed implant 51
4.8 Average number of needles vs. δ 52
4.9 Satisfactory and robust needle configurations 53
4.10 Robustness example 54

5.1 Prostate phantom 58
5.2 Candidate needle set 60
5.3 Acubot-RND 61
5.4 Acubot-RND orientation diagram 62
5.5 Post-implant photograph 66
5.6 Phantom 1: Planned vs actual needle placement 69
5.7 Phantom 2: Planned vs actual needle placement 69
5.8 Physician and Acubot-RND comparison 71

6.1 Tumor shapes 78
6.2 Helmet approximation 78
6.3 Beam approximation 80
6.4 Isodose for GK planner 82

A.1 IPIP DVH: Patient 1 105
A.2 IPIP DVH: Patient 2 105
A.3 IPIP DVH: Patient 3 106
A.4 IPIP DVH: Patient 4 106
A.5 IPIP DVH: Patient 5 107
A.6 IPIP DVH: Patient 6 107
A.7 IPIP DVH: Patient 7 108
A.8 IPIP DVH: Patient 8 108
A.9 IPIP DVH: Patient 9 109
A.10 IPIP DVH: Patient 10 109
A.11 IPIP DVH: Patient 11 110
A.12 IPIP DVH: Patient 12 110
A.13 IPIP DVH: Patient 13 111
A.14 IPIP DVH: Patient 14 111
A.15 IPIP DVH: Patient 15 112
A.16 IPIP DVH: Patient 16 112
A.17 IPIP DVH: Patient 17 113
A.18 IPIP DVH: Patient 18 113

B.1 ROBUSTNESS PLOT: Patient 1 115
B.2 ROBUSTNESS PLOT: Patient 2 115
B.3 ROBUSTNESS PLOT: Patient 3 116

B.4 ROBUSTNESS PLOT: Patient 4 116
B.5 ROBUSTNESS PLOT: Patient 5 117
B.6 ROBUSTNESS PLOT: Patient 6 117
B.7 ROBUSTNESS PLOT: Patient 7 118
B.8 ROBUSTNESS PLOT: Patient 8 118
B.9 ROBUSTNESS PLOT: Patient 9 119
B.10 ROBUSTNESS PLOT: Patient 10 119
B.11 ROBUSTNESS PLOT: Patient 11 120
B.12 ROBUSTNESS PLOT: Patient 12 120
B.13 ROBUSTNESS PLOT: Patient 13 121
B.14 ROBUSTNESS PLOT: Patient 14 121
B.15 ROBUSTNESS PLOT: Patient 15 122
B.16 ROBUSTNESS PLOT: Patient 16 122
B.17 ROBUSTNESS PLOT: Patient 17 123
B.18 ROBUSTNESS PLOT: Patient 18 123

List of Tables

2.1 RTOG-0321 dosimetric protocol 8

3.1 IPIP terms 22
3.2 Dose rate parameters 29
3.3 IPIP dosimetric criteria 30
3.4 IPIP experimental parameters 31
3.5 Median doses for IPIP, IPIP-R, IPIP-H, and IPSA 33

4.1 NPIP terms 38
4.2 NPIP experimental results 50

5.1 Clinical metrics results 67
5.2 RMS errors 70

6.1 Gamma Knife terms 75
6.2 Computational results 81

A.1 IPIP Study Patient 1 Dosimetric Results 101
A.2 IPIP Study Patient 2 Dosimetric Results 101

A.3 IPIP Study Patient 3 Dosimetric Results 101
A.4 IPIP Study Patient 4 Dosimetric Results 101
A.5 IPIP Study Patient 5 Dosimetric Results 101
A.6 IPIP Study Patient 6 Dosimetric Results 101
A.7 IPIP Study Patient 7 Dosimetric Results 102
A.8 IPIP Study Patient 8 Dosimetric Results 102
A.9 IPIP Study Patient 9 Dosimetric Results 102
A.10 IPIP Study Patient 10 Dosimetric Results 102
A.11 IPIP Study Patient 11 Dosimetric Results 102
A.12 IPIP Study Patient 12 Dosimetric Results 102
A.13 IPIP Study Patient 13 Dosimetric Results 103
A.14 IPIP Study Patient 14 Dosimetric Results 103
A.15 IPIP Study Patient 15 Dosimetric Results 103
A.16 IPIP Study Patient 16 Dosimetric Results 103
A.17 IPIP Study Patient 17 Dosimetric Results 103
A.18 IPIP Study Patient 18 Dosimetric Results 103

Chapter 1

Introduction

Radiation has been used to treat cancer for over 100 years [76] and is currently involved in approximately two-thirds of all cancer treatment regimens [6]. There are two major categories of radiation therapy: external beam therapy and brachytherapy. In external beam therapy, radiation is delivered to the tumor from a source outside the body, and in brachytherapy, the radiation source is delivered directly into the cancer site using temporarily inserted needles ("brachys" is Greek for "proximal"). Although the focus of this dissertation is on brachytherapy, external beam therapy and brachytherapy are often used in combination to treat cancer [9, 26, 30, 45, 47, 92]. There are also two types of brachytherapy: low dose rate (LDR) and high dose rate (HDR) brachytherapy. The two types of brachytherapy differ in the strength of the radioactive source used, low and high respectively, and the amount of time the radioactive source stays inside the body. In LDR brachytherapy, a set of radioactive sources are placed at the tumor site using hollow needles. After the sources are placed, the needles are removed, but the sources stay inside the body indefinitely. In HDR brachytherapy, a configuration of needles is placed at the tumor site, and a single radioactive source is sequentially threaded through the needles. Radiation, i.e. dose, is delivered to the tumor by temporarily halting the source at pre-specified dwell positions along the needles. After the procedure, the needles and the source are removed from the body.

Brachytherapy is used used to treat cancer in many areas of the body including the prostate [42,72,73], cervix [39,81,87], vagina [71,107], breast [4,97,114,115], lungs [8,44,65], and head and neck [43, 58, 98], and brachytherapy was performed on over 80,000 patients in 2004 [6]. In particular, brachytherapy has been shown to be an effective treatment for prostate cancer with a cancer-free survival rate over 90% at 12 years [113]. However HDR brachythreapy can have serious negative side effects. Several studies confirm that prostate brachytherapy can result in incontinence [29, 51] and impotence [34, 75]. These side-effects are a result of excessive radiation to the penile bulb [29,34,75], number of needles used [112] and needle penetration through sensitive structures near the penile bulb [78].

A major goal of prostate brachytherapy research is to reduce trauma and side effects without reducing its therapeutic effectiveness. Over the last 20 years, there has been a tremendous amount of effort in patient-specific dose planning [2, 18, 53, 62, 68, 104], needle planning [52,96], and improving the needle insertion process using robots [19,24,28,89]. This dissertation contributes to this body of work by developing a dose planning algorithm for directly optimizing dosimetric indices, which are the standard metrics used to define dose objectives for HDR brachytherapy [46, 48]. As a test bench, we always optimize towards the standard dose objectives, i.e. constraints on dosimetric indices, specified in the RTOG-0321 dosimetric protocol. However, our algorithm can handle other dose objectives based on constraining dosimetric indices. We also develop a needle planning algorithm for computing needle configurations that meet dose objectives, use fewer needles than standard practice, and avoid puncturing organs at risk. Finally, we demonstrate the use of these planning algorithms in the first robotic, needle-insertion experiment for HDR brachytherapy. The contributions of this dissertation towards reducing trauma and side effects of HDR brachytherapy are detailed below.

Contributions:

1. We developed Inverse Planning by Integer Program (IPIP), the first integer program which directly optimizes dosimetric indices. However, we showed that for anatomy data taken from patients previously treated at the UCSF clinic and the RTOG-0321 dose objectives, CPLEX [21], a general commercial linear and integer program solver, could not solve IPIP within 30 minutes of computing time using its default parameters.

2. We developed a heuristic algorithm, IPIP-H, which uses two linear programs to compute feasible solutions for IPIP. Thus, it is a polynomial-time heuristic algorithm for IPIP. We used IPIP-H to compute dose plans for the same patients as IPIP. We showed that IPIP-H could compute a dose plan for each patient which met all the dose objectives specified by the RTOG-0321 protocol in less than 30 seconds of computing time (avg. 13 seconds). The solutions computed from IPIP-H were always feasible for IPIP and were within 5% of the optimal solution. We compared IPIP-H to Inverse Planning Simulated Annealing (IPSA) [2, 68], a dose planning model which is clinically deployed and has been used worldwide for over a decade. IPSA was not able to compute a dose plan which met all the dose objectives for any of the patients using its standard class solution. Therefore, IPSA would require iterations of manual fine tuning of its optimization parameters until a feasible dose plan was found. IPIP-H would not require iteration.

3. We formulated the problem of positioning HDR brachytherapy needles as a spatial coverage problem: given a large candidate set of needles for insertion, anatomy data, and a user parameter, δ, find the smallest candidate needle subset such that the minimum distance between any point in the prostate and a needle in the chosen set is less than δ. We showed that this problem could be represented as a set cover integer program.

4. We developed Needle Planning by Integer Program (NPIP), an algorithm which generates a set of candidate needles represented by skew-line segments, solves an integer program which chooses a candidate needle subset that covers the prostate according to the user-parameter, δ, and verifies that the final needle configuration meets dose objectives by computing a dose plan for it using IPIP. NPIP uses a candidate needle set which is approximately 10 times larger than considered with Hyrbid Inverse Planning Optimization (HIPO) [52], the only other fully computerized needle planning system known to us. By construction, NPIP

avoids choosing needles which penetrate OAR and needles which collide with each other. We used NPIP to compute needle configurations for the same patients used for the IPIP-H study and compared the computed needle configurations to the ones implanted by the physician. NPIP could find needle configurations which met the RTOG-0321 dose objectives and used 10 or fewer needles; the physician used 16 needles. NPIP always computed a needle configuration that avoided puncturing the penile bulb; the average number of punctures made by the physician was 5. NPIP required an average of 5 minutes of computing time, but there was a wide range of run times, up to almost one hour. We also conducted a sensitivity analysis of NPIP-generated needle configurations to placement errors on the order expected from current needle insertion robots in a clinical environment, which was about 2 mm [80]. We showed that, although dose objectives could be met with 10 or fewer needles, 16 needles were required to meet dose objectives robustly.

5. We designed and implemented the first end-to-end robotic HDR brachytherapy experiment. Our experiment utilized Contributions 1 through 4, and Acubot-RND [108], a needle insertion robot specialized for needle insertion. We planned and executed NPIP-generated needle configurations in a fully equipped brachytherapy environment on two anatomically-correct gelatin phantoms. There were non-trivial placement errors between the planned needle configuration and the implanted needle configuration. We separated the error into systematic error and random error. We computed the systematic error as the rigid least squares fit between points regularly sampled along the needles in the planned and actual needle configuration [5]. The total RMS error between the planned and actual needle configuration was 3 mm for the first phantom and 5 mm for the second phantom. We computed the random error as the total RMS error between the planned and actual needle configuration after the systematic error was removed. The random error was 1.4 mm for the first phantom and 2.5 mm for the second phantom. Our random errors were close to the placement error of current needle insertion robots which have more sophisticated calibration devices [80], which is about 2 mm. Although there were discrepancies between the planned and actual needle configuration, we showed that our end-to-end robotic experiment could execute the planned needle configurations with sufficient accuracy to meet the RTOG-0321 dose objectives and avoid puncturing OAR. We compared the needle configurations executed by our robotic workflow with a needle configuration executed by a world-class brachytherapist, who also used 16 needles, met dose objectives and avoided puncturing OAR. Therefore, the needle configurations executed in our experiment are comparable to an expert physician.

This remainder of this dissertation is organized as follows. Chapter 2 gives relevant background for this dissertation including the brachytherapy clinical environment and workflow, treatment objectives, planning, and an overview of linear and integer programming. Chapter 3 presents our work on IPIP and IPIP-H, Chapter 4 presents our work on NPIP, and Chapter 5 presents our end-to-end robotic needle insertion experiment. Chapter 6 presents some initial work using developing treatment planning algorithms for a radiotherapy device called Gamma Knife Perfexion, which is used to treat tumors in the head. Chapter 7 concludes this dissertation and describes avenues of future work.

Chapter 2

Background

This chapter gives relevant background on brachytherapy, mathematical optimization and how it has been used for brachytherapy planning, robots that have been specifically designed to assist in brachytherapy needle insertion, and the limitations of current practices that are addressed in this dissertation. Since the focus of this dissertation is on prostate HDR brachytherapy, “brachytherapy” always refers to this type of brachythery unless otherwise stated.

2.1 Prostate High Dose Rate Brachytherapy

The standard clinical environment for prostate HDR brachytherapy includes an operating room, a transrectal ultrasound (TRUS) probe and associated hardware (i.e. monitor, computer, etc), hollow needles specialized for brachytherapy, a needle guide or template, a CT scanner, a computerized planning system, a radioactive source, and a pneumatic remote afterloading device which houses and moves the source. Needles are sometimes referred to as catheters when they are flexible, which is a useful property for producing non-straight needle trajectories.

The standard brachytherapy workflow consists of five main steps: (1) needle insertion, (2) anatomy digitization, (3) planning, (4) delivery, and (5) removal. These steps are described in the following paragraphs.

The procedure begins by taking the patient into an operating room. The patient undergoes general anesthesia, and a physician inserts 16-18 needles, under TRUS guidance, into the prostate through the perineum, the patch of skin between the testicles and the anus. Currently, needles are inserted with the assistance of a rigid template. A template is a rigid device which is in a locked orientation relative to the patient and constrains the movement of the needles during insertion. The needle movement is constrained by a grid of evenly-spaced, parallel holes which are drilled into the template. The tips of the needles can be observed in real time using TRUS. The holes to be used are selected according to the experience of the physician to produce a needle arrangement which has an even distribution of needles within the prostate volume and avoids puncture of nearby organs, especially the urethra which runs through the prostate. However, since the template only allows a limited set of parallel trajectories, sometimes puncturing organs, such as the sensitive structures near the penile bulb, is unavoidable.

A few clinics have developed freehand needle insertion techniques which do not require a template [57]. Freehand techniques allow the physician to produce skew-line (i.e. straight, non-parallel, non-intersecting) needle arrangements which have an even distribution of needles in the prostate and avoid puncturing nearby organs. Freehand techniques also use catheters rather than needles, which gives the physician additional freedom to control the arrangement of needles. There are clear advantages to using freehand techniques to reduce puncture trauma; however, they require years of training to master and are only taught by a limited set of physicians and clinics.

Once the needle arrangement is inserted into the patient, the patient anatomy and needles are CT scanned in parallel slices taken in the z-direction (from feet to head). The scans are sent to a computerized treatment planning station which provides a touchscreen-stylet interface for digitizing the patient anatomy and needles. In each CT image slice, the physician uses a stylet to contour the individual organs, or regions of interest. For prostate brachytherapy, the regions of interest include the prostate, urethra, bladder, and rectum. At the UCSF clinic, the penile bulb is also contoured. The physician also marks the location of the needles in each slice. Positions along each needle where a radioactive source can be halted, i.e. dwell positions, are interpolated by the planning system every 5 mm along the length of each needle, beginning at the tip. When the physician is finished, the digitized anatomy consists of a set of 3D coordinates for each of the regions of interest contoured by the physician, and a set of 3D coordinates denoting the dwell positions. From this point forward, the term “anatomy” will always refer to the digitized anatomy set unless otherwise stated. Figure 2.1 shows a digitized anatomy set.

The more time a radioactive source stays at a dwell position, the more dose it will deliver to surrounding tissue from that location. Thus the dose distribution can be controlled through the dwell times. Treatment planning can be summarized by the following problem statement. Given (1) an anatomy set, (2) dwell positions, (3) radioactive source properties, and (4) and treatment objectives, determine a dose plan, i.e. a set of source times spent at each dwell position, that will produce a radiation dose distribution which meets treatment objectives. In summary, treatment objectives generally consist of delivering a physician-specified “prescription dose” of radiation to the prostate while limiting the allowed dose to nearby healthy organs, or organs at risk (OAR). Here the prescription dose is given in units of centi-Gray (cGy), which is a unit of absorbed radiation. There have been a number of computerized treatment planning systems designed to help guide the user towards a satisfactory dose plan. A more complete description of treatment objectives and treatment planning systems is given in the following section.

Once a dose plan is computed, the needles are connected to a device called a remote afterloader, which executes the dwell times. The remote afterloader houses a radioactive source in a shielded compartment, and it can pneumatically drive the source through each

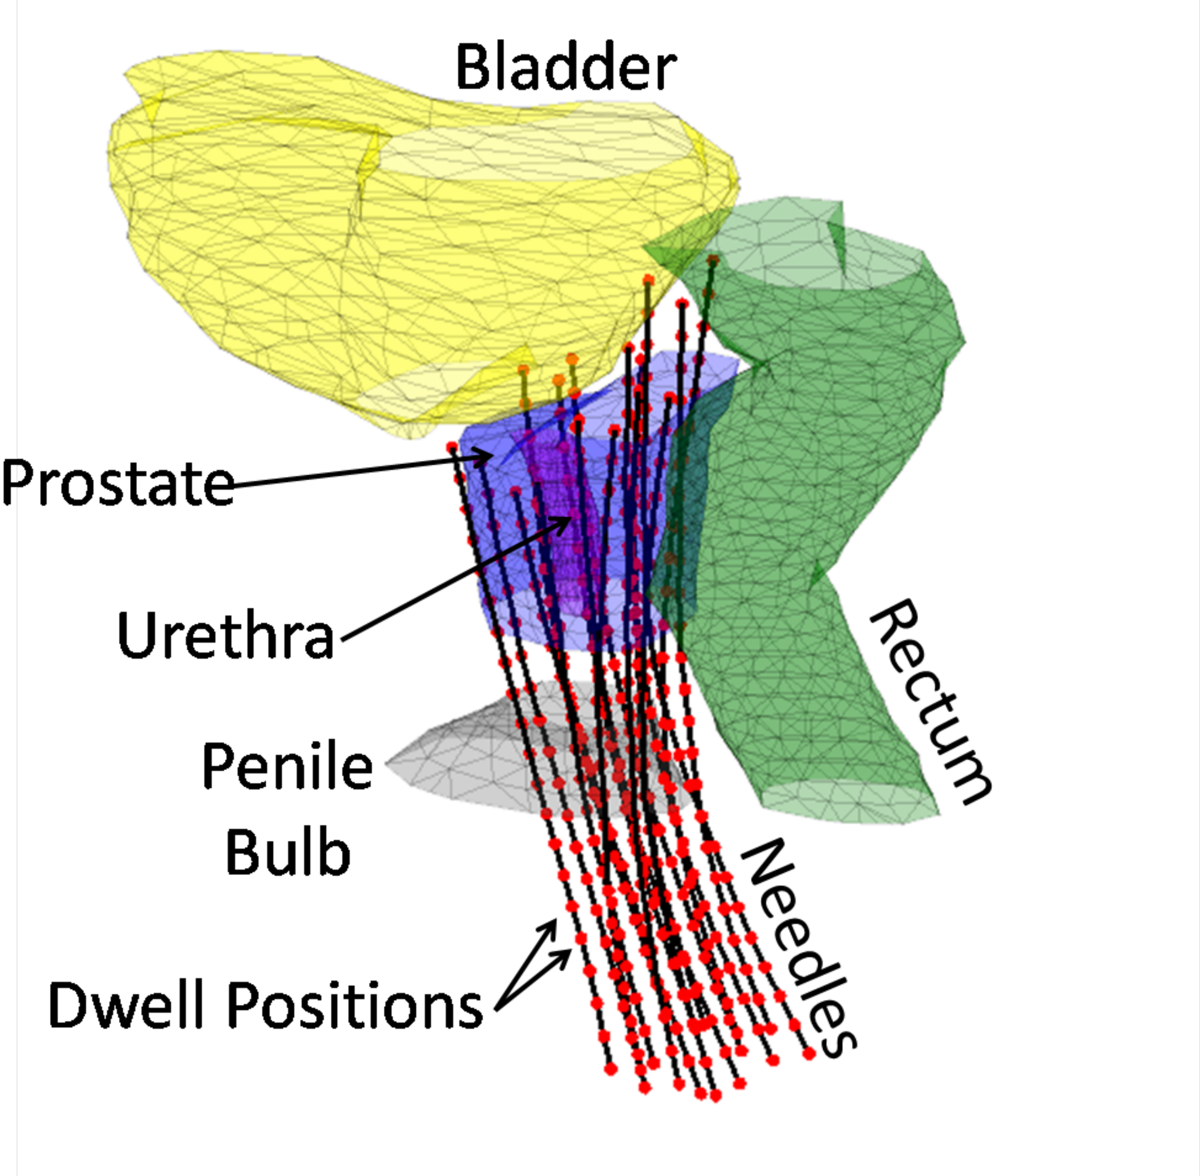

Figure 2.1: Shown are the regions of interest for prostate HDR brachytherapy: prostate, bladder, rectum, and urethra. Also shown is the penile bulb, which is contoured at the UCSF clinic. A set of implanted needles is also shown in cyan, with the dwell positions marked in magenta.

needle, halting the source at each dwell position for the specified dwell time. After the dwell times are executed, the source is returned to the afterloader, and the needles are removed from the patient.

Prostate HDR brachytherapy is sometimes fractionated [118], meaning the same dose plan is delivered to the patient multiple times with several hours of rest between deliveries. Between deliveries, the needles are disconnected from the afterloader but remain in the patient. After a few hours, the needles are reconnected to the afterloader and the dose plan is delivered again. Since healthy cells are known to have superior recovery potential to radiation exposure than cancer cells, fractionation allows healthy cells to recover during the rest period without decreasing the lethality of treatment to cancer cells. However, fractionation also consumes more time and resources. Although the work in this dissertation is generalizable to fractionated treatment, all treatments are assumed to be single fraction.

LDR brachytherapy for prostate cancer is referred to as prostate permanent-seed implant (PPI) brachytherapy. In PPI brachytherapy, the radioactive sources, or seeds, are inserted into the prostate where they are left indefinitely. The dose delivered to surrounding tissue is integrated over all the half-lives of the radioactive sources. Thus, the distribution is controlled by the configuration of seeds rather than the source time spent at dwell positions. In HDR brachytherapy, the source strength is much higher than the seeds used for PPI brachytherapy, and therefore only requires a short exposure period. The dwell times are sufficiently short compared to the half life of the source such that the radioactive properties of the material can be assumed constant during the treatment time.

HDR brachytherapy is an effective treatment for prostate cancer with over 90% survival rate at 12 years [113]. However it can also have serious negative side effects. Several studies confirm that prostate brachytherapy can result in incontinence [29,51] and impotence [34,75]. These side-effects are a result of excessive radiation to the penile bulb [29, 34, 75], number of needles used [112], and needle penetration through sensitive structures near the penile bulb [78].

This dissertation develops mathematical models for computing HDR brachytherapy dose distributions and needle configurations with the goal of reducing side effects and maintaining therapeutic effectiveness. The following two section gives relevant background to current methods of brachytherapy planning.

2.2 Inverse Planning

Before modern imaging devices such as magnetic resonance imaging (MRI) and computerized tomography (CT) in the early 1990's, only bony structures and needles were visible for planning. Since soft tissues, including the prostate, were not visible, computing dose distributions, i.e. dose planning, for both LDR and HDR brachytherapy focused on creating a homogeneous dose distribution in the region of the needles. In other words, since the exact position of the prostate was not known, the prostate region was blanketed with radiation, with special attention given to ensure that there was no volume receiving extremely

Index	Requirement
$V_{100}^{\text{Prostate}}$	$\geq 90\%$
V_{125}^{Urethra}	≤ 0.1 cm^3
V_{75}^{Rectum}	≤ 1 cm^3
V_{75}^{Bladder}	≤ 1 cm^3

Table 2.1: Dose restrictions according to the RTOG-0321 dosimetric protocol

high dose. Thus, the prostate was certain to be irradiated, but surrounding healthy organs also received significant dose. This type of dose planning is called forward planning. The improvement of medical imaging devices allowed for planning techniques which took into consideration the amount of radiation delivered to each organ - a type of planning called inverse planning. Inverse planning has superior patient outcome than forward planning [49] and is now standard in most clinics.

In inverse planning, it is necessary to define the dose delivered to an organ given a dose plan. For HDR brachytherapy, and many other radiotherapies, the dose delivered to an organ, G, from a dose plan is defined as the volume, V, receiving more than a threshold dose R. This aggregated dose is called a dosimetric index and is denoted by V_R^{G}. For each treatment, the physician specifies a prescription dose, R_x, which should be delivered to as much of the prostate volume as possible. It is customary for R to be given as a percentage of R_x when specifying dosimetric indices. For example, the prostate volume receiving at least 100% R_x is denoted by $V_{100}^{\text{Prostate}}$. This term is interchangeable with target coverage and is the dosimetric index most closely associated with therapeutic effect (i.e. killing the cancer).

The Radiation Therapy Oncology Group (RTOG) has determined a set of dosimetric index constraints which are correlated with long term treatment success [46,48]. The RTOG-0321 dosimetric protocol is given in Table 2.1. Note that dosimetric indices for the prostate are stated as a percentage of the total prostate volume, while dosimetric indices for OAR are specified in cm^3. The two main reasons for this discrepancy is that (1) the entire prostate volume is usually contoured, which is not the case for most OAR, and (2) the irradiated volume of OAR necessary to produce negative side effects does not scale with the size of the organ, and therefore absolute volume is a more appropriate metric for constraint. It is also important to mention that the results of Hsu et al., 2010 [46], only show that compliance with these dose objectives yields better results than not complying with these objectives. Specifically, there is no evidence to show that 95% coverage is better than 90% coverage with all other criteria being equal. Therefore, this dissertation focuses on compliance with dose objectives rather than the degree of compliance.

There are two qualitative tools for evaluating a dose distribution worth mentioning. The first dose evaluation tool is called a cumulative dose-volume histogram (DVH). For each organ, the x-axis of the associated DVH denotes dose, and on the y-axis is the organ volume receiving at least the dose. In other words, a DVH for an organ can be defined the set of pairs, (D, V_D^{organ}). A dose volume histogram is shown in Figure 2.2. A DVH can give

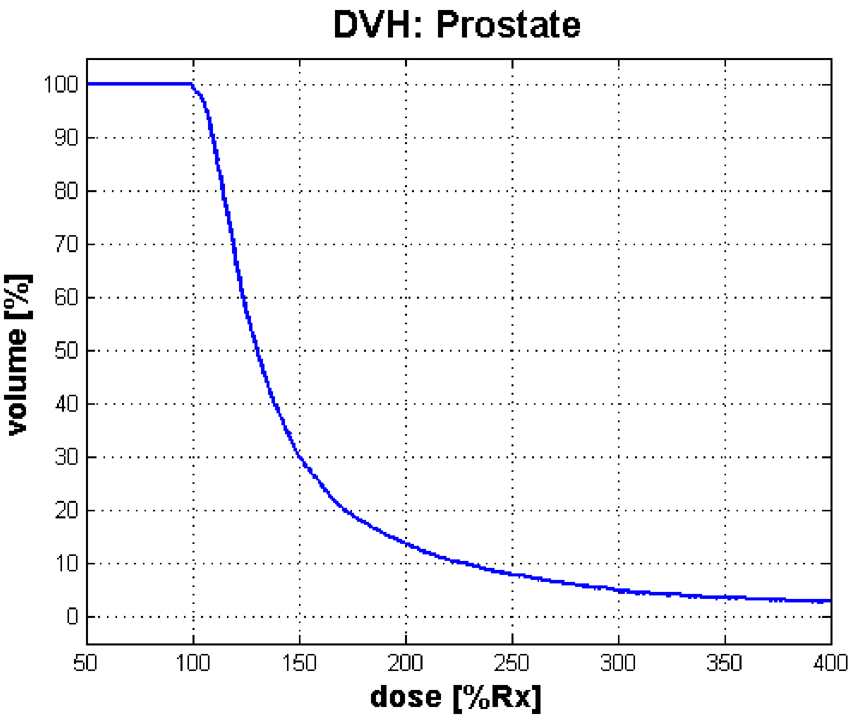

Figure 2.2: Example DVH for the prostate. On the x-axis is dose, as a percentage of the prescription dose, and on the y-axis is the organ volume receiving at least the given dose. A DVH can be thought of as all pairs, $(D, \mathrm{V}_D^{\mathrm{Prostate}})$. DVHs provide intuition for the dose distribution within the prostate, but there are no quantitative metrics to evaluate and compare them.

intuition about the distribution of dose inside an organ. For the prostate, it is desired to have the DVH stay above 90%, or as high as possible, until the prescription dose, then drop off as quickly as possible afterwards. The second is a heat map of the dose distribution shown in consecutive slices. Figure 2.3 shows such a heat map. These heat maps can give the user an intuitive sense of the coverage, exposure to OAR, and distribution of high-dose regions inside the prostate. Since there are no quantitative metrics to evaluate or compare dose distributions using DVHs or heat maps, we avoid using these qualitative tools to evaluate a dose distribution. Instead, we will evaluate dose distributions in terms of their compliance to dose objectives defined by dose-volume constraints.

There have been a number of computational inverse planning tools designed to guide the user towards a final plan for both HDR brachytherapy and PPI brachytherapy. These methods are described in the following section.

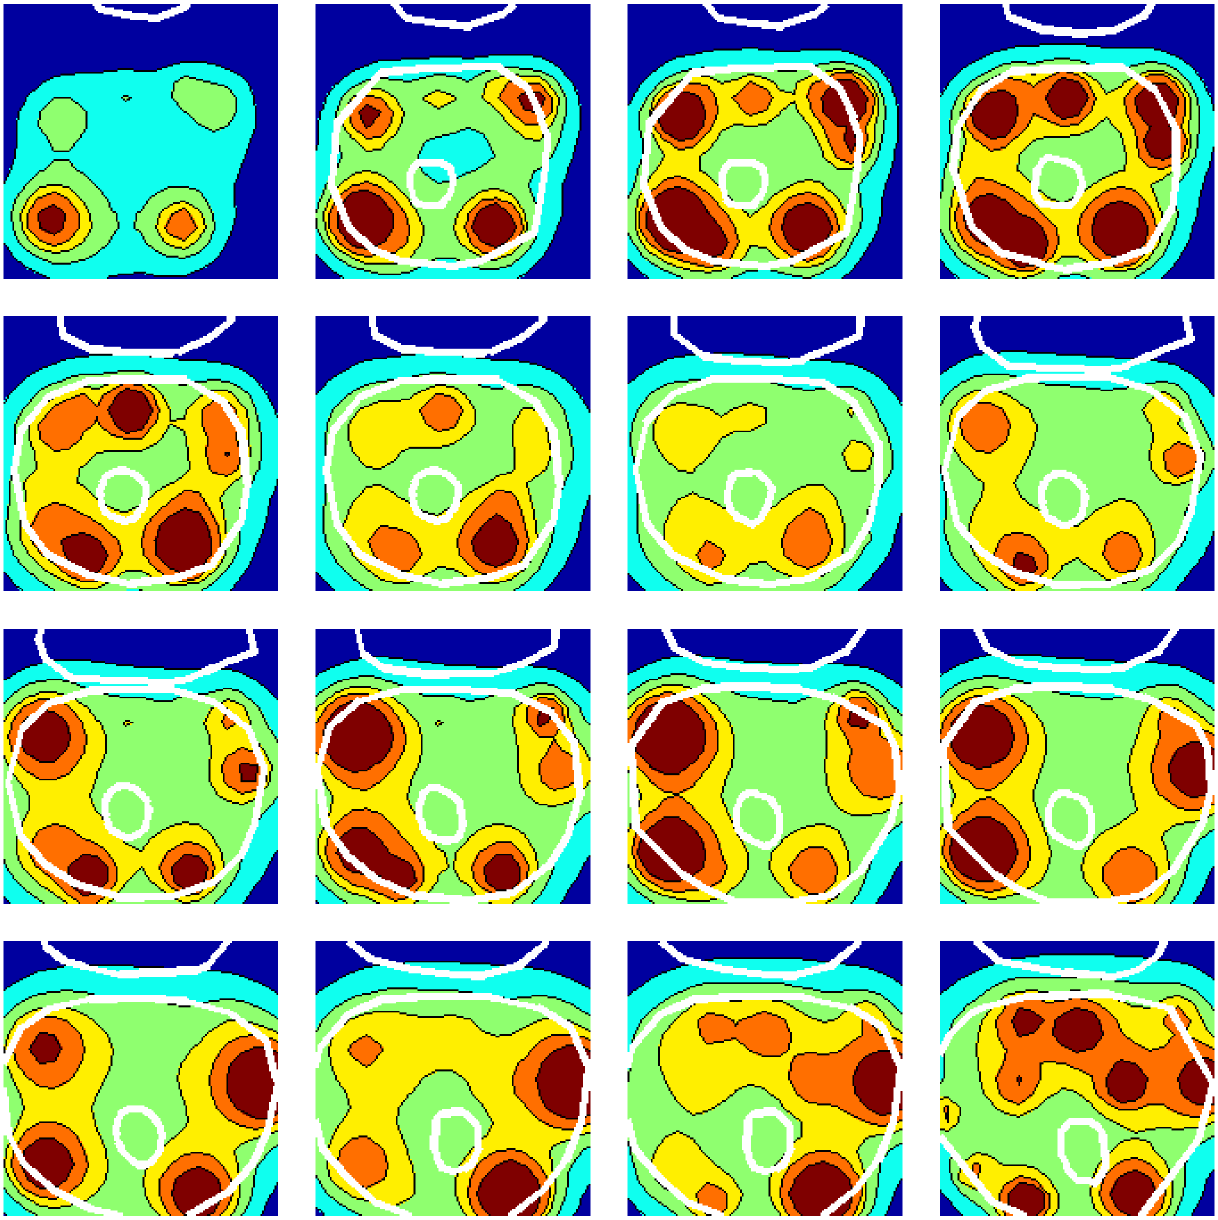

Figure 2.3: Heat map example for HDR brachytherapy dose distribution. The heat maps are shown in consecutive slices going in the positive z-direction (from feet to head). Shown are the prostate (outer contour) and inner contour (inner contour). Dose over 200% the prescription dose is shown in red, dose over 150% the prescription dose is shown in orange, dose over 125% is shown in yellow, dose over 100% is shown in green, dose over 75% the prescription dose is shown in cyan, and dose less than 75% the prescription dose is shown in blue. Ideally the entire prostate region would be green, and everything else would be blue. However, due to physical and geometric constraints, this is never possible.

2.3 Computational Inverse Planning Tools

When inverse planning was novel in the early 1990's [38], a protocol for the dose which should be delivered to organs was not known. Therefore, early computational tools for computing both PPI and HDR brachytherapy dose plans used multi-objective functions to control the dose distribution. Some tools maximized a weighted sum of user-defined dosimetric indices [20, 61, 104], positive weight for dosimetric indices which should be high, such as target coverage, and negative weight for dosimetric indices which should be low, such as those for OAR. More recent tools minimized the sum of penalty functions [18, 22, 52, 53, 66, 67]. A penalty function was assigned to each discretized organ voxels in terms of the dose received. For voxels in organs which should receive dose, such as the prostate, the penalty function would accumulate linearly below a threshold dose, i.e. the prescription dose, according to a penalty weight assigned for that organ. For voxels in OAR, the penalty function would accumulate for dose above some user-specified threshold dose for that organ. The use of penalty functions to control the dose distribution is described in further detail in Chapter 3. Thus, multi-objective methods are effective at representing the relative importance of target coverage versus healthy organ sparing by a single cost function and the relative weights associated with each dose objective.

Any planning systems always have two major components: a mathematical model and a solution algorithm for computing a plan for the model. The aforementioned multi-objective models have traditionally been solved with genetic algorithms or simulated annealing. However, these algorithms have very few mathematical guarantees on the quality of the final solution and on the computational effort required to achieve solutions of a given quality. Despite the limitations of simulated and geometric algorithms as solution algorithms, and therefore the limitations of the planning systems that use them, this dissertation focuses on the underlying mathematical models used to describe inverse planning objectives and the properties of these models, rather than comparing solutions algorithms. For example, Alterovitz et al, 2006 [2], compared a linear programming solution method with a simulated annealing solution method for the same penalty model. This dissertation focuses on comparing different models and the ways in which they can improve brachytherapy planning.

The advantage of multi-objective dose planning is that the relative importance between the different objectives can be represented by the relative weights of each objective in the cost function. By varying the objective weights, the set of dose distributions (evaluated by dosimetric indices) that are generally possible can be explored, which can eventually lead to a dosimetric protocol based on a set of desired dosimetric indices and constraints on them; for example, those specified by the RTOG-0321 dosimetric protocol.

However, multi-objective planning methods are not well suited for planning towards an established protocol because they do not necessarily guarantee compliance with dose-volume constraints on a per-patient basis. In other words, given an anatomy set and set of dose objectives, there is no mathematical guarantee that multi-objective methods will return a solution which complies with the dose objectives. Thus, multi-objective planning towards a dose plan which complies with dose-volume constraints is an iterative process. The user specifies an initial set of objective weights for each for each dose objective, and computes a dose distribution based on those penalty weights. If the dose distribution complies with the objectives, then it can be delivered to the patient. If it does not comply with the objectives, then the user must adjust the penalty weights and re-compute dose distributions until it does satisfy the objectives. Although dose objectives can usually be met in a few iterations given a good implant, manual fine-tuning of computational parameters can be a time-consuming process. This is especially true since physicians are trained to evaluate dose distributions in terms of dosimetric indices rather than system-specific optimization parameters.

Dose planning methods assume that the needles have already been inserted into the patient and digitized along with the anatomy. However, there has been very little work in computational needle planning for HDR brachytherapy [52], probably because digitized anatomy information is usually not available during needle insertion when following the standard workflow. The study by Karabis et al, 2009 [52] combined needle and dose planning into a single optimization model. Needles were selected from a candidate set which contained one needle for every hole in the template and the dose planning method was multi-objective. The model could not be solved quickly as an integer program (discussed in the following section), even given three days of computational time, for two of the three trials, but high quality solutions could be found using simulated annealing. Although needle planning is incompatible with the current HDR brachytherapy workflow, it may be possible to incorporate properties of the PPI brachytherapy workflow, which does have anatomy information during needle planning.

There are two studies that have taken skew-line needle planning into account for PPI brachytherapy that we know of. Roy et al, 1991 [96], did some computational needle planning that utilized non-parallel, skew-line needles. However, their optimization method has a significant manual component, i.e. required significant human input and decision making, and did not take needle collisions into account since in PPI brachytherapy, only one or a few needles are inserted at once. More than 15 years later, Fu et al, 2006 [40] used simulated annealing algorithm to optimize a needle configuration to the seed positions. That is, given a set of seeds and their desired location, they optimized a set of skew-lines that could implant the seed configuration.

2.4 Mathematical Programming

Many dose planning approaches draw on modeling and solution methods from the field of mathematical programming, commonly known as optimization. The work presented in this dissertation draws particularly on linear and integer programming. This section summarizes the fundamentals of these optimization models. A thorough explanation of linear programming can be found in Introduction to Linear Optimization by Bertsimas [14] and Convex Optimization by Boyd [16], and a complete overview of integer programming can be found in Integer and Combinatorial Optimization by Nemhauser and Wolsey [116].

2.4.1 Linear Programming

A linear program (LP) is a problem with the following form:

$$
\text{(LP)} \qquad \begin{array}{ll} \text{maximize} & c^T x, \\ \text{subject to:} & Ax \leq b, \\ \text{and} & x \in R^n, \end{array}
$$

where the known parameters c, A, and b are $n \times 1$, $m \times n$, and $m \times 1$ real matrices, respectively, and x is an $n \times 1$ vector of decision variables (i.e. values which can be chosen). An inequality of the form $a_i x \leq b_i$, where a_i is the i^{th} row of A and b_i is the i^{th} element of b, is called a constraint. Linear programs are usually specified in terms of constraints because the structure of the model relative to the underlying problem is more intuitive in this form. A linear program constraint form is given below.

$$
\text{(LP)} \qquad \begin{array}{ll} \text{maximize} & c^T x, \\ \text{subject to:} & a_i x \leq b, \forall i, \\ \text{and} & x \in R^n. \end{array}
$$

A feasible solution is an x that satisfies $Ax \leq b$, and the feasible set is the set of all feasible solutions. An LP is infeasible if the feasible set is empty. The term $c^T x$ is called the objective function. A feasible solution is optimal if there are no other feasible solutions with higher objective function value. Note that there may be multiple optimal solution with the same objective function value. An LP is unbounded if there are feasible solutions which can make the objective function arbitrarily large. It can be proven that a linear program is either infeasible, unbounded, or has an optimal solution.

The most popular LP solvers are Simplex and the Interior Point Method (IPM) algorithms. The feasible set of a linear program forms a mathematical object known as a polyhedron, $\{x : Ax \leq b\}$. Assuming that an LP is feasible and bounded (has a finite optimal solution), it can be proven that there is at least one optimal solution at a corner, or extreme point, of the polyhedron. Simplex works by starting from a vertex of the feasible set, and checks if there are any adjacent vertices which have better objective function values. If so, then the improving vertex becomes the current vertex of the next iteration, and the process is repeated until no improvements are possible. Because of the mathematical structure of polyhedrons, once there are no improving vertices adjacent to the current vertex, then the solution is provably optimal. Although Simplex has not been proven to be an efficient algorithm, it works extremely well in practice [15, 105]. The IPM starts from an interior point of the polyhedron, and follows the objective function to an optimal solution. IPM has been shown to be polynomial time [55], which for the purpose of this dissertation, means that LPs can be reliably proven to be infeasible, proven to be unbounded, or solved to optimality within a few minutes on a personal computer, even with hundreds of thousands of variables and constraints.

2.4.2 Integer Programming

An integer program (IP) is a linear program where some or all of the variables are restricted to be integer. An IP in general form is given below.

$$\text{(IP)} \quad \begin{array}{lll} \text{maximize} & c^T x, & \\ \text{subject to:} & Ax \leq b, & \\ \text{and} & x_i \in R, i = 1, ..., p & x_i \in Z, i = p+1, ..., n. \end{array}$$

Integer programs are an important modeling tool for solving real-life problems which require decisions to be made in discrete quantities, e.g. often an item must be bought or not bought, it cannot be partially bought. However, even though IPs are able to model scenarios which cannot be modeled by linear programs, there are no known polynomial time algorithms for the general integer programs. The following paragraphs describe some terminology associated with integer programs, and a description of a general branch and bound (B&B) method for solving them. The section closes with a brief note on the use of the word "optimal" in relation to dose planning and optimization.

A vector x is called an integer feasible solution for IP if it satisfies all the constraints and the integrality requirements of an IP. The set of all integer feasible solutions is referred to as the integer feasible lattice. A vector x is called a fractional solution if it satisfies the constraints of IP but not the integrality requirements. A variable is said to have a fractional value if it is (1) integer constrained and (2) has a non-integer value in a fractional solution.

If the integrality requirements of IP are dropped, the resulting optimization model is called the linear programming relaxation of IP, (R(IP)), because the result is a linear program. If the optimal solution of R(IP), assuming it exists, is integer feasible, then the solution is provably optimal for IP. However, the optimal solution to R(IP) is generally fractional because it is unlikely that the constraints will come together in a way that the optimal solution for R(IP) will fall exactly on the integer lattice.

If the optimal solution for R(IP) is fractional, its optimal objective function value is always superior to the optimal objective function value of IP because the feasible set for IP is a subset of the feasible set of R(IP), and therefore, the R(IP) cannot produce a worse solution than IP. As a consequence, the objective function value of R(IP) gives an upper bound on the optimal objective function value of IP.

Integer programs can be solved using branch and bound (B&B) algorithms. The first step of B&B is to take the linear programming relaxation of IP. If the optimal solution for R(IP), is integer feasible, then the solution is optimal for IP. If the solution is fractional, B&B begins constructing a tree of nodes called the branch and bound tree, and the solution for R(IP) becomes the first, or root, node.

The branch and bound tree is constructed according to the following procedure. Let $\overline{x}$ be a fractional solution at an arbitrary B&B node, the current node, and x_i be a fractional variable at the node having the value $\overline{x}_i$. Since x_i must be integer, then the optimal solution must have $x_i \leq \lfloor \overline{x}_i \rfloor$ or $x_i \geq \lceil \overline{x}_i \rceil$. Thus, the problem of solving the original IP can be divided into two subproblems: The original IP with the added constraint $x_i \leq \lfloor \overline{x}_i \rfloor$ and the original IP with the added constraint $x_i \geq \lceil \overline{x}_i \rceil$. These two subproblems, the down branch and up branch, respectively, are associated with two child nodes of the current node. The variable x_i in the current node is called the branching variable. There are many methods for choosing which fractional variable to branch on, but these methods are beyond the scope of this overview.

As the branch and bound tree grows, nodes lower in the tree accumulate constraints. With the accumulation of constraints, the linear programming relaxation of the IP with the additional constraints is either fractional, infeasible, or integer feasible. Fractional constraints continue to grow the tree. The child nodes of infeasible or integer feasible nodes do not need to be searched, and so the branch and tree does not grow from these nodes.

When an integer feasible solution is found at a node, then its objective function value can be compared with the objective function value of all the fractional solutions in the branch and bound tree. If a fractional objective function value is lower than the objective function value of the feasible solution, then the optimal solution cannot be in any child node of that fractional solution, and the node can be pruned from the tree (i.e. do not search any of its child nodes). This is a consequence of the fact that (1) the relaxation is an upper bound on the objective function value of any integer feasible solution that could be found in any child node and (2) there is already an integer feasible solution with higher objective function value. The ability to prune branches, even ones far away from the current node, based on these properties gives B&B an advantage over enumerating all the integer possibilities.

The difference, whether relative or absolute, between the largest fractional solution objective function value of a node whose branches have not been explored (the upper bound) and the largest integer feasible solution objective function value (the lower bound) is called the gap. As nodes are searched, the upper bound decreases and as integer feasible solutions are found, the lower bound increases. The node with the highest integer feasible solution objective function value is provably optimal when the gap is zero. This is easy to argue since an optimal solution implies there does not exist any other feasible solution with higher objective function which is guaranteed when the upper bound is equal to lower bound.

Linear programs can be solved quickly and there are efficient mathematical methods of using the solution of a linear program to solve the same linear program with an additional constraint, such as when branching on a fractional variable. However, the number of nodes in the tree grows exponentially, and in the worst case, every possible integer combination would have to be enumerated and checked before an optimal solution could be found and proven to be optimal. Even for an integer program with as few as 100 variables, enumerating and checking all the possibilities could easily take billions of years, even using the fastest computers in the world.

The gap gives an estimation of the quality of the best integer feasible solution found so far. For instance, if the gap is small, then we know that the best feasible solution is close to optimal. It may even be optimal but not yet proven by B&B. In many instances of integer programming, these suboptimal, or not-provably-optimal solutions are still of practical value if the gap is small.

Although integer programs are provably difficult in the general case, there has been a tremendous amount of research in developing mathematical methods for solving specific instances of integer programs and/or improving the performance of B&B. These methods can generate heuristic, high-quality feasible solutions for the general model (i.e. small gap); produce new constraints called cuts for the model which remove fractional solutions but not integer solutions, which improves the bound of the relaxation at any node; or determining branching variables which are promising for reducing the gap. The development of customized solution strategies has made integer programming a practical solution for many real-world problems.

2.4.3 Mathematical Programming for Inverse Planning

Linear and integer program have been used in inverse planning. For example, the study by Alterovitz et al, 2006, [2], created an equivalent linear program to the simulated annealing model presented by Lessard et al, 2004 [66], which provided much stronger performance guarantees on optimality and computation time. A similar linear programming reformulation was created and tested on the simulated annealing model presented by Karabis, 2005 [53]. A linear program was also used to compute the dose distribution for an external beam therapy tool called Intensely Modulated Radiation Therepy (IMRT) which constrained the mean of the dose to volume in each organ above a user-specified dose-threshold [95]. In other words, this method constrained the mean of all the dosimetric indices for an organ above a dose threshold. However despite the use of more formal optimization techniques, these methods do not enforce dose volume constraints, and therefore are still parametric and iterative.

Integer programming has been used for PPI source configuration planning [62–64] which used binary variables, variables which could only take on values 0 or 1, to model the placement or non-placement of a source at a candidate location. Lee et al, 1999, developed a custom algorithm for solving a multi-objective integer program for determining the placement of these sources which could find a solution within a few minutes. Again, the model used was multi-objective and therefore would require iteration to reach a solution which met dose-volume constraint objectives. Also, although PPI brachytherapy and HDR brachytherapy are very similar in practice, the discrete nature of placing sources versus the continuous nature of dwell times make PPI and HDR brachytherapy mathematically different. Therefore, this customized solution method for PPI brachytherapy is not usable for brachytherapy.

Finally, it is possible to use integer programming to model dose-volume constraints [31]. However, attempts to solve practical instances of these models using standard software such as CPLEX, have not yielded solutions quickly enough to be clinically viable.

2.5 Robot-Assisted Brachytherapy

There have been several robotic hardware platforms specially designed to assist in the insertion seeds for PPI brachytherapy [19, 28, 32, 33, 59, 60, 79, 90, 100, 108–111] , and they are capable of operating within a standard brachytherapy clinical environment. These robots can bring a needle to a virtual continuum of insertion points and angles, which simulataneously provides degrees of freedom that can be used to avoid OAR and precision beyond that of a template. Once a desired source configuration is computed by a dose planning system, the sources can be implanted into the patient one at a time, or several at a time if they are collinear, using these robots. Some robots can be brought to a needle entry point and insertion angle by joystick or other manual control. The needle can then be inserted a desired depth, or the physician can manually insert the needle under TRUS guidance. There have been several studies which confirm the placement accuracy of robot-assisted brachytherapy workflows [1,74,101], even in a living canine subject [24]. Thus, robot-assisted brachytherapy is a promising paradigm for improving brachytherapy, reducing side effects, and reducing the skill requirements of performing brachytherapy at an expert level. In fact as of this year, a set of guidelines is currently being developed by the American Association of Physicists in Medicin (AAPM) for image-guided robotic brachytherapy [89].

There are several limitations for applying these robotic platforms to an HDR brachytherapy workflow. HDR brachytherapy requires a configuration of needles to be simultaneous inside the patient, while PPI brachytherapy robots insert a single needle several times. Also, there has been very little work in needle configuration planning for HDR brachytherapy and therefore, the positioning of the needles would have to be decided upon by the physician, which would reintroduce skill dependence. Finally, in PPI brachytherapy, the anatomy is known when the needles are inserted, because it is required to plan the seed configuration. However, in HDR brachytherapy, the needles are inserted before the anatomy is scanned. Thus, using robots for HDR brachytherapy would require an alteration of the standard workflow.

2.6 Addressing Previous Work

This dissertation directly addresses several limitations of current HDR brachytherapy dose planning, needle planning, and robotic-assistance. We develop a fast algorithm for directly optimizing dose-volume constraints, which are the standard objectives for HDR brachytherapy. Thus, our method does not require manual fine-tuning of optimization parameters, only the specification of dose objectives. Our method can also claim very strong mathematical guarantees on the quality of the final dose distribution relative to the dose objectives, and how long it will take to acquire a solution. We also develop an algorithm for computing anatomy-specific needle configurations over a set of needles which approximates the needles which can be inserted by a robot, which is a much larger set than those used in Karabis et al, 2009 [52]. Although this algorithm is guaranteed to be tractable, it works reasonable well in practice. Finally, we modify a needle insertion robot specialized for PPI brachytherapy to perform HDR brachytherapy, and use it in an end-to-end brachytherapy workflow which includes our needle and dose planning algorithms. We show that the needle configurations we plan can be implanted with sufficient quality and accuracy to meet objectives as planned.

Chapter 3

Dose Planning with IPIP

3.1 Overview

This chapter expands on work presented in Siauw et al, 2011 [102], which developed Inverse Planning by Integer Program (IPIP), a dose planning method for HDR brachytherapy. This dose planning method has several advantages over multi-objective methods. In particular, dose-volume constraints, a standard for evaluating HDR brachytherapy dose plans, are directly expressed and enforced in IPIP. Thus, a dose plan meeting dose objectives can be met without having to iterate through system-specific parameters.

However, this method utilizes an integer program, which is known to be difficult to solve in the general case (see Background). This means that the expected solution time is long enough to interrupt the standard clinical workflow. Therefore, we also develop a heuristic for IPIP, IPIP-H, which we show can compute high quality feasible solutions given the standard RTOG-0321 clinical objectives [48]. This heuristic solves a sequence of two linear programs and as a consequence is polynomial time solvable [55].

It is important to mention that Siauw et al. 2011 [102] only presented computational results on the heuristic, and therefore the term "IPIP" was used to refer to the heuristic. Since this chapter references both the integer program and the heuristic, "IPIP" is used to refer to the integer program, and IPIP-H is used to refer to the heuristic. This chapter repeats the derivation of IPIP and IPIP-H, and also presents computational results for IPIP, IPIP-H, the relaxation for IPIP, IPIP-R, and Inverse Planning Simulated Annealing (IPSA), which is a commercially deployed dose planning system used world wide for over a decade. These studies were not included in Siauw et al., 2011 [102].

3.2 Background

For the purpose of dose planning, it useful to discretize the organ volumes into voxels. A dose control point, or dose point, is a point at the center of each voxel, and the dose delivered to the dose point is representative of the dose delivered to every point within the voxel. If the dose points are generated on a finely and evenly spaced grid which covers the entire region of interest, then a dose distribution can be approximated by the set of doses delivered to the dose points. Also, dosimetric indices can be approximated by the sum of the dose points receiving over the threshold dose times the voxel volume.

Several inverse planning tools such as Inverse Planning Simulated Annealing (IPSA) [66] and Hybrid Inverse Planning Optimization (HIPO) [53] plan the dose distribution by penalizing the dose received at each dose point. Given a dose d at a dose point, the penalty is computed using a piecewise linear penalty function with the following properties: within a dose interval $[D^{\min}, D^{\max}]$, the penalty is 0, below $D^{\min}$, the penalty is $M^{\min}(D^{\min} - d)$, and above $D^{\max}$, the penalty is $M^{\max}(d - D^{\min})$. Figure 3.1 shows an example penalty function. For compactness, we represent the penalty function for dose points in an organ using the quadruplet $(M^{\min}, [D^{\min}, D^{\max}], M^{\max})$. All dose points within the same organ have the same penalty function, and therefore, the relative importance of the dose delivered to each organ can be represented by the magnitude of the penalty weights associated with dose points in each organ. For example, if $M^{\min}$ for the prostate is 4 and $M^{\max}$ for the bladder is 1, this is analagous to delivering dose above $D^{\min}$ to the prostate being four times as important as delivering dose below $M^{\min}$ to the bladder. The goal of the solver is to find a dose plan which has the lowest total penalty.

Multi-objective dose planning methods such as IPSA and HIPO do not enforce dose-volume constraints and can have undesirable behavior. To highlight this fact, consider the following thought experiment. Let the regions of interest contain two prostate dose points, p_1, p_2, and two urethra dose points, p_3 and p_4, and let the prescription dose be 1000 cGy. For prostate dose points, let the penalty weight be defined by $(2, [1000, 1500], 2)$ and $(0, [0, 1250], 1)$ for the urethra dose points. Now consider two dose distributions, D_1 and D_2, which are defined by the set of doses $\{d_1, d_2, d_3, d_4\}$, where d_i is the dose at p_i.

$$D_1 = \{950, 960, 1200, 1300\}$$

and

$$D_2 = \{1000, 1100, 1240, 1490\}$$

The total penalty for D_1 is 230, and the total penalty for D_2 is 240. Therefore according to this set of penalty weights, D_1 is superior. However the dosimetric indices for D_1 are $\mathrm{V}_{100}^{\mathrm{Prostate}} = 0$ dose points (0% target coverage), and $\mathrm{V}_{125}^{\mathrm{Urethra}} = 1$ dose point, while D_2 has $\mathrm{V}_{100}^{\mathrm{Prostate}} = 2$ dose points (100% target coverage), and $\mathrm{V}_{125}^{\mathrm{Urethra}}$ of 1 dose point. Thus according to dosimetric indices, D_2 is superior (i.e. more coverage, same urethra dose).

Although in terms of biological effect, the two dose plans are very similar, and it is possible that a different set of penalty weights may yield a better dose plan, this experiment illustrates where penalty methods do not perform as desired. Thus, penalty methods are iterative. A user specifies a set of penalties that usually produce a good dose distribution. These initial penalties are set according to a standard class solution which has been found through clinical experience [67]. Then the user observes the dose distribution, i.e. compliance with dose objectives. If the plan does not conform to the objectives, the penalty weights are

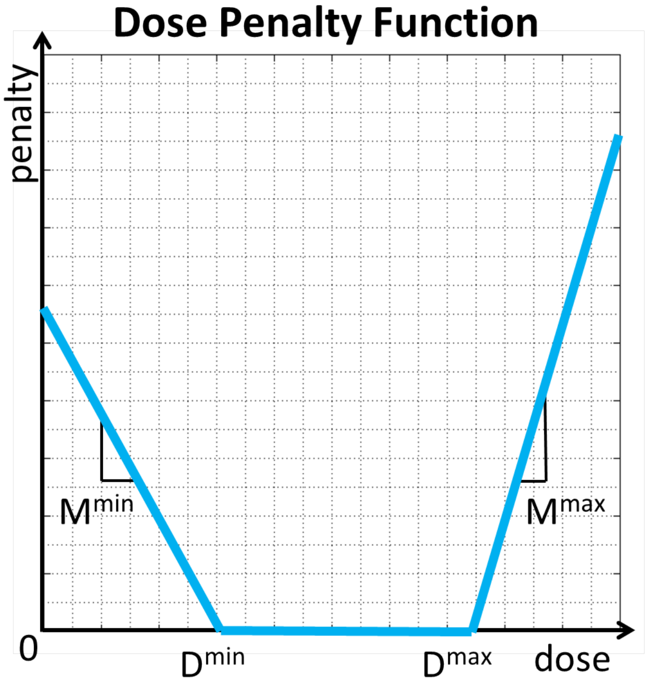

Figure 3.1: Example of piecewise linear penalty function. Within an interval, $[D^{\min}, D^{\max}]$, no penalty is accumulated. Outside the interval, penalty is accumulates linearly. Penalty functions help control the dose distribution during dose planning, but do not directly enforce dose volume constraints.

adjusted and the process is repeated until a satisfactory dose plan is achieved. Although the RTOG-0321 can usually be achieved with several iterations, the iterative process requires manual fine tuning of many parameters which are specific to the model (e.g. IPSA or HIPO), which can be a time consuming process. The dose planning method presented in this chapter directly optimizes dosimetric indices and therefore dose not require manual tuning.

3.3 Method and Materials

3.3.1 Model Formulation

In this section, Inverse Planning by Integer Program (IPIP) is developed. A description of the parameters and variables of this model are given for reference in Table 3.1. Note the difference in notation presented in this chapter and in Siauw et al. 2011. The changes were made to assist in clarity.

Let P be a set of 3D coordinates for a grid of evenly spaced dose points, I_s be the set of

Term	Description
G	Set of organs
P	Set of dose points
I_s	The set of indices of dose points for $g_s \in G$
T	The set of dwell positions
D_{ij}	The dose rate from $\tau_j \in T$ to $p_i \in P$
t_j	Dwell time at $\tau_j \in T$
d_i	Dose at $p_i \in P$
R_s	Threshold dose for $g_s \in G$
M_s	Maximum dose for $g_s \in G$
x_i	Indicator variable for $p_i \in P$
v_s	Dosimetric index for $g_s \in G$
L_s	Lower bound for v_s
U_s	Upper bound for v_s

Table 3.1: IPIP terms

dose points that belong to organ s, and T be the set of dwell positions for a given implant. It is assumed that the grid of dose points covers the entire region of interest, meaning that each organ has enough dose points to accurately count a dosimetric index for it. Also, let the dose rate parameter, $D_{ij} > 0$ [cGy/s], be the dose received at $p_i \in P$ for every second the source remains at $\tau_j \in T$, R_s [cGy] be the dose required for the volume associated with $p_i \in P$ to be counted towards the dosimetric index for $g_s \in G$, M_s [cGy] be the maximum dose allowed for dose points in g_s, and L_s and U_s be the lower and upper bounds for the dosimetric index for $g_s \in G$, respectively.

Let the optimization variables be the dwell times, t [seconds], the doses at each dose point, d [cGy], the dosimetric indices for each organ, v, and the indicator variables, x.

A dwell time, $t_j \geq 0$, is a continuous variable which denotes the time spent at $\tau_j \in T$. The total dose received at $p_i \in P$ is then the sum of the dose contributions from every dwell position:

$$d_i = \sum_{j=1}^{|T|} D_{ij} t_j$$

.

The indicator variables, x, are binary variables with the following behavior:

$$x_i = \begin{cases} 1 & \text{if } d_i \geq R_s \\ 0 & \text{otherwise} \end{cases}$$

for every $i \in I_s$. This behavior can be enforced by the constraints

$$R_s x_i \leq d_i \leq R_s + (M_s - R_s)x_i - \epsilon \text{ and } x_i \in \{0, 1\}$$

where ϵ is a small number. For a given dose at p_i, if $x_i = 1$, then the constraints reduce to $R_s \leq d_i \leq M_s - \epsilon$, and if $x_i = 0$, they reduce to $0 \leq d_i \leq R_s - \epsilon$. If $d_i < R_s$, then $x_i = 0$ to satisfy $R_s x_i \leq d_i$, and if $d_i \geq R_s$, then x_i must be 1 to satisfy $d_i \leq R_s + (M_s - R_s)x_i$. Therefore, $x_i = 1$ if and only if $d_i \geq R_s$, which is the desired behavior.

The ϵ parameter eliminates the ambiguity for the value of x_i when $d_i = R_s$. Specifically, if d_i equals R_s, then x_i should take on value 1 according to its definition. However without the ϵ term, 1 or 0 is valid for x_i in the context of the given constraints. For very small ϵ, $d_i \leq R_s + (M_s - R_s)x_i - \epsilon$ is an approximation for $x_i \leq d_i < R_s + (M_s - R_s)x_i$. If this level of precision is not required, then ϵ can be omitted from the formulation.

A dosimetric index for an organ, v'_s, can be computed as the sum of the voxel volume, V_i, multiplied by the indicator variable associated with the dose point p_i,

$$v'_s = \sum_{i \in I_s} V_i x_i.$$

Since the grid of dose points is evenly spaced, then $V_i = V$, and

$$v'_s = V \sum_{i \in I_s} x_i.$$

Finally, dose objectives can be enforced with the constraints

$$L'_s \leq v'_s \leq U'_s,$$

where L'_s and U'_s are respectively the minimum and maximum volumes for v'_s. For compactness, it is useful to drop the V parameter and express v'_s, L'_s, and U'_s as counts rather volumes. In other words, let

$$v_s = \sum_{i \in I_s} x_i,$$

and

$$L_s \leq v_s \leq U_s,$$

where $L_s = \lceil \frac{L'_s}{V} \rceil$ and $U_s = \lfloor \frac{U'_s}{V} \rfloor$.

Since the dosimetric indices are represented exactly using the aforementioned set of integer programming constraints, the objective can be the maximization or minimization of any dose objective, i.e. maximizing target covearge or minimize urethra dose. However, maximizing target coverage, v_0, is a natural choice since it represents the most desired therapeutic effect, which is the eradication of the cancer. We present the IPIP model below.

$$
(IPIP)\qquad
\begin{array}{lll}
\text{maximize} & v_0 & \\
\text{subject to:} & d_i = \displaystyle\sum_{j=1}^{|T|} D_{ij} t_j, & \forall i, \\
 & R_s x_i \le d_i \le R_s + (M_s - R_s) x_i - \epsilon, & \forall i \in I_s, \forall s, \\
 & v_s = \displaystyle\sum_{i \in I_s} x_i, & \forall s, \\
 & L_s \le v_s \le U_s, & \forall s, \\
\text{and} & x_i \in \{0,1\}, t_j \ge 0, & \forall i, j.
\end{array}
$$

IPIP can be represented more compactly by substituting the equality constraints.

$$
(IPIP)\qquad
\begin{array}{lll}
\text{maximize} & \displaystyle\sum_{i \in I_0} x_i & \\
\text{subject to:} & R_s x_i \le \displaystyle\sum_{j=1}^{|T|} D_{ij} t_j \le R_s + (M_s - R_s) x_i - \epsilon, & \forall i \in I_s, \forall s, \\
 & L_s \le \displaystyle\sum_{i \in I_s} x_i \le U_s, & \forall s, \\
\text{and} & x_i \in \{0,1\}, t_j \ge 0, & \forall i, j.
\end{array}
$$

Since the constraints of IPIP directly constrain and optimize the volume of organs receiving a certain dose, IPIP can be used to represent dose objectives such as the one in Table 2.1. However, our initial tests with IPIP on actual patient data showed that standard solvers could not reliably find high quality (i.e. near optimal) solutions for a clinically relevant number of dose points. Therefore, in the next section, a fast heuristic is developed for computing feasible solutions for IPIP.

3.3.2 IPIP Heuristic Formulation

The ability of IPIP to control the dose plan based on dosimetric criteria is dependant on an accurate count of the number of dose points which contribute to a dosimetric index. Specifically, it relies on the assumption that x_i is either 0 or 1 depending on the dose at p_i. However, when branch and bound solves IPIP, the first step is to take the linear program relaxation of IPIP, which allows x_i to take any value between 0 and 1. When the binary constraint on x_i is relaxed, dosimetric indices are no longer counted correctly.

Recall that some dosimetric indices, specifically the ones for OAR, are upper bound, i.e. $v_s \le U_s$. As a consequence, IPIP will attempt to allocate which dose points in g_s can have

$d_{i \in I_s}$ over R_s as efficiently as possible (i.e. which x_i should take value 1) to make the objective as high as possible. However, the linear programming relaxation does not induce the correct behavior. For example, let $d_i = 1.5R_s$, where p_i is in an OAR. Under these conditions x_i should take value 1. In the relaxation, the linear programming optimization will attempt to be as economical with x_i as possible, since it only has a "budget" of U_s dose points which can be above R_s. Specifically, x_i can take value

$$\frac{d_i - Rs}{M_s - R_s} + \epsilon = \frac{0.5R_s}{M_s - R_s} + \epsilon,$$

which is less than 1, and still satisfy the constraint $d_i \leq R_s + (M_s - R_s)x_i$. In fact, if $M_s >> R_s$, x_i can take on a value very close to 0 when it should be 1, which is not desirable behavior. Therefore, dosimetric indices which are upper bound are likely to be under-counted in the LP relaxation of IPIP, making the relaxed dose plans infeasible for the original integer program. Similarly, dosimetric indices which are lower bound, such as target coverage, are likely to be over counted.

Our IPIP heuristic (IPIP-H) is a fast method that can generate feasible solutions for IPIP assuming that (1) target coverage is the only dose objective which is lower bound (i.e. $L_s \leq v_s$) and (2) there is no "primary" objective, such as minimizing urethra dose, with the exception of maximizing target coverage. Note that all OAR dose objectives are upper bound constraints, i.e. $v_s \leq U_s$, because by definition, an OAR is never required to receive dose. Given these restrictions on the dose objectives, the fundamental intuition of IPIP-H is the following. An optimization model is developed which will maximize target coverage while meeting only the hard dose constraints, i.e. $d_i < M_s$. We show that this optimization model is easy to solve to optimality. However, since dosimetric indices are not constrained in this model, it is unlikely that the solution of this optimization model will be feasible for IPIP. Therefore, the infeasible dose distribution is analyzed, and new constraints are that restrict some of the OAR dose points to receive less than R_s dose. The dose points with the smallest over-dosing are chosen so that high target coverage can still be achieved. A diagram of this process is shown in Figure 3.2.

Next, we give a formal description of IPIP-H. First, a relaxation of IPIP is used to generate a dose plan. This relaxed optimization model is referred to as the heuristic relaxation (HR). In HR, the constraint $x_i \in \{0, 1\}$ is relaxed to allow continuous values of x_i. The restrictions on dosimetric indices, $L_s \leq v_s \leq U_s$ are also relaxed. The result is an LP which enforces the individual hard dose constraints $d_i \leq M_s$. HR is shown below. Note that without constraints on dosimetric indices, all x variables for OAR can be dropped.

Figure 3.2: IPIP-H works by creating a dose distribution which meets hard dose constraints (left), but ignores soft constraints. Since the objective in IPIP-H is to maximize target coverage, these dose plans are likely to overdose OAR. To address this problem, IPIP-H cools the dose plan starting from the smallest overdosed volume on the periphery (middle) until it meets the dose constraints for OAR (right).

$$
\begin{array}{llll}
 & \text{maximize} & \displaystyle\sum_{i \in I_0} x_i & \\
 & \text{subject to:} & \displaystyle d_i = \sum_{j=1}^{|T|} D_{ij} t_j & \forall i \\
(HR) & & R_0 x_i \leq d_i & \forall i \in I_0, \\
 & & d_i \leq M_s - \epsilon, & \forall s \neq 0, \forall i \in I_s, \\
 & \text{and} & 0 \leq x_i \leq 1, & \forall i \in I_0, \\
 & & t_j \geq 0, & \forall j.
\end{array}
$$

As with any relaxation, there is some loss of fidelity between the original model and the relaxation. This paragraph describes this loss. Since HR does not constrain dosimetric indices, dose objectives for the OAR are unlikely to be met in dose plans generated from HR. HR will attempt to meet the target coverage constraint, $v_0 \geq L_s$, but only through its maximization in the objective function. However, it is extremely important to note that the target coverage represented in HR is not true target coverage since the x variables which were once binary are now continuous. As a consequence, target coverage is not counted correctly. For instance, if p_i receives $0.75R_x$, then x_i can have value 0.75 and still be valid for $R_s x_i \leq d_i$, even though the value for x_i should be 0. This can have poor performance if, for instance, every prostate dose point received $.99R_x$, in which HR would report 99% coverage when the real target coverage is 0%. However, our computational results for standard dose objectives show that HR still performs well, despite this possible pathology. The advantage of HR over

IPIP is that it is a linear program, which has desirable performance guarantees over integer programming as mentioned in Chapter 2.

Let $\overline{d}_s^{\mathrm{HR}}$ be the vector of doses for dose points in organ s for the dose plan computed from HR. By construction, $\overline{d}_s^{\mathrm{HR}}$ will meet hard dose constraints specified in IPIP. Also, it is likely that target coverage will be high since it is maximized in the objective, and improving target coverage in the optimization is unincumbered by restrictions on dosimetric indices for OAR. However, it does not necessarily meet constraints on dosimetric indices specified for IPIP because they were relaxed. Specifically, more than U_s dose points will receive more than R_s dose in each organ. To make the solution feasible for IPIP, constraints must be added to $|I_s| - U_s$ dose points of the form $d_i \leq R_s - \epsilon$. In other words, only U_s dose points in organ s may have a dose upper bound of M_s; the remaining dose points must have a dose upper bound of R_s.

The selection of the dose points should have their dose upper bound reduced to R_s should be done in such a way as to still produce high target coverage. For IPIP-H, every dose point not among the hottest (i.e. most dose points) U_s dose points are constrained to receive less than R_s dose. The reason this method of selection was chosen is as follows. Without the dose restrictions to OAR, an HR dose plan will likely be too hot. Specifically, too many dose points in OAR will be receiving more than R_s dose. Therefore, the dose to some of the dose points must be reduced to make it feasible for IPIP. Our heuristic assumes that adding constraints that reduce the dose to the dose points receiving the least dose will have the least impact on target coverage in the HR dose plan. In other words, the HR dose plans is cooled as little as possible to make it feasible for IPIP. In this way, high coverage is retained.

HR is then solved with the new constraints. Since only U_s dose points in each organ have a dose upper bound of M_s, upper bound dose objectives for IPIP are guaranteed to be met in augmented HR. Thus, the solution for the augmented HR is guaranteed to be feasible for IPIP, except for target coverage, which is maximized in the objective. Our IPIP-H is summarized as follows.

(1) Solve HR and store dwell times.

(2) Using the dwell times from HR, for each organ except the CTV, let P_s* denote the U_s dose points receiving the most dose in that organ.

(3) For every $p \in P_s*$ add the constraint $\sum_{j=1}^{|T|} D_{ij} t_j \leq R_s - \epsilon$ to HR.

(4) Re-solve augmented HR to get dose plan.

We mention here that steps (1) and (4) are linear programs and step (2) involves sorting (which is also polynomial time). Therefore, IPIP-H is polynomial time.

3.3.3 Patient Data Sets

In Siauw et al., 2011 [102], 20 prostate cancer patient data sets from patients previously treated at the UCSF clinic were used to evaluate IPIP-H. In this study, a different group of 18 anonymized patients were used. These patients were chosen to have a wide range of prostate volumes ranging from 23 to 82 cm^3. The physician used a freehand technique to implant 16 catheters into the prostate under transrectal ultrasound guidance. Plastic catheters were inserted transperineally by following the tip of the catheter from the apex of the prostate to the base of the prostate using ultrasound and a stepper. A Foley catheter was inserted to visualize the urethra.

Three-millimeter-thick CT slices were collected using a spiral CT. The implanted catheters, target (prostate), and the organs at risk (OAR) were contoured using Oncentra 3.2. The OAR included the rectum, bladder, urethra, and critical structures near the bulb. No margins were added. The structures near the penile bulb were contoured as a single organ. When segmenting the bladder and rectum, the outermost mucosa surface was contoured. The urethra was defined by the outer surface of the Foley catheter, and only the urethral volume within the prostate was contoured. Only dwell positions within the prostate were used during dose planning.

Contour slices with more than 15 points were reduced to 15 points. Reducing the number of contour points is a common feature in many dose planning systems. A visual inspection was made to ensure that the organ features were preserved. This reduction improved computational performance generating dose points, which were created on a uniform grid of points spaced 2 mm in the x- and y-direction, and 3 mm in the z-direction for every region of interest except the space the between organs, which had dose points spaced 4 mm in the x- and y-direction, and 3 mm in the z-direction.

It is worth noting that the needle implants for these patients were implanted by an expert physician (or under the direct supervision of one). Since the set of possible dose plans is governed by the needle implant (i.e. the dwell positions), it is easier to find high quality dose plans when the needle implant is well done. Specifically, if the needles are evenly spaced throughout the prostate volume. Conversely if the needle implant is poor, i.e. has too few needles or badly placed needles that leave large voids in the prostate volume that do not have needles nearby, then dose planning will be much harder, if not impossible.

3.3.4 Dose Rate Calculations and Dose Objectives

When using the point-source calculation specified in the AAPM TG-43 dosimetry protocol [83], the dose-rate contribution [cGy/sec] to a dose point from a source dwell position is only a function of the distance between them and the radioactive properties of the material, which was ^{192}Ir. The dose rate between $p_i \in P$ and $\tau_j \in T$, is calculated according to

$$D_{ij} = \frac{\gamma_1\gamma_2\gamma_3\gamma_4}{3600D^2}(g_0 + g_1D + g_2D^2 + g_3D^3 + g_4D^4 + g_5D^5)$$

where γ_1 is the air kerma factor, γ_2 is the source activity, γ_3 is the dose rate constant, γ_4 is the anisotropy factor, and $g_0 + g_1D + g_2D^2 + g_3D^3 + g_4D^4 + g_5D^5$ is an approximation of the radial dose function in terms of the distance $D = \max(\|p_i - \tau_j\|, 2mm)$. The reason the distance between the points is truncated at 2 mm is to keep the dose rate function from blowing up at small distances. The parameters $g_0, g_1, g_2, g_3, g_4, g_5$ are source-specific

Parameter	Name	Value
γ_1	Air kerma factor	4.03
γ_2	Activity	8000
γ_3	Dose rate constant	1.12
γ_4	Anisotropy	0.98
g_0	-	0.989054
g_1	-	0.0081319
g_2	-	0.0035177
g_3	-	-0.00146637
g_4	-	0.000092437
g_5	-	0

Table 3.2: Dose rate parameters. Need to find units of values

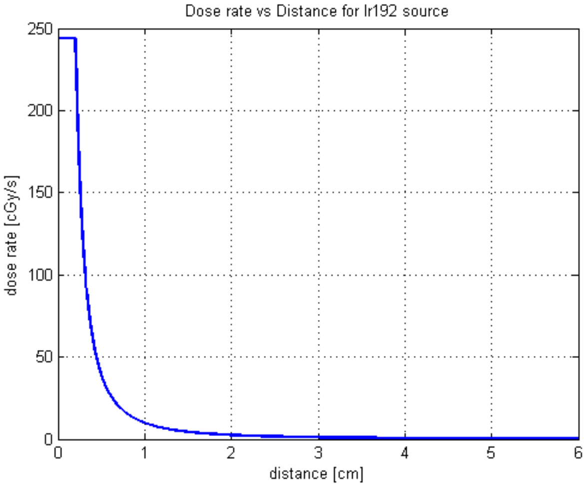

Figure 3.3: Dose rate vs. distance approximation used for ^{192}Ir.

parameters found from least squares regression. The $\frac{1}{D^2}$ term is called the geometry factor and the $\frac{1}{3600}$ converts the dose rate units from cGy/hr to cGy/sec. Table 3.2 shows the relevant values used for this study. Figure 3.3 shows the dose-rate vs. distance curve using these values.

The prescription dose used in this study was 950 cGy, and the dosimetric criteria used in this study can be found in Table 3.3. The specifications for the $\mathrm{V}_{100}^{\mathrm{Prostate}}$, $\mathrm{V}_{125}^{\mathrm{Urethra}}$, $\mathrm{V}_{75}^{\mathrm{Rectum}}$, and $\mathrm{V}_{75}^{\mathrm{Bladder}}$ are defined by RTOG-0321 [46]. RTOG specifies that the $\mathrm{V}_{125}^{\mathrm{Urethra}}$ be much less than 1 cm^3. We interpreted this to mean less than 0.1 cm^3. The $\mathrm{V}_{150}^{\mathrm{Prostate}}$ is not explicitly constrained by the RTOG-0321 protocol. The $\mathrm{V}_{150}^{\mathrm{Prostate}}$ is restricted by the homogeneity

Index	Requirement
$\mathrm{V}_{100}^{\text{Prostate}}$	$\geq 90\%$
$\mathrm{V}_{150}^{\text{Prostate}}$	$\leq 45\%$
$\mathrm{V}_{125}^{\text{Urethra}}$	≤ 0.1 cm^3
$\mathrm{V}_{150}^{\text{Urethra}}$	$= 0$ cm^3
$\mathrm{V}_{75}^{\text{Rectum}}$	≤ 1 cm^3
$\mathrm{V}_{100}^{\text{Rectum}}$	0 cm^3
$\mathrm{V}_{75}^{\text{Bladder}}$	≤ 1 cm^3
$\mathrm{V}_{100}^{\text{Bladder}}$	$= 0$ cm^3
$\mathrm{V}_{200}^{\text{Body}}$	$= 0$ cm^3

Table 3.3: IPIP dosimetric criteria

index (HI), where:

$$\mathrm{HI} = \frac{\mathrm{V}_{100}^{\text{Prostate}} - \mathrm{V}_{150}^{\text{Prostate}}}{\mathrm{V}_{100}^{\text{Prostate}}}.$$

It is generally preferred that HI ≥ 0.6. However, lower values of HI are acceptable if they allow for higher CTV coverage. For this study, we constrain HI to be greater than 50% to maintain some control over the $\mathrm{V}_{150}^{\text{Prostate}}$ while not being overly restrictive. Since we expect target coverage over 90%, this restriction on HI can be enforced with $\mathrm{V}_{150}^{\text{Prostate}} \leq 45\%$.

The restrictions that the $\mathrm{V}_{150}^{\text{Urethra}}$, $\mathrm{V}_{100}^{\text{Rectum}}$, $\mathrm{V}_{100}^{\text{Bladder}}$, and $\mathrm{V}_{200}^{\text{Body}}$ be equal to 0 are not specified by RTOG but are considered preferable when possible at our clinic. The preference of the $\mathrm{V}_{200}^{\text{Body}}$ comes from the desire to keep hot spots localized within the CTV. We also added the restriction that $\mathrm{V}_{75}^{\text{Bulb}}$ must be less than 1 cm^3, and that $\mathrm{V}_{100}^{\text{Bulb}}$ be 0 cm^3. These restrictions are not standard but were added to give some control over the dose to the bulb.

The parameters used for IPIP to reflect these dosimetric criteria can be found in Table 3.4. The values of 8 and 83 for U_s represent the number of dose points in 0.1 cm^3 and 1.0 cm^3, respectively, based on our grid spacing, which was 2mm in the x- and y-direction and 3 mm in the z-direction. The value of U_s for the $\mathrm{V}_{150}^{\text{Prostate}}$ is 45% of the number of dose points in the prostate. The dose to prostate dose points should be unrestricted so we have used an unrestrictively high number of 20000 cGy. We did this to avoid using infinity, which creates numerical problems with our optimization solver. The dosimetry in our results show no dose points received this dose level for any patient.

It is important to note that since target coverage is being maximized in the objective function, we do not explicitly enforce the constraint that target coverage be higher than 90% of the prostate volume. Relaxing this restriction (but still checking for it later), has one major advantage: if it is physically impossible to meet dose objectives, it allows IPIP and IPIP-H to return a solution with target coverage less than 90% rather than return infeasible. If the model just returns infeasible, then it is difficult to determine which objective caused the infeasibility (i.e. which objective is the most restrictive) and it does not allow one to determine how close the dose objectives are to being achievable. For instance, if 90% coverage

Organ	s	R_s	M_s	U_s
Prostate	0	950	20000	N_0
Prostate	1	1425	20000	.45 N_0
Urethra	2	1140	1425	8
Rectum	3	712	950	83
Bladder	4	712	950	83
Body	5	1900	1900	0

Table 3.4: IPIP experimental parameters

is impossible in the context of the other dose objectives, but 89% is possible, then the implant is *almost* able to meet the dose objectives, which is a useful fact to know. Conversely, if only 70% coverage is possible, then the implant is not close to meeting the dose objectives. Although meeting standard dose objectives from a well-placed implant is likely to be possible, knowing how close an implant is to meeting dose objectives is an important component of evaluating a needle configuration during needle planning, which is the topic of the next chapter.

3.3.5 Method Evaluation

We attempted to compute dose plans for each patient in our data set using using (1) IPIP, (2) the linear program relaxation of IPIP, IPIP-R, (3) IPIP-H, and (4) Inverse Planning Simulated Annealing (IPSA), a clinically deployed dose planning system used worldwide. We used a Lenova ThinkPad with an Intel i5-2410M processor and 4GB of RAM, and the Windows 7 64-bit operating system, to perform all calculations. Matlab R2011a (Mathworks Inc.) was used to perform all dose point, dose rate, and dosimetric index calculations. All linear and integer programming calculations were done using the Matlab interface for CPLEX 12.4 [21], an industrial standard for solving large scale linear and integer programs.

When optimizing the dose distribution using IPIP, CPLEX was set to terminate in 30 minutes of running time if an optimal solution was not found yet. If an optimal solution was not found, the solution with the best objective function value was returned. Since target coverage was not explicitly restricted in any of the optimization models, the zero solution, i.e. the dose plan where all dwell times are zero seconds, was always feasible and found at the root node. Therefore, a contingency for the case where a feasible solution was not found in 30 minutes was not required. Except for the runtime limitation, all CPLEX parameters were left at their default settings.

We recorded the values of the dosimetric indices which were controlled in IPIP and their compliance with the specified dose objectives. For IPIP-H, the difference between relaxed coverage and real coverage was recorded. The optimization time and optimality gap was recorded.

We compared the performance of IPIP-H with the linear program formulation by Alterovitz et al, 2006 [2] of the IPSA model. Comparison was made in terms of compliance with dose objectives and runtime. Since IPSA has been clinically validated for over a decade, it serves as a good baseline for performance standards. Each instance of IPSA was run only once for each patient. There was no manual fine tuning of optimization parameters if dose objectives were not met. However, this would be done in practice when treating patients. The penalty weights used were the standard class solution found in Lessard et al, 2004 [67].

3.4 Results

CPLEX found a non-zero feasible solution for 12/18 patient cases. Of these, only 8/12 had target coverage over 90% and only 2/12 non-zero solutions were optimal. Generally, the upper bound was close to 100% at the root and was not improved much during branch and bound. At termination, the 12 non-zero solutions had a minimum upper bound of 97%.

Dose plans computed from IPIP-R met target coverage for every patient case with a mean target coverage of 99%. However, it did not meet all the OAR constraints for any of the patient cases. This makes intuitive sense since IPIP-R only seeks to maximize target coverage, and counts OAR dosimetric indices inaccurately to its advantage, as described earlier in the second paragraph of Section 3.3.2. Since IPIP-R is a linear program, CPLEX could solve it fairly quickly - between 6 and 40 seconds.

For each patient, IPIP-H found a dose plan which satisfied all the dose objectives, including the target coverage requirement. This confirms the results for the 20 patient cases from Siauw et al. 2011. The run time ranged from 3 to 27 seconds with an average of 13 seconds. This was faster than the times recorded for IPIP-H in Siauw et al, 2011, most likely due to the substitution of Mosek [3], the optimization software used for the previous study, with CPLEX. The average target coverage was 97%. The average difference between the true target coverage and the relaxed target coverage used by IPIP-H was 3%.

IPIP-H found a higher quality solution than IPIP running for 30 minutes (i.e. higher target coverage within OAR dose constraints) in 13/18 cases. The absolute difference in target coverage between the IPIP-H solution and the upper bound found in IPIP after 30 minutes of running time was between 1 and 5% target coverage. Thus, IPIP-H was always within 5% of the optimal solution for IPIP. It is worth noting that using IPIP-H to find an initial feasible solution for branch and bound, combined with a custom method of reducing the upper bound, is a promising paradigm for quickly finding lower bound solutions for IPIP.

IPSA did not find a solution which met all the dose objectives on the first iteration for any of the trials. However, if the constraint on the $\mathrm{V}_{200}^{\mathrm{Body}}$ was ignored, then 9/18 out of the patients would have been compliant. The average run time for IPSA was 4 seconds. However, half of the cases would have needed at least one re-optimization to meet dose objectives.

The median dosimetric indices over the patient data set for each of the algorithms is presented in Table 3.5. The median doses to the penile bulb and body were omitted because they were zero for every algorithm. The table shows that the median dose to each of the organs was approximately the same when comparing IPIP-H and IPSA. However, since IPIP-

Alg.	$V_{100}^{\text{Prostate}}$	$V_{150}^{\text{Prostate}}$	V_{75}^{Bladder}	V_{75}^{Rectum}	V_{125}^{Urethra}
IPIP	91	26	0.3	0.1	0.0
IPIP-R	99	44	1.4	1.3	0.4
IPIP-H	97	36	0.8	0.7	0.1
IPSA	98	29	0.8	0.6	0.1

Table 3.5: Median dose to organs for IPIP, IPIP-R, IPIP-H, and IPSA.

H always complied with the dose objectives, it seems that IPIP-H allocated dose better than IPSA. As expected IPIP-R had the hottest plans with the highest $V_{100}^{\text{Prostate}}$ and $V_{150}^{\text{Prostate}}$. The median dose to the bladder, rectum, and urethra were all above their respective limits. A table of the dosimetric results for each patient can be found in Appendix A.1. A prostate, urethra, bladder, and rectum DVH is shown for each patient in Appendix A.2 for future reference.

3.5 Discussion

It is clear that solving IPIP is not a reliable method of dose planning. Allotting 30 minutes of planning time would be a severe interruption to the workflow (generally clinicians want a solution within one minute), and even then may not return a satisfactory dose plan. Although IPIP-R can be solved quickly, it does not comply with the specified objectives. IPIP-H, which is based on linear programming, is fast and can meet standard dose objectives. Therefore, it is a promising tool for clinical use. However, it also has several limitations which are discussed next.

Since every dose objective can be represented as a constraint in IPIP, the objective function is free to be interchanged with other objectives which the user may choose to induce a desired behavior. For example, some patients are successfully treated for prostate cancer, but later the cancer returns. In such cases, referred to as recurrences cases, the urethra requires additional protection, and it would be useful to minimize the V_{125}^{Urethra}, subject to a minimum coverage constraint. It may also be interesting to minimize the total treatment time. Both these objectives, and others, can be represented in IPIP. However, there are two reasons why the objective function cannot be interchanged in IPIP-H. First, and most obviously, IPIP-H requires the use of the objective function to meet target coverage. Second, maximizing target coverage is a fundamental mechanism that allows IPIP-H to perform well. Specifically, it relies on the fact that the first iteration of HR will produce a solution which is too hot, which may not be the case if another dosimetric index is optimized in the objective. Therefore, IPIP-H can only find high quality solutions for IPIP when target coverage is maximized, and no other dosimetric index is lower bound. Other objectives will not work.

Another limitation of IPIP, IPIP-H, and dose planning based on dosimetric indices in general is that they will tend to utilize all the dose allowance to OAR, even for small improvements to the objective. For example, a dose plan with target coverage of 99% and $\mathrm{V}_{75}^{\mathrm{Bladder}} = 1$ cc will be chosen over a dose plan with 98% coverage and $\mathrm{V}_{75}^{\mathrm{Bladder}} = 0$ cc. To catch this case, the user would have to input the constraint $\mathrm{V}_{75}^{\mathrm{Bladder}} \leq 0.5$ cc and check if the problem was feasible. If it is feasible, then $\mathrm{V}_{75}^{\mathrm{Bladder}} \leq 0.25$ cc could be checked and so on. Thus, directly optimizing according to dose-volume constraint is still, in a way, iterative. However, the iterations are done through dose volume constraints, which all users would be familiar with and have intuitive meaning, rather than system-specific optimization parameters.

Finally, IPIP-H can fail to find a solution which meets target coverage, even when one exists. Specifically, IPIP-H creates constraints which cool the first iteration of HR enough such that OAR sparing constraints are met. However, this cooling effect is heuristic and not necessarily optimal. Thus, it may add constraints in a way that cool the solution too much (i.e. not meeting target coverage) even when it is physically possible given the arrangement of dwell positions. At least with the standard objectives, with a well placed needle configuration, it seems that this case is unlikely.

The quality of possible dose plans is governed by the quality of the needle implant because dose plans are produced from the dwell positions in the needle implant. If there are dwell positions evenly spaced throughout the prostate volume, then dose planning is likely to be much easier, and dose plans of high quality are likely to exist and be easily found. If there are a regions of the prostate that do not have any needles nearby, then it may be difficult or impossible to cover that volume with the prescription dose within the limits to OAR dose. The implants used in this experiment were executed by Dr. I-Chow Hsu, a radiation oncologist with a specialization in brachytherapy and an expert in freehand needle insertion. A study should be done to determine the effectiveness of IPIP-H on needle configurations which have been placed by a less experienced physician.

This brings to the foreground the most important point of this chapter: the use of IPIP-H as an evaluation tool for needle configurations. In the patient cases used in this study, and in general practice which uses 16-18 needles, most needle configurations are likely to be satisfactory relative to standard dose objectives because the standard dose objectives are designed to be achievable in the vast majority of cases using the standard clinical workflow. However, the fewest number of needles required to meet dose objectives, and where these needles should be placed are still open questions. To answer these questions, an evaluation tool is necessary for the quality of needle configurations. Towards this goal, let a needle configuration be called satisfactory relative to a set of dose objectives if it can meet those dose objectives, and unsatisfactory otherwise. IPIP-H is an ideal tool for this purpose because it is (1) fast, (2) guaranteed to meet OAR objectives, and (3) can fail only in meeting target coverage. Specifically, if IPIP-H can find a dose plan which meets target coverage, then the needle configuration is satisfactory since all other dose constraints are guaranteed to be met. If not, then the target coverage is a measure of how close the needle configuration is to being satisfactory. For example, if target coverage is close to 90%, then the needle configuration is almost satisfactory. This kind of evaluation for needle configurations is not possible with multi-objective dose planning methods because they can fail in any of the dose criteria, and the inability to find a dose plan which meets dose criteria, even after many iterations of parameter tuning, does not guarantee that one does not exist. This will become especially problematic when evaluating needle configurations which are on the borderline between being unsatisfactory and satisfactory. The use of IPIP-H as an evaluation tool for needle configurations is an important component of the next chapter on Needle Planning.

3.6 Conclusion

We developed Inverse Planning by Integer Program, IPIP, an integer program that directly optimizes dosimetric indices. We showed that CPLEX, using the standard branch and bound parameters, could not reliably compute solutions for IPIP within 30 minutes. Assuming that maximizing target coverage is the objective function and that target coverage is the only lower bound constraint in the set of dose objectives, we developed a linear programming heuristic called IPIP-H which is fast and produces high quality (i.e. near optimal) feasible solutions using the RTOG-0321 standard dose objectives and needle configurations implanted by an expert. We showed that IPIP-H could outperform IPSA in terms of meeting dose objectives.

Since IPIP-H is fast, always produces a solution, meets OAR objectives, and can only fail in meeting target coverage, we use it as an evaluation tool for needle implants in our work on needle planning. This work is presented in the following chapter.

Chapter 4

Skew-Line Needle Planning using NPIP

4.1 Overview

There are many desirable properties for a needle configuration. The number of needles, number of collisions with organs such as the penile bulb, the difficulty of implantation, and the robustness to placement uncertainties are a few important factors. The most important property is that the needle plan should support a dose plan which meets the dose objectives set by the physician. However, optimizing both the dose distribution and the needle configuration simultaneously has been shown to be difficult to optimize using integer programming [52], even when choosing from a relatively few number of needle position choices (i.e. ≈100) and optimizing the dose distribution using dose penalties, which does not guarantee compliance with dose objectives.

This chapter expands on work presented in Siauw et al, 2012 [103], which developed Needle Planning by Integer Program (NPIP). NPIP is an optimization algorithm for computing patient-specific, feasible, needle configurations for HDR brachytherapy. NPIP uses a novel spatial coverage approach to compute a needle configuration such that the entire prostate volume is some distance, δ, from at least one needle in the configuration. Here δ is a user-specified parameter. We show that solving this spatial coverage problem can be represented as a set cover integer program.

For this study, we represented needle configurations as arrangements of skew lines (non-intersecting, non-parallel lines) rather than the parallel set enforced by rigid templates. NPIP is guaranteed to compute a needle configuration which does not penetrate any OAR and does not contain needles which intersect each other. We had NPIP choose a needle configuration from a candidate set of over 2500 needles, which is much larger than that used by Karabis et al, 2009, [52]. We showed that given the correct parameter, δ, NPIP could compute needle configurations which met dose objectives and used fewer needles than the standard number used by physicians, which is 16. We also present a sensitivity analysis which showed that

NPIP-generated needle configurations are robust to placement disturbances on the order of those expected by current robotic needle insertion devices. Limitations and future work are discussed.

4.2 Background

There has been some work in optimizing needle configurations (needle planning) for HDR brachytherapy [52]. In Karabis et al, 2009 [52], needles were chosen for insertion from a virtual set of candidate needles computed from the standard HDR template. This candidate needle set has on the order of 100 needles. An integer programming model was formulated which would simultaneous optimize the needle configuration and the dose distribution. For this study, the dose distribution was optimized using penalty weights, which does not necessarily comply with dose objectives. However, standard branch and bound could only find a feasible, but suboptimal, solution in one out of the three cases, even given 36 hours of running time.

Although our methodology also uses a candidate needle set, the set of needles considered for this study include skew-line needles at a much finer discretization of insertion points. Thus our method must handle a candidate needle set which is much larger than considered before. Also, since our candidate needle set includes non-parallel options, our method must be able to ensure that the chosen needle configuration does not contain needles which come within a needle-diameter's distance from each other (i.e. collide).

One solution to the computational difficulty of simultaneously optimizing the needle configuration and dose distribution is to break the problem into a sequence of problems: needle planning, then dose planning. This allows us to deal with two simpler problems and to take advantage of our previous work in dose planning using IPIP. However, placing needles without utilizing any knowledge of dose planning requires some approximation as to what constitutes a valid needle configuration. Our work uses a novel spatial coverage model to ensure that the entire prostate volume is near to at least one dwell position in the final needle configuration. The assumption is that volume near dwell positions has the opportunity to receive enough dose, which is the primary therapeutic objective of treatment. Clearly what constitutes "close enough" is debatable and is controlled by a user parameter. The focus of this study is on the selection of this parameter and its effect on the final needle configuration.

The main application of needle planning for this dissertation is robot-assisted needle insertion. Needle planning makes it possible to compute a needle configuration without the biases and limitations of a physician, e.g. clinical experience and placement precision, respectively. We elaborate on these ideas further in the Discussion section.

Term	Description
E	Entry zone vertices
$\mathcal{E}$	Convex hull of E
C	Set of target contour points
T	Set of projected points above target
$\mathcal{T}$	Convex hull of T
N	Set of candidate needle
$\mathcal{N}$	Chosen set of needles to be inserted
P	Coordinate of $p_i \in P$
d_j	Coordinate of j^{th} dwell position
δ	Needle coverage radius
x_k	Binary variable for $n_k \in N$
X	Set of pairs (k, ℓ) s.t. $n_k, n_\ell \in N$ collide
$K_i(\delta)$	Set of needles that cover p_i given δ
$I(\delta)$	Set of points that can be covered by N given δ

Table 4.1: NPIP terms

4.3 Method and Materials

4.3.1 System Formulation

In this section, a system for computing skew-line needle configurations is described. The system consists of a candidate needle generation component, a needle selection component, and a dose planning component. The candidate needle set is a large, virtual set of needles that is representative of the needles that are feasible for insertion. The needle selection component chooses a small subset of the candidate needles for insertion. The needle selection component should choose the needles in such a way as to be able to meet dose objectives during dose planning and such that the chosen needles do not intersect. The dose planning component is used to verify that the chosen needle configuration can meet dose objectives. Since integer programming is a central component of this system, we name it Needle Planning by Integer Program (NPIP). The components of NPIP are described in the following section. For reference, the notation used is given in Table 4.1.

4.3.1.1 Candidate Needle Set Generation

The candidate needle set generation component assumes it is possible for the user to define an entry zone where needles can be inserted into the body. In practice, this entry zone would be on the perineum and its size and shape would be governed by the desire to avoid hitting hard obstacles such as the pubic arch. Defining the entry zone in practice presents some challenges which are discussed in Section 4.5.

Let Ω denote the set of possible needles modeled by line segments which are (1) are straight, (2) originate inside the entry zone, (3) have some length within the prostate volume, and (4) do not have any length within a healthy organ volume. The candidate needle set generation component attempts to create a representative subset of Ω. We have explored two methods of generating candidate needle sets. The first is a naive method based on a manually generated needle configuration presented in Cunha et al, 2009 [23] which was designed to avoid the penile bulb. This method is computationally and geometrically simple and so convenient to generate, but does not represent Ω well. The second method generates a more representative set, but is more computationally demanding and can generate many false positives during one of the steps (i.e. generates needles which later turn out not to be in Ω). These two methods are described next. Eventually, we would like a full mathematical representation of Ω, either analytically or up to some entry point and angulation precision. We leave this to future work. The advantages of such a method are discussed.

4.3.1.1.1 Method 1: This method has a rigid method of defining the entry zone, which is not based on the anatomy of the patient. Specifically, let the entry zone be defined by a square some distance below the most inferior slice of the prostate (presumably also below the penile bulb), centered at the x-y center of mass of the prostate. Let the size of the square and the distance below the prostate be defined by two user parameters, α and β, respectively.

A candidate needle set is generated for this entry zone by taking the union of five needle sets: a parallel set and four conical sets. The parallel needle set consists of evenly spaced needles going exactly in the z-positive direction (from the perineum towards the head). The spacing of needles in the x- and y-direction is defined by another user defined parameter, ρ. Each needle begins at the entry zone and ends at the most superior slice of the prostate directly superior to the entry point. Every needle within one conical set share the same entry point at one of the corners of entry zone. Each needle begins at the entry point and ends at a point on a $\alpha \times \alpha$ square at the most superior slices of the prostate, centered at the x-y centroid of the prostate, and gridded with ρ spacing in the x- and y-direction. The conical set was inspired by the conical needle configuration from Cunha et. al, 2009, which was designed to have the majority of needles avoid the penile bulb. An example candidate needle set generated from this method is shown in Figure 4.1.

The candidate needle set is then truncated by removing all needles that intersected an organ at risk before they intersect the prostate. Needles that do not intersect the prostate (i.e. have no dwell positions within the prostate) are also excluded. Needles that intersect the prostate before intersecting an organ at risk are truncated at the last dwell position within the prostate. A needle is assumed to intersect an organ if, at any contour slice of that organ, the interpolated x-y position of that needle is contained in the slice.

This method of generating candidate needles has several limitations. First, it only captures a very small subset of Ω and is not nearly representative. Second, the method is highly parametric, with no real intuition as to how the parameters should be chosen. Finally, the

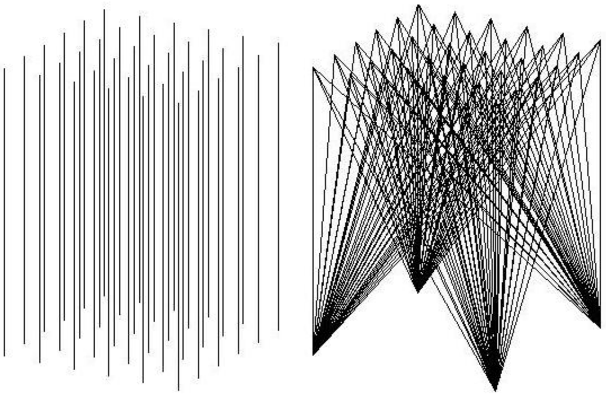

Figure 4.1: Method 1 of generating candidate needle sets used the union of a parallel needle set (left) and non-parallel needle set (right).

gridding induces artifacts in the candidate needle set which may harm the performance of the overall algorithm. For example, many needles in the conical set have the same origin, which is undesirable because it could cause collision problems when inserting the needle configuration.

4.3.1.1.2 Method 2: This method describes a more principled approach to generating the candidate needle set. In particular, we reduce the number of user-parameters, and use random sampling to avoid artifacts that may be caused by gridding.

Let $z*$ be the z-coordinate of the most superior slice of the target. Let E be the set of vertices of the entry zone, $\mathcal{E}$ be the set of points in the convex hull of E, and C be the set of contour points of the target. The candidate needle set is generated according to the anatomy of the patient in four steps. (1) Compute $T = \{t_{ij} : t_{ij} = e_i + \left(\frac{z* - e_i^z}{c_j^z - e_i^z}\right)(c_j - e_i), \forall i \in E, \forall j \in C\}$. Here, e_i^z and c_j^z are the z-components of e_i and c_j respectively, and t_{ij} is the intersection of the line going through e_i and c_j with a plane perpendicular to the z-axis at $z*$. (2) Compute $\mathcal{T}$, the convex hull of T. (3) Generate a large number of needles by uniformly sampling a point in $\mathcal{E}$ and in $\mathcal{T}$ and creating a line between them. Each line represents a single needle. Sampling inside a convex polygon is discussed below. (4) Remove needles that do not intersect the target, remove needles that intersect a healthy structure before intersecting the target, and truncate needles that intersect a healthy structure after the prostate at a point before the intersection. Figure 4.2 shows a diagram of the terms used for generating candidate needle sets.

Uniform random sampling inside a convex polygon is described in Devroye [27]. First, the convex polygon is subdivided into triangles (i.e. a triangle mesh). Then, a triangle in the mesh is randomly selected with probability proportional to its area. Finally, a point is chosen from a uniform distribution inside the triangle.

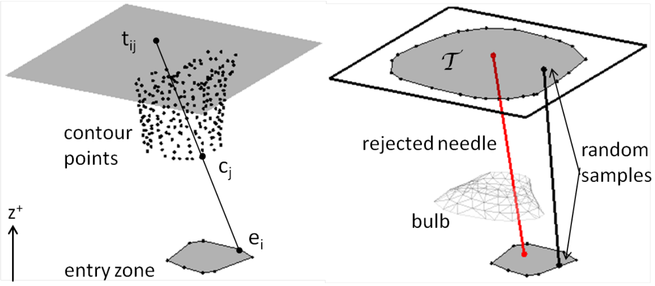

Figure 4.2: A set of candidate needles was generated according to the anatomy of the patient. Given the vertices of an entry zone, E, and the contour points of the prostate, C, a set of points was computed, T, such that $t_{ij} \in T$ was the intersection of the line containing $e_i \in E$ and $c_j \in C$ with a plane perpendicular to the z-axis at the most superior slice of the target. The convex hull of T was called $\mathcal{T}$. A needle was generated by randomly sampling a point in $\mathcal{E}$, the convex hull of E, and in $\mathcal{T}$, and generating a line connecting them. Needles that intersected a healthy structure were removed.

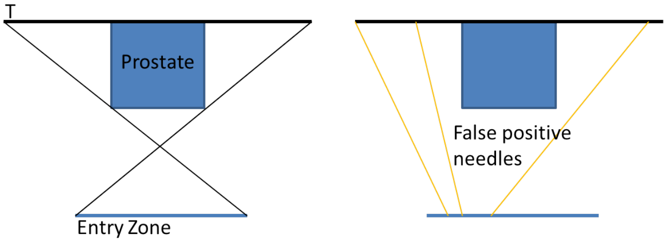

Figure 4.3: False Positive Example for Method 2 for candidate needle set generation.

This method has one parameter, the initial number of needles, and with enough needles, represents Ω much better than Method 1. However, there are still some limitations of this method worth mentioning. For instance, the T region may be very large if entry zone is fat and close to the target. Consider the two dimensional case shown in Figure 4.3. In this case, there are many possible needles which originate at the entry zone and end in T but do not intersect the prostate, i.e. many false positives. As a consequence, a huge number of initial needles would need to be generated to ensure that the needles that survive the rejection process still adequately represent Ω. Additionally, this method may also create many needles which are highly angled, which although possible, may be more difficult to implant than needles which have a majority z-direction.

As in Method 1, needles which do not intersect the prostate or intersect an OAR are removed from the candidate needle set. Needles which intersect the prostate then an OAR are truncated at the last dwell position before the OAR intersection. All needles were truncated

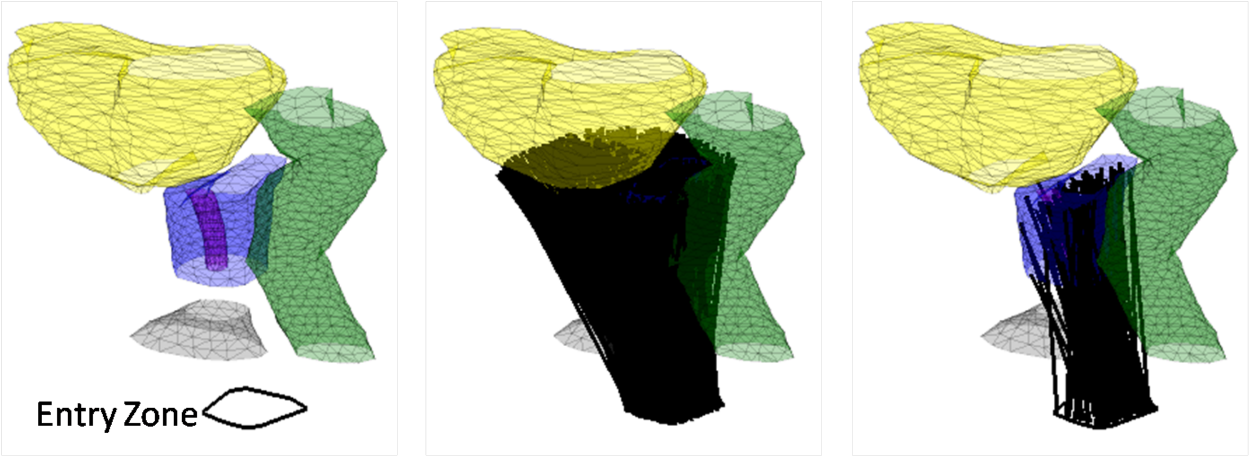

Inputs

Random Generation

Truncation

Figure 4.4: Given an anatomy data set and an entry zone for needle insertion, Method 2 generates many randomly sampled needles according to the process described in Figure 4.2, then eliminates needles which intersect an OAR or do not intersect the prostate. Needles are truncated at the last dwell position in the prostate or the last dwell position before an intersection with OAR.

at the last dwell position within the prostate. A diagram of the Method 2 candidate needle set generation scheme is shown in Figure 4.4.

4.3.1.2 Needle Selection

Let the candidate needle set be denoted by N, and let $n_k \in N$ be a needle defined by the set of dwell positions that belong to it. The goal of the needle selection component is to find an $\mathcal{N} \subset N$ that (1) meets dose objectives, (2) is collision free, and (3) uses as few needles as possible. However, meeting dose objectives is a complicated constraint to enforce, and therefore, our needle selection component only requires that $\mathcal{N}$ have good spatial coverage of the target. Some mathematical formalism to these criteria is given and it is shown that needle selection with these properties can be formulated as a set cover integer program.

First the spatial coverage of needles is formalized. Assume that the target volume has been discretized into a grid of evenly spaced points, P. Usually, a point $p_i \in P$ is referred to as a "dose points" or "dose control points" because it is used to evaluate the dose delivered to an organ voxel. However for this study the points are referred to as "cover points" to distinguish them from points used to control the dose distribution during dose planning. For a user-specified parameter, $\delta > 0$, n_k *covers* p_i if there exists a $d_j \in n_k$ such that $||p_i - d_j|| < \delta$, where d_j is a dwell position in n_k. Since δ defines the region that needles can cover, it is referred to as the needle coverage radius.

To achieve good spatial coverage of the target, it is desired that every point be covered by at least one needle in $\mathcal{N}$. However, some points may not be coverable by any needle in N, making this criterion impossible since $\mathcal{N} \subset N$. Instead, let $I(\delta)$ be the indices for the points which are coverable by at least one needle in the candidate set given δ: $I(\delta) = \{i : \exists n_k \in N \text{ that covers } p_i \text{ given } \delta\}$, and $\mathcal{N}$ is said to give good spatial coverage of the target if there is a needle in $\mathcal{N}$ which covers every point in $I(\delta)$. Clearly for this criterion to be reasonable, $I(\delta)$ must contain the vast majority of points in P. Our experiments showed that in almost every case, every point in P was indexed in $I(\delta)$, and in cases which this was not true, $I(\delta)$ contained at least 98% of the points in P.

To avoid needle collisions, constraints on the selection process are added that restrict only one needle in a colliding pair to be chosen. Formally, let L_k be the line segment between the entry point and the most superior dwell position of n_k, and let n_k and n_ℓ collide if the minimum distance between L_k and L_ℓ is less than the diameter of a needle, γ. Define X to be the set of colliding pairs, (k, ℓ), between all the needles in N. If $\mathcal{N}$ is selected such that for every $(k, \ell) \in X$, only n_k or n_ℓ (but not both) can be chosen, then $\mathcal{N}$ is guaranteed to be collision free.

Minimizing the size of $\mathcal{N}$ subject to the coverage constraints defined by δ and the collision constraints defined by γ can be formulated as a BIP. Let x_k be an indicator variable for n_k with the following behavior:

$$x_k = \begin{cases} 1 & \text{if } n_k \text{ is chosen to be in the configuration,} \\ 0 & \text{otherwise,} \end{cases}$$

and let the parameter

$$K_i(\delta) = \{k : n_k \text{ covers } p_i \text{ given } \delta\}$$

be the set of needles that cover p_i given δ. Then needle selection, N(P,N,δ), is the following integer program.

$$\begin{array}{llll} & \text{minimize} & \sum_k x_k & \\ \mathrm{N}(P,N,\delta) & \text{subject to:} & \sum_{k \in K_i(\delta)} x_k \geq 1, & \forall i \in I(\delta), \quad (4.1) \\ & & x_k + x_\ell \leq 1 & \forall (k,\ell) \in X, \quad (4.2) \\ & \text{and} & x_k \in \{0,1\}, & \forall k. \quad (4.3) \end{array}$$

The objective function ensures that the fewest number of needles are chosen. Constraint 1 ensures that every point that can be covered is covered by at least one needle and makes N(P,N,δ) a set cover integer program, constraint 2 ensures that only one needle in a colliding pair is chosen, and constraint 3 enforces the binary requirement on the definition of x_k.

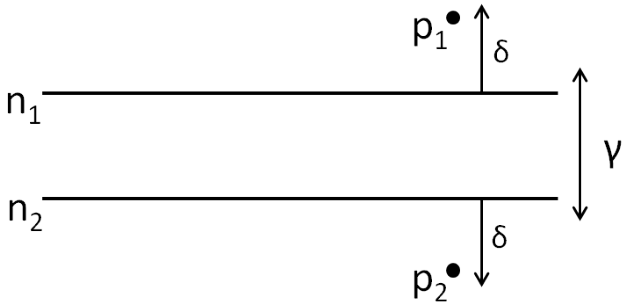

Figure 4.5: Test case where coverage and collision-free requirement cannot both be met, even if every point can be covered by at least one needle in the candidate needle set. Here, n_1 can cover p_1 and n_2 can cover p_2. However, since they are less than γ apart, only one needle can be chosen. Thus, the problem is infeasible.

Note that it is still possible for N(P,N,δ) to be infeasible, even if only the points in $I(\delta)$ are used for the coverage requirement, because it may be impossible to meet both the coverage and collision free requirement. A small test case in which this occurs is depicted in Figure 4.5. Future work may be able to identify points that cannot be covered because of this restriction and exclude them from $I(\delta)$.

4.3.1.3 Dose Planning

The primary objective of a needle configuration is to meet dose objectives. However, NPIP relaxes this requirement through a spatial coverage approximation defined by δ. Therefore, it is possible to choose δ such that dose objectives cannot be met by $\mathcal{N}$. For example, δ can be chosen large enough such that $|\mathcal{N}| = 1$, and a needle configuration with one needle will not meet reasonable dose objectives. Therefore, a dose planning tool is required to verify that $\mathcal{N}$ meets dose objectives.

Inverse Planning by Integer Program (IPIP) [102] was chosen as the dose planning tool for this system because it is guaranteed to meet all healthy-tissue sparing constraints while making target coverage as high as possible. These properties of IPIP allow needle configurations to be evaluated by a single metric, target coverage, without the need to interpret the relative value of competing dose objectives. Specifically, if IPIP produces a dose plan for a needle configuration that meets the target coverage requirement, then the needle configuration meets all dose objectives, and the needle configuration can be declared satisfactory. Otherwise, it is unsatisfactory.

4.3.2 Patient Data Sets

The same patient data sets used for the previous study were used for this study. A model of the entry zone was not available for this experiment because contouring it is not standard for HDR brachytherapy. Since an entry zone definition was critical for Method 2 of generating candidate needles, the entry zone was taken to be the convex hull of the needle locations at a plane perpendicular to the z-axis taken 2 cm below the most inferior slice of the bulb, which was an estimation of the skin surface. A 5 mm margin was added to the entry zone so that the region would not be tight to the needle locations used. The needle locations in the plane were interpolated between the closest dwell positions to the plane. This entry zone definition ensured that NPIP needle configurations only utilized the space used by the physician. The entry zone area ranged from 7 to 13 cm^2.

4.3.3 Method Evaluation

4.3.3.1 NPIP Evaluation

NPIP was used to compute needle configurations for each patient. For the candidate needle set generation component, the convex hull calculations were done using the Matlab function, `convhull`, and convex hulls were subdivided into triangles using the Matlab function, `delaunay`. The candidate needle sets were generated using Method 2 with 5,000 random samples. This number of initial candidate needles was found to produce consistent results across trials. Dwell positions were generated along each needle from the entry zone to the needle tip in 5 mm increments. Needles intersected an organ if at any contour slice of that organ, the interpolated x-y position of that needle was contained in the slice. The tip of each needle was truncated at the most superior dwell position in the target.

A needle diameter of $\gamma = 1.7$ mm was used to check collisions. This is the diameter of the needles used in our clinic.

Cover points were generated for the target using a uniform grid with 5 mm spacing in the x-y direction and 3 mm spacing in the z–direction. Since the density of the final needle configuration was controlled by the user-parameter, δ, it was expected that for a given δ, large prostates would generate more needles than small prostates. To make the results from different-sized prostates more comparable, δ was always chosen relative to the radius of a sphere with equivalent prostate volume. For conciseness, all δ values in this study were stated as a percentage of this radius. For each patient candidate needle set, a needle configuration was computed for δ values ranging from 25% to 50% in increments of 5%. Initial tests with NPIP showed that this range of δ produced an interesting range of needle configuration sizes. Specifically choosing δ less than 25% produced needle configurations with more than 25 needles, which was too many to be clinically relevant, and choosing δ more than 50% had too few needles to meet dose objectives.

Instances of N(P,N,δ) were solved using Matlab R2011a on a Lenova ThinkPad with an Intel i5-2410M processor and 4GB of RAM. The Matlab interface for CPLEX 12.4 [21] was used for the integer program optimization. There were six conditions to terminate the optimization process: (1) a provably optimal solution was found, (2) N(P,N,δ) was proven to be infeasible, (3) the optimization timer went over 10 minutes, (4) the memory usage exceeded 1 MB, (5) the number of branch and bound nodes exceeded 1000, or (6) the user aborted the process. In the third through sixth case, the feasible solution with the smallest objective function value was returned if one had been found already, and an error message was displayed otherwise. These conditions can be set in CPLEX using standard internal parameters. In practice, only a time limit needs to be specified for early termination. Other termination criteria were added for this study because there was no straightforward way to guarantee that the optimization solver would stop in 10 minutes, and more criteria were added to reduce the risk of running over time. These limitations are strictly due to the way CPLEX implements the time limit criterion. In a clinical deployment of this system, more extensive software should be written to ensure that the optimization solver stops when the timer expires. [1]

A dose plan was generated for every computed needle configuration using IPIP. The IPIP parameters and dosimetric criteria used in this study were the same as in Siauw et. al., 2011 [102], except with the added restrictions that $V_{75}^{\text{Bulb}} \leq 1$ cm^3 and $V_{100}^{\text{Bulb}} = 0$ cm^3. This criteria was added to give some control over the dose to the bulb, which was not included in the original IPIP study. The dose objectives used in this study conform to the specifications given in the RTOG-0321 dosimetric protocol [46]. Dwell positions were generated along each needle every 5 mm starting from the tip *for this study and the sensitivity analysis study.* Only dwell positions inside the prostate were used for dose planning.

For each patient, the number of candidate needles, candidate needle set generation time, number of collision pairs, and collision calculation time was recorded. For each instance of N(P,N,δ), the number of points not included in $I(\delta)$, the time to construct the CPLEX optimization model, solve time, termination status, and the optimality gap at termination were recorded. The number of needles and compliance with dose objectives were recorded for every needle configuration computed. To ensure the consistency of the results, the experiment was repeated five times.

For comparison, a dose plan was also computed from the physician implant using IPIP. The number of bulb punctures for each patient was also recorded.

4.3.3.2 Sensitivity Analysis

For the sensitivity analysis, the candidate needle set was generating using Method 1 with parameters: $\alpha = 5$ cm, $\beta = 4$ cm, and $\rho = 5$ mm. This produces an entry zone which is underneath, z-negative direction, the most inferior slice of the penile bulb for every patient case and is smaller than than the standard template for PPI brachytherapy, which is 10 cm × 10 cm. Dwell positions were generated along each needle starting from the entry zone, every 5 mm along the needle. Dwell positions outside the prostate were not used for this study.

---

[1] In the newest version of CPLEX, currently 12.4, this bug is fixed and run time restrictions are reliably enforced.

As in the previous study, δ was always chosen relative to prostate size. For this study γ was not used, or $\gamma = 0$, which means intersections were allowed. [2] The relationship between δ, the resulting number of needles, the prostate size of the patient, dosimetry, and dosimetry robustness was evaluated.

For each patient, a needle configuration was computed using NPIP for δ values ranging from 25% to 50% in increments of 2.5%. Initial tests with NPIP showed that smaller δ values produced very large needle configurations (e.g. 25 or more needles) and lower δ values gave very small configurations (e.g. 6 or fewer needles). Configurations with more than 25 needles are of no practical value given that state-of-the-art practice for HDR brachytherapy uses between 16-18 needles. Configurations with less than 6 needles were not able to meet dose objectives. The 2.5% increment was determined to be the smallest value for which subsequent increments produced differently sized needle configurations.

After computing a needle configuration for each patient-δ pair, a dose plan was computed using IPIP. For this study, we say a needle configuration meets dose objectives, if a dose plan can be computed from it, using IPIP, that meets dose objectives. A needle configuration is said to be *satisfactory* if it meets dose objectives.

The robustness of each needle configuration to placement errors was evaluated by perturbing each needle configuration 25 times. A perturbation was computed by adjusting the superior tip of each needle in the x– and y– direction according to a normal distribution with mean equal to 0 and standard deviation equal to 2.5 mm. The perturbation in the x– and y– direction was drawn independently. Since robots exist that can perform image guided needle placement within 2 mm of a target location [80], these perturbations are approximately, but greater than, the placement accuracy of existing robots. If a needle was perturbed such that it intersected an OAR, it was truncated to the last dwell position in the target. If a needle was perturbed such that it had no dwell positions in the target, then it was removed from the configuration. These assumptions are, again, very conservative because it is unlikely that, under some kind of image guidance, a robot or human-robot system would hit an OAR or miss the target completely.

New dose plans were recomputed on every perturbed needle configuration. The reasoning is as follows. In this study, perturbations represented uncertainty in needle placement. In the clinical HDR brachytherapy workflow, needle positions are determined via scanning after the needles are placed and after the placement uncertainty has been realized. Therefore, their placed positions are known during dose planning.

The robustness of needle configurations is difficult to quantify since it depends on many factors including, but not limited to, the strength and type of perturbation, patient anatomy, and the geometry of the nominal needle configuration. For this study, we say a needle configuration is *robust* if, over the perturbation set on that needle configuration, the average target coverage minus one standard deviation is more than 90%. Recall that OAR dose objectives will always be met with IPIP, regardless of the perturbed needle configuration.

For this study only, the integer program optimization was done using the Matlab interface

---

[2] At this stage of development of NPIP, we had not yet considered needle collisions.

for the Mosek Optimization Toolbox v.6 [3], a medium to large scale optimization package. The optimization solver for NPIP was set to terminate after 2000 subproblems. This corresponded to approximately one minute of running time. If an optimal solution could not be found before the termination criterion was reached, the feasible solution with the fewest needles was returned if one had been found, and an error message was thrown otherwise.

4.4 Results

4.4.1 NPIP Results

Despite the large variety of entry zone shapes, organ geometries and sizes, a needle configuration that used less needles than the physician, avoided the penile bulb, met all dose objectives, and was collision-free was computed for each patient. A needle configuration for each patient is shown in Figure 4.6. Results for each instance of N(P,N,δ) are summarized in Table 4.2. Each table entry contains the average number of needles in N(P,N,δ) over the experiments, the average target coverage of the computed needle configurations over each experiment, and the number of instances that were feasible. The number of feasible instances is marked in bold if there were any infeasible instances. Note that the dose planning component, IPIP, was guaranteed to meet all dose objectives except target coverage. Therefore if target coverage was over 90%, the needle configuration met all dose objectives.

The number of needles and the ability to meet dose objectives decreased as the needle coverage radius, δ, increased. This was expected because a higher needle coverage radius allows each chosen needle to cover more volume, which requires fewer needles to cover every point. The table suggests that in practice, 35% or 40% could be chosen for the value of δ to ensure that the final needle configuration met dose objectives. For $\delta = 35\%$, N(P,N,δ) was sometimes (but rarely) infeasible, but when feasible, could always meet dose objectives. For $\delta = 40\%$, N(P,N,δ) was always feasible and on average, always met dose objectives. However, there were a few cases where target coverage was less than 90%. The likelihood of infeasibility did not appear to be related to the size of the prostate.

After needles were discarded due to collision with organs, the number of candidate needles ranged from 510 to 2599. Generally for each patient, the standard deviation of the number of candidate needles over the experiments was within 10% of the average number of needles. The run time for generating the candidate needle set ranged from 17 to 55 seconds. Almost all the time generating candidate needle sets was spent checking for needle intersections with organs.

Collision checking needles found between 19611 and 352872 collision pairs, depending on the size of the candidate needle set. Collision checking took between 11 and 286 seconds.

There were 540 instances of N(P,N,δ) that were attempted: 18 patients $\times$ 6 δ-values $\times$ 5 experimental iterations. The number of cover points ranged from 463 to 1722. The number of points in $I(\delta)$ was less than the total number of cover points in 56/540 cases. The difference between the size of $I(\delta)$ and the size of P was between 1 and 9, and generally more

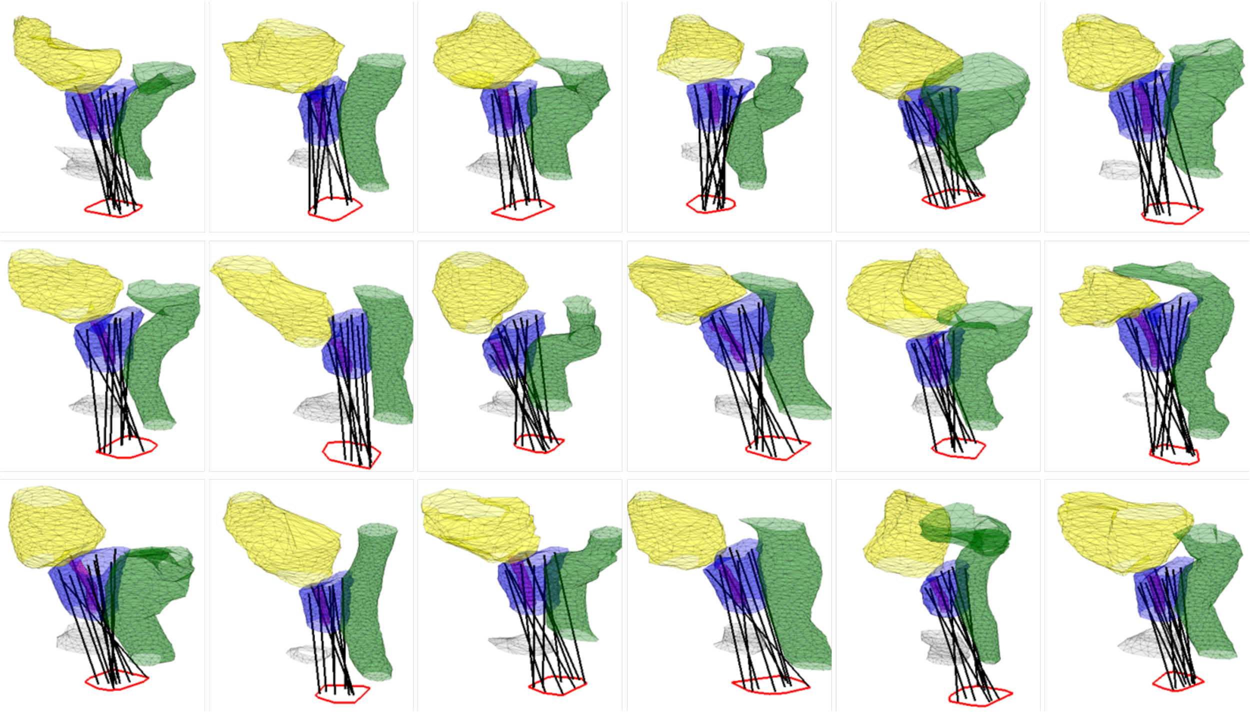

Figure 4.6: A needle configuration that used less needles than the physician, avoided the penile bulb, met all dose objectives, and was collision-free was computed for each patient. An example needle configuration is shown for each patient.

points were removed when δ was small. The number of infeasible instances of N(P,N,δ) was 40/540. Infeasibility was more likely for smaller needle coverage radii, which was expected since it was harder to meet coverage when the needle coverage radius was small. However, small needle coverage radii were also associated with needle configurations which contained more needles than were required to meet dose objectives. The runtime of N(P,N,δ) ranged from 1 second to 2329 seconds, although there was an outlier that ran for 8809 seconds.

The total end-to-end runtime ranged from 31 seconds to 2968 seconds, excluding the case where N(P,N,δ) took 8809 seconds. The average runtime was 286 seconds. For a given patient, the total NPIP running time was approximately constant across δ values.

The physician implant could always meet dose objectives. The average number of intersections with the penile bulb was 5 with a standard deviation of 3. Every patient case had at least one bulb puncture. It is important to note that the bulb is not visible to the physician under ultrasound, but can be contoured later on for dose planning. In Figure 4.7, a physician implant is shown as well as a 12 needle NPIP configuration for the same patient. The NPIP plans never intersected the penile bulb and met all dose objectives.

Given these results, needle configurations computed from NPIP have the potential to reduce side effects of HDR brachytherapy without compromising the high survival rate

Px	Target	Entry Zone	δ =25%			δ =30%			δ =35%			δ =40%			δ =45%			δ =50%		
1	23	11	23	99	5	15	99	5	12	97	5	9	95	5	8	93	5	7	90	5
2	26	9	19	99	5	13	98	5	10	96	5	8	93	5	6	86	5	6	81	5
3	27	12	20	99	5	14	98	5	10	93	5	8	91	5	6	84	5	5	75	5
4	28	10	25	99	5	17	98	5	12	95	5	9	94	5	8	90	5	6	85	5
5	28	13	22	100	5	14	99	5	10	97	5	8	94	5	7	89	5	6	84	5
6	30	7	–	–	**0**	18	98	**1**	15	97	**4**	11	95	5	9	93	5	8	87	5
7	34	9	–	–	**0**	20	98	**4**	15	95	5	11	93	5	9	91	5	7	81	5
8	34	9	22	100	5	16	99	5	11	98	5	9	94	5	8	91	5	6	87	5
9	34	12	20	99	5	15	98	5	11	96	5	8	96	5	7	89	5	6	83	5
10	38	8	33	99	**3**	22	98	**3**	15	97	5	11	95	5	9	91	5	8	87	5
11	40	12	24	100	5	16	99	5	12	97	5	9	96	5	7	90	5	6	89	5
12	42	9	32	99	**1**	20	98	5	14	97	5	11	95	5	9	92	5	7	92	5
13	43	10	29	96	5	18	94	5	13	92	5	10	90	5	8	86	5	7	86	5
14	48	12	24	96	5	17	95	5	13	94	5	10	91	5	8	87	5	7	81	5
15	49	11	–	–	**0**	19	97	5	14	96	**4**	11	95	5	9	92	5	8	88	5
16	61	10	33	99	**1**	20	97	**4**	15	96	**4**	12	95	5	10	92	5	7	87	5
17	69	13	25	97	5	16	96	5	12	95	5	10	93	5	8	91	5	7	88	5
18	85	11	27	97	5	18	96	5	13	95	5	10	92	5	8	91	5	7	86	5

Table 4.2: Table of results for N(P,N,δ). Each table entry contains (1) the average number of needles, (2) the average target coverage as a percentage of the target volume, and (3) the number of feasible instances of N(P,N,δ) over the NPIP experiments. Also included is the prostate size [cm^3] and entry zone size [cm^2]. Dose objectives could always be met by choosing $\delta = 35\%$ or 40%. These δ values corresponded to needle configurations with approximately 10-15 needles. Note IPIP guarantees healthy tissue sparing constraints. Therefore, if target coverage was met, then all dose objectives were met.

associated with this treatment modality. However, the execution of these needle configurations will probably require some form of robotic assistance. The use of robots has been been explored extensively for Prostate Permanent-seed Implant (PPI) brachytherapy [19, 28, 32, 33, 59, 60, 79, 90, 100, 108–111]. It is possible that this technology can be adapted to insert skew-line needle configurations for HDR brachytherapy.

4.4.2 Sensitivity Results

Using NPIP, 198 attempts were made to compute needle configurations for each patient-δ pair: 18 patients × 11 δ values. Although a needle configuration could be computed for every patient, 3/198 patient-δ pairs produced infeasible instances of NPIP. This condition occurred because the candidate needle set could not cover every cover point in the prostate for the smallest values of δ. Of the needle configurations computed, 67% were optimal, and smaller needle configurations were more likely to be optimal than larger ones. On average, the suboptimal solutions returned were within 2 needles of the optimal solution (i.e. gap $\leq$ 2 needles). The average runtime for NPIP was 50 seconds, and the median runtime was 20 seconds. Since the candidate needle set did not contain any needles that punctured the penile bulb, no NPIP needle configuration punctured the penile bulb or any other OAR.

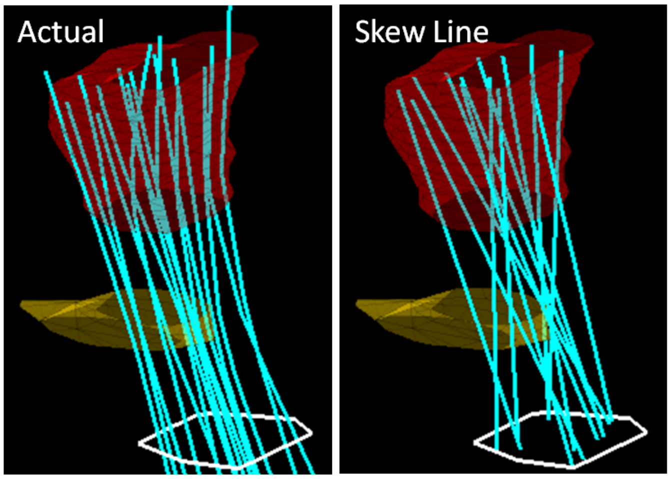

Figure 4.7: Actual implant (left) and a computationally generated skew line configuration (right). Needles are shown as cyan lines. The prostate is shown in red and the penile bulb in yellow. The entry zone for needles is contoured in white. The physician implant used 16 needles and intersected the penile bulb six times while the computed needle configuration used 13 needles and did not intersect the bulb. Both needle configurations met all dose objectives for the patient.

The relationship between the average number of needles across patients versus δ, as well as ± 1 standard deviation error bars, is shown in Figure 4.8. As expected, the number of needles decreased as δ increased. In other words, a larger spacing allowance resulted in fewer needles. Note that the relationship between the number of needles selected and δ does not contain any information about the quality of dose plans that could be generated from these needle configurations (i.e. some of these needle configurations could not produce viable dose plans). The standard deviation in the number of needles decreased as δ increased.

The relationship between the percentage of satisfactory and robust needle configurations versus δ is shown in Figure 4.9. The percentage of satisfactory needle configurations decreased as δ increased. This was expected since the number of needles was also decreasing. The largest δ value for which every needle configuration was satisfactory was $\delta = 40\%$. This corresponded to needle configurations with an average of 11 needles. However, only 11% of these needle configurations were robust. For $\delta = 32.5\%$, all but one of the needle configurations were robust. This corresponded to needle configurations with an average of 16 needles.

Unexpectedly, there were fewer robust needle configurations for $\delta = 30\%$ than $\delta = 32.5\%$, which implied that larger needle configurations were less likely to be robust. Closer inspection of the results showed three non-robust for $\delta = 30\%$ had target coverage over 89%, which is very close to the 90% cutoff. In other words given 1% leeway on the target coverage dose objective, all patients would have had robust needle configurations at $\delta = 30\%$.

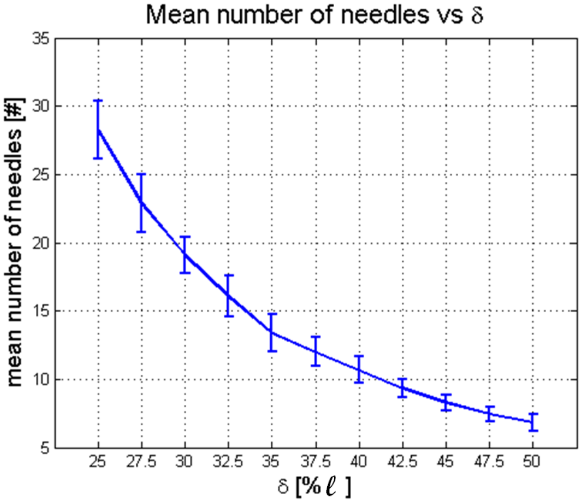

Figure 4.8: Average number of needles in NPIP needle configuration versus δ. Here δ is given as a percentage of the radius of a sphere with equivalent volume to the prostate volume of the patient. Specifically, $\ell = \left(\frac{3V}{4\pi}\right)^{\frac{1}{3}}$, where V is prostate volume. The error bars are ±1 standard deviation from the mean. A δ value of 32.5% resulted in needle configurations with the same average number of needles as is standard at our clinic, which is 16, and a standard deviation of 2 needles.

A robustness plot summarizing the results related to a patient anatomy data set (i.e. needle configuration, dose plan, and perturbation results) was created for each patient. An example is shown in Figure 4.10. On the x-axis is δ. The number of needles resulting from the given δ value is shown as a bar graph. Each bar is labeled with the number of needles in the needle configuration. Above the bar graph are filled dots signifying the target coverage for nominal needle configurations. Also above the bar graph is a plot of the mean target coverage over the perturbations of the given needle configurations vs. δ as well as ±1 standard deviation error bars. Recall that we consider a needle configuration to be robust if the mean target coverage minus one standard deviation over the perturbations is more than 90%. The robustness plot for each patient are shown in Appendix B.

In general, target coverage decreased as δ increased. The standard deviation for target coverage increased as the number of needles in a configuration decreased. This was expected since the contribution of a random large perturbation to a single needle influences the quality small needle configuration is more than a larger one. Patients with larger prostates tended to have a smaller standard deviation for target coverage than patients with smaller prostates, even for needle configurations with the same number of needles.

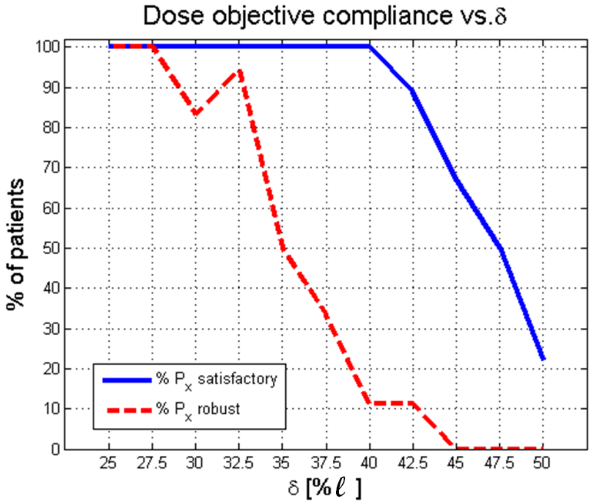

Figure 4.9: Percentage of satisfactory and robust needle configurations versus δ. A needle configuration is satisfactory if the nominal dose plan has target coverage more than 90%; it is robust if the mean target coverage over the perturbation set minus the standard deviation is more than 90%. The largest δ value for which every needle configuration was satisfactory was $\delta = 40\%$. This corresponded to needle configurations with an average of 11 needles. However, only 11% of these needle configurations were robust. For $\delta = 32.5\%$, all but one of the needle configurations were robust. This corresponded to needle configurations with an average of 16 needles.

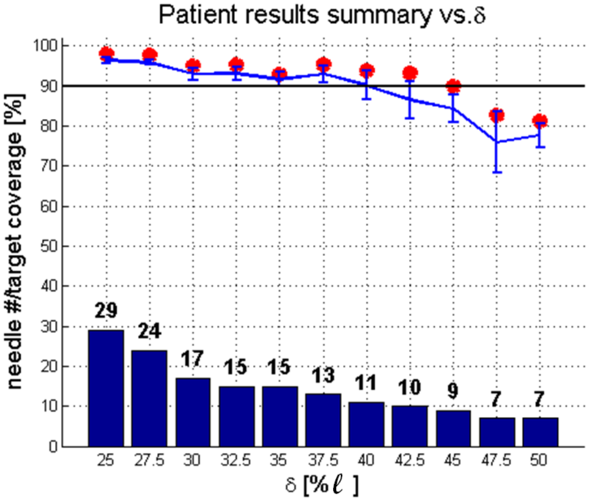

Figure 4.10: Representative summary plot for single patient . On the x-axis is δ as a percentage of prostate radius. The number of needles resulting from the given δ is shown as a bar graph. Each bar is labeled with the number of needles in the needle configuration. The target coverage for the nominal needle configuration is shown as a filled dot. There is a plot for the average target coverage of the perturbed needle configurations and ±1 standard deviation error bars. Needle configurations were perturbed by moving the inside tip of each needle in the x and y direction according to a normal distribution with 0 mean and 2.5 mm standard deviation. The plots looked very similar across all patients except that patients with larger prostates had slightly more perturbations meeting dose objectives, even for the same number of needles.

4.5 Discussion

Needle configurations could be computed for each patient in our data set that avoided the penile bulb and met dose objectives. Since the bulb was punctured at least once by the physician in every patient case, and bulb puncturing may be related to side effects, there is potential for automatically generated needle configurations to reduce side effects of HDR brachytherapy. Robot-assisted brachytherapy is a framework that can utilize a needle planning system like NPIP. However, some work is still required before HDR brachytherapy robots can become clinically relevant. Some of this work is discussed in this section.

Needle planning requires a definition of the entry zone where needles can be inserted into the body. The main challenge in defining the entry zone is that the patient is in supine position with legs closed during scanning, but legs open during needle insertion. Therefore, any entry zone definition made according to the CT scan would not be an accurate representation of the entry zone during the insertion procedure. To address this concern, the patient should either be scanned in a different position that could accommodate both the robot and the scanner, or the entry zone would need to be modified according to the change in position. Methodology to accomplish this task is beyond the scope of this paper.

NPIP requires a digitized anatomy set to compute a configuration of needles. However in the standard HDR brachytherapy workflow, the scan is taken after the needles have been inserted. This workflow inconsistency can be addressed by borrowing from the PPI brachytherapy workflow. In PPI brachytherapy, the patient is scanned and from the digitized anatomy set, a source configuration is computed and implanted by a physician. To incorporate NPIP into HDR brachytherapy, a similar approach can be taken. First, a scan is taken, then a needle configuration is computed, and finally is implanted by a physician or with the assistance of a robotic implant device. Another scan can then be taken to determine the needle locations for dose planning.

NPIP was designed to be used in conjunction with needle insertion robots. Care was taken in designing NPIP such that needle configurations computed from it could be executed by these robots. However, a study testing the implant feasibility of computationally generated skew-line needle configurations by robots should be done as a follow-up study. We suggest the following workflow for this future study: (1) scan a tissue phantom, (2) digitize the relevant anatomy, (3) compute a needle configuration with NPIP, (4) use a needle insertion robot to implant the needle configuration, (5) scan the phantom to verify the needle locations, (6) given the implanted needle configuration and patient anatomy, compute a dose plan for the patient, (7) deliver the dose plan in-silico, and (8) remove the needles. Our sensitivity analysis of dose distributions with respect to the number of needles suggests that choosing 16 needles (δ approximately 35%) will be able to meet dose objectives robustly (i.e. within the placement precision of robotic needle insertion systems).

As mentioned earlier, it would be very useful to mathematically represent Ω, the set of straight needles that can be inserted by a robot, enter through the entry zone, penetrate the prostate, and do not penetrate any OAR. This representation would be useful for several reasons. First, it could shed some light on why some cases, such as patients 6,7, and 15 in Table 4.2, are difficult to find needle configurations for, even when the same δ value can find a needle configuration for other other patients. For example, it could show that certain regions of the prostate are difficult, if not impossible, to reach without penetrating OAR. This information could be used to select candidate needles in a way that helps generate needle configurations for these patients for lower values of δ (i.e. less needles). Since a finite subsample of Ω must be always be generated prior to needle selection, a full representation of Ω could also determine regions of the prostate which can only be reached by a few needles. Thus, these needles can be included in the candidate needle set, which has advantages over relying on random sampling to find these needles.

4.6 Conclusion

We have presented Needle Planning by Integer Program (NPIP), a system for automatically generating skew line needle configurations that are patient-specific, collision-free, and avoid the penile bulb. We have shown that given the correct input parameters, NPIP can generate needle configurations which meet dose objectives and use as many or fewer needles than the current HDR brachytherapy workflow. We have also shown that, even given a crude candidate needle set, the final needle configuration is robust to minor placement perturbation. In the following chapter, we explore the use of a needle insertion robot to execute these implants, and make some comparisons to an expert brachytherapist.

Chapter 5

Towards Robot-Assisted HDR Brachytherapy

5.1 Overview

In the previous chapter, we developed NPIP, a computational system for generating patient specific HDR brachytherapy needle configurations for a patient. This chapter presents work from Garg et al., 2012 [41], which utilized NPIP and a needle insertion robot called Acubot-RND [108] to physically implant needle configurations computed from NPIP into an anatomically correct prostate gelatin model. We present experiments with a robotic HDR brachytherapy system consisting of planning software, Acubot-RND, and a human with no brachytherapy clinical experience. We show that needle configurations implanted using our robotic workflow were able to meet dose objectives without puncturing any healthy organs. Thus, the work presented in this study has the potential to successfully execute HDR brachytherapy with less side effects and without requiring a highly trained specialist.

5.2 Background

Recall that in PPI brachytherapy, a source configuration is computed from a treatment planning system and is implanted into the patient using one needle at a time, although one needle may contain several sources. A needle is loaded with the sources and inserted into the patient. The sources are left inside the patient as the needles are retracted. The needle is then reloaded with the next set of sources and the process is repeated. Currently, needle insertion robots designed for brachytherapy are specialized for PPI brachytherapy. Thus, brachytherapy robots are designed to have only one needle inserted at a time; not several simultaneously as is required for HDR brachytherapy.

To be suitable for HDR brachytherapy, a robot would need to “drop”, or let go of, an inserted needle so that the next needle could be inserted. Fortunately, the robot used for this study has the ability to drop needles, although this functionality was designed as a safety

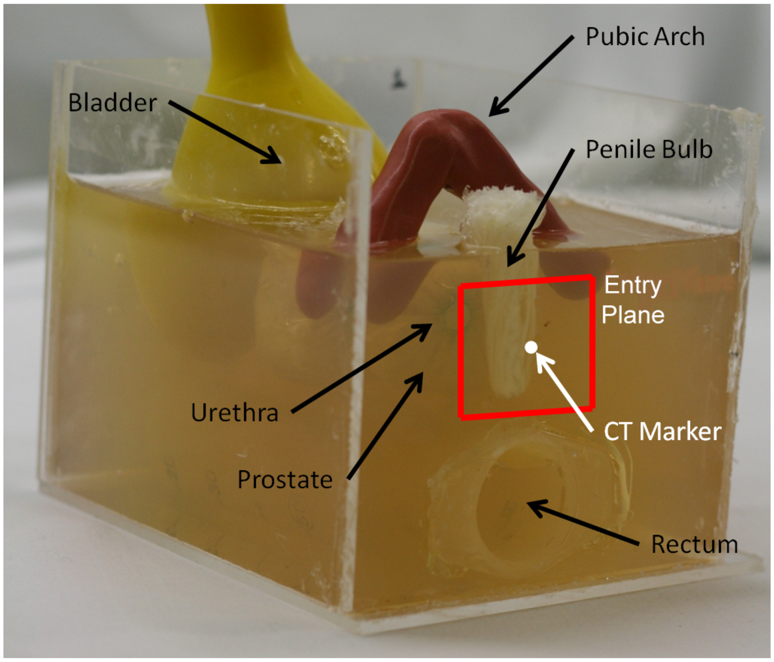

Figure 5.1: Photograph of a prostate phantom used in this study. The anatomy included was the prostate, urethra, bladder, penile bulb, pubic arch, rectum, and a CT marker for calibration with the robot and to represent the center of the perineum. The soft organs and connective tissues were gelatin-based. The pubic arch was constructed from clay putty. The Acubot-RND is registered to the CT-marker on the entry face of the phantom.

measure for both the patient and the robot. In this study, we modify the drop function of the Acubot-RND and develop custom software which allow it to perform brachytherapy autonomously.

5.3 Method and Materials

We will complete the following tasks within a brachytherapy environment: (1) compute a needle arrangement for the digitized anatomy of the patient that lies completely within the workspace of the robot, avoids non-prostate structures, uses needles economically, and meets dose objectives for the patient, (2) in a structured experimental setting, implant the needle arrangement in a tissue phantom using the robot, and (3) compute a dose plan from the implanted needle arrangement that meets treatment dose objectives. We will confirm the needle placement accuracy and the quality of dose plans using established clinical metrics.

5.3.1 Phantom

For this environment, the patient was modeled by a tissue phantom which included realistic representations (i.e. shapes and sizes) of the prostate, urethra, bladder, penile bulb, pubic arch, and rectum. We assumed that the tissue phantom was an acceptable model of the human prostate and nearby structures. Additionally the gelatin-based phantom was assumed to mimic human tissue stiffness properties. We also assumed that the anatomy remained unperturbed between scanning (anatomy digitization), and needle insertion. For this study we did not account for anatomy deformations during needle insertion, i.e. we assumed that needle insertions did not significantly alter the anatomy models.

A CT marker, a small radio-fluorescent pellet visible on CT-scan, was placed on the entry surface of the phantom. The CT marker was used to register the coordinate system of the planning algorithms with the robot space. We also defined an entry plane on the surface of the phantom for needles to enter. The entry plane was centered on the CT marker. We assumed that this entry plane defined the space where needles could be inserted into the perineum, which is the region where needles are inserted in an actual HDR brachytherapy procedure. The CT marker and entry plane location is shown on a phantom in Figure 5.1.

5.3.2 Planning

Needle arrangements were optimized using Needle Planning by Integer Program (NPIP) [103]. A candidate needle set was generated for the patient. This candidate needle set consisted of parallel lines and skew-lines. Skew-lines are non-parallel, non-intersecting lines. Although the needles intersected at four base points in the non-parallel set, the physical needles used for this study did not extend back to the point of intersection. The candidate needle set was chosen to (1) be within the workspace of the robot, (2) cover the space of needle arrangements possible with the degrees of freedom of the robot, and (3) contain mostly needles that did not intersect the penile bulb. A picture of the candidate needle set is shown in Figure 5.2. Note that this is Method 1 from the previous chapter.

Needles that intersected organs other than the prostate were removed from the set because naturally, we would like to avoid puncturing them to avoid side effects. However, organs such as the penile bulb may be punctured in practice due to placement inaccuracies of the robot. Therefore, we propose the organ volume displaced by needles as a “trauma metric” for needle arrangements. Specifically the trauma metric for organ s induced by a needle configuration is

$$T_s = \sum_k A_k L_k^s,$$

where A_k is the cross sectional area of needle k and L_k^s is the length of needle k puncturing organ s.

All dose plans were optimized to a given needle arrangement using Inverse Planning by Integer Program (IPIP). The IPIP model, RTOG dose objectives, and optimization parameters were the same as in [102]. A prescription dose of 950 cGy was applied to all planning cases.

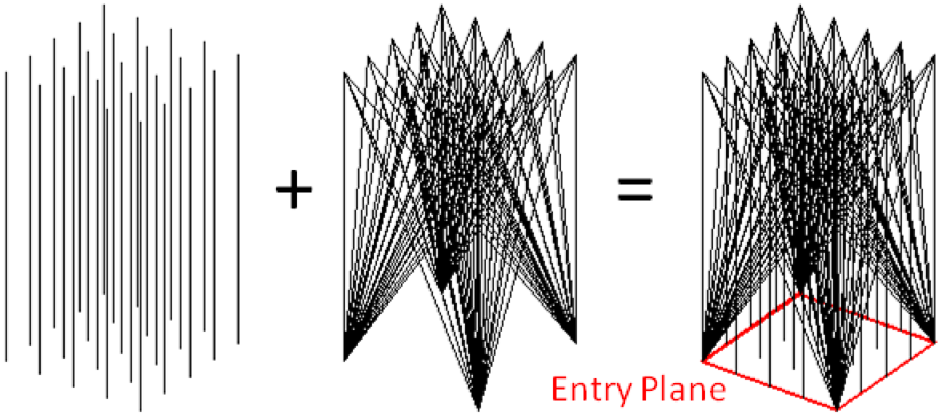

Figure 5.2: The candidate needle set was the set of needles that were available during needle planning. For this study, the candidate needle set consisted of parallel lines and skew-lines (non-parallel, non-intersecting) minus any needle that intersected a non-prostate organ. The entry plane, which represents the region of the perineum where needles can enter the phantom is also shown.

5.3.3 The Robot: Acubot-RND

This study used the Acubot-RND system developed at Johns Hopkins University [108] which is shown in Figure 5.3. Custom software was written for the Acubot-RND to automate needle insertion. The hardware and software components are described in detail in the following paragraphs.

Hardware Specifications Acubot-RND is a 7-DoF robot designed specifically for needle insertion. The 7 DoF are separated into three components of Acubot-RND: a 3-DoF Cartesian Positioning Stage (CPS), a 2-DoF Rotating Center of Motion (RCM), and a 2-DoF Rotating Needle Driver Module (RND). Figure 5.3 shows the degrees of freedom of the Acubot-RND. The RND is supported by the RCM, which sits on a passive positioning arm. The arm can be manually adjusted using tightening screws to position the RCM in a default configuration. The positioning arm is unactuated and rests on the CPS.

The CPS can move in the 200 mm along the x-axis, 50 mm along the y-axis, and 50 mm along the z-axis. The CPS has a spatial resolution of 0.01mm [32]. A trapezoid velocity profile was used for the CPS (Velocity: 10mm/s, Acceleration: 2 mm/s^2).

The RCM is a robotic orientation module that aligns the RND about two orthogonal axes while maintaining a fixed location for the tip of the outer nozzle of the needle holder (i.e. the center of motion). The inverse kinematics of the RCM joints to hold the needle tip in place while orienting it are built into the original Acubot-RND software. The default RCM joints have ball-worm motors which do not have rotation hard stop limits. Hence the RCM did not have a built-in homing sequence to return it to a fixed orientation in space. Therefore, the Acubot-RND was operated under manual joystick control. The system was modified to allow for a software homing procedure which is essential for autonomous operation of the robot.

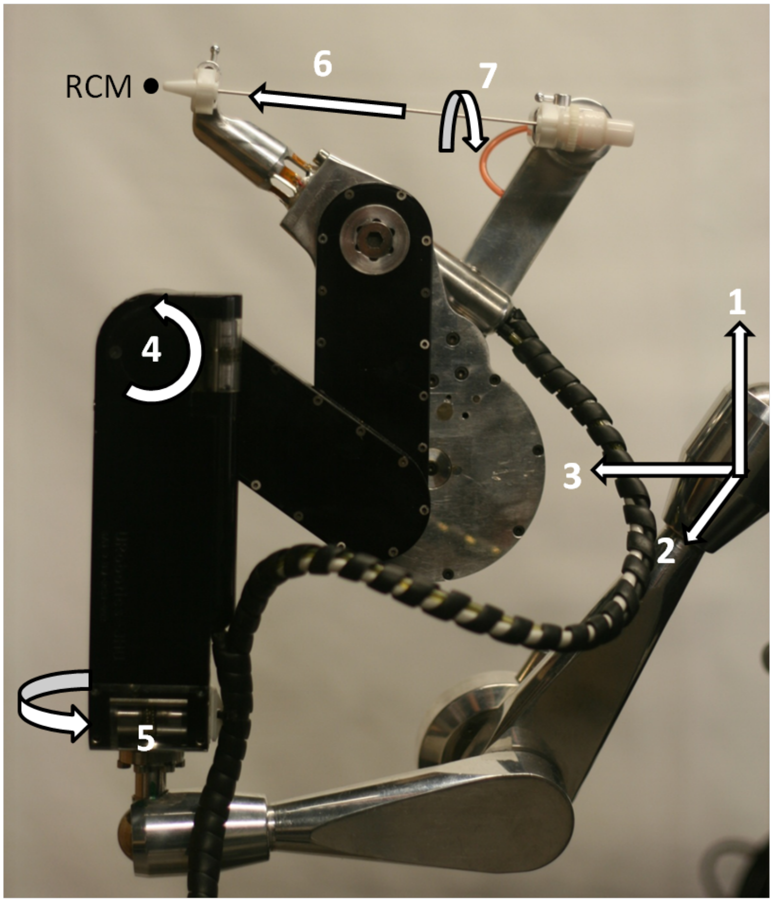

Figure 5.3: Photograph of Acubot-RND and degrees of freedom (DOFs). The Acubot-RND has 7 DOFs: a Cartesian Positioning Stage (CPS) which can move in the x-, y-, and z-direction, a Rotating Center of Motion (RCM) module which can rotate the needle around the needle tip, and a Rotating Needle Drive (RND) which can insert and rotate the needle.

The range of the RCM joints were set to allow 150°of rotation (-75°to 75°) in both axes with a resolution of 0.01°. A trapezoid velocity profile was also used for the RCM (Velocity: 8°/s, Acceleration: $2°/s^2$).

The RND can insert and rotate the needle. The RND performs insertion to a maximum depth of 120 mm using a trapezoid velocity profile with maximum velocity if 10 mm/s and acceleration of 2 mm/s^2. The needle insertion DoF has a spatial resolution of 0.01mm. Needle rotation was not used in this experiment. However, it has been shown that needle rotation improves targeting accuracy [10]. The RND has a needle release mechanism which enables the user to command it to release the needle at anytime during the procedure. The RND also is equipped with force sensors at the needle-guide. These sensors were not used in this study, but could be used in future work to improve placement precision.

In Figure 5.3, DoF 1-3 belong to the CPS, DoF 4-5 belong to the RCM, and DoF 6-7 belong to the RND.

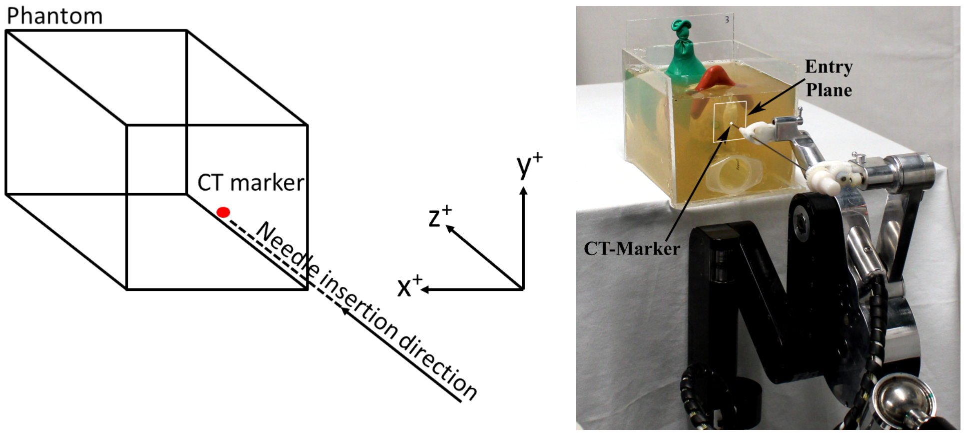

Figure 5.4: Diagram of assumed orientation of phantom, and needle insertion direction (left) and photograph of Acubot-RND in relationship to phantom and in *zero-position* (right).

Software Specifications Acubot-RND by design operates under manual joy-stick control to perform needle insertion. Specifically, each joint can be individually selected and controlled by an operator using a joy-stick interface. Autonomous operation was not previously possible because the unactuated positioning arm and the lack of stop limits on the RCM ball-worm motors did not allow a fixed reference space for the RCM. The original control system for Acubot-RND was augmented with custom software to allow automated needle insertion using an assumed hardware configuration, i.e. default orientation for each DoF and positioning arm orientation. The system takes as input a pair of cartesian coordinates: the insertion point of the needle on the entry plane and the target point inside the phantom. The input is passed from planner to the Acubot-RND as a formatted text file of coordinate pairs, one pair for each needle.

For inverse kinematic calculations, a specific orientation of the hardware was assumed. For this study, the RCM was oriented using the positioning arm such that 0 degrees in both RCM DoF resulted in the RND oriented along the z-positive axis. Furthermore, the rotation of one RCM joint resulted in a rotation of the RND end effector in only the x-z plane, Θ_{xz}, and the other RCM joint swept only in the y-z plane, Θ_{yz}. This orientation had the x-axis and z-axis of the robot space perpendicular to gravity. The Acubot-RND was brought into this orientation by carefully fixing the positioning arm using a level.

Let $P_0 = (x_0, y_0, z_0)$ and $P_f = (x_f, y_f, z_f)$ denote a cartesian coordinate input pair for our custom software. To implant the needle represented by this pair, a mapping of the 3D coordinate space to the 6D joint space of the Acubot-RND must be calculated. To simplify inverse kinematics, the insertion of the needle specified by these points was decoupled into three sequential tasks. (1) The tip of the needle was moved to P_0 using only the actuation in the CPS. (2) The RCM joints were actuated to orient the RND toward P_f. According to the physical orientation of the Acubot-RND used for this study, the angles of the two RCM joints were

$$\Theta_{xz} = \tan^{-1}\left(\frac{x_f - x_0}{z_f - z_0}\right), \text{ and } \Theta_{yz} = \tan^{-1}\left(\frac{y_f - y_0}{z_f - z_0}\right).$$

(3) The needle tip was inserted to P_f by actuating the needle insertion DoF of the RND a depth of $d = ||P_f - P_0||$.

After the needle was inserted to the target point, the robot was retracted back to the insertion point (P_0) using the RND, and the RCM was then re-orientated to *zero position*.

The software was also set to pause the needle insertion operation at critical points in the execution and wait for an operator command to continue. These stopping points were between needle insertions, before releasing the needle from the RND (discussed in Section 5.3.4), and orienting the RCM back to *zero position* after needle insertion.

Integration with Needle Planning The entry plane is defined as the planar segment where needles can be inserted. The size of the entry plane is limited by the maximum ranges of the x- and y-DoF of the CPS because the needle tip is brought to the needle insertion point on the entry plane using only these joints. The maximum ranges of these joints result in a rectangle 200 mm in the x-direction and 50 mm in the y-direction. In case of calibration or other setup errors, a conservative entry plane of 45 mm × 45 mm was used for this study. This entry plane is smaller than the standard template for PPI- and HDR brachytherapy, which is 100 mm × 100 mm. This entry plane is shown on the phantom in Figure 5.1 and relative to the candidate needle set in Figure 5.2.

5.3.4 Method Evaluation

5.3.4.1 Robotic Implant

The robot implant workflow consisted of pre-implant scanning and planning, robot setup and calibration, needle implanting, and post-implant scanning and dose planning. Each step is discussed in detail in the following paragraphs.

Pre-Implant Scanning and Planning A CT scan was performed on each phantom. The organs of the phantom and the CT marker were segmented using the Nucletron Oncentra® Dynamic Planning Environment. The 3D organ models were exported to NPIP. A candidate needle set was generated, and NPIP computed a needle arrangement for each phantom. A dose plan was also generated using IPIP to verify that the planned needle arrangement could produce a clinically acceptable dose plan. The needle insertion and target points for each needle arrangement were written to a formatted text file for our custom software.

Robot Setup and Calibration The phantom was moved to the robot work area and clamped to the table to ensure immobility during the procedure. The list of cartesian coordinate pairs specifying the needle arrangement was input to Acubot-RND.

The robot was calibrated to the position of the phantom by simultaneously (1) putting the robot in *zero position*, (2) placing the needle tip on the CT marker, and (3) placing the phantom such that the translational axis , x-DoF and y-DoF, of the CPS were parallel to the entry surface of the phantom. Parallelism of the X-Y plane of the robot space to the entry surface was determined by moving the needle tip in the x-axis and y-axis of CPS using the joy-stick controller. A visual inspection was made to determine if the needle tip was always the same distance from the phantom surface. We refer to this calibrated positioning of robot and phantom as the *initial state*. Figure 5.4 shows the Acubot-RND and phantom in *initial state*.

The Acubot-RND used an 18-gauge, 15 cm long, diamond-tip brachytherapy needle (COOK Biotech) for this procedure. The needle is a 2 mm diameter hollow sheath which houses a rigid stylet. Since Acubot-RND was not designed for multiple needle insertions in HDR brachytherapy, needle segments were cut from the stylets of regular brachytherapy needles. These needle segments were used to represent the hollow needles that would be inserted in a regular HDR brachytherapy procedure and connected to an afterloader for dose delivery.

Needle Implant Workflow To implant the input needle arrangement, the Acubot-RND was brought into *initial state*, and each needle was inserted sequentially according to the following procedure:

(1) Acubot-RND reached an insertion point in *zero position*.
(2) Acubot-RND aligned needle towards the target point.
(3) Acubot-RND inserted the needle tip to the target point.
(4) Acubot-RND halted for needle segment insertion.
(5) Operator removed the Acubot-RND stylet.
(6) Operator insertsed a needle segment to needle tip.
(7) Operator manually held the stylet.
(8) Acubot-RND released the needle.
(9) Operator pulled the needle over the stylet.
(10) Needle segment was left in place.
(11) Acubot-RND returned to *zero position*.
(12) Operator placed needle and stylet back into RND.

Note that the Steps 1, 8, and 11 were initiated on operator command.

Post-Implant-Scanning and Dose Planning After executing the implant, a second CT scan was obtained of the phantom. The metal needle segments could be identified but created significant artifacts on the CT scan. This made segmenting the anatomy rather difficult. However, since the phantom was static and provided many rigid control points, co-registration between the pre-implant and post-implant CT scans could be performed without introducing significant uncertainty. Co-registration is a standard feature in the Oncentra planning system. The needle positions and co-registered anatomy were used to compute a second dose plan using IPIP. Relevant dosimetric indices, trauma metrics, and error measurements were computed and are listed in Table 5.1.

5.3.4.2 Expert Implant

To compare our system with the performance of a physician on a similar task, an implant was performed on a tissue phantom by a Dr. I-Chow Hsu (the fifth author), a professor of radiation oncology at UCSF (expert). The expert is a certified radiation oncologist with a specialization in brachytherapy and over 18 years of clinical experience. The expert inserted 16 standard HDR brachytherapy needles into the phantom under trans-rectal ultrasound (TRUS) guidance using the UCSF-developed freehand technique [57]. A HAWK 2102 EXL TRUS system from B-K Medical was used for ultrasound imaging. Post-implant scanning and dose planning was performed using the same method as for the robot implants.

5.4 Results

There were three custom tissue phantoms constructed for this study: two for tests with the robot and one for the expert. There were two needle arrangements defined for each phantom: the pre-implant NPIP arrangement (planned) and the post-implant arrangement (actual). We label the needle arrangements computed by NPIP for the two phantoms as P1 and P2 to distinguish them from the actual implanted needle arrangements, A1 and A2. Since the expert implant was done manually, there was no pre-implant needle plan. We label the third phantom implanted by the expert as *Expert*.

The CT scans of the tissue phantoms were taken in 3 mm thick slices. The contoured prostate volumes for the three phantoms were 39 cm^3, 32 cm^3, and 37 cm^3. The total phantom volume was 750 cm^3. The regions of interest were contoured using Oncentra® version 4.1 by a medical physicist. Note that a clinician in radiotherapy, even one without brachytherapy experience, would have basic proficiency in organ contouring. Therefore, although the novice did not perform the contouring in this study, it can be assumed that a novice clinician using our robotic system would be able to do it. There was a 2 mm margin added to the contours of the penile bulb to ensure small placement errors did not result in puncturing of the bulb. Specifically, the contours of the bulb were made 2 mm larger in every direction. Adding margins is a common tool in radiation oncology dose planning and is a standard feature in Oncentra.

The number of needles in the NPIP candidate needle set after needles intersecting non-prostate organs were removed was 287 and 229 for P1 and P2, respectively. The NPIP δ parameter [103] was fine-tuned until NPIP returned a needle arrangement with 16 needles, the standard number used for HDR brachytherapy at the UCSF clinic. This corresponded to a δ value of 6.5 mm for the first phantom and 6.0 mm for the second phantom. All computational tasks were performed using Matlab R2011a on a Lenovo ThinkPad with an

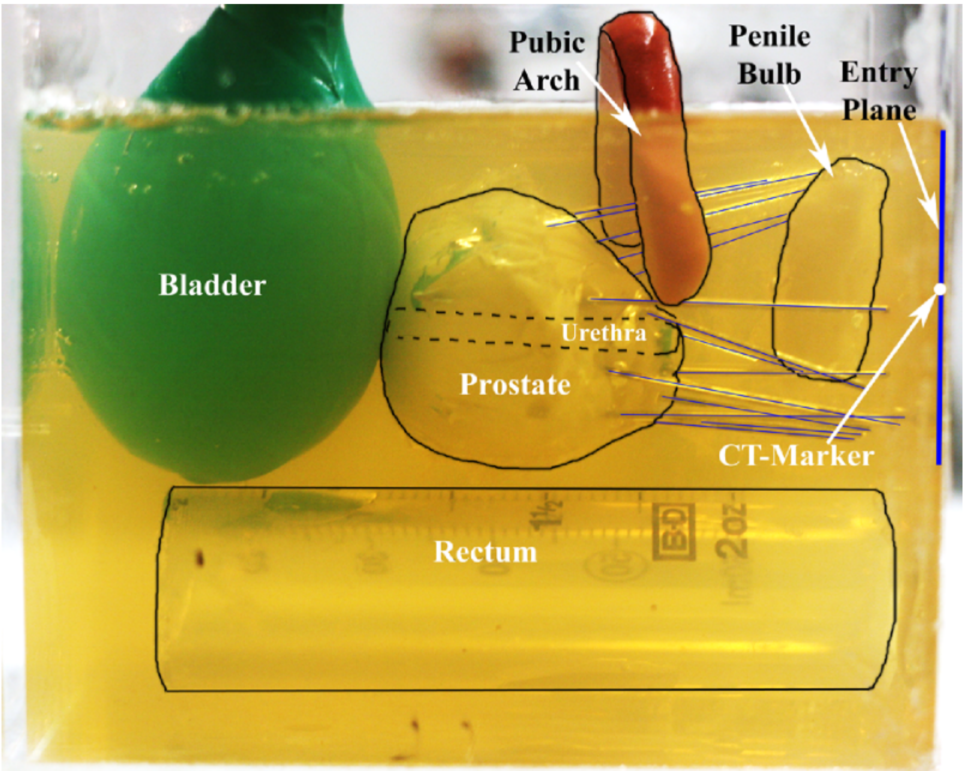

Figure 5.5: Photograph of tissue phantom with robot-implanted needle segments. The organ boundaries and needle patterns are superimposed on the figure for visual aid. No non-prostate organ was punctured in the insertion of this implant.

Intel i5-2410M processor and 4GB of RAM. The integer program optimization was done using the Matlab interface for the Mosek Optimization Toolbox v.6 [3].

The computed needle arrangements were implanted into each phantom in approximately 45 minutes plus setup time. A side view of an implanted phantom is shown in Figure 5.5. The expert implant took approximately 15 minutes to execute.

5.4.1 Clinical Analysis

Relevant dosimetric indices and trauma metrics are shown in Table 5.1.

There was a small difference between the dose plan from the planned arrangements and the dose plan from the actual arrangements for most of the indices. The one exception is the difference in V_{75}^{Bladder} values for P2 and A2 which were 0.3 cm^3 and 0.8 cm^3, respectively. The clinically acceptable limit for this criterion is 1 cm^3. This discrepancy is due to the needles not being inserted far enough into the prostate. As a result, the dwell times at the superior section of the prostate were increased by the dose planning system to cover more of the prostate with the prescription dose. The increased dwell times at the apex of the prostate resulted in slightly higher bladder dose because the bladder is above the prostate.

		Phantom 1		Phantom 2		Expert
Index	Req.	P1	A1	P2	A2	
$\mathrm{V}_{100}^{\mathrm{Prostate}}$	$\geq 90\%$	99.0	97.0	96.0	96.0	98.0
$\mathrm{V}_{150}^{\mathrm{Prostate}}$	$\leq 45\%$	39.0	40.0	40.0	37.0	37.0
$\mathrm{V}_{75}^{\mathrm{Bladder}}$	$\leq 1\ \mathrm{cm}^3$	0.00	0.00	0.30	0.80	0.30
$\mathrm{V}_{100}^{\mathrm{Bladder}}$	$= 0\ \mathrm{cm}^3$	0.00	0.00	0.00	0.00	0.00
$\mathrm{V}_{75}^{\mathrm{Bulb}}$	$\leq 1\ \mathrm{cm}^3$	0.00	0.00	0.00	0.00	0.00
$\mathrm{V}_{100}^{\mathrm{Bulb}}$	$= 0\ \mathrm{cm}^3$	0.00	0.00	0.00	0.00	0.00
$\mathrm{V}_{75}^{\mathrm{Rectum}}$	$\leq 1\ \mathrm{cm}^3$	0.06	0.00	0.00	0.00	0.00
$\mathrm{V}_{100}^{\mathrm{Rectum}}$	$= 0\ \mathrm{cm}^3$	0.00	0.00	0.00	0.00	0.00
$\mathrm{V}_{125}^{\mathrm{Urethra}}$	$\leq 1\ \mathrm{cm}^3$	0.06	0.05	0.04	0.06	0.07
$\mathrm{V}_{150}^{\mathrm{Urethra}}$	$= 0\ \mathrm{cm}^3$	0.00	0.00	0.00	0.00	0.00
$\mathrm{V}_{100}^{\mathrm{Body}}$	$= 0\ \mathrm{cm}^3$	0.00	0.00	0.00	0.00	0.00
$\mathrm{T}^{\mathrm{Bulb}}$	min	0.00	0.00	0.00	0.00	0.00

Table 5.1: Dose and trauma metrics for implants by robotic system (planned (P) and actual (A)) and for implant by expert

The trauma metric computed for this study was $\mathrm{T}^{\mathrm{bulb}}$ because it was the only organ that was at risk of being punctured in our setup. The trauma metric was computed using $A_k = \pi r^2$ mm^2 for all k using $r = 1$ mm (standard HDR brachytherapy needles have 1 mm radius) and L_k^{bulb} was computed by counting the number of contour slices the needle was within the bulb and multiplying it by the CT-Scan slice thickness, which was 3 mm. The bulb contours without the 2 mm margins were used to compute this metric. However, the trauma metric for the bulb was 0.00 mm^3 for every case because the bulb was not punctured. However, with a different experimental setup or physician, this may not be the case.

5.4.2 Placement Error Analysis

To measure the placement error, we sampled regularly spaced points on the planned and actual needles starting from the tip, every millimeter, for 6 centimeters along the length of the needles. This was approximately the length required to reach the entry face of the phantom. For a single needle, the corresponding points were the points on the planned and actual needle which were the same distance from their respective tips.

The total placement error was calculated as the RMS distance between all the planned and actual needle points. The RMS error was 2.6 mm for Phantom 1, and 4.3 mm for Phantom 2.

These placement errors were large with respect to other needle insertion experiments with robots [80], which was about 2 mm, and the internal accuracy of Acubot-RND, which is sub-millimeter. We separated the placement error into a systematic error and random error. We computed the the systematic error as the least-squares rigid translation and rotation from the implanted needle points, X, to the planned needle points, Y. Specifically, we computed the least-squares rotation matrix, R, and the translation vector, T, which minimized

$$\sum_i \|y_i - (Rx_i + T)\|,$$

where y_i was the i^{th} planned point and x_i was the i^{th} implanted point. For this purpose, we utilized the least-squares rigid transformation between point sets presented in Arun et al, 1987 [5].

We report the systematic error as the sextuple, $(\epsilon_x, \epsilon_y, \epsilon_z, \alpha, \beta, \gamma)$, where ϵ_x, ϵ_y, and ϵ_z are the least squares translations in the respective directions reported in millimeters, and α, β, and γ are the rotations in the Euler angles reported in degrees. The Euler angles were computed as

$$\alpha = \arcsin(R_{1,3}),\ \beta = \arccos\left(\frac{R_{1,1}}{\cos(\alpha)}\right),\ \text{and } \gamma = \arccos\left(\frac{R_{3,3}}{\cos(\alpha)}\right).$$

The systematic error was, for Phantom 1 and Phantom 2 respectively,

$$(1.2, 1.4, 0.81, 1.8^\circ, -.90^\circ, 1.6^\circ) \text{ and } (.90, 2.2, 2.4, 1.3^\circ, 0.83^\circ, 3.5^\circ).$$

We computed the random error as the total RMS distance between the planned points and actual points after the systematic error was removed, which was 1.4 mm for Phantom 1, and 2.4 mm for Phantom 2. These errors are similar to those reported in Muntener et al, 2008 [80], which performed multiple needle insertions into a living canine in a clinical environment.

A superposition of the planned (blue) and actual (red) needle arrangements for Phantom 1 and Phantom 2 are shown in Figure 5.6 and 5.7 (left), respectively. A superposition of the planned (blue) and actual needles after calibration error was removed (magenta) is also shown in the same figures on the right.

Table 5.2 gives numerical results for the needle placement errors. The errors are divided into “Total Error” and ”Random Error”, where the Total Error is the RMS error between the planned and actual points, and the Random Error is the RMS error between the planned points and actual points after correction for systematic error. The RMS distance, d, and the RMS distances in the x-, y-, and z-direction are given for each type of error. Also given are the minimum and maximum RMS distances over the points in each needle (i.e. individual needle errors) and the minimum and maximum RMS distances in the x-, y-, and z-directions.

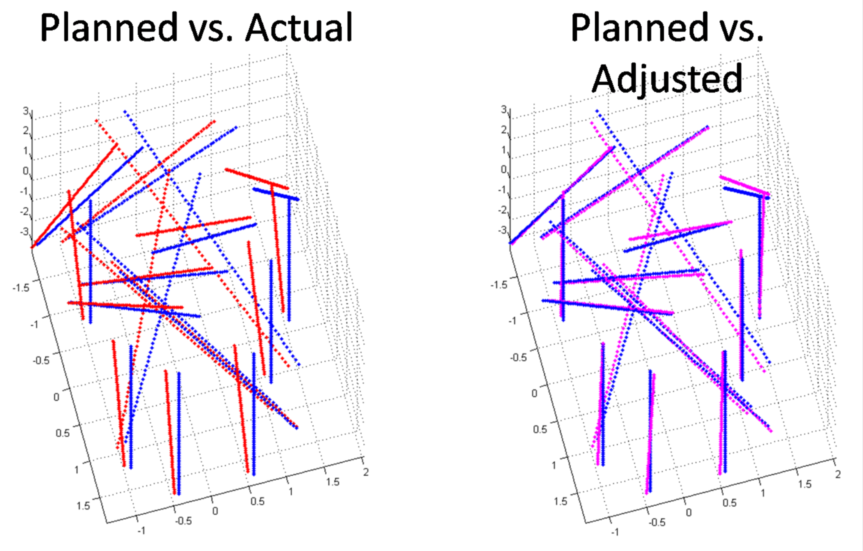
Planned vs. Actual

Planned vs. Adjusted

Figure 5.6: Superposition of planned (blue) and implanted (red) needle arrangement for Phantom 1. Although no sensitive structure was punctured in the implanted needle arrangement and all dose objectives were met, there was some placement error. The placement error was separated into systematic and random error. When the systematic error was removed from the implanted needle arrangement (magenta), the planned and adjusted points line up very well.

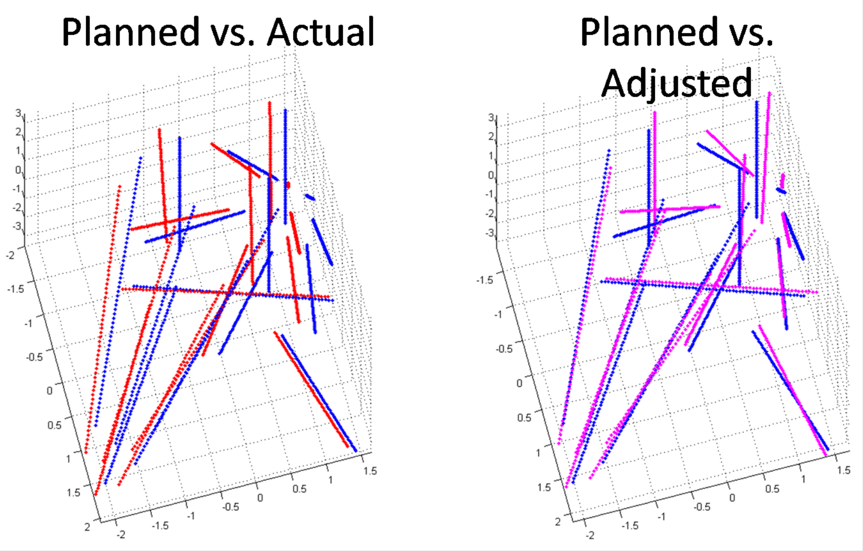
Planned vs. Actual

Planned vs. Adjusted

Figure 5.7: Superposition of planned (blue) and implanted (red) needle arrangement for Phantom 2, and planned and adjusted (magenta) needle arrangement.

	Phantom 1							
	Total				Random			
	d	x	y	z	d	x	y	z
All needle	2.6	1.4	1.6	1.5	1.4	0.5	0.5	1.2
min	1.3	0.5	0.2	0.1	0.2	0.0	0.1	0.1
max	4.1	2.5	2.5	3.0	2.5	0.8	1.1	2.3

	Phantom 2							
	Total				Random			
	d	x	y	z	d	x	y	z
All needle	4.3	1.9	2.3	3.1	2.4	1.2	1.1	1.8
min	2.0	0.8	0.7	0.9	0.8	0.1	0.1	0.0
max	6.3	2.7	3.2	5.3	5.2	2.3	1.8	5.1

Table 5.2: RMS errors

5.5 Discussion

These experiments have confirmed that computationally generated needle arrangements can be planned and executed using our robotic system with sufficient accuracy to achieve treatment objectives and avoid puncturing healthy organs in our experimental setup. However, our system had non-trivial placement errors which we have separated into systematic error and random error. It is likely that systematic error is mostly the result of calibration error between the robot and the phantom since the robot was manually calibrated and set into *zero-position* using only a leveling tool. The reason for this explanation is that the total RMS distance between the points was substantially reduced when the systematic error was removed. Assuming that Acubot-RND can execute the needles accurately according to its own internal reference frame (which previous studies have confirmed it can), a rigid translation and rotation of the robot relative to the phantom would have yielded much better results.

The random error is either the cause of operator error during needle retraction (i.e. the novice did not hold the needle exactly in place when retracting the needle over the stylet) or error in co-registration of the CT images. Since the anatomy was unperturbed during needle insertion and the rigid housing of the phantom provided many reliable reference points, it is unlikely that co-registration between the CT images contributed significantly to placement error. Operator error is most likely to manifest along the z-direction since the insertion depth for the stylet is made manually and the stopping is made visually. This was the case in our results since we had higher random error in the z-direction than in the other directions.

Both registration error and operator error may be reduced by hardware specialized for HDR brachytherapy. For example, registration error may be reduced by improving the calibration of the robot to the phantom with a custom fixture, which we will explore in future work. Additionally, a robot specialized for HDR brachytherapy needle insertion can reduce the role of the operator, which may reduce operator error. Specifically, a robot that can perform the needle insertion step, performed manually by the novice in this study, may

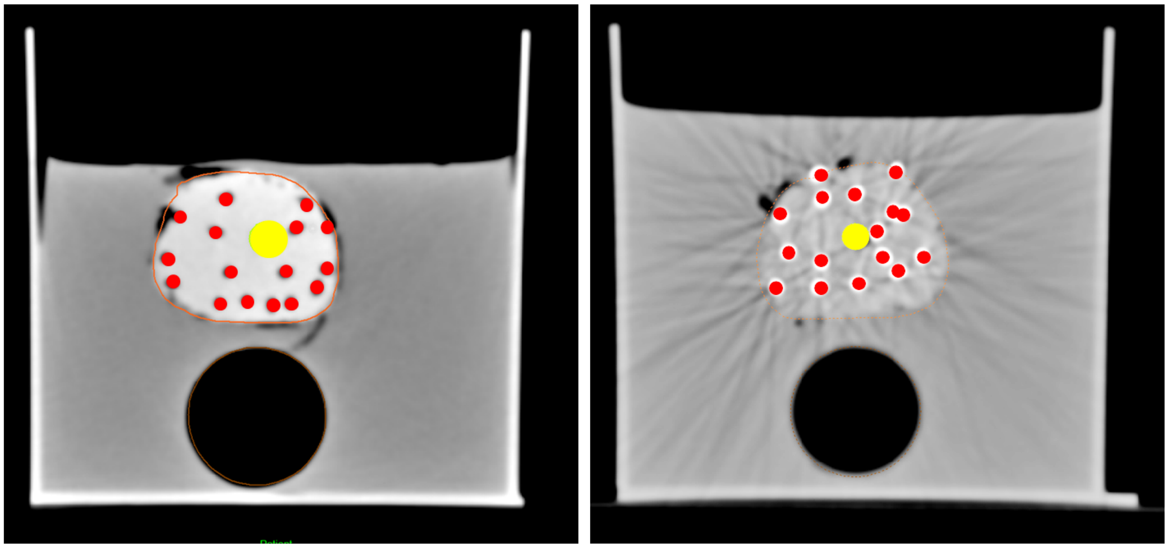

Figure 5.8: CT slice of physician implant (left) and robot implant (right). The expert implant has needles evenly distributed around the periphery and inside of the prostate, while the robot implants needles as skew lines. Our semi-autonomous system pre-computes optimal anatomy specific needle arrangements which might not be symmetrical but are able to meet treatment objectives and reduce trauma to sensitive organs at risk. Future improvements to needle insertion hardware and planning software will result in improved treatment quality, reproducibility and repeatability.

have superior performance to a human operator. According to other studies sub-millimeter placement precision is possible with robots [10, 91, 108], but not in a clinical setting.

In this study, we showed that, within our experimental environment, a novice without any brachytherapy experience could use our robotic system to execute an HDR brachytherapy needle arrangement that met treatment objectives. Figure 5.8 shows a cross-section of the needle arrangement implanted by the expert (left) and by the robot (right). Although there are some differences between the expert- and robot-executed implant, such as the lack of needles in the periphery of the prostate in the robot-executed implant, it is possible that robot-executed implants will be indistinguishable from human-executed implants in the near future. The ability to repeatedly, reliably, and safely execute needle arrangements with minimal involvement from an operator would be a significant step towards automating HDR brachytherapy.

5.6 Conclusion

We have utilized our previous work on IPIP and NPIP to demonstrate the first semi-automated robot implementation of skew-line needle arrangements for HDR brachytherapy. Results suggest that the resulting skew-line needle arrangements are comparable to those achieved by an expert human physician; both meet clinical objectives and avoid sensitive organ structures. Although there was non-trivial placement error, our results suggest that an improved calibration device and specialized robotic HDR brachytherapy hardware may be able to achieve much higher accuracy.

Chapter 6

Towards a Dose Planning Integer Program for Gamma Knife Perfexion

6.1 Overview

This chapter presents initial work in optimizing dose distributions for Gamma Knife, a radiotherapy for tumors in the head [69,84,93] including acoustic tumors [35,36,85], cerebral metastasis, [54, 56, 77], and sinus meningiomas [70, 94]. Although this work digresses from the previously described work on brachytherapy, it applies several of the same optimization principles, which shows that radiation oncology is a promising field for applied operations research.

Gamma Knife Perfexion (GKP) is the latest hardware model in the Gamma Knife series [69]. Studies have shown that it is capable of achieving higher dose conformity, i.e. delivering high dose to the target with low dose to the surrounding tissue, than its predecessors [69]. However, it also a complicated device with many degrees of freedom to control the dose distribution, which can make current manual treatment planning methods time consuming and difficult. Since tumors in the head are often immediately surrounded by healthy brain tissue, every cubic millimeter of healthy organ sparing can improve post-treatment quality of life.

This study develops a quadratic integer program for computing GKP dose plans. The dose planning philosophy is multi-objective, but it is readily extensible to dose-volume objectives. We provide computational studies on mock data since real data was not available. We leave the acquisition of real data and comparison to current treatment planning methods as immediate future work.

6.2 Background

Gamma Knife is a radiation therapy tool for tumors in the head [25], including but not limited to trigeminal neuralgia, acoustic neuromas, pituitary adenomas, arteriovenous malformations, and cerebral matastases, and has been shown to be an effective alternative to neurosurgery [25]. A Gamma Knife unit consists of three main components: a couch, a helmet, and a radiation unit. The patient is attached to the couch in supine position using a head frame which also immobilizes the head. The helmet sits over the head and collimates radiation from the radiation unit into narrow beams which intersect at the focal point of the helmet. The coach can move in the x-, y-, and z-direction with sub-millimeter accuracy to move the focal point inside the head.

The Gamma Knife Perfexion (GKP) is the latest hardware model in the Gamma Knife series. Unlike previous models, the GKP helmet is divided into 8 sectors of 24 beams (192 beams total). The beams in each sector can be set to one of four collimation diameters: 0mm (closed), 4mm, 8mm, and 16mm. A GKP treatment plan is delivered in ”shots”, and each shot consists of (1) a focal point, or ”shot center” inside the patient’s head, (2) a collimation size selection for each of the helmet sectors, and (3) a length of time, or ”shot time”, where the shot is actively delivering radiation to the patient. During planning, the shots are determined which will produce a dose distribution that will meet treatment objectives for the patient; generally, the eradication of the tumor with limited exposure to surrounding tissue.

Studies have shown that conformity of the volume receiving a lethal dose of radiation to the tumor volume is critical for eradicating the tumor and minimizing side effects [11, 17, 82, 88, 117]. Therefore, conformity of lethal dose to the target volume is the primary consideration when dose planning. However, other contributing factors are the total treatment time, which is related to patient discomfort, and the amount of time the technician can or is willing to spend creating the plan, which is related to treatment cost.

Current treatment planning systems are primarily manual. The user selects shot centers, collimation sizes for each of the sectors, and shot times. The planning system displays a heat map of the resulting dose distribution and relevant quantitative information such as the maximum dose to healthy tissue. The user makes changes to the plan according to their experience and intuition and repeats the process until a satisfactory dose distribution is achieved.

This planning workflow has three main limitations. First, the reliance on user input creates treatment disparity between experienced and non-experienced users. Second, the process consumes significant technician time, which is very expensive. Finally, and most importantly, there are little to no metrics on the quality of the final dose plan, nor are there guarantees that the final distribution is close the best distribution the hardware is capable of delivering to the patient. In other words, there may be a preferable dose distribution that is not found because of the searching limitations of the user.

Although optimization has been applied to Gamma Knife planning in the past, we could find very little work on optimization for GKP [86], with the exception of the default planning system, GammaPlan 10.0 [99].

Term	Description
J	Set of shot centers
K	Set of helmet sectors
L	Set of collimator sizes
T	Set of active shot times
G	Set of organs of interest
P	Set of dose control points
I_g	The set of points that belong to $g \in G$
d_i	Total dose to $p_i \in P$
$D_{ijk\ell}$	Dose rate parameters
$x_{jk\ell}$	Binary variable for collimators
$y_{jk\ell}$	Dummy variable for $t_j x_{jk\ell}$
T^{max}	Maximum allowed shot time

Table 6.1: Gamma Knife terms

6.3 Method and Materials

6.3.1 Model Formulation

In this section, a quadratic integer program is developed for computing GKP dose plans. A list of relevant terms is given in Table 6.1 for reference.

The model assumes that a set of shot centers have been specified by the user. Let J be the set of shot centers, K be the set of sectors, L be the set of collimator sizes, and T be the set of shot times. For brevity, we refer to a shot as j, a sector as k, and a collimator size as ℓ.

For this purpose, let G be the set of regions of interest, P be a set of uniformly spaced dose control points over the total volume of interest, and $I_g = \{i : p_i \in P \text{ belongs to } g \in G\}$. Note that p_i can only belong to one organ. To differentiate the tumor volume and healthy organs, let G^{tumor} be the tumor organ, and $G^{healthy}$ be the healthy organs.

First we develop constraints to model the dose delivered to a dose point as a function of J and T. The total dose to a dose point is the sum of the dose contribution of each shot, which is the aggregation of all the active beams. Let d_i be the total dose delivered to $p_i \in P$, $D_{ijk\ell}$be the dose delivered to p_i for every minute j is active given a configuration of collimators over the sectors in the shot, and let $x_{jk\ell}$ equal 1 if and only if collimator size ℓ is chosen for sector k during shot j, and equal 0 otherwise. Note that the individual beams in each sector can be aggregated together since the same collimator size must be chosen for each sector. Then the dose at a dose point can be computed as

$$d_i = \sum_j t_j \left(\sum_k \sum_\ell D_{ijk\ell} x_{jk\ell} \right),$$

which is a quadratic equality constraint that states that the total dose is the sum of the shot times, t_j, times the total dose rate over every sector and collimator size, $\sum_k \sum_\ell D_{ijk\ell} x_{jk\ell}$. To ensure that only one collimator size is chosen, the following constraint is also used:

$$\sum_\ell x_{jk\ell} = 1, \forall j \forall k.$$

The dose equation can be linearized using a variable with behavior $y_{jk\ell} = t_j x_{jk\ell}$. This behavior can be achieved with the following set of constraints.

$$\begin{aligned} &0 \le y_{jk\ell} \le T^{max} x_{jk\ell}, \\ &t_j - T^{max}(1 - x_{jk\ell}) \le y_{jk\ell} \le t_j, \end{aligned}$$

where T^{max} is the maximum allowed shot time, and now $d_i = \sum_j \sum_k \sum_\ell D_{ijk\ell} y_{jk\ell}$

Let $R_g^{\min}$and $R_g^{\max}$be the minimum and maximum dose desired for dose points in organ g. In general it is not possible to achieve dose within this range for every dose point because of the physical limitations of the delivery hardware. Therefore, we do not hard constraints for the doses. Instead, let c_i be the overdose or underdose for p_i. Assuming that c_i is ultimately minimized, then it can be represented using the following three inequalities:

$$\begin{aligned} &c_i \ge 0, \\ &c_i \ge R_g^{\min} - d_i, \\ &c_i \ge d_i - R_g^{\max}, \end{aligned}$$

Note that, there is no maximum dose for the tumor, nor minimum dose for healthy organs. Therefore, the constraints $c_i \ge R_g^{\min} - d_i$ are omitted for tumor dose points and the constraints $c_i \ge d_i - R_g^{\max}$ are omitted for the healthy organ dose points. However, they are retained in the formulation for completeness.

The primary clinical objective is to eradicate the tumor. That is, all of the tumor should receive more than the prescription dose. However, a small percentage of the tumor volume ($\tilde{1}$%) can receive less than the prescription dose if it will reduce high dose to large amounts of healthy tissue. To achieve this effect, our model uses a linear penalty function with high coefficient for tumor dose points, and a quadratic penalty with low coefficient for healthy organs. This objective can be expressed as:

$$\sum_{g \in G^{tumor}} \sum_{i \in I_g} M c_i + \sum_{g \in G^{healthy}} \sum_{i \in I_g} Q c_i^2,$$

and it is assumed that $Q << M$. Specifically, this objective function prioritizes the minimizing of underdose to the tumor (since M is large), except when causing large overdosing to healthy organs (since the quadratic penalty will eventually overwhelm the linear penalty).

Some preliminary values for M and Q are used for our computational study, but finding a suitable class solution, i.e. one that will in general give reasonable results, is beyond the scope of this study.

Our dose planning model for GKP is presented in full below.

$$
\begin{array}{llll}
 & \text{minimize} & \displaystyle\sum_{g\in G^{tumor}}\sum_{i\in I_g} Mc_i + \sum_{g\in G^{healthy}}\sum_{i\in I_g} Qc_i^2 & \\
 & \text{subject to:} & & \\
 & & \displaystyle d_i = \sum_j\sum_k\sum_\ell D_{ijk\ell}y_{jk\ell} & \forall i \\
 & & 0 \le y_{jk\ell} \le T^{max}x_{jk\ell} & \forall j,k,\ell \\
 & & t_j - T^{max}(1-x_{jk\ell}) \le y_{jk\ell} \le t_j, & \forall j,k,\ell \\
(GK) & & \displaystyle\sum_\ell x_{jk\ell} = 1, & \forall j,k \\
 & & c_i \ge R_g^{\min} - d_i, & \forall i \\
 & & c_i \ge d_i - R_g^{\max}, & \forall i \\
 & \text{and} & x_{jk\ell} \in \{0,1\}, & \forall j,k,\ell, \\
 & & t_j, c_i, y_{jk\ell} \ge 0 & \forall i,j,k,\ell
\end{array}
$$

Real patient, hardware, and dose rate data were not available for these initial experiment. Therefore, mock data was used in its place. They are described in the following three sections.

6.3.2 Patient Data

The total volume of interest was the upper half of the head, which was modeled as a half sphere extruded in the vertical direction. The diameter at the base was 7 cm and the height was 10 cm. These dimensions are approximately the average head size.

There were three tumor shapes in consideration for these experiments: a U-shape, a Star, and a Dumbbell. These a-spherical shapes make them complicated and difficult to manually plan. The three tumor shapes are shown in Figure 6.1.

There were no other healthy organs created for this experiment, but usually organs such as the ocular nerves are also contoured.

6.3.3 Helmet Model

The helmet inner surface was modeled as a 20 cm diameter sphere surface cut 2 and 17 cm above the sphere equator. The focal point was taken to be the center of the sphere. The helmet major axis was assumed to point along the z-axis, from the head of the patient

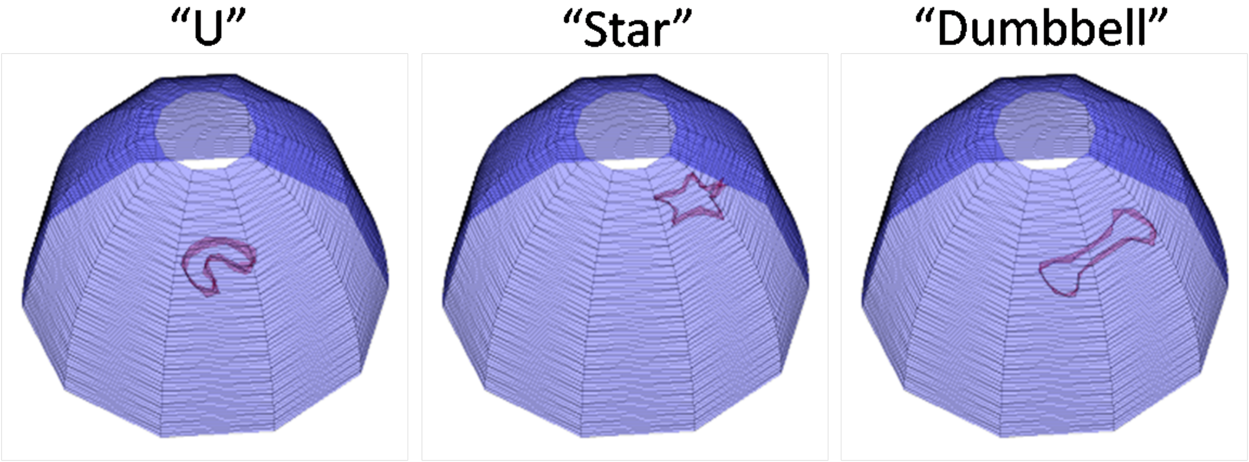

Figure 6.1: Mock tumor shapes used for this study. The tumor shapes were generated to be more complicated than a typical 'spherical' tumor. Only the tumor and the head volume was considered in this study.

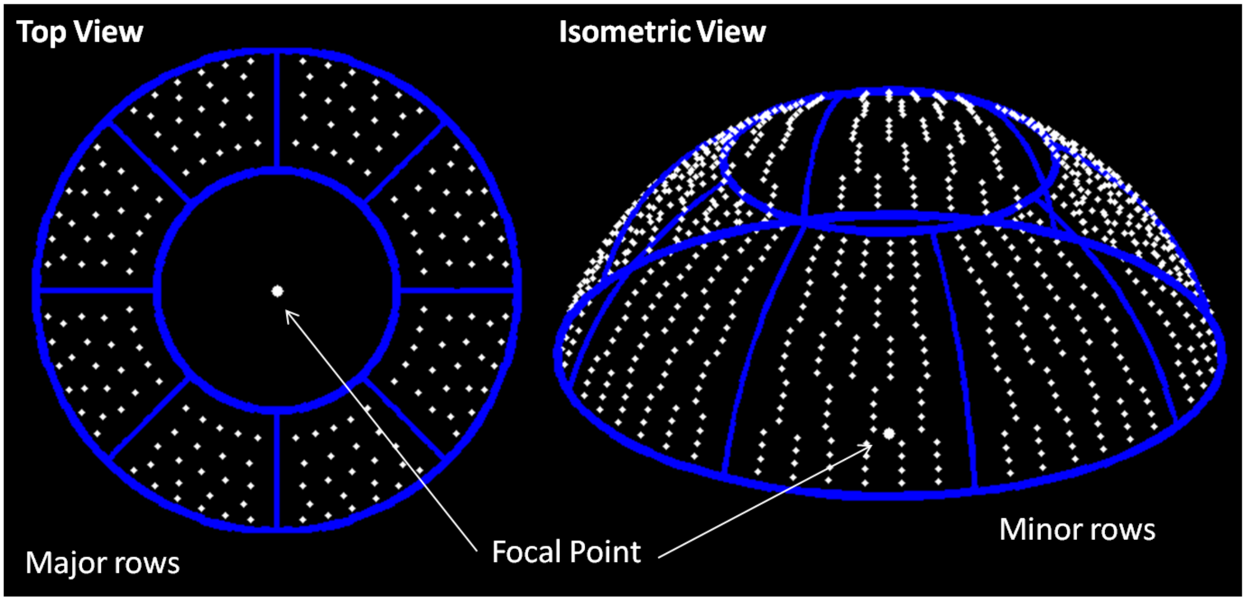

Figure 6.2: Approximation of helmet used in this study. Accurate measurements of the beam locations may be acquired manually using the GKP unit at UCSF.

towards the feet. The helmet was divided into eight equal sectors, the first sector starting along the x-axis (away from the face). Each sector contained 24 beams grouped into 5 major rows, each row going in the radial direction around the helmet. The number of beams in each row, starting from the top, was 5,4,5,4,and 6. Each major row contained 4 minor rows, one for each collimator size. The minor rows were equally spaced along the surface of the sphere with the same space between each row and the terminal ends of the helmet. The beams within each row were also equally spaced with the same space between each beam and the sides of the sector. Figure 6.2 shows the approximated helmet and focal point along with a snapshot taken from a GKP video.

6.3.4 Beam Model

The parameter $D_{ijk\ell}$is the sum of the dose rate contributions from every beam of size ℓ in sector k given shot center j. To compute this parameter, it is sufficient to have a beam model, which computes the intensity of a beam at an arbitrary point P_3, from a beam originating at P_1 pointed at P_2 with radius R given the properties of the material between P_1 and P_2. Such a beam model has been difficult to find in the literature and beam models for other radiotherapy devices involve many parameters which have to be determined experimentally on a unit-by-unit basis. For this study, an extremely approximate model was used. Our model used only geometry (no material properties) and was created to have two basic beam features: attenuation and scattering. Attenuation is the loss of intensity along the beam due to photon absorption into the material medium. Scattering is the tendency of the photons to deviate from the beam line due to photon collisions with the material medium.

Let $r_{\|}$ be length of the projection of $P_3 - P_1$ onto $P_2 - P_1$, and let $r_{\perp}$ be the length of $P_3 - P_1$ minus the projection of $P_3 - P_1$ onto $P_2 - P_1$. A depiction of $r_{\|}$ and $r_{\perp}$ are shown in Figure 6.3 (left).

The dose rate contribution, D, of a beam with radius, r, originating at P_1, pointed at P_2, at an arbitrary point was

$$D = D_0 e^{-\mu r_{\|}}$$

if $r_{\perp}$ was less than r and

$$D = D_0 e^{-\mu r_{\|}} e^{-\gamma \frac{r_{\perp} - r}{r_{\|}}},$$

where D_0 is some reference dose rate, μ is the attenuation coefficient, $e^{-\mu r_{\|}}$ is the attenuation term, γ is the scattering coefficient, and $e^{-\gamma \frac{r_{\perp} - r}{r_{\|}}}$ is the scattering term. When P_3 was inside the radius of the beam, it was assumed that the beam only attenuated. The scattering term says that scattering is exponential in the off-axis direction of the beam and more pronounced far from the beam. A slice of the beam intensity is shown in Figure 6.3 (right). The model makes intuitive sense, but is not based on any fundamental physics. The parameter, $D_{ijk\ell}$, was computed as the sum of the 24 beams in k with radius ℓ assuming shot center s.

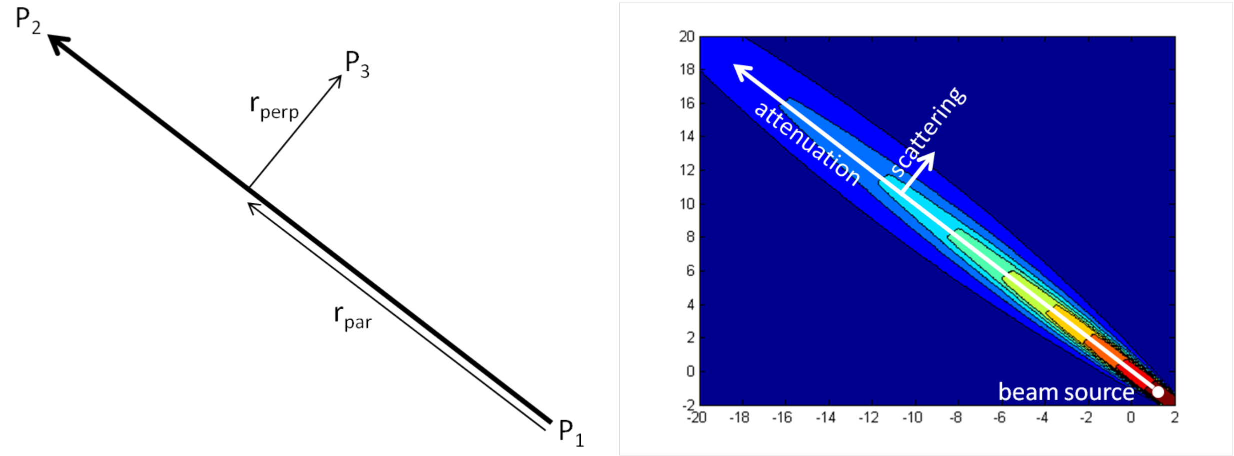

Figure 6.3: Depiction of $r_{\|}$ and $r_{\perp}$ (left). Profile of beam intensity using artificial model (right). The beam intensity makes intuitive sense, but is not based on fundamental physics. Diagram of terms used for beam model. The beam model computes the dose rate at an arbitrary point P_3, from a beam with radius r, originating at P_1 and pointed at P_2.

For this experiments $\mu = 0.0063$ was used, the attenuation coefficient of water, and $\gamma = 6.3$ was used. Although this model is extremely ad-hoc, the correct model, when found, can easily be substituted into the code base written for this experiment.

6.3.5 Method Evaluation

A uniform grid of points was generated for each organ. For the tumor volume, dose points were spaced, 2 mm in the x-, y-, and z-direction. For the head volume, dose points were spaced 8 mm in the x and y direction, and 2 mm in the z direction.

For the U-shape and Star tumor, ten shot centers were chosen randomly from the dose points in the target. For the Dumbbell shape, 15 shot centers were chosen randomly in the same manner. More shot centers were chosen for the Dumbbell because its volume is more spread out than the other shapes.

The prescription dose was set at 19 Gy, which was also R_min for the tumor. R_max for the head volume was set at 9 Gy. These are common dose limits used in practice, but since the beam model is not exact, only serve as numerical placeholders.

For this study, $M = 100$ was used for the linear coefficient and $Q = 1$ was used for the quadratic coefficient. Thus the objective function would behave as follows. Since $M >> Q$, delivering prescription dose to the tumor was prioritized over small overdosing of healthy tissue. However, delivering the prescription dose to healthy tissue, presumably lethal, would have equal weight to small underdosing of target tissue.

		U	Star	Dumbbell
Tumor dose points	[#]	321	192	371
Head dose points	[#]	7552	7563	7547
Solve time	[sec]	110	85	7539

Table 6.2: This table gives the number of dose points in the tumor and the head for each of the tumors explored in this study. Also given is the time required for CPLEX to compute the optimal solution for that tumor.

Instances of GK were solved using Matlab R2011a on a Lenova ThinkPad with an Intel i5-2410M processor and 4GB of RAM. The Matlab interface for CPLEX 12.1 was used for the integer program optimization. All default parameters were used.

6.4 Results

A brief description of the numerical results is shown in Table 6.2. Each model instance was solved to optimality. However, the solvetime for the Dumbbell tumor shape was much higher than for the other shapes. It is still unknown how adding additional healthy structures will affect solution time.

Figure 6.4 shows heat maps for the dose distributions found using our model. The rows show dose heat maps for consecutive slices of the head in descending z-axis order. The first row shows heat maps for the U-shaped tumor, the second row for the Star, and the third row for the Dumbbell. Orange shows dose exceeding 125% of the prescription dose, green for an excess of 100% of the prescription dose, cyan for an excess of 50% of the prescription dose, and blue for less than 50% of the prescription dose. It is desired for the entire tumor to be at least green, and outside the tumor to be blue.

Since this entire study is based on estimated data, these dose plans are not clinically relevant. However, the dose distribution confirms that the optimization is working toward the correct goal.

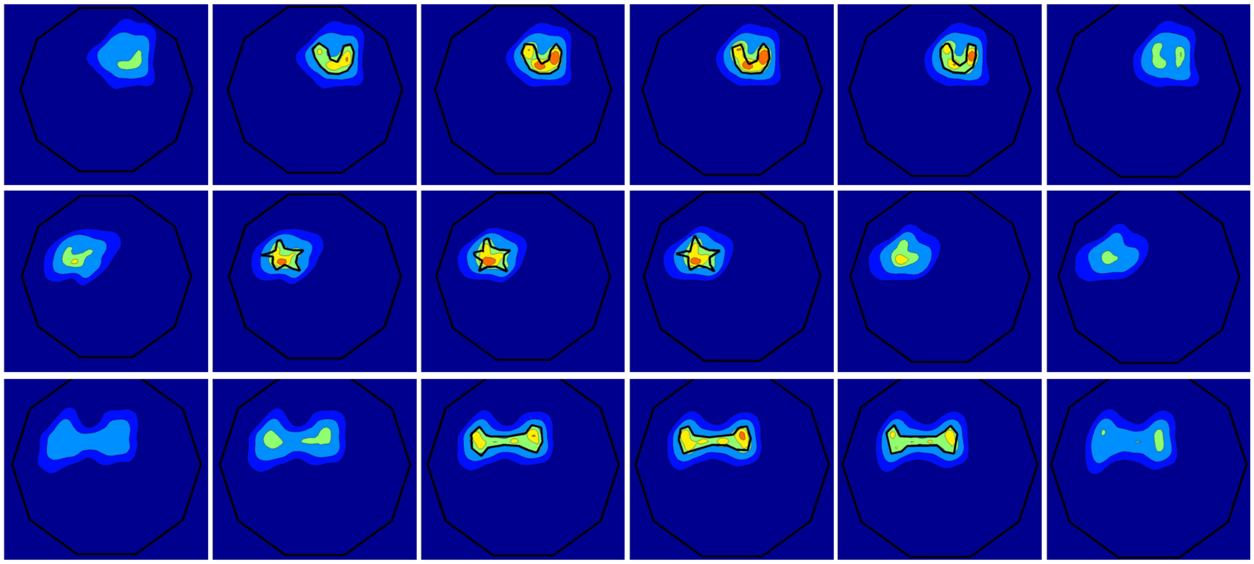

Figure 6.4: Heat maps of dose distributions over several slices around the tumor volume. Although this studied relied on artificial data, the dose distributions show that the optimization model is headed in the right direction.

6.5 Discussion

This study presented initial results for an optimization model which can compute GKP dose plans. However as a multi-objective model (i.e. dose objectives are aggregated into a single objective using penalties) and even with correct beam and helmet data, such a model would require manual fine tuning to achieve dose objectives. In particular, the user would have to tweak M for the tumor and a Q value for each healthy structure until a satisfactory dose plan was computed. Additionally, this model required that shot centers be selected by the user. This is an undesirable property since the user may choose shot centers which may exclude high quality dose plans from being found. Incorporating dose-volume constraints, such as those used in IPIP, and shot center optimization into the model is left to future work.

6.6 Conclusion

We have developed an integer program for creating Gamma Knife Perfexion dose plans. Our initial computational studies using mock data confirm that our optimization model can be used to control the dose distribution. Future work should (1) utilize more realistic helmet geometry, beam, patient, and dose objective data and (2) modify the model to include dose-volume constraints and shot center optimization, as well as custom algorithms for finding solutions to such models.

Chapter 7

Conclusion

7.1 Contributions

This dissertation has taken a major step in reducing side effects of brachytherapy, which is a major goal of the field. Specifically we have made the following contributions.

We developed Inverse Planning by Integer Program (IPIP), an integer programming which allows physicians to directly optimize the standard metrics used to evaluate HDR brachytherapy dose distributions. However, we showed that CPLEX could not solve this model within 30 minutes, which is already long enough to interrupt the clinical workflow.

To address the tractability shortcomings of IPIP, we developed a heuristic algorithm, IPIP-H, which uses two linear programs to compute feasible solutions for IPIP. Thus, it is a polynomial-time heuristic algorithm for IPIP. We showed that IPIP-H could produce a dose plan which met the RTOG-0321 dosimetric protocol in less than 30 seconds on a personal computer. We compared our results to IPSA, a clinically deployed dose planning model, which was not able to compute a dose plan meeting all objectives in a single iteration.

We formulated the problem of positioning HDR brachytherapy needles as a spatial coverage problem and showed that this problem could be represented as a set cover integer program.

We developed Needle Planning by Integer Program (NPIP), an algorithm which generates patient specific needle configurations represented by skew-lines. We used NPIP to compute needle configurations for 18 patients, and showed that these needle configurations met the RTOG-0321 dose objectives and used fewer needles than the physician. Additionally, NPIP always computed a needle configuration that avoided puncturing the penile bulb; the average number of punctures made by the physician was 5. We also conducted a sensitivity analysis of NPIP-generated needle configurations to placement errors on the order expected from current needle insertion robots in a clinical environment. We showed that, although dose objectives could be met with 10 or fewer needles, 16 needles were required to meet dose objectives robustly.

We designed and implemented the first end-to-end robotic HDR brachytherapy experiment. We planned and executed NPIP-generated needle configurations in a fully equipped brachytherapy environment on two anatomically-correct gelatin phantoms using Acubot-RND. Although there were non-trivial placement errors, we were able to implant the needle configurations with sufficient accuracy to meet dose objectives.

In summary we have developed planning tools for HDR brachytherapy needle and dose planning and established their viability in a robotic workflow. We leave the translation of these tools into clinical practice to future work, which is discussed in the following section.

7.2 Future Work

There are many exciting avenues of future work that warrant exploration. A few of them are discussed here.

7.2.1 Dose Planning

Our initial computational results for IPIP-H warrant testing them in a clinical setting in real time. Dose distributions could be computed from IPIP-H in parallel with the current dose planning system and the physician could pick the preferred dose plan. Compliance with dose objectives and total treatment planning time could be evaluated and compared.

Since IPIP-H allows the user to directly constraint dosimetric indices, alternative standards for constraining dosimetric indices could be explored that are more strict than the RTOG-0321 protocol. For example, it may in general be possible to restrict the $\mathrm{V}_{100}^{\mathrm{Prostate}} \geq 90\%$ *and* $\mathrm{V}_{75}^{\mathrm{Rectum}} \leq 0.5\ \mathrm{cm}^3$. If so, then the protocol can be updated and the rectal dose would be reduced for all HDR brachytherapy recipients. Also, additional constraints can be added, and should be explored, such as low dose to the rectum, $\mathrm{V}_{25}^{\mathrm{Rectum}}$. With a dose planning tool which directly constrain dosimetric indices, we can explore the set of possible dose distributions as they related to these metrics.

It is also worth exploring the use of IPIP-H to compute HDR brachytherapy dose plans for other cancer sites where brachytherapy is used such as the breast, cervix, head, and neck. Some of these regions like the head and neck have less developed dosimetric protocols, and IPIP-H could be a useful tool for establish effective protocols for dose distributions in these areas.

For IPIP, the underlying integer program, it would be useful to explore custom methods of reducing the upper bound, perhaps with cutting plane methods that work alone or assist in branch and bound.

There has been recent work in robust optimization, which augments linear programs with new constraints such that the new optimal solution is high quality and resistent to minor perturbations in the input parameters of the optimization model [12, 13]. In other words, robust optimization seeks to find a high quality solution that remains feasible even if the values contained in the A and b matrices are slightly perturbed, and such that the objective function value does not degrade under perturbations of the c matrix. There has been similar work on robust optimization of integer programs as well [7]. It would be interesting to apply these robust optimization techniques to IPIP to see if it can produce dose plans which are resistant to change given needle movement, contouring uncertainty, or edema of the prostate.

At the UCSF brachytherapy, it is preferred that the volume receiving very high dose (i.e. hot spots), be spread out in small "bubbles" rather than clumped up into a single, large volume. Theoretically, small spread out hot spots allow healthy tissue to recover, but is still lethal to cancerous tissue. However, the distribution of the volume receiving high dose is not represented by dosimetric indices since they only quantify the total volume receiving over a certain dose. Specifically, the $\mathrm{V}_{150}^{\mathrm{Prostate}}$ says nothing about how the volume receiving over 150% the prescription dose, only what the volume is. It would be interesting to represent this "second order" characteristic as integer programming constraints in some way and incorporate the constraints into IPIP. Thus the "bubbliness" of dose distributions could be improved.

7.2.2 Needle Planning

As mentioned in Chapter 4, it would be useful to improve the candidate needle set generation component of NPIP. Specifically, we would like a method which can characterize all the possible straight needles, i.e. a complete method, which originate at the entry zone, have some length within the prostate, and no length within an OAR. This method will be useful in two ways. First, the set of possible needles meeting the aforementioned criteria is infinite (finite, but very large if you take into account discretization of movement), and therefore, some down-sampling will be required. A complete method of characterizing the set of possible needles will allow sampling in a way that can ensure useful, but rare, needles are kept in the candidate needle set. In other words, the method would allow a sampling method which would result in an even distribution of candidate needles around the prostate volume. Second, this method would be able to tell the user that a specified coverage level is not achievable given the geometric constraints of the entry zone and organ models, and which areas cannot be covered.

From a practical point of view, this method would ensure that NPIP would never fail to return a needle configuration when one existed. Thus, a patient would never be rejected for treatment (i.e. due to pubic arch interference or similar disqualifying qualities) when a viable needle configuration existed for them.

The candidate needle set generation component, needle selection component, and dose planning component of NPIP are entirely modular (i.e. they can be replaced by other components which perform the same function). There some replacements worth exploring. For instance, it would be interesting to replace the current candidate needle set generation component with one that produced a family of curved needles to choose from. The same criteria for coverage would be used in the needle selection component, but it would be interesting to see if less needles could provide the same coverage. As curved needles have the potential to better avoid obstacles such as the penile bulb, needle configurations with less needles that meet treatment objectives may be possible with a curved candidate needle set. Also, the current needle selection component relies on the solution of a difficult integer program. It would be interesting to test a fast approach such as a greedy algorithm to see if the needle configurations of a similar quality to the integer program could be achieved, but with far less computational effort. Finally, it would be interesting to try replacing the dose planning component with a dose planning module for PPI brachytherapy, to see if needle and seed configurations can be generated which are superior to those computed from current systems.

Just as with IPIP-H, it would be interesting to do a computational study trying NPIP for other body sites where HDR brachytherapy is used.

7.2.3 Robot-Assisted Brachytherapy

A needle insertion robot, or robotic module, which is specialized for HDR brachytherapy would be extremely useful in studying the potential of NPIP for improving the implantation of needles. Such a device would require a method of dropping a needle in place without substantially disrupting its position, or a mechanism for holding many needles in place at once. The robot should also include a calibration device which registers the robot's frame to the anatomy. This calibration device would probably be scanned along with the patient anatomy. With an effective method of dropping or holding multiple needles and a calibration device, needle placement with substantially improved accuracy to our Acubot-RND experiment is probably possible

Towards the improved automation of the brachytherapy workflow, it would also be interesting to incorporate recent work in automatic segmentation of organs [50, 119] into the robotic workflow, since contouring uncertainty is known to be a factor in the dose to organs for various radiation therapies [37, 106].

7.3 Closing Remarks

Optimization has penetrated into many academic and industrial fields and has had a positive impact on the performance, efficiency, and cost of real-life systems. Radiation oncology is no exception. However, I believe that Radiation Oncology is still a ripe field for many applications of optimization. I believe that through optimization, we can make cancer therapies more effective at killing cancer, reduce side effects, reduce treatment costs, reduce skill dependance on treatment (and thus make treatments more accessible), reduce treatment time and patient discomfort, and even design new radiation therapy devices which improve the way patients are treated. My work up to this point has scratched the surface of the potential that I believe optimization has for improving all aspects of cancer therapy. It is my hope to see optimization used towards this purpose in the near and long term future.

Bibliography

[1] N. Abolhassani, R.V. Patel, and F. Ayazi. Minimization of needle deflection in robot-assisted percutaneous therapy. *The International Journal of Medical Robotics and Computer Assisted Surgery*, 3(2):140–148, 2007.

[2] R. Alterovitz, E. Lessard, J. Pouliot, I.C.J. Hsu, J.F. OBrien, and K. Goldberg. Optimization of HDR brachytherapy dose distributions using linear programming with penalty costs. *Medical physics*, 33:4012, 2006.

[3] E. Anderson. The MOSEK optimization toolbox for MATLAB version 2.5, users guide and reference manual. *World Wide Web, http://www. mosek. com (1999–2002)*, 1999.

[4] D.W. Arthur, D. Koo, R.D. Zwicker, S. Tong, H.D. Bear, B.J. Kaplan, B.D. Kavanagh, L.A. Warwicke, D. Holdford, C. Amir, et al. Partial breast brachytherapy after lumpectomy: low-dose-rate and high-dose-rate experience. *International Journal of Radiation Oncology* Biology* Physics*, 56(3):681, 2003.

[5] K.S. Arun, T.S. Huang, and S.D. Blostein. Least-squares fitting of two 3-D point sets. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, (5):698–700, 1987.

[6] American Medical Association. *Physician Characteristics and Distribution in the US, 2010.* American Medical Association, 2004.

[7] A. Atamtürk. Strong formulations of robust mixed 0–1 programming. *Mathematical Programming*, 108(2):235–250, 2006.

[8] C. Aygun, S. Weiner, A. Scariato, D. Spearman, and L. Stark. Treatment of non-small cell lung cancer with external beam radiotherapy and high dose rate brachytherapy. *International Journal of Radiation Oncology* Biology* Physics*, 23(1):127–132, 1992.

[9] F. Bachand, A.G. Martin, L. Beaulieu, F. Harel, and É. Vigneault. An eight-year experience of HDR brachytherapy boost for localized prostate cancer: biopsy and PSA outcome. *International Journal of Radiation Oncology* Biology* Physics*, 73(3):679–684, 2009.

[10] S. Badaan, D. Petrisor, C. Kim, P. Mozer, D. Mazilu, L. Gruionu, A. Patriciu, K. Cleary, and D. Stoianovici. Does needle rotation improve lesion targeting? *The International Journal of Medical Robotics and Computer Assisted Surgery*, 7(2):138–147, 2011.

[11] E.H. Balagamwala, J.H. Suh, G.H. Barnett, M.K. Khan, G. Neyman, R.S. Cai, M.A. Vogelbaum, E. Novak, and S.T. Chao. The importance of the conformality, heterogeneity, and gradient indices in evaluating Gamma Knife radiosurgery treatment plans for intracranial meningiomas. *International Journal of Radiation Oncology* Biology* Physics*, 2011.

[12] A. Ben-Tal and A. Nemirovski. Robust solutions of uncertain linear programs. *Operations research letters*, 25(1):1–13, 1999.

[13] A. Ben-Tal and A. Nemirovski. Robust solutions of linear programming problems contaminated with uncertain data. *Mathematical Programming*, 88(3):411–424, 2000.

[14] D. Bertsimas and J.N. Tsitsiklis. Introduction to linear optimization. 1997.

[15] K.H. Borgwardt. The average number of pivot steps required by the simplex-method is polynomial. *Mathematical Methods of Operations Research*, 26(1):157–177, 1982.

[16] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge University Press, 2004.

[17] L.S. Chin, L. Ma, and S. DiBiase. Radiation necrosis following gamma knife surgery: a case-controlled comparison of treatment parameters and long-term clinical follow up. *Journal of neurosurgery*, 94(6):899–904, 2001.

[18] D. Citrin, H. Ning, P. Guion, G. Li, R.C. Susil, R.W. Miller, E. Lessard, J. Pouliot, X. Huchen, J. Capala, et al. Inverse treatment planning based on MRI for HDR prostate brachytherapy. *International Journal of Radiation Oncology* Biology* Physics*, 61(4):1267–1275, 2005.

[19] K. Cleary, V. Watson, D. Lindisch, R. H. Taylor, G. Fichtinger, S. Xu, C. S. White, J. Donlon, M. Taylor, A. Patriciu, et al. Precision placement of instruments for minimally invasive procedures using a "needle driver" robot. *The International Journal of Medical Robotics and Computer Assisted Surgery*, 1(2):40–47, 2005.

[20] C. Cotrutz, M. Lahanas, C. Kappas, and D. Baltas. A multiobjective gradient-based dose optimization algorithm for external beam conformal radiotherapy. *Physics in Medicine and Biology*, 46(8):2161, 2001.

[21] I.B.M.I. CPLEX. 12.0 Optimization Studio, 2011.

[22] D.L. Craft, T.F. Halabi, H.A. Shih, and T.R. Bortfeld. Approximating convex pareto surfaces in multiobjective radiotherapy planning. *Medical physics*, 33:3399, 2006.

[23] J. A. Cunha, I. C. Hsu, and J. Pouliot. Dosimetric equivalence of nonstandard hdr brachytherapy catheter patterns. *Medical Physics*, 36:233, 2009.

[24] J.A. Cunha, I.C. Hsu, J. Pouliot, M. Roach, III, K. Shinohara, J. Kurhanewicz, G. Reed, and D. Stoianovici. Toward adaptive stereotactic robotic brachytherapy for prostate cancer: demonstration of an adaptive workflow incorporating inverse planning and an MR stealth robot. *Minimally Invasive Therapy & Allied Technologies*, 19(4):189–202, 2010.

[25] D.L. de Lunsford, J. Flickinger, G. Lindner, and A. Maitz. Stereotactic radiosurgery of the brain using the first United States 201 cobalt-60 source gamma knife. *Neurosurgery*, 24(2):151–159, 1989.

[26] I. Deutsch, M.J. Zelefsky, Z. Zhang, Q. Mo, M. Zaider, G. Cohen, O. Cahlon, and Y. Yamada. Comparison of PSA relapse-free survival in patients treated with ultra-high-dose IMRT versus combination HDR brachytherapy and IMRT. *Brachytherapy*, 9(4):313–318, 2010.

[27] L. Devroye. *Non-uniform random variate generation*, volume 4. Springer-Verlag New York, 1986.

[28] S. P. DiMaio, S. Pieper, K. Chinzei, N. Hata, S. J. Haker, D. F. Kacher, G. Fichtinger, C. M. Tempany, and R. Kikinis. Robot-assisted needle placement in open MRI: system architecture, integration and validation. *Computer Aided Surgery*, 12(1):15–24, 2007.

[29] L. Eapen, C. Kayser, Y. Deshaies, G. Perry, C. Morash, J.E. Cygler, D. Wilkins, S. Dahrouge, et al. Correlating the degree of needle trauma during prostate brachytherapy and the development of acute urinary toxicity. *International Journal of Radiation Oncology* Biology* Physics*, 59(5):1392–1394, 2004.

[30] F.M. Fang, Y.M. Wang, C.J. Wang, H.Y. Huang, and P.H. Chiang. Comparison of the outcome and morbidity for localized or locally advanced prostate cancer treated by high-dose-rate brachytherapy plus external beam radiotherapy (EBRT) versus EBRT alone. *Japanese journal of clinical oncology*, 38(7):474–479, 2008.

[31] M.C. Ferris, R.R. Meyer, and W. DSouza. Radiation treatment planning: Mixed integer programming formulations and approaches. *Handbook on modelling for discrete optimization*, pages 317–340, 2006.

[32] G. Fichtinger, E.C. Burdette, A. Tanacs, A. Patriciu, D. Mazilu, L.L. Whitcomb, and D. Stoianovici. Robotically assisted prostate brachytherapy with transrectal ultrasound guidance — phantom experiments. *Brachytherapy*, 5(1):14–26, 2006.

[33] G. Fichtinger, J.P. Fiene, C.W. Kennedy, G. Kronreif, I. Iordachita, D.Y. Song, E.C. Burdette, and P. Kazanzides. Robotic assistance for ultrasound-guided prostate brachytherapy. *Medical image analysis*, 12(5):535–545, 2008.

[34] B.M. Fisch, B. Pickett, V. Weinberg, and M. Roach. Dose of radiation received by the bulb of the penis correlates with risk of impotence after three-dimensional conformal radiotherapy for prostate cancer. *Urology*, 57(5):955–959, 2001.

[35] J.C. Flickinger, L. Dade Lunsford, M.E. Linskey, C.M. Duma, and D. Kondziolka. Gamma knife radiosurgery for acoustic tumors: multivariate analysis of four year results. *Radiotherapy and Oncology*, 27(2):91–98, 1993.

[36] R.L. Foote, R.J. Coffey, J.W. Swanson, S.G. Harner, C.W. Beatty, R.W. Kline, L.N. Stevens, T.C. Hu, et al. Stereotactic radiosurgery using the Gamma Knife for acoustic neuromas. *International Journal of Radiation Oncology* Biology* Physics*, 32(4):1153, 1995.

[37] F. Foppiano, C. Fiorino, G. Frezza, C. Greco, and R. Valdagni. The impact of contouring uncertainty on rectal 3d dose–volume data: Results of a dummy run in a multicenter trial (airopros01–02). *International Journal of Radiation Oncology* Biology* Physics*, 57(2):573–579, 2003.

[38] B.A. Fraass. The development of conformal radiation therapy. *Medical physics*, 22:1911, 1995.

[39] K.K. Fu and T.L. Phillips. High-dose-rate versus low-dose-rate intracavitary brachytherapy for carcinoma of the cervix. *International Journal of Radiation Oncology* Biology* Physics*, 19(3):791–796, 1990.

[40] L. Fu, H. Liu, W.S. Ng, D. Rubens, J. Strang, E. Messing, and Y. Yu. Hybrid dosimetry: Feasibility of mixing angulated and parallel needles in planning prostate brachytherapy. *Medical physics*, 33:1192, 2006.

[41] A. Garg, T. Siauw, D. Berenson, A. Cunha, I.C. Hsu, J. Pouliot, D. Stoianovici, and K. Goldberg. Initial experiments toward automated robotic implantation of skew-line needle arrangements for hdr brachytherapy. *Conference for Automation Science and Engineering*, 2012.

[42] P.D. Grimm, J.C. Blasko, J.E. Sylvester, R.M. Meier, and W. Cavanagh. 10-year biochemical (prostate-specific antigen) control of prostate cancer with ^{125}I brachytherapy. *International Journal of Radiation Oncology* Biology* Physics*, 51(1):31–40, 2001.

[43] J.T. Hepel, A.M.N. Syed, A. Puthawala, A. Sharma, and P. Frankel. Salvage high-dose-rate (HDR) brachytherapy for recurrent head-and-neck cancer. *International Journal of Radiation Oncology* Biology* Physics*, 62(5):1444–1450, 2005.

[44] B.S. Hilaris and N. Martini. Interstitial brachytherapy in cancer of the lung: a 20 year experience. *International Journal of Radiation Oncology* Biology* Physics*, 5(11):1951–1956, 1979.

[45] P.J. Hoskin, K. Motohashi, P. Bownes, L. Bryant, and P. Ostler. High dose rate brachytherapy in combination with external beam radiotherapy in the radical treatment of prostate cancer: initial results of a randomised phase three trial. *Radiotherapy and Oncology*, 84(2):114–120, 2007.

[46] I. Hsu, K. Bae, K. Shinohara, J. Pouliot, J. Purdy, G. Ibbott, J. Speight, E. Vigneault, R. Ivker, H. Sandler, et al. Phase II trial of combined high-dose-rate brachytherapy and external beam radiotherapy for adenocarcinoma of the prostate: preliminary results of RTOG 0321. *International Journal of Radiation Oncology* Biology* Physics*, 78(3):751–758, 2010.

[47] I. Hsu, J. Chow, A.R. Cabrera, V. Weinberg, J. Speight, A.R. Gottschalk, M. Roach, and K. Shinohara. Combined modality treatment with high-dose-rate brachytherapy boost for locally advanced prostate cancer. *Brachytherapy*, 4(3):202–206, 2005.

[48] I.C. Hsu, K. Shinohara, J. Pouliot, J. Purdy, J. Michalski, and G.S. Ibbott. RTOG 0321 protocol: Phase II trial of combined high dose rate brachytherapy and external beam radiotherapy for adenocarcinoma of the prostate. *Radiation Therapy Oncology Group, Philadelphia*, 2004.

[49] I.C.J. Hsu, E. Lessard, V. Weinberg, J. Pouliot, et al. Comparison of inverse planning simulated annealing and geometrical optimization for prostate high-dose-rate brachytherapy. *Brachytherapy*, 3(3):147–152, 2004.

[50] D.P. Huyskens, P. Maingon, L. Vanuytsel, V. Remouchamps, T. Roques, B. Dubray, B. Haas, P. Kunz, T. Coradi, R. Bühlman, et al. A qualitative and a quantitative analysis of an auto-segmentation module for prostate cancer. *Radiotherapy and oncology: journal of the European Society for Therapeutic Radiology and Oncology*, 90(3):337, 2009.

[51] S.K. Kang, R.H. Chou, R.K. Dodge, R.W. Clough, H.S.L. Kang, M.G. Bowen, B.A. Steffey, S.K. Das, S.M. Zhou, A.W. Whitehurst, et al. Acute urinary toxicity following transperineal prostate brachytherapy using a modified quimby loading method. *International Journal of Radiation Oncology* Biology* Physics*, 50(4):937–945, 2001.

[52] A. Karabis, P. Belotti, and D. Baltas. Optimization of catheter position and dwell time in prostate HDR brachytherapy using HIPO and linear programming. In *World Congress on Medical Physics and Biomedical Engineering, September 7-12, 2009, Munich, Germany*, pages 612–615. Springer, 2009.

[53] A. Karabis, S. Giannouli, and D. Baltas. 40 HIPO: A hybrid inverse treatment planning optimization algorithm in HDR brachytherapy. *Radiotherapy and Oncology*, 76:S29, 2005.

[54] B. Karlsson, C. Lindquist, and L. Steiner. Prediction of obliteration after Gamma Knife surgery for cerebral arteriovenous malformations. *Neurosurgery*, 40(3):425–431, 1997.

[55] N. Karmarkar. A new polynomial-time algorithm for linear programming. In *Proceedings of the sixteenth annual ACM symposium on Theory of computing*, pages 302–311. ACM, 1984.

[56] L. Kihlström, B. Karlsson, C. Lindquist, et al. Gamma knife surgery for cerebral metastases. implications for survival based on 16 years experience. *Stereotactic and functional neurosurgery*, 61:45, 1993.

[57] Y. Kim, I.C. Hsu, and J. Pouliot. Measurement of craniocaudal catheter displacement between fractions in computed tomography-based high dose rate brachytherapy of prostate cancer. *Journal of Applied Clinical Medical Physics*, 8(4):2415–2415, 2007.

[58] C. Kolotas, N. Tselis, M. Sommerlad, S. Röddiger, T. Schnabel, D. Baltas, A. Kalogera-Fountzila, G. Fountzilas, and N. Zamboglou. Reirradiation for recurrent neck metastases of head-and-neck tumors using CT-guided interstitial ^{192}Ir HDR brachytherapy. *Strahlentherapie und Onkologie*, 183(2):69–75, 2007.

[59] V. Lagerburg, M.A. Moerland, M.K. Konings, R.E. Van de Vosse, J.J.W. Lagendijk, and J.J. Battermann. Development of a tapping device: a new needle insertion method for prostate brachytherapy. *Physics in medicine and biology*, 51(4):891, 2006.

[60] V. Lagerburg, M.A. Moerland, M. van Vulpen, and J.J.W. Lagendijk. A new robotic needle insertion method to minimise attendant prostate motion. *Radiotherapy and oncology*, 80(1):73–77, 2006.

[61] M. Lahanas, D. Baltas, and N. Zamboglou. Anatomy-based three-dimensional dose optimization in brachytherapy using multiobjective genetic algorithms. *Medical Physics*, 26:1904, 1999.

[62] E.K. Lee, R.J. Gallagher, and M. Zaider. Planning implants of radionuclides for the treatment of prostate cancer: an application of mixed integer programming. *Optima March*, 61, 1999.

[63] E.K. Lee and M. Zaider. Mixed integer programming approaches to treatment planning for brachytherapy–application to permanent prostate implants. *Annals of Operations Research*, 119(1):147–163, 2003.

[64] E.K. Lee and M. Zaider. Operations research advances cancer therapeutics. *Interfaces*, 38(1):5–25, 2008.

[65] W. Lee, B.D.T. Daly, T.A. DiPetrillo, D.M. Morelli, A.C. Neuschatz, J. Morr, and M.J. Rivard. Limited resection for non–small cell lung cancer: observed local control with implantation of i-125 brachytherapy seeds. *The Annals of thoracic surgery*, 75(1):237–242, 2003.

[66] E. Lessard et al. Development and clinical introduction of an inverse planning dose optimization by simulated annealing (IPSA) for high dose rate brachytherapy. *Medical Physics*, 31(10), 2004.

[67] E. Lessard, S.L.S. Kwa, B. Pickett, III Mach Roach, and J. Pouliot. Class solution for inversely planned permanent prostate implants to mimic an experienced dosimetrist. *Medical physics*, 33:2773, 2006.

[68] E. Lessard and J. Pouliot. Inverse planning anatomy-based dose optimization for HDR-brachytherapy of the prostate using fast simulated annealing algorithm and dedicated objective function. *Medical physics*, 28:773, 2001.

[69] C. Lindquist and I. Paddick. The Leksell Gamma Knife Perfexion and comparisons with its predecessors. *Neurosurgery*, 61(3):130–141, 2007.

[70] R. Liščák, G. Šimonová, J. Vymazal, L. Janoušková, and V. Vladyka. Gamma Knife radiosurgery of meningiomas in the cavernous sinus region. *Acta neurochirurgica*, 141(5):473–480, 1999.

[71] C. MacLeod, A. Fowler, C. Dalrymple, K. Atkinson, P. Elliott, and J. Carter. High-dose-rate brachytherapy in the management of high-grade intraepithelial neoplasia of the vagina. *Gynecologic oncology*, 65(1):74–77, 1997.

[72] T. Martin, D. Baltas, R. Kurek, S. Röddiger, M. Kontova, G. Anagnostopoulos, T. Dannenberg, T. Buhleier, G. Skazikis, U. Tunn, et al. 3-D conformal HDR brachytherapy as monotherapy for localized prostate cancer. *Strahlentherapie und Onkologie*, 180(4):225–232, 2004.

[73] A.A. Martinez, L.L. Kestin, J.S. Stromberg, J.A. Gonzalez, M. Wallace, G.S. Gustafson, G.K. Edmundson, W. Spencer, and F.A. Vicini. Interim report of image-guided conformal high-dose-rate brachytherapy for patients with unfavorable prostate cancer: the William Beaumont phase II dose-escalating trial. *International Journal of Radiation Oncology* Biology* Physics*, 47(2):343–352, 2000.

[74] M. Meltsner. TH-B-224C-02: Performance evaluation and optimization of a novel brachytherapy robot. *Medical Physics*, 33:2264, 2006.

[75] G.S. Merrick, K. Wallner, W.M. Butler, R.W. Galbreath, J.H. Lief, and M.L. Benson. A comparison of radiation dose to the bulb of the penis in men with and without prostate brachytherapy-induced erectile dysfunction. *International Journal of Radiation Oncology* Biology* Physics*, 50(3):597–604, 2001.

[76] R.F. Mould and E. Towpik. *Radium history mosaic*. Maria Skłodowska-Curie Memorial Cancer Center and Institute of Oncology, 2007.

[77] A. Muacevic, F.W. Kreth, G.A. Horstmann, R. Schmid-Elsaesser, B. Wowra, H.J. Steiger, and H.J. Reulen. Surgery and radiotherapy compared with Gamma Knife radiosurgery in the treatment of solitary single cerebral metastases of small diameter. *Journal of neurosurgery*, 91(1):35–43, 1999.

[78] R.M. Munarriz, Q.R. Yan, A. Nehra, D. Udelson, and I. Goldstein. Blunt Trauma: the pathophysiology of hemodynamic injury leading to erectile dysfunction. *The Journal of urology*, 153(6):1831–1840, 1995.

[79] M. Muntener, A. Patriciu, D. Petrisor, D. Mazilu, H. Bagga, L. Kavoussi, K. Cleary, and D. Stoianovici. Magnetic resonance imaging compatible robotic system for fully automated brachytherapy seed placement. *Urology*, 68(6):1313–1317, 2006.

[80] M. Muntener, A. Patriciu, D. Petrisor, M. Schär, D. Ursu, D.Y. Song, and D. Stoianovici. Transperineal prostate intervention: Robot for fully automated MR imaging — system description and proof of principle in a canine model. *Radiology*, 247(2):543–549, 2008.

[81] S. Nag, B. Erickson, B. Thomadsen, C. Orton, J.D. Demanes, and D. Petereit. The american brachytherapy society recommendations for high-dose-rate brachytherapy for carcinoma of the cervix. *International Journal of Radiation Oncology* Biology* Physics*, 48(1):201–211, 2000.

[82] J.L. Nakamura, L.J. Verhey, V. Smith, P.L. Petti, K.R. Lamborn, D.A. Larson, W.M. Wara, M.W. McDermott, and P.K. Sneed. Dose conformity of Gamma Knife radiosurgery and risk factors for complications. *International Journal of Radiation Oncology* Biology* Physics*, 51(5):1313–1319, 2001.

[83] R. Nath, L.L. Anderson, G. Luxton, K.A. Weaver, J.F. Williamson, and A.S. Meigooni. Dosimetry of interstitial brachytherapy sources: recommendations of the AAPM Radiation Therapy Committee Task Group No. 43. *Energy*, 1(2):3, 1995.

[84] A. Niranjan, J. Novotny Jr, J. Bhatnagar, J.C. Flickinger, D. Kondziolka, and L.D. Lunsford. Efficiency and dose planning comparisons between the Perfexion and 4C Leksell Gamma Knife units. *Stereotactic and functional neurosurgery*, 87(3):191–198, 2009.

[85] G. Noren, D. Greitz, A. Hirsch, I. Lax, et al. Gamma Knife surgery in acoustic tumours. *Acta neurochirurgica. Supplementum*, 58:104, 1993.

[86] M.R. Oskoorouchi, H.R. Ghaffari, T. Terlaky, and D.M. Aleman. An interior point constraint generation algorithm for semi-infinite optimization with health-care application. *Operations research*, 59(5):1184–1197, 2011.

[87] F.D. Patel, S.C. Sharma, P.S. Negi, S. Ghoshal, and B.D. Gupta. Low dose rate vs. high dose rate brachytherapy in the treatment of carcinoma of the uterine cervix: a clinical trial. *International Journal of Radiation Oncology* Biology* Physics*, 28(2):335–341, 1994.

[88] P.L. Petti, S. Kunwar, and D.A. Larson. A theoretical investigation of optimal target-dose conformity in gamma knife radiosurgery. *Medical Physics*, 38:2812, 2011.

[89] T. Podder, L. Beaulieu, A. Dicker, M. Meltsner, M. Moerland, R. Nath, M. Rivard, D. Song, B. Thomadsen, and Y. Yu. AAPM-ESTRO guidelines for image guided robotic brachytherapy: Report from Task Group 192. *International Journal of Radiation Oncology* Biology* Physics*, 84(3):S129–S130, 2012.

[90] T.K. Podder, J. Sherman, D.P. Clark, E.M. Messing, D.J. Rubens, J.G. Strang, L. Liao, R.A. Brasacchio, Y. Zhang, W.S. Ng, et al. Evaluation of robotic needle insertion in conjunction with in vivo manual insertion in the operating room. In *Robot and Human Interactive Communication, 2005. ROMAN 2005. IEEE International Workshop on*, pages 66–72. IEEE, 2005.

[91] R. Pollock, P. Mozer, T.J. Guzzo, J. Marx, B. Matlaga, D. Petrisor, B. Vigaru, S. Badaan, D. Stoianovici, and M.E. Allaf. Prospects in percutaneous ablative targeting: comparison of a computer-assisted navigation system and the acubot robotic system. *Journal of Endourology*, 24(8):1269–1272, 2010.

[92] J. Pouliot, Y. Kim, E. Lessard, I.C. Hsu, D.B. Vigneron, J. Kurhanewicz, et al. Inverse planning for HDR prostate brachytherapy used to boost dominant intraprostatic lesions defined by magnetic resonance spectroscopy imaging. *International Journal of Radiation Oncology Biology Physics*, 59(4):1196–1207, 2004.

[93] J. Régis, M. Tamura, C. Guillot, S. Yomo, X. Muraciolle, M. Nagaje, Y. Arka, and D. Porcheron. Radiosurgery with the world's first fully robotized Leksell Gamma Knife PerfeXion in clinical use: a 200-patient prospective, randomized, controlled comparison with the Gamma Knife 4C. *Neurosurgery*, 64(2):346, 2009.

[94] P.H. Roche, J. Régis, H. Dufour, H.D. Fournier, C. Delsanti, W. Pellet, F. Grisoli, and J.C. Peragut. Gamma Knife radiosurgery in the management of cavernous sinus meningiomas. *Journal of neurosurgery*, 93:68–73, 2000.

[95] H.E. Romeijn, R.K. Ahuja, J.F. Dempsey, A. Kumar, and J.G. Li. A novel linear programming approach to fluence map optimization for intensity modulated radiation therapy treatment planning. *Physics in Medicine and Biology*, 48(21):3521, 2003.

[96] J.N. Roy, K.E. Wallner, S.T. Chiu-Tsao, L.L. Anderson, and C.C. Ling. CT-based optimized planning for transperineal prostate implant with customized template. *International Journal of Radiation Oncology* Biology* Physics*, 21(2):483–489, 1991.

[97] D.J. Scanderbeg, C. Yashar, R. Rice, and T. Pawlicki. Clinical implementation of a new HDR brachytherapy device for partial breast irradiation. *Radiotherapy and Oncology*, 90(1):36–42, 2009.

[98] F. Schiefke, G. Hildebrandt, S. Pohlmann, F. Heinicke, A. Hemprich, and B. Frerich. Combination of surgical resection and HDR-brachytherapy in patients with recurrent or advanced head and neck carcinomas. *Journal of Cranio-Maxillofacial Surgery*, 36(5):285–292, 2008.

[99] D.J. Schlesinger, F.T. Sayer, C.P. Yen, and J.P. Sheehan. Leksell GammaPlan version 10.0 preview: performance of the new inverse treatment planning algorithm applied to Gamma Knife surgery for pituitary adenoma. *Journal of neurosurgery*, 113(Special Supplement):144–148, 2010.

[100] C.M. Schneider, A.M. Okamura, and G. Fichtinger. A robotic system for transrectal needle insertion into the prostate with integrated ultrasound. In *Robotics and Automation, 2004. Proceedings. ICRA'04. 2004 IEEE International Conference on*, volume 1, pages 365–370. IEEE, 2004.

[101] S. Shah, A. Kapoor, J. Ding, P. Guion, D. Petrisor, J. Karanian, WF Pritchard, D. Stoianovici, BJ Wood, and K. Cleary. Robotically assisted needle driver: evaluation of safety release, force profiles, and needle spin in a swine abdominal model. *International Journal of Computer Assisted Radiology and Surgery*, 3(1):173–179, 2008.

[102] T. Siauw, A. Cunha, A. Atamtürk, I. Hsu, J. Pouliot, and K. Goldberg. IPIP: A new approach to inverse planning for HDR brachytherapy by directly optimizing dosimetric indices. *Medical Physics*, 38:4045, 2011.

[103] T. Siauw, A. Cunha, D. Berenson, A. Atamtürk, I.C. Hsu, K. Goldberg, and J. Pouliot. NPIP: A skew line needle configuration optimization system for HDR brachytherapy. *Medical physics*, 39(7):4339, 2012.

[104] R.S. Sloboda. Optimization of brachytherapy dose distributions by simulated annealing. *Medical physics*, 19:955, 1992.

[105] S. Smale. On the average number of steps of the simplex method of linear programming. *Mathematical Programming*, 27(3):241–262, 1983.

[106] W.Y. Song, B. Chiu, G.S. Bauman, M. Lock, G. Rodrigues, R. Ash, C. Lewis, A. Fenster, J.J. Battista, and J. Van Dyk. Prostate contouring uncertainty in megavoltage computed tomography images acquired with a helical tomotherapy unit during image-guided radiation therapy. *International Journal of Radiation Oncology* Biology* Physics*, 65(2):595–607, 2006.

[107] R.G. Stock, B. Mychalczak, J.G. Armstrong, J.P. Curtin, and L.B. Harrison. The importance of brachytherapy technique in the management of primary carcinoma of the vagina. *International Journal of Radiation Oncology* Biology* Physics*, 24(4):747–753, 1992.

[108] D. Stoianovici, K. Cleary, A. Patriciu, D. Mazilu, A. Stanimir, N. Craciunoiu, V. Watson, and L. Kavoussi. Acubot: a robot for radiological interventions. *Robotics and Automation, IEEE Transactions on*, 19(5):927–930, 2003.

[109] D. Stoianovici, D. Song, D. Petrisor, D. Ursu, D. Mazilu, M. Mutener, M. Schar, and A. Patriciu. "MRI Stealth" robot for prostate interventions. *Minimally Invasive Therapy & Allied Technologies*, 16(4):241–248, 2007.

[110] H. Su, M. Zervas, G.A. Cole, C. Furlong, and G.S. Fischer. Real-time MRI-guided needle placement robot with integrated fiber optic force sensing. In *Robotics and Automation (ICRA), 2011 IEEE International Conference on*, pages 1583–1588. IEEE, 2011.

[111] A.L. Trejos, A.W. Lin, M.P. Pytel, R.V. Patel, and R.A. Malthaner. Robot-assisted minimally invasive lung brachytherapy. *The International Journal of Medical Robotics and Computer Assisted Surgery*, 3(1):41–51, 2007.

[112] C. Vargas, M. Ghilezan, M. Hollander, G. Gustafson, H. Korman, J. Gonzalez, and A. Martinez. A new model using number of needles and androgen deprivation to predict chronic urinary toxicity for high or low dose rate prostate brachytherapy. *The Journal of urology*, 174(3):882–887, 2005.

[113] K. Wallner, J.C. Blasko, and M. Dattoli. *Prostate brachytherapy made complicated.* SmartMedicine Press, 2001.

[114] D.E. Wazer, L. Berle, R. Graham, M. Chung, J. Rothschild, T. Graves, B. Cady, K. Ulin, R. Ruthazer, T.A. DiPetrillo, et al. Preliminary results of a phase I/II study of HDR brachytherapy alone for T1/T2 breast cancer. *International Journal of Radiation Oncology* Biology* Physics*, 53(4):889, 2002.

[115] D.E. Wazer, D. Lowther, T. Boyle, K. Ulin, A. Neuschatz, R. Ruthazer, and T.A. DiPetrillo. Clinically evident fat necrosis in women treated with high-dose-rate brachytherapy alone for early-stage breast cancer. *International Journal of Radiation Oncology* Biology* Physics*, 50(1):107–111, 2001.

[116] L.A. Wolsey and G.L. Nehmauser. *Integer and Combinatorial Optimization.* Wiley, 1999.

[117] G. Zada, P.G. Pagnini, C. Yu, K.T. Erickson, J. Hirschbein, V. Zelman, and M.L.J. Apuzzo. Long-term outcomes and patterns of tumor progression after Gamma Knife radiosurgery for benign meningiomas. *Neurosurgery*, 67(2):322–329, 2010.

[118] M. Zaider, M.J. Zelefsky, G.N. Cohen, C.S. Chui, E.D. Yorke, L. Ben-Porat, and L. Happersett. Methodology for biologically-based treatment planning for combined low-dose-rate (permanent implant) and high-dose-rate (fractionated) treatment of prostate cancer. *International Journal of Radiation Oncology* Biology* Physics*, 61(3):702–713, 2005.

[119] Y. Zhan and D. Shen. Automated segmentation of 3d us prostate images using statistical texture-based matching method. *Medical Image Computing and Computer-Assisted Intervention-MICCAI 2003*, pages 688–696, 2003.

Appendices

Appendix A

Dose Planning Appendix

A.1 Additional Numerical Results

This appendix gives additional numerical results for each patient over the four algorithms included in the study. These were IPIP, IPIP-R, IPIP-H, and IPSA. For the sake of space, abbreviated dosimetric indices were given. Here “P” stands for Prostate, “B” for Bladder, “R” for Rectum, “U” for Urethra, ”b” for penile bulb, and “bo” for body tissue dose points. For IPIP “Gap” is the difference between the upper bound and the lower bound when CPLEX was terminated at 30 minutes. If the Gap is 0, it means IPIP was solved to optimality. For IPIP-H, the Gap was the difference between the true coverage as measured by the $V_{100}^{\text{Prostate}}$ and relaxed coverage used in the objective function of IPIP-H. A gap was not included for IPIP-R and IPSA because these algorithms are linear programs and therefore do not have a Gap. The term “HI” refers to the homogeneity index. If HI is blank, then it means it could not be computed for that particular case, e.g. $V_{100}^{\text{Prostate}}$ was 0. “Time” refers to the solve time.

Alg.	V^{P}_{100}	Gap	V^{P}_{150}	HI	V^{B}_{75}	V^{B}_{100}	V^{b}_{75}	V^{b}_{100}	V^{R}_{75}	V^{R}_{100}	V^{U}_{125}	V^{U}_{150}	V^{bo}_{200}	Time
IPIP	100	0	33	67	0.31	0.00	0.00	0.00	0.73	0.00	0.04	0.00	0.00	14
IPIP-R	100	n/a	32	68	0.42	0.00	0.00	0.00	0.95	0.00	0.32	0.00	0.00	6
IPIP-H	100	0	30	70	0.66	0.00	0.00	0.00	0.84	0.00	0.07	0.00	0.00	3
IPSA	98	n/a	25	74	0.20	0.00	0.00	0.00	0.38	0.00	0.00	0.00	0.14	3

Table A.1: IPIP Study Patient 1 Dosimetric Results

Alg.	V^{P}_{100}	Gap	V^{P}_{150}	HI	V^{B}_{75}	V^{B}_{100}	V^{b}_{75}	V^{b}_{100}	V^{R}_{75}	V^{R}_{100}	V^{U}_{125}	V^{U}_{150}	V^{bo}_{200}	Time
IPIP	100	0	34	66	0.84	0.00	0.00	0.00	0.31	0.00	0.10	0.00	0.00	1723
IPIP-R	100	n/a	50	50	1.69	0.00	0.00	0.00	0.77	0.00	0.36	0.00	0.00	9
IPIP-H	99	1	41	59	0.98	0.00	0.00	0.00	0.61	0.00	0.06	0.00	0.00	8
IPSA	98	n/a	28	71	0.32	0.00	0.00	0.00	0.28	0.00	0.07	0.00	0.38	2

Table A.2: IPIP Study Patient 2 Dosimetric Results

Alg.	V^{P}_{100}	Gap	V^{P}_{150}	HI	V^{B}_{75}	V^{B}_{100}	V^{b}_{75}	V^{b}_{100}	V^{R}_{75}	V^{R}_{100}	V^{U}_{125}	V^{U}_{150}	V^{bo}_{200}	Time
IPIP	88	10	25	71	0.54	0.00	0.00	0.00	0.35	0.00	0.08	0.00	0.00	1800
IPIP-R	96	n/a	44	54	1.76	0.00	0.00	0.00	1.25	0.00	0.29	0.00	0.00	12
IPIP-H	93	6	39	58	0.92	0.00	0.00	0.00	0.88	0.00	0.08	0.00	0.00	7
IPSA	96	n/a	31	68	1.62	0.00	0.00	0.00	0.74	0.02	0.20	0.04	0.43	4

Table A.3: IPIP Study Patient 3 Dosimetric Results

Alg.	V^{P}_{100}	Gap	V^{P}_{150}	HI	V^{B}_{75}	V^{B}_{100}	V^{b}_{75}	V^{b}_{100}	V^{R}_{75}	V^{R}_{100}	V^{U}_{125}	V^{U}_{150}	V^{bo}_{200}	Time
IPIP	94	4	29	69	0.56	0.00	0.00	0.00	0.49	0.00	0.07	0.00	0.00	1801
IPIP-R	95	n/a	44	54	0.67	0.00	0.00	0.00	1.30	0.00	0.46	0.00	0.00	22
IPIP-H	94	5	36	62	0.65	0.00	0.00	0.00	0.72	0.00	0.06	0.00	0.00	16
IPSA	98	n/a	32	67	1.19	0.36	0.00	0.00	0.00	0.00	0.10	0.01	0.05	3

Table A.4: IPIP Study Patient 4 Dosimetric Results

Alg.	V^{P}_{100}	Gap	V^{P}_{150}	HI	V^{B}_{75}	V^{B}_{100}	V^{b}_{75}	V^{b}_{100}	V^{R}_{75}	V^{R}_{100}	V^{U}_{125}	V^{U}_{150}	V^{bo}_{200}	Time
IPIP	96	3	38	61	0.29	0.00	0.00	0.00	0.61	0.00	0.02	0.00	0.00	1801
IPIP-R	97	n/a	49	49	1.36	0.00	0.00	0.00	0.77	0.00	0.42	0.00	0.00	14
IPIP-H	96	3	38	61	0.76	0.00	0.00	0.00	0.48	0.00	0.10	0.00	0.00	8
IPSA	97	n/a	34	65	0.10	0.00	0.00	0.00	0.77	0.05	0.13	0.00	0.24	3

Table A.5: IPIP Study Patient 5 Dosimetric Results

Alg.	V^{P}_{100}	Gap	V^{P}_{150}	HI	V^{B}_{75}	V^{B}_{100}	V^{b}_{75}	V^{b}_{100}	V^{R}_{75}	V^{R}_{100}	V^{U}_{125}	V^{U}_{150}	V^{bo}_{200}	Time
IPIP	99	0	31	69	0.74	0.00	0.00	0.00	0.49	0.00	0.06	0.00	0.00	1801
IPIP-R	100	n/a	51	48	1.27	0.00	0.00	0.00	2.00	0.00	0.35	0.00	0.00	28
IPIP-H	98	2	39	60	0.48	0.00	0.00	0.00	0.67	0.00	0.05	0.00	0.00	26
IPSA	98	n/a	26	74	0.34	0.00	0.00	0.00	0.70	0.00	0.01	0.00	0.14	9

Table A.6: IPIP Study Patient 6 Dosimetric Results

Alg.	V^{P}_{100}	Gap	V^{P}_{150}	HI	V^{B}_{75}	V^{B}_{100}	V^{b}_{75}	V^{b}_{100}	V^{R}_{75}	V^{R}_{100}	V^{U}_{125}	V^{U}_{150}	V^{bo}_{200}	Time
IPIP	0	100	0		0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1801
IPIP-R	99	n/a	49	51	3.04	0.00	0.00	0.00	2.40	0.00	0.37	0.00	0.00	39
IPIP-H	94	5	26	72	0.92	0.00	0.00	0.00	0.73	0.00	0.06	0.00	0.00	17
IPSA	97	n/a	23	76	2.00	0.02	0.00	0.00	0.65	0.00	0.00	0.00	0.10	8

Table A.7: IPIP Study Patient 7 Dosimetric Results

Alg.	V^{P}_{100}	Gap	V^{P}_{150}	HI	V^{B}_{75}	V^{B}_{100}	V^{b}_{75}	V^{b}_{100}	V^{R}_{75}	V^{R}_{100}	V^{U}_{125}	V^{U}_{150}	V^{bo}_{200}	Time
IPIP	0	100	0		0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1800
IPIP-R	100	n/a	48	52	2.08	0.00	0.00	0.00	1.16	0.00	0.30	0.00	0.00	12
IPIP-H	99	1	36	64	0.97	0.00	0.00	0.00	0.48	0.00	0.06	0.00	0.00	12
IPSA	98	n/a	28	72	0.76	0.00	0.00	0.00	0.00	0.00	0.10	0.00	0.29	4

Table A.8: IPIP Study Patient 8 Dosimetric Results

Alg.	V^{P}_{100}	Gap	V^{P}_{150}	HI	V^{B}_{75}	V^{B}_{100}	V^{b}_{75}	V^{b}_{100}	V^{R}_{75}	V^{R}_{100}	V^{U}_{125}	V^{U}_{150}	V^{bo}_{200}	Time
IPIP	61	38	1	99	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1836
IPIP-R	99	n/a	48	51	0.58	0.00	0.00	0.00	0.02	0.00	0.48	0.00	0.00	10
IPIP-H	98	2	43	57	0.36	0.00	0.00	0.00	0.01	0.00	0.10	0.00	0.00	8
IPSA	98	n/a	40	60	0.23	0.00	0.00	0.00	0.00	0.00	0.06	0.00	0.29	3

Table A.9: IPIP Study Patient 9 Dosimetric Results

Alg.	V^{P}_{100}	Gap	V^{P}_{150}	HI	V^{B}_{75}	V^{B}_{100}	V^{b}_{75}	V^{b}_{100}	V^{R}_{75}	V^{R}_{100}	V^{U}_{125}	V^{U}_{150}	V^{bo}_{200}	Time
IPIP	41	59	1	96	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1801
IPIP-R	99	n/a	50	49	1.36	0.00	0.00	0.00	1.30	0.00	0.40	0.00	0.00	16
IPIP-H	97	3	41	58	0.56	0.00	0.00	0.00	0.80	0.00	0.08	0.00	0.00	11
IPSA	97	n/a	38	61	0.22	0.00	0.00	0.00	0.73	0.00	0.04	0.00	0.29	5

Table A.10: IPIP Study Patient 10 Dosimetric Results

Alg.	V^{P}_{100}	Gap	V^{P}_{150}	HI	V^{B}_{75}	V^{B}_{100}	V^{b}_{75}	V^{b}_{100}	V^{R}_{75}	V^{R}_{100}	V^{U}_{125}	V^{U}_{150}	V^{bo}_{200}	Time
IPIP	0	100	0		0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1801
IPIP-R	100	n/a	43	57	2.84	0.00	0.00	0.00	2.44	0.00	0.28	0.00	0.00	34
IPIP-H	96	4	42	56	1.00	0.00	0.00	0.00	0.86	0.00	0.07	0.00	0.00	18
IPSA	98	n/a	30	70	1.75	0.00	0.00	0.00	0.42	0.00	0.17	0.00	0.19	16

Table A.11: IPIP Study Patient 11 Dosimetric Results

Alg.	V^{P}_{100}	Gap	V^{P}_{150}	HI	V^{B}_{75}	V^{B}_{100}	V^{b}_{75}	V^{b}_{100}	V^{R}_{75}	V^{R}_{100}	V^{U}_{125}	V^{U}_{150}	V^{bo}_{200}	Time
IPIP	99	1	27	73	0.80	0.00	0.00	0.00	0.91	0.00	0.01	0.00	0.00	1800
IPIP-R	100	n/a	41	59	0.89	0.00	0.00	0.00	2.32	0.00	0.36	0.00	0.00	10
IPIP-H	98	2	25	74	0.76	0.00	0.00	0.00	0.76	0.00	0.04	0.00	0.00	9
IPSA	97	n/a	26	73	0.32	0.00	0.00	0.00	0.95	0.01	0.07	0.00	0.05	3

Table A.12: IPIP Study Patient 12 Dosimetric Results

Alg.	V^{P}_{100}	Gap	V^{P}_{150}	HI	V^{B}_{75}	V^{B}_{100}	V^{b}_{75}	V^{b}_{100}	V^{R}_{75}	V^{R}_{100}	V^{U}_{125}	V^{U}_{150}	V^{bo}_{200}	Time
IPIP	0	100	0		0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1801
IPIP-R	99	n/a	42	57	3.11	0.00	0.00	0.00	2.28	0.00	0.41	0.00	0.00	34
IPIP-H	96	4	31	68	0.89	0.00	0.00	0.00	0.72	0.00	0.07	0.00	0.00	27
IPSA	98	n/a	23	77	2.40	0.01	0.00	0.00	0.58	0.00	0.04	0.00	1.58	24

Table A.13: IPIP Study Patient 13 Dosimetric Results

Alg.	V^{P}_{100}	Gap	V^{P}_{150}	HI	V^{B}_{75}	V^{B}_{100}	V^{b}_{75}	V^{b}_{100}	V^{R}_{75}	V^{R}_{100}	V^{U}_{125}	V^{U}_{150}	V^{bo}_{200}	Time
IPIP	0	100	0		0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1801
IPIP-R	100	n/a	42	58	2.58	0.00	0.00	0.00	3.04	0.00	0.35	0.00	0.00	28
IPIP-H	97	3	30	69	0.91	0.00	0.00	0.00	0.82	0.00	0.08	0.00	0.00	17
IPSA	98	n/a	24	75	1.01	0.01	0.00	0.00	0.96	0.00	0.01	0.00	0.48	11

Table A.14: IPIP Study Patient 14 Dosimetric Results

Alg.	V^{P}_{100}	Gap	V^{P}_{150}	HI	V^{B}_{75}	V^{B}_{100}	V^{b}_{75}	V^{b}_{100}	V^{R}_{75}	V^{R}_{100}	V^{U}_{125}	V^{U}_{150}	V^{bo}_{200}	Time
IPIP	99	0	26	74	0.80	0.00	0.00	0.00	0.22	0.00	0.06	0.00	0.00	1801
IPIP-R	99	n/a	37	63	1.28	0.00	0.00	0.00	1.13	0.00	0.22	0.00	0.00	6
IPIP-H	99	1	30	70	0.88	0.00	0.00	0.00	0.34	0.00	0.00	0.00	0.00	10
IPSA	99	n/a	29	70	1.19	0.00	0.00	0.00	0.25	0.00	0.00	0.00	0.10	3

Table A.15: IPIP Study Patient 15 Dosimetric Results

Alg.	V^{P}_{100}	Gap	V^{P}_{150}	HI	V^{B}_{75}	V^{B}_{100}	V^{b}_{75}	V^{b}_{100}	V^{R}_{75}	V^{R}_{100}	V^{U}_{125}	V^{U}_{150}	V^{bo}_{200}	Time
IPIP	55	45	1	98	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1801
IPIP-R	98	n/a	39	60	1.86	0.00	0.00	0.00	2.04	0.00	0.29	0.00	0.00	18
IPIP-H	94	5	27	72	0.96	0.00	0.00	0.00	0.53	0.00	0.05	0.00	0.00	11
IPSA	97	n/a	29	70	1.51	0.00	0.00	0.00	0.79	0.00	0.48	0.10	0.67	6

Table A.16: IPIP Study Patient 16 Dosimetric Results

Alg.	V^{P}_{100}	Gap	V^{P}_{150}	HI	V^{B}_{75}	V^{B}_{100}	V^{b}_{75}	V^{b}_{100}	V^{R}_{75}	V^{R}_{100}	V^{U}_{125}	V^{U}_{150}	V^{bo}_{200}	Time
IPIP	100	0	35	65	0.48	0.00	0.00	0.00	0.00	0.00	0.10	0.00	0.00	1800
IPIP-R	100	n/a	52	48	0.84	0.00	0.00	0.00	0.06	0.00	0.52	0.00	0.00	18
IPIP-H	99	1	39	60	0.65	0.00	0.00	0.00	0.29	0.00	0.06	0.00	0.00	14
IPSA	98	n/a	29	70	0.25	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.38	4

Table A.17: IPIP Study Patient 17 Dosimetric Results

Alg.	V^{P}_{100}	Gap	V^{P}_{150}	HI	V^{B}_{75}	V^{B}_{100}	V^{b}_{75}	V^{b}_{100}	V^{R}_{75}	V^{R}_{100}	V^{U}_{125}	V^{U}_{150}	V^{bo}_{200}	Time
IPIP	0	100	0		0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1800
IPIP-R	99	n/a	42	57	0.31	0.00	0.00	0.00	1.51	0.00	0.29	0.00	0.00	8
IPIP-H	98	2	40	59	0.28	0.00	0.00	0.00	0.89	0.00	0.10	0.00	0.00	9
IPSA	98	n/a	31	69	0.26	0.00	0.00	0.00	0.36	0.00	0.23	0.00	0.29	2

Table A.18: IPIP Study Patient 18 Dosimetric Results

A.2 Additional DVH Results

Although this dissertation did not use DVH's to compare dose distributions, they are a common method of characterizing a dose distribution in practice. For that reason, we give DVH for each patient and organ for the four algorithms considered in this study. Here IPIP is in black, IPIP-R in blue, IPIP-H in green, and IPSA in red. Note that the IPSA DVH is below the other algorithms' for almost every patient and organ. This is because IPIP generally produces hotter solutions than IPSA, since its objective is to maximize target coverage.

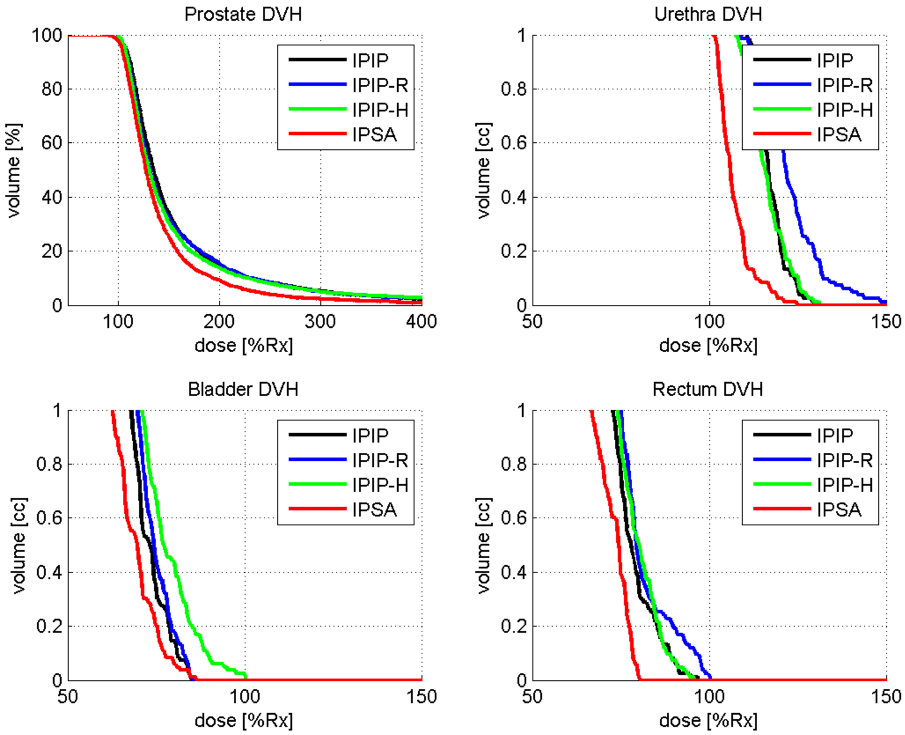

Figure A.1: IPIP DVH: Patient 1

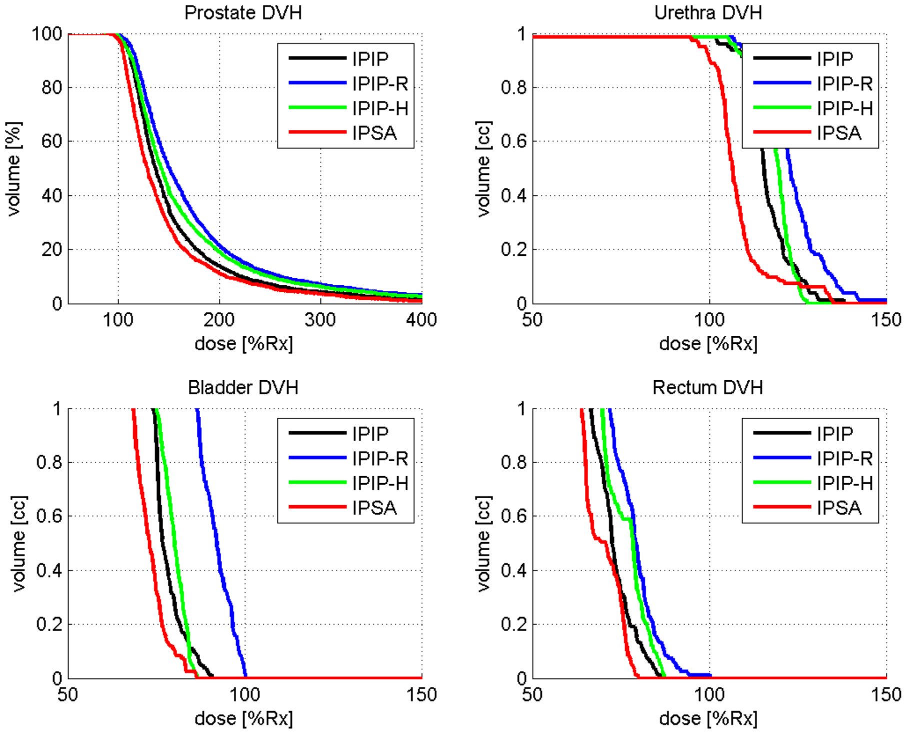

Figure A.2: IPIP DVH: Patient 2

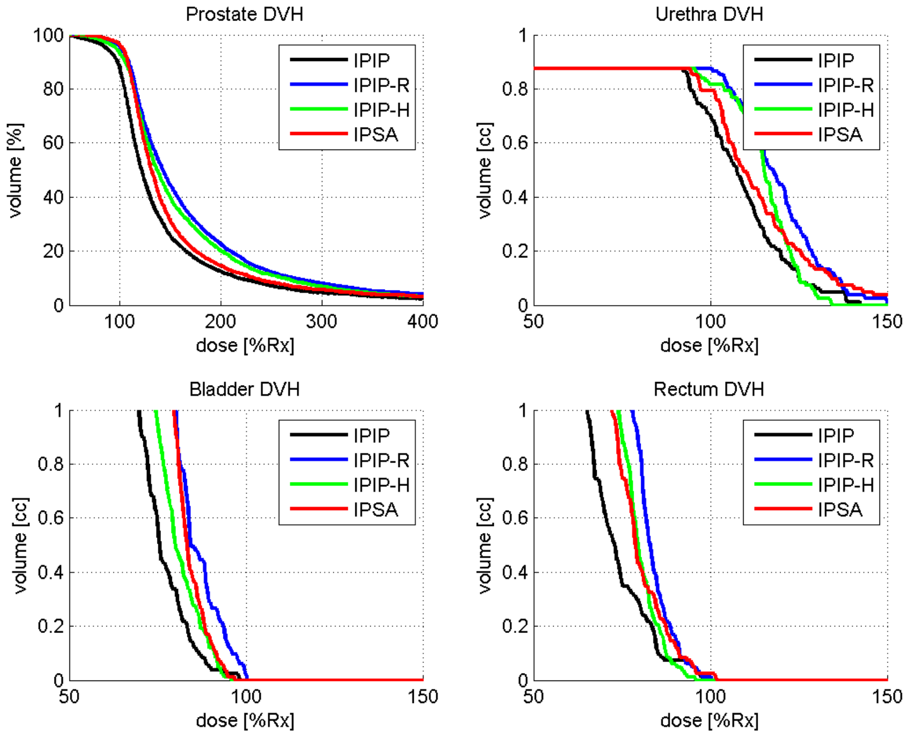

Figure A.3: IPIP DVH: Patient 3

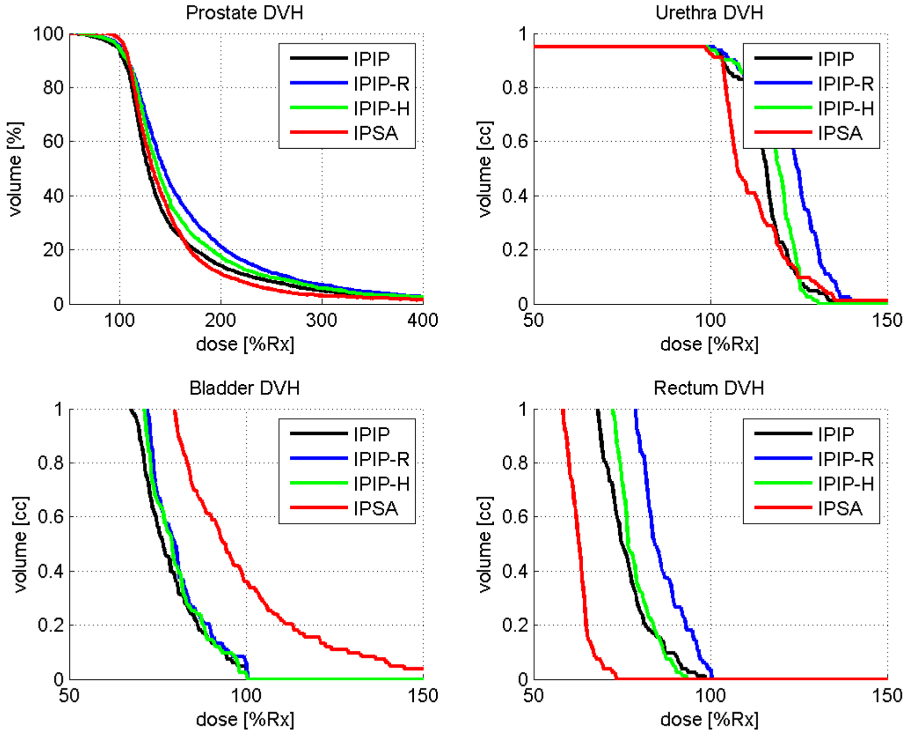

Figure A.4: IPIP DVH: Patient 4

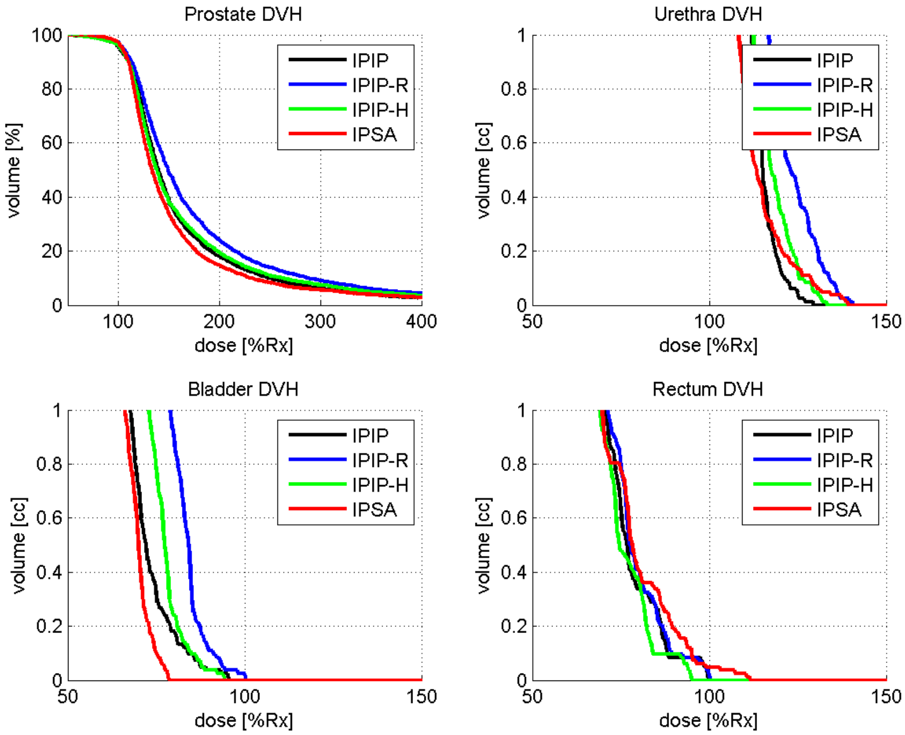

Figure A.5: IPIP DVH: Patient 5

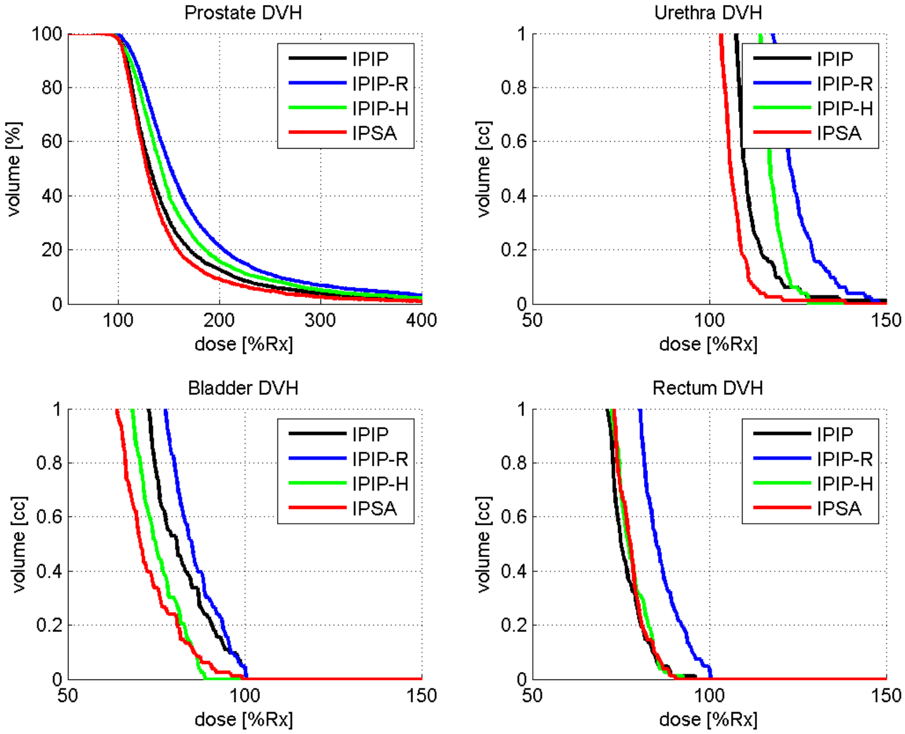

Figure A.6: IPIP DVH: Patient 6

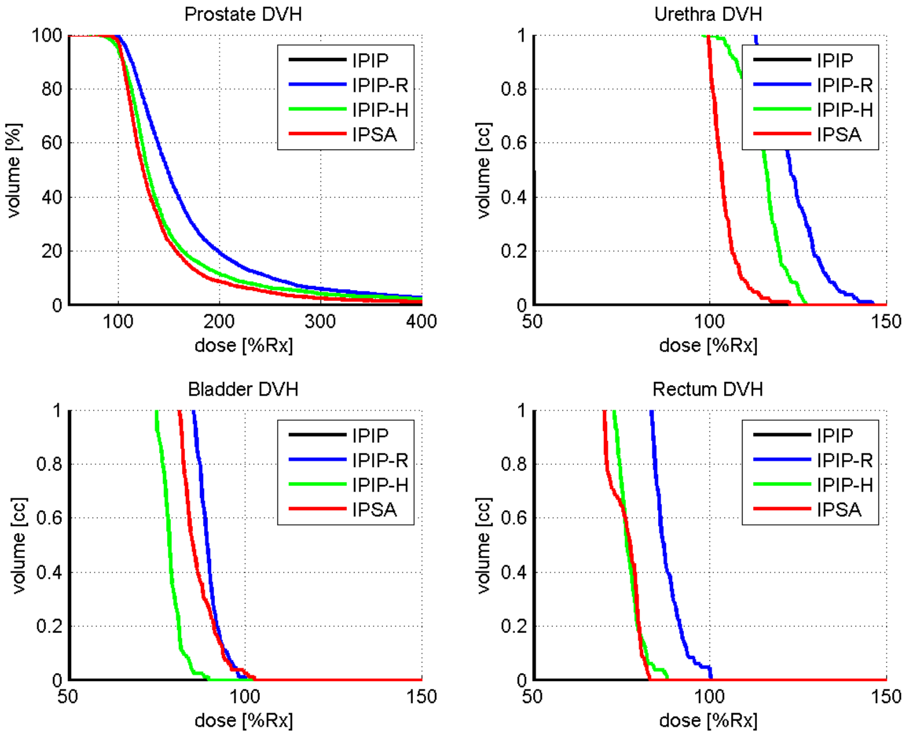

Figure A.7: IPIP DVH: Patient 7

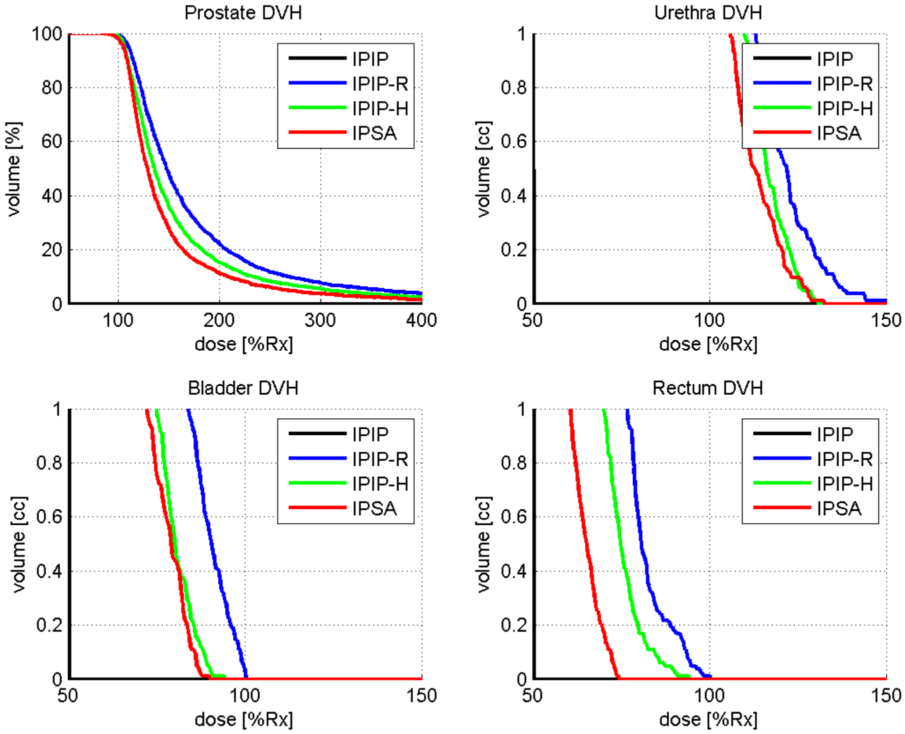

Figure A.8: IPIP DVH: Patient 8

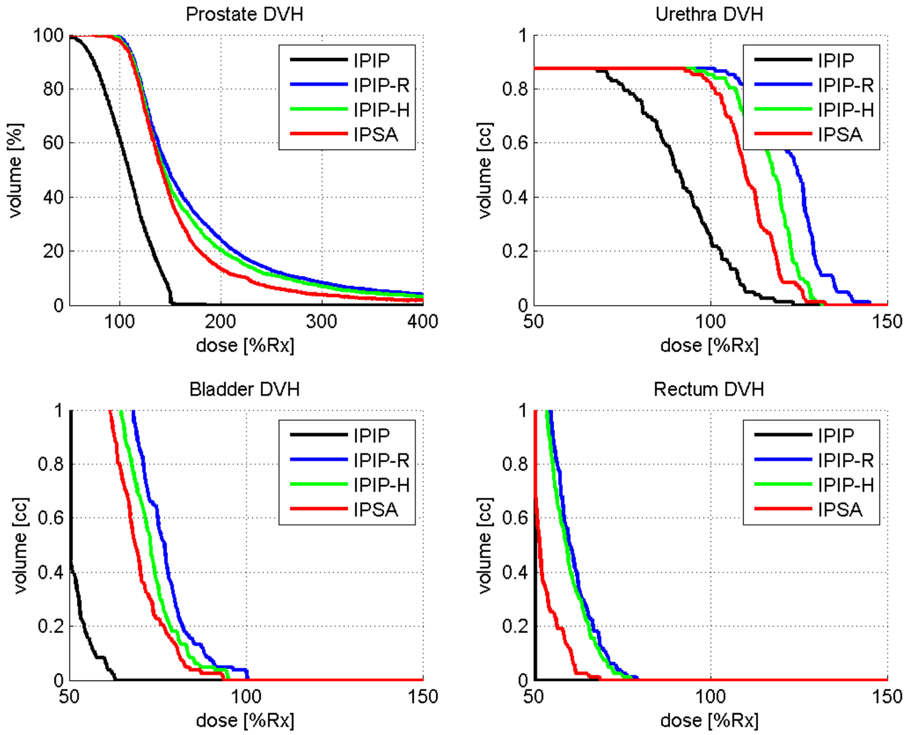

Figure A.9: IPIP DVH: Patient 9

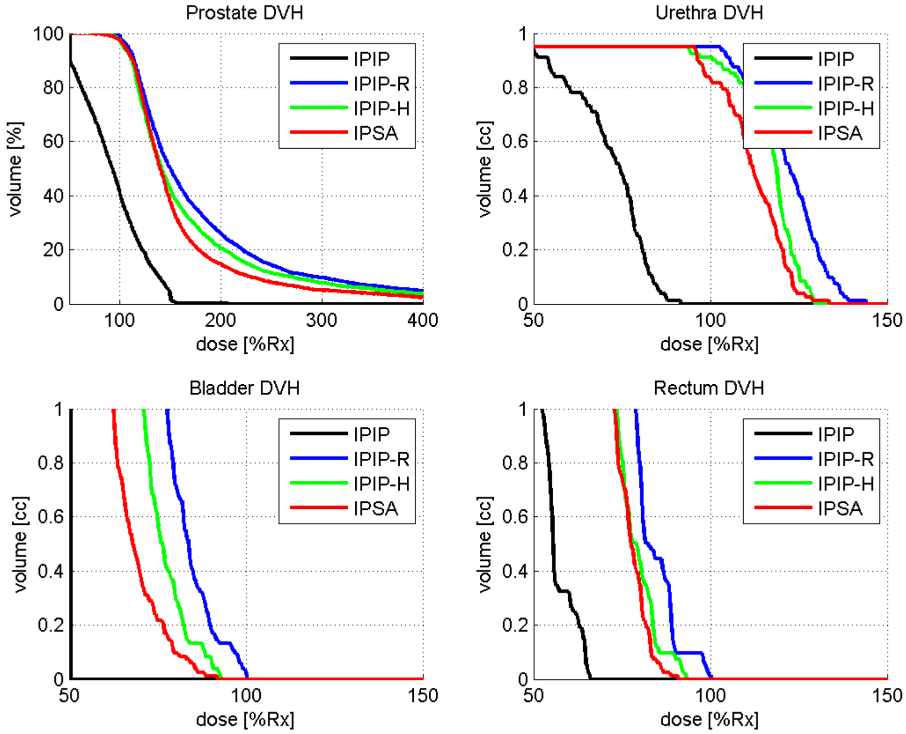

Figure A.10: IPIP DVH: Patient 10

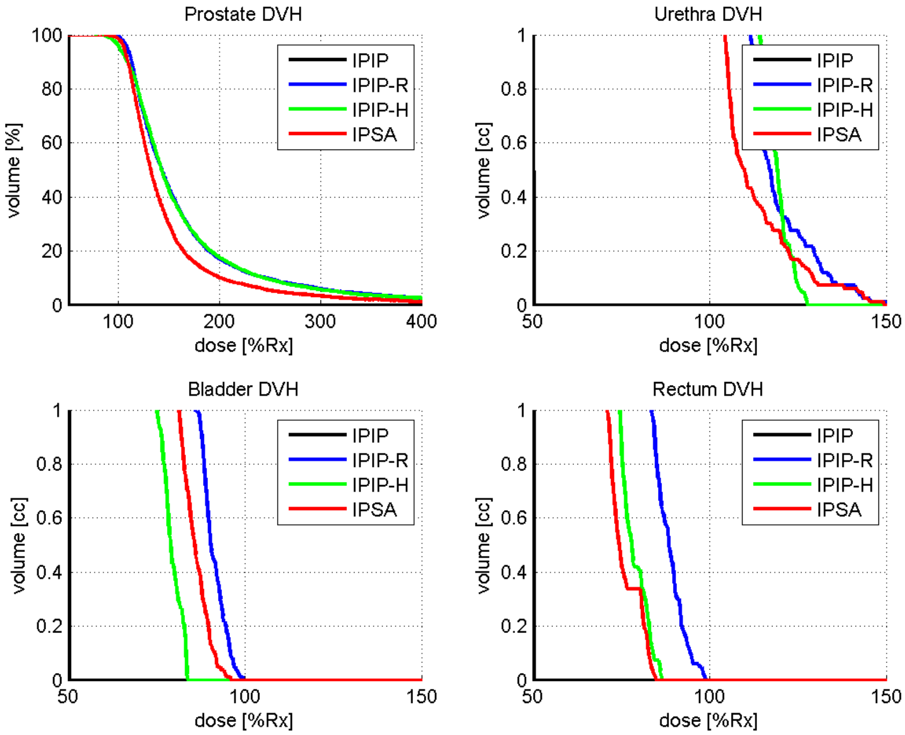

Figure A.11: IPIP DVH: Patient 11

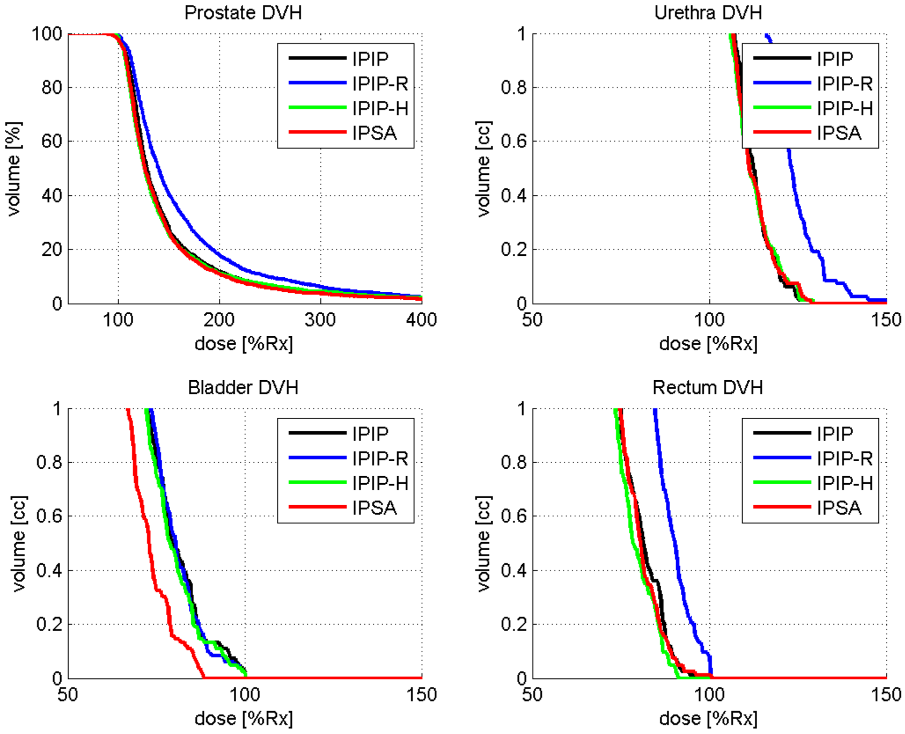

Figure A.12: IPIP DVH: Patient 12

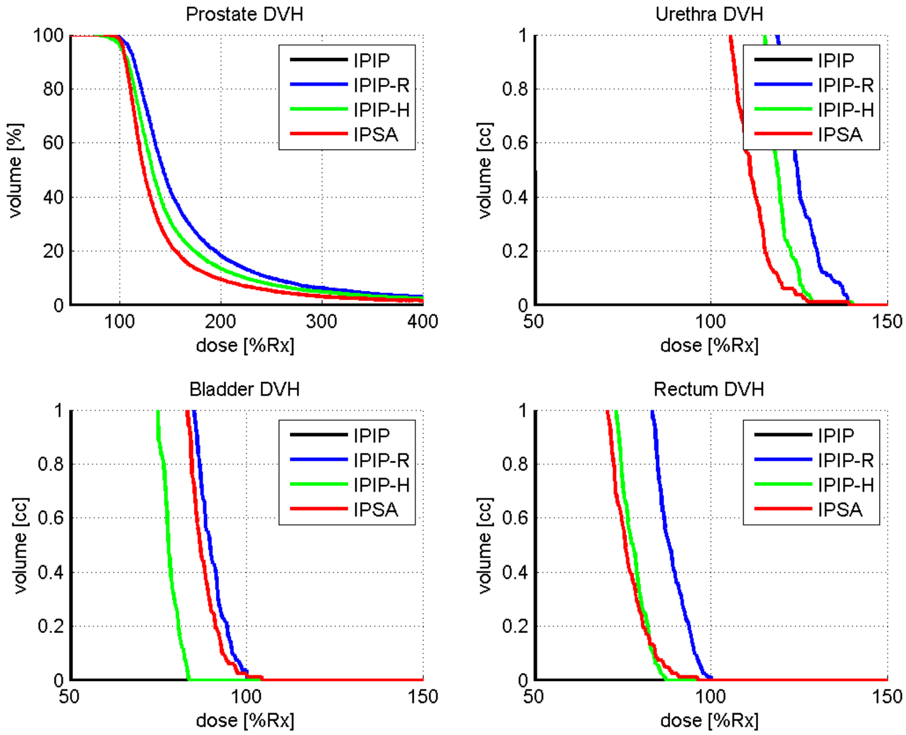

Figure A.13: IPIP DVH: Patient 13

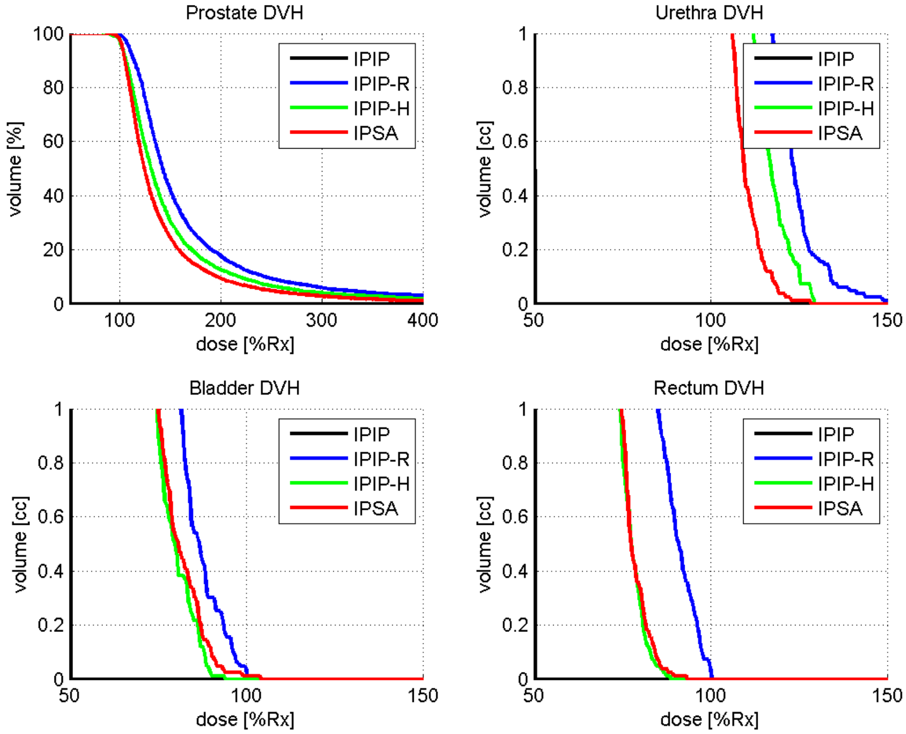

Figure A.14: IPIP DVH: Patient 14

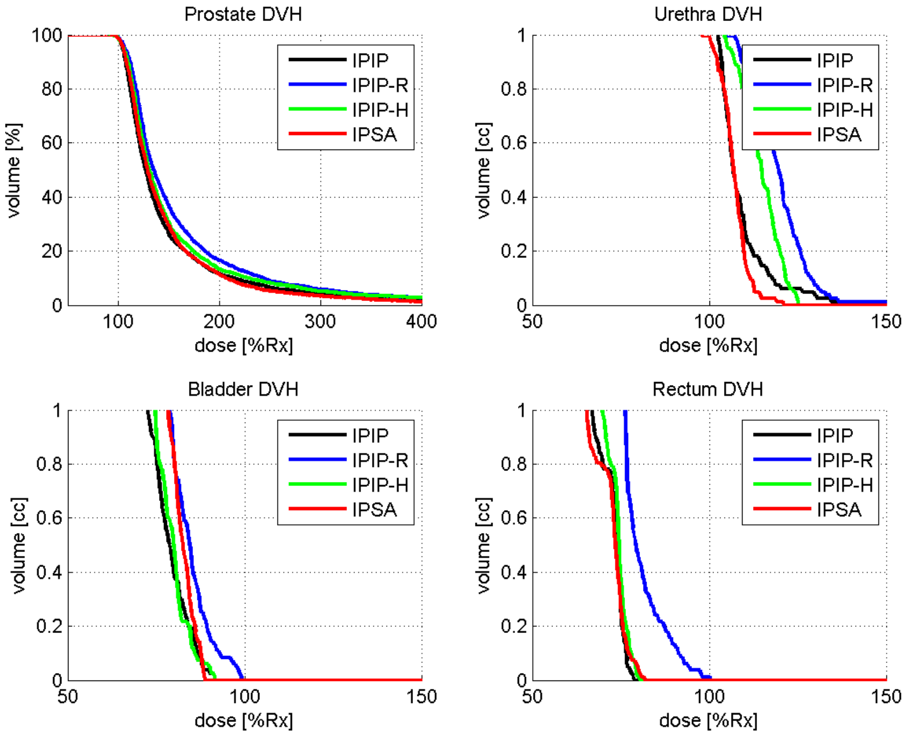

Figure A.15: IPIP DVH: Patient 15

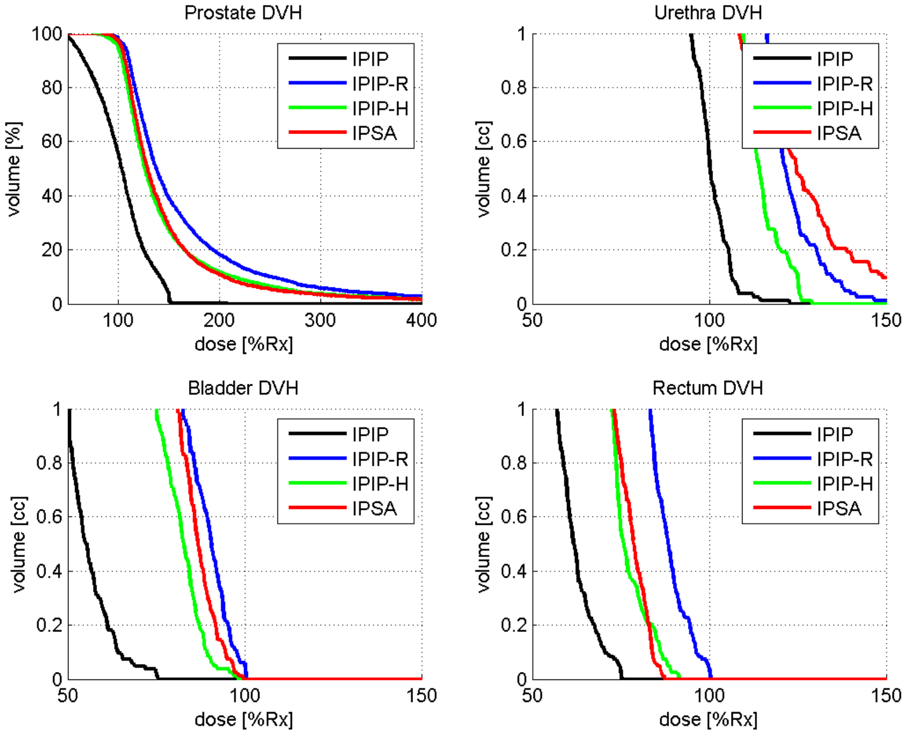

Figure A.16: IPIP DVH: Patient 16

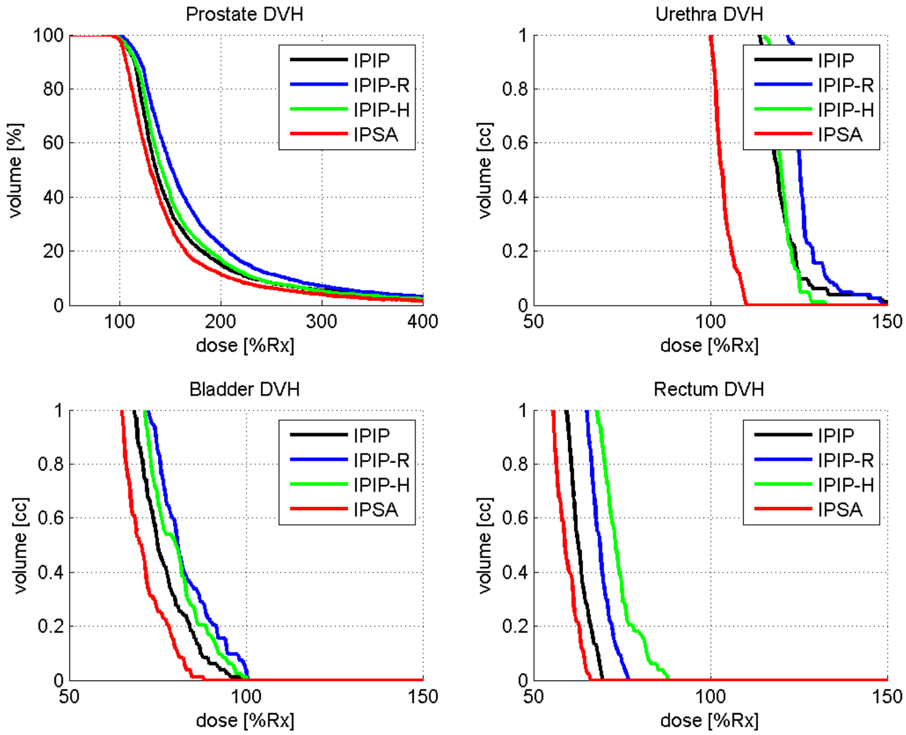

Figure A.17: IPIP DVH: Patient 17

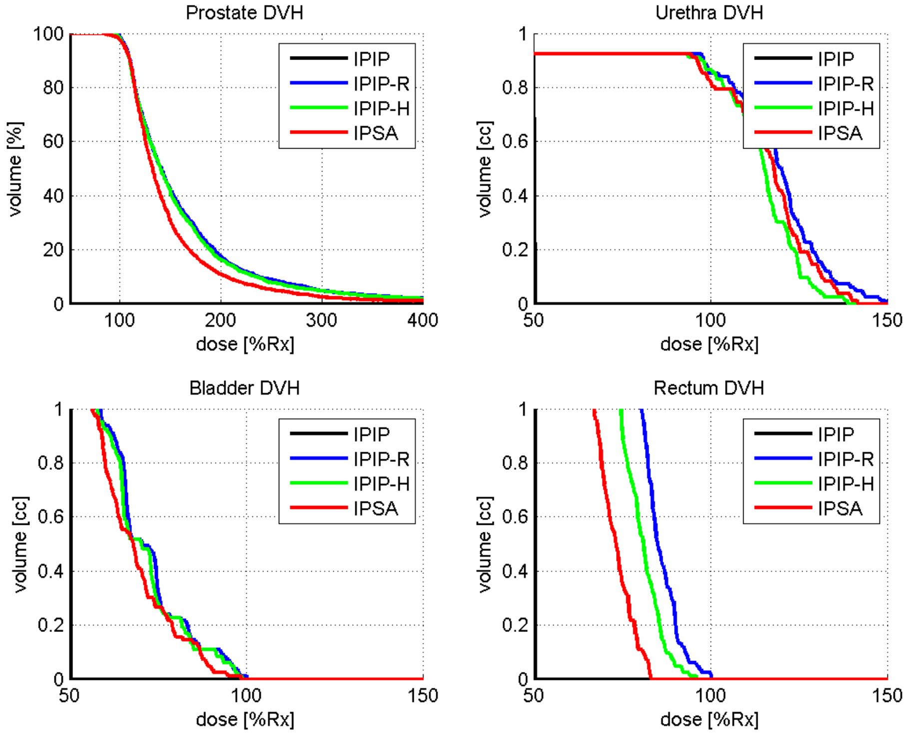

Figure A.18: IPIP DVH: Patient 18

Appendix B

Needle Planning Appendix

B.1 Additional Robustness Results

This section contains the robustness figures for each patient from the sensitivity analysis in our NPIP experiment. On the x-axis is δ. On the y-axis is the number of needles resulting from that δ (bar graph) and the target coverage for the nominal needle configuration is shown as red dots. The average target coverage over the 25 perturbations, as well as one standard deviation error bars, are shown in blue. Note that infeasible instances of NPIP register as 0 needles and 0 target coverage for these plots. For clarity, a black horizontal line is given at the 90% target coverage, the minimum cut off to meet our objectives. Note that because dose distributions were computed with IPIP, all dose constraints for OAR were always met.

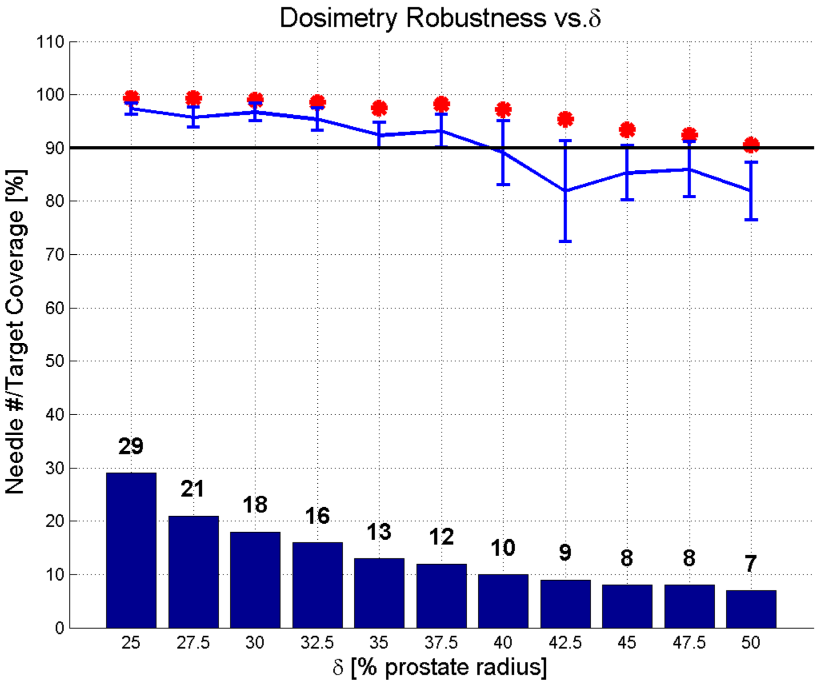

Figure B.1: ROBUSTNESS PLOT: Patient 1

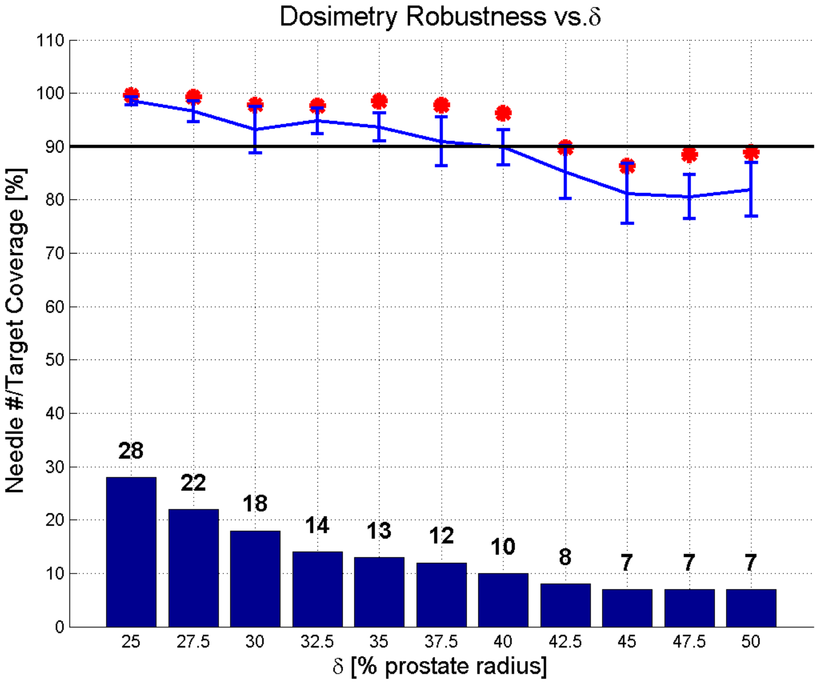

Figure B.2: ROBUSTNESS PLOT: Patient 2

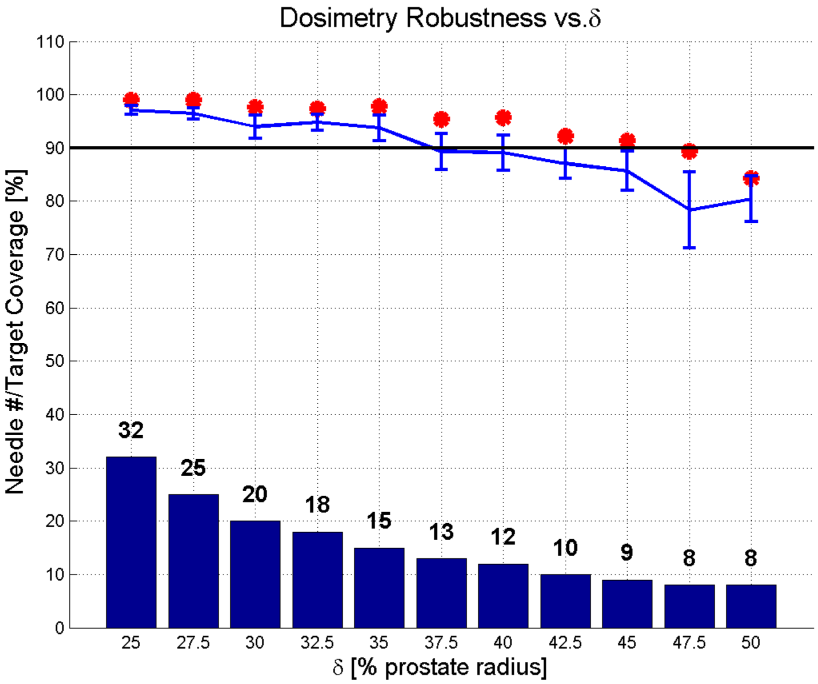

Figure B.3: ROBUSTNESS PLOT: Patient 3

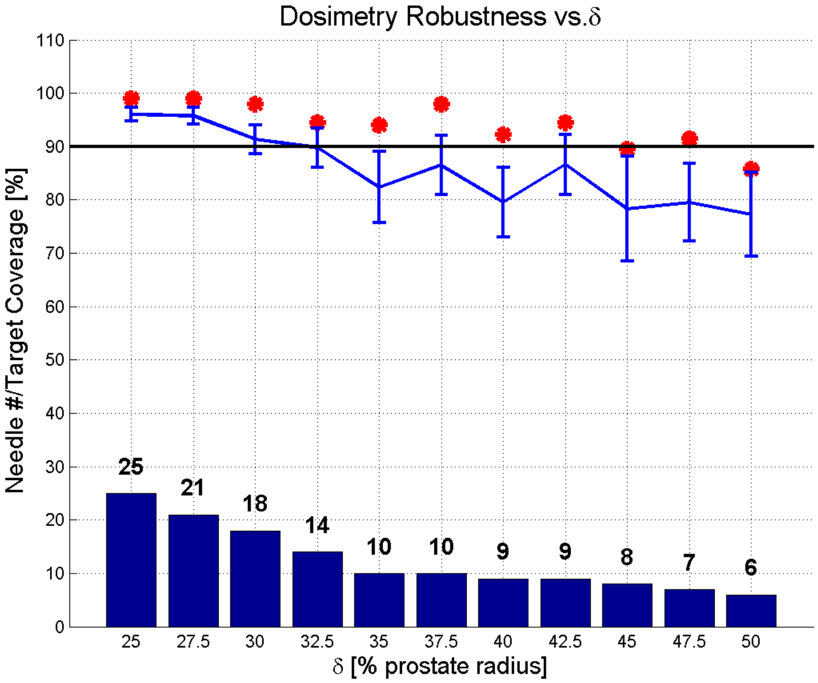

Figure B.4: ROBUSTNESS PLOT: Patient 4

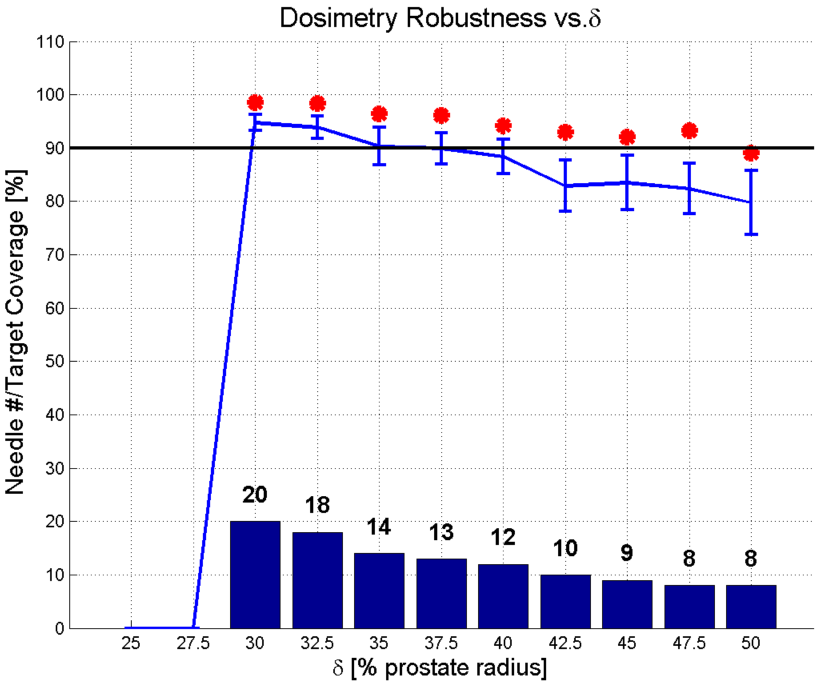

Figure B.5: ROBUSTNESS PLOT: Patient 5

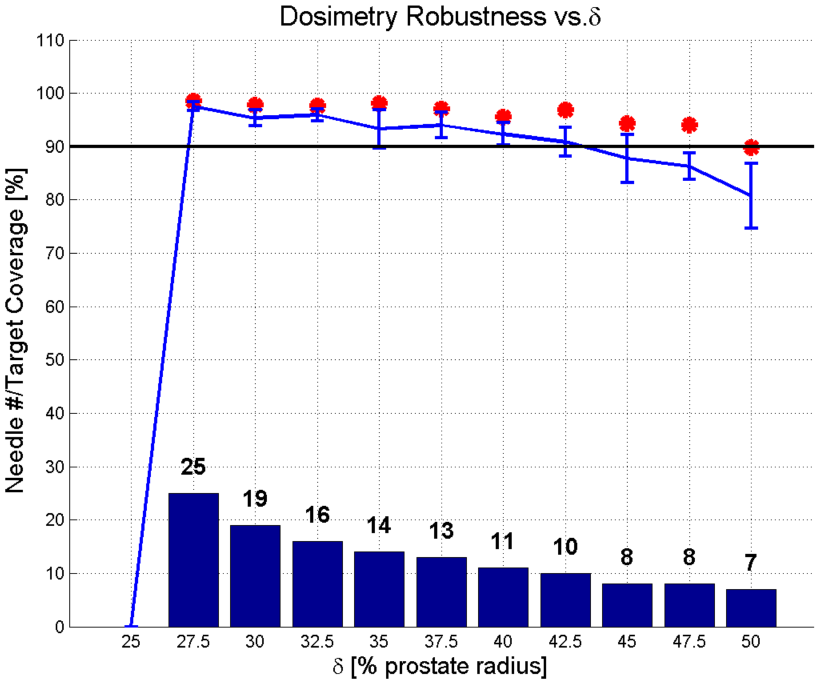

Figure B.6: ROBUSTNESS PLOT: Patient 6

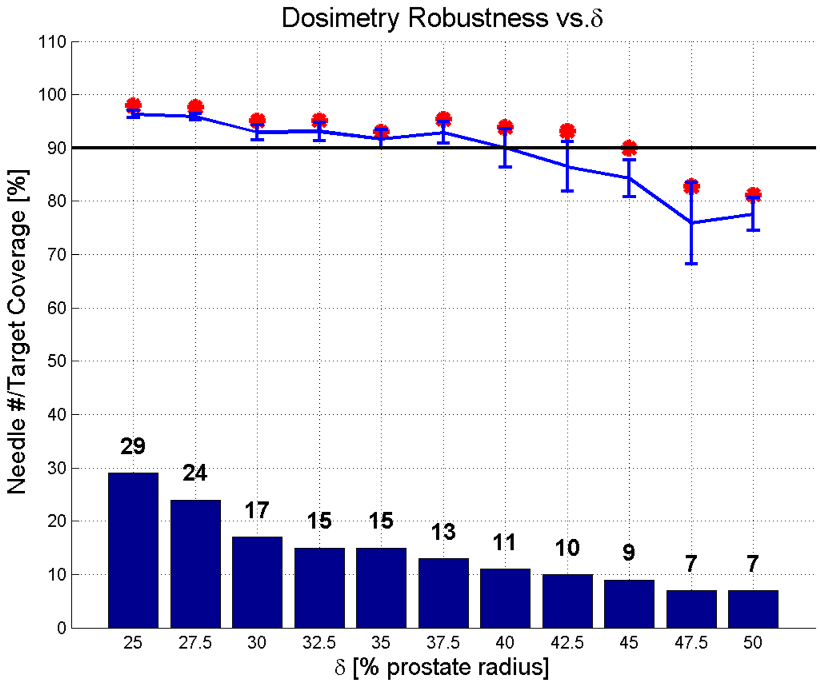

Figure B.7: ROBUSTNESS PLOT: Patient 7

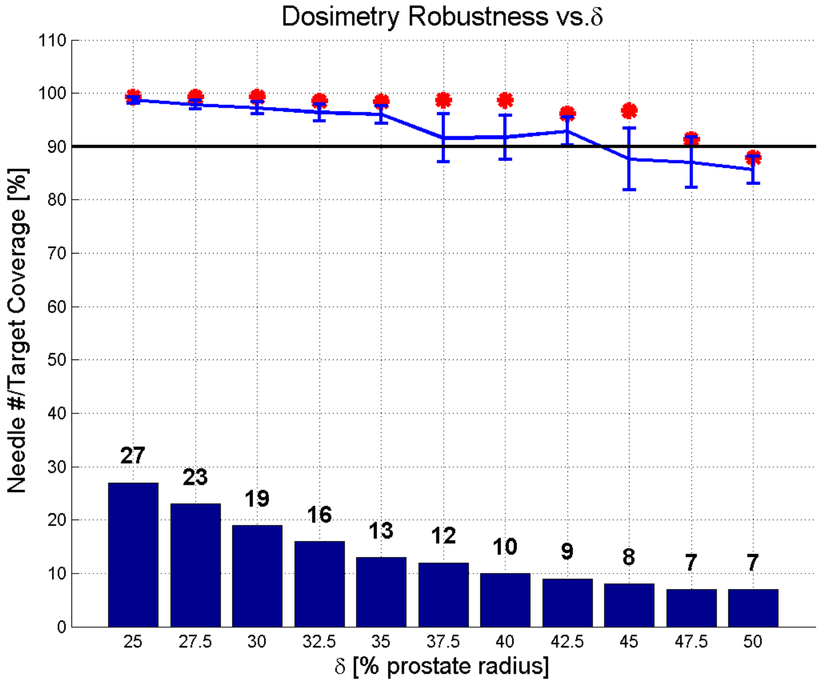

Figure B.8: ROBUSTNESS PLOT: Patient 8

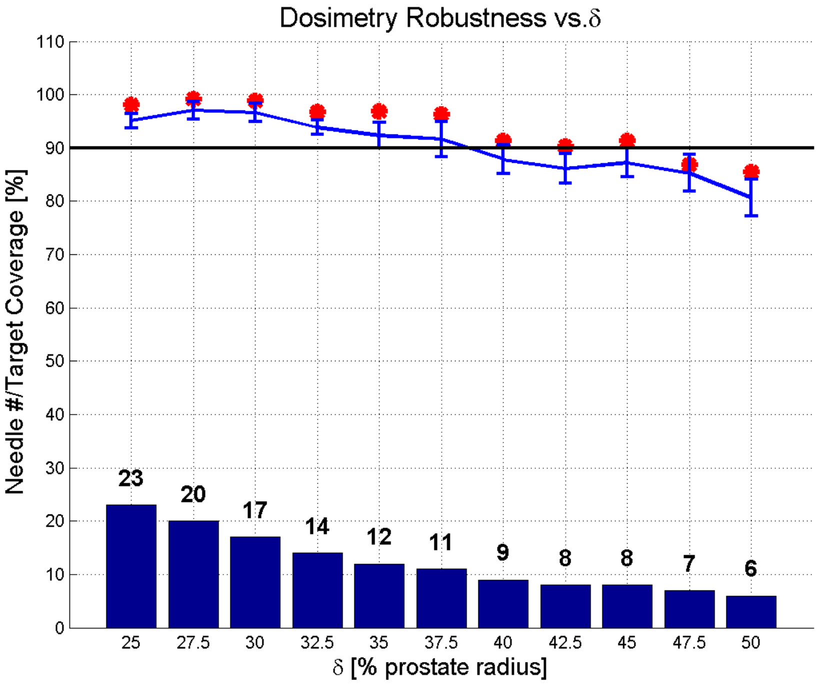

Figure B.9: ROBUSTNESS PLOT: Patient 9

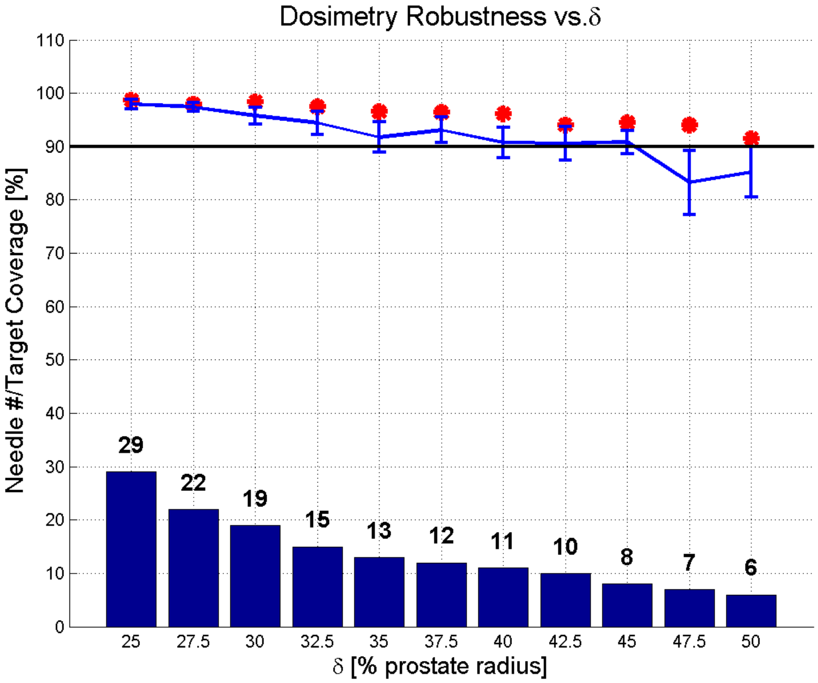

Figure B.10: ROBUSTNESS PLOT: Patient 10

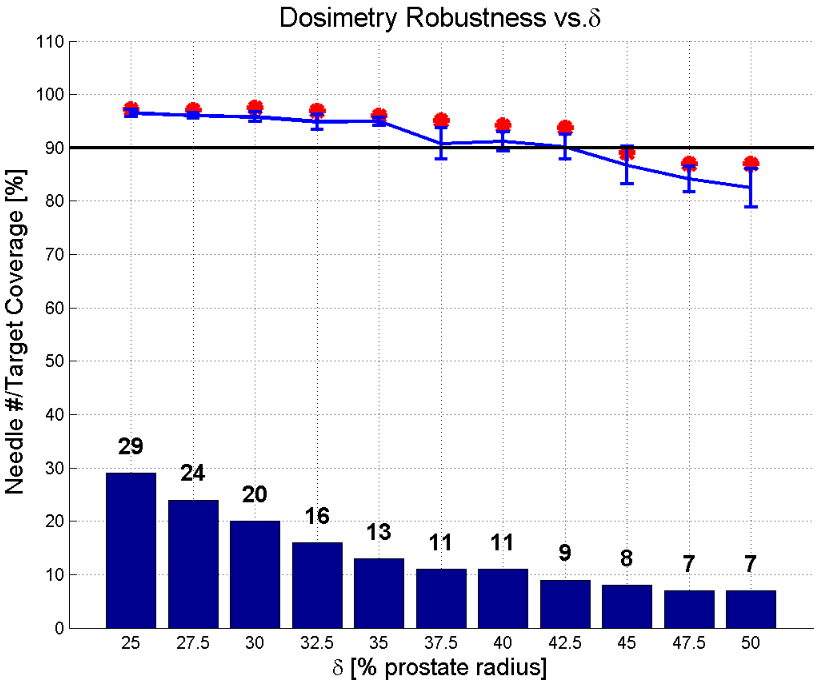

Figure B.11: ROBUSTNESS PLOT: Patient 11

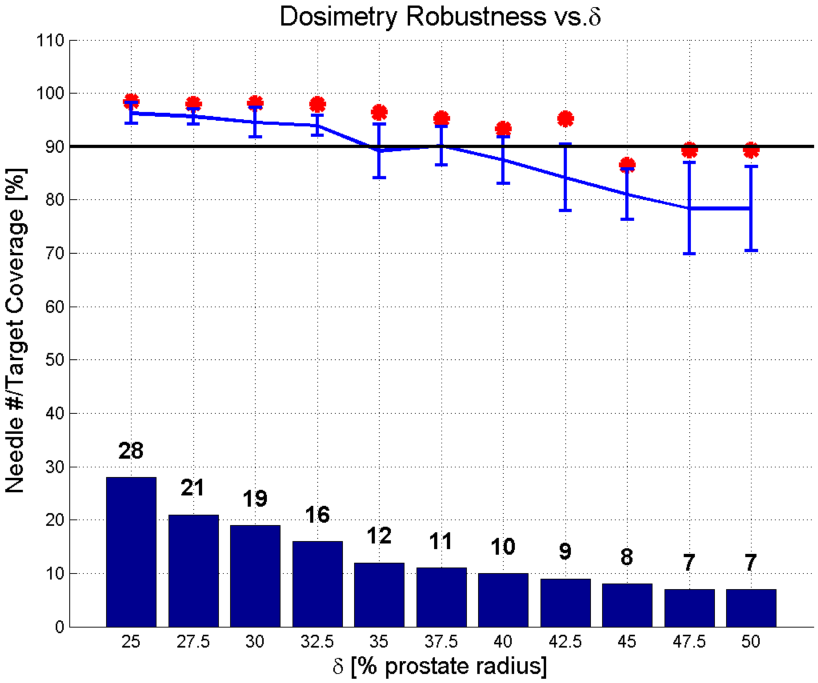

Figure B.12: ROBUSTNESS PLOT: Patient 12

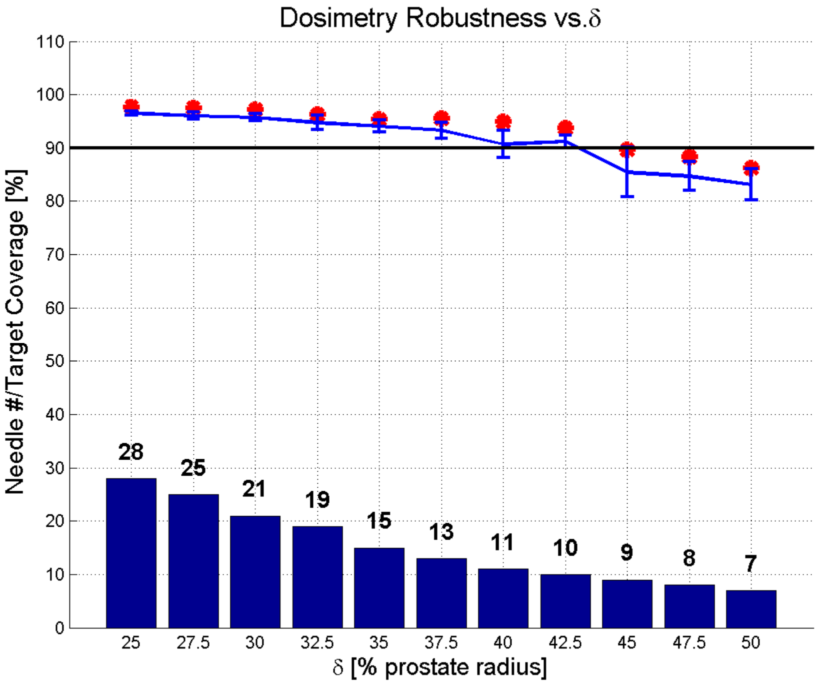

Figure B.13: ROBUSTNESS PLOT: Patient 13

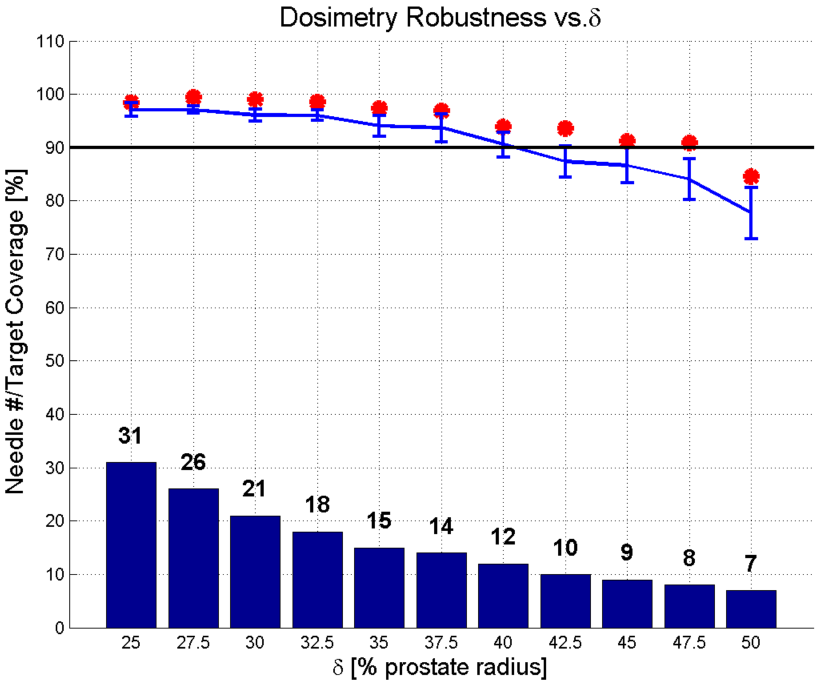

Figure B.14: ROBUSTNESS PLOT: Patient 14

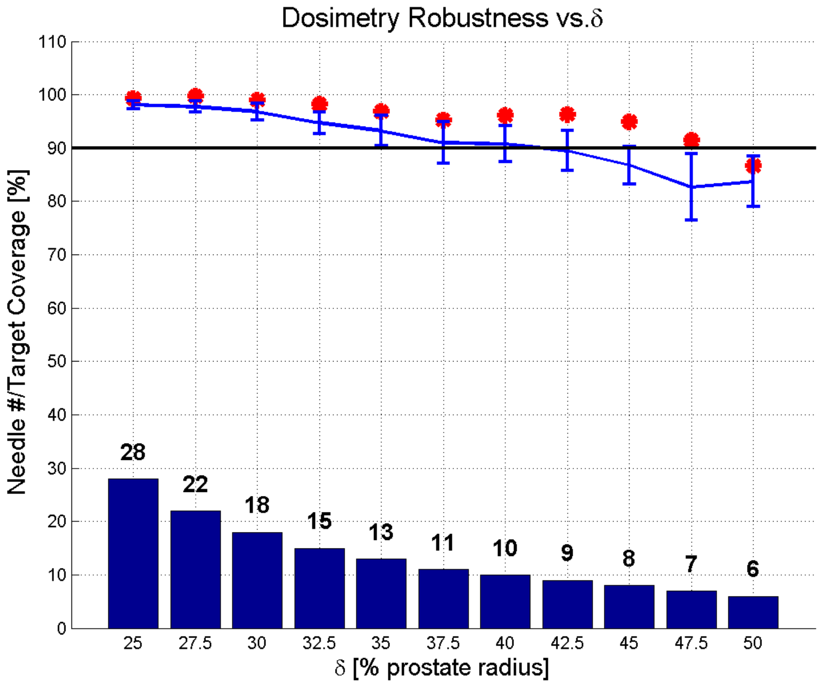

Figure B.15: ROBUSTNESS PLOT: Patient 15

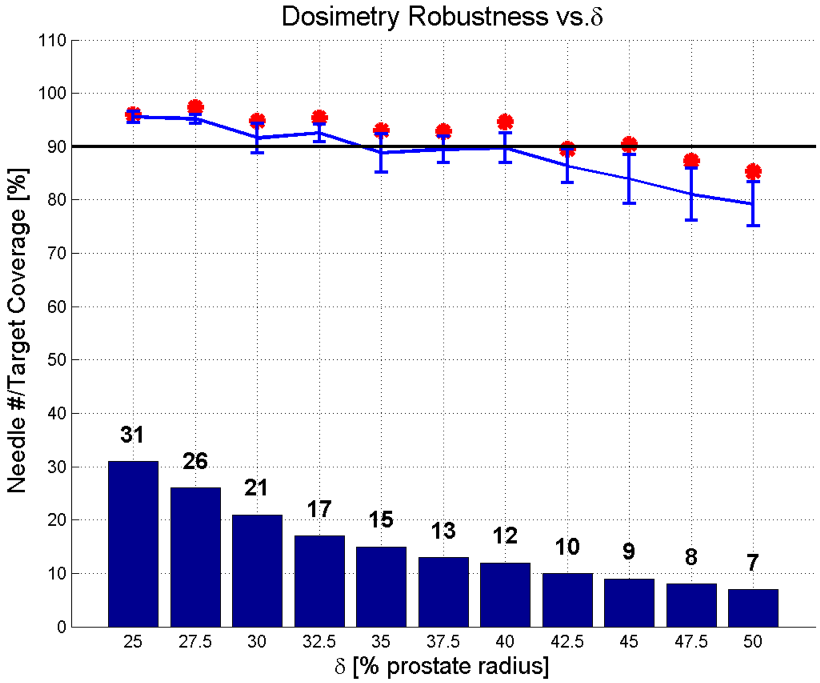

Figure B.16: ROBUSTNESS PLOT: Patient 16

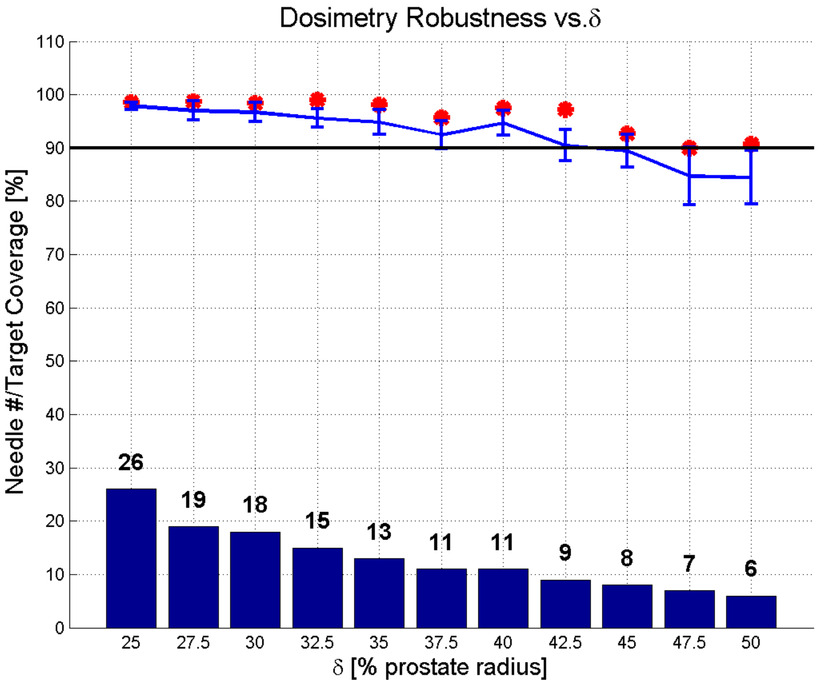

Figure B.17: ROBUSTNESS PLOT: Patient 17

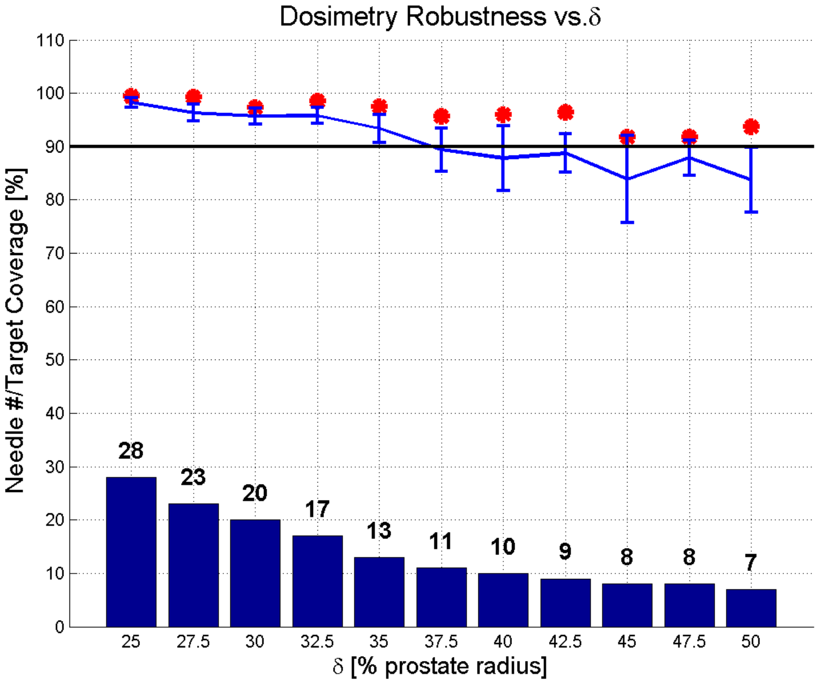

Figure B.18: ROBUSTNESS PLOT: Patient 18