

**A Mixed Integer Linear Programming Optimization Model for Capturing Expert Planning
Style in Interstitial Low Dose Rate Prostate Brachytherapy**

by

Mustafa Ege Babadagli

A thesis submitted in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

in

Engineering Management

Department of Mechanical Engineering

University of Alberta

© Mustafa Ege Babadagli, 2016

Abstract

Low dose rate (LDR) brachytherapy is a minimally invasive form of radiation therapy, used to treat prostate cancer, and it involves permanent implantation of radioactive sources (seeds) inside of the prostate gland. Treatment planning in brachytherapy involves a decision making process for the placement of radioactive sources in order to deliver an effective dose of radiation to cancerous tissue in the prostate while sparing the surrounding healthy tissue. Such a decision making process can be modeled as a mixed-integer linear programming (MILP) problem. In this thesis, we initially introduce a novel MILP optimization model for interstitial LDR prostate brachytherapy that attempts to mimic the qualities of treatment plans produced manually by expert planners. Our approach involves incorporating a unique set of clinically important constraints, called spatial constraints, into our models. These constraints are combined with several proposed data processing techniques, and collectively these methods enable us to capture the fundamental aspects of the planning style present at a local cancer clinic. While this represents the primary clinical challenge to overcome in this thesis, the engineering challenge involves producing clinically acceptable treatment plans in a matter of seconds to minutes with our automated planning system. Such a goal initially presents itself to be a difficult one to accomplish, as our large-scale models exhibit tendencies of intractability while using high-resolution data sets. We introduce pseudo high-resolution data sets and constraint-violating feasibility-based modelling in order to overcome these issues related to solution performance of our model. Through the synergistic effects of these two modelling techniques, we demonstrate the ability to produce treatment plans with solution times suitable for pre-operative and intra-operative planning, which constitute the two main forms of treatment planning in LDR brachytherapy. Specifically speaking, our solution

times range from less than a minute to roughly five minutes for prostates of varying shapes and sizes (volume-wise, the tested prostates range from 20.4 cc to 63.1 cc). Finally, with the help of a radiation oncologist, we verify the clinical acceptability of our automated plans through a pilot study involving data from twenty patients. Therefore, through this study we also confirm the clinical suitability of the several concepts introduced in this thesis; namely, these are spatial constraints, pseudo high-resolution data sets and constraint-violating feasibility-based modelling. Following these results, we consequently identify three main areas of improvements for our automated system: avoiding placement of strands outside of the prostate and seeds in the bladder, as well as lowering dose to the organs-at-risk. Moving forward, we suggest branching out and examining the clinical performance of our automated planning system in intra-operative planning, as the results obtained in thesis reflect the performance of our approach within pre-operative planning. In summary, through this thesis we demonstrate the ability to uniquely capture the planning style of expert planners at a local clinic in a fraction of the time that it currently takes expert planners to produce treatment plans.

Preface

This thesis is an original work by Mustafa Ege Babadagli. The research project, of which this thesis is a part, received research ethics approval from the University of Alberta Research Ethics Board.

Project Name:

“Mathematical Optimization Models for Practical Clinical Implementation in Brachytherapy”

Study ID:

Pro00056288

Date of Approval:

20 August 2015

The research for this thesis was conducted under the supervision of Dr. John Doucette and Dr. Ron Sloboda from the University of Alberta, and in collaboration with Dr. Nawaid Usmani from the University of Alberta.

Chapter 3 of this thesis has been submitted for publication as:

Babadagli, M.E.; Sloboda, R.; Doucette, J. “An Integer Linear Programming Optimization Model for Capturing Expert Planning Style in Interstitial Low Dose Rate Prostate Brachytherapy”. Springer, Annals of Operations Research.

Chapter 4 of this thesis has been submitted for publication as:

Babadagli, M.E.; Sloboda, R.; Doucette, J. “Enhanced Data Processing and Integer Linear Programming Techniques for Capturing Expert Planning Style in Interstitial Low Dose Rate Prostate Brachytherapy”. Springer, Annals of Operations Research.

Chapter 5 of this thesis will be submitted for publication as:

Babadagli, M.E.; Sloboda, R.; Doucette, J. “Mixed-Integer Linear Programming Modelling Techniques for Improving Treatment Plan Solution Times in Interstitial Low Dose Rate Prostate Brachytherapy”. Springer, Annals of Operations Research.

Chapter 6 of this thesis will be submitted for publication as:

Babadagli, M.E.; Sloboda, R.; Usmani, N.; Amanie, J.; Murtha, A.; Yee, D.; Jamaluddin, M.; Doucette, J. “Clinical Assessment of Center-Specific Automated Treatment Plans for Interstitial Low Dose Rate Prostate Brachytherapy”. Elsevier, Brachytherapy.

For the publications stemming from Chapters 3, 4, 5 and 6 of this thesis, I was primarily responsible for mathematical model formulation, study and experimental design, data collection and data analysis as well as manuscript composition. Dr. John Doucette assisted with mathematical model formulation, study and experimental design, data collection, data analysis and contributed to manuscript composition. Dr. Ron Sloboda assisted with mathematical model formulation, study and experimental design, data collection, data analysis and manuscript composition. Additionally, Dr. Nawaid Usmani assisted with the clinical study design and data collection of the research presented in Chapter 6 specifically.

To my family...

Acknowledgements

This interdisciplinary research would not have been possible without the support, collaboration and guidance of many individuals over the past three years.

I am extremely fortunate to have been supervised by two excellent mentors, Dr. John Doucette and Dr. Ron Sloboda. Their collective guidance, support and patience helped me appreciate the nuanced aspects of academia (both in the engineering and the clinical world), and develop myself as an independent researcher. I thank them for continuously encouraging me to strive for high levels of standard in my research work and for teaching me ways to get there. Sculpting this thesis out under their mentorship has been a wonderful learning experience.

I also would like to express my gratitude to Dr. Nawaid Usmani for his keen interest and collaboration in my work, as well as for showing me the inner clinical workings of the prostate brachytherapy operation. I also thank Heather Warkentin for explaining to me the subtleties of manual treatment planning.

I would like to thank Dr. Thomas Hillen, who initially suggested that I probe into the fascinating field of brachytherapy following a long afternoon of discussion on ways to integrate grid-based location problems within medicine.

I also extend my gratitude to my Ph.D. supervisory and examining committee members, Dr. Hooman Askari-Nasab, Dr. Armann Ingolfsson and Dr. Hossein Rouhani for their contribution to improving the overall quality of this thesis.

I would like to recognize a number of organizations for their financial support: the Killam Trusts for the Honorary Izzak Walton Killam Memorial Scholarship, Alberta Innovates for the Alberta Innovates Technology Futures (Doctoral) Scholarship, Faculty of Graduate Studies and

Research (FGSR) for the Queen Elizabeth II Graduate (Doctoral) Scholarship, and the Sadler family for the Sadler Graduate Scholarship.

I also thank Dr. Hasan Uludag for his support and guidance during my journey from engineering to medicine over the many years, especially the last few.

I extend my appreciation to my labmates and friends, with whom I have shared many discussions. Among the many, they are Wenjing Wang, Yali Wang, Roberto Gallardo, Mark Ruhl, Brody Todd, Varun Sharma, Andres Castillo Luongo, Shiya Lin, Sami Alharbi and Katarina Ondrusova. I would also like to thank Richard Groulx and Gail Dowler for the many interesting conversations and administrative help.

Finally, I am grateful to my mother Şule, my father Tayfun, my sister Hazal, and Luigi for all their strength, support, and direction. I would not be exaggerating if I stated that they too completed this Ph.D. alongside me! Thank you for exemplifying, each and every day, what unconditional love and support truly means.

Table of Contents

Chapter 1: Introduction	1
1.1 Mathematical Programming in LDR Prostate Brachytherapy	4
1.2 Challenges and Current Research Gap in LDR Brachytherapy Optimization.....	6
1.3 Research Objectives and Outline of the Thesis	11
Chapter 2: Mixed-Integer Linear Programming (MILP) & Solution Algorithms.....	12
2.1 Mixed Integer Linear Programming	12
2.2 Deterministic Solution Algorithms for Solving MILP Models	15
2.2.1 Simplex Algorithm	16
2.2.2 Branch-and-Bound Algorithm	18
Chapter 3: The Basic MILP Model Formulation	21
3.1 LDR Prostate Brachytherapy Problem Formulation.....	22
3.1.1 Grid and Coordinate System.....	23
3.2 Preprocessing of Contour Data	27
3.2.1 Discretization of Prostate Contour Data	27
3.2.2 Processing and Conversion of Anatomical Contour Data	28
3.2.3 Permissible Seed-Placement Region	29
3.2.4 Seed-Placement Limitations Due to the Urethral Curvature and Rectum	30
3.3 LDR Brachytherapy Optimization Model	34
3.3.1 General Overview	34
3.3.2 Sets, Parameters, Variables.....	34
3.3.3 Objective Function.....	36
3.3.4 1-D Point Source Dose Calculation	38
3.3.5 Dose-Volume Constraints.....	42
3.3.6 Spatial Constraints	47
3.3.6.1 Seed/Position Coupling	47
3.3.6.2 Seed/Needle Coupling	47
3.3.6.3 Avoiding Adjacent Seeds	47

3.3.6.4 Symmetrical Seed Distribution.....	48
3.3.6.5 Symmetrical Position Pairs.....	50
3.3.6.6 Asymmetrical Position Pairs	51
3.3.6.7 Maximum/Minimum Number of Seeds.....	52
3.3.6.8 Maximum/Minimum Number of Needles	53
3.3.6.9 Maximum Needle Retraction.....	53
3.3.6.10 Maximum/Minimum Number of Seeds in a Needle	55
3.3.6.11 Maximum Number of Consecutive Seeds in a Needle	55
3.4 Experimental Design.....	56
3.5 Results.....	58
3.6 Discussion	67
3.7 Conclusion	69
Chapter 4: The Enhanced MILP Model Formulation	71
4.1 Benchmark LDR Prostate Brachytherapy Problem Formulation (Base Model)	73
4.2 Improvements to the Base Model	74
4.2.1 PTV Expansion and Seed-Placement Margins	74
4.2.2 The Pubic Arch	79
4.2.3 The Urethral Pathway and Safety Measures.....	81
4.2.3.1 Removal of Column D (Column 32)	82
4.2.3.2 Separate Seed-Placement and Needle Penetration Margins	83
4.2.4 Maximum Seed Separation in a Needle.....	84
4.2.5 Minimum/Maximum Number of Seeds Per Slice	87
4.3 Experimental Design.....	88
4.4 Results.....	90
4.5 Discussion	98
4.6 Conclusion	99
Chapter 5: Pseudo High-Resolution Data Sets & Feasibility-Based Modelling.....	101
5.1 Pseudo High-Resolution Data Sets	102
5.2 Feasibility-Based Modelling in LDR Brachytherapy	112
5.2.1 Suitability of Feasibility-Based Models	112

5.2.2 Various Methods to Obtain Feasible Solutions	118
5.3 Experimental Design.....	122
5.4 Results.....	131
5.5 Discussion	148
5.6 Conclusion	151
Chapter 6: Clinical Verification of Automated Plans	153
6.1 Study Design.....	154
6.2 Results.....	158
6.3 Discussion	172
6.4 Conclusion	177
Chapter 7: Conclusion & Future Work.....	179
7.1 Pathways to Future Research	181
References	184

List of Figures

Chapter 2: Mixed-Integer Linear Programming (MILP) & Solution Algorithms.....	12
Figure 2.1 Depiction of a graphical solution involving a minimization MILP problem	14
Figure 2.2 Depiction of a solution process using the simplex algorithm.....	17
Figure 2.3 Depiction of a branch-and-bound algorithm search tree	18
 Chapter 3: The Basic MILP Model Formulation	 21
Figure 3.1 Process flow of our automated treatment planning system	25
Figure 3.2 A 0.5x0.5 cm ² coarse-resolution (x-y plane) physical grid	26
Figure 3.3 Three-dimensional coordinate system	26
Figure 3.4 Coarse-resolution grid (red) superimposed on the fine-resolution grid (pink).....	26
Figure 3.5 Discretized PTV in coarse-resolution (A) and fine-resolution (B)	28
Figure 3.6 Visualization of the physical urethral region and the physical rectal region	32
Figure 3.7 Symmetrical and asymmetrical position pairs.....	49
Figure 3.8 Legend for treatment plan isodoses and contours.	62
Figure 3.9. Comparison of treatment plans (slices z=0.0 cm to z=1.5 cm)	63
Figure 3.10 Comparison of treatment plans (slices z=2.0 cm to z=3.5 cm)	64
Figure 3.11 Comparison of treatment plans (slices z=4.0 cm to z=5.5 cm)	65
Figure 3.12 Needle loading and needle distribution in Model A.....	66
Figure 3.13 Needle loading and needle distribution in Model B.....	66
Figure 3.14 Needle loading and needle distribution in the manual treatment plan	67
 Chapter 4: The Enhanced MILP Model Formulation	 71
Figure 4.1 Seed-placement regions, produced through global encapsulation	76
Figure 4.2 Distant placement of seeds from the PTV at the base slice.....	76
Figure 4.3 The bladder appears anteriorly to the base PTV slice	77
Figure 4.4 Seed-placement regions, produced through slice-based local encapsulation.....	79
Figure 4.5 Nearby placement of seeds to the PTV at the base slice	79
Figure 4.6. Pubic arch interference shown for a mid-gland slice	80

Figure 4.7 Removal of seed-placement positions on Column D	83
Figure 4.8 Comparison of a single 0.40 cm urethral protection margin and two separate 0.40 cm and 0.20 cm protection margins.....	83
Figure 4.9 Seed separation of 4.0 cm, 4.5 cm and 5.0 cm in a needle.....	86
Figure 4.10 Legend for treatment plan isodoses and contours	93
Figure 4.11 Comparison of treatment plans (slices $z=0.0$ cm to $z=1.5$ cm)	94
Figure 4.12 Comparison of treatment plans (slices $z=2.0$ cm to $z=3.5$ cm)	95
Figure 4.13 Comparison of treatment plans (slices $z=4.0$ cm to $z=5.5$ cm)	96
Figure 4.14 Needle loading and needle distribution in the base model	97
Figure 4.15 Needle loading and needle distribution in the enhanced model	97
Figure 4.16. Needle loading and needle distribution in the manual treatment plan	98

Chapter 5: Pseudo High-Resolution Data Sets & Feasibility-Based Modelling..... 101

Figure 5.1 Comparison of low-resolution and high-resolution discretization of the urethra.....	103
Figure 5.2 Comparison of high-resolution (100% sampling) and pseudo high-resolution data (75 and 50% sampling) for the PTV, rectum and urethra.....	107
Figure 5.3 Comparison of pseudo high-resolution data (25 and 10% sampling) and low-resolution data.....	108
Figure 5.4 Visualization of the process for creating pseudo high-resolution PTV data	109
Figure 5.5 Sample process flow outlining treatment planning and seed implantation	114
Figure 5.6 Solution processes of the optimality-based model and the constraint-violating feasibility-based model	121
Figure 5.7 Histograms and cumulative distribution functions of solution times.....	141
Figure 5.8 Number of binary variables in relation to prostate volume	142
Figure 5.9 Number of constraints in relation to prostate volume.	143
Figure 5.10 Solution time in relation to prostate volume	144
Figure 5.11 Solution time in relation to number of binary variables.....	145
Figure 5.12 Solution time in relation to number of binary constraints.....	146
Figure 5.13 Ordered solution times of all 200 prostate test cases	147

Chapter 6: Clinical Verification of Automated Plans	153
Figure 6.1 Placement of seeds in off-grid positions as indicated by the red arrows in a sample manual plan.....	160
Figure 6.2 Trends in PTV V100 values of manual and automated plans	163
Figure 6.3 Trends in PTV V150 values of manual and automated plans	163
Figure 6.4 Trends in PTV V200 values of manual and automated plans	164
Figure 6.5 Trends in PTV D90 values of manual and automated plans	164
Figure 6.6 Trends in Urethra D5 values of manual and automated plans.....	165
Figure 6.7 Trends in Rectum D1cc values of manual and automated plans.....	165
Figure 6.8 Trends in asymmetrically positioned seeds in manual and automated plans	166
Figure 6.9 Trends in seeds placed in off-grid positions in manual and automated plans	166
Figure 6.10 Trends in adjacent placed seeds in manual and automated plans.....	167
Figure 6.11 Trends in strands with potential pubic arch interference in manual and automated plans.....	167
Figure 6.12 Placement of seeds to the bladder	173
Figure 6.13 Removal of seed-placement positions from input data	174

List of Tables

Chapter 3: The Basic MILP Model Formulation	21
Table 3.1 Dose-volume requirements of the Cross Cancer Institute	44
Table 3.2 Numerical features of branch-and-bound solution	59
Table 3.3 Spatial characteristics of treatment plans.....	59
Table 3.4 Summary dose-volume metric values for the three treatment	62
 Chapter 4: The Enhanced MILP Model Formulation	 71
Table 4.1 Seed-placement boundaries used in slice-based local encapsulation	90
Table 4.2 Numerical features of branch-and-bound solutions.....	91
Table 4.3 Spatial characteristics of treatment plans.....	91
Table 4.4 Summary dose-volume metric values for the three treatment plans.....	93
 Chapter 5: Pseudo High-Resolution Data Sets & Feasibility-Based Modelling.....	 101
Table 5.1 Model formulations used in our experimental design	125
Table 5.2 Details of the objective functions and input data used in each model formulation ...	125
Table 5.3 Seed-placement margins used for the differently sized and shaped prostates	129
Table 5.4 Clinical parameters set for the differently sized and shaped prostates	129
Table 5.5 Solution performance results obtained for the six model formulations.....	136
Table 5.6 Number of constraints observed in feasibility-based models	137
Table 5.7 Number of binary variables observed in feasibility-based models.....	137
Table 5.8 Number of total variables observed in feasibility-based models	137
Table 5.9 Starting objective function values observed in feasibility-based models	137
Table 5.10 Number of explored nodes observed in feasibility-based models.	138
Table 5.11 Number of simplex iterations observed in feasibility-based models	138
Table 5.12 Presolve times observed in feasibility-based models	138
Table 5.13 Branch & bound solution times observed in feasibility-based models.....	138
Table 5.14 Total solution times (minutes) observed in feasibility-based models.....	139
Table 5.15 Total solution times (seconds) observed in feasibility-based models.....	139

Table 5.16 Performance of three formulations containing pseudo high-resolution data	140
---	-----

Chapter 6: Clinical Verification of Automated Plans 153

Table 6.1 A representation of the score sheet given to radiation oncologists.....	157
---	-----

Table 6.2 Clinical analysis of manual plans, conducted by five radiation oncologists	168
--	-----

Table 6.3 Summary of the clinical analysis conducted on all twenty manual plans.	169
---	-----

Table 6.4 Clinical analysis of automated plans, conducted by five radiation oncologists	170
---	-----

Table 6.5 Summary of the clinical analysis conducted on all twenty automated plans.....	171
---	-----

List of Algorithms

Chapter 3: The Basic MILP Model Formulation	21
Algorithm 3.1 Transforming the PTV region from high-resolution to low-resolution.....	29
Algorithm 3.2 Creating the discretized permissible seed-placement region.	33
 Chapter 4: The Enhanced MILP Model Formulation	 71
Algorithm 4.1. Accounting for pubic arch, urethral curvature and rectal curvature	81
 Chapter 5: Pseudo High-Resolution Data Sets & Feasibility-Based Modelling.....	 101
Algorithm 5.1 Algorithm for creating pseudo high-resolution PTV data.....	106

Chapter 1

Introduction

Despite its fairly small size, the prostate gland plays a major role in the health of North American men. With an estimated 220,800 cases in 2015, carcinoma of the prostate is the leading type of new malignancy involving men in the United States (Siegel, Miller & Jemal 2015). Prostate cancer is also expected to cause 27,540 deaths in 2015, making it the second deadliest form of cancer among males (American Cancer Society 2015). Early diagnosis and treatment of prostate cancer can eliminate disease-related symptoms and increase the prospects of patient survival (Shore 2014). Treatment options available today depend on the stage of the disease, as well as the personal preference, age, and comorbidity of the patient (Heidenreich et al. 2008). Among the various treatment alternatives, which include androgen deprivation therapy, cryotherapy, external beam radiation therapy, and prostatectomy, one particular method stands out to be of high value for early-stage prostate cancer: brachytherapy (Gomez-Veiga et al. 2012). In its early stages, when the disease has not yet penetrated the prostate capsule and spread to surrounding tissue, brachytherapy is considered to be an extremely effective treatment method with a low mortality rate (Zelevsky et al. 2007).

Brachytherapy is a minimally invasive form of radiation therapy that can be classified under three general categories: interstitial, intracavitary, and surface brachytherapy. Interstitial brachytherapy involves the permanent or temporary surgical introduction of radioactive sources, inside a cancerous organ within the human body, such as the prostate or breasts (Lin et al. 2008). Intracavitary brachytherapy involves the temporary exposure to radioactive sources of cancerous tissue lining the inner cavities of the human body, such as the esophagus or cervix (Gorissen, Den Hertog & Hoffmann 2013). Surface brachytherapy involves the application of radioactive sources to body surfaces where cancer is present. All forms of brachytherapy treatment can be delivered using lower strength sources at low dose rate (LDR) or higher strength sources at high dose rate (HDR). Both permanent LDR and temporary HDR brachytherapy are suitable for prostate cancer, although the former is generally the more preferred option among clinicians.

In permanent interstitial LDR brachytherapy, the main objective is to employ patient-personalized treatment plans, or seed (source) distributions, that appropriately target tumors by delivering an effective dose of radiation to cancerous tissue while sparing surrounding healthy tissue. The prostate, a chestnut-sized gland oriented as an inverted pyramid, is securely suspended under the bladder and in front of the rectum, with the urethra tunneling approximately midway through the gland. Thus, a major challenge in treatment planning is the presence of these near-by critical organs-at-risk, for which the dose of radiation must be minimized or constrained. Overdosing otherwise healthy structures with radiation may lead to severe and undesired post-operative medical complications, such as incontinence. Meanwhile, care must be taken not to underdose the target prostate tissue; otherwise the treatment itself may be rendered ineffective and

the patient may require further medical intervention in the future. Considering that LDR prostate brachytherapy is designed as a one-time procedure involving the permanent implantation of seeds within a patient, post-operative remediation of treatment planning or implantation-related errors is technically challenging. Thus, it is of utmost importance, for the success of the brachytherapy treatment, that both the treatment planning and the implantation are conducted with the greatest attainable accuracy.

Treatment planning for brachytherapy can be categorized under two major classes: pre-operative and intra-operative planning. In the pre-operative approach, treatment planning is conducted days or weeks prior to the surgical implantation of seeds in a patient. Intra-operative planning involves treatment planning during the actual surgical procedure itself, while the patient is under anesthesia. Thus, there exists a noteworthy restriction on the length of time treatment planning may take with intra-operative planning, whereas this time pressure to formulate seed distributions is almost non-existent in pre-operative planning. In many cancer centers, including the Cross Cancer Institute in Edmonton, Alberta, treatment planning is carried out pre-operatively through a manual approach (Liu 2014). In this approach, a brachytherapy team, consisting of oncologists and medical physicists, utilize their professional medical expertise to address the previously mentioned factors involving brachytherapy, and through a heuristic process, make decisions regarding the placement of seeds within the prostate. In this heuristic approach, repetitive location changes are made to a starting seed distribution in an attempt to achieve dosimetric improvements, while also paying close attention to the spatial characteristics of the resulting seed

distribution. This iterative improvement process eventually leads to a final seed distribution that is both spatially and dosimetrically sound.

Such a manual planning approach typically takes between 1 and 4 hours, depending on the anatomical complexity of a given patient's prostate as well as the expertise of the brachytherapy team. Despite the length of time associated with it, the manual treatment planning method can result in treatment plans that are of good clinical quality. In other words, these manually produced treatment plans typically satisfy all clinical requirements necessary for subsequent implementation in patients.

1.1 Mathematical Programming in LDR Prostate Brachytherapy

Generally, a more sophisticated approach, based on mathematical programming or operations research, may also be utilized to produce brachytherapy treatment plans. Through the use of mixed integer linear programming (MILP) optimization models, the previously mentioned fine balance between underdosing and overdosing requirements may be achieved with an automated computational approach. Fundamentally, MILP models are devised to solve the problem of optimally locating seeds in order to produce acceptable radiation distributions within a target cancerous tissue or organ. Such an automated approach associates a mathematical guarantee to treatment planning, as this method essentially provides a theoretical verification that any given treatment plan, that is produced through a mathematical optimization approach, numerically offers the most optimal personalized treatment plan for a patient within the solution domain of the programming model. Furthermore, due to the automated nature of the process and its ability to

exploit the computational power of modern day computer systems, such an approach also offers an increase in the speed with which clinicians can obtain access to treatment plans.

As we are interested in mathematical models involving permanent implantation of seeds rather than temporary dwell times, our focus will be on LDR brachytherapy in this thesis, rather than HDR brachytherapy or its variants such as medium dose rate (MDR) and pulse dose rate (PDR) brachytherapy. Thus far, MILPs and other mathematical programming techniques have been employed to a limited extent in order to attempt to solve the LDR brachytherapy problem, with the existing literature largely focusing on two areas: 1) achieving clinically acceptable solutions with greater reliability through refined dose-volume constraints and 2) achieving faster solution times through various algorithms. Perhaps one of the more prominent studies in literature was devised by the combined efforts of Lee et al. (1999), as they were among the first to introduce a basic MILP model specifically designed for interstitial LDR prostate brachytherapy. Another notable addition to LDR brachytherapy optimization literature was the contribution of D'Souza et al. (2001), who developed a sequential solution method through which they were able to simplify their large-scale linear programming models and thereby achieve considerably quick solution times using a branch-and-bound algorithm. Other interesting solution approaches have also been investigated, including the use of genetic algorithms (Yu et al. 1996), simulated annealing (Pouliot et al. 1996), a greedy heuristic based on adjoint functions (Yoo et al. 2003) and a matching pursuit inspired algorithm (Guthier et al. 2015). However, despite the available academic research as well as the commercial optimization tools provided by established biomedical and life-science companies (VariSeed Volume Optimization/Inverse Planning Module, Varian Medical Systems, Palo Alto, CA)

(Oncentra Seeds Planning Module, Elekta AB, Stockholm, Sweden), we observe that the use of optimization techniques is not widely adopted by clinicians in cancer centers for interstitial LDR prostate brachytherapy. We believe that the root cause of this situation is a significant, and overlooked, research gap that exists in today's brachytherapy optimization literature.

1.2 Challenges and Current Research Gap in LDR Brachytherapy Optimization

Perhaps the most valuable quality of optimization models is their capability to model real-world problems and their ability to significantly simplify and speed up any given decision-making process. These characteristics render mathematical optimization an attractive option. However, optimization models are only as intelligent, in making decisions, as the set of instructions they are given. In other words, one needs to be diligent in providing specific and ample guidance to optimization models in order to obtain useful results that are of desirable quality. The utility of a linear programming optimization model is fundamentally dependent on three factors: 1) the input data that is supplied to a specific model, 2) the linear programming model itself, 3) the solution method or algorithm used to solve the model. These factors are not mutually exclusive, and only when considered collectively can they produce powerful optimization solutions for brachytherapy.

Mathematically modeling clinical problems is a challenge. First and foremost, it is difficult to construct realistic linear programming models that incorporate all of the necessary parameters of a real-world clinical procedure such as brachytherapy. Care must be taken in specifically sculpting input data with a consideration for the technical limits of brachytherapy, as this data provides the initial set of information for an optimization model to work with. The objective function and constraints used in a linear programming model have to also be intricately crafted in order to guide

the model towards producing solutions desired by clinicians. Secondly, attempts at capturing realism usually result in extremely large-scale models that are computationally strenuous to solve due to their substantial size. In other words, the solution of these models may take days, weeks or even months. For purposes of practicality and utility, unique decomposition methods or solution algorithms must be implemented in order to attempt to compress the solution of these large-scale models into a reasonable or acceptable amount of time. Thirdly, and most importantly, considering that the first two goals are accomplished, the results produced by such optimization models need to be clinically relevant and applicable. In other words, it is not enough for treatment plans produced by optimization models to be theoretically suitable on paper. These treatment plans not only have to be beneficial on a per-patient basis, but also have to adhere to the specific clinical protocols followed by individual cancer centers.

Although some research centers have devised brachytherapy optimization models and have achieved limited success in attempting to automate the treatment planning procedure (Lee, Zaider 2003), a significant number of cancer centers worldwide still resort to the manual planning approach. The reason behind this is that while there are some commercial optimization software programs available today for clinical use, many medical centers find these to produce treatment plans, or seed distributions, that are unsatisfactory from a clinical perspective. This situation stems in large part from the extensive focus that has been placed on achieving acceptable dosage both in research-based and commercial-based optimization approaches, and the relative lack of focus that has been placed on satisfying other general clinically important requirements. Although it is important to achieve dosimetric values that are satisfactory, this accomplishment alone does not

guarantee that a given treatment plan will be deemed as clinically acceptable. In other words, dose-volume metrics are only half of the story; the quality of treatment plans depends on both quantitative factors, such as dose-volume metrics, and equally importantly, on qualitative parameters. The other half of the story is captured by qualitative parameters, in the form of spatial considerations, that brachytherapy teams place considerable emphasis on during manual planning. These spatial considerations essentially focus on the geometric arrangement of seeds, both in the context of individual seeds as well as the seed distribution as a whole. These constraints constitute a vital portion of the clinical protocols followed by brachytherapy teams everywhere and play a significant part in the quality of treatment plans created for patients. For example, among a set of dosimetrically similar treatment plans, the deciding factor on the clinically most acceptable plan will often come down to the “appropriateness” of the spatial arrangement of seeds and needles. In this context, “appropriateness” is essentially analogous to the ability of a seed distribution to geometrically capture the planning style of expert planners at a cancer center. We believe that one of the primary reasons behind the incapability of past research-based or commercial-based optimization approaches to achieve widespread use is their inability to produce “center-specific” results. This issue truly becomes apparent when considering that regardless of how minor or substantial the variations may be, every cancer center follows their own unique planning philosophy for LDR prostate brachytherapy. A clinic’s planning philosophy is usually developed and refined over years of clinical practice to a point where it is proven to deliver consistently high quality results. Brachytherapy teams grow accustomed to and comfortable with using their native planning style in the treatment of their patients, and therefore are understandably reluctant to deviate from it when

offered the option to utilize a foreign automated planning approach that produces solutions of little resemblance to theirs. It is through our design of spatial constraints that we hope to formulate a mathematical optimization model that possesses the functionality to capture planning styles common in clinics worldwide, and therefore the ability to produce results that brachytherapy teams will find familiar and be comfortable in using for the treatment of their patients.

Prior to our work, a focused and explicit attempt to capture the planning style of expert planners through the design of spatial constraints has not existed in LDR brachytherapy literature. While the use of gEUD (or dose uniformity constraints), as presented in the works of Yoo et al. (2007) (for LDR brachytherapy) and Giantsoudi et al. (2013) (for HDR brachytherapy), as well as the use of geometric optimization, as presented in the study by Raben et al. (2007), have influenced spatial distribution of dose, we see that such approaches do not enable the geometric planning style of a clinic to be explicitly embodied within given constraints, as do the set of spatial constraints that we will be introducing in this thesis. This latter feature is one of the defining characteristics of our study and signifies the added value of our work within the context of existing brachytherapy optimization literature. In the three studies by Karabis et al. (2009), Gorissen et al. (2013) and Holm (2013), constraints are introduced to guide catheter positioning in HDR brachytherapy, however similar constraints do not exist for LDR brachytherapy. While a few clinically common spatial considerations are mentioned in a recent article by McGeachy et al. (2015), the authors have not provided explicit mathematical formulations of the constraints required to achieve such clinical considerations. Fu et al. (2005) have examined spatial relationships of seeds and needles with their unique conical implantation approach, however the scope of their study differs significantly from

our spatial propositions which center around the use of rectilinear templates and parallel needle implantation. Although, in their efforts to produce treatment plans, Lessard et al. (2006) have attempted to mimic an experienced dosimetrist, their line of work also substantially differs from our study as we attempt to mimic an expert's planning style through constraint-based customized solutions rather than class solutions.

Thus, in an effort to fill the current gap that exists in today's literature, this thesis will introduce a novel modeling philosophy for interstitial LDR prostate brachytherapy that incorporates the previously mentioned spatial constraints. In doing so, our premise is that our optimization model will bridge the disconnect that exists between the often disappointing solutions that current optimization approaches offer and the clinically acceptable results that brachytherapy teams desire from such optimization approaches. Simply put, the fundamental goal of this thesis is to lay the foundation for a new MILP modeling approach that will enable automated construction of treatment plans that are of equivalent clinical quality to manually-produced treatment plans created by expert planners.

1.3 Research Objectives and Outline of the Thesis

We will be attempting to overcome numerous clinical and engineering challenges in this thesis. In summary, our main objectives are three-fold:

- a. Formulate a novel mixed-integer linear programming model as well as practical data processing techniques for LDR prostate brachytherapy, in order to achieve a realistic mathematical representation of the treatment planning style that is predominantly utilized at the Cross Cancer Institute.
- b. Develop and employ unique linear programming modelling techniques for simplifying the solution of the resulting large-scale model, in order to be able to produce treatment plans with short solution times that are suitable for both pre-operative and intra-operative planning.
- c. Verify the clinical acceptability of our automated treatment plans based on the clinical standards of the Cross Cancer Institute and develop pathways of further improvement for our automated planning system.

We will attempt to achieve these objectives progressively throughout this thesis. Chapter 2 will introduce the fundamentals of mixed-integer linear programming and discuss solution algorithms of interest. Chapters 3 and 4 will dive into depth on spatial constraints and data processing techniques for capturing treatment planning style at the Cross Cancer Institute. Several modelling techniques for shortening solution time will be detailed and examined in Chapter 5. Finally, we will investigate the real-world clinical utility and implications of our treatment plans in Chapter 6. This thesis will be concluded with suggestion for future avenues of research work.

Chapter 2

Mixed-Integer Linear Programming (MILP) & Deterministic Solution Algorithms

In this section, we will give a brief overview of mixed-integer linear programming (MILP). Furthermore, we will discuss the two deterministic solution algorithms of primary interest in this thesis: simplex and branch-and-bound algorithms.

2.1 Mixed-Integer Linear Programming

As mentioned earlier, mixed-integer linear programming (MILP) is a highly attractive method for modeling real-world decision making processes or systems and achieving the best outcome given a set of requirements. The basic structure of a typical mixed-linear programming model consists of decision variables (continuous, integer or binary), a linear objective function (that is either minimized or maximized), and a set of linear constraints (generally in the form of equalities and inequalities). Mathematically, a mixed-integer linear program may be described as in equations (2.1)-(2.6). Here, the integer decision variables are represented by vector x , the linear objective function by vector c , and the matrix equation $(Ax \geq b)$ designates the linear constraints that variable x is subjected to. Finally, vectors l and u represent the lower and upper limits (or bounds) on variable x .

The main difference between linear programming and mixed-integer linear programming is that where as in the former decision variables are continuous in nature, in the latter some variables are specifically designed to be discrete in form. In other words, decision variables only exist as continuous entities in LP models, while MILP models typically contain integer or binary variables in addition to continuous variables. Such integrality allows us to capture the discrete nature of decisions like the placement or non-placement of a seed (represented by a binary variable that is restricted to a 0 or 1 respectively). Regardless of the model type, decision variables are generally assumed to be non-negative.

$$\text{minimize: } c^T x + d^T y \quad (2.1)$$

$$\text{subject to: } Ax + By \geq b \quad (2.2)$$

$$l \leq y \leq u \quad (2.3)$$

$$x, y \geq 0 \quad (2.4)$$

$$x \in Z \quad (2.5)$$

$$y \in R \quad (2.6)$$




Figure 2.1 Depiction of a graphical solution involving a minimization MILP problem.

Taking an initial look at our LDR brachytherapy problem, the main decision variables may be selected as final seed locations. Furthermore, the objective function may involve the minimization of dose to the organs-at-risk, and the linear constraints may reflect the clinical criteria (spatial and dosimetric in nature) that need to be satisfied in each treatment plan. In addition to these basic components, parameters may be included to uniquely tweak the model in order to produce patient-personalized results. Finally, input data (representing anatomical imaging data) is also provided so that through our model we can make good decisions given available clinical imaging information.

The constraints of a MILP model essentially define the boundaries of a feasible region that houses feasible solutions for the problem at hand. A feasible solution is essentially a set of values for a given set of decision variables satisfying all of the linear constraints of the model. Thus, the feasible region may also be defined as the area within which solutions should be

searched for while solving a MILP model. Furthermore, solutions falling on the boundaries defined by the constraints of the model are also included in this search space. The boundaries of the feasible region are important, as they generally contain the optimal solutions generally sought after in optimization models. Typically, distinct integer solutions exist exactly on (or close to) some of the intersection points found on the boundaries of the feasible region. As seen in Figure 2.1, one of these extreme points houses the single best solution (assuming that the problem has a unique optimum). Such an optimal solution not only satisfies the entire set of constraints in the model, but also results in the lowest attainable value for a minimizing objective function. Of course, the extent of minimization in the objective function value is dependent on various factors, some of which are user-specified limits on permitted solution time, and the desired percentage of optimality that the final integer solution needs to meet. We will discuss these concepts in the next few paragraphs.

2.2 Deterministic Solution Algorithms for Solving MILP Models

In this thesis, our desire is to use a deterministic approach (the simplex method in combination with the branch-and-bound algorithm) to solve our MILP brachytherapy model. The main reason behind this decision is due to the relative unpopularity of the simplex method (and the branch-and-bound algorithm) in comparison to more extensively researched stochastic and heuristic methods in LDR and HDR brachytherapy optimization literature (De Boeck et al. 2014). In the next two sections, we will give a brief overview of the simplex and branch-and-bound algorithms.

2.2.1 Simplex Algorithm

The simplex algorithm was formulated by Dantzig (1951), and to this day it continues to be one of the most effective and widely used methods for solving linear programming problems. The simplex method is based on an iterative process in which the goal is to find better solutions to the linear programming model in each successive step. The simplex algorithm outputs a final optimal solution when it is determined that the current best solution cannot be further improved. The degrees of improvement in successive solutions are determined through their resulting objective function values.

The procedure for the simplex algorithm begins at a vertex (a corner point) of the feasible region. As we have previously mentioned, an intersection holds a single solution out of many feasible ones, and therefore possesses a particular objective function value. As shown in Figure 2.2, the algorithm then attempts to locate and move to another neighboring vertex that may hold a solution with an improved objective function value (a lower value if the objective function is minimized, or a higher value if the objective function is maximized in the model). The search for vertices (containing improved solutions) is conducted along the boundary edges of the feasible region. This sort of a search is made possible through the mathematical property that if at a given vertex the objective function does not possess the maximum possible value, there still exists an edge (stemming out from that vertex) along which the objective function value grows smaller for minimization problems or larger for maximization problems. Thus, the simplex algorithm is guided towards improved solutions in each progressive step, regardless of whether it is finding solutions for a minimization or a maximization problem. Depending on the model size, the

feasible region may be a polygon (when the model has two variables), a polyhedron (when the model has three variables), or an n -dimensional polytope, where the number n of variables could be in the thousands or millions. In more technical terms, this search is carried out through iterative pivoting of the starting simplex tableau, which leads to new candidate solutions entering the basis. Considering that for any given polygon or polyhedron the number of vertices and edges is finite, the number of pivots that need to be executed (in order to introduce new solutions into the basis) is also finite. Thus, the simplex algorithm is guaranteed to eventually find an optimal solution to a linear programming model given enough time.




Figure 2.2 Depiction of a solution process using the simplex algorithm.

2.2.2 Branch-and-Bound Algorithm

The branch-and-bound algorithm was first proposed by Land and Doig (1960). The algorithm can typically be categorized as a form of tree search, and is generally effective for finding optimal solutions to mixed-integer linear programming problems. The reason behind its efficiency is that the branch-and-bound algorithm aims to find successively better solutions in the search space of candidate solutions while concurrently eliminating large sections of the search space based on past and present bounds on the values of variables in a MILP model.

In order to explain the basic operation of the branch-and-bound algorithm, we will now take a look at how the algorithm handles a minimization MILP problem that contains binary variables (which should only be assigned a value of 0 or 1). The branch-and-bound tree consists of nodes and branches, as shown in Figure 2.3.




Figure 2.3 Depiction of a branch-and-bound algorithm search tree.

The initial node that the entire search tree stems out of is called the root node (which we will denote as n_R). A relaxed version of the MILP model is solved in the root node (relaxed meaning all integrality restrictions on binary variables are removed). If the relaxed solution still happens to satisfy all of the original integrality restrictions on the binary variables, then an appropriate solution has been found to the original problem. However, this is usually not the case. If a relaxed solution is obtained in the root node where a binary variable x is not assigned either a 0 or 1, but a value of 0.45 for example, then additional work needs to be done. The algorithm has to branch out from the root node to two new nodes (n_1 and n_2) in which two specifically bound instances of the relaxed MILP problem are solved (a new bound, or restriction, of $x \leq 0$ is applied in node n_1 and a bound of $x \geq 1$ is implemented in node n_2). Through these two newly created sub-problems in nodes n_1 and n_2 , the algorithm must determine whether it would be more appropriate to assign a 0 or a 1 to variable x . As we mentioned before, the branch-and-bound algorithm attempts to eliminate searches that are considered futile. If it is proven in node n_1 that an integer solution of $x=0$ leads to an objective function value of at least Y , while an integer solution of $x=1$ in node n_2 leads to an objective function value of at most Z (with $Z < Y$), then it is clear that it would be ineffective to further investigate instances where x is set to 0. Node n_1 would be “pruned”, meaning that any further search down that node is halted and therefore no additional branching out of node n_1 is necessary (nodes may also be pruned for other reasons, such as when a bound, or restriction, on a variable leads to an infeasible solution). The lower bound on Y and the upper bound on Z directs the branch-and-bound algorithm towards the right path to optimality (successively lower objective function values for integer

solutions) by eliminating large sections of the search space where x would be set to 0. Thus, the search tree is selectively grown out of the most promising nodes. Although we only examined how branching-and-bounding works for a single binary variable, this process is carried out for all binary variables in a MILP model. When all integrality requirements of a problem are satisfied, we have our first feasible integer solution to the original MILP problem. Linear programming solvers, such as Gurobi or CPLEX, utilize the simplex algorithm to find integer and relaxed solutions at each node of the branch-and-bound tree.

The branch-and-bound algorithm proves the optimality of a given solution based on the objective function value of the current best (incumbent) integer solution and the objective function value of the best bound relaxed solution. It should be noted that the incorporation of integrality requirements within a relaxed problem typically leads to higher objective function values for a minimization problem. Therefore, any given relaxed solution acts as a lower bound to an integer solution in the case of a minimization problem. During its tree search, the branch-and-bound algorithm leads to successively better feasible integer solutions (with decreasing objective function values) and successively tighter relaxed solutions (with increasing objective function values). These two objective function values converge until a desired level of optimality gap is reached, resulting in an optimal solution. The optimality gap between the best integer and the best bound solution is calculated as shown in equation (2.7).

$$\text{Optimality Gap \%} = \left(\frac{|\text{Best Bound} - \text{Best Integer}|}{\text{Best Integer}} \right) * 100 \quad (2.7)$$

Chapter 3¹

The Basic MILP Model Formulation

In this chapter, we will focus on defining the characteristics of the mathematical problem at hand, introducing basic data processing concepts for brachytherapy and formulating our base MILP model. Our ideas will then be demonstrated through a comparison of our results to those obtained from a benchmark model as well as a manual planning approach. However, it is important to note that the basic ideas introduced in this chapter will be enough to partially capture, through mathematical means, the ideas that exist inside the mind of an expert planner. Chapter 4 will further develop the concepts proposed in this chapter and will present more complex constraints and data processing techniques that will enable us to more fully capture the planning style of an expert planner. MILP solution performance, such as solution speed or time, will not constitute a major focus of either of these two studies, and will instead be comprehensively analyzed in a separate future investigation. This future study will also propose a shift of perspective away from solution strategies that have already been implemented in

¹*A version of this chapter has been submitted for publication: Babadagli, M.E.; Sloboda, R.; Doucette, J. “An Integer Linear Programming Optimization Model for Capturing Expert Planning Style in Interstitial Low Dose Rate Prostate Brachytherapy”. Springer, Annals of Operations Research (submitted November 21, 2016).*

brachytherapy MILP models, by introducing a novel modeling approach that enables rapid automated construction of treatment plans. The real-world utility of our MILP model (in other words, its ability to consistently produce high quality solutions for prostates of any shape and size) will also not be a major focus of our current study, and will instead be thoroughly examined in a future twenty-patient study.

Our proposed system, outlined in Figure 3.1, is not meant to replace an expert planner (medical physicist) in the treatment planning process. Instead, our eventual aim is to provide computational assistance to experts, in both pre-operative and intra-operative settings, by developing an intuitive automated planning approach that integrates speed and versatility with clinically germane and high quality solutions.

3.1 LDR Prostate Brachytherapy Problem Formulation

In its essence, the brachytherapy problem is based on a complex supply-demand relationship that is in ways similar to those observed in facility-location problems (Aikens, 1985). On the one side, there exists a demand for radiation by prostate tissue, as determined by the radiation dose that is prescribed by the oncologist. On the opposing side, there exists a limitation in the supply of radiation, as determined by the strength and quantity of radioactive seeds that can be placed in the prostate in a clinically viable manner. Such supply-demand relationships can be expressed as grid-based location problems (GBLPs) (Noor-E-Alam et al. 2012). We further build on that work by adding a third dimension to their two-dimensional (2-D) GBLP models and applying fundamental components of these models to LDR prostate brachytherapy. Our formulation of the prostate LDR brachytherapy problem, as presented here,

is specific to the clinical protocol followed at a local cancer centre, but the formulation itself may be generalized or modified to accommodate the requirements of any cancer centre.

3.1.1 Grid and Coordinate System

In our three-dimensional (3-D) GBLPs, anatomical structures of interest, such as the prostate, urethra and rectum, are partitioned into three-dimensional matrices containing voxels of small cells. Essentially, two types of grids need to be considered: 1) i - j - k indexed grids used for modeling the three anatomical structures, and 2) a x - y - z indexed grid used for modeling seed-placement locations. Considering the rapid dose fall-off that occurs with radioactive seeds used in brachytherapy, a fine-resolution grid of $0.1 \times 0.1 \times 0.5 \text{ cm}^3$ resolution is chosen for grids associated with anatomical structures. On the other hand, due to the spacing present in the physical template grids used in the surgical procedure itself, as seen in Figure 3.2, a coarse-resolution grid of $0.5 \times 0.5 \times 0.5 \text{ cm}^3$ resolution is suitable for modeling potential seed locations. The x - y - z (or similarly, the i - j - k) coordinate system that we utilize is oriented as shown in Figure 3.3. In this coordinate system, the x or i direction varies medial-laterally, the y or j direction varies posterior-anteriorly and the z or k direction varies superior-inferiorly relative to the volumetric center of the prostate. 0.5 cm spacing is utilized for the z or k direction in both grids because the prostate data is imaged, via ultrasound, in transaxial slices spaced at 0.5 cm intervals. Thus, relatively speaking, the superior-most slice, where $z=0.0 \text{ cm}$, corresponds to the base of the prostate. For the specific patient prostate we will be examining in this chapter, columns (i direction) and rows (j direction) are numbered from 0-64 with increments of 1 (representing a length of 6.5 cm with a spacing of 0.1 cm) in the fine-resolution anatomical grids.

In the coarse-resolution seed-placement grid, columns (x direction) and rows (y direction) are numbered from 2-62 with increments of 5 (representing a length of 6.0 cm with a spacing of 0.5 cm). For both types of grids, slices in z (or k direction) are numbered from 0-55 with increments of 5 (representing a length of 5.5 cm with a spacing of 0.5 cm). This system is depicted in Figure 3.4 for a sample slice, where a fully populated coarse-resolution grid (represented by red voxels) is superimposed on a fully populated fine-resolution grid (represented by pink voxels). In order to underline the power of the spatial constraints that we are introducing in this chapter, we will use coarse grids of $0.5 \times 0.5 \times 0.5 \text{ cm}^3$ resolution for modeling both the anatomical structures as well as the seed-placement locations. This approach will demonstrate that, through the implementation of spatial constraints, clinically viable results are achievable despite the use of a coarse-resolution grid for dosimetric calculations. In future studies, we will transition to using a high-resolution grid and demonstrate pragmatic methods to overcome issues of intractability that occur with the consequent large-scale MILP model.

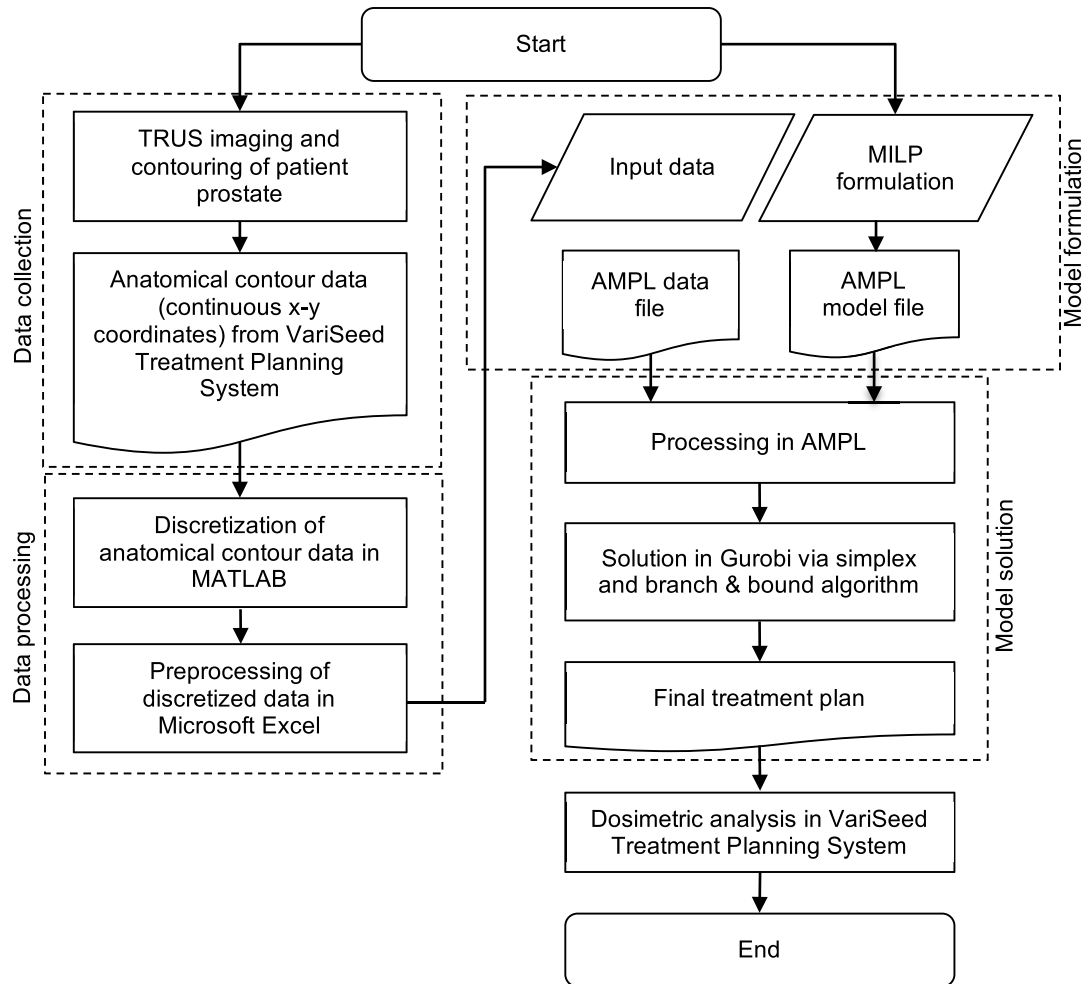


Figure 3.1. Process flow of our automated treatment planning system.




Figure 3.2. A 0.5×0.5 cm² coarse-resolution (x-y plane) physical grid for prostate brachytherapy.




Figure 3.3. Three-dimensional coordinate system, displaying needle insertion points (based on the coarse-resolution grid).




Figure 3.4. Coarse-resolution grid (red) superimposed on the fine-resolution grid (pink).

3.2 Preprocessing of Contour Data

3.2.1 Discretization of Prostate Contour Data

Anatomical contour data is obtained from the VariSeed (version 8.0, Varian Medical Systems Inc., Palo Alto, CA) treatment planning system (TPS) in the form of x - y coordinates of contour points for each imaged anatomical slice. We utilize MATLAB (MathWorks Inc., 2015) to further discretize this continuous contour data. Thus, a 65x65 grid (0.1×0.1 cm² resolution in the x - y plane) is superimposed onto the contour data, and cells (or pixels) that fall within the polygon representing a contour in each slice are given a value of 1, while pixels that fall out of the polygon are given a value of 0. MATLAB's built-in function, `poly2mask()`, is utilized to automate the decision on whether a given voxel is deemed to be within or outside of a polygon region, based on where the mid-point of each cell falls relative to polygon boundaries. Essentially, this function subdivides each cell in the 65x65 grid into a further 5x5 grid, and checks whether the middlemost square in each cell's associated 5x5 grid falls within the boundaries of a given polygon. Before discretization, a 0.3 cm margin is added to the prostate contour in all directions, except the posterior, to construct the planning target volume (PTV) contour. The PTV is required in order to account for prostate contour location uncertainty. The resulting discretized data for a PTV slice, at $z=3.0$ cm, is shown in Figure 3.5 with the white voxels representing voxels that fall within the PTV contour. Image (A) represents the discretized PTV data in coarse-resolution while image (B) represents the discretized PTV data in fine-resolution.

3.2.2 Processing and Conversion of Anatomical Contour Data

For the purposes of this chapter, anatomical data is converted from a grid of $0.1 \times 0.1 \times 0.5$ cm^3 resolution to a grid of $0.5 \times 0.5 \times 0.5$ cm^3 resolution for the PTV, urethra and rectum. The PTV contour is used for calculating prostate dose requirements, while the urethra and rectum contours are used to calculate dose to these respective tissues. The initial discretized PTV data has the corresponding physical urethra regions removed from it on each slice, before being converted from high resolution to low-resolution data. Algorithm 3.1 is utilized to execute the conversion from high resolution to low-resolution data. Algorithm 3.1 is utilized to execute the conversion from fine-resolution to coarse-resolution for the PTV region, $T_{(i,j,k)}$, itself. A procedure similar to Algorithm 3.1 (excluding step 2a) is executed to convert the urethra, $URE_{(i,j,k)}$, and rectum, $REC_{(i,j,k)}$, regions from high-resolution to low-resolution.

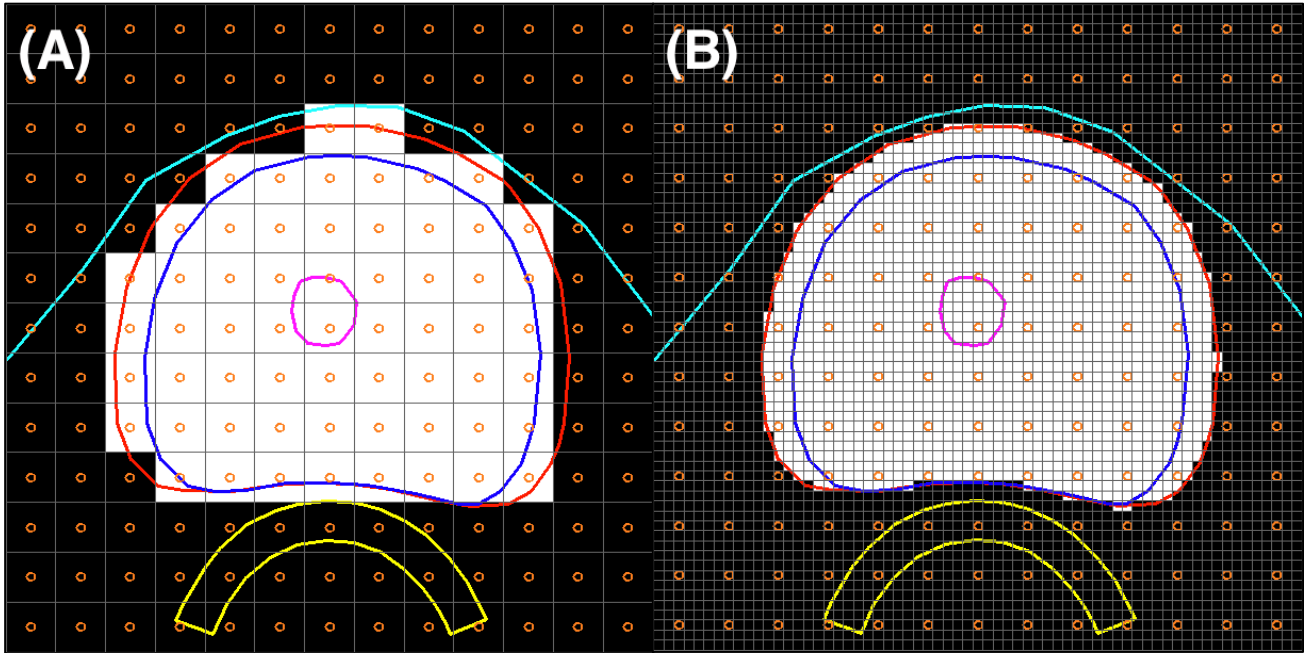


Figure 3.5. Discretized PTV in coarse-resolution (A) and fine-resolution (B) in slice $z=3.0$ cm. The prostate is represented by the dark blue line, the planning target volume (PTV) by the red line, the urethra by the magenta line, the rectum by the yellow line, the pubic arch by the cyan line and the seed-placement grid by the orange circles.

```

1  1. Start with the highest numbered (apical) slices of discretized (fine-resolution) PTV region
2  (RawPTV[k]), urethra region (Urethra[k]), fully populated coarse-resolution grid (CRGrid[k]),
3  and three empty grids (ProPTV[k], NegPTV[k], T[k]):
4  2. k = apical slice coordinate
5  WHILE k ≥ 0:
6      a. FOR each voxel ProPTV[i,j,k] in slice ProPTV[k]:
7          Assign ProPTV[i,j,k] = MAX(0, (RawPTV[i,j,k] - Urethra[i,j,k]))
8      b. FOR each voxel NegPTV[i,j,k] in slice NegPTV[k]:
9          IF ProPTV[i,j,k] + CRGrid[i,j,k] > 1:
10             Assign NegPTV[i,j,k] = 0
11          ELSE:
12             Assign NegPTV[i,j,k] = ProPTV[i,j,k]
13      c. FOR each voxel T[i,j,k] in slice T[k]:
14          Assign T[i,j,k] = MAX(0, (ProPTV[i,j,k] - NegPTV[i,j,k]))
15      d. Perform k=k-1.
16  3. End with (coarse-resolution) PTV region indexed as T[i,j,k]

```

Algorithm 3.1. Algorithm for transforming the discretized PTV region from high-resolution to low-resolution.

3.2.3 Permissible Seed-Placement Region

Defining the permissible seed-placement region involves a number of considerations. Placement of seeds outside of the PTV is not uncommon in clinical practice, and helps to achieve sufficient dose coverage within the PTV, especially at the base and the apical slices. One simple approach to achieve this, as is briefly described in the study by McGeachy et al. (2015), is to select the largest PTV contour (usually a mid-gland contour) and superimpose this contour as the permissible seed-placement region for all slices. We call this the “cylinder approach” or “global encapsulation”, as the seed-placement boundary is uniform throughout all of the slices and essentially encapsulates the entire prostate. In order to assess how this approach fares, we will apply it to our data with a slight modification. Our modification involves taking the largest two mid-gland PTV contours, slices $z=2.5$ cm and $z=3.0$ cm (out of a set of $z=0.0-5.5$ cm), and using these to act as the permissible seed-placement regions for the superior ($z=0.0-2.5$ cm) and

inferior ($z=3.0-5.5$ cm) halves of the prostate respectively. The resulting data will act as our initial seed-placement region.

Considering shape differences among slices, the PTV contour of a single mid-gland slice may not always sufficiently encapsulate all other neighboring mid-gland slices, and therefore may leave out a number of seed-placement positions that may prove to be necessary. On the other hand, the PTV contour of a single mid-gland slice may also prove to be excessively large when used as a seed-placement region for relatively small slices, such as the base or apex. This situation may potentially lead to the placement of seeds that are significantly far from the PTV contour at the base and the apex. To overcome these issues, we will introduce “local encapsulation” in future studies, which results in a tighter, or more balanced fit, around the PTV.

3.2.4 Seed-Placement Limitations Due to the Urethral Curvature and Rectum

Generally, the urethra follows a curved path through the prostate, entering at the anterior base, dipping posteriorly, and rising anteriorly towards the apex. Due to this dip, oncologists must be mindful not to penetrate or cause other forms of trauma to the urethra with needles during seed implantation. Therefore, a specific region of the prostate, principally delineated by the entrance and exit points as well as the curvature and the diameter of the urethra, is off limits to needles. Thus, seeds cannot practically be placed in this region. We implement an additional protection zone around the urethral curvature, to prevent any seeds from being placed within 0.4 cm of this region. This is a technique common in clinical practice that effectively decreases the likelihood of the urethra being penetrated with a needle. In order to account for these restrictions in the input data that we provide to our linear programming models, the initial seed-placement

region is processed so that the shadow or superimposition of the urethral contours (with the 0.4 cm protection zone) from each slice (starting with the inferior-most slice) is removed or subtracted from seed-placement regions that fall superior to it, before being converted from high resolution to low-resolution data. We also extend this idea to the less pronounced, yet still present, rectal curvature as well (which also includes a 0.4 cm protection zone). Algorithm 3.2 is utilized for creating the final seed-placement region that will be used in our model. Generally, the shadow of the urethral curvature has a more pronounced effect on the removal of potential seed positions in slices that are closer to the base of the prostate, as observed in a sample slice of PTV ($z=0.5$ cm) in Figure 3.6. In image (A) only the corresponding physical urethral region is accounted for in this particular slice, while in image (B) the shadow of the urethral curvature, obtained from inferior slices, is also taken into consideration in addition to the corresponding physical urethral region. The PTV region, urethra region, rectum and the seed-placement grid are represented by pink, yellow, light blue and dark red voxels respectively in both images.

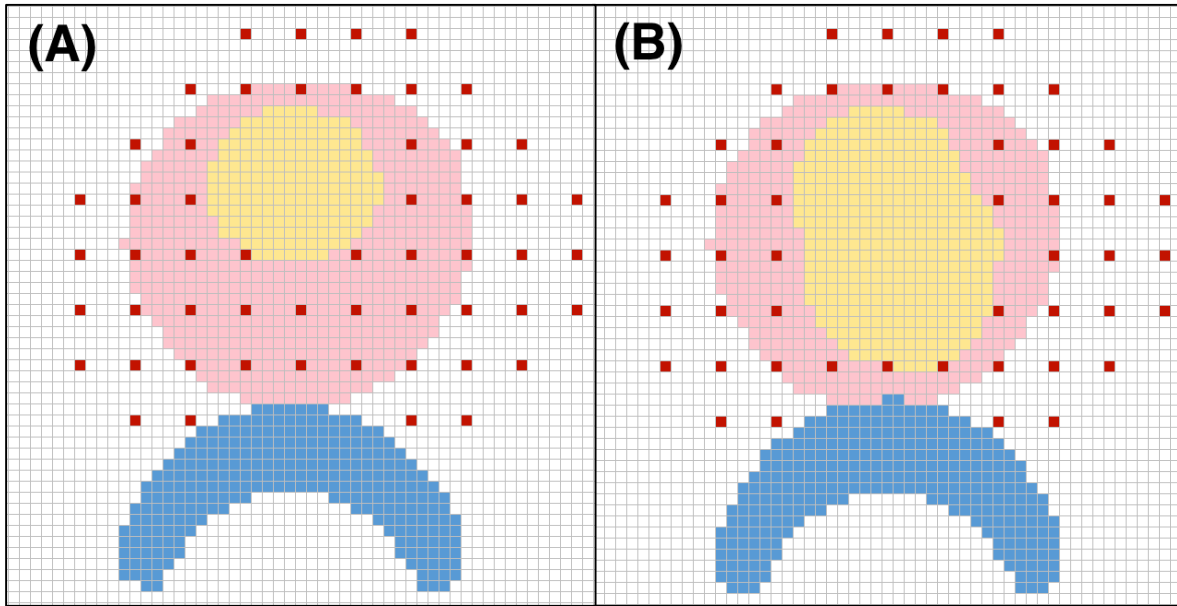


Figure 3.6. Visualization of the physical urethral region (yellow voxels) and the physical rectal region (light blue voxels) (A), both with a 0.4 cm protection zone, as well as the shadow of the urethral curvature and the shadow of the rectal curvature (B) along with the permissible seed grid (dark red voxels) and PTV (pink voxels) in slice $z=0.5$ cm.

```

1  1. Start with the highest numbered (apical) slices of discretized (fine-resolution) initial seed-
2  placement region (InitSP[k]), rectum region with 0.4 cm margin (RECMargin[k]), urethra region
3  with 0.4 cm margin (UREMargin[k]), fully populated coarse-resolution grid (CRGrid[k]), and
4  four empty fine-resolution grids (ProSP[k], ProCurveSP[k], NegSP[k], SP[k])
5  2. k = apical slice coordinate
6  m=k
7  WHILE k ≥ 0:
8      a. FOR each voxel ProSP[i,j,k] in slice ProSP[k]:
9          Assign ProSP[i,j,k] = MAX(0, (InitSP[i,j,k] - SUM(RECMargin[i,j,m], ... ,
10             RECMargin[i,j,k])))
11      b. Perform k=k-1.
12  3. k = apical slice coordinate
13  m = k
14  WHILE k ≥ 0:
15      b. FOR each voxel ProCurveSP[i,j,k] in slice ProCurveSP[k]:
16          Assign ProCurveSP[i,j,k] = MAX(0, (ProSP[i,j,k] - SUM(UREMargin[i,j,m], ... ,
17             UREMargin[i,j,k])))
18      c. Perform k=k-1.
19  4. k = apical slice coordinate
20  WHILE k ≥ 0:
21      a. FOR each voxel NegSP[i,j,k] in slice NegSP[k]:
22          IF ProCurveSP[i,j,k] + CRGrid[i,j,k] > 1:
23              Assign NegSP[i,j,k] = 0
24          ELSE:
25              Assign NegSP[i,j,k] = ProCurveSP[i,j,k]
26      c. FOR each voxel SP[i,j,k] in slice SP[k]:
27          Assign SP[i,j,k] = MAX(0, (ProCurveSP[i,j,k] - NegSP[i,j,k]))
28      d. Perform k=k-1.
29  5. End with finalized (coarse-resolution) permissible seed-placement region indexed as SP[x,y,z]
30

```

Algorithm 3.2. Algorithm to account for the shadow of urethral curvature and rectal curvature while creating the discretized permissible seed-placement region.

3.3 LDR Brachytherapy Optimization Model

3.3.1 General Overview

Our model consists of four main components: an objective function, dose-calculation constraints, dose-volume constraints and spatial constraints. The objective function is used to guide the model towards a specific goal, in our case minimizing the average dose to the urethra and the rectum. The dose-calculation constraints are used to calculate dose delivered to each tissue type. Dose-volume constraints are used to keep track of dose delivered to each voxel and ensure clinical dose limits are satisfied. Finally, spatial constraints are used to ensure that the final geometric distribution of seeds and needles matches the style of planning present at the local cancer clinic. Each of these components will be discussed in detail in the next few sections. Our model consists primarily of binary variables, with the exception of the variables used to calculate total dose, which are continuous.

3.3.2 Sets, Parameters, Variables

The sets used to formulate our model are given as follows:

U is the set of all x-coordinates $\{0,1,2...64\}$ in the anatomical grid, indexed by i

V is the set of all y-coordinates $\{0,1,2...64\}$ in the anatomical grid, indexed by j

W is the set of all z-coordinates $\{0,5,10...55\}$ in the anatomical grid, indexed by k

Q is the set of all x-coordinates $\{2,7,12...62\}$ in the seed-placement grid, indexed by x

R is the set of all y-coordinates $\{2,7,12...62\}$ in the seed-placement grid, indexed by y

S is the set of all z-coordinates $\{0,5,10...55\}$ in the seed-placement grid, indexed by z

The parameters used to formulate our model are given as follows:

$T_{(i,j,k)}$ identifies a voxel in the PTV (target) region

$URE_{(i,j,k)}$ identifies a voxel in the urethra (organ-at-risk) region

$REC_{(i,j,k)}$ identifies a voxel in the rectum (organ-at-risk) region

$SP_{(i,j,k)}$ identifies a voxel in the permissible seed-placement region

α is the preference coefficient for the average dose to the urethra (set to 1)

β is the preference coefficient for the average dose to the rectum (set to 1)

$ptvDoseV100$ is the dose limit for PTV V100 (set to 145 Gy)

$ptvDoseV150$ is the dose limit for PTV V150 (set to 217.5 Gy)

$ptvDoseV200$ is the dose limit for PTV V200 (set to 290 Gy)

$ptvMaxDoseD90$ is the upper dose limit for PTV D90 (set to 189 Gy)

$ptvMinDoseD90$ is the lower dose limit for PTV D90 (set to 181.5 Gy)

$urethraDoseD90$ is the dose limit for Urethra D5 (set to 215 Gy)

$rectumDoseD1cc$ is the dose limit for Rectum D1cc (set to 145 Gy)

$ptvVolumeV100$ is the percentage volume limit for PTV V100 (set to 0.98)

$ptvVolumeV150$ is the percentage volume limit for PTV V150 (set to 0.65)

$ptvVolumeV200$ is the percentage volume limit for PTV V200 (set to 0.25)

$ptvVolumeD90$ is the percentage volume limit for PTV D90 (set to 0.90)

$urethraVolumeD5$ is the percentage volume limit for Urethra D5 (set to 0.05)

$rectumMaxVoxelD1cc$ is the voxel count limit for Rectum D1cc (set to 8)

$upperLimitSeeds$ is the maximum permissible number of seeds (set to 125)

$lowerLimitSeeds$ is the minimum permissible number of seeds (set to 75)

$upperLimitNeedles$ is the maximum permissible number of needles (set to 25)

$lowerLimitNeedles$ is the minimum permissible number of needles (set to 15)

$maxNeedleRetraction$ is the maximum permissible needle retraction (set to 15)

$minSeedsPerNeedle$ is the minimum permissible number of seeds per needle (set to 2)

$maxSeedsPerNeedle$ is the maximum permissible number of seeds per needle (set to 10)

$maxConsecSeeds$ is the maximum permissible number of consecutive seeds in a needle (set to 4)

r is the radius, or Euclidean distance, from a voxel (indexed by i, j, k) to a seed (indexed by x, y, z)

S_K is the air kerma strength (set to 0.502 U)
 Λ is the dose rate constant (set to $0.965 \text{ cGy} \cdot \text{hr}^{-1} \cdot \text{U}^{-1}$)
 $T_{1/2}$ is the half-life (set to 59.4 days)
 $g_p(r)$ is the point-source radial dose function
 $g_L(r)$ is the line-source radial dose function
 $\phi_{an}(r)$ is the anisotropy factor
 M is a large constant (set to 50,000)

The variables used to formulate our model are given as follows:

$DoseP_{(i,j,k)}$ is the total dose delivered to a PTV voxel (continuous)
 $DoseU_{(i,j,k)}$ is the total dose delivered to a urethra voxel (continuous)
 $DoseR_{(i,j,k)}$ is the total dose delivered to a rectum voxel (continuous)
 $PTVLessV100_{(i,j,k)}$ checks if a PTV voxel receives a dose less than $ptvDoseV100$ (binary)
 $PTVExceedV150_{(i,j,k)}$ checks if a PTV voxel receives a dose greater than $ptvDoseV150$ (binary)
 $PTVExceedV200_{(i,j,k)}$ checks if a PTV voxel receives a dose greater than $ptvDoseV200$ (binary)
 $PTVExceedD90_{(i,j,k)}$ checks if a PTV voxel receives a dose greater than $ptvMaxDoseD90$ (binary)
 $PTVLessD90_{(i,j,k)}$ checks if a PTV voxel receives a dose less than $ptvMinDoseD90$ (binary)
 $UrethraExceedD5_{(i,j,k)}$ checks if a urethra voxel receives a dose greater than $urethraDoseD5$ (binary)
 $RectumExceedD1cc_{(i,j,k)}$ checks if a rectum voxel receives a dose greater than $rectumDoseD1cc$ (binary)
 $Seed_{(x,y,z)}$ represents the use of a seed (binary)
 $Needle_{(x,y)}$ represents the use of a needle (binary)

3.3.3 Objective Function

The decision of selecting which variables to include in the objective function is essentially a clinical one, and can be coordinated in a manner that accommodates local clinical practice. As such, most of the past approaches have been formulated as multi-objective problems. For the purposes of utility and simplicity, our objective function, shown in equation

(3.1), has been formulated to include the average dose delivered to urethra voxels, based on the work of Meyer et al. (2003). We further extend this approach to the rectum voxels as well.

$$\min : \alpha * \left(\frac{\sum_{i \in U} \sum_{j \in V'} \sum_{\substack{k \in W \\ URE_{(i,j,k)}=1}} DoseU_{(i,j,k)}}{\sum_{i \in U} \sum_{j \in V'} \sum_{\substack{k \in W \\ URE_{(i,j,k)}=1}} URE_{(i,j,k)}} \right) + \beta * \left(\frac{\sum_{i \in U} \sum_{j \in V'} \sum_{\substack{k \in W \\ REC_{(i,j,k)}=1}} DoseR_{(i,j,k)}}{\sum_{i \in U} \sum_{j \in V'} \sum_{\substack{k \in W \\ REC_{(i,j,k)}=1}} REC_{(i,j,k)}} \right) \quad (3.1)$$

The specific weight to assign to each set of variables in the objective function is a highly subjective decision to make, and in operations research it is generally advisable to determine these values either through trial-and-error or an automated fine-tuning process. However, a posteriori adjustments of weights may prove to be exceptionally arduous and time-consuming, even for the case of a single patient (considering that the entire model has to be solved multiple times in order to reveal an appropriate set of weights that results in a desirable outcome for that patient). Furthermore, the performance of any given set of weights is highly specific to a particular patient's prostate data and is usually not applicable to prostate data from other patients without further modification or adjustment. These realities render the entire fine-tuning process impractical in a high-paced clinical setting, especially during intra-operative planning.

Due to the manner in which average dose is calculated (i.e. total dose divided by the total voxel count), our choice of objective function variables allows us to bypass any balancing coefficients required to level the two variables since their values are already normalized. Coefficients α and β are implemented simply to act as preference coefficients, in order to offer the end-user a choice to place greater emphasis on minimizing one of these two variables.

Similar to the choice of variables to include in the objective function, the choice of values to use for the preference coefficients essentially rests on a clinical decision, and can be set to accommodate local clinical practice. Considering that our model will be including dose-volume constraints to set strict dose limits (which have been tuned for each tissue type following years of clinical practice), it becomes redundant to further set weights in the objective function in order to fine-tune the task of minimizing average dose to the organs-at-risk. For the sake of practicality, it was deemed suitable to simplify this decision making process by setting the preference coefficients as $\alpha = 1, \beta = 1$ (hence, the relative importance of minimizing average dose to both the urethra and the rectum is considered to be equal for our local cancer center).

3.3.4 1-D Point Source Dose Calculation

There are several options, provided by the American Association of Physicists in Medicine (Rivard et al. 2004), for calculating dose distributions for Amersham model 6711 Iodine-125 seeds that we assumed in our model: 1) the 1D formalism, based on a point source approximation, and 2) the 2D formalism, based on a cylindrically symmetric line source. In our MILP models, we use the general 1D isotropic point-source approximation for dose distribution calculation, as outlined in equation (3.2).

Even though the use of the 2D formalism is preferred for stranded seeds, the 1D formalism is still clinically acceptable as the differences between the two methods are relatively minor (Yu et al. 1999). In order to concur with local clinical practice, which uses stranded seeds as opposed to loose seeds during implantation, we implement the 1D formalism for the

automated construction of treatment plans, and the 2D formalism for the post-optimization dosimetric analysis.

Although the 1D formalism only approximates the more complex and fine-tuned dose distribution offered by the 2D formalism (Rivard et al. 2004), there are major advantages to using the 1D point source formalism in a MILP model. One advantage is that the dose calculated through the 1D formalism essentially varies based on a single parameter, distance r , through a predominantly inverse square relationship. Thus, the 1D formalism provides simplicity and practicality, as it can be readily utilized in our MILP models without the need of further linearization when used as a constraint. On the other hand, the 2D formalism requires the determination of the source longitudinal axis orientation, which introduces additional parameters, such as polar angle, into the dose calculation (Rivard et al. 2007). The addition of these parameters inevitably requires linearization of the 2D dose calculation when used as a constraint, and such linearization can add considerable strain and complexity to the solution process of a MILP model.

In comparison to the 2D formalism, the 1D formalism outlined in equation (3.2) is generally known to be susceptible to dose calculation inaccuracies for distances less than $r_{min} = 0.5$ cm, which is the minimum distance for which data is listed for the anisotropy factor, $\phi_{an}(r)$. To account for potential inaccuracies, equation (3.3), which incorporates aspects of the line-source approximation, is utilized specifically for distances less than r_{min} , as suggested by Rivard et al. (2004) in the Update of AAPM Task Group No. 43 (TG-43) Report.

$$\dot{D}(r) = S_K * \Lambda * \left(\frac{r_o}{r} \right)^2 * g_p(r) * \phi_{an}(r) \quad r \geq r_{min} \quad (3.2)$$

$$\dot{D}(r) = S_K * \Lambda * \left(\frac{r_o}{r} \right)^2 * \left(\frac{g_L(r)}{g_L(r_{min})} \right) * g_p(r_{min}) * \phi_{an}(r_{min}) \quad r < r_{min} \quad (3.3)$$

Equations (3.2) and (3.3) represent initial dose rate $\dot{D}(r)$ with units of $\text{cGy} \cdot \text{hr}^{-1}$ as calculated for distances ranging from 0.10 cm to 15.00 cm. In these equations, S_K is the air kerma strength of a seed source, set to 0.502 U. The dose rate constant Λ is set to $0.965 \text{ cGy} \cdot \text{hr}^{-1} \cdot \text{U}^{-1}$, as recommended by the AAPM Update. Consensus data values, provided by the Update of the TG-43 report, are utilized for the point-source and line-source radial dose functions, $g_p(r)$ and $g_L(r)$ respectively. These data values have been interpolated from 0.10 cm to 15.00 cm, with increments of 0.01 cm, using shape-preserving piecewise cubic Hermite interpolation (PCHIP) in MATLAB. Similarly, point-source anisotropy function data values have been interpolated from 0.50 cm to 15.00 cm, using the same incremental values and the same interpolation method. The total dose over time can be calculated by equation (3.4), where $T_{1/2}$ represents the half-life of 59.4 days for Iodine-125.

$$D(r) = \int_0^{\infty} \dot{D}(r) e^{\left(\frac{-0.693t}{T_{1/2}} \right)} dt = \dot{D}(0) \left(\frac{T_{1/2}}{\ln(2)} \right) \quad (3.4)$$

Using the above equation, we can calculate the dose contribution from any radioactive seed to any surrounding voxel. However, in deriving the total dose received by a voxel, we have to be mindful of several cases. Since we are using a different equation for dose rates calculated at

distances of $r < 0.5$ cm and $r \geq 0.5$ cm, for any given voxel, we need to firstly compute the dose contribution from seeds that fall within a radius of $r < 0.5$ cm. Secondly, we need to calculate the dose contribution from seeds that are at a distance of $r \geq 0.5$ cm. Lastly, we need to be attentive of a case in which we may be required to calculate the dose received by a seed-containing voxel from the seed that the voxel itself holds. Theoretically, due to the inverse square relationship present in the dose rate equation, one would presume that the dose contribution at a distance of $r = 0$ cm away from the seed would be infinite. However, dividing by 0, or assigning an infinite dose, can lead to potential complications in linear programming models. One could bypass such a situation by assigning an arbitrarily high value (above 290 Gy) to the dose contribution at $r = 0$ cm, which may provide a practical workaround considering that the highest dose to be taken into consideration by our models is twice the prescription dose (for the dose-volume metric PTV V200), which is 290 Gy. However, to be more consistent, we assign the dose contribution calculated at the shortest possible distance away from a seed in the $0.1 \times 0.1 \times 0.5$ cm³ resolution grid, which is located at $r = 0.10$ cm. The radiation dose value that corresponds to for this specific distance is ~1001 Gy, but for the sake of simplicity, we round that number to an easily identifiable value of ~1000 Gy. Thus, the dose calculation approach that we use for the PTV in our MILP model consists of summing doses from the previously mentioned three scenarios involving different seed distances from a voxel, as in equation (3.5). Equations (3.6) and (3.7), used for calculating dose to the urethra and the rectum respectively, do not contain dose contribution at $r = 0$ cm, as seeds are not permitted to be placed within the urethra or the rectum.

$$DoseP_{(i,j,k)} = \sum_{x \in Q} \sum_{y \in R} \sum_{z \in S} D(r)_{r < r_{min}} * Seed_{(x,y,z)} + \sum_{x \in Q} \sum_{y \in R} \sum_{z \in S} D(r)_{r \geq r_{min}} * Seed_{(x,y,z)}, \forall i \in U, j \in V, k \in W \mid \quad (3.5)$$

$$+ \sum_{x \in Q} \sum_{y \in R} \sum_{z \in S} D(r)_{r=0} * Seed_{(x,y,z)} \quad T_{(i,j,k)} = 1$$

$$DoseU_{(i,j,k)} = \sum_{x \in Q} \sum_{y \in R} \sum_{z \in S} D(r)_{r < 5} * Seed_{(x,y,z)} + \sum_{x \in Q} \sum_{y \in R} \sum_{z \in S} D(r)_{r \geq 5} * Seed_{(x,y,z)}, \forall i \in U, j \in V, k \in W \mid \quad (3.6)$$

$$URE_{(i,j,k)} = 1$$

$$DoseR_{(i,j,k)} = \sum_{x \in Q} \sum_{y \in R} \sum_{z \in S} D(r)_{r < 5} * Seed_{(x,y,z)} + \sum_{x \in Q} \sum_{y \in R} \sum_{z \in S} D(r)_{r \geq 5} * Seed_{(x,y,z)}, \forall i \in U, j \in V, k \in W \mid \quad (3.7)$$

$$REC_{(i,j,k)} = 1$$

Despite the simplicity of the 1D formalism, the total dose equation given above produces acceptably accurate calculations of the dose contribution for all distances r for the purpose of seed distribution optimization (Rivard et al. 2004).

3.3.5 Dose-Volume Constraints

Clinicians evaluate the dosimetric quality of treatment plans based on a collection of specific dose-volume limits set for each tissue type. Several methods have been proposed, in current optimization literature, to achieve satisfactory radiation dosage in treatment plans produced through automated planning. In our view, these methods are essentially divided into two categories: 1) “crude dose-management” methods that are not designed to explicitly incorporate, within the MILP model, the dose-volume metric criteria used by clinicians; 2) “fine dose-management” methods that are designed to explicitly incorporate such dose-volume metric criteria within the model. Implementations of each of these two methods are found in the study by D'Souza et al. (2001).

“Crude dose-management” methods involve underdosing or overdosing constraints (containing continuous dose-related variables), through which the goal is to minimize

underdosing of the target organ or overdosing of organs-at-risk in the objective function. Many of the past studies have adopted this method, perhaps due to the fact that such constraints are simple to design, as well as easier to implement since the use of continuous variables places little to no burden on the solution of a MILP model. The downside of this approach is that treatment plans are essentially produced using overly simplistic dose-management rules that are fundamentally different from the dose-volume metric criteria used by expert planners. Thus, with this approach, it is hoped that the dosimetric values of resulting treatment plans generally align with, or at least do not significantly deviate from, the dose-volume metric criteria used by clinicians during post-optimization analysis.

Only a handful of past studies have adopted “fine dose-management” methods. Such methods involve the design of more complex constraints, called dose-volume constraints, which contain binary dose-related variables. The benefit of this approach is that it allows direct implementation of the exact dosimetric goals of expert planners within the MILP model, hence allowing precise control over dose limits unlike the “crude dose-management” method. Although not perfect, treatment plans produced through the “fine dose-management” method are more likely to align with the dose-volume metric criteria used by clinicians during post-optimization analysis. Perhaps the reason for this method’s relative unpopularity in past studies is due to the inclusion of binary variables in dose-volume constraints, which add significant burden on the solution of MILP models and lead to intractability in large-scale models. As such, even studies that have incorporated dose-volume constraints, such as D’Souza et al. (2001), have only done so for relatively small structures, like the urethra, perhaps in an attempt to limit the number of

associated binary variables that would need to be included in the model. Since it enables our solutions to achieve increased clinical relevance, in this chapter we implement dose-volume constraints for all tissue types (PTV, urethra, rectum).

The structure of dose-volume constraints has been introduced for the “V_{xxx}” family (V100, V150, etc...) previously. We modify and apply these constraints to the “D_{xxx}” family (D90, D5, etc...) of dose-volume metrics as well. The specific dose limits set by the Cross Cancer Institute for each tissue type are given in Table 3.1. We should note that some differences are expected between the dose limits set in our optimization models, and the final values observed during post-optimization dosimetric analysis using VariSeed. This is due in minor part to the 1D formalism and in major part to the rough delineation of anatomical structures achievable with coarse-resolution grids (VariSeed uses the 2D formalism and high-resolution grids). Such differences will be negligible once we transition to high-resolution grids in future studies.

Table 3.1. Dose-volume requirements of the Cross Cancer Institute.

Dose Parameter	Objective
PTV Prescribed Dose	145 Gy
PTV V100	> 98%
PTV V150	≤ 65%
PTV V200	≤ 25%
PTV D90	174-189 Gy
Urethra D5	< 215 Gy
Rectum D1cc	< 145 Gy

V100 is the percentage of the PTV volume that receives at least the prescribed dose of 145 Gy. While V100 acts as a lower bound, V200 (percentage of the PTV volume that receives at least 200% of the prescribed dose) behaves as an upper bound on the allowable dose to the PTV. The V150 constraint is utilized as a method of establishing consistency among treatment plans created for different patients. D90 is the dose, in our constraints set between a range of 181.5-189 Gy, which needs to be delivered to 90% of the PTV volume. Urethra D5 and rectum D1cc constraints follow a similar train of thought to PTV D90. Constraints used to accommodate the dose-volume requirements are given in equations (3.8) to (3.21).

PTV V100:

$$ptvDoseV150 - DoseP_{(i,j,k)} \leq M * PTVExceedV100_{(i,j,k)}, \forall i \in U, j \in V, k \in W \mid T_{(i,j,k)} = 1 \quad (3.8)$$

$$\sum_{i \in U} \sum_{j \in V} \sum_{\substack{k \in W \\ T_{(i,j,k)}=1}} PTVExceedV100_{(i,j,k)} \leq (1 - ptvVolumeV100) * \sum_{i \in U} \sum_{j \in V} \sum_{\substack{k \in W \\ T_{(i,j,k)}=1}} T_{(i,j,k)} \quad (3.9)$$

PTV V150:

$$DoseP_{(i,j,k)} - ptvDoseV150 \leq M * PTVExceedV150_{(i,j,k)}, \forall i \in U, j \in V, k \in W \mid T_{(i,j,k)} = 1 \quad (3.10)$$

$$\sum_{i \in U} \sum_{j \in V} \sum_{\substack{k \in W \\ T_{(i,j,k)}=1}} PTVExceedV150_{(i,j,k)} \leq ptvVolumeV150 * \sum_{i \in U} \sum_{j \in V} \sum_{\substack{k \in W \\ T_{(i,j,k)}=1}} T_{(i,j,k)} \quad (3.11)$$

PTV V200:

$$DoseP_{(i,j,k)} - ptvDoseV200 \leq M * PTVExceedV200_{(i,j,k)}, \forall i \in U, j \in V, k \in W \mid T_{(i,j,k)} = 1 \quad (3.12)$$

$$\sum_{i \in U} \sum_{j \in V} \sum_{\substack{k \in W \\ T_{(i,j,k)}=1}} PTVExceedV200_{(i,j,k)} \leq ptvVolumeV200 * \sum_{i \in U} \sum_{j \in V} \sum_{\substack{k \in W \\ T_{(i,j,k)}=1}} T_{(i,j,k)} \quad (3.13)$$

PTV D90:

$$ptvMinDoseD90 - DoseP_{(i,j,k)} \leq M * PTVLessD90_{(i,j,k)}, \forall i \in U, j \in V, k \in W \mid T_{(i,j,k)} = 1 \quad (3.14)$$

$$\sum_{i \in U} \sum_{j \in V} \sum_{\substack{k \in W \\ T_{(i,j,k)}=1}} PTVLessD90_{(i,j,k)} \leq (1 - ptvVolumeD90) * \sum_{i \in U} \sum_{j \in V} \sum_{\substack{k \in W \\ T_{(i,j,k)}=1}} T_{(i,j,k)} \quad (3.15)$$

$$DoseP_{(i,j,k)} - ptvMaxDoseD90 \leq M * PTVExceedD90_{(i,j,k)}, \forall i \in U, j \in V, k \in W \mid T_{(i,j,k)} = 1 \quad (3.16)$$

$$\sum_{i \in U} \sum_{j \in V} \sum_{\substack{k \in W \\ T_{(i,j,k)}=1}} PTVExceedD90_{(i,j,k)} \leq ptvVolumeD90 * \sum_{i \in U} \sum_{j \in V} \sum_{\substack{k \in W \\ T_{(i,j,k)}=1}} T_{(i,j,k)} \quad (3.17)$$

Urethra D5:

$$DoseU_{(i,j,k)} - urethraDoseD5 \leq M * UrethraExceedD5_{(i,j,k)}, \forall i \in U, j \in V, k \in W \mid URE_{(i,j,k)} = 1 \quad (3.18)$$

$$\sum_{i \in U} \sum_{j \in V} \sum_{\substack{k \in W \\ URE_{(i,j,k)}=1}} UrethraExceedD5_{(i,j,k)} \leq urethraVolumeD5 * \sum_{i \in U} \sum_{j \in V} \sum_{\substack{k \in W \\ URE_{(i,j,k)}=1}} URE_{(i,j,k)} \quad (3.19)$$

Rectum D1cc:

$$DoseR_{(i,j,k)} - rectumDoseD1cc \leq M * RectumExceedD1cc_{(i,j,k)}, \forall i \in U, j \in V, k \in W \mid REC_{(i,j,k)} = 1 \quad (3.20)$$

$$\sum_{i \in U} \sum_{j \in V} \sum_{\substack{k \in W \\ REC_{(i,j,k)}=1}} RectumExceedD1cc_{(i,j,k)} \leq rectumMaxVoxelD1cc \quad (3.21)$$

3.3.6 Spatial Constraints

3.3.6.1 Seed/Position Coupling

Equation (3.22) enables us to ensure that any seed, used to construct a treatment plan, is placed only in one of the permissible seed positions as determined by the $SP_{(x,y,z)}$ parameter. This constraint, in combination with equation (3.23), ensures that all needles are also confined to locations that are permitted by the $SP_{(x,y,z)}$ parameter.

$$Seed_{(x,y,z)} \leq SP_{(x,y,z)}, \forall x \in Q, y \in R, z \in S \quad (3.22)$$

3.3.6.2 Seed/Needle Coupling

Equation (3.23) is a common constraint borrowed from D'Souza et al. (2001), and it allows keeping track of the specific needles that are in use based on the seeds that have been placed within the seed-placement boundaries.

$$Seed_{(x,y,z)} \leq Needle_{(x,y)}, \forall x \in Q, y \in R, z \in S \quad (3.23)$$

3.3.6.3 Avoiding Adjacent Seeds

In treatment planning, one of the key issues to pay attention to is avoiding placement of adjacent seeds in the x - y plane (or the transverse plane) and to a lesser extent in the y - z plane (sagittal plane). Clumping of adjacent seeds in the x - y plane can lead to regions of high radiation exposure within the prostate tissue, regions that are also known as “hot spots”. Such hot spots are generally of concern when considering the V200 dosimetric criterion. V200 acts to limit the volume of the prostate that receives a dose greater than twice the prescribed dose. Staggering of seeds diagonally results in a separation distance of $\sqrt{0.5}$ cm between the closest pair of seeds, rather than the 0.5 cm separation distance observed with seeds placed in adjacent rows or

columns. The greater seed separation achieved in the x - y or transverse plane (through staggering of seeds) acts as a form of buffer against the formation of hot spots in the prostate.

In addition to creating hot spots, placing adjacent seeds in clusters or clumps can also result in the opposite dosimetric scenario in other areas of the prostate, and this can lead to regions known as “cold spots”. Cold spots occur when a region of the prostate does not contain sufficient radioactive sources in its vicinity, and therefore, does not receive the prescribed dose. Thus, cold spots can be of concern when considering the V100 dosimetric criterion, which acts to limit the volume of the prostate receiving less than the prescribed dose. The presence of both hot spots and cold spots inevitably affects the homogeneity of the dose delivered to the prostate gland as a whole, and this can lead to inconsistent treatment plans from patient to patient. The staggering or checkering of seeds allows the above dosimetric issues to be largely avoided, and the prescribed dose to be distributed throughout the prostate in a more uniform or homogeneous fashion. Constraints (24) through (26) can be utilized to avoid placement of adjacent seeds in the prostate:

$$Seed_{(x,y,z)} + Seed_{(x+5,y,z)} \leq 1 \quad , \forall x \in Q, y \in R, z \in S \mid x \leq 57 \quad (3.24)$$

$$Seed_{(x,y,z)} + Seed_{(x,y+5,z)} \leq 1 \quad , \forall x \in Q, y \in R, z \in S \mid y \leq 57 \quad (3.25)$$

$$Seed_{(x,y,z)} + Seed_{(x,y,z+5)} \leq 1 \quad , \forall x \in Q, y \in R, z \in S \mid z \leq 50 \quad (3.26)$$

3.3.6.4 Symmetrical Seed Distribution

Establishing symmetry in the x - y plane for a given seed distribution is a desired requirement in treatment planning for many brachytherapy teams. The prostate is usually imaged with a consideration for aligning the implant template grid's Column D (or Column 32,

approximately the center, bisecting column of the grid) in a manner so that it bisects the prostate into two halves, left and right, in the transverse imaging plane. When imaging is aligned in such a fashion, interestingly enough, Column D also generally coincides with the urethral path inside the prostate. Column D is highlighted in green in Figure 3.7. Thus, when attempting to establish a symmetric seed distribution, clinicians pay special attention to placing seeds in a symmetrical fashion Column D in most cases. The resulting treatment plan consists of a prostate that contains a distribution of seeds in the left half that is a mirror image of the distribution of seeds in the right half. The vast majority of prostates are fairly symmetrical in the x - y plane, thus it is a valid anatomical consideration to create symmetrical seed distributions during treatment planning.




Figure 3.7. Symmetrical (dark blue) and asymmetrical (orange) position pairs outlined in the permissible seed-placement region (dark red voxels) for slice $z=3.0$ cm (PTV, shadow of the urethral curvature, shadow of the rectal curvature and Column D (Column 32) represented by pink, yellow, light blue and green voxels respectively).

Symmetrical seed distributions also result in symmetrical needle insertions. This symmetry simplifies the verification of needle and seed placements during surgery, as an operative oncologist will not be required to keep track of irregular or odd placements of needles.

Two factors need to be kept in mind when attempting to achieve symmetric seed distributions. These two factors are 1) the presence of symmetrical position pairs in and around the prostate, and 2) the presence of asymmetrical seed positions in the seed-placement region (seeds that do not contain an accompanying partner on the opposing side of Column D). A typical prostate possesses a roughly symmetrical geometry; therefore it is logical to use symmetrical position pairs for the majority of seeds placed in the gland. Slight irregularities in gland shape may be accommodated with the placement of seeds in asymmetrical positions. These two factors play a role in the amount of flexibility clinicians would like to have at their disposal in establishing symmetrical seed distributions. In some cases, a brachytherapy team might want a purely symmetrical distribution, while in other cases they may allow individual un-paired seeds to be placed when dealing with an irregularly shaped prostate in order to meet dosimetric requirements. Thus, it makes logical sense to handle these two types of position pairs separately, each with specific constraints of their own.

3.3.6.5 Symmetrical Position Pairs

In order to achieve a symmetrical distribution, seeds must be implanted in symmetrical position pairs that exist within the seed-placement boundaries. Symmetrical position pairs arise when a given seed position has a corresponding mirror image seed position that exists on the other half of the prostate in the transverse (x-y) plane, considering that the prostate is bisected by

Column D, as previously explained. This condition is depicted in Figure 3.7, with a single symmetrical position pair highlighted in blue circles in the seed-placement grid for a sample PTV slice, and Column D identified by a green zone.

Constraint (3.27) can be utilized to establish symmetry in the transverse plane. This constraint only searches for symmetrical seed positions in the exploration space that exists within the seed-placement boundaries. Furthermore, searching only the left half of the PTV (or equivalently, only the right half) is sufficient to find all of the symmetrical position pairs within the seed-placement boundaries.

$$Seed_{(x,y,z)} = Seed_{(64-x,y,z)}, \forall x \in Q, y \in R, z \in S \mid SP_{(x,y,z)} + SP_{(64-x,y,z)} = 2, x \leq 32 \quad (3.27)$$

3.3.6.6 Asymmetrical Position Pairs

Asymmetrical position pairs arise when a given seed position does not have a corresponding mirror image seed position that exists in the other half of the prostate in the transverse (x-y) plane. This situation is also depicted in Figure 3.7, with asymmetrical position pairs highlighted in orange circles.

Constraint (3.28) is utilized to deal with placement of seeds in asymmetrical position pairs. This constraint only searches for asymmetrical seed positions in the exploration space within the seed-placement boundaries. The inclusion of this constraint forces a given solver to avoid placing seeds in asymmetrical position pairs. The exclusion of the constraint allows the solver to place seeds in asymmetrical position pairs if deemed fit. Thus, this constraint on its own offers some level of flexibility to brachytherapy teams when dealing with irregularly shaped prostates.

$$Seed_{(x,y,z)} = 0, \forall x \in Q, y \in R, z \in S \mid SP_{(x,y,z)} + SP_{(64-x,y,z)} = 1 \quad (3.28)$$

It should be noted that $Seed_{(x,y,z)} + Seed_{(64-x,y,z)} = 0$ may also be used in place of $Seed_{(x,y,z)} = 0$ in equation (3.28), however both formulations would lead to the same result of preventing seeds from being placed on asymmetrical position pairs. As an aside, one may handle both symmetrical and asymmetrical position pairs in the same family of constraints, rather than in separate sets of constraints as we have done. Such a feat may be accomplished by removing $SP_{(x,y,z)} + SP_{(64-x,y,z)} = 2, x \leq 32$ from equation (3.27) and entirely omitting constraint (3.28) from the model. This search method provides potential simplification of the model at hand, as clinicians would only have to control a single set of constraints (placement/non-placement of symmetrical seed pairs) in a graphical user interface (GUI) environment, whereas including equation (3.28) would also provide a clinician control over the placement/non-placement of unpaired, asymmetrically positioned seeds. However, it also offers less flexibility to the end user, as it only allows placement of seeds in symmetrical position pairs and explicitly prohibits placement of seeds in asymmetrical position pairs.

3.3.6.7 Maximum/Minimum Number of Seeds

A common goal is to implant 75 to 125 seeds of strength 0.5 U to treat a prostate of a given size and shape. To accommodate this guideline, in addition to the placement of the $Seed_{(x,y,z)}$ variable in a minimizing objective function, as is usually done, constraints (3.29) and (3.30) have been developed.

$$\sum_{x \in Q} \sum_{y \in R} \sum_{z \in S} Seed_{(x,y,z)} \leq upperLimitSeeds \quad (3.29)$$

$$\sum_{x \in Q} \sum_{y \in R} \sum_{z \in S} Seed_{(x,y,z)} \geq lowerLimitSeeds \quad (3.30)$$

3.3.6.8 Maximum/Minimum Number of Needles

When considering the number of needles used to implant seeds, an upper limit of 25 needles is preferred for a given patient, although there is flexibility associated with this value as large prostates may require more needles. The associated constraints are shown in equations (3.31) and (3.32).

$$\sum_{x \in Q} \sum_{y \in R} Needle_{(x,y)} \leq upperLimitNeedles \quad (3.31)$$

$$\sum_{x \in Q} \sum_{y \in R} Needle_{(x,y)} \geq upperLimitNeedles \quad (3.32)$$

3.3.6.9 Maximum Needle Retraction

During brachytherapy, seeds are implanted into the prostate using needles. The depth to which a needle penetrates the prostate, along the z direction, is an important aspect of the surgical procedure to consider. Every penetration made with a needle causes trauma to the prostate tissue, which induces swelling. Such trauma might result in intra-operative complications for the oncologist conducting the seed implantation, because the subsequent swelling might distort the expected geometry of the prostate and hinder the surgeon's ability to implant seeds in their correct positions as determined in the pre-plan study. Furthermore, trauma to the prostate and any associated swelling may also cause post-operative discomfort or complications for the brachytherapy patient, such as the restriction of the urethra.

Thus, it is desired by most clinicians that every needle that is utilized in a typical brachytherapy procedure have a minimum needle penetration distance into the prostate, from the

inferior or apical end of the prostate along the sagittal z -axis. Such a minimum penetration distance can also be represented as a maximum needle retraction, as measured from the base of the prostate, which essentially takes the opposite reference point. The justification for using a maximum needle retraction constraint as opposed to a minimum needle penetration constraint is based on the procedure followed during transrectal ultrasound imaging of the prostate. During sequential transverse imaging, all prostate slices are imaged from the base to the apex with a 0.5 cm retraction increment, or step-size, using a transrectal ultrasound (TRUS) probe. Thus, the imaging reference point is the base of the prostate at $z=0.0$ cm rather than the apex. Based on this, the apex may be assigned varying slice numbers depending on the shape or size of a particular prostate. To take advantage of this imaging reference point, it is preferable to use a maximum needle retraction constraint, which is based on the base of the prostate, rather than a minimum penetration distance constraint, which is based on the apex.

Maximum needle retraction essentially represents the maximum distance that the tip of a needle is allowed to be located away from the base of the prostate. In our models, this distance is set to be 1.5 cm to conform to the clinical inclinations of the brachytherapy team we worked with. Equation (3.33) represents the maximum needle retraction constraint.

$$Needle_{(x,y)} \leq \sum_{\substack{z \in S \\ z \leq \max NeedleRetraction}} Seed_{(x,y,z)}, \forall x \in Q, y \in R \quad (3.33)$$

The use of the above constraint ensures economical use of needles in brachytherapy, as it prevents penetrations of the prostate solely in order to deliver seeds close to the surface of the gland near and around the apex, which is proximal to the implanter during the surgical

procedure. By preventing excessive puncturing of the prostate, unwarranted trauma to the gland is avoided.

3.3.6.10 Maximum/Minimum Number of Seeds in a Needle

To continue along with the theme of economical use of needles in brachytherapy, a minimum number of seeds per needle constraint is also incorporated into our models. The minimum acceptable number of seeds, to warrant the use of a needle, has been set at 2 seeds per needle. The constraint in equation (3.34) represents the minimum number of seeds in a needle. Similarly, the constraint in equation (3.35) represents the maximum number of seeds allowable in a needle, set at 8 seeds per needle.

$$minSeedsNeedle * Needle_{(x,y)} \leq \sum_{z \in S} Seed_{(x,y,z)}, \forall x \in Q, y \in R \quad (3.34)$$

$$\sum_{z \in S} Seed_{(x,y,z)} \leq maxSeedsNeedle * Needle_{(x,y)}, \forall x \in Q, y \in R \quad (3.35)$$

3.3.6.11 Maximum Number of Consecutive Seeds in a Needle

The constraint for avoiding adjacent seeds in the z -direction may prove to be too arduous to implement especially in small prostates. Thus, a similar yet more flexible constraint, for minimizing adjacent seeds in the z -direction, is formulated. The maximum number of consecutive seeds in a needle constraint essentially helps to prevent excessive clustering of trains of seeds in a needle. Usually a constraint to place an upper limit on the number of seeds loaded into a needle is counterintuitive to the notion of establishing economical use of seeds and needles. However, this strategy can also result in some needles being fully loaded with seeds from needle tip to base. Generally, brachytherapy teams prefer to have spacers between seeds

that are loaded into needles, both for ease of seed implantation and to avoid the previously mentioned issue of adjacent seeds leading to hot spots. For our purposes, the maximum number of consecutive seeds is set to 4, as seen in constraint (3.36). Beyond this number of seeds, a spacer is required to be placed on both sides of any given sub-train of seeds.

$$Seed_{(x,y,z)} + \sum_{\substack{m \in S \\ 5 \leq m \leq 20}} Seed_{(x,y,z-m)} \leq maxConsecSeeds, \forall x \in Q, y \in R, z \in S \mid z \geq 20 \quad (3.36)$$

3.4 Experimental Design

For the purposes of demonstration, we solve two instances of our brachytherapy model (benchmark and base), for a single medium-sized (42.5cc) prostate, in order to illustrate the improvement in treatment plan quality that is achievable after the implementation of spatial constraints and data processing techniques. Further test cases involving prostates of different sizes, shapes and geometries will be investigated in Chapter 5. The benchmark model contains the objective function given in equation (3.1), dose calculation constraints, as well as the complete set of dose-volume constraints associated with Table 3.1. With the exception of constraints (3.22)-(3.23) and (3.29)-(3.32), the rest of the spatial constraints are not applied to this model in order to illustrate the results that are attainable without them. For the base model, in addition to the objective function, dose calculation constraints and dose-volume constraints, we include the spatial constraints given in equations (3.22) to (3.36) with the exception of equations (3.26) and (3.35). These two equations are not included because they are not deemed as essential for capturing the specific planning style of our local cancer centre. For the benchmark model, the seed-placement region is confined to the PTV, while for the base model the seed-placement region is constructed through the “global encapsulation” method. For both

models, urethral curvature is accounted for; however, 0.4 cm protection zones around the urethra and rectum are only included in the base model, not in the benchmark model. Although our proposed base model may also be compared to other research-based or commercial automated planning approaches, the results of such comparisons may be less revealing due to the black-box design of these alternative systems. On the other hand, our experimental setup allows us to precisely control the internal algebraic structure of the benchmark and base models tested in our study. Therefore, we are able to visualize the direct impact of the inclusion (or non-inclusion) of the spatial constraints and input data processing techniques that we propose in this chapter.

We use AMPL (Fourer et al. 2002) to construct and Gurobi 6.5.0 (Gurobi Optimization 2014) to solve our models on a 20-core ACPI multiprocessor X64-based PC with Intel Xeon® CPU E5-2650 running at 2.3 GHz with 128 GB RAM. There were several reasons behind our decision to utilize AMPL and Gurobi: 1) the ease of use and high functionality of AMPL as a modeling language, 2) the consistently high performance results achieved by Gurobi, as a solver, for a wide range of optimization problems, 3) the extensive use of both AMPL and Gurobi in academia and industry, 4) our desire to use a deterministic approach (the simplex method in combination with the branch-and-bound approach) to solve our MILP brachytherapy model. Our rationale for this final point is based on the relative unpopularity of the simplex method (and the branch-and-bound algorithm) in comparison to more extensively researched stochastic and heuristic methods in LDR and HDR brachytherapy optimization literature (De Boeck et al. 2014). For both models, the following parameters are set: **presolve = 10** in AMPL and **mipgap = 0.0001**, **timelim = 3600**, **presolve = 2**, **heurfrac = 0.25** in Gurobi. The given **presolve** settings

in AMPL and Gurobi allow an aggressive presolve strategy to be implemented, which helps to effectively clean and simplify the models from redundant constraints and variables before the branch-and-bound solution process is started. The time limit has been set at one hour, to reflect the lower limit of the time that an expert planner takes to produce a treatment plan in the pre-operative treatment planning approach at the local clinic. Since it is generally preferable to allow linear programming models to solve to low optimality gaps, we have set our optimality gap at 0.01%. Gurobi has a built-in Relaxation Induced Neighborhood Search (RINS) heuristic (Danna et al. 2005), that enables its branch-and-bound search method to more effectively find improved feasible solutions. The Gurobi parameter **heurfrac** determines the fraction of the solution time and resources that should be allocated for using the RINS heuristic. After testing a range of values, including the default value of 0.05, we have set this parameter to 0.25, as it achieves a good balance between solution time and the number of feasible solutions found during the search. Keeping the **mipgap** and **timelim** parameters in mind, Gurobi halts its branch-and-bound search once one of these two limits is reached and provides the best integer solution, currently at hand, as the final result of the model.

3.5 Results

The treatment plans, or seed distributions, from the benchmark model and the base model are compared with each other, as well as with a manually-produced treatment plan for the same prostate. Dose distribution results for the benchmark model and the base model are displayed in the left-most and middle columns, respectively, of Figures 3.9-3.11 and the manually-produced treatment plan is displayed in the right-most column. For all plans, the final (post-optimization)

dose distribution was calculated in the VariSeed LDR Treatment Planning System using the TG-43 2D (line source) formalism (Rivard et al. 2004). The isodose legend for these treatment plans is given in Figure 3.8. Furthermore, Figures 3.12-3.14 provide the resulting needle distribution and needle loading in the benchmark model, the base model and the manually-produced treatment plan, respectively. While detailed solution performance results are not a focus of this chapter, Table 3.2 provides the relevant numerical features of both models (following presolve) for demonstrative purposes. Table 3.3 provides the spatial characteristics of the two treatment plans that were produced.

Table 3.2. Numerical features of branch-and-bound solution.

Treatment Plan	Constraints	Variables		Solution				
		Integer	Total	Obj. Func.	Best Bound	Optimality	Iterations	Time (s)
Benchmark Model	3151	2751	3210	249.460	193.267	22.5%	6642485	3600
Base Model	3666	2576	2856	236.150	208.074	11.9%	7250215	3600

Table 3.3. Spatial characteristics of treatment plans.

Treatment Plan	Seeds			Needles			
	Number	X-Y Adjacency	X-Y Symmetry	Number	Less than 2 seeds/needle	Greater than 4 consecutive seeds	Greatest Retraction
Benchmark Model	91	37/91	22/85	24	2	8	5.0 cm
Base Model	98	0/98	88/88	25	0	0	1.5 cm
Manually-produced	104	0/104	102/104	22	0	0	1.5 cm

The solution of both models was terminated at the one-hour mark, with benchmark model and the base model resulting in 22.5% and 11.9% optimality gaps respectively. An interesting aspect to note is that the base model contained fewer variables but a greater number of constraints than the benchmark model. The lower number of variables in the base model was due to the introduction of spatial constraints enforcing the use of symmetrical position pairs for seed-placement, which led to the restriction, or refinement, of the permissible seed-placement grid. The 0.4 cm protection zone added around the urethra and the rectum most likely led to the base model's lower objective function value of 236.150 Gy (in comparison to the benchmark model's value of 249.460 Gy) and hence the differences in their optimality gaps. However, the difference in the two objective function values is not of critical importance, as the urethral and rectal dose-volume metrics for both models align with clinical requirements as seen in Table 3.4.

Interestingly, both the benchmark model and the base model produced treatment plans that utilized a lower number of seeds and a higher number of needles than the manual plan: 91 seeds and 24 needles for the benchmark model, 98 seeds and 25 needles for the base model. Although minimizing seed count was not a goal that was specifically set in our objective function, it is evident that Gurobi produced plans for the benchmark model and the base model with lower seed counts than the manual plan, which contained 104 seeds. The lower seed count might be attributed to the solver's attempt to minimize average dose to the urethra and rectum, in combination with the effects of the PTV V200 constraint and the limited availability of seed and needle positions for potential use. There were clear differences in the spatial arrangement of seeds and needles between the two models. In terms of avoiding adjacent seeds, the benchmark

model performed poorly compared to the base model, as it contained 37 adjacent seeds in the x - y plane (40.7% of total seeds in the benchmark model), while none of the seeds in the base model were adjacent. In terms of symmetry, the benchmark model was once again subpar, as only 22 of its seeds (25.9%) were symmetric in the x - y plane, while this value was 100% for the base model. Seeds that were placed in Column D (Column 32) were taken out of this percentage calculation for symmetry (6 seeds for the benchmark model, 10 seeds for the base model), as these seeds technically cannot form symmetrical pairs. In other words, the base model was able to achieve a fully symmetric seed distribution in every prostate slice and was also able to avoid all possible cases of adjacency in the x - y plane. In regards to needle usage, the benchmark model encountered the exact potential issues that were mentioned earlier in this chapter. Namely, 8 needles in the benchmark model contained more than 4 consecutive seeds; of particular interest was the needle positioned at e3.0 (coordinates e on the x -axis, 3.0 on the y -axis) containing 7 consecutive seeds as see in Figure 3.13. Furthermore, the benchmark model utilized 2 needles loaded with only 1 seed, which amounts to an inefficient use of needles. The benchmark model also had difficulties with needle retraction, as 6 of its needles had a retraction of greater than 1.5 cm. Specifically, needles utilized at locations c2.5 and c3.5 had a retraction of 5.0 cm. Considering that these two needles were also loaded with only two seeds each, this gives rise to a rather impractical, and perhaps unjustifiable, delivery method involving multiple implantations of peripherally located seeds. With the implementation of the spatial constraints, the base model encountered none of these issues that were present in the benchmark model. Needles utilized in the base model had a maximum retraction of 1.5 cm. Additionally, there were no cases of single-

seed loading in the treatment plan of the base model, nor any cases of needles containing greater than four consecutive seeds. Thus, the base model was able to satisfy all of the technical spatial requirements, in a similar manner to the manual plan (which consisted of 98.1% symmetrical seeds, 0% adjacent seeds, no cases of needles with single-seed loading, no cases of needles containing greater than 4 consecutive seeds and no cases of needles with greater than 1.5 cm retraction).









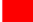





Isodose Legend Gy [% of Prescription Dose]		290.00 [200.00%]		145.00 [100.00%]		72.50 [50.00%]
		217.50 [150.00%]		130.50 [90.00%]		36.20 [24.97%]
		159.50 [110.00%]		116.00 [80.00%]		
Anatomy/Landmark Legend		Prostate		Rectum		Urethra_dosim
		Urethra		PTV		Pubic Arch

Figure 3.8. Legend for treatment plan isodoses and contours.

Table 3.4. Summary dose-volume metric values for the three treatment plans.

Dose Parameter	Objective	Treatment Plan		
		Manual	Benchmark Model	Base Model
PTV Prescribed Dose	145 Gy	145 Gy	145 Gy	145 Gy
PTV V100	> 98%	99.6%	98.7%	98.3%
PTV V150	≤ 65%	64.2%	55.0%	52.7%
PTV V200	≤ 25%	22.8%	25.7%	21.3%
PTV D90	174 – 189 Gy	186.2 Gy	174.4 Gy	171.2 Gy
Urethra D5	< 215 Gy	199.7 Gy	210.7 Gy	191.1 Gy
Rectum D1cc	< 145 Gy	110.1 Gy	102.8 Gy	95.0 Gy




Figure 3.9. Comparison of treatment plans (slices $z=0.0$ cm to $z=1.5$ cm).




Figure 3.10. Comparison of treatment plans (slices $z=2.0$ cm to $z=3.5$ cm).

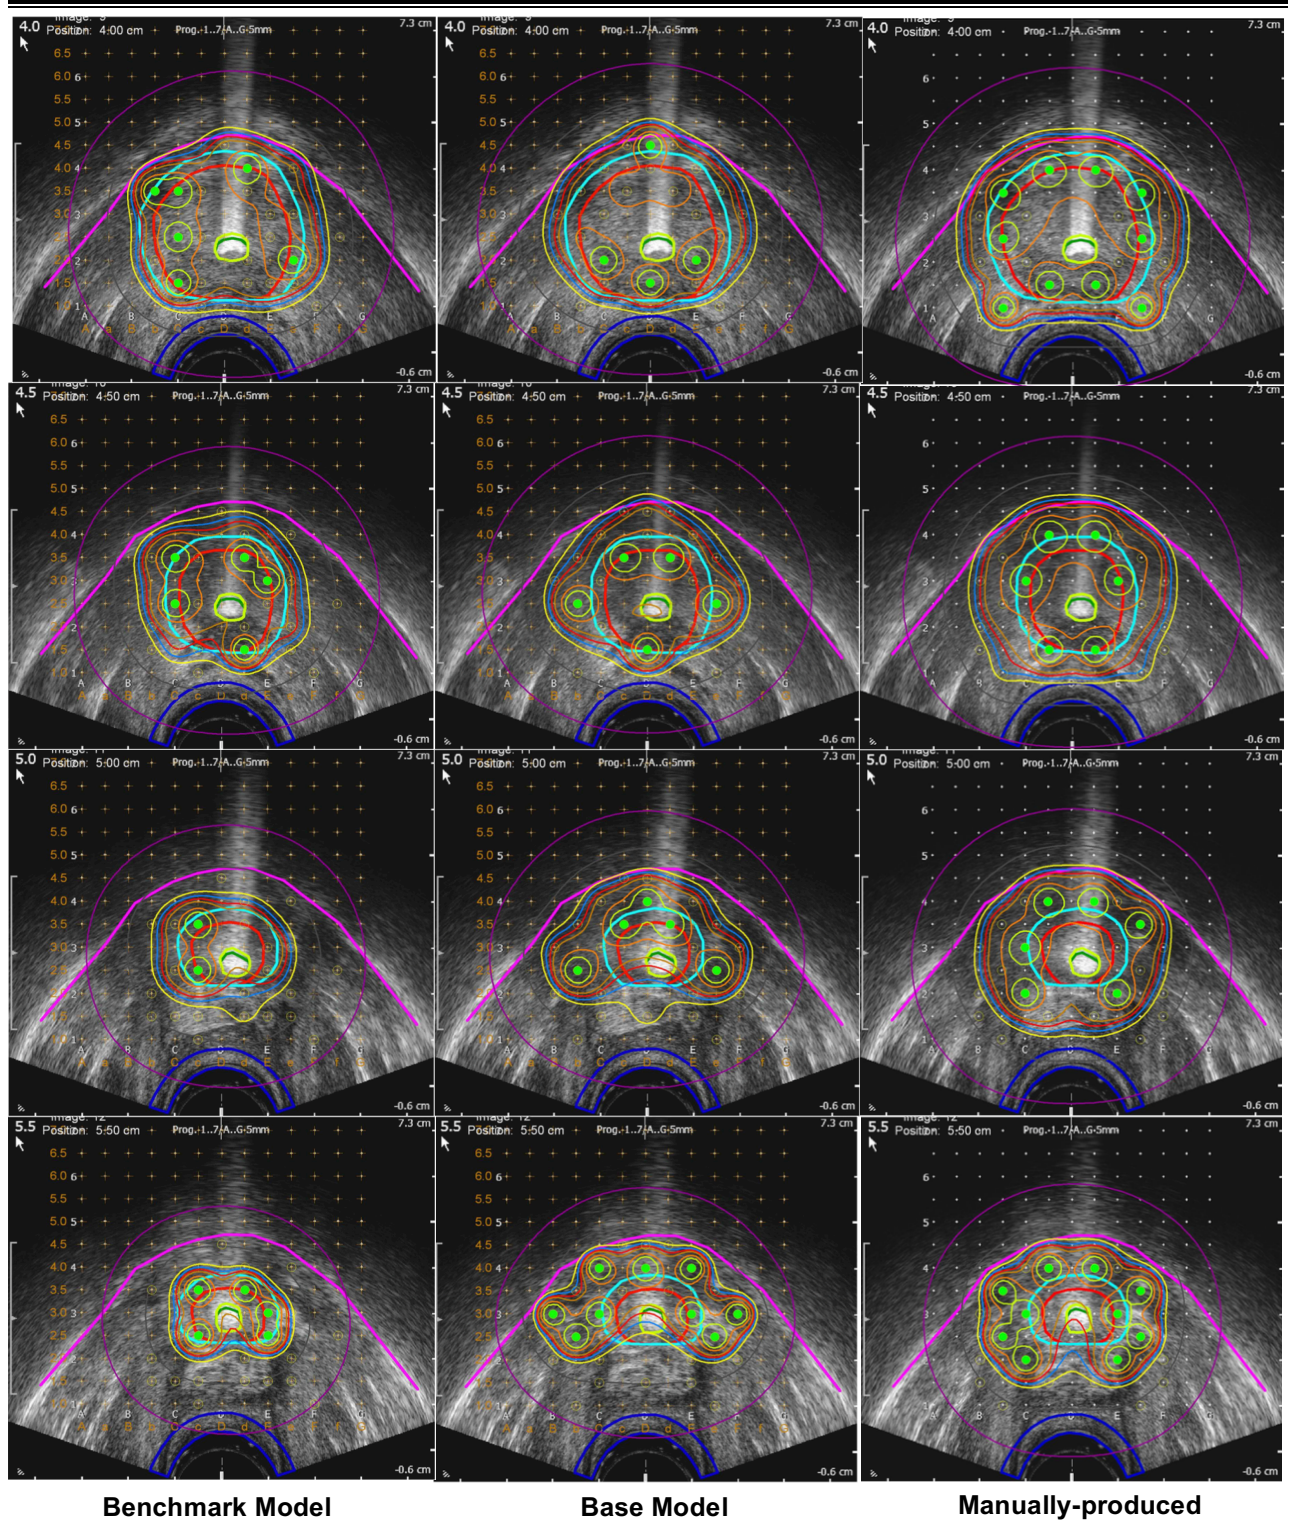


Figure 3.11. Comparison of treatment plans (slices $z=4.0$ cm to $z=5.5$ cm).

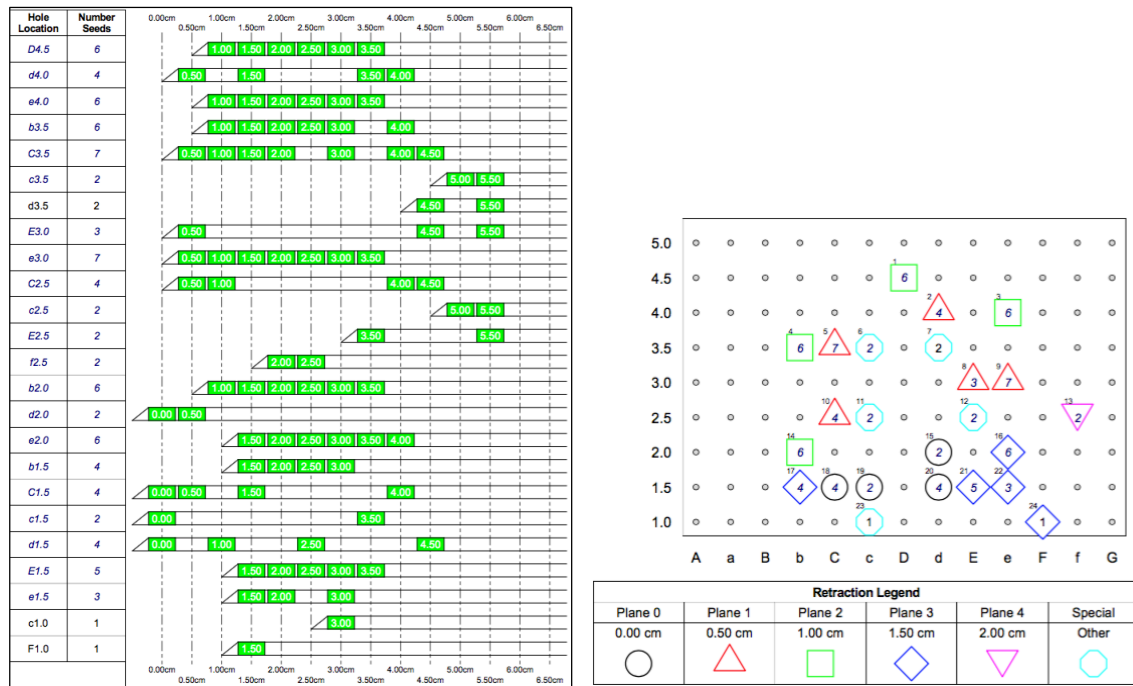


Figure 3.12. Needle loading (left) and needle distribution (right) in the benchmark model treatment plan.




Figure 3.13. Needle loading (left) and needle distribution (right) in the base model treatment plan.

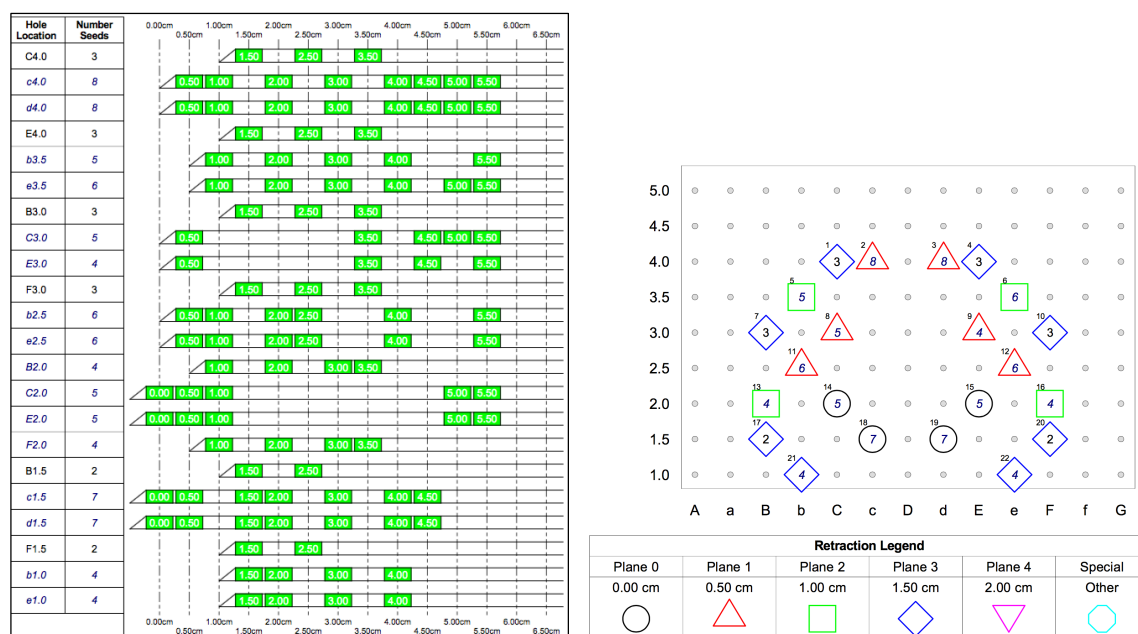


Figure 3.14. Needle loading (left) and needle distribution (right) in the manually-produced treatment plan.

3.6 Discussion

In this section, treatment plans created through the benchmark and base models are clinically analyzed by an experienced planner (medical physicist) at the local cancer center. The benchmark model plan, which is based on the dose-volume requirements in Table 3.1 but was created without any spatial constraints on seed placement, does not exhibit the spatial characteristics of a clinically acceptable seed distribution and therefore would not be used to treat a patient. As the dose-volume data in Table 4 demonstrate, the benchmark model plan meets all of the dose-volume objectives with the minor exception of PTV V200. This situation, in which the dose-volume requirements are met but the spatial distribution of seeds is unacceptable, occurs because the seed placement problem formulated using only the requirements in Table 3.1 is under-determined, and so admits many solutions.

The base model plan exhibits many of the features of a seed distribution that would be deemed clinically acceptable at the Cross Cancer Institute. It meets all but one of the specified dose-volume objectives, only missing narrowly for PTV D90. The addition of spatial constraints, equations (3.22)-(3.36), to the dose-volume requirements has had the desired effect of largely eliminating non-viable solutions from the solution space. That being acknowledged, further refinements to the base model plan are certainly desirable from a clinical perspective. With reference to Figures 3.9-3.11 and Table 3.4, these are as follows:

- Improved dose coverage at the prostate base and apex (positions 0.0, 5.0, 5.5): The 1.3% reduction in V100 relative to the manual plan is mostly due to reduced base and apex coverage.
- Improved dose coverage at the posterior surface of the prostate near mid-gland (positions 2.5, 3.0): Although this coverage lapse is relatively minor, eliminating it would definitely improve the overall quality of the plan.
- Seed placement closer to the PTV (positions 0.0, 0.5, 5.0, 5.5): Recent work by Liu et al. (2015) has shown that for some patients, seeds placed too far from the PTV do not follow the movements of the mobile prostate, which can lead to reduced dose coverage at the prostate surface. For this reason, manual plans typically avoid placing seeds more than a few millimeters outside the PTV whenever possible.
- Seed placement avoiding template Column D (positions 0.0-1.5, 4.0-5.5): Many manual planners avoid placing seeds in Column D as a matter of conservative practice to avoid any unintended intersections of a needle path with the urethra. However when the urethra is

visualized using aerated gel (as it is here), it could be argued that such avoidance is unnecessary.

3.7 Conclusion

This chapter proposes a new mixed-integer linear programming modeling approach for solving the interstitial LDR prostate brachytherapy problem through mathematical optimization. The novelty that is brought forward in this study rests in the spatial considerations of treatment plans. These spatial constraints, largely overlooked by previous optimization studies, constitute a vital aspect of the clinical protocols followed by brachytherapy teams generally. We demonstrate that with the inclusion of basic spatial constraints and input data processing techniques in our models, a result similar to that deemed clinically acceptable can be achieved even with simple dose calculation formalisms, such as the TG-43 1D isotropic point-source approximation, and a grid system based on a coarse $0.5 \times 0.5 \times 0.5 \text{ cm}^3$ resolution. The treatment plan produced through our new approach is able to largely capture the qualities and characteristics of the manual plan created by an expert planner. Nevertheless, it is clear that further improvements are necessary in order to produce treatment plans that more closely mimic manually created plans. To achieve these results, we will be introducing more complex data processing and mathematical modeling techniques in a follow-up study.

The benefits of our linear programming optimization approach have the potential to impact both healthcare practitioners and cancer patients. Implementation of our optimization models will potentially result in streamlining of treatment planning in brachytherapy, and therefore efficient use of resources (such as time, equipment and personnel) as well as cost

savings for healthcare organizations. Most importantly, cancer patients can be provided with consistent and high quality treatment plans, regardless of whether they are treated in a specialized cancer center, with a highly-experienced brachytherapy team, or a community hospital, which may have less brachytherapy experience. The availability of high quality treatment plans should also result in patients experiencing less severe post-operative complications or side effects. The power of this study lies in the fact that the concepts proposed in our optimization model are also applicable within a broader context, both in HDR prostate brachytherapy as well as in the brachytherapy treatment of other forms of cancer, including breast, cervical, lung, and head and neck cancer.

Chapter 4²

The Enhanced MILP Model Formulation

Despite attempts (and successes) at improving solution speed and dosimetric reliability of treatment plans over the years, the use of optimization techniques for interstitial LDR prostate brachytherapy is still not commonplace in cancer clinics. This may be due to the fact that although it is important to achieve fast solution times and satisfactory dosimetric values, these accomplishments alone do not guarantee that a given treatment plan will be deemed clinically acceptable. The inability to produce “center-specific” results stands out as the main culprit behind the incapability of optimization techniques to achieve extensive use in clinics. Center-specificity essentially rests on spatial considerations that expert planners pay particular attention to during manual planning. Such spatial considerations focus on the geometric arrangement of seeds and these considerations define the treatment planning style of a cancer clinic. Considering that many cancer centers have developed and refined their own unique style of treatment planning, spatial characteristics play a crucial role in the perceived clinical quality of treatment

²*A version of this chapter has been submitted for publication: Babadagli, M.E.; Sloboda, R.; Doucette, J. “Enhanced Data Processing and Integer Linear Programming Techniques for Capturing Expert Planning Style in Interstitial Low Dose Rate Prostate Brachytherapy”. Springer, Annals of Operations Research (submitted April 8, 2016).*

plans used in each clinic.

In Chapter 3, we put forth a focused attempt to capture the planning style of expert planners through the design of spatial constraints, since in prior work spatial considerations had been largely overlooked in the LDR brachytherapy optimization literature. The main objective of our previous work was to lay the groundwork for a new mathematical modeling approach that would enable automated construction of treatment plans possessing clinical quality similar or equivalent to that observed in manually-produced plans. In order to validate our unique angle on brachytherapy modeling, we descriptively compared and contrasted our particular approach to various studies including those by Lessard et al. (2006), Yoo et al. (2007) and McGeachy et al. (2015).

The basic concepts introduced in Chapter 3 were sufficient to generally capture the planning style of expert planners at a local cancer centre. In this chapter, we further develop those ideas by exploring relatively more complex data processing techniques and spatial constraints in order to improve upon our ability to realistically reproduce the planning style of expert planners. We believe that the ability to produce treatment plans that brachytherapy teams will find familiar to their own is a key feature for the practical integration of our automated planning approach in real-world clinics. The new concepts introduced in this chapter will be demonstrated through a comparison of our results to those obtained via a base model, formulated in Chapter 3, as well as a manual planning approach for a single prostate test case. MILP solution performance, such as solution speed or time, will not be a major focus of this study, and will instead be comprehensively evaluated in a future investigation. Similarly, the ability of our

automated planning approach to consistently produce high quality solutions for multiple prostates of different shapes and sizes will be analyzed in a future study.

4.1 Benchmark LDR Prostate Brachytherapy Problem Formulation (Base Model)

The mathematical structure and characteristics of the LDR prostate brachytherapy problem formulation established in Chapter 3 will act as the foundation, which we will build upon in this chapter. This base model includes the grid and coordinate system (both fine, $0.1 \times 0.1 \times 0.5 \text{ cm}^3$, resolution and coarse, $0.5 \times 0.5 \times 0.5 \text{ cm}^3$, resolution), input data processing techniques and algorithms, sets, parameters, variables, objective function, dose calculation constraints (based on the TG-43 1D isotropic point-source approximation), dose-volume constraints, and spatial constraints with the exception of equations (3.26) and (3.35). Equations (3.26) and (3.35) are omitted as they are not deemed to be critical for capturing the planning style of our local cancer centre. The flowchart given in Figure 3.1 outlines the entire process followed in Chapter 3, consisting of data collection, data processing, model formulation, model solution and post-solution analysis. The sets, parameters, constraints and objective function used to construct the enhanced model in this chapter are provided in the next few sections. The specific formulation of the prostate LDR brachytherapy problem presented in this chapter is tailored to the clinical protocol followed at a local cancer centre. However, this formulation may be modified to satisfy the requirements of other cancer centres as well.

4.2. Improvements to the Base Model

In the following sections, we propose various improvements to the base model introduced in Chapter 3 that allow us to achieve an increasingly realistic representation of the planning style present at our local cancer centre. Some of these improvements are in the form of enhanced data processing techniques that are applied to the discretized anatomical data before it is passed onto our MILP model (as in the cases for determining expanded seed-placement margins, accounting for pubic arch interference as well as establishing urethral pathway protection margins). Other improvements are in the form of new linear constraints that are used in the MILP model itself (as in the cases for removal of seeds from the geometric midline of the PTV, enforcement of maximum seed separation distances in needles and implementation of minimum number of seeds per slice in the PTV).

4.2.1 PTV Expansion and Seed-Placement Margins

During treatment planning, achieving sufficient dose coverage and satisfying the entire set of dose-volume metric requirements may prove to be difficult if seed placement is restricted to the PTV. This is particularly the case for base and apical slices of most prostates, and may also be true for slices closer to the mid-gland in small prostates. Dose coverage issues may be attributed to the fact that prostate slices that are located near the mid-gland have dose contributions from seeds placed in slices superior and inferior to those particular slices. However, base and apical slices are not subjected to the same dose contributions because they define the ends of the prostate and lack neighboring superior and inferior slices, respectively. Placement of seeds outside of the PTV helps to achieve sufficient dose coverage within the PTV,

especially at the two longitudinal extreme ends of the prostate. “Global encapsulation” or the “cylinder approach” is a method that has previously been utilized by McGeachy et al. (2015) to permit placement of seeds outside of the PTV. Global encapsulation involves superimposing the largest PTV contour, which is usually a mid-gland contour, as the uniform permissible seed-placement boundary for all slices in the prostate. In Chapter 3, we applied global encapsulation with a slight modification to create an initial seed-placement region prior to further urethral pathway related data processing of this region. Our modification involved using the largest two mid-gland PTV contours, slices $z=2.5$ cm and $z=3.0$ cm (out of a set of $z=0.0-5.5$ cm), to act as the permissible seed-placement regions for the superior ($z=0.0-2.5$ cm) and inferior ($z=3.0-5.5$ cm) halves of the prostate respectively. The global encapsulation method was carried out during the preprocessing stage of our input data, as seen in Figure 3.1, before the branch and bound solution process commenced.

Although advantageous due to its simplicity, this approach possesses one major critical flaw. The PTV contour of a mid-gland slice may not always be an appropriate fit, both shape and positioning wise, when superimposed over other neighboring mid-gland slices. Therefore, the global encapsulation method may inadvertently omit a few seed-placement positions of high dosimetric utility on slices that it is being applied to. On the contrary, a mid-gland PTV may also prove to be too large a fit when superimposed over the base or the apex as seen in Figure 4.1. Such an imbalanced fit may not be restrictive enough to prevent seeds from being placed at distant locations from the PTV contour as depicted by green circles in Figure 4.1 and green dots in Figure 4.2. This clinically unacceptable scenario is especially problematic for the base slice,

upon which the bladder snugly rests superiorly, and possibly anteriorly as depicted in Figure 4.3. Although not explicitly contoured during transrectal ultrasound (TRUS) imaging, expert planners are still acutely aware of the presence of the bladder in the base slice ultrasound image. This is because distant placement of seeds from the PTV in the base slice may potentially increase the likelihood of puncturing the bladder with a needle during the operation, which may further result in seed migration into this organ. Unless additional explicit direction is given to automated planners to avoid the bladder, utilizing global encapsulation to define seed-placement regions offers no specific measures to safeguard the bladder.

A second, minor flaw of global encapsulation is that, as seen in Figure 4.1, it may not effectively limit the number of seed-placement binary decision variables for relatively small PTV slices, and therefore may lead to the creation of MILP models that are unnecessarily large and burdensome on a MILP solver.




Figure 4.1. Seed-placement regions (dark red voxels), produced through global encapsulation, for the prostate apex (left) and base (right). PTV, urethra and rectum are represented by light pink, light yellow and light blue voxels respectively.




Figure 4.2. Distant placement of seeds from the PTV at the base slice.

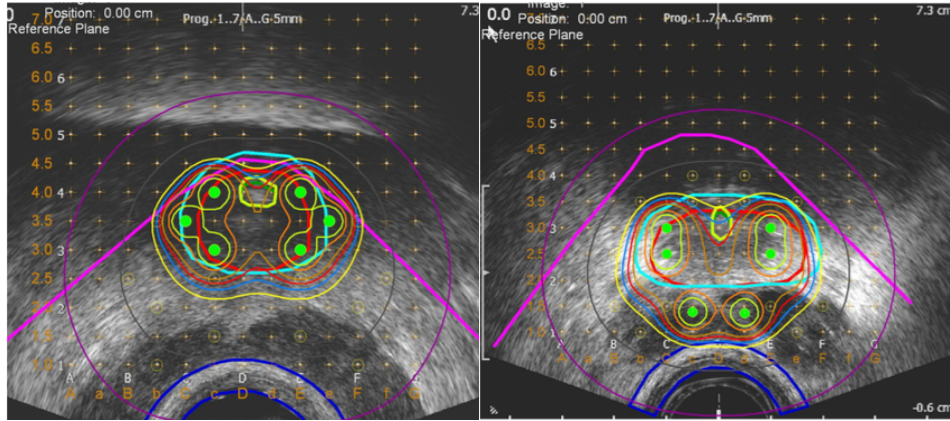


Figure 4.3. The bladder appears (as a dark region) anteriorly to the base PTV slice (turquoise contour) in the right ultrasound (US) image, while it does not in the left US image for two different prostates.

We now propose a unique approach for creating initial seed-placement regions in the prostate: slice-based local encapsulation. In this technique, the PTV contour of each individual slice is expanded by a specific amount to act as the seed-placement boundary for that given slice. The slice-based local encapsulation method is implemented during the preprocessing stage of our input data, as seen in Figure 3.1, before the branch and bound solution process begins. PTV contour expansions are calculated as follows:

- The PTV contour on the largest slice (usually a mid-gland slice) is expanded up to 0.1 cm. Having a 0.0 cm expansion may prove useful for ensuring that at least a segment of all stranded trains of seeds are physically embedded within the PTV, as strand anchoring decreases the chances of seed migration (Lin et al. 2007) (Reed et al. 2007). It may prove practical to apply a 0.1 cm expansion in extremely small prostates for the purpose of gaining extra seed-placement positions of high dosimetric utility.

- The PTV contours on the base and apical slices are expanded 0.3-0.5 cm, or as deemed clinically appropriate by the end-user (expert planner).
- The PTV contours of all slices in between the base and the selected mid-gland slice are expanded by amounts that are calculated as linear interpolations of the expansions applied at these two slices, and rounded to the nearest millimeter (as most commercial brachytherapy software programs permit only integer millimeter expansions of anatomical contours). The same procedure is carried out for the slices between the apical slice and the mid-gland slice. The above methodology is applicable for a single chosen mid-gland slice, or alternatively a selection of two mid-gland slices (which may be useful for prostates consisting of an even number of slices). For the latter case, the local encapsulation procedure is separately carried out for the base and the apex, and their corresponding mid-gland slices (the closer of the two mid-gland slices).

Slice-based local encapsulation results in a tighter, or more balanced, fit around the PTV as seen for two sample slices (apex and base) in Figure 4.4. This leads to placement of seeds that are relatively closer to the PTV, as depicted by green circles in Figure 4.4 and green dots in Figure 4.5. This is especially important to avoid potentially puncturing the bladder at the base slice during the brachytherapy operation. Due to the more conformal expansion of the originally PTV-bound seed-placement regions in each slice, seed-placement boundaries produced via local encapsulation also limit the number of seed-placement decision variables in the resulting MILP model.




Figure 4.4. Seed-placement regions (dark red voxels), produced through slice-based local encapsulation, for the prostate apex (left) and base (right). PTV, urethra and rectum are represented by light pink, light yellow and light blue voxels respectively.




Figure 4.5. Nearby placement of seeds to the PTV at the base slice.

4.2.2 The Pubic Arch

The pubic arch is positioned anteriorly to the prostate gland and is usually delineated during transrectal ultrasound imaging. Generally there exists ample clearance between the prostate and the pubic arch, thus the entirety of the gland (and therefore the PTV) is comfortably accessible by needles. On occasion however, the clearance may not be enough and the anterior portion of the prostate (usually in mid-gland slices) may rest hidden behind the pubic arch as viewed transaxially. Attempting to implant seeds in these inaccessible portions of the prostate may prove to be futile, as needles used in the process are physically obstructed by the pubic arch and therefore are inhibited from travelling into the gland. Pubic arch interference may either occur due to the natural anatomy of the patient or due to an enlarged prostate as a result of disease progression (Kono et al. 2010). Restricted access to relatively large portions of the anterior prostate, as seen in Figure 4.6, may turn seed implantation into a difficult task, and therefore treatment plans need to be mindfully designed in order to accommodate such seed-placement limitations. In Figure 4.6, the pubic arch and PTV are represented by magenta and

turquoise coloured contours (respectively) in the US slice on the left, and by the light green and light pink voxels (respectively) in the discretized image on the right. The inaccessible seed-placement positions are outlined by orange circles and the portion of the PTV that falls behind the pubic arch is represented by dark green voxels.

Pubic arch interference transforms treatment planning into a more complex task, as dose-volume metric requirements (specifically PTV V100 and PTV D90) become harder to satisfy. Anterior portions of the prostate still require adequate doses of radiation, however after sections of the permissible seed-placement region are removed in order to account for pubic arch interference, a planner has fewer positions to potentially place seeds in. For the prostate that is examined in this chapter, step 2 of Algorithm 4.1 accounts for pubic arch interference during the preprocessing stage of the initial seed-placement region input data, as seen in the flowchart given in Figure 3.1.

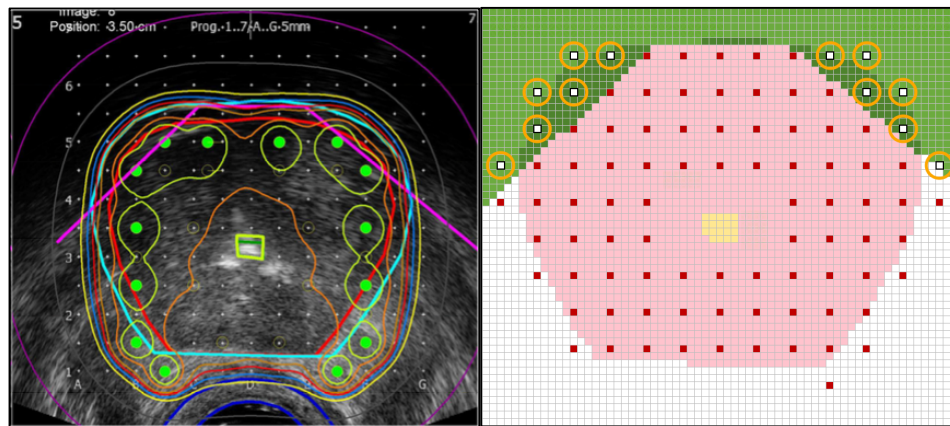


Figure 4.6. Pubic arch (magenta coloured contour) interference shown for a mid-gland slice in US image format (left) and discretized image format (right).

```

1  1. Start with the highest numbered (apical) slices of discretized (fine-resolution) initial seed-placement
2    region (InitSP[k]), pubic arch region (PubicArch[k]), rectum region with 0.4 cm margin (RECMargin[k]),
3    urethra region with 0.4 cm margin (UREMargin[k]), urethral shadow region with 0.2 cm margin
4    (UREShadow[k]), fully populated coarse-resolution grid (CRGrid[k]), and four empty fine-resolution
5    grids (ProSP[k], ProCurveSP[k], NegSP[k], SP[k])
6  2. k = apical slice coordinate
7    m = k
8    WHILE k ≥ 0:
9      b. FOR each voxel ProCurveSP[i,j,k] in slice ProCurveSP[k]:
10         Assign ProCurveSP[i,j,k] = MAX(0, (ProSP[i,j,k] - SUM(PubicArch[i,j,m], ... ,
11          PubicArch[i,j,k])))
12      c. Perform k=k-1.
13  3. k = apical slice coordinate
14    m = k
15    WHILE k ≥ 0:
16      a. FOR each voxel ProSP[i,j,k] in slice ProSP[k]:
17         Assign ProSP[i,j,k] = MAX(0, (InitSP[i,j,k] - SUM(RECMargin[i,j,m], ... , RECMargin[i,j,k])))
18      b. Perform k=k-1.
19  4. k = apical slice coordinate
20    m = k
21    WHILE k ≥ 0:
22      b. FOR each voxel ProCurveSP[i,j,k] in slice ProCurveSP[k]:
23         Assign ProCurveSP[i,j,k] = MAX(0, (ProSP[i,j,k] - SUM((UREShadow[i,j,m], ... ,
24          UREShadow[i,j,k]), UREMargin[k])))
25      c. Perform k=k-1.
26  5. k = apical slice coordinate
27    WHILE k ≥ 0:
28      a. FOR each voxel NegSP[i,j,k] in slice NegSP[k]:
29         IF ProCurveSP[i,j,k] + CRGrid[i,j,k] > 1:
30            Assign NegSP[i,j,k] = 0
31         ELSE:
32            Assign NegSP[i,j,k] = ProCurveSP[i,j,k]
33      c. FOR each voxel SP[i,j,k] in slice SP[k]:
34         Assign SP[i,j,k] = MAX(0, (ProCurveSP[i,j,k] - NegSP[i,j,k]))
35      d. Perform k=k-1.
36  6. End with finalized (coarse-resolution) permissible seed-placement region indexed as SP[x,y,z]

```

Algorithm 4.1. Algorithm to account for the pubic arch, shadow of urethral curvature (with separate margins) and rectal curvature while creating the discretized permissible seed-placement region.

4.2.3 The Urethral Pathway and Safety Measures

Due to the urethral curvature, a specifically delineated region of the PTV is considered to be off-limits for needle penetration and seed placement. In Chapter 3, we considered the simplest case of this region, marked by the entrance point, exit point, the curvature and the diameter of the urethra. Additionally, we implemented a 0.4 cm protection zone around the urethral pathway as a safety measure. In order to further assist oncologists in avoiding the urethra during seed implantation, we propose two improvements in the processing of input data.

4.2.3.1 Removal of Column D (Column 32)

In Chapter 3, we allowed seeds to be placed in Column D, however clinical consultation revealed that such placement of seeds should be avoided when possible. The prostate is generally imaged in a manner that aligns Column D (Column 32) on the implant template grid with the urethral path inside the gland. As such, seed-placement positions that fall on Column D may be removed from the input data as a safety measure. Column D is highlighted in green and the removal of seed-placement positions is outlined by dashed black circles in Figure 4.7. Removal of seeds may be carried out directly on the input data, during the preprocessing stage before the data is passed onto a MILP model, but may also be implemented via a hard constraint within the model itself, as seen in equation (4.1). The benefit of using a hard constraint is real-time flexibility; while making changes to the input data permanently removes seed positions on Column D, using a constraint to accomplish the same feat allows for these seed positions to be conveniently included or excluded as desired by the end-user. A soft constraint, implemented by adding a penalty term to the objective function as given in equation (4.2), may also be used in cases where avoiding Column D is preferable, but not absolutely critical. The same idea may be extended to essentially any other column or row, for example in situations where the urethral path does not fall on Column 32, but perhaps on a column immediately to the left or right of it.

$$Seed_{(x,y,z)} = 0 \quad \forall x \in Q, y \in R, z \in S \mid x = 32 \quad (4.1)$$

$$\text{minimize: } \sum_{x \in Q} \sum_{y \in R} \sum_{z \in S \mid x=32} Seed_{(x,y,z)} \quad (4.2)$$




Figure 4.7. Removal of seed-placement positions on Column D.

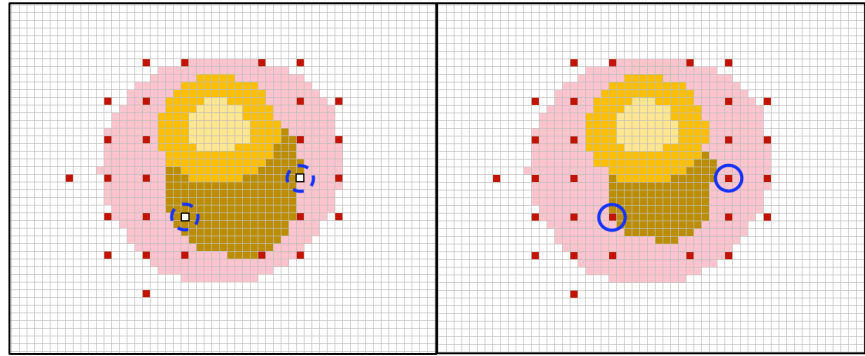


Figure 4.8. Comparison of a single 0.40 cm urethral protection margin (left) and two separate 0.40 cm and 0.20 cm protection margins (right), and the subsequent gain of seed-placement positions (solid blue circles).

4.2.3.2 Separate Seed-Placement and Needle Penetration Margins

When accounting for the urethral curvature, the additional 0.4 cm protection zone restricts both seed-placement and needle implantation around the urethra. While this approach helps lower the likelihood of the urethra from being penetrated with a needle, it may prove to be too aggressive in the consequent removal of potential seed positions. Such removal of seed positions may not be a major issue with medium to large sized prostates (30 cc to 80 cc), which contain a relatively large number of potential seed positions on each slice. However, such a protection zone may prove to be very restrictive for smaller prostates (below 25 cc), which contain a relatively small number of potential seed positions. For these small prostates, achieving sufficient coverage in the prostate tissue surrounding the urethra, and specifically satisfying the PTV V100 and PTV D90 dose limits, may prove to be arduous.

Thus, we suggest an alternative technique to establish a more reasonable balance between urethral protection and removal of seed positions. Instead of utilizing a single margin (0.4 cm) to define the protection zone for both seed-placement and needle penetration, we apply individual

margins. For each PTV slice, a 0.4 cm urethral protection margin is implemented to avoid seed-placement, however, when accounting for the urethral curvature as observed from a given slice, a less restrictive 0.2 cm protection margin is utilized to avoid needle penetration. Algorithm 4.1 outlines the data processing steps that are implemented to transform the initial seed-placement region to the final seed-placement region that is passed onto our MILP model. Step 3 of Algorithm 4.1 accounts for the 0.4 cm rectal protection margin that was introduced in Chapter 3. Step 4 of Algorithm 4.1 is executed in order account for the new separate seed-placement and needle penetration margins discussed in this section.

The technique of using separate margins is clinically permissible at our local cancer clinic, as implanters are comfortable in navigating needles within 0.2 cm of the urethra visualized with aerated gel. Through this approach, we gain a few critically located seed-placement positions around the urethra and the midline of the prostate, which effectively allows us to more easily overcome the dose coverage issues encountered with the previously implemented protection zone (based on a single 0.4 cm margin). The gain of seed-placement positions is outlined by dashed and solid blue circles in Figure 4.8. The urethra, seed-placement margins and the needle implantation margins are represented by light yellow, yellow and dark yellow voxels respectively.

4.2.4 Maximum Seed Separation in a Needle

While it is important to control occurrence of consecutive seeds in a needle, it is also important to limit the spacing or separation of successive seeds within a needle. Historically, limiting maximum seed separation has been a more pressing concern for loose (non-stranded)

seeds than for stranded seeds, as this rule helps to ensure that the distance between seeds remains true to its planned value. For loose seeds separated by several fixed 0.45 cm length spacers, it can be difficult to deposit the spacer portion of the train uniformly as bunching of spacers may occur. Stranded seed trains mitigate both individual seed migration and neighboring seed distance variability observed with loose seeds (Sloboda et al. 2012). Thus, the maximum seed separation rule is more of a stylistic planning choice for stranded seeds where the distance between (even greatly separated) seeds is maintained by the stranding itself as well as the internal spacer materials. Furthermore, needles containing seeds that are placed excessively far apart may behave dosimetrically similarly to single seed loaded needles, especially in the case of the three needles depicted in Figure 4.9 (with seed separation distances of 4.0 cm, 4.5 cm and 5.0 cm). While at first glance these seed configurations fulfill all other needle related requirements of the local cancer centre (minimum retraction is less than or equal to 1.5 cm, the needle contains 2 or more seeds, and the needle does not consist of more than 4 consecutive seeds), these needles do not satisfy the maximum seed separation requirement, which is set to less than 4.0 cm between any two successive seeds.

Constraints (4.3)-(4.5) are designed to limit maximum seed separation in a needle. Constraint (4.3) checks whether any two seeds in the needle are separated by 4.0 cm, and if so, enforces an additional seed to be placed anywhere between these two seeds. Constraints (4.4) and (4.5) are based on a similar line of thought, for separation distances of 4.5 cm and 5.0 cm respectively. One noteworthy aspect of these constraints is the right hand side. In each case, the zone that will house additional seeds (defined by m) is designed in a manner that guides the

placement of intermediate seed(s) to be a maximum of 4.0 cm away from the already existing two seeds. For example, in constraint (4.5), only the zone $15 \leq m \leq 40$ is able to satisfy the maximum seed placement rule, as highlighted by the orange sections in Figure 4.9.

$$(Seed_{(x,y,z)} + Seed_{(x,y,z+45)}) - 1 \leq \sum_{\substack{m \in S \\ 5 \leq m \leq 40}} Seed_{(x,y,z+m)} \quad \forall x \in Q, y \in R, z \in S \mid z \leq 10 \quad (4.3)$$

$$(Seed_{(x,y,z)} + Seed_{(x,y,z+50)}) - 1 \leq \sum_{\substack{m \in S \\ 10 \leq m \leq 40}} Seed_{(x,y,z+m)} \quad \forall x \in Q, y \in R, z \in S \mid z \leq 5 \quad (4.4)$$

$$(Seed_{(x,y,z)} + Seed_{(x,y,z+55)}) - 1 \leq \sum_{\substack{m \in S \\ 15 \leq m \leq 40}} Seed_{(x,y,z+m)} \quad \forall x \in Q, y \in R, z \in S \mid z = 0 \quad (4.5)$$




Figure 4.9. Seed separation of 4.0 cm (needle c4.0), 4.5 cm (needle D4.0) and 5.0 cm (d4.0). Permissible seed-placement positions are highlighted in orange.

The characteristics of constraints (4.3)-(4.5) may also be captured in a single general constraint, given in equation (4.6). Here, *maxNumSeeds* is equal to 3 and represents the total possible number of seeds that may be placed in position $z=4.5$ cm and beyond, until $z=5.5$ cm (the apical end of the prostate). Since we are attempting to design a single constraint that essentially encapsulates the governing rules behind constraints (4.3)-(4.5), we must be mindful of the choice of the right hand side. We see that only the zone $15 \leq m \leq 40$ is capable of producing results that satisfy a maximum seed separation of 4.0 cm for the three cases of seeds separated by 4.0 cm, 4.5 cm, and 5.0 cm.

$$(Seed_{(x,y,z)} + (\sum_{\substack{n \in S \\ 45 \leq n \leq (55-z)}} Seed_{(x,y,z+n)}) / maxNumSeeds) - 1 \leq M * \sum_{\substack{m \in S \\ 15 \leq m \leq 40}} Seed_{(x,y,m)} \quad (4.6)$$

4.2.5 Minimum/Maximum Number of Seeds Per Slice

The design of constraints to enforce a minimum number of seeds per slice helps to alleviate a number of minor issues that could potentially occur with automated planning approaches. MILP models may produce treatment plans containing slices that house a very small number of seeds (specifically zero to two seeds). While having zero to two seeds in a single slice may still lead to dosimetrically sound results, as seeds placed in neighboring slices may adequately supplement the dose distribution for the given slice, such a style of planning is not clinically preferable as it makes poor use of the permissible seed-placement region. Designing constraint (4.8), which globally enforces a minimum number of seeds to be placed in every slice, might help produce solutions that are more clinically acceptable. In connection with the above point, MILP models may produce treatment plans that contain cold spots at the base or apex of the prostate (due to limited dose contribution from neighboring slices of the base and apex). Thus, it might prove beneficial to design constraint (4.10), which locally enforces a minimum number of seeds to be placed at the geometrical extreme ends of the prostate in order to ensure adequate dose coverage. In order to achieve a global minimum number of seeds in all slices, equation (4.7) is developed to keep track of individual slices that house single or multiple seeds.

$$Seed_{(x,y,z)} \leq Slice_{(z)} \quad \forall x \in Q, y \in R, z \in S \quad (4.7)$$

Although the idea of a minimum number of seeds per slice has been mentioned in D’Souza et al. (2001), the authors have not introduced a mathematical formulation of the

constraints required to achieve this rule. Furthermore, the authors' stated reasoning behind such a rule is to “avoid overloading of one of two adjacent slices when the target is of comparable size” (D’Souza et al. 2001), in other words to overcome issues that arise specifically out of their solution process, which is based on iterative optimization of individual prostate slices. Our implementation of the minimum number of seeds rule arises out of the previously mentioned clinical reasons and solution-related trends that we have observed in our own automated planning approach. Despite being less of a concern clinically, equations (4.9) and (4.11) enable the end-user to implement global and local enforcement of maximum number of seeds per slice respectively.

$$\text{minSeedsPerSlice} * \text{Slice}_{(z)} \leq \sum_{x \in Q} \sum_{y \in R} \text{Seed}_{(x,y,z)} \quad \forall z \in S \quad (4.8)$$

$$\text{maxSeedsPerSlice} * \text{Slice}_{(z)} \geq \sum_{x \in Q} \sum_{y \in R} \text{Seed}_{(x,y,z)} \quad \forall z \in S \quad (4.9)$$

$$\text{minSeedsPerSlice} \leq \sum_{x \in Q} \sum_{y \in R} \sum_{\substack{z \in S \\ z=0,55}} \text{Seed}_{(x,y,z)} \quad (4.10)$$

$$\text{maxSeedsPerSlice} \geq \sum_{x \in Q} \sum_{y \in R} \sum_{\substack{z \in S \\ z=0,55}} \text{Seed}_{(x,y,z)} \quad (4.11)$$

4.3 Experimental Design

For purposes of demonstration and continuity, we solve two instances of our brachytherapy model for a single medium-sized (42.5 cc) prostate: the base model and the enhanced model. The base model is obtained directly from Chapter 3 as described in Section 2, and the enhanced model is constructed with additional spatial constraints and data processing techniques presented in this chapter in order to illustrate the further improvement in treatment

plan quality that is achievable. The base model contains the objective function given in equation (1), the complete set of dose calculation constraints (equations (3.2)-(3.7)), dose-volume constraints (equations (3.8)-(3.21)), as well as the spatial constraints given in equations (3.22) to (3.35) (with the exception of equations (3.26) and (3.35)). The enhanced model contains the exact same mathematical formulation (objective function, dose calculation constraints, dose volume constraints and spatial constraints) as the base model, with the addition of equation (4.1), equations (4.3)-(4.4) and two instances of equation (4.10) (to enforce a minimum of 4 and 6 seeds for the base and apical slices respectively) from this chapter. For the base model the seed-placement region is constructed through the global encapsulation method, while for the enhanced model the slice-based local encapsulation method is used. As seen in Table 4.1, an encapsulation margin of 0.50 cm is used for the base and apical slices ($z=0.0$ cm and $z=5.5$ cm respectively), while a margin of 0.10 cm is used for the two mid-gland slices ($z=2.5$ cm and $z=3.0$ cm). The urethral curvature is accounted for in each model, however a single 0.4 cm protection zone is utilized for the base model to protect against seed-placement and needle incursions, while two separate margins are used in the enhanced model (0.40 cm for seed-placement incursion, 0.20 cm for needle incursion). Furthermore, a 0.40 cm rectal protection margin is implemented in both models. Finally, the pubic arch is accounted for in the enhanced model, while it is not in the base model. The grid system utilized in both models is based on a $0.5 \times 0.5 \times 0.5$ cm³ resolution in order to illustrate that clinically viable results are attainable even with the use of a coarse resolution grid for dosimetric calculations.

Table 4.1. Seed-placement boundaries used in slice-based local encapsulation.

	Slice z (cm)	PTV Expansion (cm)	PTV Exp. Rounded (cm)
Base	0.0	0.50	0.50
	0.5	0.42	0.40
	1.0	0.34	0.30
	1.5	0.26	0.30
	2.0	0.18	0.20
Mid-gland	2.5	0.10	0.10
Mid-gland	3.0	0.10	0.10
	3.5	0.18	0.20
	4.0	0.26	0.30
	4.5	0.34	0.30
	5.0	0.42	0.40
Apex	5.5	0.50	0.50

AMPL (Fourer et al. 2002) is used for the formulation and Gurobi 6.5.0 (Gurobi Optimization 2014) is used for the solution of our models on a 20-core ACPI multiprocessor X64-based PC with Intel Xeon® CPU E5-2650 running at 2.3 GHz with 128 GB RAM. The specific reasons behind our choice of modeling language and solver are discussed in Chapter 3. The following parameters are set for each model: **presolve = 10** in AMPL and **mipgap = 0.0001**, **timelim = 3600**, **presolve = 2**, **heurfrac = 0.25** in Gurobi. We provide an extensive explanation for the use of these parameters and settings in Chapter 3. Gurobi stops its branch-and-bound solution process once either the **mipgap** or **timelim** limit is reached. The best current integer solution is then extracted as the final result for each model.

4.4 Results

Treatment plans produced through the base model, enhanced model and an expert planner are compared in Figures 4.11-4.13, and the corresponding post-optimization dose distribution results, calculated in the VariSeed LDR treatment planning system, are displayed in the left,

middle and right columns respectively. The associated isodose legend is given in Figure 4.10. Needle loading and distribution results for the base model, the enhanced model and the manually-produced treatment plan are presented in Figures 4.14-16 respectively. Post-presolve numerical features of the base model and the enhanced model are provided in Table 4.2, while the spatial characteristics of the three treatment plans are given in Table 4.3.

Table 4.2. Numerical features of branch-and-bound solutions.

		Variables		Solution				
Treat. Plan	Constraints	Integer	Total	Obj. Func.	Best Bound	Optimality	Iterations	Time (s)
Base Model	3666	2576	2856	236.150	208.074	11.9%	7250215	3600
Enhanced Model	3216	2444	2722	243.417	213.297	12.4%	18423151	3600

Table 4.3. Spatial characteristics of treatment plans.

		Seeds		Needles			
Treat. Plan	Number	X-Y Adjacency	X-Y Symmetry	Number	Less than 2 seeds/needle	Greater than 4 consecutive seeds	Greatest Retraction
Base Model	98	0/98	88/88	25	0	0	1.5 cm
Enhanced Model	100	0/100	100/100	24	0	0	1.5 cm
Manual	104	0/104	102/104	22	0	0	1.5 cm

It can be seen from Table 4.2 that the base model and the enhanced model resulted in relatively similar optimality gaps, 11.9% and 12.4% respectively. Since the urethral and rectal dose-volume metric values for the base and enhanced model adhere to clinical requirements as presented in Table 4.4, the difference in the objective function values for the two models is not of critical importance.

There were similarities in the spatial arrangements of seeds and needles between the two models. Both the base model and the enhanced model produced treatment plans that contained nearly identical numbers of seeds and needles: 98 seeds and 25 needles for the base model, 100 seeds and 24 needles for the enhanced model. The solutions for both the base model and the enhanced model completely avoided placement of adjacent seeds and consisted entirely of symmetrical distributions of seeds. Furthermore, both models contained needles with a maximum retraction of 1.5 cm, a minimum loading of two seeds and a maximum of four consecutive seeds. In addition, the two models avoided placement of seeds in inaccessible positions behind (superior to) the pubic arch, however this is expected as pubic arch interference was not a major concern for the specific prostate examined. There were also distinct differences between the two treatment plans. The base model contained 10 seeds in Column D, which is highly undesirable as the urethral pathway lies on this column in the prostate studied. Additionally, numerous seeds were placed significantly far from the PTV contour in the base model, specifically at positions b2.5 and e2.5 (coordinates b and e on the x-axis, 2.5 on the y-axis) in the base, and B3.0 and F3.0 in the apex. In terms of needle usage, the base model encountered issues adhering to the maximum seed separation rule. Two needles positioned at C3.0 and E3.0 contained seeds placed at $z=0.5$ cm and $z=5.5$ cm. Thus, the maximum seed separation in these two needles was 5.0 cm, which clearly violates the 4.0 cm limit of the local clinic.

With the implementation of additional spatial constraints and data processing techniques, the enhanced model did not encounter any of the issues that existed in the base model. While all

seeds were placed relatively closer to the PTV in the enhanced model, none were placed on Column D. All needles used in the enhanced model had a maximum seed separation of 4.0 cm. Thus, the enhanced model was able to satisfy all of the technical spatial requirements in a similar manner to the manual plan (which consisted of 98.1% symmetrical seeds, 0% adjacent seeds, no cases of seeds on Column D, no cases of seeds placed significantly far from PTV contours, no cases of needles with single-seed loading, no cases of needles containing greater than 4 consecutive seeds, no cases of needles with greater than 1.5 cm retraction, and no cases of needles with a seed separation beyond 4.0 cm).









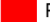





Isodose Legend Gy [% of Prescription Dose]		290.00 [200.00%]		145.00 [100.00%]		72.50 [50.00%]
		217.50 [150.00%]		130.50 [90.00%]		36.20 [24.97%]
		159.50 [110.00%]		116.00 [80.00%]		
Anatomy/Landmark Legend		Prostate		Rectum		Urethra_dosim
		Urethra		PTV		Pubic Arch

Figure 4.10. Legend for treatment plan isodoses and contours.

Table 4.4. Summary dose-volume metric values for the three treatment plans.

Dose Parameter	Objective	Treatment Plan		
		Manual	Base Model	Enhanced Model
PTV Prescribed Dose	145 Gy	145 Gy	145 Gy	145 Gy
PTV V100	> 98%	99.6%	98.3%	98.3%
PTV V150	≤ 65%	64.2%	52.7%	57.5%
PTV V200	≤ 25%	22.8%	21.3%	20.6%
PTV D90	174 – 189 Gy	186.2 Gy	171.2 Gy	176.2 Gy
Urethra D5	< 215 Gy	199.7 Gy	191.1 Gy	190.7 Gy
Rectum D1cc	< 145 Gy	110.1 Gy	95.0 Gy	106.9 Gy


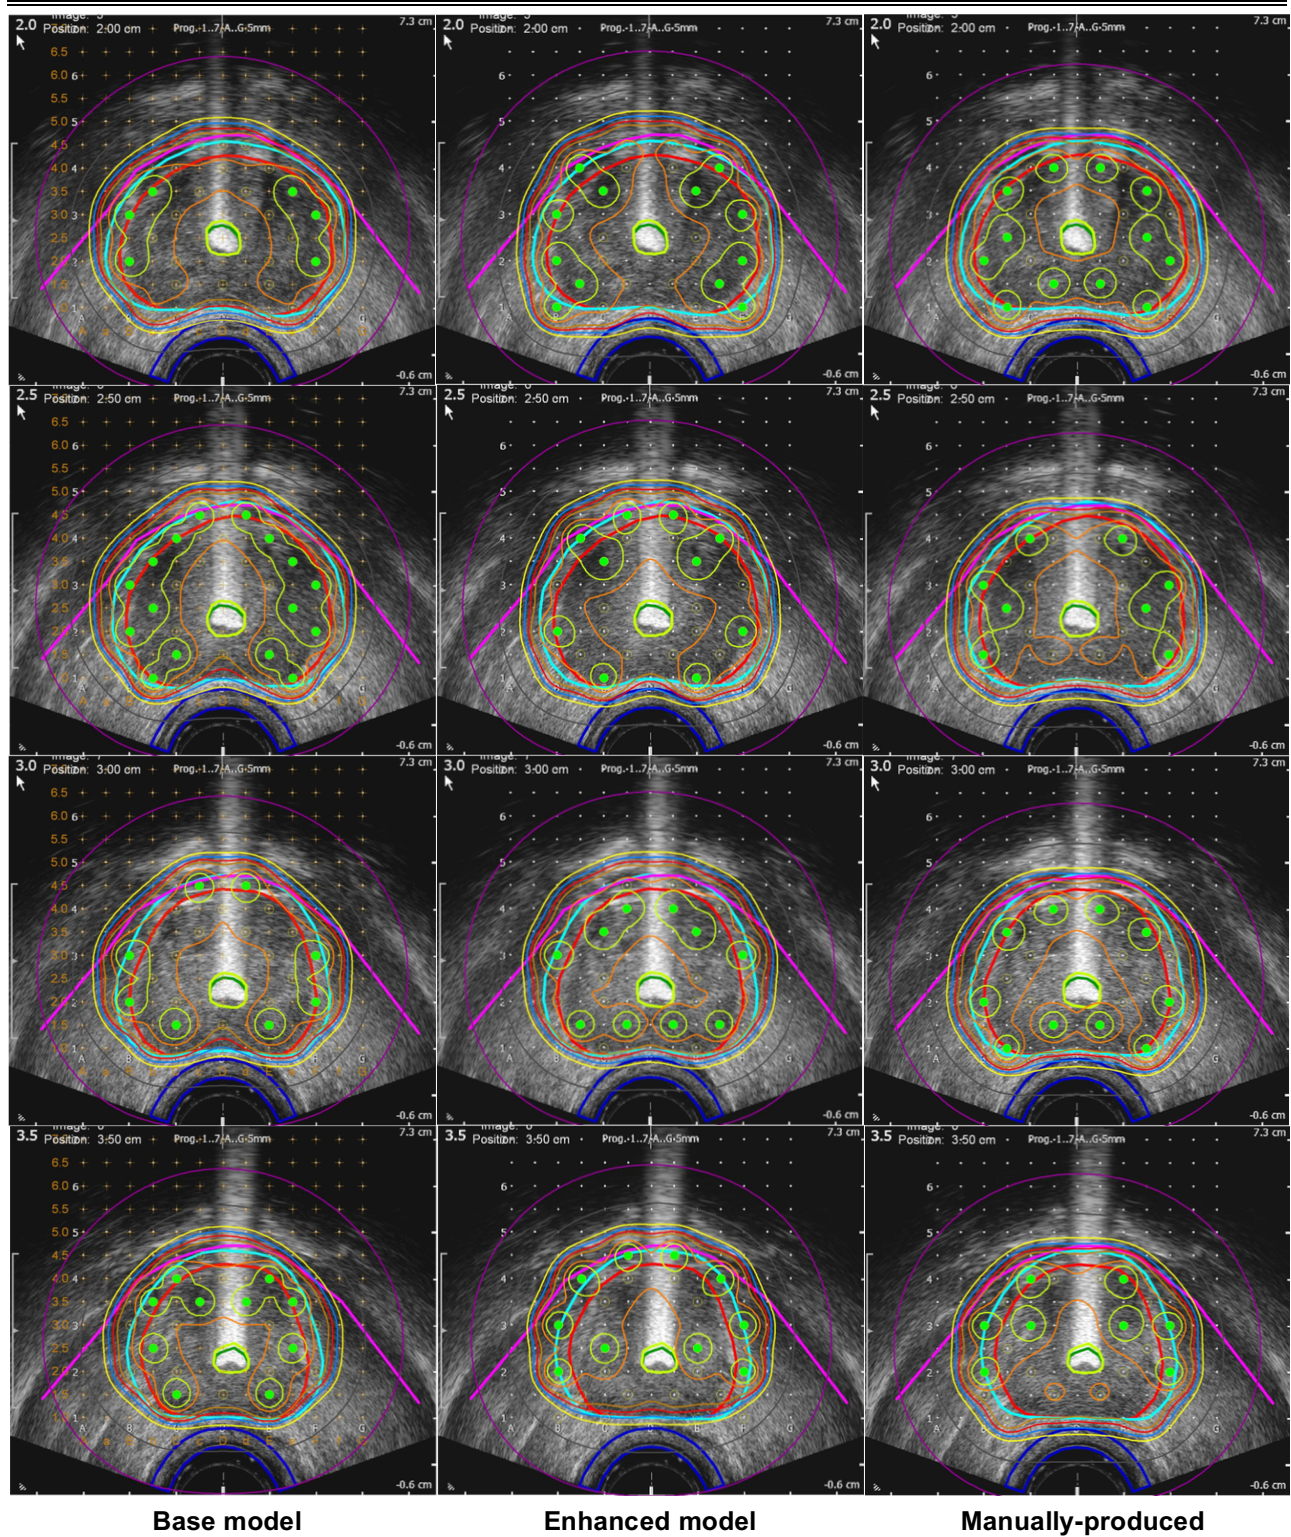


Figure 4.11. Comparison of treatment plans (slices $z=0.0$ cm to $z=1.5$ cm)



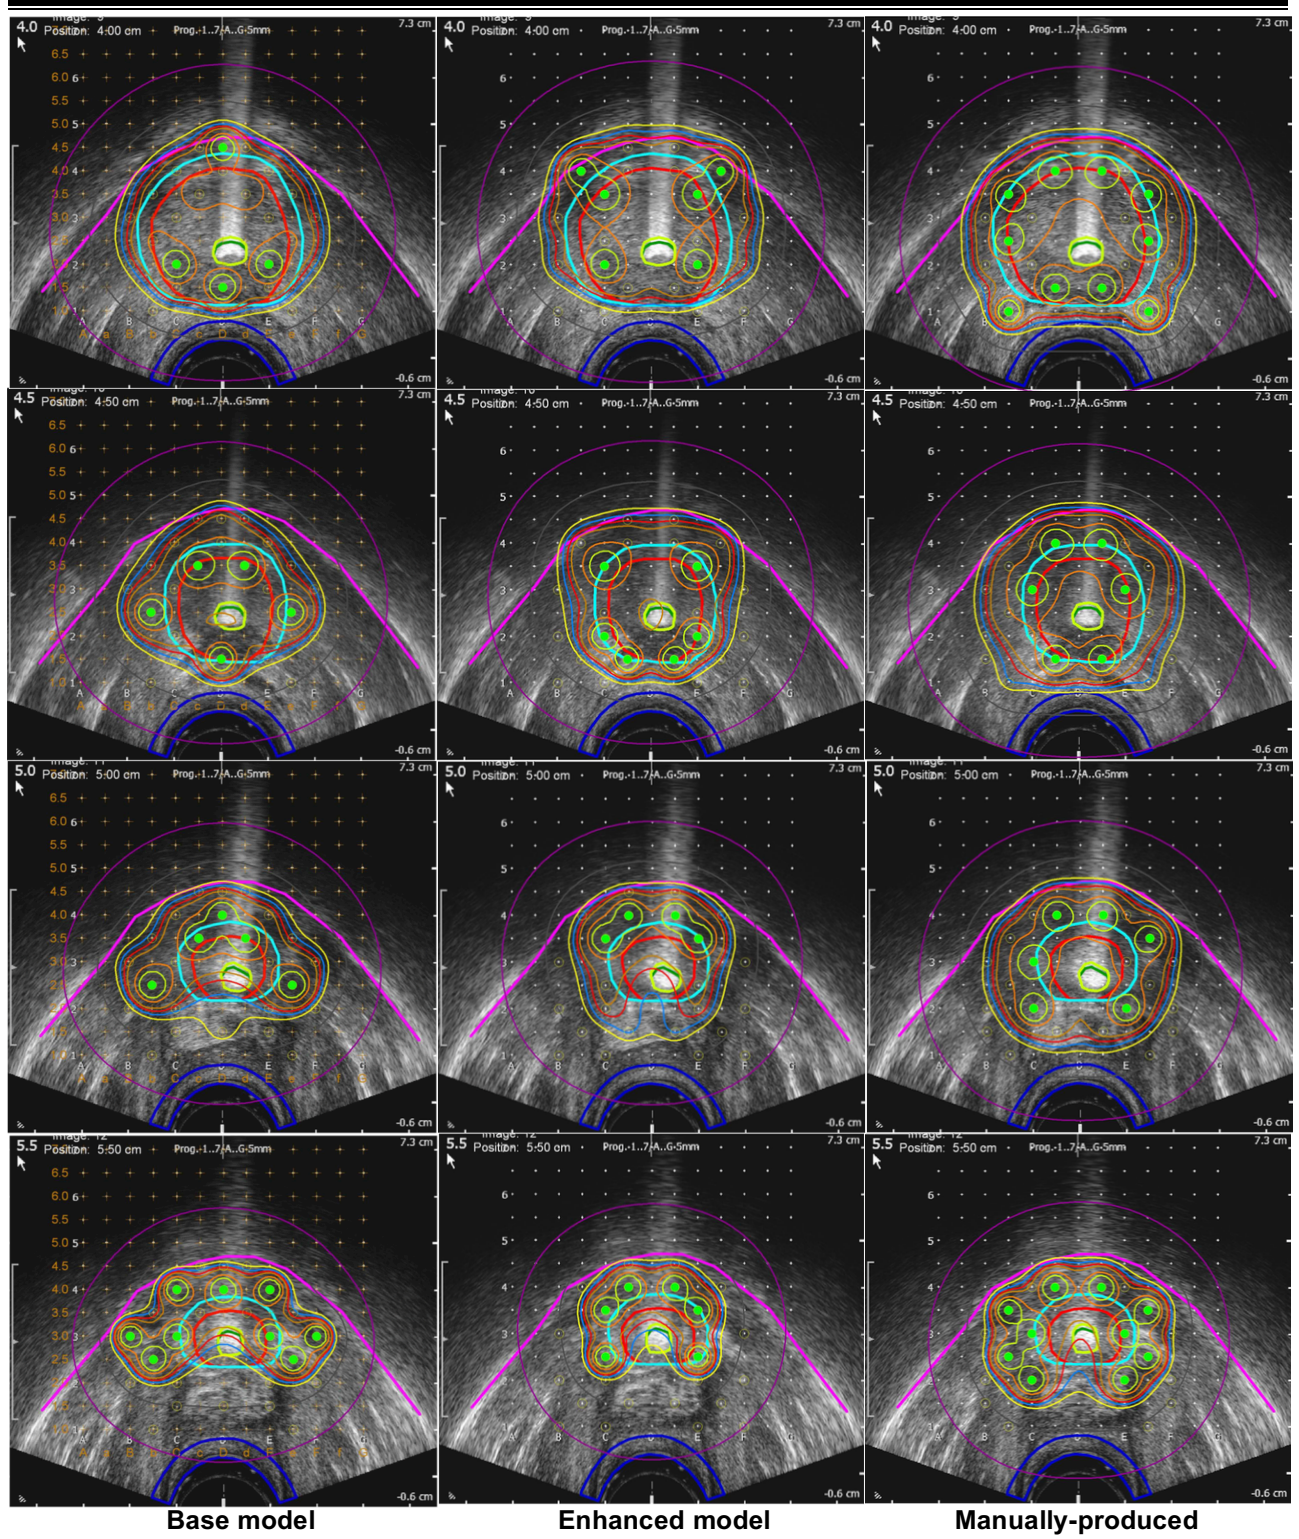


Figure 4.13. Comparison of treatment plans (slices $z=4.0$ cm to $z=5.5$ cm).




Figure 4.14. Needle loading (left) and needle distribution (right) in the base model.

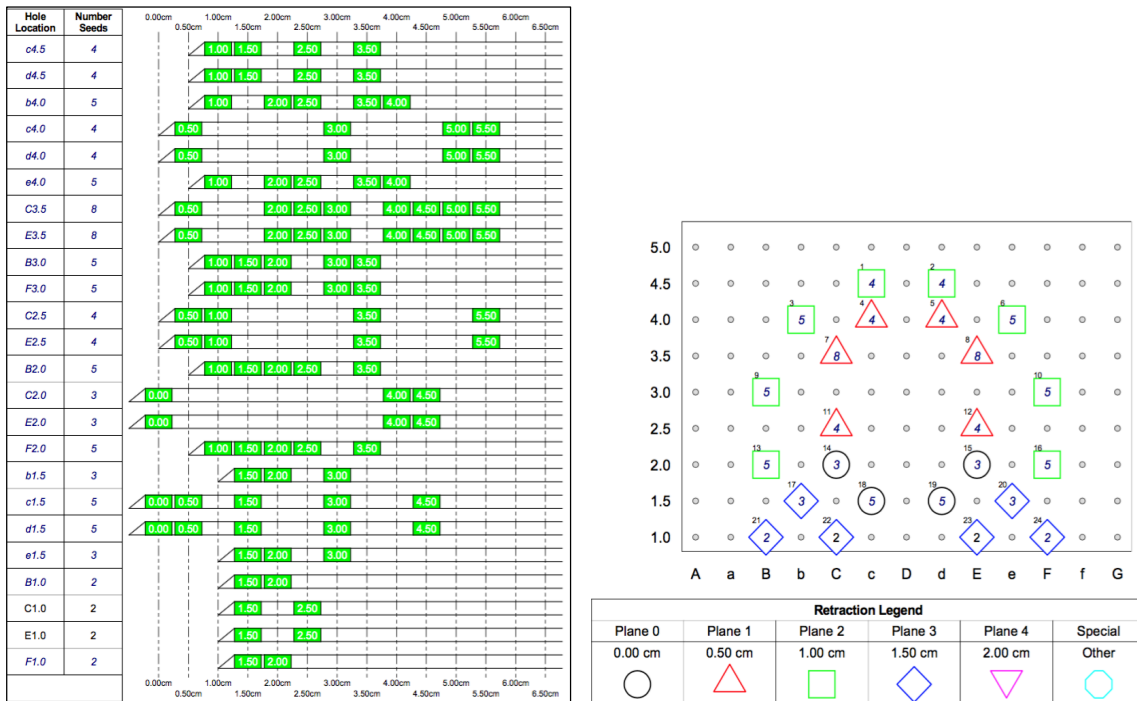
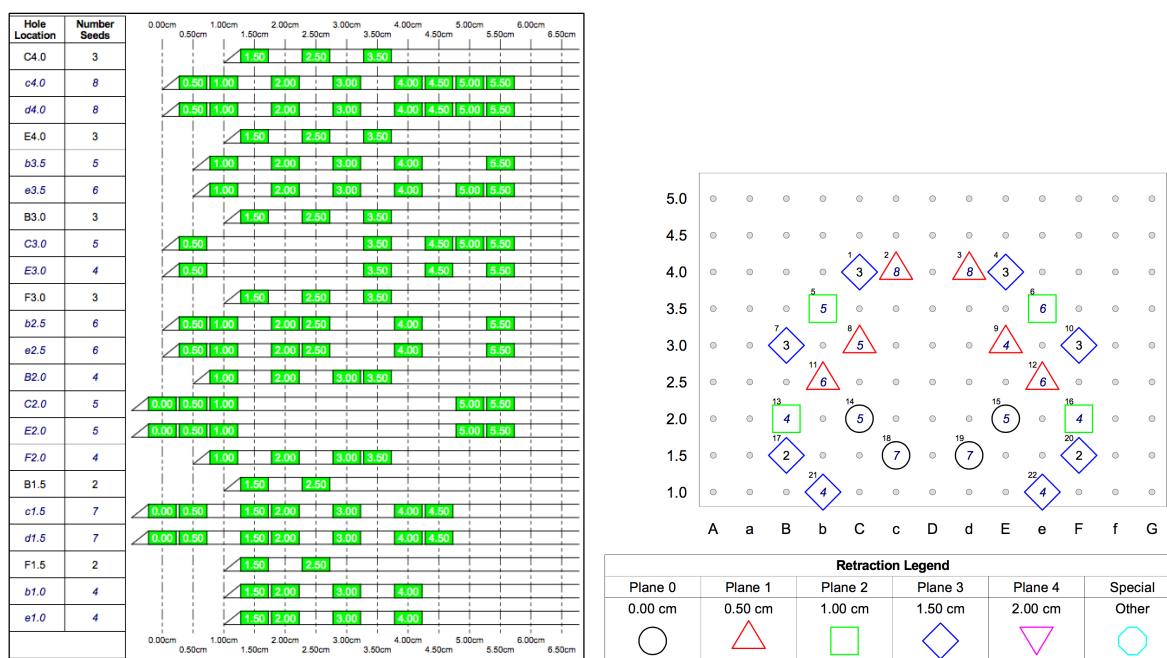


Figure 4.15. Needle loading (left) and needle distribution (right) in the enhanced model.



4.5 Discussion

In the discussion section, treatment plans created through the base and enhanced models are analyzed for their clinical quality by an experienced planner (medical physicist) at the local cancer center. As also observed in Chapter 3, the base model treatment plan exhibits many of the features of a seed distribution that would be deemed clinically acceptable at the Cross Cancer Institute. It meets all but one of the specified dose-volume objectives, only missing narrowly for PTV D90. That being so, further refinements to the base model plan are desirable from a clinical perspective. With reference to Figures 4.11-4.16, these were previously identified to be:

1. Improved dose coverage at the prostate base and apex (positions 0.0, 5.0, 5.5).
2. Improved dose coverage at the posterior surface of the prostate near mid-gland (positions 2.5, 3.0).

3. Seed placement closer to the PTV (positions 0.0, 0.5, 5.0, 5.5).
4. Seed placement avoiding template Column D (positions 0.0-1.5, 4.0-5.5).

The enhanced model plan meets all of the dose-volume objectives and exhibits the above-itemized refinements with one notable exception: dose coverage at the prostate apex (positions 5.0, 5.5) still bears improvement. This localized deficiency is readily observed by a manual planner, but is not specifically addressed by the seed placement algorithm because the dose coverage objective is stated with respect to a global metric, ie. $PTV\ V100 > 98\%$. Rectification of this issue can be achieved by increasing the objective value to $PTV\ V100 > 99\%$. Alternatively, implementing underdosing constraints exclusively for the apical and base slices (with underdosing calculated from the PTV prescribed dose of 145 Gy) may also help improve dose coverage at these locations. In general however, the enhanced model was better able to capture (in comparison to the base model) the planning style underlying the manually produced plan.

4.6 Conclusion

This chapter develops upon the mixed-integer linear programming modeling approach that was introduced in Chapter 3 for solving the interstitial LDR prostate brachytherapy problem. More complex data processing techniques and spatial constraints are introduced in order to address the shortcomings of our previous work and enhance our initial model. These features enable the ability to comprehensively capture the treatment planning style of our local cancer center, which is highly valuable for the practical and seamless integration of our automated planning approach within the pre-operative treatment planning process used there. Through our

results, we verify that even simple dose calculation formalisms and coarse resolution grid systems enable conception of a treatment plan similar to that deemed clinically acceptable can be conceived using our proposed automated planning approach. The treatment plan produced through the enhanced model (presented this chapter) is able to more fully capture the qualities and characteristics of the manual plan created by an expert planner (in comparison to the base model introduced in Chapter 3). Nevertheless, as only a single prostate case was examined in this chapter, further investigation is needed to comprehensively test the ability of our automated planning approach to consistently produce high quality solutions for prostates of varying shapes and sizes. Therefore, real-world utility of will be assessed in a follow-up twenty-patient study. Additionally, the concepts proposed here for low dose rate prostate brachytherapy may also prove beneficial for automated planning in other variants of prostate brachytherapy (high dose rate, medium dose rate and pulse dose rate), as well as in the brachytherapy treatment of breast, cervical, lung, and head and neck cancer.

Chapter 5³

Pseudo High-Resolution Data Sets & Feasibility-Based Modelling

As mentioned in Chapter 3, treatment planning for brachytherapy can be categorized under two general classes: pre-operative and intra-operative planning. In the pre-operative approach, treatment planning is conducted days or weeks prior to the surgical implantation of seeds in a patient. Intra-operative planning involves treatment planning being carried out during the actual surgical procedure itself, while the patient is under anesthesia. The major difference between the two approaches is the amount of time allotted for the formulation of a treatment plan; whereas with pre-operative planning it may be permissible to produce a treatment plan in a matter of hours, the expectation in the intraoperative setting is that treatment plans are produced in a matter of minutes, or perhaps even seconds. The main reason for such a time difference is to minimize the risk for the patient during intra-operative planning: the less time a patient spends under anesthesia while treatment planning is carried out, the less risk (and potentially complicating consequences)

³*A version of this chapter will be submitted for publication: Babadagli, M.E.; Sloboda, R.; Doucette, J. "Mixed-Integer Linear Programming Modelling Techniques for Improving Treatment Plan Solution Times in Interstitial Low Dose Rate Prostate Brachytherapy". Springer, Annals of Operations Research (to be submitted December 31, 2016).*

he is exposed to. Thus, in order to design our automated planning system so that it is useful to experts in both pre-operative and intra-operative planning, it is vital that our system consistently achieves solutions times in the vicinity of minutes or seconds. In this chapter, we will investigate two techniques in hopes of achieving such a feat: 1) the use of pseudo high-resolution data sets, 2) reformulation of our currently optimality-based model into a feasibility-based model.

5.1 Pseudo High-Resolution Data Sets

Despite leading to relatively accurate results with reasonably quick solution times in previous chapters, low-resolution anatomical data has its shortcomings. The main disadvantage of using low-resolution data is the inability to precisely delineate anatomical structures. For example, despite being a relatively small structure, the urethra is considered to be the most critical organ-at-risk in LDR brachytherapy. Therefore each individual voxel associated with the urethra is critical in correctly defining its borders. However, due to the characteristic 0.5 cm voxel spacing of the low-resolution system, we have fewer and larger voxels at our disposal to discretize anatomical structures. Thus, it is not uncommon to run into difficulties in realistically representing the shape of the urethra, as depicted in Figure 5.1. Furthermore, there exists a significant dose fall-off associated with radioactive sources the further one moves away from the source. In order to appropriately calculate the dose distribution, it is vital to accurately account for dose contributed by a source to voxels located at distances smaller than 0.5 cm. This is not possible in the low-resolution system since the minimum distance r for which the radiation distribution can be calculated is 0.5 cm. This crude dose calculation system, coupled with crude delineation of anatomical structures, as previously seen in Figure 3.5, may lead to difficulties in achieving

sufficient dose coverage at the PTV periphery and also while avoiding the critical organs-at-risk. Perhaps the clearest example of such a scenario is evident in the treatment plans produced in Chapter 3. Despite achieving reasonably acceptable results, we noted some differences between PTV D90 dose limits set in our optimization models, and the final values observed during post-optimization dosimetric analysis using VariSeed. In other words, there were slight deficiencies in the accuracy of our dose coverage calculations in our coarse resolution models.

Crude dose calculation and structure delineation in the low-resolution system may also prove to be problematic for the bladder neck. Without the use of protective margins, some radioactive sources may be placed too close to the urethra during automated planning due to the low-resolution depiction of the organ. Such seed-placement is clinically unacceptable, as the high levels of radiation emanating from seeds placed too close to the bladder neck may cause stiffening of this region and lead to post-operative incontinence in patients.

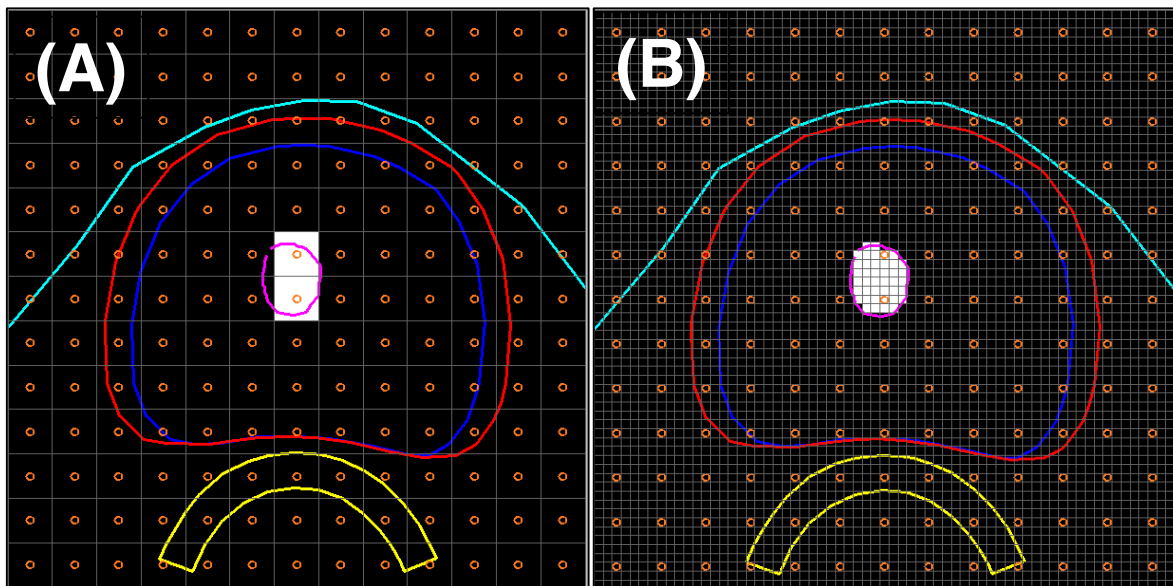


Figure 5.1 Comparison of low-resolution (A) and high-resolution (B) discretization of the urethra.

Therefore, the use of high-resolution anatomical data is ideally preferred to low-resolution data in treatment planning for LDR prostate brachytherapy. The high-resolution system is based on a $0.1 \times 0.1 \times 0.5$ cm grid (in the x - y - z coordinate system), as opposed to the $0.5 \times 0.5 \times 0.5$ cm grid of the low-resolution system. The improvement in delineating the urethra using high-resolution data is noticeable in Figure 5.1. However, there exists a significant problem in making the transition to high-resolution data, and that is the computational difficulty of solving the resulting large-scale linear programming models that involve a large number of binary variables. Whereas MILP models using low-resolution data may consist of anywhere from 2,000-3,000 binary variables for a typical medium sized prostate, these numbers increase to 50,000-60,000 binary variables in the high-resolution system. Due to the specific formulation of our model presented in Chapter 3, each voxel (data-point) that represents an anatomical structure possesses associations with often multiple binary dose-related variables. These binary variables are necessary for incorporating clinical dose-volume metrics into our mathematical model, and for achieving high reliability in dose-related results. However, as mentioned previously, large numbers of binary variables lead to large-scale optimization models that are exponentially harder to solve than if the problem consisted only of integer or continuous variables. The end result is that such large-scale problems are intractable and even finding an initial feasible solution using solvers based on the simplex and branch-and-bound algorithms (such as Gurobi and CPLEX) may turn into a nontrivial task. In our models, the true culprit of the spike in binary variables is PTV voxel data, as the gland constitutes as the largest anatomical structure in our data set and therefore is associated with the largest number of binary variables. The rectum also plays a part in the increase of binary variables,

but to a lesser extent than the PTV data. The increase in urethra voxel count may be ignored, in parts due to the much lower voxel tally but more so due to the criticality of the organ.

Various mathematical decomposition methods have been developed in the past to simplify the solution of large-scale linear programming models, such as Dantzig-Wolfe Decomposition (Dantzig & Wolfe 1960), Bender's Decomposition (Geoffrion 1972), and Cross Decomposition (Van Roy 1986). These techniques generally exploit a single, specific mathematical structure within the constraint matrix of a linear programming model. In doing so, these techniques are able to pacify complicating, or hard-to-solve, constraints or variables so that the original problem is decomposed into a simpler one and a standard implementation of the simplex or branch-and-bound algorithm is able to solve such large-scale models with increased relative ease. However, due to the plethora of structures present in the matrix structure of our LDR brachytherapy model, these methods (individually or in any given combination) were deemed not to be applicable or practical for the simplification of our mathematical model.

To overcome these clinical and computational issues, we have developed a unique work-around called "pseudo high-resolution" data, based on the simple concept of data sampling as displayed in Figures 5.2 and 5.3. This technique enables us to effectively decrease the size of our large-scale model (containing high-resolution data) and transform it into a medium-scale model (containing pseudo high-resolution data). The premise behind such a transformation is to reduce computational loads encountered during the solution of our model, and therefore decrease solution times. As outlined in Algorithm 5.1, our technique involves applying filters to our high-resolution discretized anatomical data during the pre-processing stage of our input data. Each filter begins as

an “unprocessed filter”, which is essentially a 0.1x0.1 cm resolution 65x65 grid of randomly distributed binary data, populated using a relatively simple approach defined in Steps 2 and 3 of Algorithm 5.1. A fully populated low-resolution (0.5x0.5 cm) grid is then incorporated into the unprocessed filter, resulting in the “processed filter”, which is applied to each anatomical structure. This approach ensures that we are able to not only capture every detail (data-point) originally provided by the low-resolution system, but are able to further supplement or enhance these details through the sprinkling of random high-resolution data points. The incorporation of the low-resolution grid also ensures that our anatomical data sets contain all of the voxels located at potential seed-placement positions (which is required for our dose calculation constraints to function properly).

```
1  1. Start with the highest numbered (apical) slices of discretized (fine-resolution) processed
2  PTV region (ProPTV[k]) fully populated coarse-resolution grid (CRGrid[k]), and three empty
3  fine-resolution grids, one for the random number matrix (RandNum[k]), one for the sampling
4  filter (SampFilt[k]) and one which will house the sampled PTV region (SampPTV[k]).
5  2. k = apical slice coordinate
6     WHILE k ≥ 0:
7       b. FOR each voxel RandNum[i,j,k] in slice RandNum[k]:
8         Assign RandNum[i,j,k]= RANDOM(0...1)
9       c. Perform k=k-1.
10 3. k = apical slice coordinate
11   sampRate = 0.10
12   WHILE k ≥ 0:
13     a. FOR each voxel SampFilt[i,j,k] in slice SampFilt[k]:
14       IF RandNum[i,j,k] ≤ sampRate:
15         Assign SampFilt[i,j,k] = 1
16       ELSE:
17         Assign SampFilt[i,j,k] = 0
18     b. Perform k=k-1.
19 4. k = apical slice coordinate
20   WHILE k ≥ 0:
21     b. FOR each voxel SampPTV[i,j,k] in slice SampPTV[k]:
22       Assign SampPTV[i,j,k]= SUM(CRGrid[i,j,k], SampFilt[i,j,k]) * ProPTV[i,j,k]
23     c. Perform k=k-1.
24 5. k = apical slice coordinate
25   WHILE k ≥ 0:
26     FOR voxel SampPTV[i,j,k] in SampPTV[k]:
27       IF SampPTV[i,j,k] > 0:
28         Assign SampPTV[i,j,k] = 1.
29 6. End with finalized (pseudo high-resolution) processed PTV region indexed as T[i,j,k]
30
```

Algorithm 5.1. Algorithm for creating pseudo high-resolution PTV data.

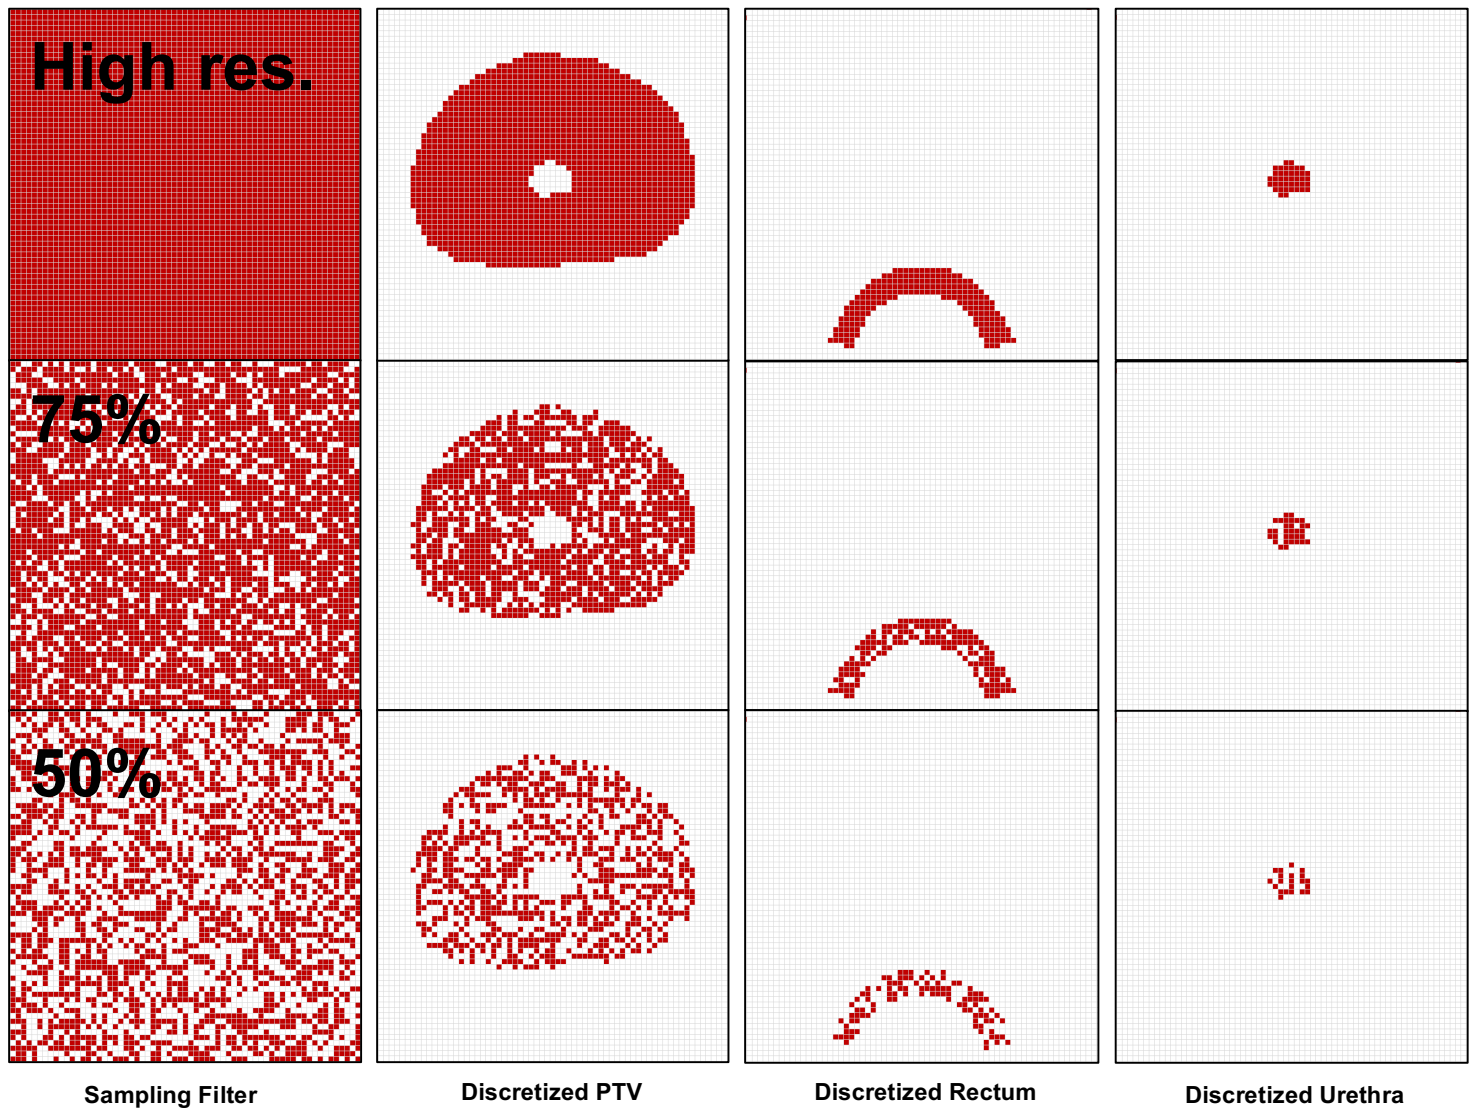


Figure 5.2 Comparison of high-resolution (100% sampling) and pseudo high-resolution data (75 and 50% sampling) for the PTV, rectum and urethra.

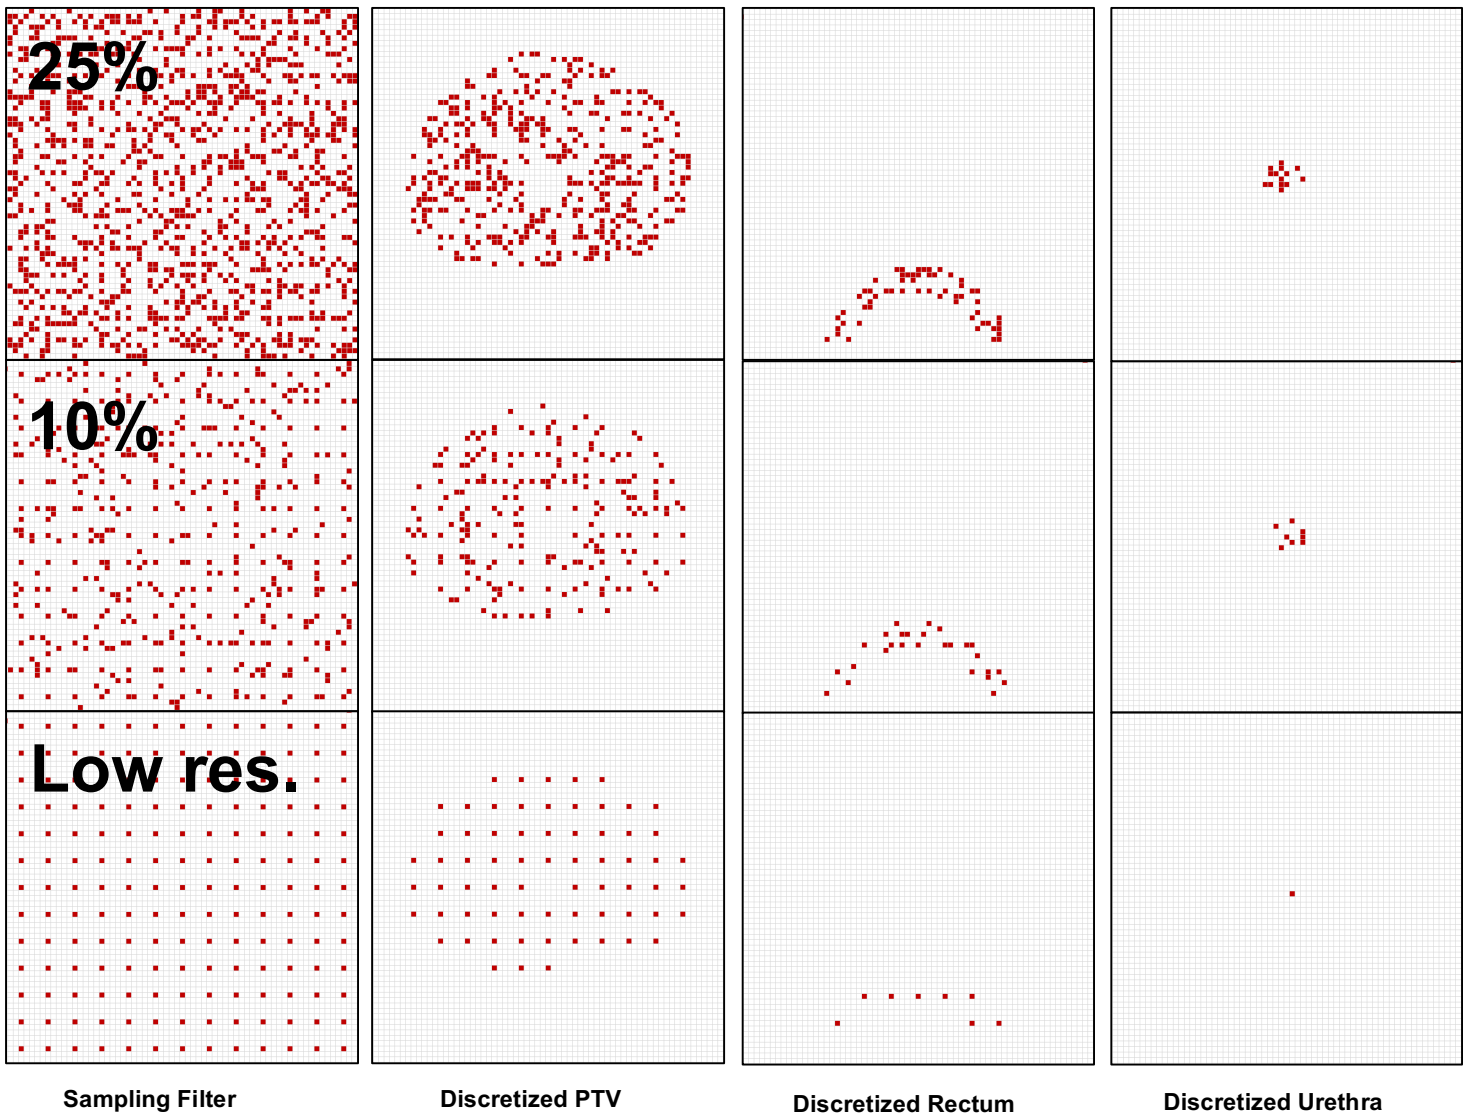


Figure 5.3 Comparison of pseudo high-resolution data (25 and 10% sampling) and low-resolution data.

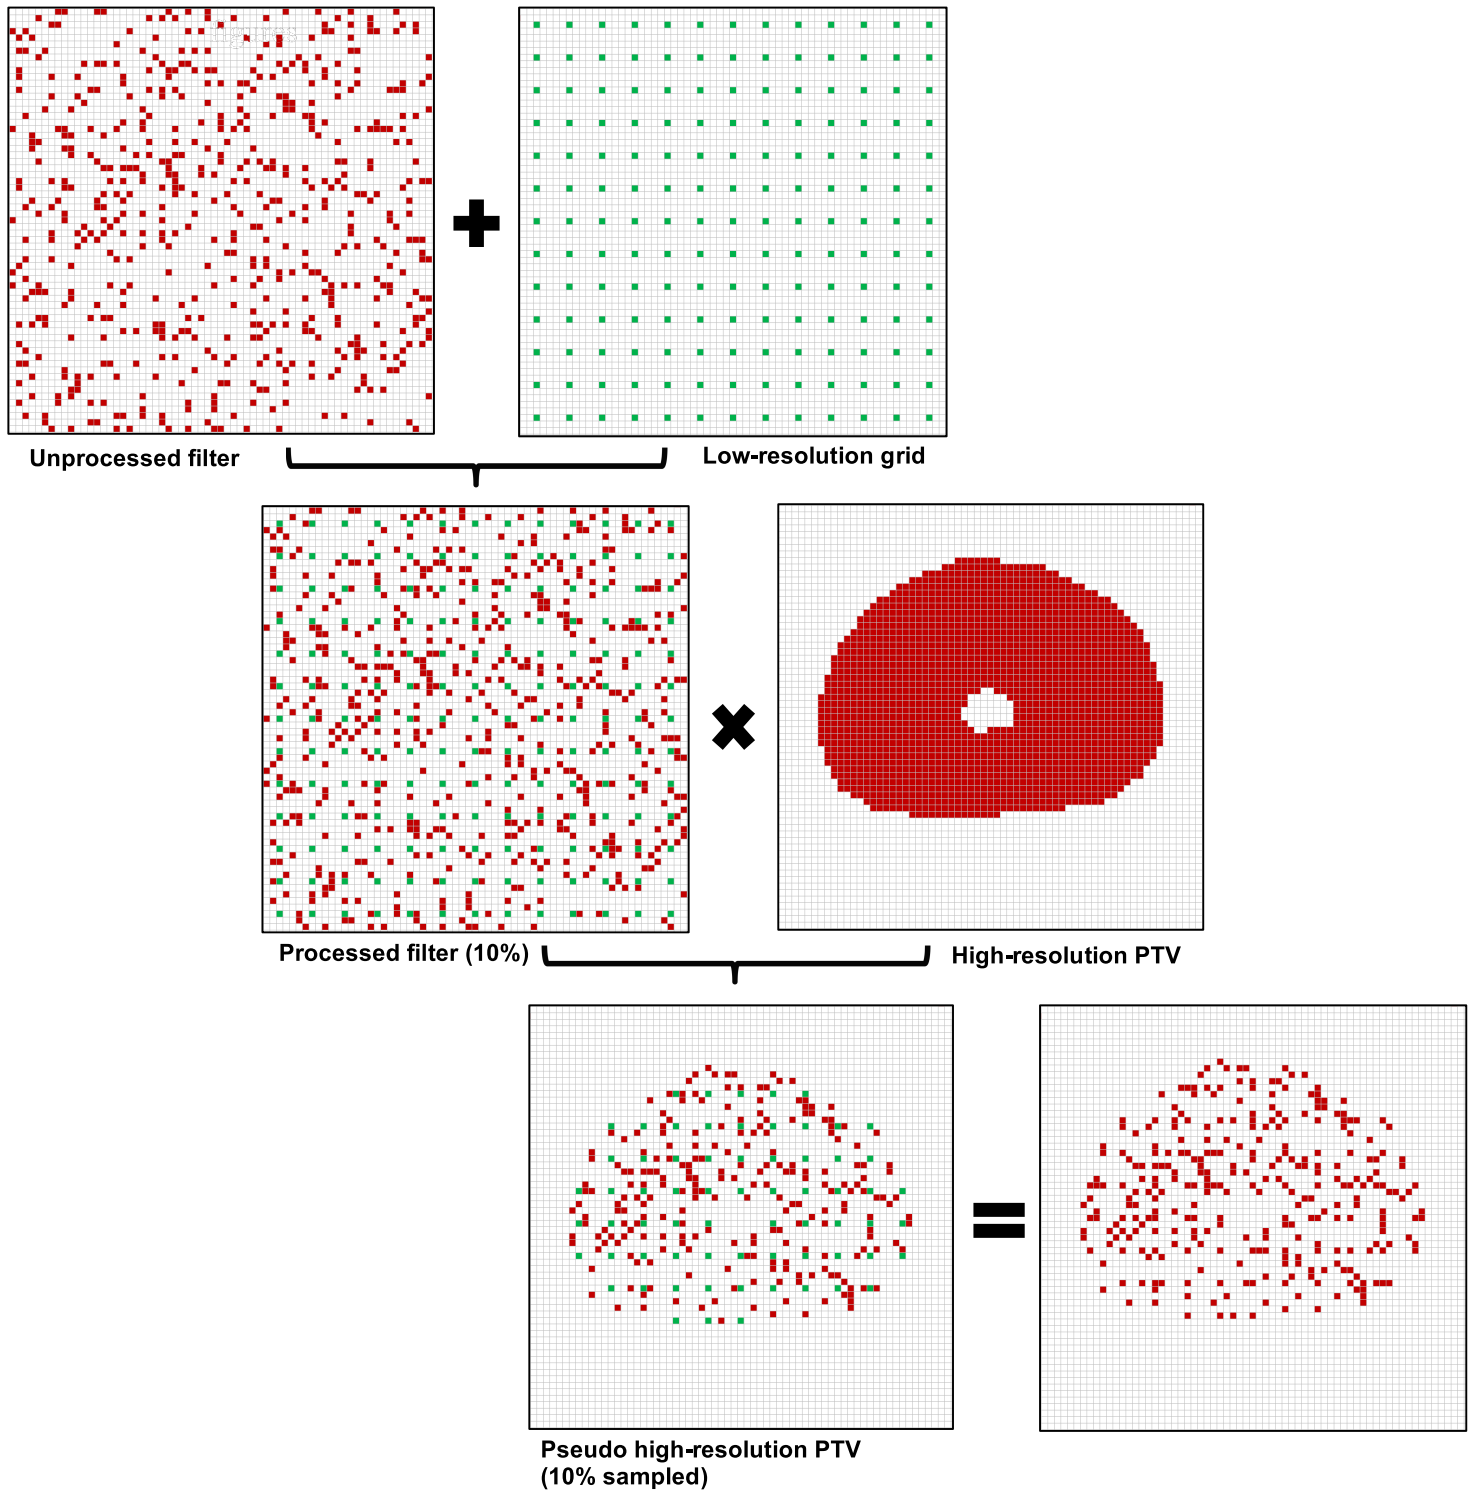


Figure 5.4 Visualization of the process for creating pseudo high-resolution PTV data.

A sampling rate is specified to control the desired ratio of 1s and 0s in the unprocessed filter and therefore control the sampling of the final discretized anatomical structure. A unique binary filter is applied to each slice of anatomical data in order to avoid introducing any form of bias or pattern into our data. A comparison of high-resolution (100% sampled), pseudo high-resolution (75%, 50%, 25%, 10% sampled) and low-resolution data is given for the PTV, rectum and urethra in Figures 5.2 and 5.3. Our approach enables anatomical data to be represented within a new fuzzy zone of the data spectrum that exists between low-resolution and high-resolution data (which are the two extreme ends of said spectrum and the sole focus of previous brachytherapy optimization literature). As seen in Figure 5.4, our pseudo high-resolution data sets are produced by combining a sense of order (present in low-resolution data sets) with elements of disorder (present in unprocessed filters). In the model-ready data sets, voxels are both randomly and deterministically distributed within a given anatomical structure at varying separation distances, rather than the structure being uniformly populated solely with voxels adhering to a single separation distance of 0.1 cm (as in high-resolution data) or 0.5 cm (as in low-resolution data). Such a scenario reveals the power of sampling, as we see that despite containing roughly a quarter of the information provided in its high-resolution counterpart, the 25% sampled PTV in Figure 5.3 is still distinguishable and its associated data still reflects the essential characteristics of the original high-resolution data. The reason for the latter point is that with pseudo high-resolution data, the radiation distribution is still required to be calculated down to very small distances r from a source (specifically 0.1, 0.2, 0.3, 0.4 and 0.5 cm), similar to what occurs in high-resolution data. This is not the case with low-resolution data, where the minimum distance r for which the radiation

distribution can be calculated is 0.5 cm. The advantage of our approach is that a significantly fewer number of these dose calculations (as well as other decisions) have to be made during the solution of our model (in comparison to the use of high-resolution data), and this eases the computational loads encountered during the solution process without compromising solution quality (since we are still capturing the essential characteristics of each structure in a manner similar to high-resolution data). Pseudo high-resolution data sets enable us to effectively transform our originally 30,000-85,000 binary variable large-scale MILP models (ranging from small to large prostates respectively) down to medium-scale models consisting of 3,500-10,000 binary variables.

As evident, our pseudo high-resolution data may be moved to either end of the data spectrum (low-resolution or high-resolution) depending on the sampling rate utilized. Thus, there exists a tradeoff in this system, as too low a sampling rate may potentially lead to less accurate results due to loss of detail in anatomical structures and too high a sampling rate may not offer enough of a reduction in computational load. After trialing various values, sampling rates of 10% for the PTV, 30% for the rectum and 100% for the urethra (as calculated based on the processed filters) were deemed to provide the most appropriate balance between computational load and preservation of anatomical detail. One of the main advantages of our unique work-around is that pseudo high-resolution data is a modelling technique that is algorithm-independent. In other words, although we are primarily interested in the solution of our model through a combination of the simplex and branch-and-bound algorithms in this thesis, pseudo high-resolution data sets may be used within any deterministic, stochastic or heuristic solution paradigm as a practical method to decrease the size and solution times of large-scale models.

5.2 Feasibility-Based Modelling in LDR Brachytherapy

5.2.1 Suitability of Feasibility-Based Models

In many cancer centers, including the Cross Cancer Institute in Edmonton, Alberta, treatment planning is carried out pre-operatively through a manual approach (Liu 2014). After anatomical contour data is collected via TRUS imaging, an expert planner (medical physicist) starts the treatment planning process. As mentioned previously, an expert planner begins with an initial arrangement of seeds, and modifies this arrangement iteratively in order to improve the dosimetric quality of the plan (while also adhering to the spatial or geometric seed-placement rules that govern the expert's planning style). This iterative improvement process eventually leads to a final seed distribution that is both spatially and dosimetrically sound. This particular stage of the process may be considered as the point of feasibility. With manual planning, this particular solution is suitable enough to be given to an oncologist for further review. In other words, the expert planner does not continue the iterative heuristic process in order to reach a point of optimality, in hopes of finding the single best (optimal) solution out of possibly millions of feasible ones. This scenario will soon become highly relevant in the rest of this chapter.

There are a number of instances in which modifications may be made to this unapproved (finalized) treatment plan, as seen in Figure 5.5. In regards to the pre-operative manual planning approach, this amounts to roughly three instances after the expert has finalized his or her plan. Once the expert planner is satisfied with the dosimetric and spatial qualities of the finalized seed distribution, the treatment plan is transferred to an oncologist for review. This represents an instance in which preemptive modifications may be suggested by the oncologist for the still

unapproved treatment plan. Once oncologist-suggested improvements are implemented by the expert planner, the plan is re-reviewed and finally approved by the oncologist. Radioactive sources are then ordered for the scheduled operation.

During the brachytherapy operation itself, there are two additional instances in which modifications (voluntary and involuntary) may also be made to the approved treatment plan. The most obvious, and unavoidable, modification to the approved treatment plan involves placement inaccuracies associated with the radioactive sources. Although a skilled operative oncologist can generally place most seeds very close to their suggested coordinate locations, implantation inaccuracies do occur due to a combination of imaging uncertainties, prostate deformation during needle penetration as well as gland swelling due to trauma. In order to correct minor changes in the dose distribution due to implantation inaccuracies, subsequent seeds may be implanted in slightly modified positions in relation to their original positions (adaptive modification). After all seeds have been implanted, a final check is carried out during which additional seeds that are not part of the original plan may be implanted to alleviate cold-spots in the prostate (corrective modification). Thus, seed placement modifications are not uncommon, and a single plan may go through several revisions during the treatment planning stage and the brachytherapy operation itself.

In practice, our automated planning system could be utilized to produce a treatment plan in place of the expert planner as indicated in Figure 5.5. This treatment plan would initially be reviewed and finalized by the expert planner (with possible refinements). The unapproved treatment plan would then be transferred to an oncologist for further review and approval,

following the previously described process. Assuming that the initial solution produced by our automated system is globally optimized based on a given, fixed set of clinical goals, any changes made to this plan (however minor or major) would theoretically pull the solution out of optimality. The opportunities that exist for potential modifications by clinicians (four with the inclusion of the expert planner review) increase the likelihood that any given optimal solution will be progressively reshaped in numerical terms (even though these solutions may be progressively improved in a clinical sense). Such a scenario entirely defeats the purpose of numerical optimality, and raises questions in regards to whether spending computational time and resources to search for optimal solutions in LDR brachytherapy is truly necessary or appropriate.

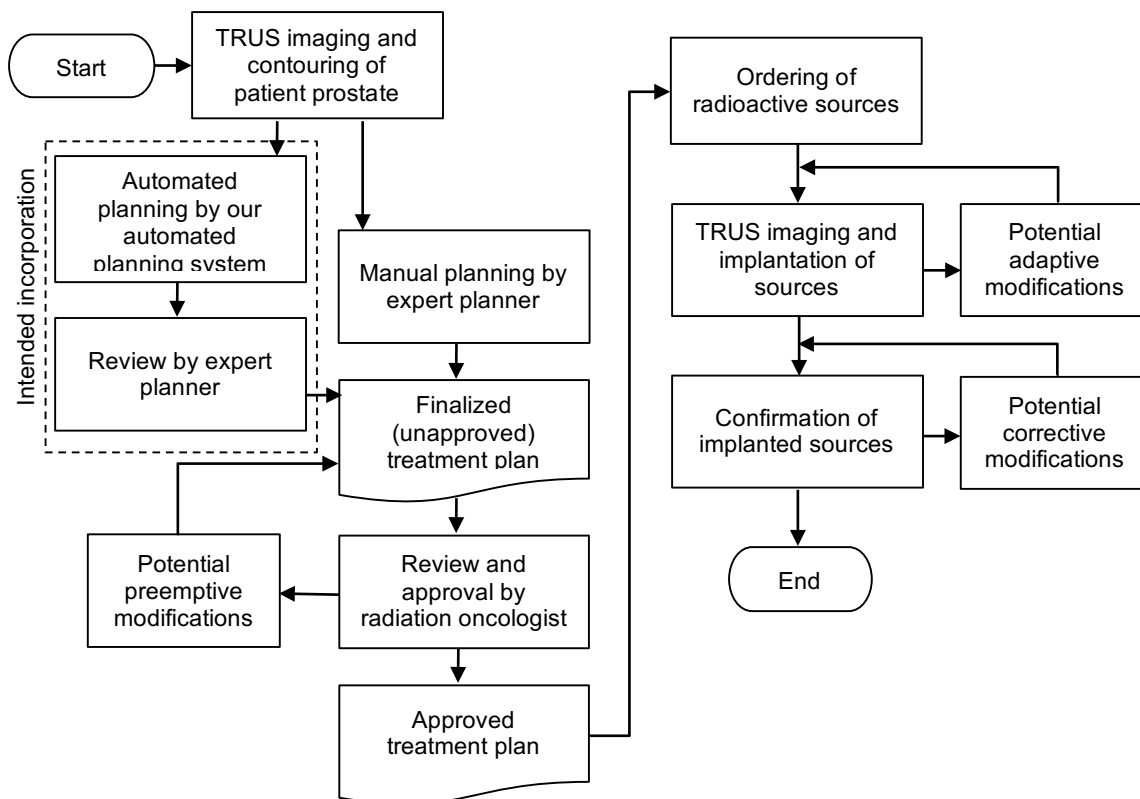


Figure 5.5 Sample process flow outlining treatment planning and seed

The objective function is the primary tool employed in the quest for optimality. As mentioned in Chapter 3, the decision of selecting which variables to include in the objective function is essentially a clinical one, and can be coordinated in a manner that accommodates local clinical practice. As such, most of the past automated planning approaches have been formulated as single or multi-objective problems seeking to achieve single or multiple clinical goals. However, it needs to be questioned whether such objective function formulations are truly leading to favourable clinical outcomes. We say this because what may be considered numerical achievements on the optimization modeling side of the brachytherapy problem may not translate to clinical achievements in the eyes of clinicians.

To be more specific, the usual approach until now has been to treat the brachytherapy problem as one that strives for an optimal solution. The constraints of a linear programming problem define the “playing field”, or the limits within which a solver should search for a solution, and therefore provide general guidance towards the final solution. On the other hand, the objective function has a much more precise and direct influence on the form that the final solution will take.

For example, one of the main goals set for objective functions in previous brachytherapy literature has been to minimize either the number of needles or the number of seeds in a particular treatment plan. While numerically this may be an important goal, perhaps to showcase the power of using automated planning in brachytherapy as opposed to manual methods, clinically it may prove to carry less weight. Many factors are involved in rendering a treatment plan as a high quality one and while keeping the number of needles and seeds at reasonable levels may be one of these factors, the treatment plan with a lower number of needles and seeds will not necessarily be

regarded as a more clinically acceptable one. Usually, a plan with perhaps a higher number of seeds and needles but one that achieves better dose coverage throughout the PTV may be considered to be clinically superior. This scenario captures one of the main potential pitfalls of using automated methods to solve clinical problems.

This brings us to the objective function that we introduced in equation (3.1). As explained in Chapter 3, for the purposes of utility and simplicity, this objective function was formulated to minimize the average dose delivered to urethra and rectum voxels. However, although engineers may consider achieving lower radiation levels in the urethra and the rectum as a favorable numerical goal, this may not necessarily translate to an ideal goal for medical experts. Clinical goals for the organs-at-risk are a bit more complex than simply minimizing the average dose to the urethra and the rectum. For example, at the Cross Cancer Institute, the brachytherapy team actually desires urethral V100 to be close to 100% so that any peri-urethral disease is adequately treated, and urethral V150 to be close to 0% to limit complications associated with radiation treatment. So once again, while a solution with a urethral V100 below 100% might lead an optimization solver to conclude that it has produced a better numerical solution than what was achieved manually by medical experts, such a solution might be viewed as sub-par in the eyes of medical experts based on the reasoning given previously. Usually dose-volume constraints are designed to be tight or specific enough to guide the solution towards clinically acceptable radiation levels, so much so that using an objective function to minimize average dose to the urethra and rectum may prove to be too crude a dose-management method in comparison to the use of dose-volume constraints.

This final point reveals one of the main reasons for the frequent utilization of a minimizing objective function to control dose in previous optimization literature, and also illustrates why such an objective function is not a necessity in our MILP model. As mentioned in Chapter 3, many of the past studies have adopted what we refer to as “crude dose-management” methods that are not designed to explicitly incorporate, within their MILP models, the dose-volume metric criteria used by clinicians. Such “crude dose-management” methods involve underdosing or overdosing constraints (containing continuous dose-related variables), through which the goal is to minimize underdosing of the target organ or overdosing of organs-at-risk in the objective function. In other words, an objective function is almost a necessity in order to manage dose within these models. Through the use of dose-volume constraints, we have adopted the “fine dose-management” method. This method is more likely to produce treatment plans that align with the dose-volume metric criteria used by clinicians during post-optimization analysis. Thus, an objective function that minimizes under or overdosing of tissues does not present itself to be a necessity in our case.

Considering the extensive influence of the objective function on the final solution, for an item to be included in an objective function, it needs to possess overbearing importance over other requirements that also need to be satisfied in a MILP model. In other words, whereas variables included in linear constraints are permitted to take on a range of predetermined values, the only truly ‘permissible’ value for a variable included in a minimizing objective function is the value for which the objective function reaches its lowest attainable value. While a number of clinical goals were initially considered for the objective function of our model (such as minimizing the number of seeds and needles, or minimizing average dose to the urethra and rectum), the brachytherapy

team of our local cancer clinic did not identify any of these objectives to be of overbearing criticality to the quality of treatment plans. This means that it is acceptable for many of the clinical requirements, in our model, to take on a range of permissible values, as long as the dose-volume requirements are satisfied and the final spatial arrangement is clinically sound. When considering the nature of the LDR brachytherapy problem, this renders the use of a generic objective function unnecessary and in this case justifies the notion that a feasible solution may be equally acceptable to an optimal one. Therefore, it makes logical sense to transform the end-goal of our brachytherapy problem from one of optimality to one of feasibility.

One final motivating factor is the lack of investigation into the suitability of feasible solutions for LDR brachytherapy in past optimization literature. While there have been a couple of instances where feasibility-based models have been mentioned in literature (Meyer et al. 2003), these cases were not followed by further inspection or analysis. For this reason, we deemed it suitable to investigate feasibility-based models for LDR brachytherapy in depth in this chapter.

5.2.2 Various Methods to Obtain Feasible Solutions

There exist numerous methods to seek an initial feasible solution in our MILP model. The first technique of interest involves maintaining the objective function presented in equation (3.1) and solving our model as was done in preceding chapters (a solution strategy in which a cascade of feasible solutions eventually leads to a final optimal solution). However in this case, the first integer solution that is found during the branch-and-bound search is extracted as our feasible solution (irrespective of its optimality gap).

The second technique involves transforming the objective function into a constant. In our case, a constant value of “0” is utilized for our objective function. The solution that is output at an optimality gap of 0.00% (also happens to be the first integer solution) is extracted as the feasible solution of our model.

While the previous two techniques are more obvious methods of obtaining feasible solutions, the third technique, which is the central point of this chapter, involves restructuring the mathematical formulation of our linear programming model. The focus of the problem is shifted from searching through a cascade of feasible solutions that eventually result in an optimal solution as was done in the original model, to searching through a cascade of infeasible solutions that eventually result in a feasible solution in the restructured model. The solution process for the original and restructured models are outlined in Figure 5.6. It should be noted that while the final solution obtained for the restructured model is still technically the optimal solution of this particular formulation of the model, such an optimal solution behaves as a feasible solution for the original model. This is accomplished by first introducing artificial variables to the left-hand side of hard-to-solve or difficult constraints (constraints that normally cause an optimization solver to stall or take a considerable amount of time to progress throughout its branch-and-bound search). These variables essentially act as slack variables and they have opposite signs to the right-hand side of the constraint they are added to. Such artificial variables cause hard constraints to become relaxed, or violated, so that these constraints become easier to solve. Secondly, the original objective function is replaced with one that contains a summation of the said artificial variables. This new objective function is structured to minimize the values of its artificial variables, or the

violations of the relaxed constraints, down to zero. Once an integer solution (resulting in an objective function value of 0 and an optimality gap of 0.00%) is reached, this solution is extracted as the initial feasible solution to our original model. This technique is similar in nature to those utilized in goal programming (Lee 1972), as well as “Phase I” of the two-phase simplex method, which is commonly executed by simplex-based solvers to find an initial solution and commence the simplex solution process (Winston, Venkataramanan, & Goldberg 2003).

Dose-volume constraints that contain binary variables are considered to be the difficult constraints in our brachytherapy models. These constraints are coupled to the fine-resolution grids used for anatomical structures (as opposed to the coarse resolution grid used for seed-placement and therefore spatial constraints). Thus, dose-volume constraints occur in large quantities and hinder progress through the branch-and-bound tree in our original model formulation. In accordance with this, each dose-volume constraint is modified as shown in equations (5.2) and (5.3), which is specifically for PTV V100. The objective function is also restructured as given in equation (5.1). Our method of seeking feasible solutions through the violation of dose-volume constraints stands out as novel, since such an approach has not been previously investigated in brachytherapy literature.

$$\text{minimize: } \textit{ArtifV100} + \textit{ArtifV150} + \textit{ArtifV200} + \textit{ArtifD90} + \textit{ArtifD5} + \textit{ArtifD1cc} \quad (5.1)$$

subject to:

$$\textit{ptvDoseV150} - \textit{DoseP}_{(i,j,k)} \leq M * \textit{PTVExceedV100}_{(i,j,k)}, \forall i \in U, j \in V, k \in W \mid T_{(i,j,k)} = 1 \quad (5.2)$$

$$\sum_{i \in U} \sum_{j \in V} \sum_{\substack{k \in W \\ T_{(i,j,k)} = 1}} \textit{PTVExceedV100}_{(i,j,k)} - \textit{ArtifV100} \leq (1 - \textit{ptvVolumeV100}) * \sum_{i \in U} \sum_{j \in V} \sum_{\substack{k \in W \\ T_{(i,j,k)} = 1}} T_{(i,j,k)} \quad (5.3)$$

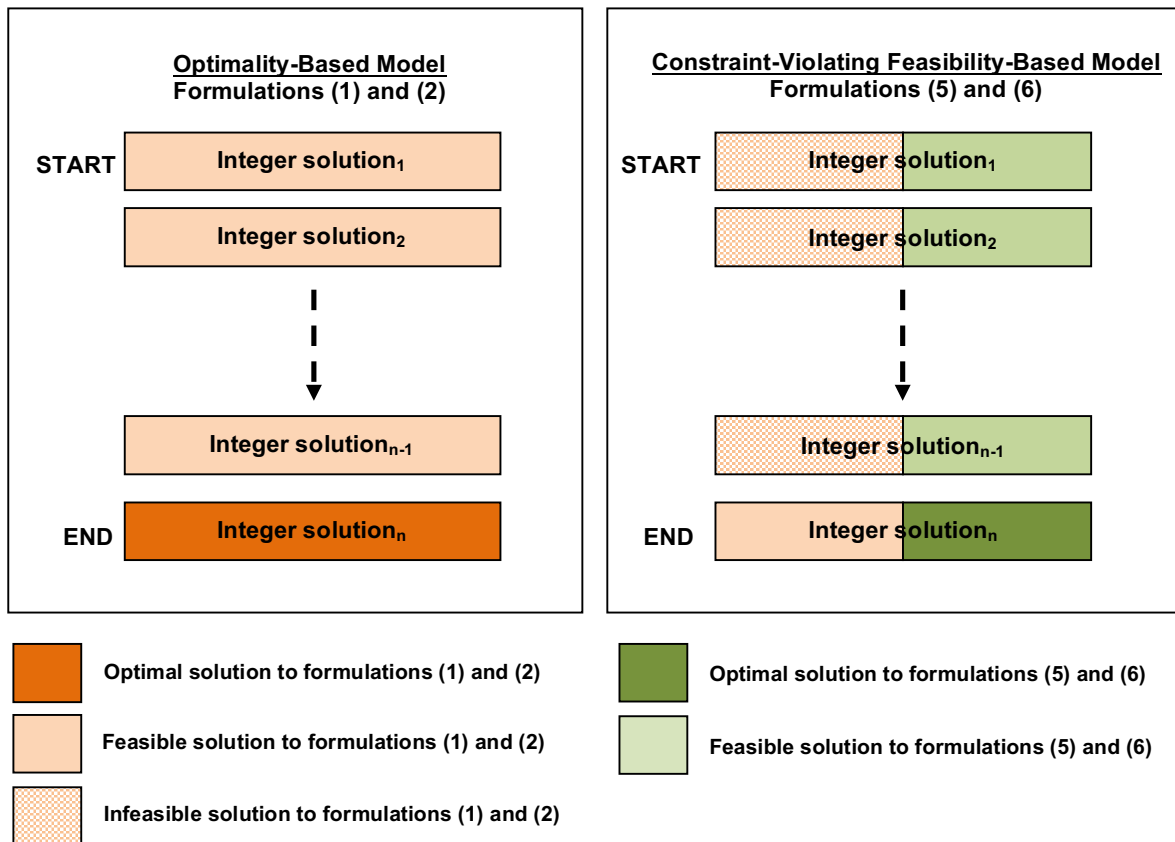


Figure 5.6 Solution processes of the optimality-based model and the constraint-violating feasibility-based model

5.3 Experimental Design

In order to demonstrate the power of the concepts we have introduced in this chapter, we make use of pseudo high-resolution data sets and constraint-violating feasibility-based modelling to create treatment plans for four prostates of varying shapes and sizes. Specifically, our test cases include a small prostate (20.4 cc), medium prostate (42.5 cc), large prostate (63.1 cc) and an irregularly shaped prostate (36.1 cc). The anatomical contour data for these prostates are obtained from four patients previously treated at the local cancer clinic. Due to their diverse anatomical characteristics, these prostates collectively pose a range of clinical and computational challenges (in terms of treatment planning and the numerical solution of our models respectively). Therefore, these test cases enable us to examine the capabilities as well as the limitations of our proposed modelling approach in producing clinically acceptable plans with solution times suitable for pre-operative and intra-operative planning.

For purposes of continuity, the prostate previously investigated in Chapters 3 and 4 has been selected as the medium prostate for the study in this chapter. The irregular prostate used in this study possesses urethral pathway deviation, asymmetrical tendencies in the PTV shape and bilateral midline, as well as significant pubic arch interference. These characteristics render this prostate a particularly difficult one to plan manually, thus making it an ideal candidate to test our automated planning method. Furthermore, in order to assess how well our automated system performs when accommodating one of the more demanding manual planning styles employed in several clinics worldwide, we impose entirely symmetric seed distributions upon the irregular prostate in our automated plans.

Similarly, both the small and large prostates offer unique challenges of their own. As mentioned in previous chapters, manually planning for a small prostate, while adhering to the clinical spatial requirements of the local cancer centre, can be quite laborious. Given the relatively small volume throughout which seed placement is permitted, it becomes a fairly arduous task to avoid adjacent seeds, establish a symmetrical seed distribution, all the while ensuring safety margins for the urethra and the rectum. While the small prostate poses mainly clinical challenges, the large prostate poses a computational challenge to our automated planning system, as the size of this prostate has the potential to result in models that contain relatively larger numbers of binary variables, which may consequently result in longer solution times.

A major experimental design factor to consider in our investigation is the selection of appropriate benchmarks or reference points for the numerical results of our models. It is important to be mindful of the fact that the concepts proposed in this chapter are based on novel MILP modelling strategies, rather than algorithmic solution strategies, to achieve improved (faster) solution times. Furthermore, as explained in Chapter 3, the simplex and branch-and-bound algorithms, in combination, are the sole solution strategies of interest in this thesis. Thus, it makes logical sense to maintain the use of this particular deterministic solution strategy for all of our models, rather than vary the solution approach by conducting comparisons with alternative algorithmic strategies (deterministic, stochastic or heuristic) that have already been investigated in LDR brachytherapy optimization literature. It is more fitting (and relevant) to our investigation that we vary the modelling approach instead, in order to reveal the solution performance

improvements that are possible through our new modelling methods (all within the realm of our chosen deterministic solution strategy).

Since our proposed approach involves changes in the use of input data (a switch from the use of high-resolution data to pseudo high-resolution data) and modelling objective (a switch from optimality-based modelling to feasibility-based modelling), it is important to select appropriate benchmark models that exhibit meaningful points of comparison. Therefore, we solve several model formulations for each prostate size, as outlined in Table 5.1. For ease of interpretation, each formulation is associated with a unique designation. In each designation, “X” represents prostate size or shape and may appropriately be replaced with “S”, “I”, “M” or “L” to represent small, irregular, medium or large prostates respectively.

Table 5.1 Model formulations used in our experimental design.

	Formulation	Designation
(1)	optimality-based model (objective function (3.1)) + high-resolution data	X.O.H
(2)	optimality-based model (objective function (3.1)) + pseudo high-resolution data	X.O.PH
(3)	feasibility-based model (constant objective function of “0”) + high-resolution data	X.FC.H
(4)	feasibility-based model (constant objective function of “0”) + pseudo high-resolution data	X.FC.PH
(5)	constraint-violating feasibility-based model (objective function (5.1)) + high-resolution data	X.FV.H
(6)	constraint-violating feasibility-based model (objective function (5.1)) + pseudo high-resolution data	X.FV.PH

Table 5.2 Details of the objective functions and input data used in each model formulation.

		Model Objective			Input Data Resolution		Count
Designation		Optimality (objective func. (3.1))	Feasibility (constant obj. func. “0”)	Feasibility (objective func. (5.1))	High	Pseudo High	
Formulation	X.O.H	+			+		1
	X.O.PH	+				+	1
	X.FC.H		+		+		1
	X.FC.PH		+			+	1
	X.FV.H			+	+		1
	X.FV.PH			+		+	50

Such an experimental design allows us to identify any improvements that our pseudo high-resolution data sets may project on the solution of formulations (1) and (3). Furthermore, this setup also enables us to orient or rank the solution performance of formulation (6) (which uses pseudo high-resolution data in combination with constraint-violating feasibility-based modelling) in relation to other model formulations. For this purpose, we focus primarily on the length of time taken to produce a feasible solution with each formulation. Several techniques exist for the extraction of a feasible solution. Formulation (1) allows us to maintain the objective function (3.1) used previously, and extract the first integer solution out of many that are sequentially offered (the resulting optimality gap may take on any value). Formulation (3) represents a common technique where the objective function is transformed into a constant and we extract the first integer solution that is offered. Interestingly, this solution generally tends to be the only integer solution produced for the model (resulting in an optimality gap of 0.00%). With formulations (2) and (4), we investigate how the use of pseudo high-resolution data may improve the solution times achieved with formulations (1) and (3) respectively. Moving on, with formulation (5) we examine how our constraint-violating feasibility-based model is able to handle the same high-resolution data sets used in formulations (1) and (3). Finally, we examine the synergistic effects of mixing pseudo high-resolution data with constraint-violating feasibility-based modelling in formulation (6).

Our experimental setup allows for a clear and direct comparison of results to be made as model formulations and input data resolutions are varied. This design also enables us to first numerically establish the common intractability issues encountered during the solution of large-scale LDR brachytherapy models (when solved using the simplex and branch-and-bounds

algorithms), and then demonstrate the ways in which our newly proposed modeling approach is able to overcome such difficulties.

Since feasibility-based models containing dose-volume constraint violations and pseudo high-resolution input data are the primary focus of this chapter, only one instance of formulations (1) to (4) and fifty instances of formulation (5) are solved for each prostate size (the single test cases used for formulations (1) to (4) are randomly selected out of the fifty test cases used for formulation (5)). For each prostate size, model formulations (AMPL model and data files) remain identical throughout the set of fifty feasibility-based test cases, and the only differences that exist among these test cases are the uniquely sampled PTV and rectum input data (sampled at 10% and 30% respectively). While clinical parameters are kept identical among different test cases of the same prostate size, these parameters do slightly differ among prostates of different sizes in order to personalize the automated plans to individual prostate anatomies. Expert-produced plans used in the past treatment of the four associated patients are utilized as guidelines for determining the clinical parameters in Table 5.4.

Each feasibility-based model containing dose-volume constraint violations (formulations (5) and (6)) consists of the objective function given in equation (5.1), the complete set of dose calculation constraints (equations (3.2)-(3.7)), modified dose-volume constraints (similar to the pair of equations (5.2) and (5.3)), the spatial constraints given in equations (3.22) to (3.34) (with the exception of equations (3.26) and (3.35)) and finally equation (4.1). Constraints enforcing lower limits on the number of seeds and needles, as well as those enforcing maximum seed separation in needles are not included in our model formulations since these constraints were not

deemed to be as critical in capturing the spatial requirements of treatment plans as the other spatial constraints. The slice-based local encapsulation method, explained in Chapter 4, is utilized to establish seed-placement margins in all models. Encapsulation margins used for each prostate size are given in Table 5.3 (the base, mid-gland and apical slices are highlighted in orange, green and blue respectively). Two separate margins are used to account for the urethral curvature for each prostate size (0.40 cm for seed-placement incursion, 0.20 cm for needle incursion). Furthermore, a 0.40 cm rectal protection margin is implemented in all models. Finally, the pubic arch is accounted for in each model as well.

The optimality-based models (formulations (1) and (3)) are almost identical in form to the feasibility-based models (formulations (5) and (6)). The only differences that exist are the use of objective function (3.1) instead of (5.1), and the inclusion of the original dose volume constraints (3.10)-(3.21) instead of their relaxed counterparts. Similarly, the structures of the feasibility-based models, outlined in formulation (3) and (4), are identical to those of the optimality-based models. The sole difference between these models is the use of a constant objective function value of “0” in formulations (3) and (4), rather than objective function (3.1).

Table 5.3 Seed-placement margins used for the differently sized and shaped prostates.

Slice z (cm)	Small (20.4 cc) (cm)	Irregular (36.1 cc) (cm)	Medium (42.5 cc) (cm)	Large (63.1 cc) (cm)
0.0	0.50	0.50	0.50	0.50
0.5	0.40	0.40	0.40	0.40
1.0	0.30	0.30	0.30	0.30
1.5	0.20	0.20	0.30	0.30
2.0	0.10	0.10	0.20	0.20
2.5	0.10	0.20	0.10	0.10
3.0	0.20	0.30	0.10	0.10
3.5	0.30	0.40	0.20	0.20
4.0	0.40	0.50	0.30	0.30
4.5	0.50		0.30	0.30
5.0			0.40	0.40
5.5			0.50	0.50

Table 5.4 Clinical parameters set for the differently sized and shaped prostates.

		Small (20.4 cc)	Irregular (36.1 cc)	Medium (42.5 cc)	Large (63.1 cc)
Spatial requirements	Upper seed limit	74	93	108	134
	Upper needle limits	16	21	24	28
	Max. needle retraction	1.5 cm	1.5 cm	2.0 cm	1.5 cm
	Min. seeds per needle	2	2	2	2
	Max. consecutive seeds in needle	4	4	4	4
Dose-volume requirements	Prescribed PTV Dose	145 Gy	145 Gy	145 Gy	145 Gy
	PTV V100	> 98%	> 98%	> 98%	> 98%
	PTV V150	≤ 65%	≤ 65%	≤ 65%	≤ 65%
	PTV V200	≤ 25%	≤ 25%	≤ 25%	≤ 25%
	PTV D90	174-189 Gy	174-189 Gy	174-189 Gy	174-189 Gy
	Urethra D5	< 215 Gy	< 215 Gy	< 215 Gy	< 215 Gy
	Rectum D1cc	< 145 Gy	< 145 Gy	< 145 Gy	< 145 Gy

AMPL (Fourer et al. 2002) is used for the formulation and Gurobi 6.5.0 (Gurobi Optimization 2014) is used for the solution of our models on a 20-core ACPI multiprocessor X64-based PC with Intel Xeon® CPU E5-2650 running at 2.3 GHz with 128 GB RAM. The specific reasons behind our choice of modeling language and solver are discussed in Chapter 3. The following parameters are set for each feasibility-based model (formulations (2), (4), (5) and (6)): **presolve = 10** in AMPL and **mipgap = 0.0001**, **timelim = 86400**, **presolve = 2**, **mipfocus = 1**, **heurfrac = 0.95** in Gurobi. The following parameter differs for the optimality-based model (formulations (1) and (3)): **mipgap = 0.50**. We provide an extensive explanation for the use of these parameters and settings in Chapter 3. The reason for using **mipgap = 0.0001** for feasibility based models and **mipgap = 0.50** for optimality-based models is due to the fact that Gurobi stops its branch-and-bound solution process once either the **mipgap** or **timelim** limit is reached. The best current integer solution is then extracted as the final result for each model. Our specific parameter values allow us to extract a solution that results in a 0.00% optimality gap in feasibility-based models, and to extract the first (initial) solution that is obtained in optimality-based models (from experience, this initial solution usually results in an optimality gap below 50.00%). Both of these extracted solutions are considered to be equivalent in nature (in other words, they are an initial feasible solution to the optimality-based model formulation). The parameter **heurfrac** determines the fraction of the solution time and resources that should be allocated for using the RINS heuristic, as mentioned in Chapter 3. The parameter **mipfocus** determines the branch-and-bound search strategy followed by Gurobi; while **mipfocus=0** (the default value) strikes a balance between finding good feasible solutions and proving optimality, **mipfocus=1** focuses more extensively on

finding good feasible solutions. After testing a range of values, including the default value of **heurfrac** = 0.05, the parameter values **mipfocus** = 1 and **heurfrac** = 0.95 were selected as they consistently resulted in the fastest solution times for all model formulations.

5.4 Results

In this section, detailed solution performance results, obtained during the creation of treatment plans for the four prostate test cases, are compared among the six model formulations. For demonstration purposes, a single solution is obtained using formulations (1)-(5) for each prostate size. On the other hand, heavy emphasis is placed on formulation (6) as this particular modelling approach constitutes the main focus of this chapter. The numerical features of the branch-and-bound solutions of these formulations are provided in Table 5.5. This table has been colour coded to ease the interpretation of results. Each family (pair) of formulations possesses its own unique colour code: formulations (1)-(2) are coloured red, formulations (3)-(4) are coloured blue and formulations (5)-(6) are coloured green. In each pairing, the darker shade of color represents a formulation that utilizes high-resolution data, while its lighter counterpart represents a formulation utilizing pseudo high-resolution data. Colour coding enables the results of all high-resolution models (the darker shades) to be easily singled out and compared to each other at first glance. This is also the case for the pseudo high-resolution models (lighter shades). Further details of the numerical results obtained for the fifty test cases of formulation (6) are presented in Tables 5.6-5.15 and visualizations of these results are provided in Figures 5.7-5.13.

Numerous trends may be highlighted among the scores of results presented in this section. The first and most obvious observation is that high resolution models (formulations (1), (3), (5))

generally result in unacceptably long total solution times. Out of these three approaches, formulation (3) seems to result in relatively shorter solution times, notably reaching values as low as 13.58 minutes for the small prostate, while formulation (5) generally leads to the longest solution times, with a value of 583.33 minutes (or 9.72 hours) for the large prostate in particular. These results are expected, and essentially establish the problematic initial conditions which we initially hoped to overcome through the use of pseudo high-resolution data. It is clear that the use of pseudo high-resolution data in formulations (2), (4), (6) has effectively managed to decrease the unsuitable total solution times that were initially observed with high-resolution data in formulations (1), (3) and (5). In particular, the total solution time of 583.33 minutes in formulation (5) has experienced an exceptional decrease to an average of 5.73 minutes in formulation (6). Furthermore, this reduction in solution time is consistent across all prostate sizes for presolve time as well as branch-and-bound time. This is a noteworthy achievement, considering that presolve times exhibited unreliable behavior, occasionally equaling or surpassing branch-and-bound times in high-resolution models. With pseudo high-resolution models, presolve times are consistently lower than branch-and-bound times.

Perhaps the most significant observation one can make throughout our results is that formulation (6), which combines pseudo high-resolution data with constraint-violating feasibility-based modelling, results in the fastest solution times for all prostates tested (achieving average values of 0.96, 0.76, 1.55 and 5.73 minutes for small, irregular, medium and large prostates respectively). When all prostates sizes are collectively taken into consideration, the average solution time amounts to 2.25 minutes. These numbers exhibit a high likelihood of suitability for

intra-operative planning. While formulation (4) does come close to the solution times of formulation (6) for the irregular and medium prostate (with solution times of 3.42, and 1.05 minutes respectively), the same cannot be said in regards to the 10.67 and 17.25 minutes obtained for the small and large prostates respectively. This scenario in a way reveals the unreliability of results achieved through feasibility-based models using a constant objective function. This observation is further exemplified through the branch-and-bound solution times obtained for optimality-based models (as opposed to total solution times). While branch-and-bound solution times of 238.43, 130.55 and 288.30 minutes are obtained for the small, irregular and large prostates, a solution time of 1.60 minutes is achieved for the medium-sized prostate. This trend in the branch-and-bound solution times is also reflected in the total solution times obtained through optimality-based models. While an obvious criticism of our results may be aimed at the lack of extensive testing presented for formulations (1)-(5), such criticism would be misguided as to the true intent behind our results. The purpose of these results is not necessarily to provide a comprehensive account of all solutions attainable, but to reveal abnormalities (and perhaps extremities) in the solution time performance of these model formulations. In order to make our case, we point out the maximum solution time values of formulation (6). The corresponding solution times of 1.52, 1.62 and 9.20 minutes obtained for small, irregular and large prostates respectively are well below those obtained for the next fastest method, formulation (4) (the only exception to this trend is observed with a maximum solution time of 3.05 minutes for the medium prostate). The potential power of formulation (6) is further demonstrated through the minimum

solution times of 17 seconds, 17 seconds, 39 seconds and 127 seconds obtained for small, irregular, medium and large prostates respectively.

It is important to note that it may be common for expert planners to slightly alter clinical parameters used in treatment plans on a per-patient case, in order to ensure positive clinical outcomes through increased dose coverage for patients that present with extensive spread of disease in the prostate. As such, the solution time response of formulations (2), (4), and (6) to variations in several clinical parameters are further investigated and the results are displayed in Table 5.16. Formulations (1), (3), and (5) are not tested, as the primary focus of this investigation is the solution time performance of formulations involving pseudo high-resolution data sets. Once again, a single instance is solved for formulations (2) and (4), while fifty instances are solved for formulation (6) (with test cases that identical to those that were previously solved to obtain the results in Table 5.5, with the exception of several modified clinical parameters). Table 5.16 includes two columns outlined in red: the left column highlights solution times obtained using the original clinical parameters (with solution time values migrated from Table 5.5), while the right column highlights solution times obtained for the modified clinical parameters. The modified clinical parameters were as follows:

- PTV V100 increased from > 98% to > 99% for all prostate sizes
- PTV D90 increased from 174-189 Gy to 181.5-189 Gy for all prostate sizes
- Maximum needle retraction decreased from 1.5 cm to 1.0 cm for small and large prostates

For formulation (2), the modified clinical parameters have led to an increase in solution times from 11.17 minutes to 32.87 minutes for the small prostate, from 1.13 hours to 9.14 hours

for the irregular prostate, and from 5.90 minutes to 3.81 hours for the medium prostate, while leading to a decrease in solution time from 2.82 hours to 1.99 hours for the large prostate. Likewise for formulation (4), the modified clinical parameters have resulted in an increase in solution times from 10.67 minutes to no-solution after 24 hours for the small prostate, from 3.42 minutes to no-solution after 24 hours for the irregular prostate, from 1.05 minutes to 37.92 minutes for the medium prostate, and from 17.25 minutes to 17.58 minutes for the large prostate. These unpredictable solution time results obtained in response to variations in clinical parameters render the optimality-based model and the constant objective feasibility-based model inappropriate for use in intraoperative planning, and non-ideal for use in pre-operative planning. Formulation (6), on the other hand, has only experienced a slight increase in solution times with the modified clinical parameters in comparison to formulations (2) and (4): from 0.96 minutes to 1.92 minutes for the small prostate, from 0.76 minutes to 3.07 minutes for the irregular prostate, from 1.55 minutes to 4.32 minutes for the medium prostate, and from 5.73 minutes to 8.97 minutes for the large prostate.

Considering the extensive testing carried out on formulation (6), these results not only reveal the efficiency but also the reliability with which we are able to produce treatment plans through our automated approach. These two characteristics render our automated planning system highly attractive for use in a high-paced clinical setting.

Table 5.5 Solution performance results obtained for the six model formulations.

	Designation	Constraints	Variables		Solution Values			Solution Time			
			Integer	Total	Object. Function	Best Bound	Optim. Gap	Presolve (mins)	B&B (mins)	Total (mins)	Total (hours)
Small (20.4 cc)	S.O.H	34062	30500	33728	236.902	196.066	17.2%	3.42	238.43	241.85	4.03
	S.O.PH	4440	3602	4106	254.342	201.186	20.9%	0.08	11.08	11.17	0.19
	S.FC.H	34062	30500	33728	0.000	0.000	0.0%	5.08	8.50	13.58	0.23
	S.FC.PH	4440	3602	4106	0.000	0.000	0.0%	0.12	10.55	10.67	0.18
	S.FV.H	34062	30500	33735	0.000	0.000	0.0%	4.90	42.68	47.58	0.79
	S.FV.PH*	4558	3702	4237	0.000	0.000	0.0%	0.13	0.84	0.96	0.02
Irregular (36.1 cc)	I.O.H	56469	50633	56040	274.679	204.560	25.5%	13.00	130.55	143.55	2.39
	I.O.PH	7135	5854	6706	250.507	207.154	17.3%	0.22	67.45	67.67	1.13
	I.FC.H	56469	50633	56040	0.000	0.000	0.0%	17.42	16.68	34.10	0.57
	I.FC.PH	7135	5854	6706	0.000	0.000	0.0%	0.28	3.13	3.42	0.06
	I.FV.H	56469	50633	56047	0.000	0.000	0.0%	17.11	327.64	344.75	5.75
	I.FV.PH*	7249	6827	5951	0.000	0.000	0.0%	0.24	0.53	0.76	0.01
Medium (42.5 cc)	M.O.H	63279	56612	62702	263.288	166.094	36.9%	43.68	1.60	45.28	0.75
	M.O.PH	8518	6945	7941	265.647	188.214	29.1%	0.28	5.62	5.90	0.10
	M.FC.H	63279	56612	62702	0.000	0.000	0.0%	34.72	29.10	63.82	1.06
	M.FC.PH	8518	6945	7941	0.000	0.000	0.0%	0.48	0.57	1.05	0.02
	M.FV.H	63279	56612	62709	0.000	0.000	0.0%	32.72	84.53	117.25	1.95
	M.FV.PH*	8431	6871	7861	0.000	0.000	0.0%	0.44	1.11	1.55	0.03
Large (63.1 cc)	L.O.H	95183	85365	94140	267.005	186.233	30.3%	152.62	288.30	440.92	7.35
	L.O.PH	12776	10259	11733	262.102	187.065	28.6%	0.90	156.58	157.48	2.62
	L.FC.H	95183	85365	94140	0.000	0.000	0.0%	133.68	189.32	323.00	5.38
	L.FC.PH	12776	10259	11733	0.000	0.000	0.0%	1.82	15.43	17.25	0.29
	L.FV.H	95183	85365	94147	0.000	0.000	0.0%	146.83	436.50	583.33	9.72
	L.FV.PH*	12607	10111	11571	0.000	0.000	0.0%	1.54	4.19	5.73	0.10

*mean values calculated from 50 test cases

Table 5.6 Number of constraints observed in constraint-violating feasibility-based models.

Prostate	Designation	Constraints			
		Minimum	Mean	Maximum	Standard Dev.
Small (20.4 cc)	S.FV.PH	4339.00	4557.12	4782.00	104.73
Irregular (36.1 cc)	M.FV.PH	6921.00	7248.96	7568.00	156.86
Medium (42.5 cc)	I.FV.PH	8107.00	8430.64	8864.00	169.31
Large (63.1 cc)	L.FV.PH	12259.00	12606.72	12867.00	140.17

Table 5.7 Number of binary variables observed in constraint-violating feasibility-based models.

Prostate	Designation	Integer (Binary) Variables			
		Minimum	Mean	Maximum	Standard Dev.
Small (20.4 cc)	S.FV.PH	3515.00	3701.56	3892.00	88.87
Irregular (36.1 cc)	M.FV.PH	5670.00	5950.76	6206.00	131.85
Medium (42.5 cc)	I.FV.PH	6594.00	6871.02	7247.00	144.46
Large (63.1 cc)	L.FV.PH	9815.00	10111.26	10322.00	117.30

Table 5.8 Number of total variables observed in constraint-violating feasibility-based models.

Prostate	Designation	Total Variables			
		Minimum	Mean	Maximum	Standard Dev.
Small (20.4 cc)	S.FV.PH	4012.00	4236.66	4455.00	102.65
Irregular (36.1 cc)	M.FV.PH	6499.00	6826.96	7146.00	156.86
Medium (42.5 cc)	I.FV.PH	7537.00	7860.64	8294.00	169.31
Large (63.1 cc)	L.FV.PH	11223.00	11570.72	11831.00	140.17

Table 5.9 Starting objective function values observed in constraint-violating feasibility-based models.

Prostate	Designation	Starting Objective Function Value			
		Minimum	Mean	Maximum	Standard Dev.
Small (20.4 cc)	S.FV.PH	1112.96	1181.17	1253.96	32.81
Irregular (36.1 cc)	M.FV.PH	1880.00	1979.79	2083.04	49.38
Medium (42.5 cc)	I.FV.PH	2133.80	2227.20	2361.28	52.83
Large (63.1 cc)	L.FV.PH	3205.40	3328.39	3412.20	44.40

Table 5.10 Number of explored nodes observed in constraint-violating feasibility-based models.

Prostate	Designation	Explored Nodes			
		Minimum	Mean	Maximum	Standard Dev.
Small (20.4 cc)	S.FV.PH	0.00	419.54	1376.00	341.84
Irregular (36.1 cc)	M.FV.PH	0.00	56.68	581.00	100.15
Medium (42.5 cc)	I.FV.PH	0.00	146.12	876.00	222.45
Large (63.1 cc)	L.FV.PH	0.00	76.66	317.00	80.04

Table 5.11 Number of simplex iterations observed in constraint-violating feasibility-based models.

Prostate	Designation	Simplex Iterations			
		Minimum	Mean	Maximum	Standard Dev.
Small (20.4 cc)	S.FV.PH	10868.00	103952.06	288466.00	64024.46
Irregular (36.1 cc)	M.FV.PH	11587.00	29023.52	194063.00	28396.02
Medium (42.5 cc)	I.FV.PH	15664.00	54650.12	282405.00	55944.09
Large (63.1 cc)	L.FV.PH	31195.00	72865.70	243414.00	45134.85

Table 5.12 Presolve times observed in constraint-violating feasibility-based models.

Prostate	Designation	Presolve Time (min)			
		Minimum	Mean	Maximum	Standard Dev.
Small (20.4 cc)	S.FV.PH	0.10	0.13	0.17	0.02
Irregular (36.1 cc)	M.FV.PH	0.20	0.24	0.28	0.02
Medium (42.5 cc)	I.FV.PH	0.38	0.44	0.53	0.04
Large (63.1 cc)	L.FV.PH	1.23	1.54	1.85	0.14

Table 5.13 Branch & bound solution times observed in constraint-violating feasibility-based models.

Prostate	Designation	Branch & Bound Time (min)			
		Minimum	Mean	Maximum	Standard Dev.
Small (20.4 cc)	S.FV.PH	0.17	0.84	1.38	0.29
Irregular (36.1 cc)	M.FV.PH	0.08	0.53	1.37	0.35
Medium (42.5 cc)	I.FV.PH	0.22	1.11	2.58	0.63
Large (63.1 cc)	L.FV.PH	0.77	4.19	7.48	1.43

Table 5.14 Total solution times (minutes) observed in constraint-violating feasibility-based models.

Prostate	Designation	Total Solution Time (min)			
		Minimum	Mean	Maximum	Standard Dev.
Small (20.4 cc)	S.FV.PH	0.28	0.96	1.52	0.30
Irregular (36.1 cc)	M.FV.PH	0.32	0.76	1.62	0.36
Medium (42.5 cc)	I.FV.PH	0.65	1.55	3.05	0.63
Large (63.1 cc)	L.FV.PH	2.12	5.73	9.20	1.48
All prostates	-	0.28	2.25	9.20	2.19

Table 5.15 Total solution times (seconds) observed in constraint-violating feasibility-based models.

Prostate	Designation	Total Solution Time (sec)			
		Minimum	Mean	Maximum	Standard Dev.
Small (20.4 cc)	S.FV.PH	17.00	57.80	91.00	17.67
Irregular (36.1 cc)	M.FV.PH	17.00	45.74	97.00	21.57
Medium (42.5 cc)	I.FV.PH	39.00	92.76	183.00	37.92
Large (63.1 cc)	L.FV.PH	127.00	343.74	552.00	88.51
All prostates	-	17.00	135.01	552.00	131.66

Table 5.16 Solution performance results of the three model formulations containing pseudo high-resolution data.

	Designation	Constraints	Variables		Solution Time (original parameters)		Solution Time (modified parameters)	
			Integer	Total	Total (mins)	Total (hours)	Total (mins)	Total (hours)
Small (20.4 cc)	S.O.PH	4440	3602	4106	11.17	0.19	32.87	0.55
	S.FC.PH	4440	3602	4106	10.67	0.18	> 14400.00	> 24.00
	S.FV.PH*	4558	3702	4237	0.96	0.02	1.92	0.03
Irregular (36.1 cc)	I.O.PH	7135	5854	6706	67.67	1.13	548.42	9.14
	I.FC.PH	7135	5854	6706	3.42	0.06	> 14400.00	> 24.00
	I.FV.PH*	7249	6827	5951	0.76	0.01	3.07	0.05
Medium (42.5 cc)	M.O.PH	8518	6945	7941	5.90	0.10	228.40	3.81
	M.FC.PH	8518	6945	7941	1.05	0.02	37.92	0.63
	M.FV.PH*	8431	6871	7861	1.55	0.03	4.32	0.07
Large (63.1 cc)	L.O.PH	12776	10259	11733	157.48	2.62	119.20	1.99
	L.FC.PH	12776	10259	11733	17.25	0.29	17.58	0.29
	L.FV.PH*	12607	10111	11571	5.73	0.10	8.97	0.15

*mean values calculated from 50 test cases

Original parameters:

- PTV V100 > 98%
- PTV D90 = 174-189 Gy
- Needle retraction:
(S) 1.5 cm (I) 1.5 cm
(M) 2.0 cm, (L) 1.5 cm

Modified parameters:

- PTV V100 > 99%
- PTV D90 = 181.5-189 Gy
- Needle retraction:
(S) 1.0 cm (I) 1.5 cm
(M) 2.0 cm, (L) 1.0 cm






Figure 5.8 Number of binary variables in relation to prostate volume, as observed in constraint-violating feasibility-based models.




Figure 5.9 Number of constraints in relation to prostate volume, as observed in constraint-violating feasibility-based models.




Figure 5.10 Solution time in relation to prostate volume, as observed in constraint-violating feasibility-based models.




Figure 5.11 Solution time in relation to number of binary variables, as observed in constraint-violating feasibility-based models.




Figure 5.12 Solution time in relation to number of binary constraints, as observed in constraint-violating feasibility-based models.

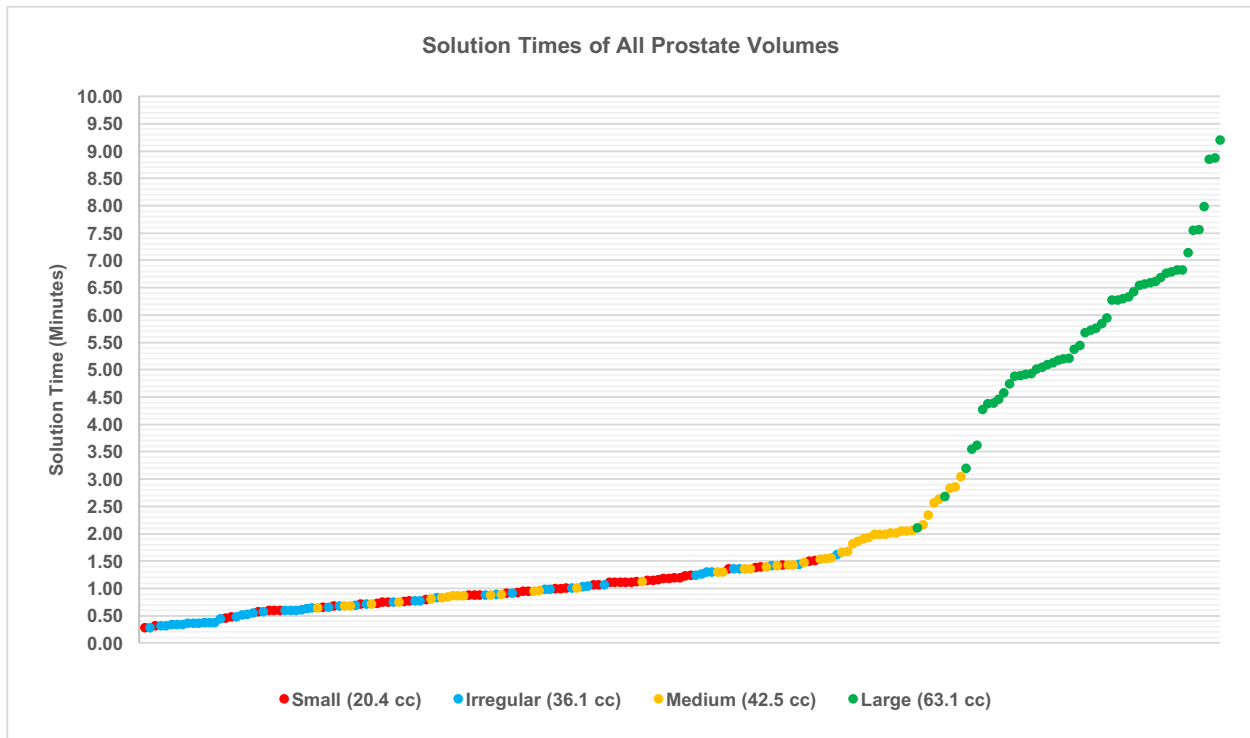


Figure 5.13 Ordered solution times of all 200 prostate test cases, as observed in constraint-violating feasibility-based models.

5.5 Discussion

Several key inferences may be made from our results. Firstly, the improvement in solution time that is possible through the use of pseudo high-resolution data is apparent when comparing the results of formulations (2) and (4) with formulations (1) and (3). Our proposed modelling technique effectively advances the solution speeds of these formulations, which represent more obvious methods of obtaining feasible solutions.

More importantly, the combination of using pseudo high-resolution data and feasibility-based modelling has had the desired effect of decreasing the solution times of our models from unsuitable levels (as seen in many instances of formulations (1)-(5)) to those appropriate for intra-operative planning. With solution times averaging a mere 2.25 minutes for the wide range of prostates tested (using the feasibility-based models given in formulation (6)), our automated planning approach has demonstrated its ability to comfortably meet the standards of our local cancer center and presumably other clinics worldwide.

Our feasibility-based models based on formulation (6) do exhibit some variance in solution time for each size of prostate, as observed in Figures 5.7-5.13. This is expected, considering the variance that also exists for the numbers of constraints and binary variables for each prostate size, although these numbers generally display tighter clustering in comparison to solution time. While not consistent in their form, the spreads (deviation from the mean) of our solution time results do resemble normal distributions, especially for the small and large prostates. These results are, on occasion, skewed towards faster solution times (in the case of the irregular and medium sized prostate). Additionally, our solution times exhibit increasing spread with increasing prostate size.

This trend is particularly exemplified during the transition from the medium sized prostate (42.5 cc) to the large sized prostate (63.1 cc). Also, while model size does seem to play a general role in solution speed, with increasingly larger prostates resulting in longer solution times, it is apparent that prostate (or model) size is not the sole factor involved in solution speed. Perhaps the most direct evidence of this is revealed through the solution times obtained for the small prostate (20.4 cc) and the larger, irregularly shaped prostate (36.1 cc). Despite containing relatively more constraints and binary variables on average, the irregularly shaped prostate solves in a shorter period of time (on average) than the small sized prostate. This situation may be attributed to the previously mentioned clinical challenges faced during treatment planning for small prostates. Although the irregular prostate provides unique clinical challenges of its own, it appears as though the limited volume within which planning has to be carried out for the small prostate plays a more influencing role in solution time. Taking an alternative view, it is also surprising to see just how efficiently our automated planning system is able to produce fully symmetrical seed distributions for the irregular prostate, considering that the expert planner who had originally carried out manual planning for this particular prostate was forced to make extensive use of asymmetrical seed positions in order to achieve an appropriate dose distribution. These observations reveal that while certain clinical challenges faced by manual planners may translate to automated planning (in the case of the small prostate), other challenges may not (in the case of the irregularly shaped prostate).

A second major finding is that while it may initially appear to the contrary, it is not necessarily a trivial task to obtain an initial (feasible) solution in large-scale models (consisting, in our classification, of more than 15,000 binary variables) or even medium-scale models (between

1,000 and 15,000 binary variables). As discussed previously, this is apparent for high-resolution models given in formulations (1), (3) and (5) using the original set of clinical parameters. However, the results obtained using the modified set of clinical parameters indicate that intractability is also a major issue for formulations (2) and (4), which utilize pseudo high-resolution data; formulation (4) results in solution times of 9.14 hours, 3.81 hours, and 1.99 hours for the irregular, medium, and large prostates, while formulation (2) is not able to produce an initial feasible solution even after 24 hours for the small and irregular prostates. These results reveal the unreliable nature of both the optimality-based model and the constant objective feasibility-based model. On the contrary, the constraint-violating feasibility based model (formulation (6)) does not experience such issues of intractability when coupled with pseudo high-resolution data sets.

Our findings here demonstrate that pseudo high-resolution data is an exceptionally effective method for decreasing model size, and therefore solution time. It should be mentioned that the results produced through our constraint-violating feasibility-based model (formulation (6)) are generally deemed to meet the treatment planning standards of our local cancer center (as based on expert planner analysis). Although verification of clinical acceptability will be detailed in Chapter 6, feasibility-based modelling presently shows promise as a method for creating clinically acceptable results.

5.6 Conclusion

This chapter builds upon the mixed-integer linear programming modeling approach introduced in Chapter 3 and further enhanced in Chapter 4 for solving the interstitial LDR prostate brachytherapy problem. The low-resolution model utilized in these preceding chapters is first transformed into its high-resolution counterpart, which enables us to more accurately delineate anatomical structures within discrete data sets and more precisely calculate the radiation distribution to these structures. Collectively, these improvements increase the clinical reliability of treatment plans produced through our automated planning system. Furthermore, we demonstrate the common intractability issues associated with the deterministic solution of our large-scale models that arise from the use of high-resolution data sets. We propose two effective solutions to overcome such issues: the use of pseudo high-resolution data and the transformation of our originally optimality-based model to one of feasibility. The results obtained from comprehensive testing on prostates of varying shapes and sizes highlight our ability to produce treatment plans satisfying specified dose-volume and clinical constraints with solution times that range from less than a minute to roughly five minutes on average. Such a significant achievement renders our automated planning system relevant not only in pre-operative planning, but also in intra-operative planning. The speed with which we are able to obtain high quality treatment plans is highly valuable for the practical integration of our automated planning approach within real-world clinics.

The concepts proposed in this chapter for low dose rate prostate brachytherapy may be applicable for achieving fast solutions in other variants of prostate brachytherapy (high dose rate, medium dose rate and pulsed dose rate), as well as in the brachytherapy treatment of breast,

cervical, lung, and head and neck cancers. Furthermore, our proposed approach may prove beneficial for increasing solution speed not only in the realm of deterministic solution algorithms, but also for alternative stochastic or heuristic methods as well. Lastly, the modelling techniques described may be utilized in solving other grid-based location problems, specifically in the resource, oil and gas exploration, telecommunications, and emergency service dispatch fields.

Chapter 6⁴

Clinical Assessment of Automated Plans

The main focus of this chapter is to compare the quality and clinical acceptability of I-125 prostate treatment plans created manually by experienced planners with plans created through our automated planning approach. The motivating factor behind this pilot study is to gain an initial understanding of the general clinical performance of our automated planning system and verify that it produces clinically approved results. More specifically, we are interested in deciphering the clinical requirements our automated planning approach excels in fulfilling, as well as those that our approach may fall short of satisfying. These observations will subsequently enable us to define new pathways of improvement in the design of our automated planning system. Furthermore, it is not our intention to surpass the clinical quality of manual plans, although this study does perhaps present an opportunity for it. Our main goal is to determine if our automated plans fare roughly equally to manual plans in particular areas or as a whole.

⁴*A version of this chapter will be submitted for publication: Babadagli, M.E.; Sloboda, R.; Usmani, N.; Amanie, J.; Murtha, A.; Yee, D.; Jamaluddin, M.; Doucette, J. “Clinical Assessment of Center-Specific Automated Treatment Plans for Interstitial Low Dose Rate Prostate Brachytherapy”. Elsevier, Brachytherapy (to be submitted June 1, 2016).*

In doing so, we also have a chance to assess the clinical acceptability and the real-world utility of the numerous concepts and techniques discussed in previous chapters. The concepts of primary interest are spatial constraints, pseudo high-resolution data sets, and feasibility-based modelling. The final purpose of this investigation is to reveal potential areas of disconnect or misalignment between the engineering and the clinical worlds. More specifically, despite our extensive attempt to approach treatment planning through the minds of expert planners, our hope is that this study reveals to us the more subtle considerations that clinicians utilize in assessing treatment plan quality, which may not have been initially obvious to us as engineers. This final point may potentially facilitate significant leaps in LDR brachytherapy optimization by enabling us to uncover previously unrecognized gaps in present-day automated planning and encouraging us to fill said gaps to conceive a system that is highly applicable in real-world clinics.

6.1 Study Design

Treatment plans (involving I-125 sources) for twenty patients previously implanted at the Cross Cancer Institute were selected to span a wide range of prostate sizes, shapes and geometries. In terms of size, these prostates ranged from 20.4 cc to 85.4 cc, with an average volume of 42.2 cc. In regards to shape, both bilaterally symmetrical (regularly shaped) and asymmetrical (irregularly shaped) prostates were included. The irregularly shaped prostates possessed such features as urethral pathway deviation, asymmetrical tendencies in gland shape and bilateral midline, as well as pubic arch interference. The model formulation (6) detailed in Chapter 5 was utilized to replicate the planning approach practiced at the Cross Cancer Institute (both in terms of spatial and dose-volume considerations) and was utilized to independently generate twenty

corresponding plans from the same clinical data as the manual plans. The resulting forty plans were anonymized and randomized for individual evaluation by five radiation oncologists. A scoresheet, provided in Table 6.1, was created to facilitate clinical evaluation of these plans.

For logistical reasons the study was conducted in two parts, each part involving the evaluation of twenty anonymized and randomized plans. In both parts of the study and for each plan, a VariSeed report in the standard Cross Cancer Institute clinical format was provided at regular intervals to the participating radiation oncologists, along with a request for evaluation. In each part of the study, the sequence of plans provided to each of the oncologists was different, and all twenty plans in that part of the study were continuously available for detailed review on the VariSeed Treatment Planning System. The oncologists were requested to:

- evaluate each plan in a timely manner, as for pending clinical implants;
- make use of the “Comments” section of the scoresheet to identify specific concerns;
- address any questions or concerns about a plan to the expert planner managing the study;
- not accumulate a large number of plans for batch evaluation at one sitting;
- not discuss evaluations of plans with colleagues participating in the study;
- submit completed scoresheets to the expert planner managing the study.

The scoresheet was composed of seven specific criteria for clinically evaluating the quality of treatment plans that were identified prior to the study by the participating radiation oncologists. These criteria are summarized as follows:

- Has the plan resulted in adequate dose coverage of the PTV?
- Has the plan avoided overdosing of critical structures?

- Has the plan avoided pubic arch interference?
- Have strands been anchored within the prostate?
- Have seeds been placed in the bladder?
- Has the plan made efficient use of needles?
- Has an acceptable distribution of strands and seeds been achieved?

The scoring or rating system for each criterion was based on a five-point scale. An additional category was included to allow radiation oncologists to assess and provide their perception of the overall quality of each plan. Finally, a response to a concluding question regarding the clinical acceptability of each treatment plan was requested of the evaluator. This question not only functioned as the main metric of concern in this study, but also allowed us to effectively avoid ambiguities associated with the rating system used for the previously mentioned eight criteria; such ambiguities had to do with the fact that the existence of a single critical flaw might be sufficient to render a treatment plan clinically unacceptable, regardless of the ratings received in the eight categories. The question regarding the clinical acceptability of each treatment plan confirmed the presence or absence of such flaws and presented the final decision of each radiation oncologist in clear terms.

Ethics approval for this study was obtained from the Health Research Ethics Board at the University of Alberta (Study ID: Pro00056288).

Table 6.1 A representation of the score sheet given to radiation oncologists.

Plan #: _____ **Radiation Oncologist:** _____

Date: _____

I-125 Plan Score Sheet

Criterion to assess plan	Unacceptable (1)	Sub-optimal (2)	Indeterminate (3)	Acceptable (4)	Excellent (5)
Adequate dose coverage of PTV					
Critical structures not overdosed					
No pubic arch interference					
Strands placed within the prostate					
No seeds placed in the bladder					
Efficient use of needles					
Acceptable distribution of strands/seeds					
Overall Plan Quality					

Plan is clinically acceptable? ☐ Yes ☐ No

Comments:

6.2 Results

In this section we compare the numerical features of the manual and automated plans, as well as their clinical analysis results for the twenty patient test cases. The associated dose-volume metric results are highlighted in Figures 6.2-6.7. For all plans, the final dose distribution was calculated in the VariSeed LDR Treatment Planning System using the TG-43 2D (line source) formalism (Rivard et al. 2004). The red lines, in Figures 6.2-6.7, represent the clinical dose-limits set by the Cross Cancer Institute and the red shaded area of each graph represents range of values that are acceptable for each dose-volume requirement. The spatial (or geometric) characteristics of all plans are displayed in Figures 6.8-6.11. Clinical analysis results from the five radiation oncologists are provided in Tables 6.2-6.5 for the manual and automated plans.

It is evident that both the manual and automated plans largely satisfied the dose-volume metric requirements of the local cancer clinic, as observed in Figures 6.2-6.7. In almost all cases, the dosimetric values of both types of plans remained within the set dose-limits (the red shaded region of each graph), with a few exceptions. These exceptions include the manual plan PTV V150 value for Patient 9, the manual plan PTV V200 values for Patients 2, 5, 8, 9, 10, 11, 13, 14, 15, 18 and 20, the automated plan PTV V200 value for Patient 19 and the automated plan Urethra D5 values for Patients 7 and 17. We believe the two Urethra D5 violations in automated plans are due to differences in dose-calculation formalisms between our models (1-D point-source) and Variseed Treatment Planning System (2-D line-source), the piece of software used to calculate the dose-volume results post-optimization. We base this conclusion on the fact that dose-volume constraints are entirely satisfied within our models. While these few instances technically constitute as

violations of dose-limits set by the local cancer clinic, they are minor in nature and therefore generally still are acceptable. It is interesting to note that such dose-limit violations are more frequent among manual plans rather than automated plans, which highlights the reliability of dosimetric results obtained through the incorporation of dose-volume constraints in our MILP model.

Having said that, there exist additional trends in our dosimetric results. For the large part, the PTV V150 values are lower in automated plans than in manual plans. This trend is also apparent for PTV V200 values between the automated and manual plans. These results tend to suggest that our automated planning approach generally produces treatment plans with “cooler” radiation distribution throughout the PTV. On the other hand, Urethra D5 values seem to be slightly higher in automated plans as opposed to their manual counterparts. A similar, but less pronounced theme is also present for Rectum D1cc values. The clinical implications of these trends are uncertain, considering that despite leading to slightly “hotter” radiation distribution around organs-at-risk, the automated plans still satisfy the dose-limits set for Urethra D5 and Rectum D1cc.

The lower urethral and rectal radiation doses achieved by the manual plans may be attributed to the extensive use of seed positions falling outside of the regular positions (holes) of the physical template grid. This situation is captured in Figure 6.1, and is frequently observed around the urethral pathway and on occasion near the rectum in manual plans. The outcome is that entire strands, and in effect their associated seeds, are pulled slightly away from their theoretical positions (which would normally be located closer to the organs-at-risk). This scenario occurs in 9 out of the 20 manual plans, with 22 seeds being placed out of position for Patient 9 in particular.

While radiation oncologists at the local clinic are comfortable placing strands in positions that do not perfectly conform to the physical template grid, such irregular coordinate points cannot be incorporated into the grid system used in our model. Therefore, none of the automated plans in this study contain seeds that were placed out of position with respect to the coordinate system of the physical template grid.

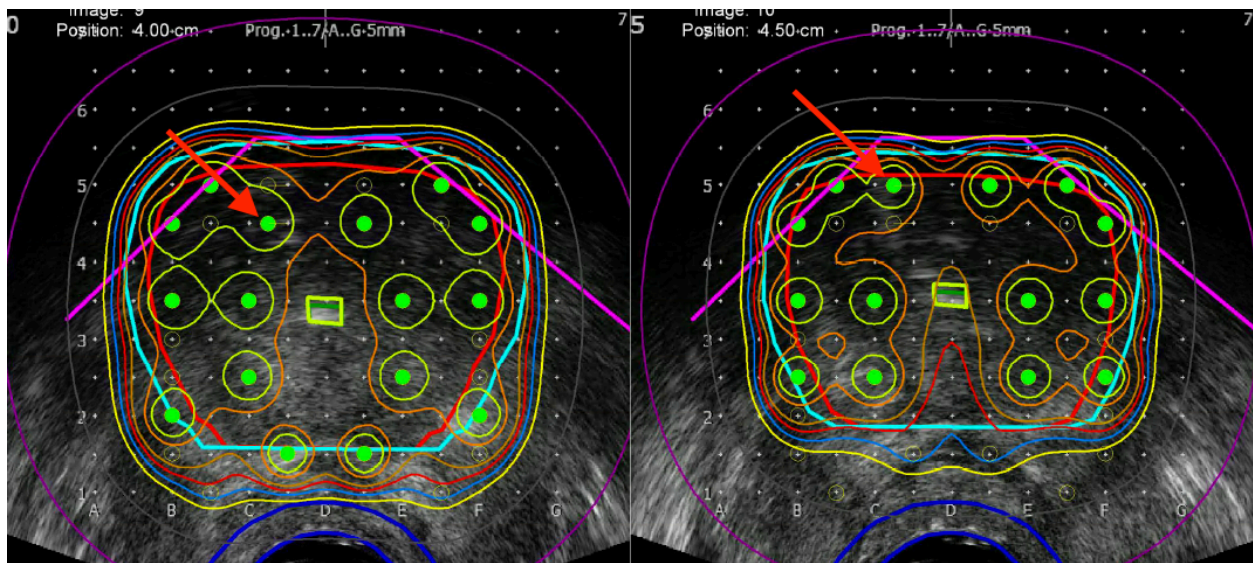


Figure 6.1 Placement of seeds in off-grid positions as indicated by the red arrows in a sample manual plan.

While less of a concern clinically, the manual plans do also possess an abundance of adjacent as well as asymmetrically positioned seeds. Specifically, 15 out of the 20 manual plans contain adjacent seeds, with Patient 8 plan notably possessing 36 adjacent seeds out of 113 total seeds. Furthermore, 11 out of the 20 manual plans exhibit asymmetrical tendencies within their seed distribution, with the plans for Patients 19 and 20 possessing 35 and 41 asymmetrically positioned seeds in particular (out of 89 and 99 total seeds respectively). On the other hand, only a single instance of adjacent seeds is encountered among the automated plans while no instances

of asymmetrically positioned seeds are reported. These results serve to highlight the fact that the use of adjacent and asymmetrically positioned seeds are acceptable at the local clinic in order to accommodate geometrical variation in patient prostates. Furthermore, these results also accentuate the power of our automated planning approach in formulating entirely symmetrical seed distributions (with minimal use of adjacent seed positions) for prostates in which expert planners are forced to make extensive use of asymmetrical and adjacent seed positions to achieve appropriate dose coverage. This feature of our automated planning system may be of interest to brachytherapy teams that desire highly symmetrical seed distributions in their treatment plans.

In terms of pubic arch interference, both the manual and automated plans fared relatively equally. While potential pubic arch interference was observed in 6 out of the 20 automated plans, this number was 5 out of 20 for the manual plans. It should be noted that any form of contact between a planned strand location and the pubic arch contoured by the implanting radiation oncologist was deemed sufficient enough to result in potential pubic arch interference in our assessment.

Moving onto the clinical analysis of the manual and automated plans (as conducted by the five radiation oncologists involved in this study), we observe a mixture of results. For the most part, the manual plans consistently achieved results that were considered “Acceptable (4)” or “Excellent (5)” across the various categories that were used for assessing treatment plan quality, as seen in Tables 6.2 and 6.3. In particular, manual plans made efficient use of needles, comfortably achieved adequate dose coverage of the PTV, avoided overdosing of critical structures and steered clear of pubic arch interference (with average ratings of 4.26, 4.23, 4.22, and 4.18 out

of 5.00 respectively, across the twenty patient prostates examined). Manual plans also performed exceptionally well in achieving an acceptable distribution of strands and seeds, as well as avoiding placement of seeds in the bladder (with average ratings of 4.15 and 4.07 out of 5.00, respectively). Not strictly limiting strand placement to within the prostate was the only indicated source of clinician dissatisfaction with the manual plans (average rating of 3.88 out of 5.00). On average, the manual plans were considered to be clinically acceptable in 90% of the analyses conducted by the five radiation oncologists, with an average treatment plan quality rating of 3.93 out of 5.00.

The assessment involving automated plans produced slightly varied outcomes, as observed in Tables 6.4 and 6.5. Similar to manual plans, automated plans excelled in avoiding pubic arch interference and achieving adequate dose coverage of the PTV (with average ratings of 4.05 and 4.00 out of 5.00 respectively). In regards to economical use of needles, critical structure overdosing and acceptable distribution of strands and needles, automated plans were evaluated to slightly lower average ratings (3.96, 3.90 and 3.69 out of 5.00 respectively) in comparison to their manual counterparts. Finally, the two key problematic areas for automated plans stood out as seed placement in the bladder and strand placement outside of the PTV (resulting in average ratings of 3.36 and 3.19 out of 5.00 respectively). The automated plans were considered to be clinically acceptable in 51% of the analyses carried out by the oncologists, with an average treatment plan quality rating of 3.09 out of 5.00.




Figure 6.2 Trends in PTV V100 values of manual and automated plans.




Figure 6.3 Trends in PTV V150 values of manual and automated plans.




Figure 6.4 Trends in PTV V200 values of manual and automated plans.




Figure 6.5 Trends in PTV D90 values of manual and automated plans.




Figure 6.6 Trends in Urethra D5 values of manual and automated plans.




Figure 6.7 Trends in Rectum D1cc values of manual and automated plans.




Figure 6.8 Trends in asymmetrically positioned seeds in manual and automated plans.




Figure 6.9 Trends in seeds placed in off-grid positions in manual and automated plans.




Figure 6.10 Trends in adjacent placed seeds in manual and automated plans.




Figure 6.11 Trends in strands with potential pubic arch interference in manual and automated plans.

Table 6.2 Details of the clinical quality analyses conducted on manual plans by five radiation oncologists.

Criterion to assess plans	Patient No.																			
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Adequate dose coverage of PTV	4.20	4.20	4.20	3.60	4.40	4.60	4.60	4.20	4.20	4.20	4.20	4.20	4.20	4.20	4.20	4.00	4.20	4.20	4.60	4.20
Critical structures not overdosed	4.20	3.80	4.20	4.20	4.60	4.60	4.40	3.80	4.40	3.80	4.20	4.20	4.20	4.20	4.20	4.20	4.20	4.40	4.40	4.20
No pubic arch interference	4.20	4.40	3.40	4.40	4.60	4.60	4.60	4.20	3.60	4.20	4.60	4.40	4.40	4.20	4.40	3.00	3.80	4.20	4.00	4.40
Strands placed within the prostate	4.00	3.40	3.60	3.80	4.40	4.80	4.20	4.00	3.80	3.80	4.00	3.60	4.00	4.00	3.80	3.40	4.00	3.60	3.60	3.80
No seeds placed in the bladder	3.60	3.40	4.00	4.40	4.40	4.60	4.40	4.20	4.00	3.80	4.40	4.20	4.00	4.20	3.40	3.80	4.40	4.00	4.00	4.20
Efficient use of needles	4.00	4.20	4.25	4.20	4.40	4.60	4.40	4.20	4.20	4.20	4.20	4.40	4.20	4.20	4.40	4.20	4.20	4.20	4.40	4.20
Acceptable distribution of strands/seeds	4.00	4.20	3.80	4.20	4.60	4.60	4.40	4.20	4.20	4.20	3.80	4.20	4.00	4.20	4.20	3.20	4.20	4.20	4.40	4.20
Maximum	4.20	4.40	4.25	4.40	4.60	4.80	4.60	4.20	4.40	4.20	4.60	4.40	4.40	4.20	4.40	4.20	4.40	4.40	4.60	4.40
Mean	4.03	3.94	3.92	4.11	4.49	4.63	4.43	4.11	4.06	4.03	4.20	4.17	4.14	4.17	4.09	3.69	4.14	4.11	4.20	4.17
Minimum	3.60	3.40	3.40	3.60	4.40	4.60	4.20	3.80	3.60	3.80	3.80	3.60	4.00	4.00	3.40	3.00	3.80	3.60	3.60	3.80
Standard Deviation	0.20	0.38	0.31	0.28	0.10	0.07	0.13	0.15	0.26	0.20	0.24	0.25	0.14	0.07	0.34	0.45	0.18	0.24	0.32	0.17
Overall plan quality	3.60	3.50	3.20	4.20	4.40	4.25	4.25	4.20	3.80	3.80	4.00	4.20	4.00	4.25	4.20	2.50	4.20	3.75	4.00	4.20
Plan is clinically acceptable? (%)	80	80	60	80	100	100	100	100	80	100	80	100	100	100	100	60	100	100	80	100

Rating Colour Scale:

	(5) Excellent
	(4) Acceptable
	(3) Indeterminate
	(2) Sub-optimal
	(1) Unacceptable

Table 6.3 Summary of the clinical analyses conducted on manual plans by five radiation oncologists.

Criterion to assess plans	Minimum	Mean	Maximum	Standard Deviation	Percentage (%) of plans* receiving a rating of "Unacceptable (1)"
Adequate dose coverage of PTV	3.60	4.23	4.60	0.21	0
Critical structures not overdosed	3.80	4.22	4.60	0.22	0
No pubic arch interference	3.00	4.18	4.60	0.42	1
Strands placed within the prostate	3.40	3.88	4.80	0.32	0
No seeds placed in the bladder	3.40	4.07	4.60	0.33	0
Efficient use of needles	4.00	4.26	4.60	0.13	0
Acceptable distribution of strands/seeds	3.20	4.15	4.60	0.30	0
Overall plan quality	2.50	3.93	4.40	0.44	0

*based on a total of 100 plans analyzed by oncologists

	Percentage (%) of plans receiving a "Yes"
Plan is clinically acceptable?	90

Rating Colour Scale:

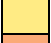
	(5) Excellent
	(4) Acceptable
	(3) Indeterminate
	(2) Sub-optimal
	(1) Unacceptable

Table 6.4 Details of the clinical quality analyses conducted on automated plans by five radiation oncologists.

Criterion to assess plans	Patient No.																			
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Adequate dose coverage of PTV	4.00	4.40	3.80	4.20	4.00	4.20	4.20	4.00	3.40	3.80	3.80	4.00	4.00	4.20	3.80	4.40	4.20	4.40	4.20	3.00
Critical structures not overdosed	3.80	3.80	3.40	4.20	3.80	4.20	4.20	4.20	4.00	3.60	3.60	4.20	4.20	4.20	3.80	3.20	3.60	4.60	3.60	3.80
No pubic arch interference	3.40	4.40	2.40	4.40	4.40	4.40	4.40	4.40	3.20	3.60	3.80	4.40	4.60	4.20	4.20	3.20	4.40	4.60	4.20	4.40
Strands placed within the prostate	3.00	1.80	2.80	3.20	3.40	3.80	3.80	2.60	2.80	3.60	2.80	3.60	3.00	3.00	3.20	3.80	3.20	3.60	3.60	3.20
No seeds placed in the bladder	3.40	2.80	2.80	2.60	3.80	3.60	3.80	2.20	2.60	3.40	3.20	4.00	4.20	3.80	2.80	3.20	4.20	4.20	3.60	3.00
Efficient use of needles	3.80	3.20	3.60	4.20	4.20	4.00	4.20	3.60	3.60	4.20	4.20	4.20	4.20	3.80	4.20	4.20	4.20	4.40	4.20	3.00
Acceptable distribution of strands/seeds	3.50	3.00	3.20	3.80	3.80	4.00	4.00	3.00	3.00	4.00	3.40	4.20	4.40	3.80	4.00	3.80	4.20	4.40	3.60	2.60
Maximum	4.00	4.40	3.80	4.40	4.40	4.40	4.40	4.40	4.00	4.20	4.20	4.40	4.60	4.20	4.20	4.40	4.40	4.60	4.20	4.40
Mean	3.56	3.34	3.14	3.80	3.91	4.03	4.09	3.43	3.23	3.74	3.54	4.09	4.09	3.86	3.71	3.69	4.00	4.31	3.86	3.29
Minimum	3.00	1.80	2.40	2.60	3.40	3.60	3.80	2.20	2.60	3.40	2.80	3.60	3.00	3.00	2.80	3.20	3.20	3.60	3.60	2.60
Standard Deviation	0.31	0.87	0.46	0.61	0.30	0.25	0.21	0.78	0.45	0.26	0.42	0.24	0.48	0.40	0.49	0.46	0.40	0.32	0.30	0.56
Overall plan quality	3.60	3.50	3.20	4.20	4.40	4.25	4.25	4.20	3.80	3.80	4.00	4.20	4.00	4.25	4.20	2.50	4.20	3.75	4.00	4.20
Plan is clinically acceptable? (%)	40	0	0	60	40	100	100	40	20	60	60	80	60	60	0	20	80	100	80	20

Rating Colour Scale:

	(5) Excellent
	(4) Acceptable
	(3) Indeterminate
	(2) Sub-optimal
	(1) Unacceptable



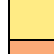
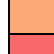

Table 6.5 Summary of the clinical analyses conducted on automated plans by five radiation oncologists.

Criterion to assess plans	Minimum	Mean	Maximum	Standard Deviation	Percentage (%) of plans* receiving a rating of "Unacceptable (1)"
Adequate dose coverage of PTV	3.00	4.00	4.40	0.33	0
Critical structures not overdosed	3.20	3.90	4.60	0.34	4
No pubic arch interference	2.40	4.05	4.60	0.58	1
Strands placed within the prostate	1.80	3.19	3.80	0.48	4
No seeds placed in the bladder	2.20	3.36	4.20	0.59	4
Efficient use of needles	3.00	3.96	4.40	0.37	0
Acceptable distribution of strands/seeds	2.60	3.69	4.40	0.50	0
Overall plan quality	2.00	3.09	4.25	0.67	8

*based on a total of 100 plans analyzed by oncologists

	Percentage (%) of plans receiving a "Yes"
Plan is clinically acceptable?	51

Rating Colour Scale:

	(5) Excellent
	(4) Acceptable
	(3) Indeterminate
	(2) Sub-optimal
	(1) Unacceptable

6.3 Discussion

The results that we have presented in the previous section may be interpreted in numerous ways. The main focus of this section will be automated plans, as the manual plans have produced excellent results overall and present little room for further improvement. It is important to keep in mind that these manual plans were previously used to treat patients at the local clinic, meaning they were already approved by the implanting clinicians. This fact renders the results of the manual plans less surprising. Moving on, there appear to be three key areas to concentrate on with our automated plans. Judging by the ratings received, these three areas of concern are placement of seed strands outside of the prostate, placement of seeds in the bladder and projection of higher-than-desired doses to the urethra and on occasionally the rectum.

Firstly, it may be concluded that our spatial constraints have, in large part, had the desired effect in producing treatment plans with a distinct style that captures what is current practice at the Cross Cancer Institute. Having said that, there are two key opportunities of improvement for our models. The first opportunity concerns the relatively higher number of strands placed outside of the prostate in automated plans. Such an outcome may have been due to our seed-placement limits (and in the correlated strand placement limits) utilized for most of the prostates. Seed placement was generally confined to the PTV (or in many cases to extra margins between 0.1-0.4 cm extending from the PTV, created using slice-based local encapsulation as discussed in Chapter 4), which reflects our initial understanding that placement of strands slightly exterior to the PTV was clinically acceptable. While this appeared to be the case for some prostates, it obviously was not for others as evident from the clinical analysis results obtained. However, the strand placement

issue can be easily addressed by requiring all needles to intersect the prostate within our models. This planning technique may be executed directly through a slight modification of the input data.

The second, and perhaps more critical issue concerning the spatial arrangement of seeds in the automated plans involves the bladder. Several of our automated plans were deemed clinically unacceptable primarily due to the presence of seeds within the bladder (as indicated by multiple ratings of 1.00 for this criterion in some of the individual clinician scoresheets). Figure 6.12 demonstrates a situation in which seed-placement occurs unacceptably close to the bladder in an automated plan, as indicated by the orange arrows.




Figure 6.12 Placement of seeds in the bladder (indicated by the arrows). The prostate is represented by the red contour.




Figure 6.13 Removal of seed-placement positions immediately anterior to the base slice PTV in image (A) creates input data, seen in image (B), that can better guide our MILP model to avoid the bladder.

Such an outcome was not expected, as we initially believed our attempt to limit seed-placement to a tight margin around the PTV in the base slice (through the use of our local-encapsulation method) would be sufficient to protect the bladder. Evidently, this was not always the case. However, before painting this as a deficiency of our model, we must remember that a linear programming model is only as intelligent as the direction or guidance provided to it. As mentioned previously in this thesis, this guidance may be passed on in the form of linear constraints or input data. Considering that the bladder is neither contoured during TRUS imaging by radiation oncologists, nor delineated afterwards by expert planners, our models could not benefit from anatomical data associated with the bladder. Without a clinical source to identify the location of the bladder at our disposal, it is not surprising that our automated planning system was imprecisely guided at times, leading to seeds in the bladder across several plans. As mentioned in Chapter 4, although not explicitly contoured during TRUS imaging, expert planners are still acutely aware of the presence of the bladder in the base slice ultrasound image. Since the possibility of acquiring clinical contour data for the bladder is uncertain, an alternative approach may be followed to safeguard the bladder. Considering that the bladder snugly rests superiorly and anteriorly to the prostate, the simple removal of seed-placement positions immediately anterior to the prostate contour in the base slice (as demonstrated in Figure 6.13) and the neighboring inferior slice would be effective in avoiding the bladder. Once again, this planning technique may be achieved through a simple modification of our input data.

Having looked at the general performance of our spatial constraints and data processing techniques, we now examine how well our use of pseudo high-resolution data fared. Immediately,

we see that our unique technique was highly effective in allowing us to discretely represent anatomical structures with sufficient accuracy. Considering the clinical analysis conducted on the twenty automated plans, the key metric to focus on would be whether these plans achieved adequate dose coverage of the PTV. It is clear that the average rating for PTV dose coverage is high (4.00 out of 5.00), indicating that our approach had the desired effect. Pseudo high-resolution data enabled us to discretely model anatomical structures with enough detail (originally lacking in low-resolution data) without resulting in larger scale models that exhibit intractability. As observed in previous chapters, the use of low-resolution data occasionally resulted in relatively less reliable dose distribution results, both due to its characteristic coarse dose calculation and rough anatomical delineation. Considering the excellent dose coverage achieved through the use of pseudo high-resolution data, it seems as though such issues of reliability have been resolved. This finding suggests that high-resolution data sets may not be an absolute necessity to achieve clinically acceptable results.

Moving on to the performance of our feasibility-based modelling (as opposed to conventional optimality-based approaches for the LDR brachytherapy problem), we can deduce that our approach of searching for feasible solutions is in fact clinically acceptable. This is evident from the numerous automated plans that were deemed acceptable following clinical analysis. Furthermore, it is doubtful if we can shift the blame for the clinical unacceptability of other plans onto the use of feasibility-based modelling. As previously mentioned, several issues arising out of our spatial constraints were key culprits in many of the rejected plans. Strand-placement outside of the prostate and seed-placement in the bladder are input-data related issues and possess no

connection to our decision to not seek a generic clinical goal in the objective function of our model. Perhaps one understandable criticism of our feasibility-based modelling approach may point out the urethral and rectal dose levels, which were relatively higher in the automated plans than the manual plans. On occasion, some of the radiation oncologists indicated that the dose to the critical structures (in the automated plans of specific patients) was relatively higher than the dose levels the clinicians were individually comfortable with. As a result, the associated automated plans were deemed to be clinically unacceptable in the eyes of those oncologists. However, one has to be mindful of the fact that the urethral and rectal dosages, in the majority of our automated treatment plans, were well within the dose-volume requirements of the local cancer centre. As such, overdosing of organs-at-risk was given an average clinical rating of 3.90 for our automated plans (in comparison to 4.22 for manual plans). The clinical concern surrounding the dose delivered to organs-at-risk in the previously mentioned select few automated plans may be that these values are higher than what is normally expected and observed in manual plans by radiation oncologists. Perhaps minimizing average dose to the urethra and rectum, as was done in previous chapters, may have resulted in slightly more clinically satisfactory results. However, it is important to re-iterate the frequent use of seed-placement positions off of the physical template grid, a phenomenon encountered frequently in the manual plans. This technique may have extensively helped lower dose to the organs-at-risk by pulling sources away from their originally intended positions. Considering that our automated plans could not make use of such a work-around, it is uncertain if a fair comparison and critique may be made of the dose delivered to the organs-at-risk between the two planning methods. As a possible remedy, enlarging the protective margins around the

urethra and the rectum, or lowering the dose limits (D5 and D1cc respectively) to these organs may help to improve this aspect of our automated plans.

6.4 Conclusion

The clinical quality and acceptability of I-125 treatment plans, created manually by experts as well as through our automated planning system, were compared in this study. Our results reveal that, unsurprisingly, manual plans used to treat patients fared exceptionally well during clinical analysis. Our results also reveal that it is clinically suitable to create treatment plans through a modelling approach that does not involve optimizing a single (or multiple) overbearing clinical goal(s). The nine clinically accepted automated plans in this study serve as validation of this finding. Furthermore, while demonstrating that feasibility based modelling leads to clinically acceptable solutions, we also observe that plans that fail to do so do not necessarily falter as a result of the search for feasibility in our models. As such, while the use of spatial constraints and data processing techniques have been largely successful in capturing expert planning style at the local clinic, there exist a few areas of further improvement. On the other hand, the use of pseudo high-resolution data has led to clinically acceptable dose coverage of the PTV in automated plans. This is an exciting finding, as such data sets allow us to produce high quality brachytherapy treatment plans without the need of solving much larger scale models.

A final note to point out is that the automated treatment plans examined in this study did not receive any expert planner review or modification prior to the analysis conducted by radiation oncologists. It is estimated that roughly 5-10% of treatment plans created by expert planners receive further modification by oncologists at our local cancer clinic. Therefore, gaining clinical

acceptability in 51% of the plans sourced straight from our automated planning system may be deemed a formidable achievement in our initial attempt at capturing expert planning style. With the suggested improvements implemented in the design of our automated planning system, we believe these clinical acceptability levels may experience exceptional increase. Finally, regardless of the extensivity of features that may be incorporated into a model, it cannot avoid producing treatment plans that may not always satisfy clinicians. Indeed, even plans created by expert planners are not always found to be clinically acceptable without modification. This is expected, as expert planners and clinicians within a clinic may also exhibit different preferences within the overall planning style they are comfortable with. Considering that in a real-world setting our automated plans would receive multiple reviews by expert planners and oncologists before being approved, we believe that slight dissatisfactions may be easily addressed through additional minor modifications.

Chapter 7

Conclusion & Future Work

In this thesis, we set out to introduce a novel modeling philosophy for interstitial LDR prostate brachytherapy in an effort to fill a gap that currently exists in today's literature. Prior to our work, spatial constraints and data processing techniques had not been comprehensively investigated in an attempt to explicitly capture expert planning style. The premise was that our research would bridge the disconnect that currently exists between the often disappointing solutions that available optimization approaches offer and the clinically acceptable results that brachytherapy teams desire from such automated approaches. The fundamental goal of this thesis was to lay the foundation for a new MILP modeling approach that would enable automated construction of treatment plans that were of equivalent clinical quality to manually-produced treatment plans created by expert planners. Furthermore, we also hoped to create such treatment plans in a matter of seconds or minutes, so that our automated planning system would be suitable for both pre-operative and intra-operative planning. However, as previously stated in the introduction, we will remind ourselves of the objectives of this thesis in more specific terms:

- a. Formulate a novel mixed-integer linear programming model as well as practical data processing techniques for LDR prostate brachytherapy, in order to achieve a realistic

mathematical representation of the treatment planning style that is predominantly utilized at the Cross Cancer Institute.

- b. Develop and employ unique linear programming modelling techniques for simplifying the solution of the resulting large-scale model, in order to be able to produce treatment plans with short solution times that are suitable for both pre-operative and intra-operative planning.
- c. Verify the clinical acceptability of our automated treatment plans based on the clinical standards of the Cross Cancer Institute and develop pathways of further improvement for our automated planning system.

We can see that each of these objectives were progressively accomplished throughout this thesis and our main contributions to brachytherapy optimization literature can be summarized as follows:

- In Chapters 3 and 4, we introduced a novel mixed-integer linear programming model with an extensive catalogue of spatial constraints and data processing techniques. Through the results obtained in these chapters, we demonstrated the ability to capture the treatment planning style of a cancer clinic (specifically, the Cross Cancer Institute).
- In Chapter 5, we introduced pseudo high-resolution data sets and feasibility-based modeling (based on the violation of dose-volume constraints). Through the synergistic effects of these two modelling techniques (which were not previously investigated in brachytherapy optimization literature), we demonstrated the ability to produce treatment plans with

solution times suitable for pre-operative and intra-operative planning (specifically, ranging from less than a minute to roughly five minutes for prostates of varying shapes and sizes).

- Finally in Chapter 6, we verified the clinical acceptability of our automated plans through a pilot study involving data from twenty patients, which was the first of its kind in brachytherapy optimization literature. Through this study, we also identified three main areas of improvements to enhance the clinical utility of our automated planning system.

Thus, it can be concluded that both the clinical as well as the engineering goals of this thesis were met. We will now consider several opportunities for further research that may branch out from this thesis.

7.1 Pathways to Future Research

Now we ask ourselves the question of where to take this research from this point forward. Perhaps the most obvious path would be to involve more radiation oncologists in the clinical analysis of the same set of twenty automated and manual plans examined in this thesis. A bigger study may further confirm the initial clinical performance results for our automated planning system, or may reveal newer trends. Following such a study, all necessary improvements (both mentioned in this thesis and also concluded from the bigger study) may then be implemented in our linear programming model.

Solving more test cases of formulations (1)-(6) from Chapter 5, but with different clinical parameter values might shed light on the solution performance of these models and their ability to handle input variation. Although not reported in this thesis, we have previously observed numerous cases of intractability with formulations (1)-(4) after PTV D90 dose levels where changed from

174-189 Gy (as was used in Chapter 5) to 181.5-189 Gy (as was used in Chapters 3 and 4). This was also the case for a change in PTV V100 from >98% to >99%. Formulations (5)-(6) were considerably more reliable in producing feasible solutions using these new dose-limits. Once again, formulation (6) may be made the primary focus of such a study, as performance data concerning solution speed and response to input variation may further emphasize our automated system's suitability for pre-operative and intra-operative planning.

We mentioned in Chapter 3 that treatment plans consisting of symmetrical seed distributions are not uncommon in clinical practice. In some centers, symmetric distributions are highly desired, in part due to their apparent dosimetric robustness. As discussed previously in this thesis, slight seed implantation errors do occur during the brachytherapy operation, arising from a combination of imaging uncertainties, prostate deformation during needle penetration as well as gland swelling due to trauma. Based on clinical experience, symmetrical seed distributions are considered to be dosimetrically more robust in comparison to seed distributions that are highly conformal to the shape of the prostate. In other words, post-operative analysis of dose-volume metrics generally suggest that symmetrical seed distributions are less prone to undesirable dose deviations due to misplacement of seeds, in comparison to highly-conformal seed distributions. However, we believe a study (whether clinical or theoretical in nature) concerning such a phenomenon does not exist in literature. Therefore, a study comparing the dosimetric robustness of automated plans containing symmetrical and asymmetrical seed distributions may be carried out. Final seed positions may be varied in both types of plans through a probabilistic numerical method. The resulting dose

distributions may then be compared between the two families of plans, allowing for inferences to be made on their dosimetric performance.

Now that we have taken an extensive look at how brachytherapy is pre-operatively planned at the Cross Cancer Institute, the next logical step would be to examine intra-operative planning. Tom Baker Cancer Centre in Calgary may be a suitable candidate for this purpose. Collaborating with experts at this clinic would allow us to determine how well our automated planning system performs intra-operatively, both in terms of the clinical quality of treatment plans and their solution times. Furthermore, our model may also be further enhanced through modifications necessary to accommodate the general planning style of Tom Baker Cancer Centre.

Lastly, as mentioned several times in this thesis, the modelling techniques we introduced for low dose rate prostate brachytherapy may be applicable for other variants of prostate brachytherapy (high dose rate, medium dose rate and pulse dose rate), as well as in the brachytherapy treatment of breast, cervical, lung, and head and neck cancer. Thus, branching out towards these areas of brachytherapy may be a potential option. Finally, as our proposed modelling techniques are relatively algorithm-independent, we may also examine the solution performance of our linear programming model as solved through alternative deterministic, stochastic or heuristic solution algorithms. Considering the plethora of algorithms that exist for the solution of mixed-integer linear programming models, we may potentially uncover better performance results with other solution methods.

References

Aikens, C.H. 1985, "Facility location models for distribution planning", *European Journal of Operational Research*, vol. 22, no. 3, pp. 263-279.

American Cancer Society 2015, "Cancer Facts & Figures 2015". Available from:

<<http://www.cancer.org/acs/groups/content/@editorial/documents/document/acspc-044552.pdf>>. [14 April 2015].

D'Souza, W.D., Meyer, R.R., Thomadsen, B.R. & Ferris, M.C. 2001, "An iterative sequential mixed-integer approach to automated prostate brachytherapy treatment plan optimization", *Physics in Medicine and Biology*, vol. 46, no. 2, pp. 297-322.

Danna, E., Rothberg, E. & Le Paper, C. 2005, "Exploring relaxation induced neighborhoods to improve MIP solutions", *Mathematical Programming*, vol. 102, no. 1, pp. 71-90.

Dantzig, G.B., Orden, A. & Wolfe, P. 1955, "The generalized simplex method for minimizing a linear form under linear inequality restraints", Pacific Journal of Mathematics, vol. 5, no. 2, pp. 183-195.

Dantzig, G.B. & Wolfe, P. 1960, "Decomposition principle for linear programs", Operations Research, vol. 8, no. 1, pp. 101-111.

De Boeck, L., Beliën, J. & Egyed, W. 2014, "Dose optimization in high-dose-rate brachytherapy: A literature review of quantitative models from 1990 to 2010", Operations Research for Health Care, vol. 3, no. 2, pp. 80-90.

Elekta 2015, "LDR Real-Time Prostate Solution", Available from:

<<https://www.elekta.com/dms/elekta/elekta-assets/Elekta-Brachytherapy/real-time-prostate-solutions/pdfs/888-00335-MKT--02--LDR-Real-time-Prostate-Brochure/LDR%20Real-time%20Prostate%20Solution%20brochure.pdf>>. [20 August, 2015]

Fourer, R., Gay, D.M. & Kernighan, B.W. 2002, "AMPL: A Modeling Language for Mathematical Programming", Duxbury Press, Cole Publishing Co.

- Fu, L., Ng, W.S., Liu, H., O'Dell, W., Rubens, D., Strang, J., Schell, M.C., Brasacchio, R., Liao, L., Messing, E. & Yu Y. 2005, "Bouquet brachytherapy: feasibility and optimization of conically spaced implants", *Brachytherapy*, vol. 5, no. 1, pp. 59-63.
- Geoffrion, A.M. 1972, "Generalized Benders decomposition", *Journal of Optimization Theory and Applications*, vol. 10, no. 4, pp. 237-260.
- Giantsoudi, D., Baltas, D., Karabis, A., Mavroidis, P., Zamboglou, N., Tselis, N., Shi, C. & Papanikolaou, N. 2013, "A gEUD-based inverse planning technique for HDR prostate brachytherapy: feasibility study", *Medical Physics*, vol. 40, no. 4, pp. 04170
- Gomez-Veiga, F., Mariño, A., Alvarez, L., Rodriguez, I., Fernandez, C., Pertega, S. & Candal, A. 2012, "Brachytherapy for the treatment of recurrent prostate cancer after radiotherapy or radical prostatectomy", *BJU International*, vol. 109, no.1, pp. 17-21.
- Gorissen, B.L., Den Hertog, D. & Hoffmann, A.L. 2013, "Mixed integer programming improves comprehensibility and plan quality in inverse optimization of prostate HDR brachytherapy", *Physics in Medicine and Biology*, vol. 58, no. 4, pp. 1041-1057.
- Gurobi Optimization. 2014, "Gurobi Optimizer Reference Manual", Gurobi Optimization, Inc.

- Guthier, C., Aschenbrenner, K.P., Buerge, D., Ehmann, M., Wenz, F. & Hesser, J.W. 2015, "A new optimization method using a compressed sensing inspired solver for real-time LDR-brachytherapy treatment planning", *Physics in Medicine and Biology*, vol. 60, no. 6, pp. 2179-2194.
- Heidenreich, A., Aus, G., Bolla, M., Joniau, S., Matveev, V.B., Schmid, H.P. & Zattoni, F. 2008, "EAU Guidelines on Prostate Cancer", *European Urology*, vol. 53, no. 1, pp. 68-80.
- Holm, A. 2013, "Mathematical Optimization of HDR Brachytherapy", Ph.D Dissertation, Linköping University.
- Karabis, A., Belotti, P. & Baltas, D. 2009, "Optimization of catheter position and dwell time in prostate HDR brachytherapy using HIPO and linear programming", *IFMBE Proceedings*, vol. 25, no. 1, pp. 612–615.
- Kono, Y., Kubota, K., Aruga, T., Ishibashi, A., Morooka, M., Ito, K., Itami, J., Kanemura, M., Minowada, S. & Tanaka, T. 2010. "Swelling of the Prostate Gland by Permanent Brachytherapy May Affect Seed Migration", *Japanese Journal of Clinical Oncology*, vol. 40, no. 12, pp. 1159-1165.

- Land, A.H. & Doig, A.G., 1960. "An automatic method of solving discrete programming problems", *Econometrica: Journal of the Econometric Society*, vol. 23, no. 8, pp.497-520.
- Lee, E.K. & Zaider, M. 2003, "Mixed Integer Programming Approaches to Treatment Planning for Brachytherapy - Application to Permanent Prostate Implants", *Annals of Operations Research*, vol. 119, no. 1-4, pp. 147-163.
- Lee, E.K., Gallagher, R.J., Silvern, D., Wu, C.-. & Zaider, M. 1999, "Treatment planning for brachytherapy: An integer programming model, two computational approaches and experiments with permanent prostate implant planning", *Physics in Medicine and Biology*, vol. 44, no. 1, pp. 145-165.
- Lee, S. M. 1972, "Goal programming for decision analysis", Philadelphia, PA: Auerbach Publications.
- Lessard, E., Kwa, S.L., Pickett, B., Roach, M. 3rd. & Pouliot, J. 2006, "Class solution for inversely planned permanent prostate implants to mimic an experienced dosimetrist", *Medical Physics*, vol. 33, no. 8, pp. 2773-82.

- Lin, K., Lee, S.P., Cho, J.S., Reiter, R.E., DeMarco, J.J. & Solberg, T.D. 2007. "Improvements in prostate brachytherapy dosimetry due to seed stranding." *Brachytherapy*, vol. 6, no. 1, pp. 44-48.
- Lin, L., Patel, R.R., Thomadsen, B.R. & Henderson, D.L. 2008, "The use of directional interstitial sources to improve dosimetry in breast brachytherapy", *Medical Physics*, vol. 35, no. 1, pp. 240-247.
- Liu, D. M. C. 2014, "Permanent prostate brachytherapy dosimetry: critical assessments and advancements", Ph.D dissertation, University of Alberta.
- Liu, D., Meyer, T., Usmani, N., Kay, I., Husain, S., Angyalfi, S. & Sloboda, R. 2015, "Implanted brachytherapy seed movement reflecting transrectal ultrasound probe-induced prostate deformation", *Brachytherapy*, vol. 14, no.6, pp. 809-817.
- MathWorks. 2015, "MATLAB Primer", Available from:
<https://www.mathworks.com/help/pdf_doc/matlab/getstart.pdf>. [20 April 2015].
- McGeachy, P., Madamesila, J., Beauchamp, A. & Khan, R. 2015, "An open-source genetic algorithm for determining optimal seed distributions for low-dose-rate prostate brachytherapy", *Brachytherapy*, vol. 45, no. 6, pp. 692-702.

- Meyer, R.R., D'Souza, W.D., Ferris, M.C. & Thomadsen, B.R. 2003, "MIP models and BB strategies in brachytherapy treatment optimization", *Journal of Global Optimization*, vol. 25, no. 1, pp. 23-42.
- Noor-E-Alam, M., Mah, A. & Doucette, J. 2012, "Integer linear programming models for grid-based light post location problem", *European Journal of Operational Research*, vol. 222, no. 1, pp. 17-30.
- Pouliot, J., Tremblay, D., Roy, J. & Filice, S. 1996, "Optimization of permanent 125I prostate implants using fast simulated annealing", *International Journal of Radiation Oncology Biology Physics*, vol. 36, no. 3, pp. 711-720.
- Raben, A., Sammons, S., Sim, S., Chen, H., Hanlon, A., Sarkar, A., Donavanik, V., Grebler, A., Geltzeiler, J., Benge, B., Glick, A., Jacob, D. & Koprowski, P. 2007, "Initial comparison of inverse optimization, modified peripheral technique, and geometric optimization as real-time intraoperative computer planning options for permanent seed implantation of the prostate", *Brachytherapy*, vol. 6, no. 4, pp. 238-245.
- Reed, D.R., Wallner, K.E., Merrick, G.S., Arthurs, S., Mueller, A., Cavanagh, W., Butler, W.B., Ford, E. & Sutlief, S.G. 2007. "A prospective randomized comparison of stranded vs. loose 125I seeds for prostate brachytherapy." *Brachytherapy*, vol. 6, no. 2, pp. 129-34.

Rivard, M.J., Butler, W.M., DeWerd, L.A., Huq, M.S., Ibbott, G.S., Meigooni, A.S., Melhus, C.S., Mitch, M.G., Nath, R. & Williamson, J.F. 2007, "Supplement to the 2004 update of the AAPM Task Group No. 43 Report", Medical Physics, vol. 34, no. 6, pp. 2187-2205.

Rivard, M.J., Coursey, B.M., DeWerd, L.A., Hanson, W.F., Huq, M.S., Ibbott, G.S., Mitch, M.G., Nath, R. & Williamson, J.F. 2004, "Update of AAPM Task Group No. 43 Report: A revised AAPM protocol for brachytherapy dose calculations", Medical Physics, vol. 31, no. 3, pp. 633-674.

Shore, N. 2014, "Management of early-stage prostate cancer", The American Journal of Managed Care, vol. 20, no. 12, pp. S260-S272.

Siegel, R.L., Miller, K.D. & Jemal, A. 2015, "Cancer statistics, 2015", CA: A Cancer Journal for Clinicians, vol. 65, no. 1, pp. 5-29.

Sloboda, R., Usmani, N., Monajemi T. T. & Liu D. M. 2012, "Impact of edema and seed movement on the dosimetry of prostate seed implants", Journal of Medical Physics, vol. 37, no. 2, pp. 81-89.

Van Roy, T.J., 1986, "A cross decomposition algorithm for capacitated facility location", Operations Research, vol. 34, no. 1, pp. 145-163.

- Varian Medical Systems 2015, “VariSeed Treatment Planning System Feature Sheet”, Available from:<https://www.varian.com/sites/default/files/resource_attachments/VariSeedFeatureSheet_RAD4257.pdf>. [20 August, 2015].
- Winston, W.L., Venkataramanan, M. and Goldberg, J.B., 2003. “Introduction to mathematical programming (Vol. 1)”. Pacific Grove, CA: Thomson/Brooks/Cole Publishing.
- Yoo, S., Kowalok, M.E., Thomadsen, B.R. & Henderson, D.L. 2003, "Treatment planning for prostate brachytherapy using region of interest adjoint functions and a greedy heuristic", *Physics in Medicine & Biology*, vol. 48, no. 24, pp. 4077-4090.
- Yoo, S., Kowalok, M.E., Thomadsen, B.R. & Henderson, D.L. 2007, "A greedy heuristic using adjoint functions for the optimization of seed and needle configurations in prostate seed implant", *Physics in Medicine & Biology*, vol. 52, no. 3, pp. 815-528.
- Yu, Y. & Schell, M.C. 1996, "A genetic algorithm for the optimization of prostate implants", *Medical Physics*, vol. 23, no. 12, pp. 2085-2091.

Yu, Y., Anderson, L.L., Li, Z., Mellenberg, D.E., Nath, R., Schell, M.C., Waterman, F.M., Wu, A. & Blasko, J.C. 1999, "Permanent prostate seed implant brachytherapy: report of the American Association of Physicists in Medicine Task Group No. 64", Medical Physics, vol. 26, no. 10, pp. 2054-76.

Zelefsky, M.J., Kuban, D.A., Levy, L.B., Potters, L., Beyer, D.C., Blasko, J.C., Moran, B.J., Ciezki, J.P., Zietman, A.L., Pisansky, T.M., Elshaikh, M. & Horwitz, E.M. 2007, "Multi-institutional analysis of long-term outcome for stages T1–T2 prostate cancer treated with permanent seed implantation", International Journal of Radiation Oncology Biology Physics, vol. 67, no. 2, pp. 327-333.