Generalized Inverse Optimization with Application to Cancer Therapy

by

Taewoo Lee

A thesis submitted in conformity with the requirements
for the degree of Doctor of Philosophy
Graduate Department of Mechanical and Industrial Engineering
University of Toronto

Abstract

Generalized Inverse Optimization with

Application to Cancer Therapy

Taewoo Lee

Doctor of Philosophy

Graduate Department of Mechanical and Industrial Engineering

University of Toronto

2015

Inverse optimization has recently received a growing amount of attention as a data-driven approach to determining parameter values for an optimization problem. In this thesis, we study new inverse optimization methodologies that generalize the traditional method of solving inverse optimization problems and accommodate data that makes the standard method ill-posed. We apply the proposed methodologies to prostate cancer therapy data and provide novel insights for radiation therapy treatment planning.

In the first part of the thesis, we briefly review recent theoretical development of inverse optimization and discuss how this thesis contributes to the literature. In the second part of the thesis, we describe a motivational example in radiation therapy treatment planning and illustrate the problem settings and clinical data.

In the third part of the thesis, we develop generalized inverse linear optimization models. We characterize the relationship between the generalized models and the standard model in the literature. By building on the models in general inverse multiobjective optimization problems, we establish a new connection between inverse optimization and existing multiobjective optimization techniques. We show how our methods can be used for determining objective function weights for radiation therapy treatment planning.

Next, we present a clinical application of the results from the previous part of the thesis, by proposing a statistical model that relates the objective function weights to

a patient's anatomical characteristics. Using the statistical relationship, we propose a prediction model that infers objective function weights from patient anatomy and provide a proof of concept of automated, knowledge-based weight determination for radiation therapy treatment planning.

Finally, we extend the theory of the second part of the thesis by developing generalized inverse convex optimization models. In the multiobjective optimization framework, we propose inverse convex optimization models that preserve the preference ordering among different objectives that is encoded by given data, and compare our models to existing inverse optimization models using prostate cancer therapy data. We present a unifying framework that encompasses many of the inverse optimization models in the literature.

Acknowledgements

No word can describe how grateful I am to my advisor Timothy Chan for his incredible mentorship, support, patience, and encouragement. From day one where we first met in the small room writing the dual problem on the white board, he has always been extremely passionate about my research and supportive of my career, and offered numerous opportunities for me to mature as a researcher and as a person. I learned from you so many things about research, teaching, and life. Thank you for making these years the happiest and most growing time in my life.

I am extremely grateful to Michael Sharpe and Tim Craig for their support with clinical data, for putting up with all my questions, and for always encouraging me. Meetings with you have been so much fun and inspiring, always full with insightful jokes and discussions that have led to many interesting research ideas. I would like to thank Dionne Aleman and Roy Kwon for their guidance and advice that have improved my thesis and widened my view of research, and for being great teaching mentors. I would like to thank Wei Yu for his insightful comments and suggestions and his help for Chapter 6 of this thesis, and Andrew Schaefer for his thorough feedback on my thesis and great encouragement which made me love my research even more. I am also indebted to Chi-Guhn Lee for his support which helped me get off the ground at the U of T. I would like thank Chul B. Park for the valuable life lessons he gave me, Daniel Frances for giving me opportunities with this course where I had so much fun with undergraduate students, and Michael Carter for accommodating me when I was lab-less during our move to RS308. I am also very thankful to Brenda Fung and Donna Liu for their patience and help in so many occasions and their kindness.

My colleagues and friends have been the biggest fans of my research. First I would like to thank my labmates Sarina Turner, Houra Mahmoudzadeh, Brendan Eagen, Philip Mar, Daria Terekhov, Justin Boutilier, Derya Demirtas, Velibor Misic, Auyon Siddiq, Ali Goli, Heyse Li, Mark Michael, Chris Sun, Yifang Liu, Muhannad Hammad, and Deepak

Subramani for all the fun time we have had from the basement of MC to the third floor of RS. Outside the Applied Optimization Laboratory, I would like to thank Jim Kuo, Peter Zhang, Jonathan Li, Mike Kim, Kimia Ghobadi, Jason Lee, Shefali Kulkarni-Thaker, Curtiss Luong, Vahid Roshanaei, Albin Fredriksson, and Rasmus Bokrantz for being my alter egos (either offline or online), for never giving up dragging me out for drinks, and for giving me invaluable advice and encouragement when I was starting out, which I have never forgotten and will never forget. Thank you all.

Lastly, I would like to give my everlasting thanks to my parents and my brother Brian for their boundless love, support, patience and all the opportunities they have given me. I would have never made this journey without their unconditional love and sacrifice.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

Inverse optimization has recently received a growing amount of attention as a data-driven approach to determining the values of modeling parameters for an optimization problem. Traditional inverse optimization typically assumes that the functional form of the optimization problem used in the actual system from which data is obtained is known *a priori*. In this thesis, we develop new inverse optimization methodologies that accommodate data obtained from an unknown system, and demonstrate that our methodologies lead to novel insights for cancer therapy treatment planning.

1.1 Inverse optimization

An inverse problem takes observations from a certain system as input and determines parameter values for a model that describe the system dynamics and are consistent with the original observations. Its history dates back at least to the 19th century when an inverse problem for the calculus of variations was first posed by Darboux (1894). Since then, inverse problems have found a wide range of applications including geophysical science and medical imaging. A well-known inverse problem faced by geophysicists is reconstructing the unobservable transmission time of the seismic waves within a certain area of the Earth's interior using past earthquake patterns (Tarantola, 1987). In medical

imaging, the inverse problem is used to reconstruct an image of an organ inside the human body from noninvasive measurements such as x-ray absorption (Arridge, 1999).

In the last two decades, there has been an increasing amount of interest in inverse optimization problems in the operations research community. The inverse optimization problem seeks to determine the values of the parameters of an optimization problem – e.g., the cost coefficients, right-hand-side vector, and constraint matrix – that make a given feasible solution optimal. The original optimization problem for which such parameters are determined is referred to as the "forward" optimization problem. General frameworks for inverse linear optimization and inverse convex optimization were presented by Ahuja and Orlin (2001) and Iyengar and Kang (2005), respectively. A review of inverse optimization theory is provided in Chapter 2.

One common assumption in the majority of the literature on inverse optimization is that the functional form of the objective function and the feasible region of the forward formulation that was used to generate a given solution are known *a priori*. Since an inverse optimization problem is directly derived from the optimality conditions for the underlying forward optimization formulation, this assumption implies that the corresponding inverse optimization problem can always find some cost vector that makes the given solution exactly optimal. However, in practice, there is no guarantee that such information is available – one can only guess the functional form of the forward formulation, which may be an approximation to the actual system that generated the given solution. Moreover, the given solution may have been adjusted post-hoc before implementation in a practical application and no longer be optimal for the forward formulation. The mismatch between the assumed and actual forward formulations and the noise or errors in a given solution cause the inverse optimization problem to be ill-posed or infeasible. In this thesis, we address this issue by generalizing the traditional approach to solving inverse optimization problems, so that even if the data does not "fit" the assumed forward model, the inverse optimization problem would still be able to return meaningful

solutions.

1.2 Intensity-modulated radiation therapy treatment planning

Throughout the thesis, we apply the proposed generalized inverse optimization methods to prostate cancer therapy data. We specifically use data from intensity-modulated radiation therapy (IMRT). IMRT is an advanced cancer treatment technology that uses external beams of high energy x-rays to deliver radiation to a tumour, and is one of the primary treatment techniques for prostate cancer (Foroudi et al., 2003). In IMRT, each radiation beam is modeled as a set of many small beamlets, whose intensities can be optimized to construct a dose of radiation that conforms to the shape of the tumour. The resulting beamlet intensity profile is converted into a collection of deliverable aperture shapes using a multileaf collimator (MLC). In IMRT treatment planning, optimization plays a crucial role in determining various treatment configurations such as the number of radiation beams, the angles of the beams, and the intensity of the beamlets that constitute a beam (Shepard et al., 1999). In our applications, we assume the number and the angles of beams are fixed and focus on optimizing the beamlet intensities. For a review of IMRT, see Bortfeld (2006).

Since treatments need to balance conflicting goals such as delivering a high dose to the tumour while keeping healthy tissue dose low, IMRT treatment planning is typically modeled as a multiobjective optimization problem. The standard approach to solving these problems is to construct a single composite objective function by taking a weighted sum of the individual objectives (Webb, 1994; Xing et al., 1999). The weight values assigned to the objective functions are referred to as "importance factors" in the medical literature (Webb, 1994; Xing et al., 1999). The effect of varying objective function weights can be quite significant, with different sets of weights resulting in significantly different

dose distributions (Shepard et al., 1999). Despite the critical role of weights in optimizing a radiation therapy treatment, the process of determining appropriate weight values is typically done in a trial-and-error manner in practice, with minimal scientific guidance. Moreover, treatment planners often do not know which objectives are most influential to the final treatment quality, which leads to a system with a large number of objectives, i.e., a large parameter space to search over. A more thorough description of the treatment planning system can be found in Chapter 3.

We apply our inverse optimization methodologies to historical IMRT treatment plans to reverse-engineer the most influential objectives and the corresponding weights. We show that the use of inverse optimization can support the development of knowledge-based treatment planning, which is an increasingly popular planning paradigm that uses a database of historical treatments to derive dose objectives and plan parameters for a new patient (Wu et al., 2011; Moore et al., 2011; Chanyavanich et al., 2011).

1.3 Contributions

We view our contributions as follows:

1. We generalize the standard method of solving the inverse linear and convex optimization problems in situations that would otherwise lead to an ill-posed or infeasible inverse problem. Our generalized inverse optimization models retain the complexity of the corresponding forward optimization models as well as the standard inverse optimization models, and specialize to the standard models when the given data perfectly "fits" the underlying forward model. We perform theoretical analysis of the generalized inverse process and characterize the relationship between our generalized models and the standard models in the literature.

2. We develop the first inverse optimization methodology for linear and convex multi-objective optimization. We incorporate the idea of preserving the decision maker's

initial preference ordering among different objectives into the general inverse optimization framework. We establish a connection between inverse optimization and existing multiobjective optimization techniques.

3. We present the first application of inverse optimization to radiation therapy treatment planning. We show how inverse optimization can be used to identify the most critical objectives as well as the corresponding objective function weights for prostate cancer treatment planning.

4. We develop the first statistical model that relates objective function weights for radiation therapy treatment planning to a patient's anatomical characteristics. We build a prediction model that infers the weights for prostate cancer treatment planning from patient anatomy and provide a proof of concept of automated, knowledge-based weight determination.

1.4 Overview of thesis

The structure of the thesis is as follows. In Chapter 2, we review the literature on inverse optimization and elucidate its connection to inverse problems in other fields of research as well as other optimization techniques. In Chapter 3, we describe a motivational example in radiation therapy treatment planning and illustrate the clinical problem settings and data.

Chapter 4 is an expanded version of the paper "Generalized inverse multiobjectve optimization with application to cancer therapy," which has been published in *Operations Research*, Vol. 62. No. 3, 680–695, 2014. In this chapter, we generalize the standard inverse linear optimization approach. We characterize the relationship between our generalized inverse linear optimization models and the standard model of Ahuja and Orlin (2001). We develop our models in the multiobjective framework and establish a new connection between inverse optimization and existing multiobjective optimization

techniques. The multiobjective aspect was motivated by the problem of determining objective function weights for IMRT treatment planning, which we use for demonstrating the proposed inverse models.

Chapter 5 is an expanded version of the paper "Predicting objective function weights from patient anatomy in prostate IMRT treatment planning," *Medical Physics*, Vol. 40, No. 12, 121706, 2013. In this chapter, we present a direct clinical application of the results from Chapter 4. We propose a statistical model that relates the optimal objective function weights for prostate cancer treatment planning determined in Chapter 4 to a patient's anatomical characteristics. We then use this statistical relationship to build a weight prediction model for future patients that uses patient anatomy as a predictor. We demonstrate that inverse optimization can be used for knowledge-based treatment planning.

Chapter 6 is an expanded version of the paper "Preference preservation in inverse multiobjective convex optimization," which is currently under review. In this chapter, we extend the method presented in Chapter 4 by developing generalized inverse convex optimization methods. We develop a generalized inverse convex optimization model that itself is convex, and create a linear approximation to the model as well as apply a successive linear programming algorithm based on the approximation to solve the model. Using data from prostate IMRT treatment planning, we compare our models to existing inverse convex optimization models and examine how the preferences encoded by a given solution are maintained by the different models in the multiobjective optimization framework. Lastly, we present a unifying framework that encompasses many of the inverse optimization models in the literature.

In Chapter 7, we summarize the contributions of the thesis, state future work, and conclude.

Chapter 2

Review of inverse optimization

Since the survey paper on inverse combinatorial optimization by Heuberger (2004), which captured the early stage of development of inverse optimization methodologies, many different extensions and applications of inverse optimization have been studied. In this chapter, we review recent theoretical development of inverse optimization and discuss how this thesis contributes to the literature.

2.1 Preliminaries

We categorize inverse optimization problems into two classes: inverse continuous optimization and inverse discrete optimization. Although some network flow problems are actually continuous linear programming problems, here we classify them as discrete optimization problems because most of the inverse combinatorial or network flow optimization problems have been solved by algorithms using combinatorial arguments rather than pure linear programming approaches. Since the focus of the thesis is inverse linear and convex optimization, in this chapter we pay more attention to the review of inverse continuous optimization.

First, we provide a general, compact representation of the inverse optimization problem. We define our forward optimization problem as FOP($\mathbf{c}$): $\underset{\mathbf{x}\in\mathbf{X}}{\text{minimize}}\ f(\mathbf{x};\mathbf{c})$, where

$\mathbf{c} \in \mathbb{R}^K$ denotes a cost vector and $\mathbf{X}$ denotes the feasible region of $\mathbf{x}$.

Given a solution $\hat{\mathbf{x}} \in \mathbf{X}$, the standard inverse optimization problem can be written as follows:

$$\begin{aligned} \underset{\mathbf{c}}{\text{minimize}} \quad & \phi(\mathbf{c}; \hat{\mathbf{c}}) \\ \text{subject to} \quad & f(\hat{\mathbf{x}}; \mathbf{c}) = \min_{\mathbf{x} \in \mathbf{X}} \; f(\mathbf{x}; \mathbf{c}), \end{aligned} \tag{2.1}$$

where $\hat{\mathbf{c}}$ is an estimated cost vector to which the distance from $\mathbf{c}$ is minimized, measured by $\phi : \mathbb{R}^K \to \mathbb{R}$. We assume $\phi(\mathbf{c}; \hat{\mathbf{c}})$ is minimized when $\mathbf{c} = \hat{\mathbf{c}}$. The objective function $\phi(\mathbf{c}; \hat{\mathbf{c}})$ is used for finding a single inverse solution $\mathbf{c}$ that is closest to the initial belief, when there exist multiple possible $\mathbf{c}$ vectors that satisfy constraints of (2.1).

Problem (2.1) has been studied with different forms of $\phi(\mathbf{c}; \hat{\mathbf{c}})$. Ahuja and Orlin (2001) considered $\sum_{k=1}^{K} |\hat{c}_k - c_k|$ (l_1 norm) and $\max_{k=1,\ldots,K} |\hat{c}_k - c_k|$ (l_∞ norm) which are now standard choices for inverse linear optimization, while Burton and Toint (1992) and Iyengar and Kang (2005) considered $\sum_{k=1}^{K} (\hat{c}_k - c_k)^2$ (l_2 norm). Weighted versions of these norms, e.g., $\sum_{k=1}^{K} w_k |\hat{c}_k - c_k|$ (weighted l_1 norm), have also been used (Ahuja and Orlin, 2001). The Hamming distance objective function has recently received increasing attention, for example, given an indicator function $I_k(\mathbf{c}; \hat{\mathbf{c}}) = 1$ if $c_k \neq \hat{c}_k$ and $I_k(\mathbf{c}, \hat{\mathbf{c}}) = 0$ if $c_k = \hat{c}_k$, Duin and Volgenant (2006) considered $\phi(\mathbf{c}; \hat{\mathbf{c}}) = \sum_{k=1}^{K} w_k I_k(\mathbf{c}; \hat{\mathbf{c}})$. The choice of $\mathbf{c}$ is often restricted by some constraint such as $\mathbf{e}'\mathbf{c} = 1$ where $\mathbf{e}$ denotes the vector of ones, or $\boldsymbol{\ell} \leq \mathbf{c} \leq \mathbf{u}$ where $\boldsymbol{\ell}$ and $\mathbf{u}$ represent appropriate lower and upper bounds, respectively (Heuberger, 2004).

2.2 Inverse continuous optimization

2.2.1 Inverse linear optimization

Inverse linear optimization was first studied by Zhang et al. (1995) in the context of combinatorial optimization. The authors proposed a column generation method for the inverse shortest path problem which can also be generally used for inverse linear opti-

mization. Zhang and Liu (1996) and Zhang and Liu (1999) used the linear programming optimality conditions to derive an inverse formulation for the shortest path problem, which has quickly become a central idea of general inverse linear optimization. Ahuja and Orlin (2001) provided a comprehensive survey on inverse linear optimization.

Let $\mathbf{A} \in \mathbb{R}^{m \times n}, \mathbf{b} \in \mathbb{R}^m$, $\mathbf{c} \in \mathbb{R}^n$, and $\mathbf{x} \in \mathbb{R}^n$. Following the convention of Ahuja and Orlin (2001), the forward linear program is written as follows:

$$\begin{aligned} \underset{\mathbf{x}}{\text{minimize}} \quad & \mathbf{c}'\mathbf{x} \\ \text{subject to} \quad & \mathbf{A}\mathbf{x} \geq \mathbf{b}, \\ & \mathbf{x} \geq \mathbf{0}. \end{aligned} \tag{2.2}$$

A solution $\hat{\mathbf{x}} \in \mathbf{X}$ is optimal for problem (2.2) if and only if there exists $\boldsymbol{\pi} \in \mathbb{R}^m$ such that:

(1) $\mathbf{A}'\boldsymbol{\pi} \leq \mathbf{c}$ and $\boldsymbol{\pi} \geq \mathbf{0}$,

and either

(2a) $(\mathbf{c} - \mathbf{A}'\boldsymbol{\pi})'\hat{\mathbf{x}} = 0$ and $(\mathbf{A}\hat{\mathbf{x}} - \mathbf{b})'\boldsymbol{\pi} = 0$,

or

(2b) $\mathbf{c}'\hat{\mathbf{x}} = \mathbf{b}'\boldsymbol{\pi}$.

The conditions (1), (2a), and (2b) state dual feasibility, complementary slackness and strong duality, respectively. Using complementary slackness, the inverse linear optimization problem that finds $\mathbf{c}$ that renders $\hat{\mathbf{x}}$ an optimal solution can be written as

follows:

$$\begin{aligned} \underset{\boldsymbol{\pi},\mathbf{c}}{\text{minimize}} \quad & ||\mathbf{c}-\hat{\mathbf{c}}||_p \\ \text{subject to} \quad & \sum_{j\in\mathcal{B}} a_{ji}\pi_j \leq c_i, \forall i \text{ such that } \hat{x}_i = 0, \\ & \sum_{j\in\mathcal{B}} a_{ji}\pi_j = c_i, \forall i \text{ such that } \hat{x}_i > 0, \\ & \pi_j \geq 0, \forall j \in \mathcal{B}, \end{aligned} \tag{2.3}$$

where $\mathcal{B} = \{j \mid \sum_{i=1}^{n} a_{ji}\hat{x}_i = b_j\}$. Formulation (2.3) can be formulated as a linear program when $p = 1$ or $p = \infty$.

Inverse linear optimization has exploited the structure of the dual of an inverse formulation due to its similar structure to the original forward formulation. Let y_i for i such that $\hat{x}_i = 0$ be the dual variable associated with the first set of constraints in formulation (2.3), and y_i for i such that $\hat{x}_i > 0$ be the dual variable associated with the second set of constraints in formulation (2.3). With some appropriate substitutions and additional dual variables (Ahuja and Orlin, 2001), the dual of problem (2.3) when $p = 1$ is

$$\begin{aligned} \underset{\mathbf{y}}{\text{minimize}} \quad & \hat{\mathbf{c}}'\mathbf{y} \\ \text{subject to} \quad & \sum_{i=1}^{n} a_{ji}y_i \geq 0, \forall j \in \mathcal{B}, \\ & 0 \leq y_i \leq 1, i \in \{1,...,n\} \text{ such that } \hat{x}_i = 0, \\ & -1 \leq y_i \leq 1, i \in \{1,...,n\} \text{ such that } \hat{x}_i > 0. \end{aligned} \tag{2.4}$$

If we replace the $\mathbf{y}$ with $\mathbf{x} - \hat{\mathbf{x}}$, then formulation (2.4) can be written as

$$\begin{aligned} \underset{\mathbf{x}}{\text{minimize}} \quad & \hat{\mathbf{c}}'\mathbf{x} \\ \text{subject to} \quad & \sum_{i=1}^{n} a_{ji}x_i \geq b_j, \forall j \in \mathcal{B}, \\ & 0 \leq x_i \leq 1, i \in \{1,...,n\} \text{ such that } \hat{x}_i = 0, \\ & \hat{x}_i - 1 \leq x_i \leq \hat{x}_i + 1, i \in \{1,...,n\} \text{ such that } \hat{x}_i > 0. \end{aligned} \tag{2.5}$$

Note that problem (2.5) is very similar to the original forward linear program (2.2).

Coming back to the primal inverse formulation (2.3), $(\boldsymbol{\pi}, \mathbf{c}) = (\mathbf{0}, \mathbf{0})$ is a feasible solution for (2.3). If $\hat{\mathbf{x}}$ is an interior point of $\mathbf{X}$, then $(\boldsymbol{\pi}, \mathbf{c}) = (\mathbf{0}, \mathbf{0})$ is the only feasible and thus the unique optimal solution to (2.3). Although it mathematically makes sense that any feasible solution is an optimal solution if a cost vector is a zero vector (e.g., a feasibility problem), the zero vector is certainly not an interesting solution. In practice, a solution given as input to the inverse problem is not always guaranteed to be an optimal solution to the underlying forward problem, i.e., there may not be a nonzero $\mathbf{c}$ that can make it optimal. This issue has largely been overlooked in the literature.

Troutt et al. (2006) and Troutt et al. (2008) studied problems that can be seen as approximate inverse linear optimization problems that address the issue described above. These papers considered a variation of the standard inverse linear program, which is to determine $\mathbf{c}$ that "best-fits" multiple feasible solutions. While it may not be possible to find $\mathbf{c}$ that renders all the given solutions optimal simultaneously, this method employs the concept of the decisional regret principle to determine $\mathbf{c}$. For each input solution $\mathbf{x}^t, t = 1, \ldots, T$, the decisional regret incurred by a certain cost vector $\mathbf{c}$ is defined as $\mathbf{c}'\hat{\mathbf{x}}^t - \mathbf{c}'\mathbf{x}^*(\mathbf{c})$ in Troutt et al. (2006) and $\mathbf{c}'\hat{\mathbf{x}}^t / \mathbf{c}'\mathbf{x}^*(\mathbf{c})$ in Troutt et al. (2008), where $\mathbf{x}^*(\mathbf{c})$ denotes the actual optimal solution that would have been obtained from the forward problem with $\mathbf{c}$ as a cost vector. Given multiple data points, this type of problem aims to find $\mathbf{c}$ such that the average or maximum decisional regret over all input solutions is minimized. If $T = 1$, i.e., if there is only one input solution, the models of Troutt et al. (2006) and Troutt et al. (2008) become a direct extension to the model of Ahuja and Orlin (2001).

A new, generalized inverse linear optimization method that accommodates any input solution and produces a meaningful inverse solution is one of the major contributions of this thesis, and will be presented in Chapter 4. Our study is different from the approaches of Troutt et al. (2006) and Troutt et al. (2008) in that we build on this approach in general

inverse multiobjective linear optimization problems where objective function weights are not precisely known. We focus on characterizing the structure of the generalized inverse solutions and providing geometrical insight into the mechanism of the proposed models.

2.2.2 Inverse conic/convex optimization

Inverse conic optimization was studied by Iyengar and Kang (2005). Following the convention of Iyengar and Kang (2005), the forward conic programming formulation is written as follows:

$$\begin{aligned} \underset{\mathbf{x}}{\text{minimize}} \quad & f(\mathbf{x};\mathbf{c}) \\ \text{subject to} \quad & \mathbf{g}(\mathbf{x}) \succeq_{\mathcal{K}} \mathbf{0}, \\ & \mathbf{A}\mathbf{x} \geq \mathbf{b}, \\ & \mathbf{x} \geq \mathbf{0}, \end{aligned} \tag{2.6}$$

where $f(\mathbf{x};\mathbf{c}) : \mathbb{R}^n \to \mathbb{R}$ is a linear or quadratic objective function and $\mathbf{g} : \mathbb{R}^n \to \mathbb{R}^s$ is a differentiable, concave vector-valued function. The partial order $\succeq_{\mathcal{K}}$ is defined in a cone $\mathcal{K} \in \mathbb{R}^s$.

Let the Lagrangian be $L(\mathbf{x}, \boldsymbol{\mu}, \boldsymbol{\pi}, \boldsymbol{\nu}) = f(\mathbf{x};\mathbf{c}) - \mathbf{g}(\mathbf{x})'\boldsymbol{\mu} - (\mathbf{A}\mathbf{x} - \mathbf{b})'\boldsymbol{\pi} - \mathbf{x}'\boldsymbol{\nu}$, where $\boldsymbol{\mu} \in \mathbb{R}^s_+$, $\boldsymbol{\pi} \in \mathbb{R}^m_+$, and $\boldsymbol{\nu} \in \mathbb{R}^n_+$. Assume that problem (2.6) satisfies the constraint qualification. Then a given solution $\hat{\mathbf{x}} \in \mathbf{X}$ is an optimal solution to (2.6) if and only if there exist $\boldsymbol{\mu} \succeq_{\mathcal{K}^*} \mathbf{0}$, where $\mathcal{K}^*$ denotes the dual cone of $\mathcal{K}$, and $\boldsymbol{\pi} \in \mathbb{R}^m_+$ such that

(1) $\nabla_{\mathbf{x}} L(\hat{\mathbf{x}}, \boldsymbol{\mu}, \boldsymbol{\pi}) = \nabla_{\mathbf{x}} f(\hat{\mathbf{x}};\mathbf{c}) - \sum_{l=1}^{s} \mu_l \nabla_{\mathbf{x}} g_l(\hat{\mathbf{x}}) - \mathbf{A}'\boldsymbol{\pi} - \boldsymbol{\nu} = \mathbf{0}$,

(2) $\mathbf{g}(\hat{\mathbf{x}})'\boldsymbol{\mu} = 0$,

(3) $(\mathbf{A}\hat{\mathbf{x}} - \mathbf{b})'\boldsymbol{\pi} = 0$, and

(4) $\hat{\mathbf{x}}'\boldsymbol{\nu} = 0$.

The inverse optimization problem of Iyengar and Kang (2005) is to find $\mathbf{c}$ that satisfies the above system of equations with respect to a given solution $\hat{\mathbf{x}} \in \mathbf{X}$ while minimizing

$||\mathbf{c} - \hat{\mathbf{c}}||_p$ for some p. While the focus of Iyengar and Kang (2005) is restricted to conic programs, we note that this model can also be applied to general convex programs. We refer readers who are interested in inverse optimization for more specialized instances of convex programming to Xiao et al. (2009) (inverse semi-definite programming), Zhang et al. (2010) (inverse quadratic programming), and Zhang and Xu (2010) (inverse separable convex programming).

Similar to the discussion in the previous section, an arbitrary $\hat{\mathbf{x}}$ need not be a candidate optimal solution to the forward problem, in which case the inverse problem of Iyengar and Kang (2005) returns a zero cost vector as the unique solution. Keshavarz et al. (2011) presented a problem that can be seen as a generalization of Iyengar and Kang (2005). In the model of Keshavarz et al. (2011), given $\hat{\mathbf{x}} \in \mathbf{X}$, the KKT conditions are relaxed with residuals as follows:

(1) $\nabla_{\mathbf{x}} f(\hat{\mathbf{x}}; \mathbf{c}) - \sum_{l=1}^{s} \mu_l \nabla_{\mathbf{x}} g_l(\hat{\mathbf{x}}) - \mathbf{A}'\boldsymbol{\pi} - \boldsymbol{\nu} = \boldsymbol{\gamma}^{(1)}$,

(2) $g_l(\hat{\mathbf{x}})\mu_l = \gamma_l^{(2)}, l = 1, \ldots, s$

(3) $(\sum_{i=1}^{n} a_{ji}\hat{x}_i - b_j)\pi_j = \gamma_j^{(3)}, \quad \forall j = 1, \ldots, m$, and

(4) $\hat{x}_i \nu_i = \gamma_i^{(4)}, i = 1, \ldots, n$,

which are used as constraints while the objective is to minimize $\phi(\boldsymbol{\gamma}^{(1)}, \boldsymbol{\gamma}^{(2)}, \boldsymbol{\gamma}^{(3)}, \boldsymbol{\gamma}^{(4)})$, where $\phi(\boldsymbol{\gamma}^{(1)}, \boldsymbol{\gamma}^{(2)}, \boldsymbol{\gamma}^{(3)}, \boldsymbol{\gamma}^{(4)}) = 0$ if and only if all the residuals are zeros. The original version of the formulation in Keshavarz et al. (2011) accommodates multiple input data points each with their own residuals, but here we simply illustrate their model with a single data point $\hat{\mathbf{x}}$.

In Chapter 6, we consider a different version of generalized inverse convex optimization. We develop our models in the multiobjective framework again. We introduce a notion of a preference ordering among different objectives, and develop generalized inverse convex optimization models that accommodate any given solution as input and

determine a cost vector that preserves the preference ordering encoded by the input solution. We then compare our models to the model of Keshavarz et al. (2011) and provide a unifying framework for inverse convex optimization.

2.3 Inverse discrete optimization

2.3.1 Inverse combinatorial optimization

Most studies in inverse combinatorial optimization have focused on whether an inverse problem for a certain combinatorial problem can be transformed into another combinatorial problem for which a solution algorithm is well-established. Among a myriad number of different inverse combinatorial optimization problems, here we look into two major applications and provide a list of other applications. An extensive review of inverse combinatorial optimization can be found in Heuberger (2004).

Inverse shortest path problem

Burton and Toint (1992) solved the inverse shortest path problem under the ℓ_2 objective function via a quadratic programming algorithm. Ahuja and Orlin (2001) showed that the inverse shortest path problem under ℓ_1 norm can be solved simply by solving the original forward shortest path problem – existing well-known algorithms such as Dijkstra's algorithm can be used. Similarly, the inverse shortest path problem under the weighted ℓ_1 objective function is equivalent to a general minimum cost flow problem and therefore can be solved by existing efficient algorithms. Duin and Volgenant (2006) proposed an algorithm that solves the inverse shortest path problem under the min-max weighted Hamming distance in polynomial time.

Inverse minimum spanning tree problem

Sokkalingam et al. (1999) showed that the inverse minimum spanning tree problem under ℓ_1 norm is the dual of an assignment problem. Similarly, the weighted ℓ_1 problem is also the dual of a transportation problem. The authors also proposed an algorithm that solves the ℓ_∞ problem in polynomial time. Hochbaum (2003) presented more efficient algorithms for general inverse minimum spanning tree problems. Inverse minimum spanning tree problems under the Hamming distance have been studied by He et al. (2005), Zhang et al. (2006) and Duin and Volgenant (2006).

Other inverse combinatorial optimization problems

Other inverse combinatorial optimization problems that have been studied include inverse assignment problems (Yang, 1997; Zhang and Liu, 1999; Ahuja and Orlin, 2001), inverse minimum cost flow problems (Zhang and Liu, 2002; Jiang et al., 2010), inverse minimum cut/maximum flow problems (Yang et al., 1997; Zhang and Cai, 1998; Ahuja and Orlin, 2002; Liu and Zhang, 2006), inverse traveling salesman problems (Chung and Damange, 2008), and inverse center location problems (Cai et al., 1999; Zhang et al., 2005; Alizadeh et al., 2009; Alizadeh and Burkard, 2011).

2.3.2 Inverse integer/mixed integer programming

Integer and mixed-integer programming problems have been relatively less studied in terms of inverse optimization compared to other areas reviewed above. Schaefer (2009) studied inverse optimization for general integer programming problems. Schaefer (2009) formulated the problem via superadditive integer programming duality. The author showed that the inverse integer progam is NP-hard, and proposed algorithms to solve the problem. Huang (2005) solved an inverse optimization problem for an integer programming problem with a fixed number of constraints by reformulating the problem as an integer shortest path problem.

Wang (2009) presented a cutting plane method for inverse mixed-integer programs. We note that the cutting plane algorithm is also applicable to general inverse linear optimization problems. In fact, a column generation approach, which is a dual approach to the cutting plane algorithm, was one of the earliest inverse linear/combinatorial optimization techniques (Zhang et al., 1995; Yang and Zhang, 1999). To address some issues around computational efficiency and convergence that arose in Wang (2009), Duan and Wang (2011) proposed a new heuristic algorithm for the cutting plane method.

2.4 Connection to other areas

Among many different types of inverse problems studied in various research areas, the framework that is closest to inverse optimization is inverse reinforcement learning, or inverse optimal control, in the computer science community. Inverse reinforcement learning seeks to recover a reward function of a Markov decision process (MDP) model that renders an observed behaviour of an expert an optimal solution. Many variations and extensions to inverse reinforcement learning have been studied in the last decade (Ratliff et al., 2006; Ramachandran and Amir, 2007; Ziebart et al., 2008). Inverse reinforcement learning is also used for apprenticeship learning, which is a type of machine learning technique that uses a reward function approximated by inverse reinforcement learning to teach an agent (Abbeel and Ng, 2004). As it assumes that the expert is acting optimally in a Markov decision process (MDP), inverse reinforcement learning can be seen as an inverse MDP problem. The inverse MDP problem was only recently studied in the operations research community by Erkin et al. (2010) where the underlying MDP was formulated as a linear program. Although MDPs have been extensively studied in the operations research community, we conjecture from the lists of citations in inverse optimization studies that the existence of inverse reinforcement learning has rarely been exposed to the community. Interestingly, both inverse reinforcement learning and inverse

optimization began to emerge in the early 2000's.

Inverse optimization also has a close connection to De Novo programming studied in Zeleny (1986) and Zeleny (2010). De Novo programming is a technique that is used to adjust the right-hand-side vector (i.e., **b** vector) of an optimization formulation under some restriction so that an optimal solution to the modified problem is closest to a specific desirable solution. For example, in a multiobjective optimization problem, given an ideal point – a vector of the best objective function values for each objective independently, which is in general not achievable, – De Novo programming can be used to adjust the **b** vector such that the adjusted optimization problem produces a solution that is closest to the ideal point. We note that this problem can be seen as a generalized inverse optimization problem of recovering the **b** vector given the ideal point as input. If information of the dual solution corresponding to the ideal point is available, then the problem can be reformulated as the traditional inverse problem where the **b** vector is now a cost vector.

Inverse optimization has also turned out to be closely related to multiobjective optimization. Wei et al. (2000) showed that the inverse optimization problem for data envelopment analysis – determining additional inputs such that the current efficiency level of a system remains the same given some increase in output – is equivalent to a forward multiobjective optimization problem. Also, in Chan et al. (2014), which is a part of Chapter 4 of this thesis, the inverse multiobjective optimization problem was shown to be the dual of a special type of the well-known Benson's formulation (Benson, 1978). In general, inverse multiobjective optimization and existing multiobjective optimization techniques seem to share common interests because they both basically search over the normals of supporting hyperplanes of the set of feasible objective function values.

2.5 Summary of contributions to the literature

Our approach to generalizing the standard inverse optimization framework is similar to Troutt et al. (2006), Troutt et al. (2008), and Keshavarz et al. (2011) in that a given solution is near-optimal with respect to the inversely determined cost vector. We develop the first inverse optimization methods for multiobjective linear and convex optimization. Also, by focusing on theoretical analysis of the structure of the proposed models and resulting inverse solutions, our study provides a better understanding of how the proposed generalized models are related to the standard models and how the original preference of a given solution is retained by the inversely determined cost vector, which has not been considered in the literature.

This thesis also provides the first applications of inverse optimization to radiation therapy treatment planning. Erkin et al. (2010) is the only previous application of inverse optimization in healthcare. In Chapter 4, we demonstrate the application of inverse optimization to historical prostate cancer treatment to impute optimal treatment parameters. Then, in Chapter 5, we provide the first study in the medical physics literature that uses the results of inverse optimization for predicting treatment parameters from patient anatomy and creating treatment plans for de novo patients.

Throughout this thesis, we strive to make connections to other research fields and other existing optimization techniques in the literature. In Chapter 4, we elucidate a relationship between inverse optimization and existing multiobjective optimization techniques. In Chapter 6, we propose a unified framework for inverse optimization that encompasses many of the existing inverse optimization models as well as our proposed models. In Chapter 5, we demonstrate how machine learning techniques can build on the results of our inverse optimization methods. To elaborate more on this point, the recent development of inverse reinforcement learning was initially motivated by the fact that it is often rather easy to demonstrate some examples of desirable behaviour and let the agent learn a reward function from the examples, instead of having the agent blindly learn the

expert's behaviour. This motivation for inverse reinforcement learning was actually very similar to what drove us to apply inverse optimization to radiation therapy treatment planning, where there are many factors that make the radiation therapy treatment planning process unpredictable and subjective but there is historical treatment data we can make use of.

Chapter 3

Clinical motivations and background

In this chapter, we describe the clinical motivation for our work, which is a prostate cancer radiation therapy treatment planning problem. We illustrate the current approaches to formulating a radiation therapy treatment plan optimization problem and the associated challenges. We also provide detailed information about our clinical data and the initialization process that is required before the development of the inverse optimization framework for radiation therapy treatment planning.

3.1 Prostate cancer and radiation therapy

Prostate cancer is the most commonly diagnosed cancer among men in North America (American Cancer Society, 2014; Canadian Cancer Society's Advisory Committee on Cancer Statistics, 2014). Radiation therapy is one of the most common forms of treatment for prostate cancer (Foroudi et al., 2003). While survival from prostate cancer post-IMRT is generally high (Zelefsky et al., 2006), radiation-induced toxicity of the healthy organs near the tumour, referred to as organs-at-risk (OARs), is still an important concern. The ability to better differentiate between OAR objectives and encourage treatments to focus on the most critical ones has the potential to reduce secondary toxicity while maintaining tumour dose levels, which is the primary clinical challenge in treating prostate cancer

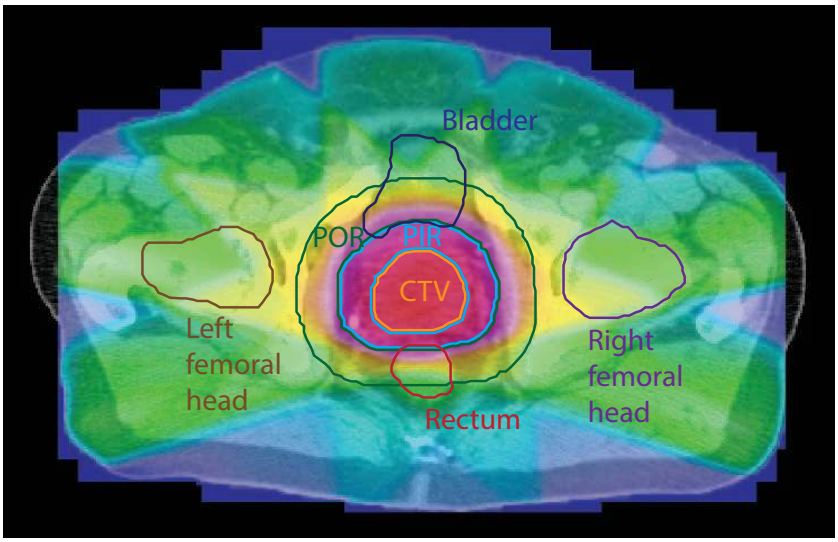

Figure 3.1: Anatomy of prostate site.

currently.

Figure 3.1 depicts the typical anatomical geometry surrounding the prostate. The clinical target volume (CTV) is the volume suspected to contain the primary disease – the prostate in this case. The planning target volume (PTV) encompasses the CTV and accounts for uncertainties such as possible organ movement or variations in the position of the patient on the treatment unit. The PTV inner ring (PIR) is a shell around the CTV defined as the set difference between the PTV and CTV (i.e., PIR = PTV\CTV). The PTV outer ring (POR) is a shell outside the PTV that is used to promote conformality of the dose around the target. Around the prostate, there are multiple healthy organs that need to be spared including the bladder, rectum, and left and right femoral heads (the tops of the femur bones).

Certain clinical criteria must be met in order for a radiation therapy treatment to be acceptable. The clinical criteria are expressed in the form of dosimetric conditions such as "the minimum dose to the tumour needs to be at least x Gy" or "at most y% of the bladder can receive more than x Gy". The latter is known as a partial dose-volume constraint, which can be formulated exactly using binary variables to choose the volume elements that meet the dose threshold (Lee et al., 2003; Preciado-Walters et al.,

Table 3.1: Objective functions and weights in the clinical treatment planning formulation for a prostate case.

Obj. #	Structure	Metric	Dose (Gy)	Percent	Weight
1	CTV	Min dose	78.65		50
2	CTV	Max dose	81.90		38
3	PIR	Min dose	74.75		50
4	PIR	Max dose-volume	78.00	2	38
5	POR	Max dose	74.10		38
6	Bladder	Max dose-volume	47.45	50	1
7	Rectum	Max dose-volume	47.45	50	1
8	Bladder	Max dose-volume	59.15	29	1
9	Rectum	Max dose-volume	59.15	29	1
10	Left femur	Max dose-volume	52.00	4	1
11	Right femur	Max dose-volume	52.00	4	1
12	Bladder	Mean dose	42.90		1
13	Rectum	Mean dose	44.20		1
14	Bladder	Max dose	78.00		1
15	Rectum	Max dose	78.00		1

2004). In this thesis, we address such criteria using a more tractable alternative, as outlined in Sections 4.4 and 6.4. Table 3.1 outlines the objectives used to plan a prostate IMRT treatment at Princess Margaret Cancer Centre in Toronto, Canada. Each row corresponds to an objective function that is combined into the composite objective using the weight indicated in the last column. The "Metric" column indicates the functional form of each criterion. For example, the first two criteria require the dose to every part of the CTV to be between 78.65 and 81.90 Gy. Criteria listed as "Max dose-volume" indicate partial dose-volume criteria. For example, the eighth criterion requires that no more than 29% of the bladder receive more than 59.15 Gy. While the criteria shown in Table 3.1 read like constraints, they are formulated as objectives by penalizing violations of the criteria. For example, objective 1 is formulated as an one-sided penalty, charging a positive penalty to parts of the CTV that receive less than 78.65 Gy, proportional to the square of the difference between the delivered and prescription dose. The tumour objectives receive the highest priority. However, there is no differentiation between the importance of the OAR objectives, even though it is well-known that different organs

respond differently to radiation (Marks, 1996). In addition, the weight values seem to be chosen subjectively.

3.2 The treatment planning process

In practice, an iterative, trial-and-error process is used to determine a final, acceptable treatment plan. The optimization problem represented by Table 3.1 is solved once. Next, the corresponding dose distribution is evaluated to see if it meets the clinical criteria prescribed by the oncologist. If the plan is unacceptable, the treatment planner re-optimizes the plan with different parameter settings. This back and forth process continues until an acceptable treatment plan is found. For complex cases, this process may take days to complete, potentially delaying the start of treatment. Even though many modifications are possible during the re-optimization step (e.g., add new objective functions, modify dose limits, etc.), our focus in this thesis is on the determination of the weights. Here we briefly review the literature on using weights for treatment planning.

Iterative approaches to weight determination (*a priori* approaches). Most research in radiation therapy treatment planning with a weighted objective function views the weights as a tunable parameter. With the goal of creating "good" dose distributions, weights are iteratively updated in a forward manner. For example, Yu (1997), Xing et al. (1999), Cotrutz and Xing (2002), and Wu et al. (2003) present algorithms that adjust weights in an inner loop while an outer loop evaluates various scoring functions associated with the dose distribution generated by those weights.

Pareto treatments (*a posteriori* approaches). The use of weights in IMRT optimization is also prevalent in the construction of sets of Pareto optimal treatments (Cotrutz et al., 2001; Hamacher and Küfer, 2002; Romeijn et al., 2004; Craft et al., 2006). Instead of generating one acceptable treatment, a set of Pareto optimal solutions is obtained by solving the multiobjective IMRT optimization problem

with many different sets of weights. A final treatment is chosen based on the clinical expertise of the treatment planner as s/he navigates the Pareto surface.

Note that, in the medical physics community, the optimization approach to determining the optimal beamlet intensities and dose distribution is called "inverse planning." In this thesis, we use the terms "inverse optimization" and "inverse planning" to refer to two distinct problems: inverse optimization determines weights from a given dose distribution while inverse planning determines a beamlet intensity profile (and thus a dose distribution) from a given set of weights. We also refer to inverse planning as forward optimization.

3.3 Clinical data

All of the computational experiments in this thesis are based on treatment data for prostate cancer patients who had previously received radiation therapy at Princess Margaret Cancer Centre. The patient datasets each had a CTV comprising the prostate gland ($\pm$ proximal 10 mm of seminal vesicles) and a PTV defined as the CTV plus a 10 mm margin (7 mm posteriorly). All treatments were delivered with seven 6 MV step-and-shoot intensity-modulated x-ray fields at angles 40°, 80°, 110°, 250°, 280°, 310°, and 355°. We excluded patients who had previous prostatectomy, pelvic lymph node irradiation, and atypical anatomic features including proximal small bowel, pelvic kidneys, and prosthetic hips. We used CERR (Computational Environment for Radiotherapy Research) to read and analyze the data (Deasy et al., 2003). Treatment data was exported from the Philips Pinnacle treatment planning system in DICOM (Digital Imaging and COmmunications in Medicine) and RTOG (Radiation Therapy Oncology Group) formats, and read into MATLAB via CERR.

To represent a patient's anatomy algebraically, the anatomy is discretized into volume elements called voxels. As an example, patient #1 had $18,549$ voxels in the CTV, $54,964$

in the PIR, 294,075 in the POR, 6,926 in the bladder wall, 8,959 in the rectum wall, 47,061 in the left femoral head, and 47,102 in the right femoral head. The voxel grid resolution of the CT image, and thus our data, was 1 mm × 1 mm × 2 mm. The voxel grid resolution used for prostate treatment planning at Princess Margaret Cancer Centre is 2 mm × 2 mm × 2 mm. To improve tractability, while accounting for the clinical planning resolution, we sampled voxels at a rate of 1:4 in all structures. We found the difference in the resulting dose distributions to be negligible. To further speed up the computation, we considered sampling the POR at a rate of 1:10, since it was the largest structure by far. Again, we found negligible differences in the resulting dose distributions. A total of 354 beamlets were used for patient #1 across the seven beams. The resolution of each beamlet was 5 mm × 5 mm.

3.4 A multiobjective forward formulation

We formulate a multiobjective forward optimization problem that emulates the clinical IMRT treatment planning formulation presented in Table 3.1. Since the tumour dose criteria are given very high weights and typically only a small amount of violation is allowed, we model them as hard constraints. Putting only the OARs in the objective function allows for better differentiation between the OAR criteria.

Let $\mathcal{B}$ be the set of beamlets and w_b be the intensity delivered by beamlet $b \in \mathcal{B}$. As mentioned above, the patient's anatomy is discretized into volume elements called voxels. We denote by $D_{v,b}$ the dose deposited to voxel v from unit intensity of beamlet b. Let $\mathcal{K}$ be the set of all objectives. For any $k \in \mathcal{K}$, let $\mathcal{O}_k$ be the set of voxels in the OAR associated with objective k. We also let $\mathcal{T}$ and $\mathcal{V}$ be the sets of voxels in the target structures (i.e., the CTV and PIR) and the whole anatomy, respectively. Lastly, let $\alpha_k \geq 0$ denote the weight assigned to objective k. The complete forward formulation

is

$$\begin{aligned}
\underset{\mathbf{w}}{\text{minimize}} \quad & \sum_{k\in\mathcal{K}} \alpha_k f_k(\mathbf{w}), \\
\text{subject to} \quad & \sum_{b\in\mathcal{B}} D_{v,b} w_b \geq \ell_v, \quad \forall v \in \mathcal{T}, \\
& \sum_{b\in\mathcal{B}} D_{v,b} w_b \leq u_v, \quad \forall v \in \mathcal{V}, \\
& \frac{\beta_1}{|\mathcal{B}|}\sum_{b'\in\mathcal{B}} w_{b'} \leq \; w_b \leq \frac{\beta_2}{|\mathcal{B}|}\sum_{b'\in\mathcal{B}} w_{b'}, \quad \forall b \in \mathcal{B}, \\
& w_b \geq 0, \quad \forall b \in \mathcal{B},
\end{aligned} \tag{3.1}$$

where ℓ_v is a lower bound on the dose to voxel $v \in \mathcal{T}$ and u_v is an upper bound on the dose to voxel $v \in \mathcal{V}$. The third set of constraints forces every beamlet intensity to be within a certain multiple of the average beamlet intensity ($\beta_1 < 1 < \beta_2$), which discourages a highly heterogeneous intensity map. These constraints act as a proxy for the smoothing of the intensity map that clinical hardware and software enforce. Such a simple smoothing mechanism seems to work well for simple geometries like the prostate, but may be insufficient for complex cases where heterogeneity can be useful, especially in the case of motion. Smoothing methods like constraining/optimizing the total variation or sum-of-positive gradients may be more applicable in other cases (Zhu et al., 2008; Craft et al., 2007).

For the objective functions $f_k(\mathbf{w}), k \in \mathcal{K}$, in Chapters 4 and 5 we consider penalty-based linear objective functions, and in Chapter 6 we consider penalty-based convex objective functions. More sophisticated CVaR constraints (Romeijn et al., 2006), general piece-wise linear objective functions (Craft et al., 2007), and linear EUD objectives (Thieke et al., 2002) can also be incorporated. However, our computational results suggest that formulation (3.1) with objective functions used in this thesis is sufficient to replicate clinical-quality plans.

In all instances, Table 3.1 was used to guide the parameter settings. We set $\ell_v = 78.0$ for $v \in$ CTV and $\ell_v = 74.1$ for $v \in$ PIR. We set $\mathcal{V}$ to be the union of the PTV and all

OARs, as the remaining unclassified tissue had little impact on the final dose distribution. We let $u_v = 81.9$ for $v \in \mathcal{V} \setminus \text{POR}$, and let $u_v = 78.0$ for $v \in \text{POR}$. Lastly, we chose $\beta_1 = 0.5$ and $\beta_2 = 1.5$, which were guided by discussions with medical physicists at Princess Margaret Cancer Centre.

3.5 Initialization

Two pieces of data are needed to formulate the inverse optimization models: the delivered beamlet intensities, and the patient-dependent clinical dose influence matrix, $\mathbf{D}^{clin}$ (see formulation (3.1)) . The clinical plans exported by Pinnacle did not contain either of these quantities, but instead had a delivered dose value for each voxel (i.e., the dose distribution, denoted $\mathbf{d}^{clin}$), anatomical information such as contours and three dimensional voxel coordinates, and the beam angles used for the treatment.

To overcome the lack of $\mathbf{D}^{clin}$, we used CERR's IMRTP function to generate a new dose influence matrix, $\mathbf{D}^{CERR}$. All the structure information and beam environment parameters were set to match the clinical plan. However, because CERR uses a different method to calculate dose influence matrices than a commercial treatment planning system, $\mathbf{D}^{CERR}$ is different from $\mathbf{D}^{clin}$ (Jeraj et al., 2002). As a result, the exact clinical dose distribution obtained using $\mathbf{D}^{clin}$ may not be "achievable" using $\mathbf{D}^{CERR}$. That is, there may not exist $\mathbf{w}$ such that $\mathbf{d}^{clin} = \mathbf{D}^{CERR}\mathbf{w}$. To find such a $\mathbf{w}$, we solve an auxiliary optimization problem to find a dose distribution that is close to $\mathbf{d}^{clin}$ and that is achievable using $\mathbf{D}^{CERR}$:

$$\begin{aligned} \underset{\mathbf{w}}{\text{minimize}} \quad & \sum_{v \in \mathcal{V}} \left(\sum_{b \in \mathcal{B}} D^{CERR}_{v,b} w_b - d^{clin}_v \right)^2 \\ \text{subject to} \quad & \text{constraints from problem (3.1).} \end{aligned} \tag{3.2}$$

A common way to evaluate an IMRT treatment plan is via a dose-volume histogram

(DVH). A DVH shows what fraction of a particular structure receives a certain level of dose or higher. Figure 3.2 shows fairly good agreement between the DVHs from a clinical treatment plan (solid lines) and the corresponding achievable plan resulting from solving (3.2) (dashed lines) for patient #1.

It is not critical that these two dose distributions be identical – in fact, it is probably impossible to make them so. The goal of this initialization process is simply to generate a dose distribution that is similar to the clinical one and that can be achieved using $\mathbf{D}^{CERR}$. The achievable dose distributions satisfy all clinical criteria that the clinical dose distributions are measured against. Going forward, we treat the achievable dose distribution and $\mathbf{D}^{CERR}$ as being the clinical dose distribution and dose influence matrix, respectively.

In this chapter, we described the clinical motivation for our work and formulated the forward treatment plan optimization problem. In the computational results of the subsequent chapters, we address the problem of finding objective function weights for the forward problem (3.1) using historical treatment plans as input.

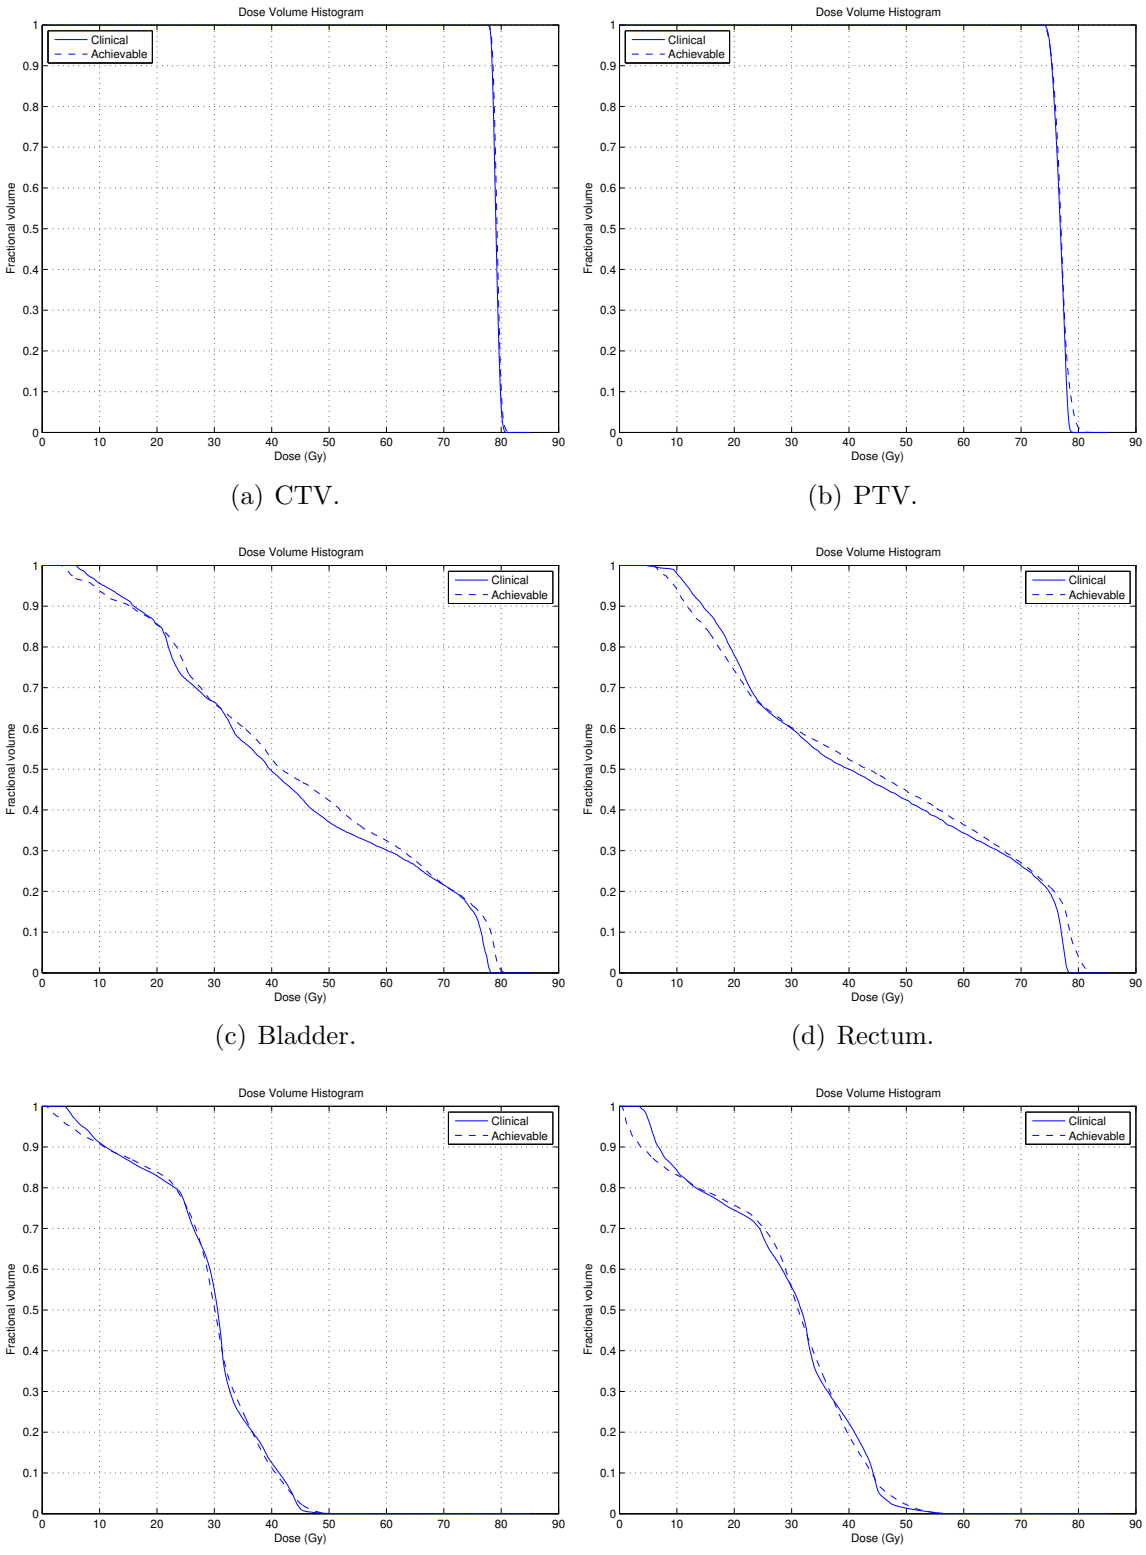

(a) CTV.

(b) PTV.

(c) Bladder.

(d) Rectum.

(e) Left femoral head.

(f) Right femoral head.

Figure 3.2: Comparison of DVHs of the clinical treatment plan and achievable treatment plan.

Chapter 4

Generalized inverse multiobjective linear optimization

4.1 Introduction

In this chapter, we develop a generalized inverse linear optimization model that returns a meaningful, nonzero solution to the inverse problem for any initial $\hat{\mathbf{x}}$, regardless of whether $\hat{\mathbf{x}}$ is on the boundary of the forward feasible region, is an interior point, or is even an infeasible solution. Instead of formulating the inverse optimization problem using strong duality as a constraint as was done in Ahuja and Orlin (2001), we introduce a duality gap that is to be minimized. By relaxing the strong duality constraint, we simultaneously relax the assumption that $\hat{\mathbf{x}}$ must be on the boundary of the forward feasible region.

This framework generalizes many existing inverse optimization models and, we believe, opens the door for much broader and more effective application of inverse optimization in practice. There is no guarantee that a given input $\hat{\mathbf{x}}$ for an inverse problem will be on the boundary of the corresponding forward feasible region. Previously, one might focus effort on modifying the constraints of the forward problem so that the boundary of

its feasible region contains $\hat{\mathbf{x}}$. Instead, our general framework does not require any modification of the forward problem to accommodate the input vector $\hat{\mathbf{x}}$. Another attractive feature of our generalized inverse optimization model is that it retains the underlying complexity of the forward problem, like the standard inverse linear (Ahuja and Orlin, 2001) and conic (Iyengar and Kang, 2005) optimization models. Furthermore, it specializes to the standard inverse optimization model when $\hat{\mathbf{x}}$ is on the boundary of the forward feasible region. Thus, there is little downside to adopting this more generalized framework when solving inverse optimization problems.

We develop this inverse methodology in a multiobjective linear optimization context, which has not been considered previously, and develop connections with established multiobjective optimization literature, especially Pareto surface approximation techniques. Both the multiobjective and generalization aspects of this research are motivated by a practical problem in radiation therapy treatment planning.

As an application of our generalized inverse optimization methodology, we demonstrate how to determine appropriate objective function weights in prostate cancer IMRT. Given a historical treatment plan, our method determines objective function weights that make the given plan nearly optimal with respect to a given set of objective functions, which are not necessarily identical to the ones used to create the plan. Clinical treatment planning formulations have a large number of objective functions and vary from institution to institution. As a result, a treatment plan from one institution is not guaranteed to be optimal (or even feasible) for another, as measured by the latter's formulation. We demonstrate that with intelligently chosen weights, treatments of clinical quality can be created using many fewer objective functions than currently used in practice. We also show how our methodology can determine which objective functions most heavily influence a final treatment. While our approach is motivated by radiation therapy applications, the methodology we develop is intended to solve general inverse (multiobjective) optimization problems.

Our specific contributions in this chapter are as follows:

1. We generalize the standard inverse optimization framework by allowing an arbitrary solution to the forward problem, even an infeasible one, to be used as input. We characterize the theoretical properties and geometry of the resulting inverse solutions.

2. We develop the first inverse optimization approach for multiobjective linear optimization. We elucidate connections between inverse optimization and established multiobjective optimization theory, especially with Pareto surface approximation techniques.

3. We develop the first inverse optimization methodology applied to IMRT treatment planning. Determining the most critical objective functions and optimized starting values for their weights may provide new insight into the design of future treatments.

4.2 Inverse optimization methodology

The derivation of our inverse optimization methodology is based on the canonical multiobjective linear programming formulation (Zeleny, 1974). Let $\mathbf{x} \in \mathbb{R}^n, \mathbf{A} \in \mathbb{R}^{m \times n}, \mathbf{b} \in \mathbb{R}^m$, and $\mathcal{K}$ be the set of objective functions. Then the multiobjective optimization problem can be written as

$$\begin{aligned} \text{FOP}(\boldsymbol{\alpha}): \quad \underset{\mathbf{x}}{\text{minimize}} \quad & \boldsymbol{\alpha}'\mathbf{C}\mathbf{x} \\ \text{subject to} \quad & \mathbf{A}\mathbf{x} = \mathbf{b}, \\ & \mathbf{x} \geq \mathbf{0}, \end{aligned} \tag{4.1}$$

where each row of $\mathbf{C} \in \mathbb{R}^{|\mathcal{K}| \times n}$ represents a different linear objective function associated with $\mathbf{x}$ and $\boldsymbol{\alpha} \in \mathbb{R}^{|\mathcal{K}|}$, assumed nonnegative, denotes the weight vector. Without loss of generality, we assume $\mathbf{C}\mathbf{x} \in \mathbb{R}_+^{|\mathcal{K}|}$ for every feasible $\mathbf{x}$ (Ehrgott, 2005).

4.2.1 Preliminaries

We define notation and illustrate some geometrical properties of the forward optimization problem (4.1). Let $\mathbf{X} = \{\mathbf{x} \in \mathbb{R}^n \,|\, \mathbf{A}\mathbf{x} = \mathbf{b}, \mathbf{x} \geq \mathbf{0}\}$ be the set of feasible solutions to formulation (4.1), assumed non-empty. The set of attainable objective function values will be denoted $\mathbf{Z} = \{\mathbf{C}\mathbf{x} \in \mathbb{R}_+^{|\mathcal{K}|} \,|\, \mathbf{x} \in \mathbf{X}\}$. Following the convention of Ehrgott (2005), $\mathbf{X}$ is a subset of the *decision space* and $\mathbf{Z}$ is a subset of the *criterion space*. Let $\mathbf{Z}^+ = \{\mathbf{z} \in \mathbb{R}_+^{|\mathcal{K}|} \,|\, \mathbf{z} \geq \hat{\mathbf{z}}, \hat{\mathbf{z}} \in \mathbf{Z}\}$. A point $\mathbf{C}\hat{\mathbf{x}} \in \mathbf{Z}$ is a *non-dominated point* of $\mathbf{Z}$ if there is no $\mathbf{C}\mathbf{x} \in \mathbf{Z}$ satisfying $(\mathbf{c}^k)'\mathbf{x} \leq (\mathbf{c}^k)'\hat{\mathbf{x}}$ for all $k \in \mathcal{K}$ with at least one k such that $(\mathbf{c}^k)'\mathbf{x} < (\mathbf{c}^k)'\hat{\mathbf{x}}$. A point $\mathbf{C}\hat{\mathbf{x}} \in \mathbf{Z}$ is a *weakly non-dominated point* of $\mathbf{Z}$ if there is no $\mathbf{C}\mathbf{x} \in \mathbf{Z}$ such that $(\mathbf{c}^k)'\mathbf{x} < (\mathbf{c}^k)'\hat{\mathbf{x}}$ for all $k \in \mathcal{K}$. A (weakly) non-dominated point of $\mathbf{Z}^+$ is defined similarly. We denote the set of non-dominated points and the set of weakly non-dominated points of $\mathbf{Z}$ ($\mathbf{Z}^+$) by $\mathbf{Z}_N$ and $\mathbf{Z}_{WN}$ ($\mathbf{Z}_N^+$ and $\mathbf{Z}_{WN}^+$), respectively. Note that $\mathbf{Z}_{WN}^+ \supset \mathbf{Z}_{WN} \supset \mathbf{Z}_N^+ = \mathbf{Z}_N$ (see Figure 4.1). A solution $\hat{\mathbf{x}} \in \mathbf{X}$ corresponding to a (weakly) non-dominated point of $\mathbf{Z}$ is an *(weakly) efficient solution* to problem (4.1). The following is a fundamental result of multiobjective linear programming (see Ehrgott (2005)).

Proposition 4.2.1 *In a multiobjective linear program, a vector $\mathbf{x} \in \mathbf{X}$ is a weakly efficient solution if and only if there exists a nonzero $\boldsymbol{\alpha} \geq \mathbf{0}$ such that $\mathbf{x}$ is an optimal solution to* FOP($\boldsymbol{\alpha}$).

Proposition 4.2.1 suggests that if $\mathbf{Z}$ is full-dimensional and $\hat{\mathbf{z}} \in \text{int}(\mathbf{Z})$, there is no nonzero $\boldsymbol{\alpha} \geq \mathbf{0}$ such that $\hat{\mathbf{x}}$ is optimal for FOP($\boldsymbol{\alpha}$) and $\hat{\mathbf{z}} = \mathbf{C}\hat{\mathbf{x}}$. If $\hat{\mathbf{z}} \notin \mathbf{Z}$, there will also be no $\boldsymbol{\alpha}$ that makes the corresponding $\hat{\mathbf{x}}$ optimal, because $\hat{\mathbf{x}}$ is infeasible.

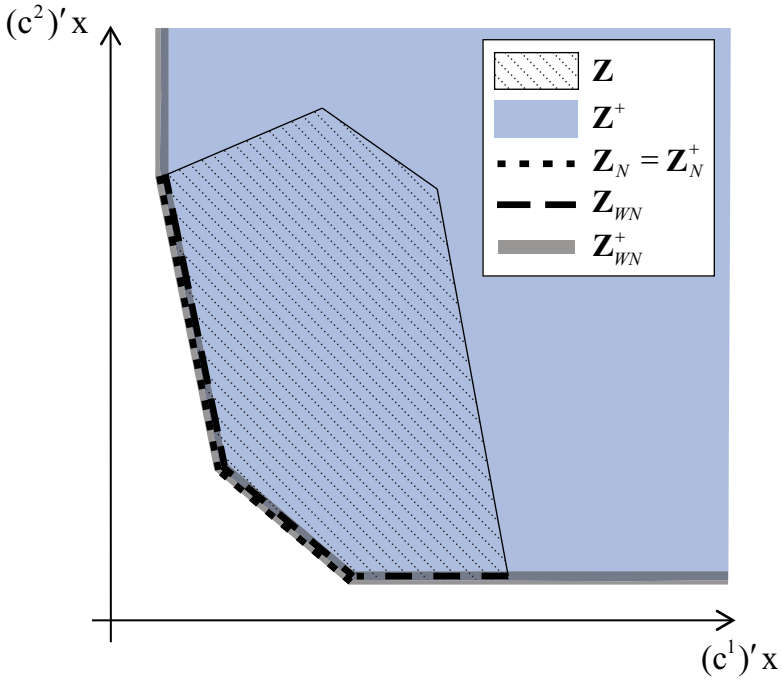

Figure 4.1: Geometry of criterion space.

4.2.2 An inverse linear optimization problem

Given $\hat{\mathbf{x}} \in \mathbf{X}$, we formulate the inverse problem associated with the forward problem (4.1) as

$$\begin{aligned} \text{IOP}(\mathbf{C}, \hat{\mathbf{x}}): \quad \underset{\boldsymbol{\alpha},\mathbf{p}}{\text{minimize}} \quad & 0 \\ \text{subject to} \quad & \mathbf{A}'\mathbf{p} \leq \mathbf{C}'\boldsymbol{\alpha}, \\ & \boldsymbol{\alpha}'\mathbf{C}\hat{\mathbf{x}} = \mathbf{b}'\mathbf{p}, \\ & \boldsymbol{\alpha}'\mathbf{e} = 1, \\ & \boldsymbol{\alpha} \geq \mathbf{0}. \end{aligned} \tag{4.2}$$

The first and second constraints correspond to dual feasibility and strong duality, respectively. We add the third constraint to prevent the trivial solution $(\boldsymbol{\alpha}, \mathbf{p}) = (\mathbf{0}, \mathbf{0})$ from being a feasible solution. We note that other constraints, such as $\mathbf{b}'\mathbf{p} = 1$ or $\boldsymbol{\alpha}'\mathbf{C}\hat{\mathbf{x}} = 1$ would also serve this purpose. However, for the generalized inverse optimization formulations we present later, using $\boldsymbol{\alpha}'\mathbf{e} = 1$ in formulation (4.2) facilitates the analysis. Omitting an objective function simplifies the exposition without fundamentally affecting

the structure of the inverse problem. As we will see in Corollary 4.2.2, it is unlikely that a specific objective function is needed.

4.2.3 A generalized inverse linear optimization problem

By Proposition 4.2.1, IOP($\mathbf{C}, \hat{\mathbf{x}}$) is feasible if and only if $\hat{\mathbf{x}} \in \mathbf{X}$ is a weakly efficient solution. We now modify IOP($\mathbf{C}, \hat{\mathbf{x}}$) to ensure that it remains feasible, given any $\hat{\mathbf{x}}$, be it a weakly efficient solution, a solution corresponding to an interior point of $\mathbf{Z}$, or even an infeasible solution. We will assume that $\hat{\mathbf{x}}$ is such that $\mathbf{C}\hat{\mathbf{x}} \in \mathbb{R}^{|\mathcal{K}|}_{+}$.

Since $\mathbf{Cx} \in \mathbb{R}^{|\mathcal{K}|}_{+}$ for all $\mathbf{x} \in \mathbf{X}$, the forward problem (4.1) has a finite optimal solution for any nonzero $\boldsymbol{\alpha} \geq \mathbf{0}$. Therefore, the set of feasible solutions to the dual formulation is non-empty for any nonzero $\boldsymbol{\alpha} \geq \mathbf{0}$. However, the strong duality constraint $\boldsymbol{\alpha}'\mathbf{C}\hat{\mathbf{x}} = \mathbf{b}'\mathbf{p}$ in formulation (4.2) holds if and only if a given solution $\hat{\mathbf{x}} \in \mathbf{X}$ is weakly efficient. So if we have a feasible $\hat{\mathbf{x}}$ that is not weakly efficient, instead of enforcing strong duality, we can allow a duality gap and aim to find a feasible solution to the corresponding inverse problem with as small a duality gap as possible. This idea is also appropriate when $\hat{\mathbf{x}}$ is an infeasible solution, as will be shown later. We consider two different types of duality gaps: (i) a relative duality gap $\epsilon_r > 0$, leading to the constraint $\boldsymbol{\alpha}'\mathbf{C}\hat{\mathbf{x}} = \epsilon_r\, \mathbf{b}'\mathbf{p}$, and (ii) an absolute duality gap ϵ_a, leading to the constraint $\boldsymbol{\alpha}'\mathbf{C}\hat{\mathbf{x}} = \mathbf{b}'\mathbf{p} + \epsilon_a$. We denote the two corresponding generalized inverse optimization problems (GIOP) by GIOP$_r$ and GIOP$_a$, respectively.

Given any $\hat{\mathbf{x}} \in \mathbb{R}^n$ satisfying $\mathbf{C}\hat{\mathbf{x}} \in \mathbb{R}^{|\mathcal{K}|}_+$ and $\epsilon_r > 0$, the *relative* GIOP is

$$\begin{aligned} \text{GIOP}_r(\mathbf{C}, \hat{\mathbf{x}}, \epsilon_r): \quad \underset{\boldsymbol{\alpha},\mathbf{p}}{\text{minimize}} \quad & 0 \\ \text{subject to} \quad & \mathbf{A}'\mathbf{p} \leq \mathbf{C}'\boldsymbol{\alpha}, \\ & \boldsymbol{\alpha}'\mathbf{C}\hat{\mathbf{x}} = \epsilon_r\, \mathbf{b}'\mathbf{p}, \\ & \boldsymbol{\alpha}'\mathbf{e} = 1, \\ & \boldsymbol{\alpha} \geq \mathbf{0}. \end{aligned} \tag{4.3}$$

In formulation (4.3), $\epsilon_r > 0$ is a parameter. Formulation (4.3) can be solved repeatedly with varying ϵ_r until the smallest ϵ_r, denoted ϵ_r^*, is found such that (4.3) is feasible. A univariate search technique such as the bisection or golden section algorithm may be used. The following result provides an upper bound on ϵ_r^*, which can be used as an initial value of ϵ_r in the search.

Proposition 4.2.2 *Let* $\mathbf{z}^I = (z_1^I, ..., z_{|\mathcal{K}|}^I)$, *where* $z_k^I = \min_{\mathbf{z}\in\mathbf{Z}}\{z_k\}$ *for all* $k \in \mathcal{K}$. *Assume* $z_k^I > 0$ *for at least one* $k \in \mathcal{K}$. *If* $\mathbf{C}\hat{\mathbf{x}} \in \mathbf{Z}^+$, $\epsilon_r^* \leq \min_{\{k\in\mathcal{K}|z_k^I>0\}} \left\{ \frac{(\mathbf{c}^k)'\hat{\mathbf{x}}}{z_k^I} \right\}$. *If* $\mathbf{C}\hat{\mathbf{x}} \notin \mathbf{Z}^+, \epsilon_r^* \leq 1$.

Proof: Let $\mathbf{C}\hat{\mathbf{x}} \in \mathbf{Z}^+$ and $\mathbf{C}\mathbf{x}^* = \mathbf{C}\hat{\mathbf{x}}/\epsilon_r^*$. By Theorem 4.2.1, $\mathbf{C}\mathbf{x}^* \in \mathbf{Z}^+_{WN}$. By the definition of $\mathbf{z}^I$, $\mathbf{z}^I \leq \mathbf{C}\mathbf{x}$ for any $\mathbf{C}\mathbf{x} \in \mathbf{Z}$. Therefore, $\mathbf{z}^I \leq \mathbf{C}\mathbf{x}^* \leq \mathbf{C}\hat{\mathbf{x}}$. For any $k \in \mathcal{K}$ such that $z_k^I > 0$, $\epsilon_r^* = (\mathbf{c}^k)'\hat{\mathbf{x}}/(\mathbf{c}^k)'\mathbf{x}^* \leq (\mathbf{c}^k)'\hat{\mathbf{x}}/z_k^I$. If $\mathbf{C}\hat{\mathbf{x}} \notin \mathbf{Z}^+$ then by Theorem 4.2.1 $\mathbf{C}\mathbf{x}^* \leq \mathbf{C}\hat{\mathbf{x}}$, and therefore $\epsilon_r^* \leq 1$. □

Note that instead of solving formulation (4.3) multiple times with varying ϵ_r, by writing $\epsilon_r = \boldsymbol{\alpha}'\mathbf{C}\hat{\mathbf{x}}/\mathbf{b}'\mathbf{p}$, a single linear program can be solved to find ϵ_r^*:

$$\begin{aligned} \underset{\boldsymbol{\alpha},\mathbf{p}}{\text{minimize}} \quad & \boldsymbol{\alpha}'\mathbf{C}\hat{\mathbf{x}} \\ \text{subject to} \quad & \mathbf{A}'\mathbf{p} \leq \mathbf{C}'\boldsymbol{\alpha}, \\ & \mathbf{b}'\mathbf{p} = 1, \\ & \boldsymbol{\alpha} \geq \mathbf{0}. \end{aligned} \tag{4.4}$$

In formulation (4.4), we have to omit the constraint $\boldsymbol{\alpha}'\mathbf{e} = 1$ when introducing the constraint $\mathbf{b}'\mathbf{p} = 1$. Thus, an optimal solution to formulation (4.4) must be normalized post-hoc. If we let $\bar{\boldsymbol{\alpha}}$ be an optimal solution to formulation (4.4), then $\epsilon_r^* = \bar{\boldsymbol{\alpha}}'\mathbf{C}\hat{\mathbf{x}}$. Thus, this formulation simultaneously generates the optimal weight vector and the minimal duality gap. Going forward, we will use relative GIOP formulations (4.3) and (4.4) interchangeably, depending on which is more convenient for the context. In particular, the theoretical analysis will focus on formulation (4.3), while the computational results will be based on formulation (4.4).

Given any $\hat{\mathbf{x}} \in \mathbb{R}^n$ satisfying $\mathbf{C}\hat{\mathbf{x}} \in \mathbb{R}_+^{|\mathcal{K}|}$, the *absolute* GIOP is

$$\begin{aligned} \text{GIOP}_a(\mathbf{C}, \hat{\mathbf{x}}): \quad & \underset{\boldsymbol{\alpha}, \mathbf{p}, \epsilon_a}{\text{minimize}} && \epsilon_a \\ & \text{subject to} && \mathbf{A}'\mathbf{p} \leq \mathbf{C}'\boldsymbol{\alpha}, \\ & && \boldsymbol{\alpha}'\mathbf{C}\hat{\mathbf{x}} = \mathbf{b}'\mathbf{p} + \epsilon_a, \\ & && \boldsymbol{\alpha}'\mathbf{e} = 1, \\ & && \boldsymbol{\alpha} \geq \mathbf{0}. \end{aligned} \tag{4.5}$$

Let ϵ_a^* be the optimal value of problem (4.5). If $\epsilon_r^* = 1$ or $\epsilon_a^* = 0$, then the two GIOP formulations become the standard inverse formulation (4.2). In other words, the GIOP formulations can identify a weakly efficient solution $\hat{\mathbf{x}}$ to the forward problem. Otherwise, as we show next, ϵ_r^* and ϵ_a^* measure how far a given solution $\mathbf{C}\hat{\mathbf{x}}$ is from the set of weakly non-dominated points, with respect to two particular distance metrics induced by the two duality gaps considered.

4.2.4 Structure of optimal inverse solutions

If $\hat{\mathbf{x}}$ is not weakly efficient, then solving any GIOP formulation will return an $\boldsymbol{\alpha}^*$ such that FOP($\boldsymbol{\alpha}^*$) generates a weakly efficient solution $\mathbf{x}^* \neq \hat{\mathbf{x}}$. Recall that if $\mathbf{x}^*$ is optimal for FOP($\boldsymbol{\alpha}^*$), then IOP($\mathbf{C}, \mathbf{x}^*$) is feasible. Therefore, solving a GIOP formulation with $\hat{\mathbf{x}}$

as input can be viewed as solving IOP$(\mathbf{C}, \mathbf{x}^*)$ for an $\mathbf{x}^*$ that has been suitably perturbed from $\hat{\mathbf{x}}$. In this section, we study the relationship between the $\mathbf{x}^*$ that generates $\boldsymbol{\alpha}^*$ from IOP$(\mathbf{C}, \mathbf{x}^*)$ and an $\hat{\mathbf{x}}$ that generates the same $\boldsymbol{\alpha}^*$ from a GIOP formulation. Standard inverse optimization models typically assume that the given $\hat{\mathbf{x}}$ is a feasible solution to the forward problem. First, Proposition 4.2.3 shows that even if $\hat{\mathbf{x}}$ is infeasible, the inverse problem can be feasible.

Proposition 4.2.3 *Let* $\tilde{\mathbf{x}} \notin \mathbf{X}$ *and* $\mathbf{C}\tilde{\mathbf{x}} \in \mathbb{R}_+^{|\mathcal{K}|}$. IOP$(\mathbf{C}, \tilde{\mathbf{x}})$ *is feasible if and only if* $\mathbf{C}\tilde{\mathbf{x}} \in (\mathbb{R}_+^{|\mathcal{K}|} \setminus \mathbf{Z}^+) \cup (\mathbf{Z}_{WN}^+ \setminus \mathbf{Z}_{WN})$.

Proof: ($\Leftarrow$) Let $\mathbf{y}, \pi$, and σ be the dual variables corresponding to the first three sets of constraints, respectively, of IOP$(\mathbf{C}, \tilde{\mathbf{x}})$. With suitable sign changes, the following formulation is equivalent to the dual of IOP$(\mathbf{C}, \tilde{\mathbf{x}})$:

$$\begin{aligned} -\underset{\mathbf{y},\pi,\sigma}{\text{minimize}} \quad & \sigma \\ \text{subject to} \quad & \mathbf{A}\mathbf{y} = \pi\mathbf{b}, \\ & \mathbf{C}\mathbf{y} \leq \pi\mathbf{C}\tilde{\mathbf{x}} + \sigma\mathbf{e}, \\ & \mathbf{y} \geq \mathbf{0}. \end{aligned} \tag{4.6}$$

To show that IOP$(\mathbf{C}, \tilde{\mathbf{x}})$ is feasible, it suffices to show that formulation (4.6) has a finite optimal solution. Moreover, if formulation (4.6) has a finite optimal solution, its optimal cost is zero by strong duality. Note that $(\mathbf{y}, \pi, \sigma) = (\mathbf{0}, 0, 0)$ is feasible for formulation (4.6). We can assume $\pi \geq 0$ because the strong duality constraint in IOP$(\mathbf{C}, \tilde{\mathbf{x}})$ can be written as a less than or equal to inequality (weak duality is already guaranteed). Consider the change of variables $\mathbf{y} = \pi\mathbf{x}$. The second set of constraints in formulation (4.6) becomes

$$\sigma \geq \pi \max_{k \in \mathcal{K}} ((\mathbf{c}^k)'\mathbf{x} - (\mathbf{c}^k)'\tilde{\mathbf{x}}). \tag{4.7}$$

If $\mathbf{C}\tilde{\mathbf{x}} \in \mathbb{R}_+^{|\mathcal{K}|} \setminus \mathbf{Z}^+$, there exists $\mathbf{Cx} \in \mathbf{Z}_{WN}$ such that $(\mathbf{c}^k)'\tilde{\mathbf{x}} < (\mathbf{c}^k)'\mathbf{x}$ for at least one $k \in \mathcal{K}$. Therefore, $\max_{k\in\mathcal{K}}((\mathbf{c}^k)'\mathbf{x}-(\mathbf{c}^k)'\tilde{\mathbf{x}}) > 0$. If $\mathbf{C}\tilde{\mathbf{x}} \in \mathbf{Z}_{WN}^+ \setminus \mathbf{Z}_{WN}$, there exists $\mathbf{Cx} \in \mathbf{Z}_{WN}$ such that $(\mathbf{c}^k)'\tilde{\mathbf{x}} = (\mathbf{c}^k)'\mathbf{x}$ for at least one $k \in \mathcal{K}$, so $\max_{k\in\mathcal{K}}((\mathbf{c}^k)'\mathbf{x} - (\mathbf{c}^k)'\tilde{\mathbf{x}}) \geq 0$. In both cases, σ is bounded below by 0, as required.

($\Rightarrow$) If IOP$(\mathbf{C}, \tilde{\mathbf{x}})$ is feasible, then formulation (4.6) is also feasible with optimal cost $\sigma^* = 0$. Suppose there exists $\mathbf{Cx} \in \mathbf{Z}$ such that $(\mathbf{c}^k)'\mathbf{x} - (\mathbf{c}^k)'\tilde{\mathbf{x}} < 0$, $\forall k \in \mathcal{K}$. It follows from equation (4.7) that σ can become arbitrarily negative for sufficiently large π, which is a contradiction. Hence, there cannot be $\mathbf{Cx} \in \mathbf{Z}$ satisfying $(\mathbf{c}^k)'\mathbf{x} < (\mathbf{c}^k)'\tilde{\mathbf{x}}$, $\forall k \in \mathcal{K}$. As $\tilde{\mathbf{x}} \notin \mathbf{X}$, it must be that $\mathbf{C}\tilde{\mathbf{x}} \in (\mathbb{R}_+^{|\mathcal{K}|} \setminus \mathbf{Z}^+) \cup (\mathbf{Z}_{WN}^+ \setminus \mathbf{Z}_{WN})$, as required. □

The next result draws an equivalence between the IOP model and the GIOP models for a related pair of input vectors.

Proposition 4.2.4 *(a) A solution* $(\boldsymbol{\alpha}, \mathbf{p})$ *is a feasible solution to* GIOP$_r(\mathbf{C}, \hat{\mathbf{x}}, \epsilon_r)$ *if and only if* $(\boldsymbol{\alpha}, \mathbf{p})$ *is a feasible solution to* IOP$(\mathbf{C}, \mathbf{x}^*)$ *where* $\mathbf{x}^*$ *is such that* $\mathbf{Cx}^* = \mathbf{C}\hat{\mathbf{x}}/\epsilon_r$.

(b) Suppose $\epsilon_a \leq \min_{k\in\mathcal{K}}(\mathbf{c}^k)'\hat{\mathbf{x}}$. *A solution* $(\boldsymbol{\alpha}, \mathbf{p}, \epsilon_a)$ *is a feasible solution to* GIOP$_a(\mathbf{C}, \hat{\mathbf{x}})$ *if and only if* $(\boldsymbol{\alpha}, \mathbf{p})$ *is a feasible solution to* IOP$(\mathbf{C}, \mathbf{x}^*)$ *where* $\mathbf{x}^*$ *is such that* $\mathbf{Cx}^* = \mathbf{C}\hat{\mathbf{x}} - \epsilon_a\mathbf{e}$.

Proof: For part (a), $(\boldsymbol{\alpha}, \mathbf{p})$ satisfies $\boldsymbol{\alpha}'\mathbf{C}\hat{\mathbf{x}} = \mathbf{b}'\mathbf{p}\epsilon_r$ if and only if it satisfies $\boldsymbol{\alpha}'\mathbf{Cx} = \mathbf{b}'\mathbf{p}$ for $\mathbf{Cx} = \mathbf{C}\hat{\mathbf{x}}/\epsilon_r$. Since all other constraints are the same between the formulations GIOP$_r(\mathbf{C}, \hat{\mathbf{x}}, \epsilon_r)$ and IOP$(\mathbf{C}, \mathbf{x})$, the statement holds. The proof for part (b) is similar and omitted. □

Corollary 4.2.1 *(a) For* $\epsilon_r > 0$ *sufficiently large,* GIOP$_r(\mathbf{C}, \hat{\mathbf{x}}, \epsilon_r)$ *is feasible.*

(b) GIOP$_a(\mathbf{C}, \hat{\mathbf{x}})$ *is always feasible.*

Proof: Consider GIOP$_r(\mathbf{C}, \hat{\mathbf{x}}, \epsilon_r)$. For ϵ_r sufficiently large, $\mathbf{0} \leq \mathbf{C}\hat{\mathbf{x}}/\epsilon_r \leq \mathbf{Cx}$ for some $\mathbf{Cx} \in \mathbf{Z}_{WN}^+$. Thus, $\mathbf{C}\hat{\mathbf{x}}/\epsilon_r \in \mathbb{R}_+^{|\mathcal{K}|} \setminus \mathbf{Z}^+$. Let $\tilde{\mathbf{x}}$ satisfy $\mathbf{C}\tilde{\mathbf{x}} = \mathbf{C}\hat{\mathbf{x}}/\epsilon_r$. By Proposition 4.2.3,

IOP($\mathbf{C}, \tilde{\mathbf{x}}$) is feasible and then by Proposition 4.2.4, GIOP$_r(\mathbf{C}, \hat{\mathbf{x}}, \epsilon_r)$ is feasible, as required. The proof for part (b) is similar and omitted. □

Corollary 4.2.1 suggests that if $\mathbf{C}\hat{\mathbf{x}} \in \mathbf{Z}$ and ϵ_r is sufficiently large, then $\mathbf{C}\hat{\mathbf{x}}/\epsilon_r \in \mathbb{R}_+^{|\mathcal{K}|} \setminus \mathbf{Z}^+$ and solving GIOP$_r(\mathbf{C}, \hat{\mathbf{x}}, \epsilon_r)$ will return a vector $\boldsymbol{\alpha}$ such that $\boldsymbol{\alpha}'\mathbf{C}\hat{\mathbf{x}}/\epsilon_r \leq \boldsymbol{\alpha}'\mathbf{C}\mathbf{x}$ for all $\mathbf{C}\mathbf{x} \in \mathbf{Z}$. In other words, $\boldsymbol{\alpha}$ defines a separating hyperplane between the point $\mathbf{C}\hat{\mathbf{x}}/\epsilon_r$ and the nonempty, closed, convex set $\mathbf{Z}$. However, since $\boldsymbol{\alpha}$ is generated from a point that is an infeasible solution to the forward problem (that is, the pre-image of $\mathbf{C}\hat{\mathbf{x}}/\epsilon_r$ is infeasible in decision space), there is no way to recover the pre-image of the point $\mathbf{C}\hat{\mathbf{x}}/\epsilon_r$ by solving FOP($\boldsymbol{\alpha}$).

The main result of this section states that when the duality gap is minimized, an optimal solution $\boldsymbol{\alpha}^*$ to either the relative or absolute inverse problem defines a supporting hyperplane of $\mathbf{Z}^+$, and is also an optimal solution to the standard original inverse optimization problem for a suitably perturbed input vector.

Theorem 4.2.1 *(a) If* $(\boldsymbol{\alpha}^*, \mathbf{p}^*)$ *is an optimal solution to* GIOP$_r(\mathbf{C}, \hat{\mathbf{x}}, \epsilon_r^*)$*, then it is an optimal solution to* IOP$(\mathbf{C}, \mathbf{x}^*)$ *where* $\mathbf{C}\mathbf{x}^* = \mathbf{C}\hat{\mathbf{x}}/\epsilon_r^*$*, and* $\mathbf{C}\mathbf{x}^* \in \mathbf{Z}_{WN}^+$.

(b) If $(\boldsymbol{\alpha}^*, \mathbf{p}^*, \epsilon_a^*)$ *is an optimal solution to* GIOP$_a(\mathbf{C}, \hat{\mathbf{x}})$*, then it is an optimal solution to* IOP$(\mathbf{C}, \mathbf{x}^*)$ *where* $\mathbf{C}\mathbf{x}^* = \mathbf{C}\hat{\mathbf{x}} - \epsilon_a^*\mathbf{e}$*, and* $\mathbf{C}\mathbf{x}^* \in \mathbf{Z}_{WN}^+$.

Proof: For part (a), given any ϵ_r, there is a one-to-one correspondence between the feasible solutions of GIOP$_r(\mathbf{C}, \hat{\mathbf{x}}, \epsilon_r)$ and IOP$(\mathbf{C}, \mathbf{x}^*)$, by Proposition 4.2.4. Thus, if $(\boldsymbol{\alpha}^*, \mathbf{p}^*)$ is optimal for GIOP$_r(\mathbf{C}, \hat{\mathbf{x}}, \epsilon_r^*)$, it is optimal for IOP$(\mathbf{C}, \mathbf{x}^*)$. Lastly, recall that ϵ_r^* is the smallest value of ϵ_r such that GIOP$_r(\mathbf{C}, \hat{\mathbf{x}}, \epsilon_r)$ remains feasible. The case of $\epsilon_r^* = 1$ is trivial, so assume that $\epsilon_r^* \neq 1$. Suppose to the contrary that $\mathbf{C}\mathbf{x}^* = \mathbf{C}\hat{\mathbf{x}}/\epsilon_r^* \notin \mathbf{Z}_{WN}^+$. There are two possibilities: either $\mathbf{C}\mathbf{x}^* \in \text{int}(\mathbf{Z}^+)$ or $\mathbf{C}\mathbf{x}^* \in \mathbb{R}_+^{|\mathcal{K}|} \setminus \mathbf{Z}^+$. In the first case, IOP$(\mathbf{C}, \mathbf{x}^*)$ cannot be feasible by Proposition 4.2.3. In the second case, there must exist $\bar{\epsilon}_r$ such that $\mathbf{C}\mathbf{x}^* = \mathbf{C}\hat{\mathbf{x}}/\epsilon_r^* < \mathbf{C}\hat{\mathbf{x}}/\bar{\epsilon}_r \leq \mathbf{C}\mathbf{x}$ for some $\mathbf{C}\mathbf{x} \in \mathbf{Z}_{WN}^+$. Thus

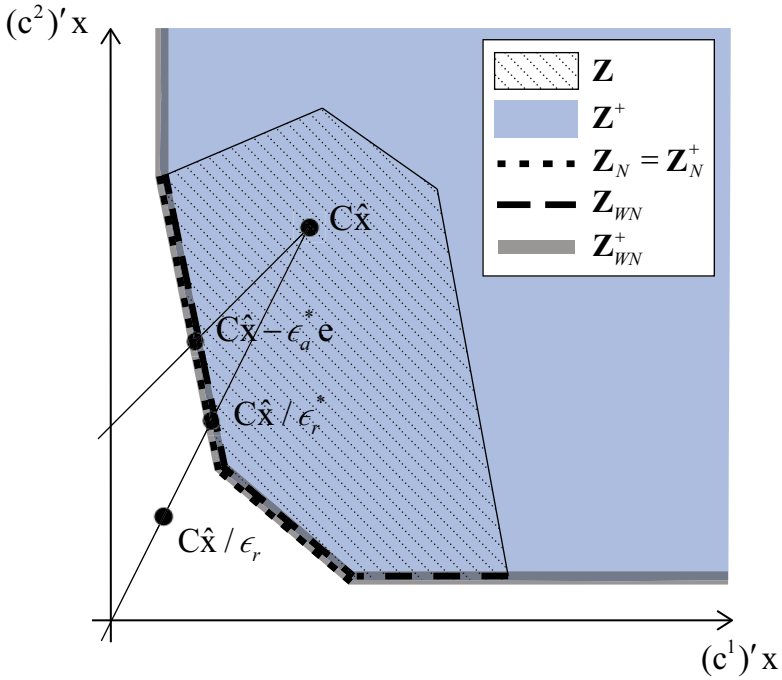

Figure 4.2: Adjustments of an input point for GIOP to a non-dominated point.

$\mathbf{C}\bar{\mathbf{x}} = \mathbf{C}\hat{\mathbf{x}}/\bar{\epsilon}_r \in \mathbb{R}_+^{|\mathcal{K}|} \setminus \mathbf{Z}^+$ and IOP$(\mathbf{C}, \bar{\mathbf{x}})$ is feasible by Proposition 4.2.3, which contradicts the minimality of ϵ_r^*. The proof for part (b) is similar and is omitted. □

The significance of this result is that there is an exact relationship between the two generalized inverse optimization problems with $\mathbf{C}\hat{\mathbf{x}}$ as an input, and the original inverse optimization problem with a modified input $\mathbf{C}\mathbf{x}^* \in \mathbf{Z}_{WN}^+$ (see Figure 4.2). For GIOP$_r$, if we connect the point $\mathbf{C}\hat{\mathbf{x}} \in \mathbf{Z}^+$ with a line segment to the origin (i.e., along the direction vector $-\mathbf{C}\hat{\mathbf{x}}$), then $\mathbf{C}\hat{\mathbf{x}}/\epsilon_r^*$ is where the line segment intersects $\mathbf{Z}_{WN}^+$. Moreover an optimal solution to GIOP$_r(\mathbf{C}, \hat{\mathbf{x}}, \epsilon_r^*)$ is also an optimal solution to IOP$(\mathbf{C}, \mathbf{x}^*)$ where $\mathbf{C}\mathbf{x}^* = \mathbf{C}\hat{\mathbf{x}}/\epsilon_r^*$. Similarly, for GIOP$_a$, if we draw a line moving away from a point $\mathbf{C}\hat{\mathbf{x}} \in \mathbf{Z}^+$ in the direction of $-\mathbf{e}$, then where this line intersects $\mathbf{Z}_{WN}^+$ is the point $\mathbf{C}\hat{\mathbf{x}} - \epsilon_a^*\mathbf{e}$ and an optimal solution to GIOP$_a(\mathbf{C}, \hat{\mathbf{x}})$ is also optimal for IOP$(\mathbf{C}, \mathbf{x}^*)$ where $\mathbf{C}\mathbf{x}^* = \mathbf{C}\hat{\mathbf{x}} - \epsilon_a^*\mathbf{e}$. If $\mathbf{C}\hat{\mathbf{x}} \notin \mathbf{Z}^+$, then the direction of the adjustment is $+\mathbf{C}\hat{\mathbf{x}}$ in the relative case and $+\mathbf{e}$ in the absolute case. Moving away from $\mathbf{C}\hat{\mathbf{x}}$ in a direction other than $\pm\mathbf{C}\hat{\mathbf{x}}$ ($\pm\mathbf{e}$) may result in a closer non-dominated point (e.g., in ℓ_2 distance), but the relative (absolute) duality gaps

have the nice interpretation in the multiobjective context of adjusting all the objective function values by the same relative (absolute) amount. The optimal relative and absolute gaps are thus a measure of the distance between $\mathbf{C}\hat{\mathbf{x}}$ and $\mathbf{Z}^+_{WN}$ along these two direction vectors. Note that if $\mathbf{C}\hat{\mathbf{x}} \in \text{int}(\mathbf{Z}^+)$, then $\epsilon^*_r > 1$ and $\epsilon^*_a > 0$. If $\mathbf{C}\hat{\mathbf{x}} \in \mathbb{R}^{|\mathcal{K}|}_+ \setminus \mathbf{Z}^+$, then $\epsilon^*_r < 1$ and $\epsilon^*_a < 0$.

In multiobjective linear optimization, both $\mathbf{Z}$ and $\mathbf{Z}^+$ are polyhedra (Zeleny, 1974). If a point $\mathbf{Cx} \in \mathbf{Z}^+_{WN}$ is a vertex of $\mathbf{Z}^+$, then many weight vectors $\boldsymbol{\alpha}$ can generate $\mathbf{Cx}$. The same is true if two or more constraints defining $\mathbf{Z}^+$ are active at $\mathbf{Cx}$. This observation leads to the following result, whose proof is omitted.

Corollary 4.2.2 *An optimal solution $\boldsymbol{\alpha}^*$ to* IOP$(\mathbf{C}, \mathbf{x}^*)$ *is unique if and only if $\mathbf{Cx}^* \in \mathbf{Z}^+_{WN}$ lies on the relative interior of a facet of $\mathbf{Z}^+$.*

Starting from an arbitrary point $\mathbf{C}\hat{\mathbf{x}}$ and traveling in the direction $\pm\mathbf{C}\hat{\mathbf{x}}$ or $\pm\mathbf{e}$, the likelihood of intersecting $\mathbf{Z}^+_{WN}$ at anything other than the relative interior of a facet of $\mathbf{Z}^+$ is low, which is verified in our computational results. A practical implication is that an objective function to minimize $||\boldsymbol{\alpha} - \hat{\boldsymbol{\alpha}}||_1$, for example, is generally not needed in either GIOP formulation.

4.2.5 Validating optimal inverse solutions

Let $\boldsymbol{\alpha}^*$ be an optimal solution to GIOP$_r(\mathbf{C}, \hat{\mathbf{x}}, \epsilon^*_r)$. Because $\mathbf{Z}^+$ is a polyhedron and $\mathbf{C}\hat{\mathbf{x}}/\epsilon^*_r$ may lie on the relative interior of a facet of $\mathbf{Z}^+$, solving FOP$(\boldsymbol{\alpha}^*)$ is not guaranteed to recover $\mathbf{C}\hat{\mathbf{x}}/\epsilon^*_r$ – it is only guaranteed to return an $\mathbf{x}^* \in \mathbf{X}$ such that $(\boldsymbol{\alpha}^*)'\mathbf{Cx}^* = (\boldsymbol{\alpha}^*)'\mathbf{C}\hat{\mathbf{x}}/\epsilon^*_r$. The same is true for GIOP$_a(\mathbf{C}, \hat{\mathbf{x}})$. Thus, solving FOP$(\boldsymbol{\alpha}^*)$ to validate an optimal solution $\boldsymbol{\alpha}^*$ from GIOP$_r$ or GIOP$_a$ may not be sufficient. Instead, dual information can be used to validate $\boldsymbol{\alpha}^*$.

Proposition 4.2.5 *(a) Let $\boldsymbol{\alpha}^*$ be an optimal solution to GIOP$_r$ formulation (4.4), $\epsilon^*_r = \boldsymbol{\alpha}^{*\prime}\mathbf{C}\hat{\mathbf{x}}$, and $\mathbf{y}^{dual}_r$ be the optimal dual variables associated with the constraints*

$\mathbf{A}'\mathbf{p} \leq \mathbf{C}'\boldsymbol{\alpha}$. *Then the vector* $\mathbf{x}_r^{dual} = \mathbf{y}_r^{dual}/\epsilon_r^*$ *satisfies* $\mathbf{C}\mathbf{x}_r^{dual} \in \mathbf{Z}_{WN}$, $(\mathbf{c}^k)'\mathbf{x}_r^{dual} = (\mathbf{c}^k)'\hat{\mathbf{x}}/\epsilon_r^*$ *for* k *such that* $\alpha_k^* > 0$ *and* $(\mathbf{c}^k)'\mathbf{x}_r^{dual} \leq (\mathbf{c}^k)'\hat{\mathbf{x}}/\epsilon_r^*$ *for* k *such that* $\alpha_k^* = 0$.

(*b*) *Let* $\boldsymbol{\alpha}^*$ *be an optimal solution to* GIOP$_a$ *formulation* (4.5) *and* $\mathbf{x}_a^{dual}$ *be the optimal dual variables associated with the constraints* $\mathbf{A}'\mathbf{p} \leq \mathbf{C}'\boldsymbol{\alpha}$. *Then the vector* $\mathbf{x}_a^{dual}$ *satisfies* $\mathbf{C}\mathbf{x}_a^{dual} \in \mathbf{Z}_{WN}$, $(\mathbf{c}^k)'\mathbf{x}_a^{dual} = (\mathbf{c}^k)'\hat{\mathbf{x}} - \epsilon_a^*$ *for* k *such that* $\alpha_k^* > 0$ *and* $(\mathbf{c}^k)'\mathbf{x}_a^{dual} \leq (\mathbf{c}^k)'\hat{\mathbf{x}} - \epsilon_a^*$ *for* k *such that* $\alpha_k^* = 0$.

Proof: For part (a), let γ be the dual variable associated with the first set of constraints in formulation (4.4). Then the dual of formulation (4.4) can be written as:

$$\begin{aligned} \underset{\mathbf{y},\gamma}{\text{maximize}} \quad & \gamma \\ \text{subject to} \quad & \mathbf{A}\mathbf{y} = \gamma\mathbf{b}, \\ & \mathbf{C}\mathbf{y} \leq \mathbf{C}\hat{\mathbf{x}}, \\ & \mathbf{y} \geq \mathbf{0}. \end{aligned} \tag{4.8}$$

Let $(\mathbf{y}_r^{dual}, \gamma^*)$ be an optimal solution to formulation (4.8). By strong duality, $\gamma^* = \epsilon_r^*$. Then from formulation (4.8), the vector $\mathbf{x}_r^{dual} = \mathbf{y}_r^{dual}/\epsilon_r^*$ must satisfy

$$\epsilon_r^* \leq \min_{k \in \mathcal{K}} \frac{(\mathbf{c}^k)'\hat{\mathbf{x}}}{(\mathbf{c}^k)'\mathbf{x}_r^{dual}}. \tag{4.9}$$

By complementary slackness, $\epsilon_r^* = (\mathbf{c}^k)'\hat{\mathbf{x}}/(\mathbf{c}^k)'\mathbf{x}_r^{dual}$ for $k \in \mathcal{K}$ such that $\alpha_k > 0$ and $\epsilon_r^* \leq (\mathbf{c}^k)'\hat{\mathbf{x}}/(\mathbf{c}^k)'\mathbf{x}_r^{dual}$ for $k \in \mathcal{K}$ such that $\alpha_k = 0$. Lastly, since $\mathbf{C}\mathbf{x}_r^{dual} \leq \mathbf{C}\hat{\mathbf{x}}/\epsilon_r^* \in \mathbf{Z}_{WN}^+$ (Theorem 4.2.1) and $\mathbf{x}_r^{dual} \in \mathbf{X}$ (from formulation (4.8)), $\mathbf{C}\mathbf{x}_r^{dual} \in \mathbf{Z}_{WN}$.

The proof for part (b) is similar and omitted. □

Proposition 4.2.5 implies that once a GIOP formulation is solved with a given $\mathbf{C}\hat{\mathbf{x}}$, the optimal solution $\boldsymbol{\alpha}^*$ can be validated by looking up the appropriate optimal dual variables. Note that for both relative and absolute cases, if $\boldsymbol{\alpha}^* > \mathbf{0}$, then $\mathbf{C}\mathbf{x}^{dual} \in \mathbf{Z}_N$.

4.3 Benson's method and Pareto surface approximation techniques

By minimizing the duality gaps in the relative and absolute GIOP formulations, we are taking a point $\mathbf{C}\hat{\mathbf{x}}$ in criterion space and projecting it on to $\mathbf{Z}^{+}_{WN}$ along the two associated direction vectors. Note that by projecting $\mathbf{C}\hat{\mathbf{x}}$ on to $\mathbf{Z}^{+}_{WN}$, the GIOP formulations can also be used to find the normals of the hyperplanes that define $\mathbf{Z}^{+}_{WN}$, which is the goal of many Pareto surface approximation techniques. In this section, we briefly describe some of these techniques and highlight a connection with our GIOP formulations.

Benson's method (Benson, 1978) is a well-known method to identify efficient solutions in a multiobjective optimization problem. Given $\mathbf{C}\hat{\mathbf{x}} \in \mathbf{Z}$, Benson's formulation is

$$\begin{aligned} \underset{\mathbf{x},\boldsymbol{\epsilon}}{\text{maximize}} \quad & \boldsymbol{\epsilon}'\mathbf{e} \\ \text{subject to} \quad & \mathbf{C}\mathbf{x} \leq \mathbf{C}\hat{\mathbf{x}} - \boldsymbol{\epsilon}, \\ & \mathbf{A}\mathbf{x} = \mathbf{b}, \\ & \mathbf{x} \geq \mathbf{0}, \quad \boldsymbol{\epsilon} \geq \mathbf{0}. \end{aligned} \tag{4.10}$$

Formulation (4.10) attempts to find a point on the efficient frontier that (weakly) dominates $\hat{\mathbf{x}}$. Let $(\mathbf{x}^*, \boldsymbol{\epsilon}^*)$ be an optimal solution to formulation (4.10). The point $\hat{\mathbf{x}}$ is efficient if and only if the optimal cost $\boldsymbol{\epsilon}^{*\prime}\mathbf{e}$ equals 0. If $0 < \boldsymbol{\epsilon}^{*\prime}\mathbf{e} < \infty$, then $\mathbf{x}^*$ is an efficient solution and (weakly) dominates $\hat{\mathbf{x}}$. If problem (4.10) is unbounded, then no efficient solutions exist.

Variants of formulation (4.10) have been widely used in Pareto surface approximation techniques (Das and Dennis, 1998; Benson, 1998; Shao and Ehrgott, 2008; Ehrgott et al., 2011). For example, given an interior point $\mathbf{C}\hat{\mathbf{x}}$ of $\mathbf{Z}^{+}$ and a vertex $\hat{\mathbf{z}}$ of $\mathbf{S} \supset \mathbf{Z}^{+}$, Benson (1998) first used a univariate search technique to find the point $\mathbf{C}\mathbf{x}^*$ where the line segment connecting $\mathbf{C}\hat{\mathbf{x}}$ and $\hat{\mathbf{z}}$ intersects $\mathbf{Z}^{+}_{WN}$. Then, the dual of formulation (4.10)

was solved to find a supporting hyperplane of $\mathbf{Z}^+$ at $\mathbf{Cx}^*$. Shao and Ehrgott (2008) improved the method in Benson (1998) by using a variant of formulation (4.10) to find $\mathbf{Cx}^* \in \mathbf{Z}^+_{WN}$, instead of using univariate search. A second LP was still used to find the supporting hyperplane of $\mathbf{Z}^+$ at $\mathbf{Cx}^*$. Next, we draw a connection between our GIOP formulations and formulation (4.10), which allows us to find both $\mathbf{Cx}^*$ and the associated supporting hyperplane by solving a single LP.

The dual of GIOP$_a(\mathbf{C}, \hat{\mathbf{x}})$ formulation (4.5) can be written as

$$\begin{aligned} \underset{\mathbf{x},\sigma}{\text{maximize}} \quad & \sigma \\ \text{subject to} \quad & \mathbf{Cx} \leq \mathbf{C}\hat{\mathbf{x}} - \sigma\mathbf{e}, \\ & \mathbf{Ax} = \mathbf{b}, \\ & \mathbf{x} \geq \mathbf{0}. \end{aligned} \tag{4.11}$$

Notice that formulations (4.10) and (4.11) produce equivalent optimal solutions if $\boldsymbol{\epsilon} = \sigma\mathbf{e}$. In other words, the GIOP$_a$ method is the dual of Benson's method, in the case where the components of $\boldsymbol{\epsilon}$ are constrained to be equal. Similarly, the GIOP$_r$ method is the dual of Benson's method, in the case where the components of $\boldsymbol{\epsilon}$ are constrained to be proportional to the components of $\mathbf{C}\hat{\mathbf{x}}$. While the GIOP methods attempt to find a weakly non-dominated point $\mathbf{Cx}^*$ in a particular direction from $\mathbf{C}\hat{\mathbf{x}}$, Benson's method tries to find a non-dominated point that is furthest from $\mathbf{C}\hat{\mathbf{x}}$ as measured by the 1-norm, and that (weakly) dominates $\mathbf{C}\hat{\mathbf{x}}$. Theorem 4.3.1, whose proof is straightforward and omitted, formalizes these ideas.

Theorem 4.3.1 *Let* $\hat{\mathbf{x}} \in \mathbf{X}$.

(*a*) *Let* $\boldsymbol{\epsilon} = \sigma\mathbf{C}\hat{\mathbf{x}}$ *in formulation* (4.10) *and let* $\boldsymbol{\lambda}^*_r$ *be an optimal dual vector associated with the first constraint in formulation* (4.10). *Let* $\boldsymbol{\alpha}^*_r$ *be an optimal solution to GIOP*$_r$ *formulation* (4.4). *Then a solution is optimal for FOP*$(\boldsymbol{\lambda}^*_r)$ *if and only if it is optimal for FOP*$(\boldsymbol{\alpha}^*_r)$.

(*b*) *Let* $\boldsymbol{\epsilon} = \sigma\mathbf{e}$ *in formulation* (4.10) *and let* $\boldsymbol{\lambda}_a^*$ *be an optimal dual vector associated with the first constraint in in formulation* (4.10)*. Let* $\boldsymbol{\alpha}_a^*$ *be an optimal solution to GIOP*$_a$ *formulation* (4.5)*. Then a solution is optimal for FOP*$(\boldsymbol{\lambda}_a^*)$ *if and only if it is optimal for FOP*$(\boldsymbol{\alpha}_a^*)$*.*

Theorem 4.3.1 suggests that the two steps – the first step to find a point on the non-dominated frontier and the second step to find a supporting hyperplane at that point – used in some of the previous Pareto surface approximation techniques can be combined into a single step, by solving formulation (4.10) with a given direction vector and simply evaluating its dual variables. While ideas from inverse optimization may lead to new results in Pareto surface approximation, the main purpose of the preceding discussion is to elucidate a connection between these two areas.

4.4 Computational results

In this section, we demonstrate the use of our GIOP formulations in the context of radiation therapy treatment planning. We consider 12 treatment plans for prostate cancer patients who had previously received radiation therapy at Princess Margaret Cancer Centre.

4.4.1 A multiobjective forward formulation

We consider the forward multiobjective problem (3.1) with two types of objective functions for the OARs: (a) a linear penalty objective function that penalizes delivering dose to any part of the OAR above a certain dose threshold level, and (b) an objective that minimizes the maximum dose delivered to the OAR. These two types of OAR objective functions are motivated by the classification of organs into *parallel* and *serial* organs (Dale and Olsen, 1997; Thieke et al., 2002). For a parallel organ, a high dose to a small volume can typically be tolerated if the rest of the organ is protected, and

therefore minimizing the total (or average) dose delivered to the overall organ is an appropriate clinical objective. For serial organs, a high dose to a small portion of the organ can result in dysfunction, and therefore it is important to minimize the maximum dose delivered (objective function type (b)). As organs are typically somewhere in between these extremes, an objective function that penalizes dose above some threshold dose is often used (type (a)).

Let $\mathcal{I}$ be the set of OAR objectives of type (a), $\mathcal{J}$ be the set of OAR objectives of type (b), and $\mathcal{K} := \mathcal{I} \cup \mathcal{J}$ be the set of all objectives. Then the objective function of (3.1) can be written as:

$$\sum_{k\in\mathcal{K}} \alpha_k f_k(\mathbf{w}) = \sum_{i\in\mathcal{I}} \frac{\alpha_i}{|\mathcal{O}_i|} \sum_{v\in\mathcal{O}_i} \max\left\{0, \sum_{b\in\mathcal{B}} D_{v,b} w_b - \theta_v^i\right\} + \sum_{j\in\mathcal{J}} \alpha_j \max_{v\in\mathcal{O}_j} \sum_{b\in\mathcal{B}} D_{v,b} w_b, \quad (4.12)$$

where θ_v^i denotes a dose threshold on voxel $v \in \mathcal{O}_i, i \in \mathcal{I}$ above which overdosing is linearly penalized.

The forward optimization formulation (3.1) with the objective function (4.12) can be linearized in a standard fashion by introducing auxiliary variables and constraints

(Bertsimas and Tsitsiklis, 1997):

$$
\begin{aligned}
\underset{\mathbf{w},\boldsymbol{\xi},\mathbf{z}}{\text{minimize}} \quad & \sum_{i\in\mathcal{I}} \frac{\alpha_i}{|\mathcal{O}_i|} \sum_{v\in\mathcal{O}_i} \xi_{i,v} + \sum_{j\in\mathcal{J}} \alpha_j z_j \\
\text{subject to} \quad & \xi_{i,v} \geq \sum_{b\in\mathcal{B}} D_{v,b} w_b - \theta^i_v, \quad \forall v \in \mathcal{O}_i, i \in \mathcal{I}, \\
& z_j \geq \sum_{b\in\mathcal{B}} D_{v,b} w_b, \quad \forall v \in \mathcal{O}_j, j \in \mathcal{J}, \\
& \sum_{b\in\mathcal{B}} D_{v,b} w_b \geq \ell_v, \quad \forall v \in \mathcal{T}, \\
& \sum_{b\in\mathcal{B}} D_{v,b} w_b \leq u_v, \quad \forall v \in \mathcal{V}, \\
& \frac{\beta_1}{|\mathcal{B}|} \sum_{b'\in\mathcal{B}} w_{b'} \leq \ w_b \leq \frac{\beta_2}{|\mathcal{B}|} \sum_{b'\in\mathcal{B}} w_{b'}, \quad \forall b \in \mathcal{B}, \\
& \xi_{i,v} \geq 0, \quad \forall v \in \mathcal{O}_i, i \in \mathcal{I}, \\
& z_j \geq 0, \quad j \in \mathcal{J}, \\
& w_b \geq 0, \quad \forall b \in \mathcal{B}.
\end{aligned}
\tag{4.13}
$$

Note that, if we let $\mathbf{x} = (\mathbf{w}, \boldsymbol{\xi}, \mathbf{z})$ and

$$
\mathbf{C} = \begin{bmatrix} \mathbf{c}^{1\prime} \\ \vdots \\ \mathbf{c}^{|\mathcal{K}|\prime} \end{bmatrix} = \left[\begin{array}{c|c|c} \mathbf{0}_{|\mathcal{I}|\times|\mathcal{B}|} & \begin{matrix} \frac{\mathbf{e}'_{|\mathcal{O}_1|}}{|\mathcal{O}_1|} & & \\ & \ddots & \\ & & \frac{\mathbf{e}'_{|\mathcal{O}_{|\mathcal{I}|}|}}{|\mathcal{O}_{|\mathcal{I}|}|} \end{matrix} & \mathbf{0}_{|\mathcal{I}|\times|\mathcal{J}|} \\ \hline \mathbf{0}_{|\mathcal{J}|\times|\mathcal{B}|} & \mathbf{0}_{|\mathcal{J}|\times\sum_{i\in\mathcal{I}}|\mathcal{O}_i|} & \mathbf{I}_{|\mathcal{J}|} \end{array}\right],
$$

then with appropriately defined $\mathbf{A}$ and $\mathbf{b}$, formulation (4.13) can be written as FOP($\boldsymbol{\alpha}$) formulation (4.1).

Let $\mathbf{p}, \mathbf{q}, \mathbf{r}, \mathbf{s}, \mathbf{t}^1$ and $\mathbf{t}^2$ be the dual variables associated with the six primary constraints in formulation (4.13) ($\mathbf{t}^1$ and $\mathbf{t}^2$ correspond to the left and right parts of the fifth

constraint, respectively). Given $\hat{\mathbf{w}}$,

$$\hat{\xi}_{v,i} = \max\left\{0, \sum_{b\in\mathcal{B}} D_{v,b}\hat{w}_b - \theta_v^i\right\}, \quad \forall v \in \mathcal{O}_i, i \in \mathcal{I} \tag{4.14}$$

and

$$\hat{z}_j = \max_{v\in\mathcal{O}_j} \sum_{b\in\mathcal{B}} D_{v,b}\hat{w}_b, \quad \forall j \in \mathcal{J} \tag{4.15}$$

are uniquely determined by $\hat{\mathbf{w}}$ and used as inputs in the inverse formulation.

Then the corresponding relative inverse formulation represented as $\mathrm{GIOP}_r(\mathbf{C}, \hat{\mathbf{x}})$ is written as follows:

$$\begin{aligned}
\underset{\mathbf{p},\mathbf{q},\mathbf{r},\mathbf{s},\mathbf{t}^1,\mathbf{t}^2,\boldsymbol{\alpha}}{\text{minimize}} \quad & \sum_{i\in\mathcal{I}} \frac{\alpha_i}{|\mathcal{O}_i|} \sum_{v\in\mathcal{O}_i} \hat{\xi}_{v,i} + \sum_{j\in\mathcal{J}} \alpha_j \hat{z}_j \\
\text{subject to} \quad & -\sum_{i\in\mathcal{I}}\sum_{v\in\mathcal{O}_i} D_{v,b}p_{v,i} - \sum_{j\in\mathcal{J}}\sum_{v\in\mathcal{O}_j} D_{v,b}q_{v,j} + \sum_{v\in\mathcal{T}} D_{v,b}r_v - \sum_{v\in\mathcal{V}} D_{v,b}s_v \\
& \qquad + t_b^1 - \frac{\beta_1}{|\mathcal{B}|}\sum_{b'\in\mathcal{B}} t_{b'}^1 - t_b^2 + \frac{\beta_2}{|\mathcal{B}|}\sum_{b'\in\mathcal{B}} t_{b'}^2 \le 0, \quad \forall b \in B, \\
& p_{v,i} \le \frac{\alpha_i}{|\mathcal{O}_i|}, \quad \forall v \in \mathcal{O}_i, \forall i \in \mathcal{I}, \\
& \sum_{v\in\mathcal{O}_j} q_{v,j} = \alpha_j, \quad \forall j \in \mathcal{J}, \\
& -\sum_{i\in\mathcal{I}}\sum_{v\in\mathcal{O}_i} \theta_v^i p_{v,i} + \sum_{v\in\mathcal{T}} \ell_v r_v - \sum_{v\in\mathcal{V}} u_v s_v = 1, \\
& \alpha_k \ge 0, \quad \forall k \in \mathcal{K}, \\
& p_{v,i} \ge 0, \quad \forall v \in \mathcal{O}_i, \forall i \in \mathcal{I}, \\
& q_{v,j} \ge 0, \quad \forall v \in \mathcal{O}_j, \forall j \in \mathcal{J}, \\
& r_v \ge 0, \quad \forall v \in \mathcal{T}, \quad s_v \ge 0, \quad \forall v \in \mathcal{V}, \\
& t_b^1 \ge 0, \quad \forall b \in \mathcal{B}, \quad t_b^2 \ge 0, \quad \forall b \in \mathcal{B}.
\end{aligned} \tag{4.16}$$

We consider three different instances of forward formulation (4.12), each with a different number of objectives: (A) four objectives, (B) six objectives, and (C) 18 objectives.

In model (A), only one objective per OAR is included. It has been shown that one objective per OAR is sufficient in generating clinically acceptable treatments at low computational overhead (Craft, 2011). Because of the low volume tolerance for the femoral head criteria, we modeled those objectives using a maximum dose objective (type (b)). The bladder and rectum objectives were modeled using the linear penalty objective function (type (a)), with thresholds $\theta^i = 50$ for both. In all models, we assume all voxels in OAR objective i (of type (a)) have a common dose threshold $\theta^i = \theta^i_v$ for all $v \in \mathcal{O}_i$. In model (B), we add a mean dose objective, which can be modeled with a linear penalty function and $\theta^i = 0$, for both the bladder and rectum to the four objectives from model (A). In model (C), we introduce a linear penalty objective for the bladder and rectum for each of the values of $\theta^i \in \{0, 10, \ldots, 70\}$. The objective with $\theta^i = 70$ can be viewed as an approximate maximum dose objective. The femoral head objectives remain the same as in models (A) and (B).

4.4.2 Reproducing a clinical plan with fewer objectives

We consider model (A) and demonstrate the use of the GIOP models in recreating a plan of clinical quality using inversely optimized weights for patient #1. All optimization problems were solved using CPLEX 12.0 on a computer with a quad-core 2.66 GHz Intel Xeon W3520 processor and 6 GB of memory. Note that all the clinical specific formulations discussed above can be written as compact formulations presented in Section 4.2. To facilitate the interpretation of the clinical results and their connection to our theoretical results, our discussion in this section will be based on the compact inverse formulations in Section 4.2 with $\hat{\mathbf{x}}$ representing the input vector. We solved formulation (4.4) for GIOP$_r$, and (4.5) for GIOP$_a$. The sizes of the FOP, GIOP$_r$, and GIOP$_a$ formulations for patient #1 are shown in Table 4.1 along with solution (CPU) times.

Table 4.2 shows the objective function weights, $\boldsymbol{\alpha}^*_r$, that result from solving GIOP$_r$ for patient #1. As discussed in Section 4.2.5, using dual information to validate the

Table 4.1: Problem sizes and solution times for the FOP and GIOP formulations (model (A)) for patient #1.

	FOP	GIOP$_r$	GIOP$_a$
Number of variables	4,326	98,219	98,220
Number of constraints	98,215	4,327	4,328
Solution time (s)	211	318	303

Table 4.2: Results from GIOP$_r$ for model (A) for patient #1.

OAR	$\boldsymbol{\alpha}_r^*$	ϵ_r^*	$\mathbf{C}\hat{\mathbf{x}}$	$\mathbf{C}\mathbf{x}_r^*$	$\mathbf{C}\mathbf{x}_r^{dual}$	$\mathbf{C}\hat{\mathbf{x}}/\mathbf{C}\mathbf{x}_r^*$	$\mathbf{C}\hat{\mathbf{x}}/\mathbf{C}\mathbf{x}_r^{dual}$
Bladder	0.248	1.098	8.280	7.539	7.540	1.098	1.098
Rectum	0.746		9.625	8.765	8.765	1.098	1.098
L.Fem	0.002		51.611	47.127	47.001	1.095	1.098
R.Fem	0.004		57.398	52.320	52.271	1.097	1.098

weights is the most accurate method. In Table 4.2, using the definition of $\mathbf{x}_r^{dual}$ from Proposition 4.2.5, the component-by-component ratio $\mathbf{C}\hat{\mathbf{x}}/\mathbf{C}\mathbf{x}_r^{dual}$ is precisely ϵ_r^*, as expected (Proposition 4.2.4 and Theorem 4.2.1). Similar results are seen in Table 4.3 for GIOP$_a$. The weights are similar but not identical between the two GIOP formulations, which suggests that $\mathbf{C}\mathbf{x}_r^{dual}$ and $\mathbf{C}\mathbf{x}_a^{dual}$ may be sitting on nearby facets of $\mathbf{Z}^+$. We also solved the GIOP formulations with objective functions of the form $||\boldsymbol{\alpha} - \hat{\boldsymbol{\alpha}}||_1$, for many different $\hat{\boldsymbol{\alpha}}$. Each time, the same $\boldsymbol{\alpha}^*$ was generated, which suggests that the solutions $\mathbf{C}\mathbf{x}_r^{dual}$ and $\mathbf{C}\mathbf{x}_a^{dual}$ lie in the relative interiors of facets of $\mathbf{Z}^+$. If we validate $\boldsymbol{\alpha}_r^*$ ($\boldsymbol{\alpha}_a^*$) by solving FOP($\boldsymbol{\alpha}_r^*$) (FOP($\boldsymbol{\alpha}_a^*$)), then we obtain the results in the columns associated with $\mathbf{C}\mathbf{x}_r^*$ ($\mathbf{C}\mathbf{x}_a^*$). In this case, the component-wise ratios $\mathbf{C}\hat{\mathbf{x}}/\mathbf{C}\mathbf{x}_r^*$ (and differences $\mathbf{C}\hat{\mathbf{x}} - \mathbf{C}\mathbf{x}_a^*$) are very similar but not identical. This observation again reinforces the idea that $\mathbf{C}\mathbf{x}_r^{dual}$ and $\mathbf{C}\mathbf{x}_a^{dual}$ lie in the relative interiors of facets of $\mathbf{Z}^+$. Moreover, the similarity in the component-wise ratio/difference is indicative of the facets being very small. The duality

Table 4.3: Results from GIOP$_a$ for model (A) for patient #1.

OAR	$\boldsymbol{\alpha}_a^*$	ϵ_a^*	$\mathbf{C}\hat{\mathbf{x}}$	$\mathbf{C}\mathbf{x}_a^*$	$\mathbf{C}\mathbf{x}_a^{dual}$	$\mathbf{C}\hat{\mathbf{x}} - \mathbf{C}\mathbf{x}_a^*$	$\mathbf{C}\hat{\mathbf{x}} - \mathbf{C}\mathbf{x}_a^{dual}$
Bladder	0.326	0.846	8.280	7.434	7.434	0.846	0.846
Rectum	0.670		9.625	8.779	8.779	0.846	0.846
L.Fem	0.002		51.611	50.787	50.765	0.824	0.846
R.Fem	0.002		57.398	56.545	56.552	0.853	0.846

gap values $\epsilon_r^* > 1$ and $\epsilon_a^* > 0$ indicate that the given solution was an interior point of $\mathbf{Z}^+$.

While dual information is more accurate in validating the inversely optimized weights, solving the forward problem with the optimal weights does provide additional information that is clinically relevant and useful in validating the GIOP results. For example, by solving FOP($\boldsymbol{\alpha}_r^*$) we can generate an entire dose distribution (as opposed to just the objective function values) and compare it to the clinically achievable one used as input to GIOP$_r$. The treatment plan and dose distribution that result from solving FOP($\boldsymbol{\alpha}_r^*$) will be referred to as the inversely optimized plan and inversely optimized dose distribution, respectively. Figure 4.3 plots the DVHs of the CTV, PTV, and four OARs for both the clinically achievable and inversely optimized plans using weights from GIOP$_r$ for patient #1.

The DVHs match fairly well for the bladder and rectum, indicating that the dose distributions are similar in the clinical and inversely optimized plans for these organs. For the femoral heads, the DVHs of the inversely optimized plan mostly dominate the clinical plan. Overall, the inversely optimized treatment plan seems to be at least as good as the clinical one, which suggests that the clinical plan is feasible but not on the efficient frontier of the forward problem. The DVH comparison between the clinically achievable and inversely optimized plans for patients #2 and #3 is shown in Figures 4.4 and 4.5, respectively, which show similar results. Figure 4.6 shows the similar-looking clinical and inversely optimized dose distributions on a CT image for patient #1.

Lastly, we present a quantitative comparison in Table 4.4. The criterion $Vx \leq y\%$ ($Vx \geq y\%$) is a partial-volume metric, evaluating whether the volume of the structure that receives x Gy or more is at most (at least) $y\%$. It can be seen that, for patient #1, the clinically achievable and inversely optimized plans satisfy all clinical requirements. In fact, the inversely optimized plan dominates the clinically achievable one in the bladder and rectum metrics. The inversely optimized plan is equivalent to the clinical one in the femoral head criteria, though a reduction in the maximum dose delivered is evident in

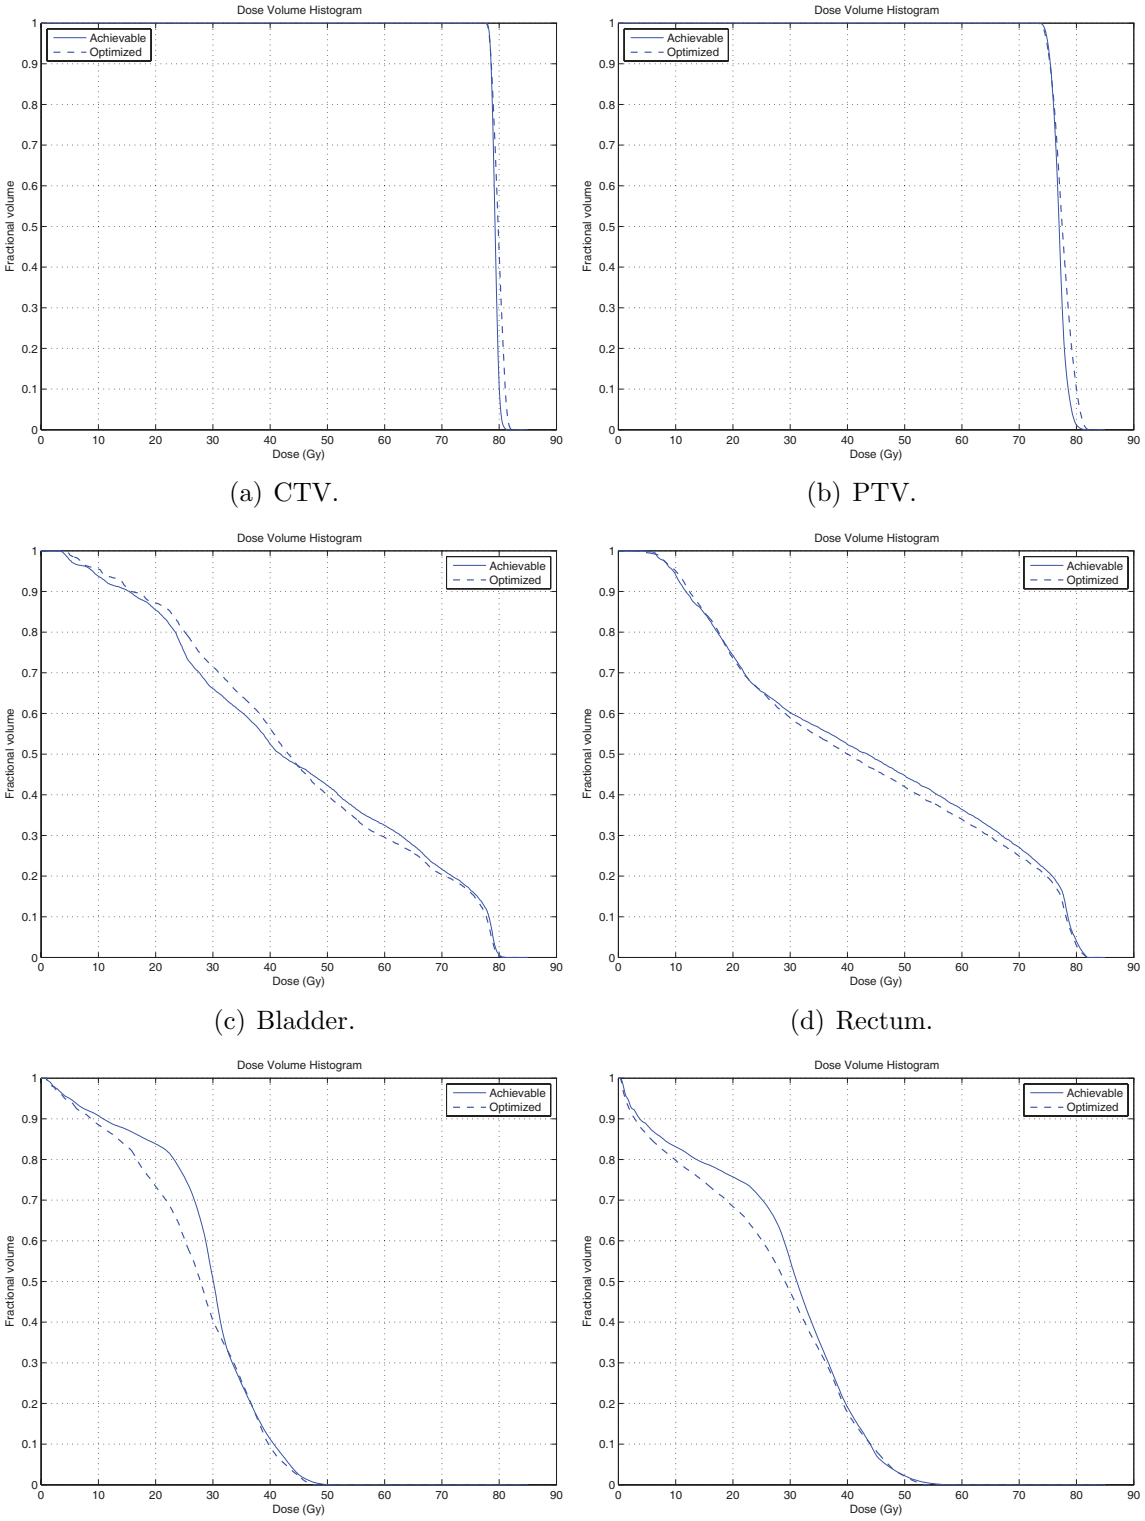

(a) CTV.

(b) PTV.

(c) Bladder.

(d) Rectum.

(e) Left femoral head.

(f) Right femoral head.

Figure 4.3: DVHs of the clinically achievable and inversely optimized plans for patient #1 using GIOP$_r$ Model (A).

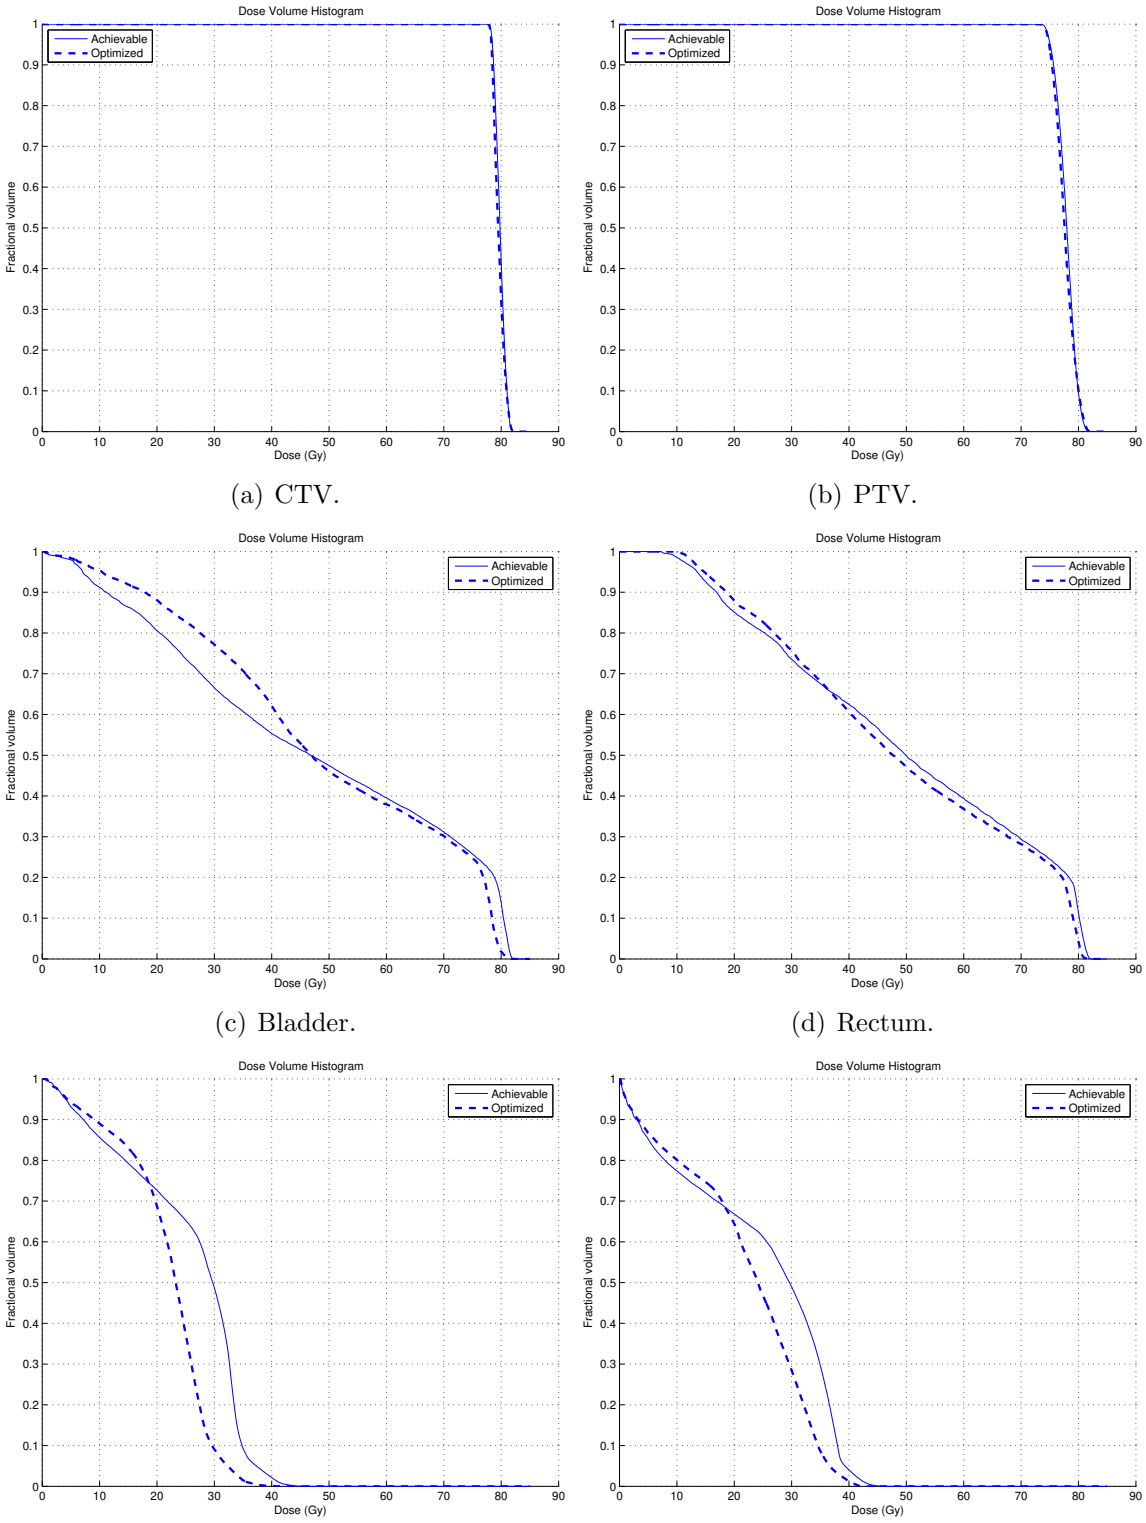

(a) CTV.

(b) PTV.

(c) Bladder.

(d) Rectum.

(e) Left femoral head.

(f) Right femoral head.

Figure 4.4: DVHs of the clinically achievable and inversely optimized plans for patient #2 using GIOP$_r$ Model (A).

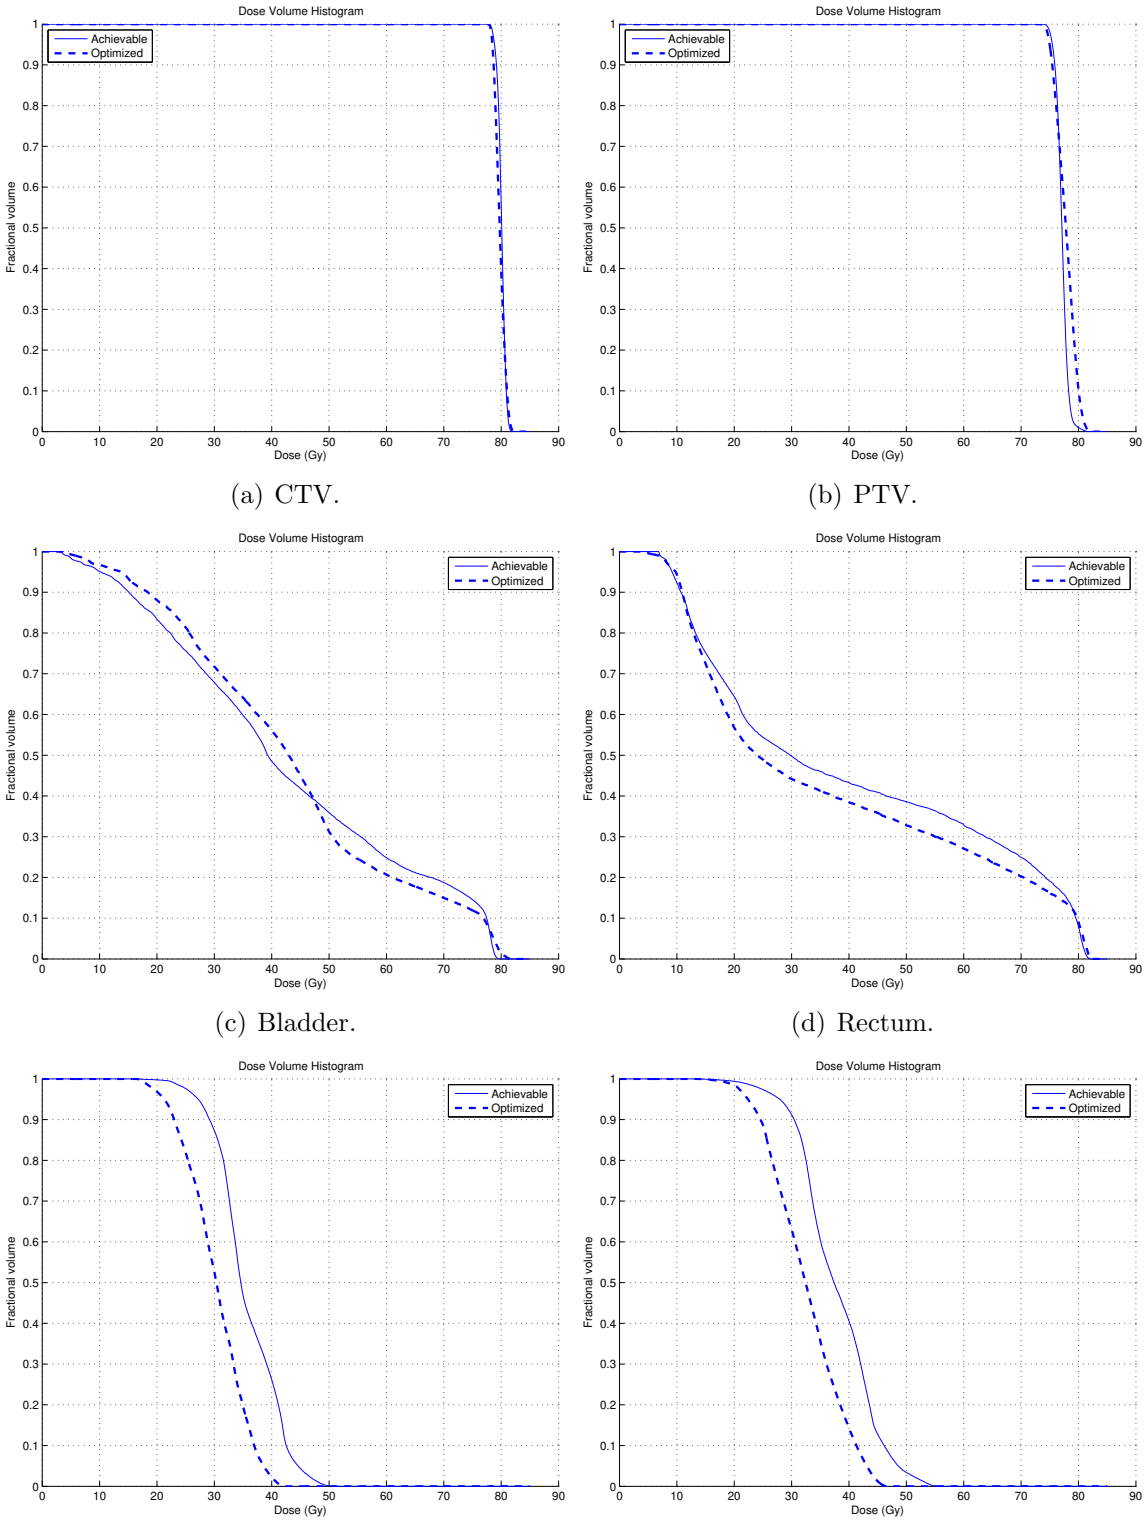

(a) CTV.

(b) PTV.

(c) Bladder.

(d) Rectum.

(e) Left femoral head.

(f) Right femoral head.

Figure 4.5: DVHs of the clinically achievable and inversely optimized plans for patient #3 using GIOP$_r$ Model (A).

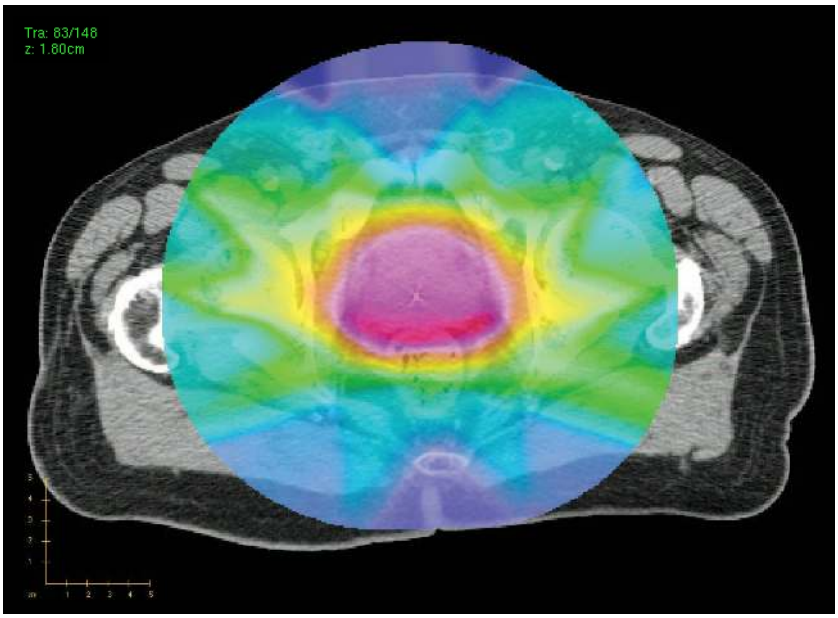

(a) Clinically achievable dose distribution.

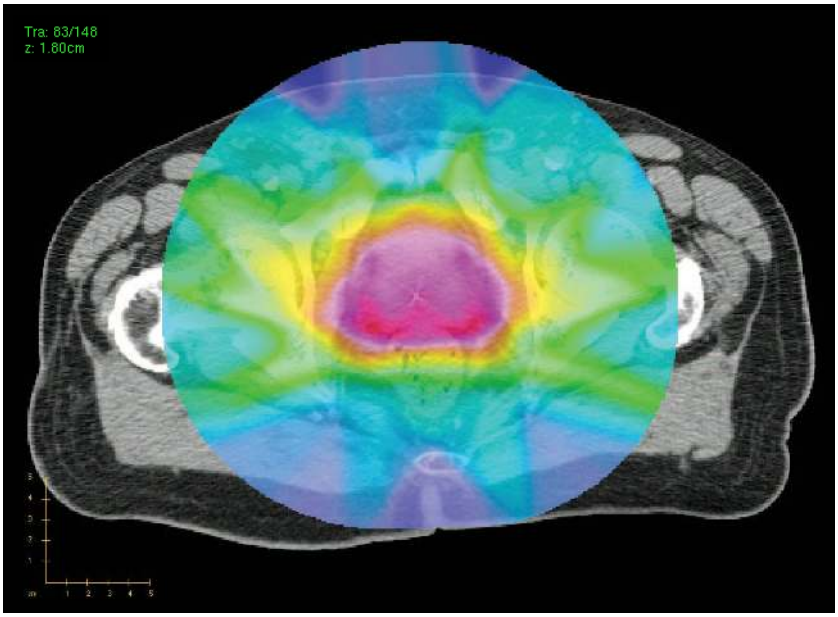

(b) Inversely optimized dose distribution.

Figure 4.6: Dose distributions of the clinically achievable and inversely optimized plans for patient #1 using GIOP$_r$.

Table 4.4: Dose metrics computed from inversely optimized plans for patient #1.

Structure	Metric	Target level (%)	Clinical plan (%)	Inversely optimized plan (%)		
				Model (A)	Model (B)	Model (C)
CTV	V78.0	≥ 99	100.00	100.00	100.00	100.00
PTV	V74.1	≥ 99	100.00	100.00	100.00	100.00
PTV	V81.9	$\leq$100	100.00	100.00	100.00	100.00
Bladder	V70.0	≤ 30	21.72	20.53	21.13	21.27
Bladder	V54.3	≤ 50	36.90	34.54	36.03	36.09
Rectum	V70.0	≤ 30	26.79	24.65	26.26	26.29
Rectum	V54.3	≤ 50	41.15	38.31	39.97	39.87
L. femoral head	V52.0	≤ 5	0.00	0.00	0.00	0.00
R. femoral head	V52.0	≤ 5	0.01	0.01	0.02	0.02

Table 4.2. Similar results to those shown in this section were obtained for the remaining patients. A summary of the GIOP$_r$ results for all 12 patients is provided in Table 4.5. The GIOP$_a$ results were omitted.

4.4.3 Identifying important objective functions

Treatment planners often do not know which objective functions influence the optimization the most and which are the most crucial in designing a high quality, clinically acceptable treatment for a particular patient. As a result, clinical treatment planning formulations may have many more objectives than needed and cause a large parameter space to search over when iteratively designing a treatment. In this section, we demonstrate how our GIOP methodology can identify important (or redundant) objectives from a large family of candidate objectives.

Consider the forward formulation (C). The results from solving the associated GIOP$_r$ for each of the 12 patients is shown in Table 4.6. Notice that the values of ϵ_r^* in Table 4.6 are less than their corresponding values in Table 4.5, which is expected because the objective functions used in model (C) contain the ones used in model (A). Analogous to a regression, there are simply more "explanatory variables" in model (C). Therefore we expect the duality gap to be smaller as more of the variation in the dose distribution

Table 4.5: Summary of results from GIOP$_r$ model (A) for all patients.

Pat.	ϵ_r^*	$\boldsymbol{\alpha}_r^*$				Time
		Blad	Rect	L.Fem	R.Fem	(s)
1	1.098	0.248	0.746	0.002	0.004	318
2	1.087	0.616	0.380	0.001	0.003	347
3	1.205	0.088	0.906	0.004	0.002	463
4	1.107	0.007	0.989	0.001	0.003	238
5	1.067	0.963	0.023	0.012	0.002	844
6	1.108	0.972	0.011	0.006	0.011	465
7	1.073	0.297	0.695	0.006	0.002	366
8	1.071	0.672	0.315	0.007	0.006	329
9	1.091	0.980	0.010	0.007	0.003	363
10	1.231	0.833	0.147	0.010	0.010	238
11	1.164	0.956	0.040	0.002	0.002	491
12	1.063	0.918	0.066	0.006	0.010	388

Gray cells indicate a weight greater than 0.10.

Table 4.6: Summary of results from GIOP$_r$ model (C) for all patients.

Pat.	ϵ_r^*	$\theta^i = 0$		$\theta^i = 10$		$\theta^i = 20$		$\theta^i = 30$		$\theta^i = 40$		$\theta^i = 50$		$\theta^i = 60$		$\theta^i = 70$		Max		Time
		Bl	Re	Bl	Re	Bl	Re	Bl	Re	Bl	Re	Bl	Re	Bl	Re	Bl	Re	LF	RF	(s)
1	1.027	0.58	0.06	-	-	-	-	-	-	-	-	0.06	0.01	-	-	0.27	-	0.01	0.01	1094
2	1.027	0.54	0.01	-	-	-	-	-	-	0.08	-	0.26	0.09	-	-	-	-	0.01	0.01	1702
3	1.092	0.41	0.01	-	-	-	-	-	-	-	-	-	-	0.05	-	0.51	∼0.00	0.01	0.01	1813
4	1.092	0.01	∼0.00	-	-	-	0.01	-	-	-	-	-	-	-	0.98	-	-	∼0.00	∼0.00	1188
5	1.038	0.18	-	-	-	-	-	-	-	-	-	0.66	-	0.12	0.02	-	∼0.00	0.01	0.01	2973
6	1.047	0.73	-	-	-	-	-	-	-	-	-	0.23	-	-	-	-	-	0.01	0.03	2445
7	1.063	0.08	-	-	-	-	-	-	-	-	-	0.37	0.47	-	0.07	-	-	0.01	∼0.00	1570
8	1.033	0.62	0.02	-	-	-	-	-	-	0.31	-	-	0.01	-	-	-	-	0.02	0.02	1488
9	1.068	0.15	∼0.00	-	-	-	-	-	-	-	-	0.52	-	-	∼0.00	0.32	-	0.01	∼0.00	1785
10	1.076	0.94	-	-	-	-	-	-	-	0.02	-	-	-	-	-	-	0.01	0.02	0.01	1111
11	1.041	0.62	-	-	-	-	-	-	-	-	-	0.36	-	-	-	-	∼0.00	0.01	0.01	1863
12	1.034	0.23	-	-	-	-	-	-	-	-	-	0.07	0.03	0.64	-	-	-	0.01	0.02	1564

Bl = bladder, Re = rectum, LF = left femoral head, RF = right femoral head. Gray cells indicate a weight greater than 0.10. Dashes indicate a weight of zero.

is explained by the increased number of objective functions. However, having more objective functions and a smaller duality gap does not necessarily imply that the final treatment using the inversely optimized weights will result in better clinical performance. Table 4.4 shows the dose criteria corresponding to the inversely optimized plan for patient #1 using the 18 objective functions (model (C)), as well as the four objective functions from the previous section (model (A)). In this case, we see that model (A) performs slightly better than model (C), which we believe is due to the fact that model (C) is "overfitting" patient #1's input. Having more objective functions makes it easier to fit the entire dose distribution, but may actually put undue emphasis on parts of the dose distribution that are less relevant clinically. For example, given that the primary bladder and rectum criteria are partial dose-volume criteria above 54.3 and 70.0 Gy, including additional objectives may confound the optimization.

Figures 4.7, 4.8, and 4.9 show that the DVHs for the clinically achievable and inversely determined plans using Model (C) are very similar. Compared to the DVH comparisons when Model (A) was used (i.e., Figures 4.3, 4.4, and 4.5), the plans generated via Model (C) look closer to the clinical ones than those generated via Model (A) are. However, again, Model (C) seems to put undue emphasis on the DVH regions that are not so much relevant to the clinical acceptability criteria (e.g., regions that correspond to dose less than 50.0 Gy), i.e., having plans more similar to the clinical plans does not necessarily imply better treatment quality.

Table 4.5 shows that in model (A), roughly 98-99% of the objective function weights are placed on the bladder and rectum objectives. When additional bladder and rectum objectives are included, Table 4.6 shows that most of the weights remain on the bladder and rectum objectives (95-99%). Among the bladder and rectum objectives, only a few receive nonzero weights and the top two account for 80-99% of the weight. For almost all patient cases, the two most heavily weighted objective functions for the bladder and rectum are the mean dose objective ($\theta^i = 0$) and one of the ones corresponding to $\theta^i =$

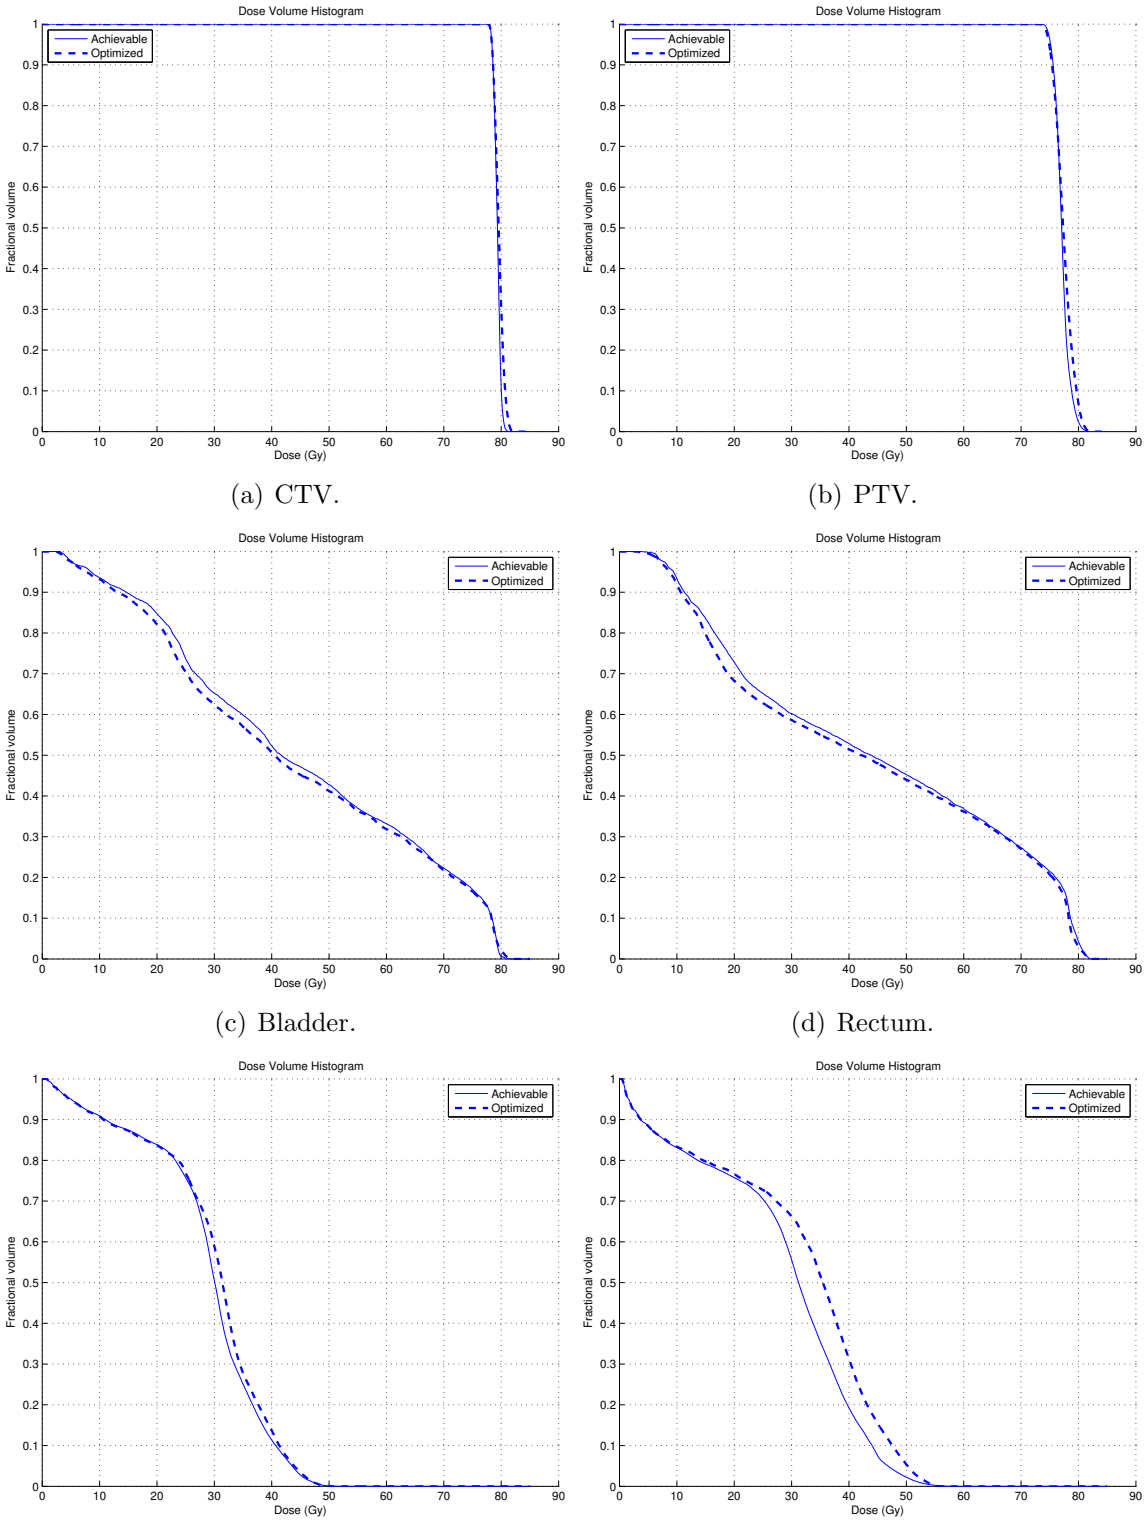

(a) CTV.

(b) PTV.

(c) Bladder.

(d) Rectum.

(e) Left femoral head.

(f) Right femoral head.

Figure 4.7: DVHs of the clinically achievable and inversely optimized plans for patient #1 using GIOP$_r$ Model (C).

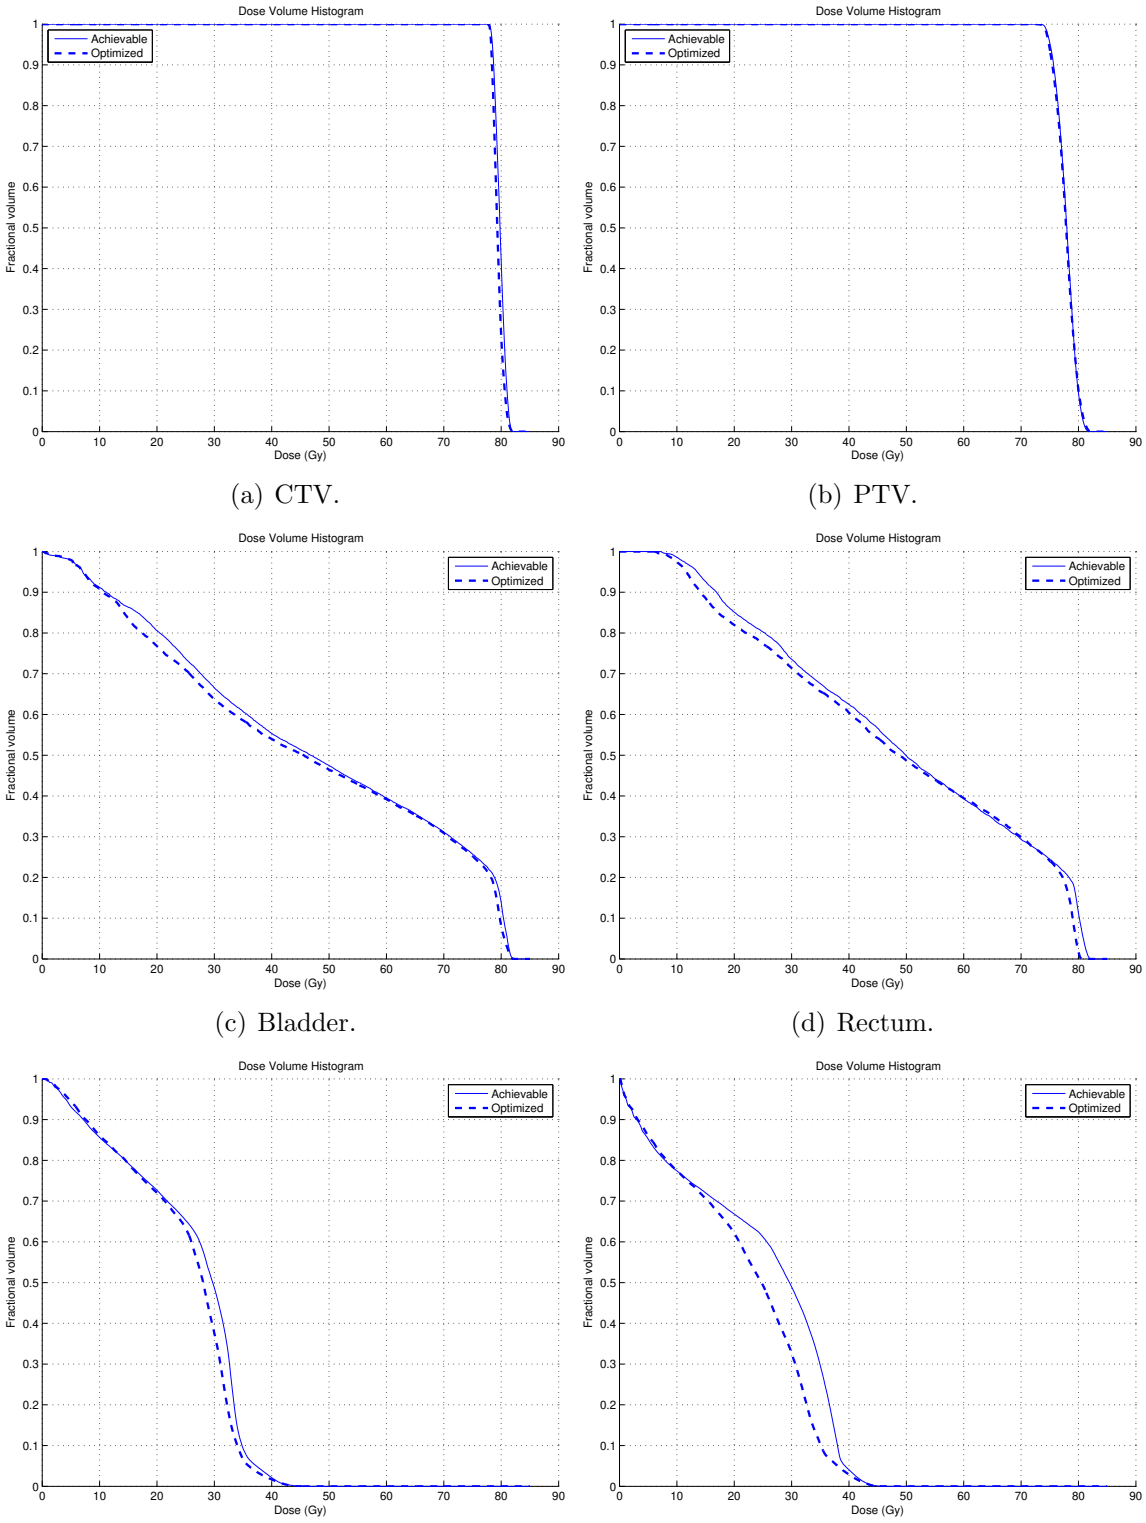

(a) CTV.

(b) PTV.

(c) Bladder.

(d) Rectum.

(e) Left femoral head.

(f) Right femoral head.

Figure 4.8: DVHs of the clinically achievable and inversely optimized plans for patient #2 using GIOP$_r$ Model (C).

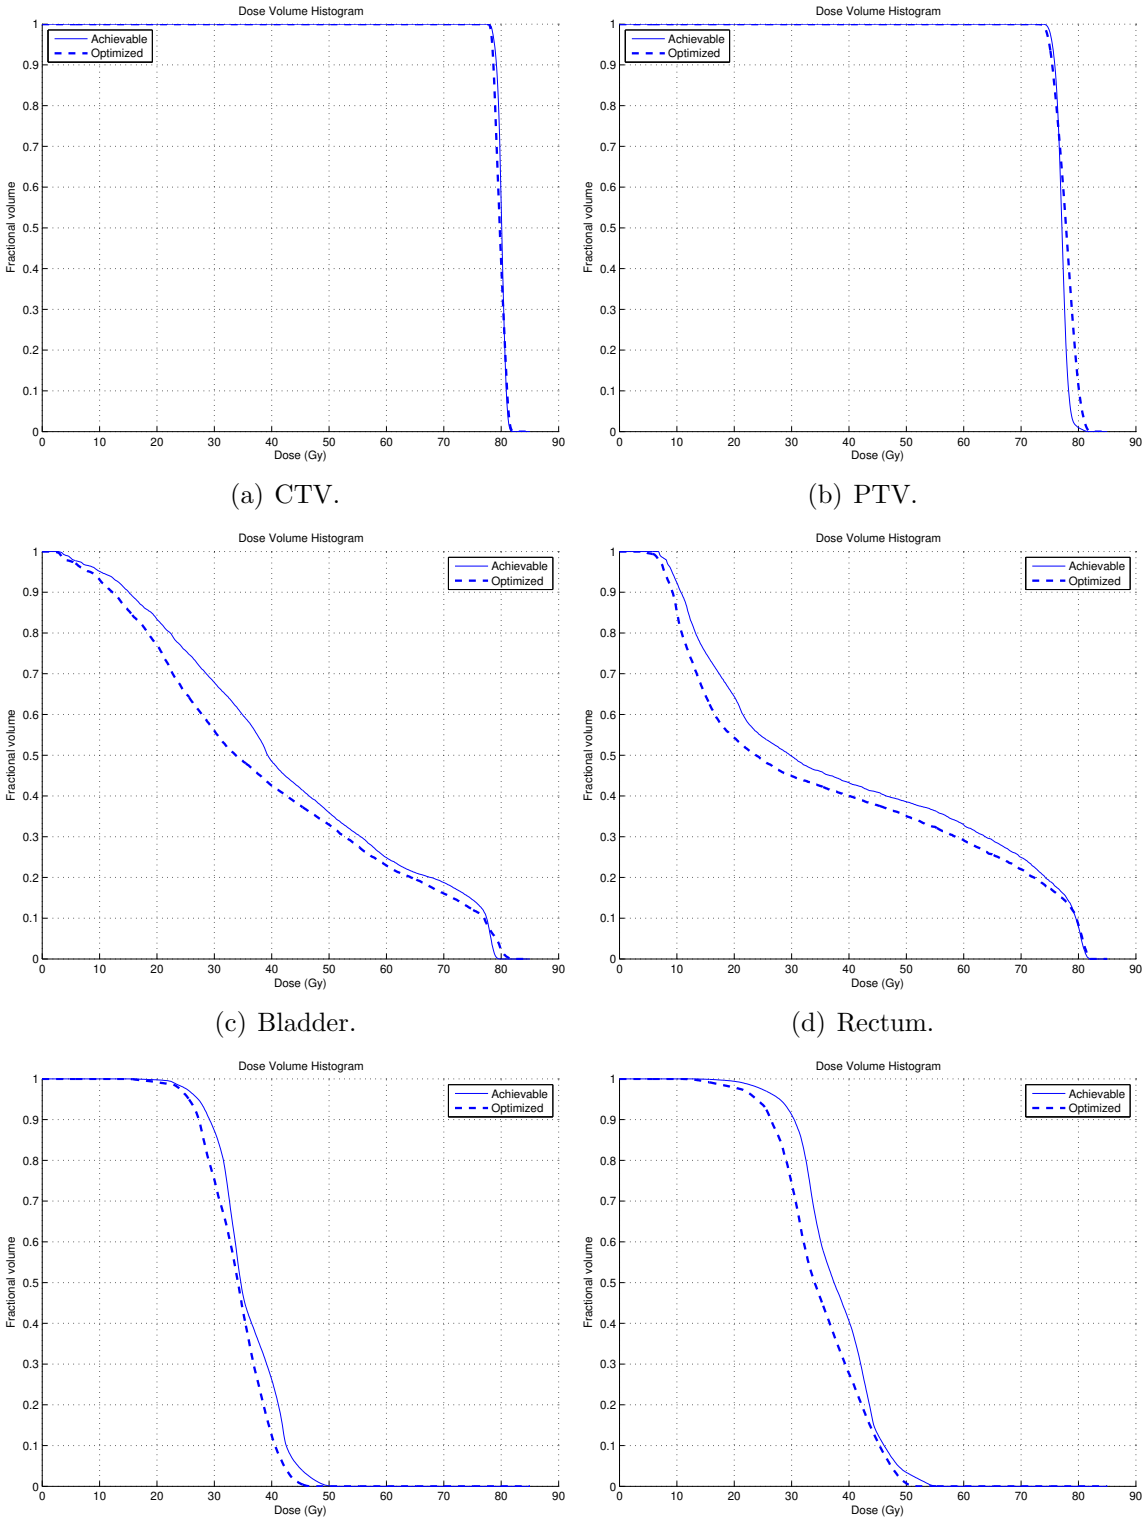

(a) CTV.

(b) PTV.

(c) Bladder.

(d) Rectum.

(e) Left femoral head.

(f) Right femoral head.

Figure 4.9: DVHs of the clinically achievable and inversely optimized plans for patient #3 using GIOP$_r$ Model (C).

Table 4.7: Summary of results from GIOP$_r$ model (B) for all patients.

Pat.	ϵ_r^*	$\theta^i = 0$		$\theta^i = 50$		Max		Time
		Blad	Rect	Blad	Rect	L.Fem	R.Fem	(s)
1	1.027	0.704	0.073	0.192	0.007	0.012	0.012	426
2	1.027	0.518	0.002	0.371	0.093	0.006	0.010	517
3	1.093	0.710	0.013	0.258	-	0.011	0.008	527
4	1.105	0.030	0.003	-	0.961	0.002	0.004	286
5	1.039	0.198	-	0.784	0.005	0.010	0.003	990
6	1.047	0.730	-	0.231	-	0.010	0.029	651
7	1.063	0.087	-	0.393	0.510	0.007	0.003	494
8	1.033	0.737	0.014	0.199	0.009	0.020	0.021	478
9	1.068	0.213	0.005	0.766	0.002	0.011	0.003	519
10	1.076	0.968	-	-	-	0.017	0.015	349
11	1.041	0.622	-	0.369	-	0.003	0.006	511
12	1.035	0.366	-	0.553	0.044	0.014	0.023	502

Gray cells indicate a weight greater than 0.10. Dashes indicate a weight of zero.

$50, 60,$ or 70. The objectives corresponding to $\theta^i = 10, 20, 30,$ and 40 generally receive little to no weight, and therefore do not play a central role in driving the optimization. These results are aligned with the choice of objectives for the bladder and rectum in the clinical treatment planning formulation (cf. Table 3.1).

To further explore this issue, we modified the forward problem to only include six objective functions (model (B)). Inversely optimized weights derived from solving the corresponding GIOP$_r$ for all 12 patients are shown in Table 4.7. Note that the values of ϵ_r^* from Table 4.7 are very similar to those in Table 4.6, indicating that the six objectives in model (B) are capturing almost all the explanatory power of the 18 objectives from model (C). We performed further validation by solving the forward problem with the inversely optimized weights. Results for patients #1, 2, and 3 are shown in Figures 4.10, 4.11, and 4.12, respectively. For all of them, even with a different number of objective functions, the dose distributions from model (B) and (C) are virtually identical. The dose criteria achieved by the two models are also nearly identical (Table 4.4).

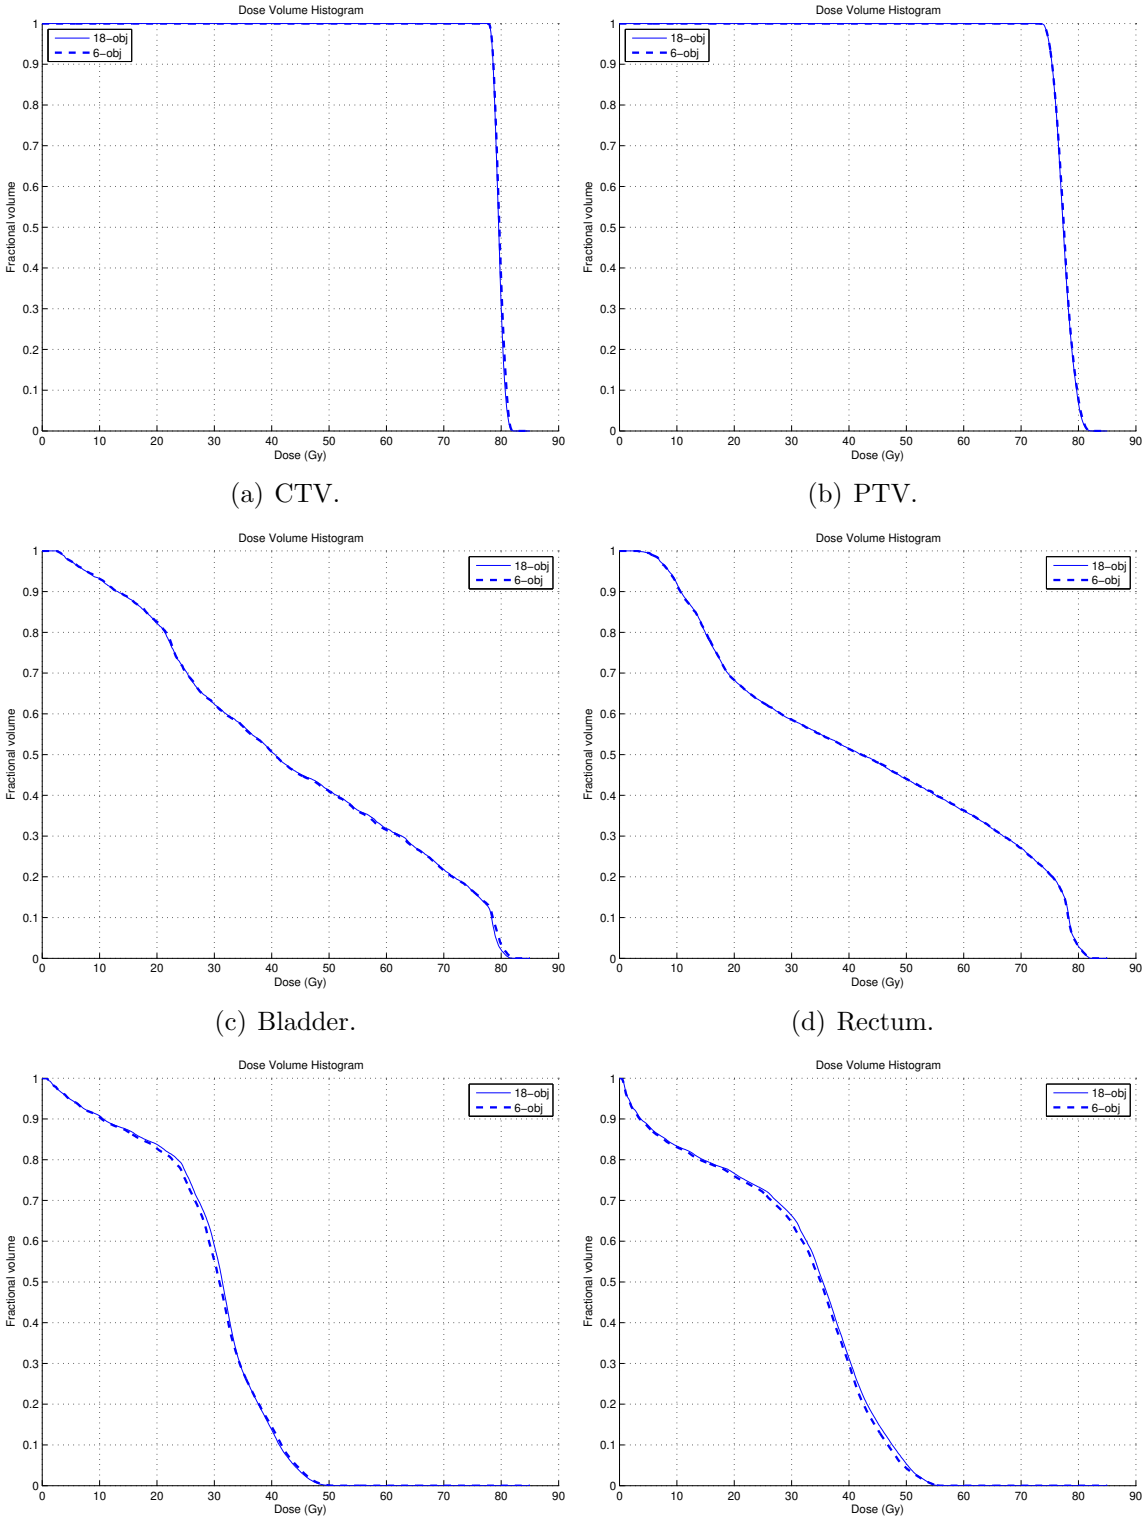

(a) CTV.

(b) PTV.

(c) Bladder.

(d) Rectum.

(e) Left femoral head.

(f) Right femoral head.

Figure 4.10: DVHs of inversely optimized plans from model (B) (six objectives) and model (C) (18 objectives) for patient #1.

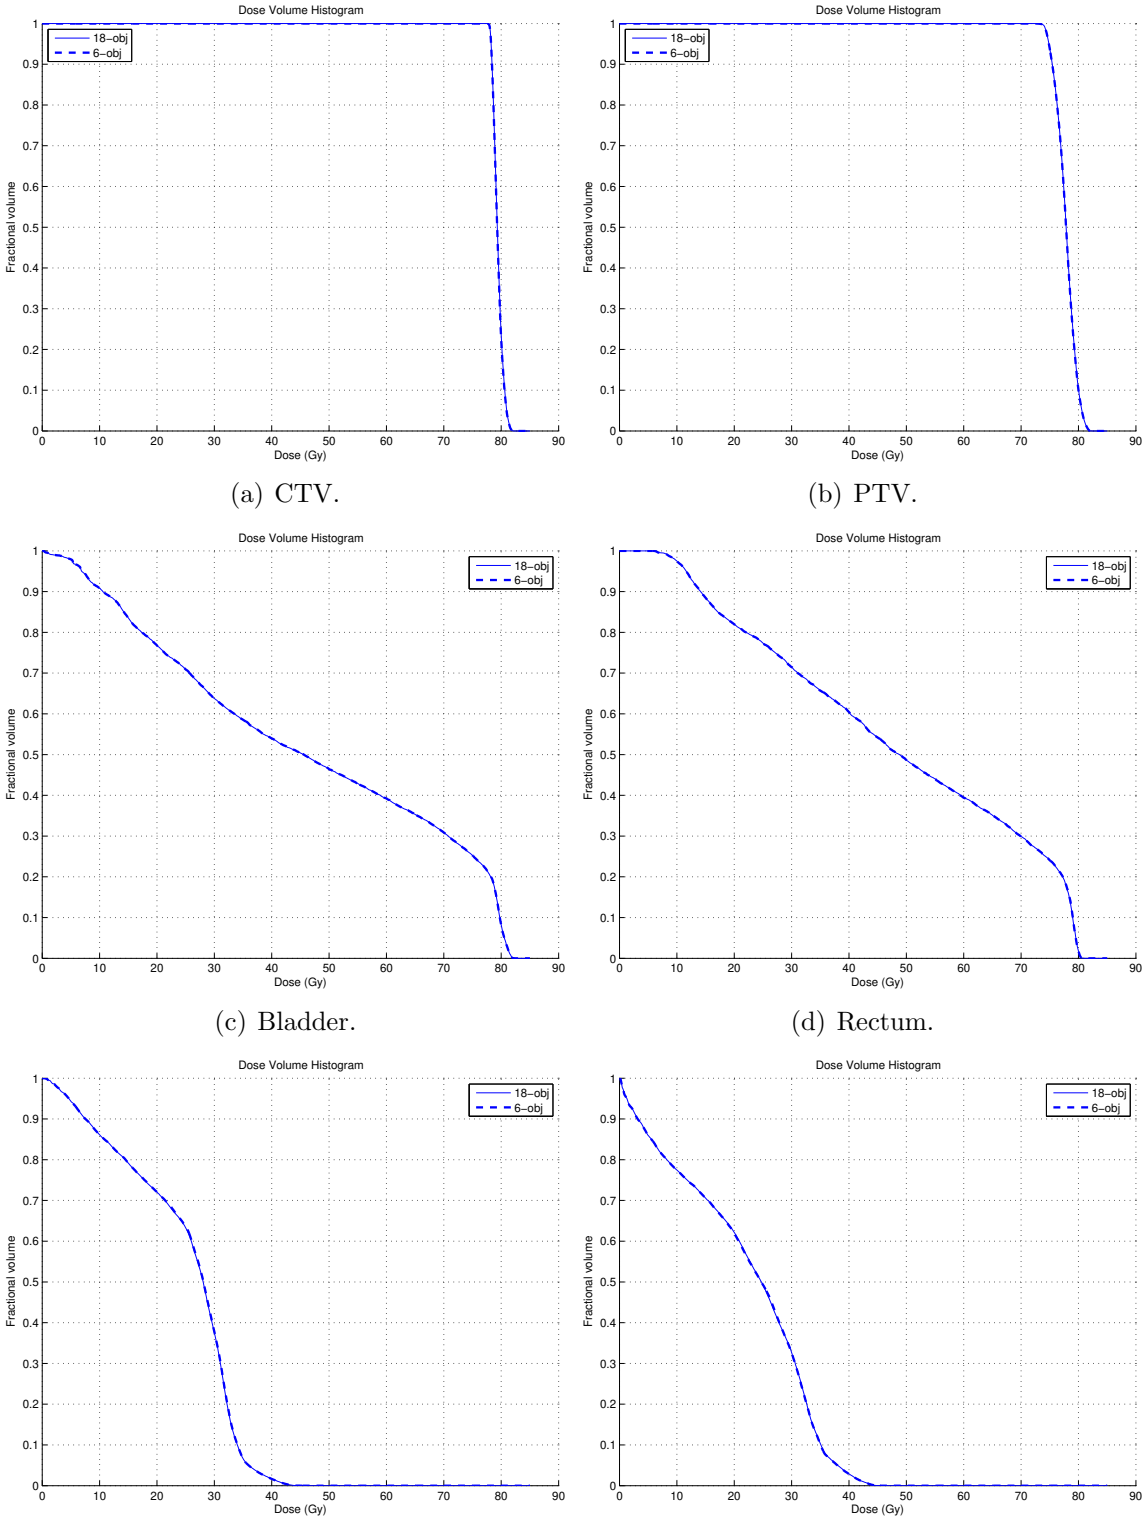

(a) CTV.

(b) PTV.

(c) Bladder.

(d) Rectum.

(e) Left femoral head.

(f) Right femoral head.

Figure 4.11: DVHs of inversely optimized plans from model (B) (six objectives) and model (C) (18 objectives) for patient #2.

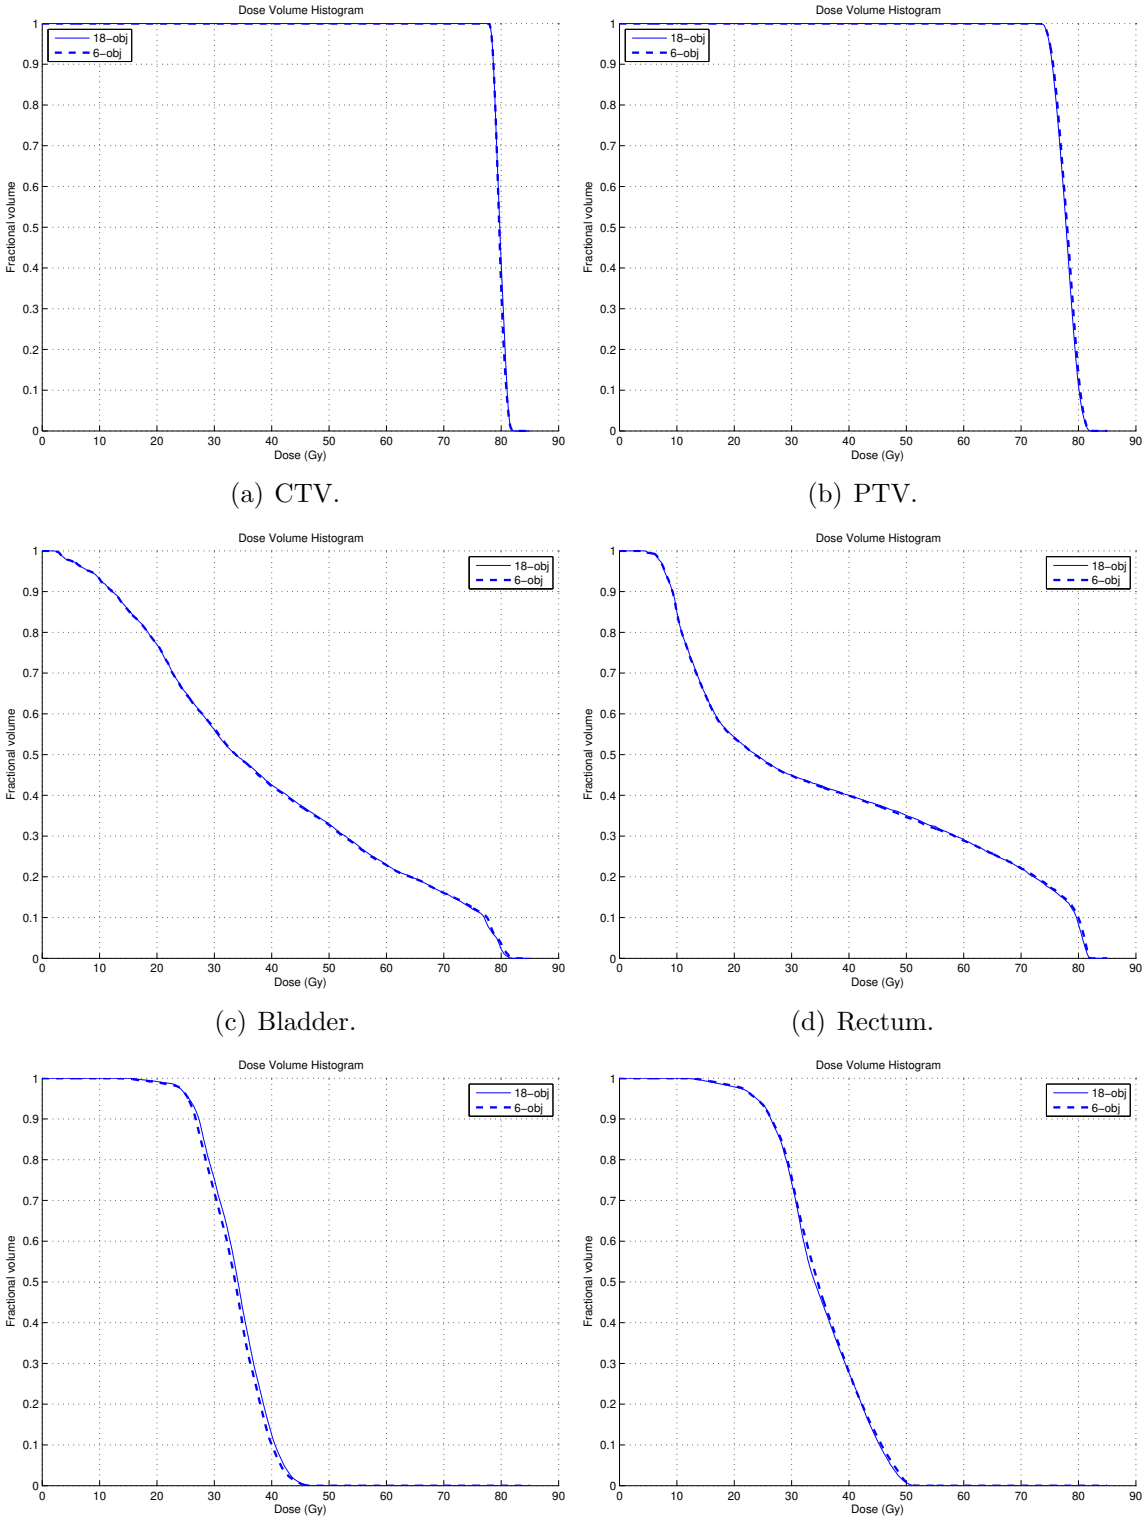

(a) CTV.

(b) PTV.

(c) Bladder.

(d) Rectum.

(e) Left femoral head.

(f) Right femoral head.

Figure 4.12: DVHs of inversely optimized plans from model (B) (six objectives) and model (C) (18 objectives) for patient #3.

4.4.4 Implications for prostate cancer treatment design

As shown in the previous section, almost all objective function weights are placed on bladder and rectum objectives. This result suggests that these objectives are the most important ones in determining the final clinical treatment plan, and therefore, should not receive the same weight as the femoral head objectives. However, whether the bladder or rectum is more important depends on the patient, which reinforces the idea that no one set of weights will be universally applicable to each patient. Classifying patients based on their anatomical geometry (e.g., the size of the rectum and bladder or their proximity to the prostate) may be an appropriate way to stratify patients. The inversely optimized weights can be thought of as being proportional to the level of difficulty in achieving the corresponding criteria, reflecting the anatomical geometry inherent in prostate cancer cases. As shown in Figure 3.1, the bladder and rectum are in close proximity to the target, overlapping the PTV, whereas the femoral heads are located farther from the target. Determining a statistical relationship between objective function weights and geometric quantities such as the distance between the prostate and OAR centroids, and the volume of overlap between the prostate and OARs, would provide insight into the impact of different geometries on treatment planning.

As an immediate consequence of the concentration of weight values in a few objectives, it seems that current clinical treatment planning formulations may have more objective functions than needed to generate acceptable treatment plans. This observation has potential efficiency and effectiveness implications. From an operational efficiency point of view, having a large number of objectives may result in a large parameter space to search through when designing treatments. Because of the iterative nature of treatment planning, it may take longer to design a treatment. Simplification of the treatment planning process would be useful for new institutions starting prostate cancer treatments and new treatment planners being trained in treatment planning. The benefit of faster planning is amplified for adaptive radiation therapy (Yan et al., 1997), an increasingly

popular paradigm where a plan is re-optimized regularly over the course of a multi-week treatment (instead of the current paradigm of creating a single treatment that is delivered daily over many weeks). Furthermore, having a smaller model may allow us to take uncertainty into account (e.g., through stochastic programming or robust optimization) without the resulting model being overly large. Also, a smaller number of objectives makes it easier to construct and visualize the Pareto frontier, which is becoming a popular approach clinically.

It could be argued that having a large number of objectives may compensate for the sensitivity in the choice of weights, and thus it would be easier to find appropriate weights compared to a simpler formulation. We found through computational experiments that having more objectives did not reduce the sensitivity of the formulation to the objective function weights, and could even increase its sensitivity. We generated weight vectors by randomly perturbing each component of an inversely-optimized weight vector plus or minus 0.1 (ensuring the resulting value was between 0 and 1), and then re-normalizing the weight vector. We made this adjustment 100 times for each patient, solved models (A), (B), and (C) with these new weights, and plotted the resulting DVHs. Results for one patient are shown in Figures 4.13, 4.14, 4.15, and 4.16, for the bladder, rectum, left femoral head, and right femoral head, respectively.

Each line represents one forward optimization run using a simulated $\boldsymbol{\alpha}$ vector. The thickness of the cloud of lines indicates how sensitive the dose distribution of the organ is to the choice of weights. The cloud of lines for model (A) seems to be about the same thickness or even thinner than the corresponding clouds for models (B) and (C). Similar results were seen for the other patients. Here, placing even a small non-zero weight on some of the objectives that were excluded from model (A) seems to result in more sensitivity of the final treatment. We believe this observation reinforces the idea that having more objective functions may result in undue emphasis being placed on objectives that are not as critical to the ultimate quality of the treatment. However, it is also worth

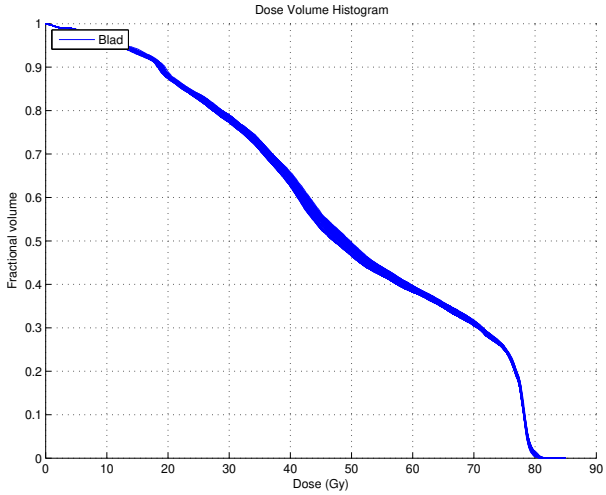

(a) Bladder, model (A) (4-obj)

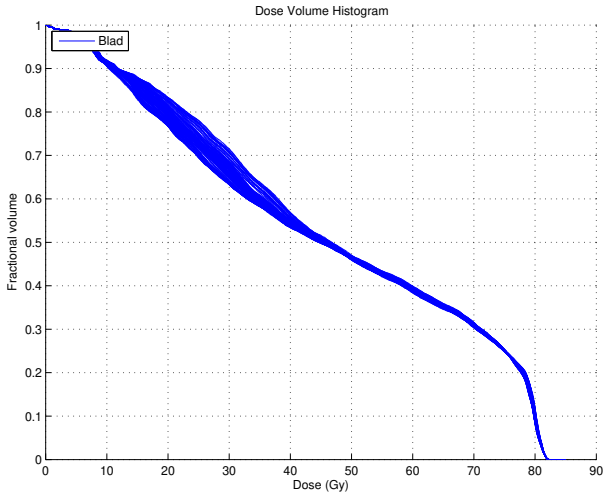

(b) Bladder, model (B) (6-obj)

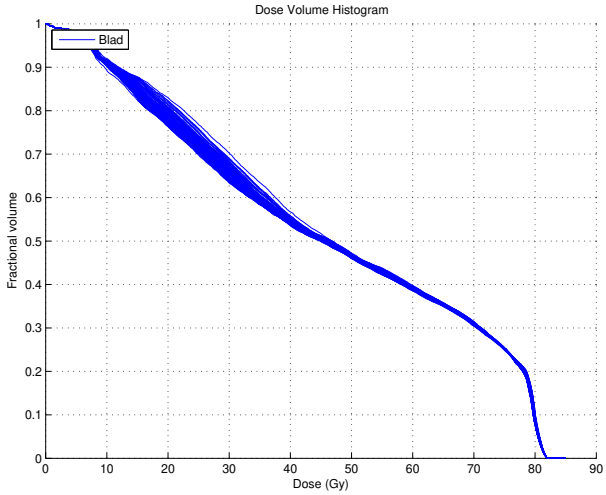

(c) Bladder, model (C) (18-obj)

Figure 4.13: DVH clouds for the bladder from 100 sets of weights in the neighborhood of the inversely optimized weights.

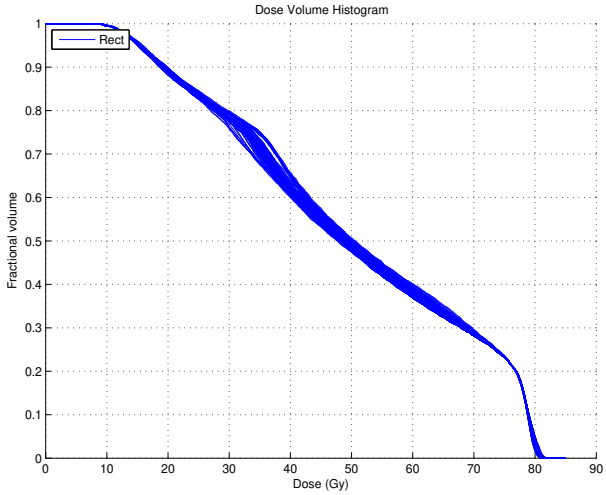

(a) Rectum, model (A)

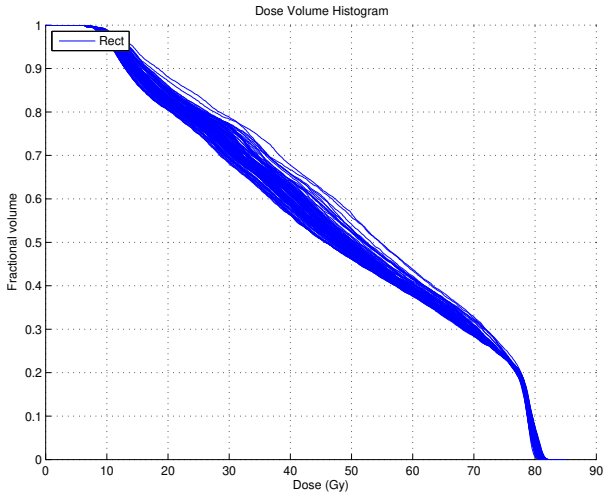

(b) Rectum, model (B)

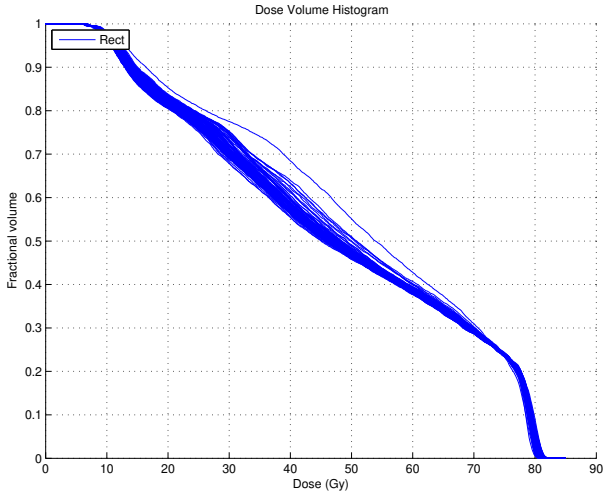

(c) Rectum, model (C)

Figure 4.14: DVH clouds for the rectum from 100 sets of weights in the neighborhood of the inversely optimized weights.

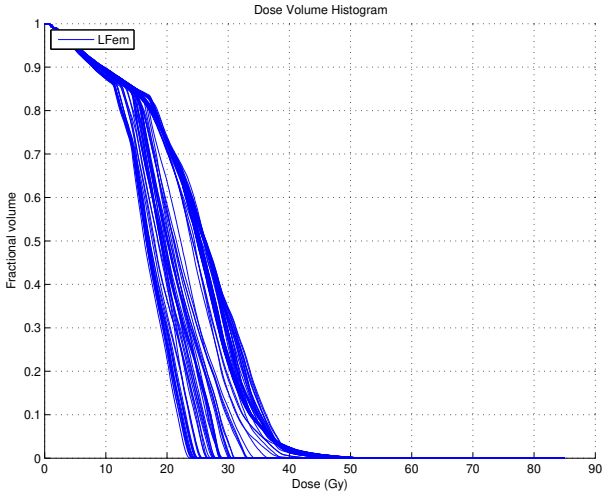

(a) Left femur, model (A)

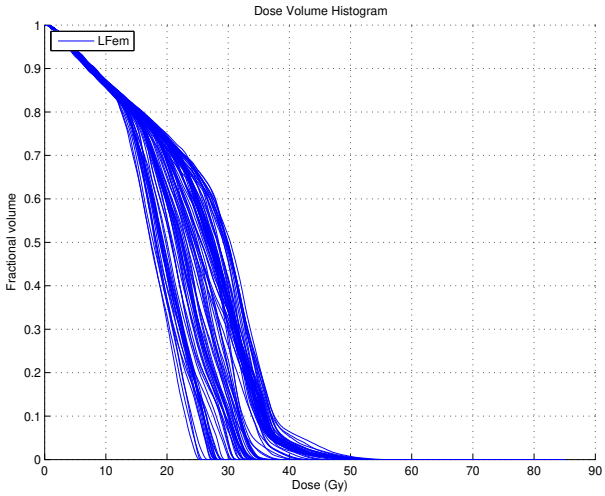

(b) Left femur, model (B)

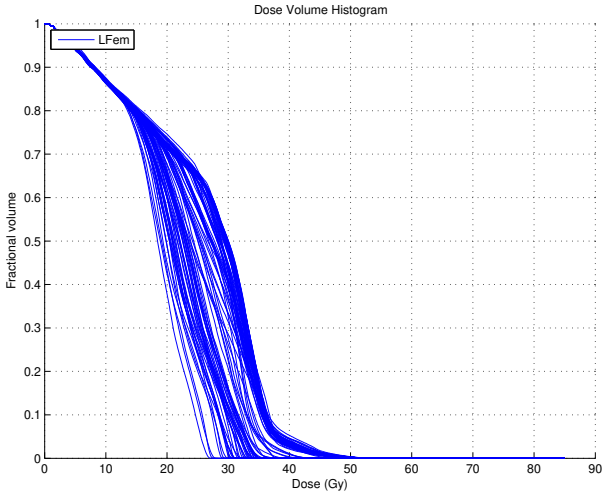

(c) Left femur, model (C)

Figure 4.15: DVH clouds for the left femoral head from 100 sets of weights in the neighborhood of the inversely optimized weights.

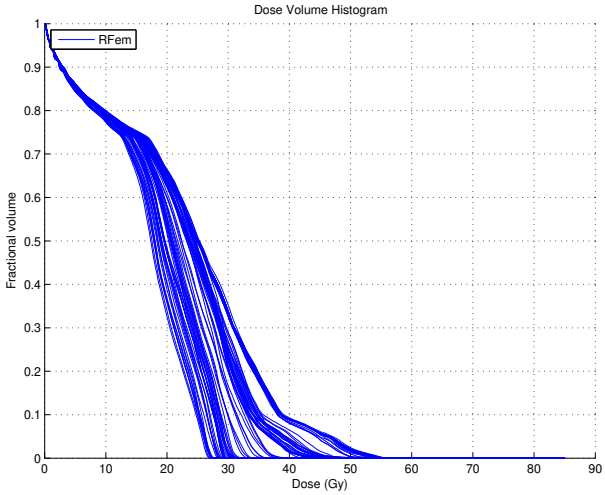

(a) Right femur, model (A)

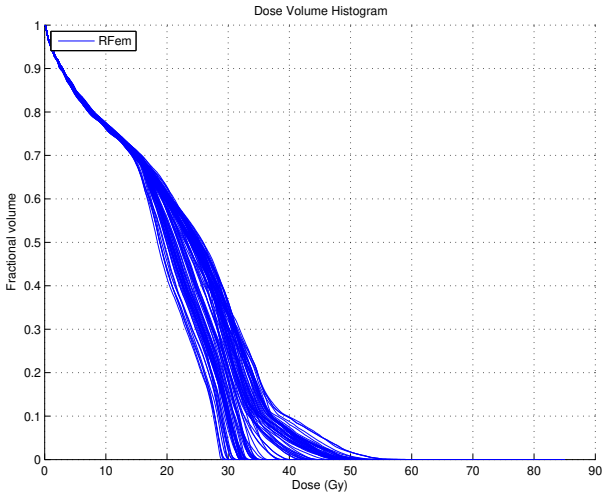

(b) Right femur, model (B)

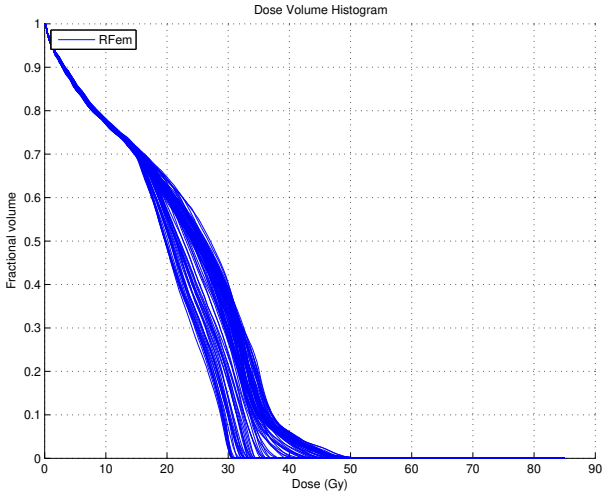

(c) Right femur, model (C)

Figure 4.16: DVH clouds for the right femoral head from 100 sets of weights in the neighborhood of the inversely optimized weights.

noting that when comparing models (B) and (C), model (B) seems to be slightly more sensitive. Overall, it seems that the number of objectives may not be an ideal indicator of the sensitivity of the solution to the choice of weights. From the clinical effectiveness point of view, having more objective functions could lead to a treatment that is less able to satisfy clinical criteria, as shown in the previous section.

The sensitivity of the bladder and rectum, especially in the higher dose region, seems to be much lower than that of the femoral heads, which is likely due to the proximity of the bladder and rectum to the PTV. That is, the requirement to deliver a certain dose to the PTV provides less flexibility to reduce dose to nearby organs (or organs that overlap the PTV). Sensitivity is the subject of ongoing research. The results in this chapter should be taken as a starting point for investigating potential improvements in the treatment planning process, but more research is needed to clearly demonstrate clinical impact in this area.

The results in Section 4.4.3 also highlight a connection with the equivalent uniform dose (EUD) concept (Niemierko, 1997). The idea behind EUD is to encapsulate in a single number a measure of the biological impact of a heterogeneous dose distribution. The EUD is the equivalent dose that if delivered uniformly to a structure would have the same biological impact as some heterogeneous dose distribution. Mathematically, given a vector $\mathbf{d}$ whose components d_v represent the dose to voxel v in some structure $\mathcal{O}$, the EUD is

$$\text{EUD}(\mathbf{d}; a) = \left(\frac{1}{|\mathcal{O}|} \sum_{v \in \mathcal{O}} {d_v}^a \right)^{\frac{1}{a}}, \tag{4.17}$$

where a is a structure-specific parameter that describes the biological dose-volume effect. If $a < 1$, then lower doses are given greater weight and the EUD function is used for the tumour and other target structures. If $a \geq 1$, the EUD function is used for OARs. Note that if $a = 1$, then the EUD is simply the mean dose. As $a \to \infty$, the EUD approaches the maximum dose.

Consider the following linearization of the EUD function (4.17) for OARs (Thieke

et al., 2002):

$$\text{EUD}_L(\mathbf{d};\lambda) = \lambda d^{mean} + (1-\lambda)d^{max}. \tag{4.18}$$

Here, d^{mean} and d^{max} represent the mean and maximum dose delivered to the OAR, respectively, and $\lambda \in [0,1]$ is an OAR-specific parameter. It has been shown that by optimizing EUD_L instead of the original EUD function (4.17), comparable treatments can be created (Thieke et al., 2002). Note that for a tumour or target structure, d^{max} would be replaced by d^{min}, the minimum dose to the structure. The results in Section 4.4.3 show that by generalizing the idea of the EUD_L beyond mean, minimum and maximum doses, and allowing a convex combination of the mean dose objective and some linear penalty objective, clinical quality treatments can still be generated.

4.5 Conclusion

This chapter adds to the growing literature on inverse optimization by developing a new methodology to address the situation where the input data renders the inverse problem infeasible, a topic that has received little attention to date. In the context of a multiobjective forward optimization problem, given a feasible solution that is not weakly efficient, our method determines objective function weights that make the given solution approximately optimal (weakly efficient) with as small a duality gap as possible. Our method generalizes the standard inverse optimization methodology and retains the underlying complexity of the forward problem. Thus, there is little downside to adopting the more generalized approach when solving inverse optimization problems in practice. We also elucidate a connection between our generalized inverse optimization approach and Pareto surface approximation techniques from the literature.

We applied our generalized inverse optimization method to historical prostate cancer radiation therapy treatments in order to determine the objective functions and corresponding weights that most heavily influenced the optimization of the treatments. Being

able to demonstrate value in prostate IMRT provides a springboard to address more complex sites such as head and neck cancers, where the number and importance of nearby OARs increase significantly. Applying inverse optimization to historical data can help identify a small, but relevant set of objectives for multiobjective IMRT planning methods, such as Pareto surface navigation methods. Overall, we believe that inverse optimization has potential to quantify the implicit preferences of decision makers in clinical settings using historical data.

Many directions for future research exist. As an immediate future research direction from the results of Section 4.4.3, incorporating the ability of determining the minimum number of objectives that can capture the real model of data in the inverse optimization framework may generalize the idea of stepwise regression, which aims to find independent variables in a regression model that represent the data. In fact, topics of connecting statistics and optimization have recently received a growing amount of attention in the optimization community (e.g., Bertsimas and Mazumder (2014)). Also, the solution of an inverse multiobjective optimization problem induces a prioritization of the objectives, which may be used in a prioritized optimization approach such as preemptive or lexicographic goal programming. Methods to efficiently explore or approximate the Pareto surface can be augmented using inverse optimization to identify search directions that focus on only the most relevant parts of the Pareto surface. Lastly, our method in this chapter will be extended to general convex optimization, which will open the door for much broader application.

Chapter 5

Predicting objective function weights from patient anatomy in prostate IMRT treatment planning

5.1 Introduction

Using clinical prostate treatment plans as input, in Chapter 4, the generalized inverse optimization model determined that the vast majority of the weight was placed on objectives corresponding to the bladder and rectum, though the exact values varied from patient to patient. Some patients were heavily rectum-weighted while others were heavily bladder-weighted. This observation led to the question of whether there is a quantifiable metric that differentiates patients who are more "rectum-weighted" and patients who are more "bladder-weighted". In particular, assuming other factors are controlled for (same beam angles, same treatment planning team, etc.), it was hypothesized that the underlying anatomy is responsible for the difference in weights. Thus, in this chapter we aim to develop a quantifiable relationship between anatomy surrounding the prostate and optimal objective function weight values.

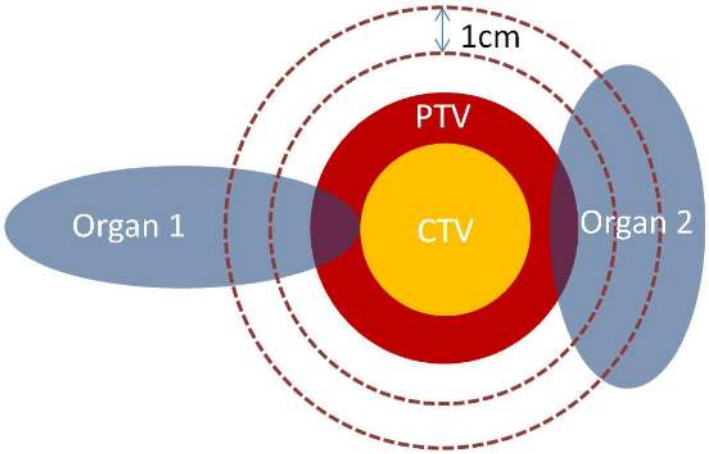

Figure 5.1: Illustration of two OARs of which relative spatial information can be obtained via OVHs.

Recent studies have considered using patient anatomical information to control or improve IMRT treatment quality, with the ultimate goal of automating the treatment planning process (Lu et al., 2006; Wu et al., 2009, 2011; Moore et al., 2011; Chanyavanich et al., 2011; Zhu et al., 2011; Yang et al., 2013). This type of treatment planning is often referred to as "knowledge-based treatment planning". These studies focused on the relationship between the anatomy and achievable dose in order to analyze "what could have been done" in the past treatments and "what will be possible" for future patients. The overlap volume histogram (OVH) has recently been proposed to quantify the spatial relationship between OARs and the planning target volume (PTV) (Wu et al., 2009). The OVH computes fractional overlap volume of an OAR with PTV expansions of different amounts and provides a way to compare the relative spatial information between different patients. For example, an OAR that has more volume around the PTV will result in a steeper OVH, which implies that the organ is at a higher risk and requires more careful treatment than other organs (see Figures 5.1 and 5.2). In recent applications of the OVH, given a new patient, the lowest dose delivered to historical patients who have OVHs as low as the new patient was found for each OAR, and used as a dose objective for the new patient (Wu et al., 2009, 2011; Yang et al., 2013).

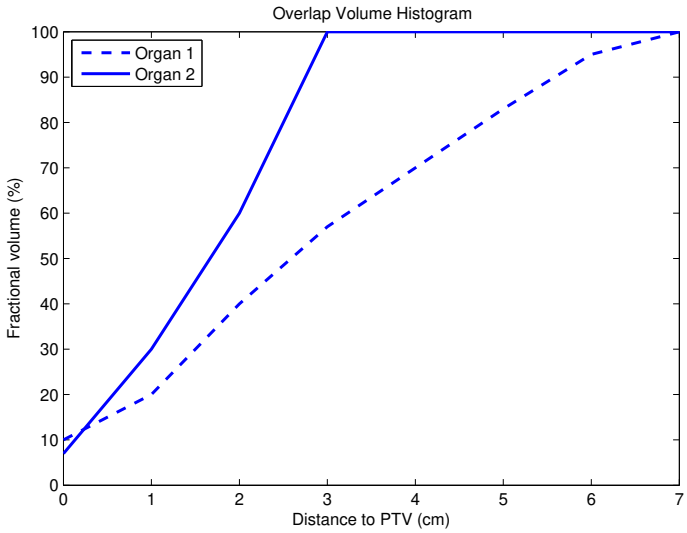

Figure 5.2: OVHs of the two organs shown in Figure 5.1.

In this chapter, we demonstrate the use of patient anatomy for determining objective function weights for treatment planning as a proof of concept. We develop and validate a statistical model that relates patient anatomy derived from the overlap volume of the bladder and rectum with the expanded PTV to the optimal weight values derived from the inverse optimization model. We then apply this model to determine how well it predicts weights for different (out of sample) patients given their anatomy. We measure the objective values and dosimetric properties of the treatment plans generated from these predicted weights compared with the inversely optimized weights for each patient and the weights averaged over all patients.

5.2 A weight prediction model

5.2.1 Inversely optimized weights

We consider the inverse optimization model (IOM) based on the forward optimization formulation (3.1) with four objectives – the bladder, rectum, and left and right femoral heads – described in Section 4.4 (i.e., Model (A)). We applied the IOM to the 24 patient

Table 5.1: Weights for the four objectives determined by the IOM.

Pat.	IOM weights			
	Bladder	Rectum	L.Fem	R.Fem
1	0.256	0.679	0.029	0.036
2	0.616	0.380	0.001	0.003
3	0.118	0.879	0.002	0.001
4	0.975	0.006	0.005	0.014
5	0.848	0.143	0.005	0.005
6	0.197	0.798	0.003	0.004
7	0.618	0.376	0.004	0.002
8	0.200	0.773	0.017	0.010
9	0.845	0.126	0.016	0.012
10	0.939	0.052	0.004	0.004
11	0.324	0.639	0.014	0.017
12	0.088	0.906	0.004	0.002
13	0.248	0.746	0.002	0.004
14	0.007	0.989	0.001	0.003
15	0.573	0.412	0.010	0.005
16	0.964	0.023	0.012	0.002
17	0.973	0.011	0.006	0.011
18	0.297	0.695	0.006	0.002
19	0.672	0.314	0.007	0.006
20	0.980	0.010	0.007	0.003
21	0.815	0.147	0.018	0.020
22	0.431	0.545	0.011	0.012
23	0.926	0.061	0.004	0.008
24	0.892	0.091	0.011	0.006
Average	0.565	0.418	0.009	0.008

datasets and obtained 24 different sets of weights, which we refer to as "IOM weights". Table 5.1 shows the IOM weights for all 24 patients as well as the average weight vector. For the detailed information about the inverse optimization formulation and clinical data, we refer the readers to Chapter 4.

5.2.2 Quantifying patient geometry

For each patient, we generated OVHs for the bladder and rectum at PTV expansion distances from 0 cm to 2 cm. We then computed the ratio of the rectum overlap to bladder overlap at a particular PTV expansion, which we hypothesize is correlated with

the IOM weights in Table 5.1. We denote the overlap volume ratio at PTV expansion distance X cm as "OV_X". Figures 5.3 and 5.4 show scatterplots of the rectum and bladder weights, respectively, versus OV_X at different values of X.

5.2.3 A regression model

Next, we developed a regression model that predicts the rectum weight given an OV_X value. Because of the consistently low IOM femoral head weights (see Table 5.1), we fixed the left and right femoral head weights at 0.01. Due to its strong linear relationship to the corresponding rectum weight (see Figure 5.5), the bladder weight was obtained by subtracting the sum of the rectum weight and two femoral head weights from 1. We assume that 0.01 is the minimum weight for both the bladder and rectum objectives. Thus, the maximum weight for the bladder and rectum objectives is 0.97. An S-shaped logistic function was used for the regression model, ensuring that the predicted rectum weight is between 0.01 and 0.97. The functional form of the regression equation is

$$\alpha_R = 0.01 + \frac{0.96}{1 + e^{(-\theta_1 OV_X + \theta_2)}}, \tag{5.1}$$

where α_R denotes the rectum weight.

Since the OV_X values are correlated in X, we choose to use only a single value of OV_X, $X = 0$, 0.5, 1, 1.5, or 2.0, in the model. For each value of OV_X and each patient, we determined the optimal values of parameters θ_1 and θ_2 using the `nls` function in R with the data from the other 23 patients (R Core Team, 2013). Then, for the given OV_X value and patient, we computed the sum of squared errors between the predicted weights and the weights determined from the IOM. We repeated this computation for all 24 patients and computed the total error. Using this leave-one-out approach, the sums of squared errors for the models with predictors OV_0, $OV_{0.5}$, OV_1, $OV_{1.5}$, and OV_2 were 1.31, 0.57, 0.42, 1.01, and 1.77, respectively. As a result, we chose OV_1 as the independent variable

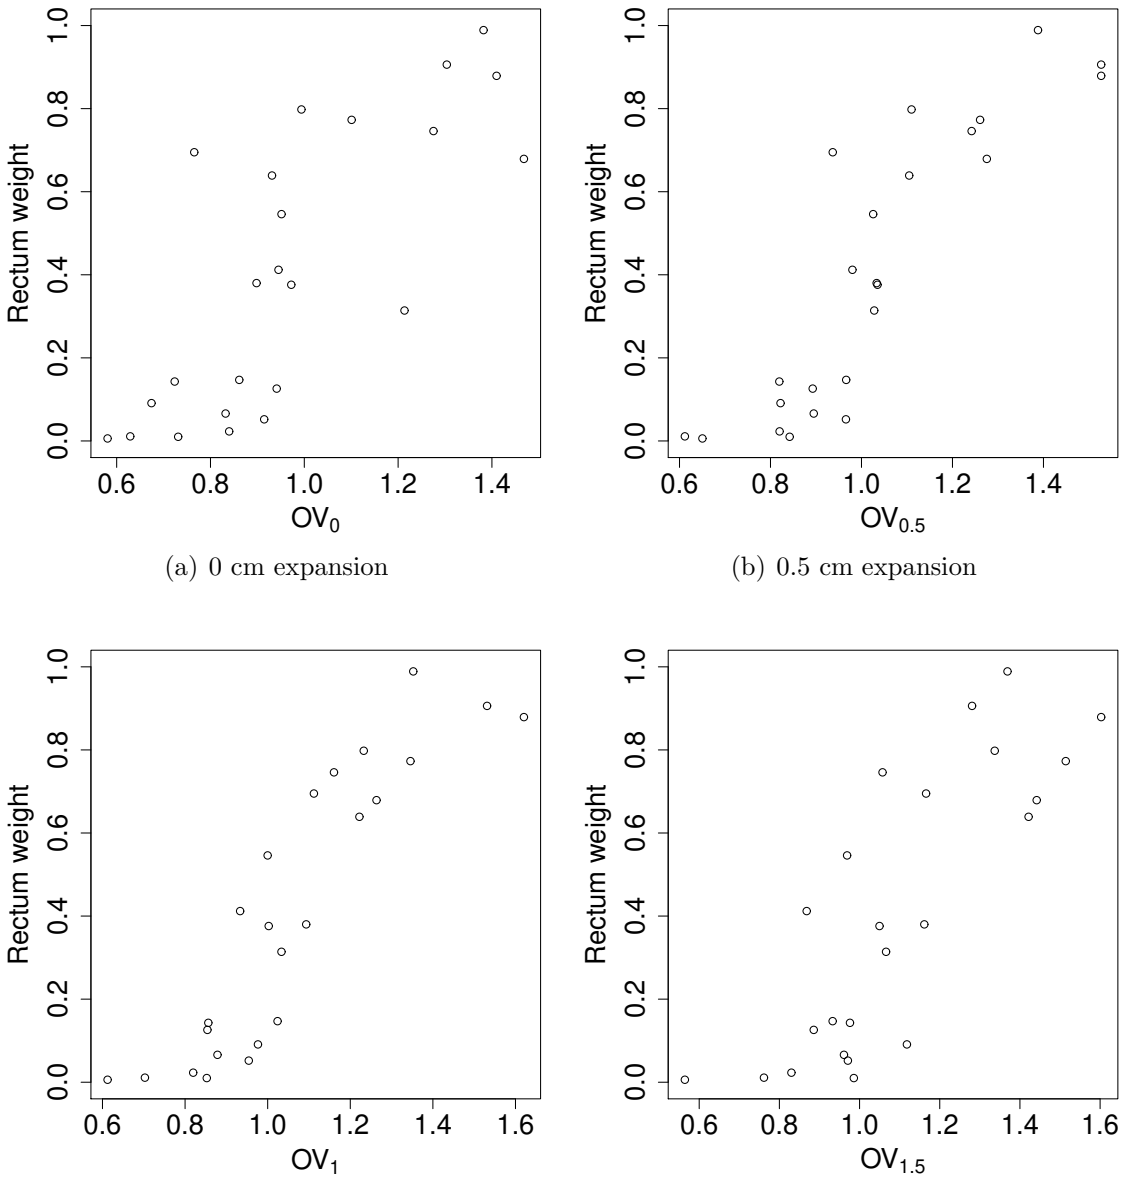

(a) 0 cm expansion

(b) 0.5 cm expansion

(c) 1 cm expansion

(d) 1.5 cm expansion

Figure 5.3: Scatterplots of the rectum weight and rectum-bladder OVH point ratios at different PTV expansion distances.

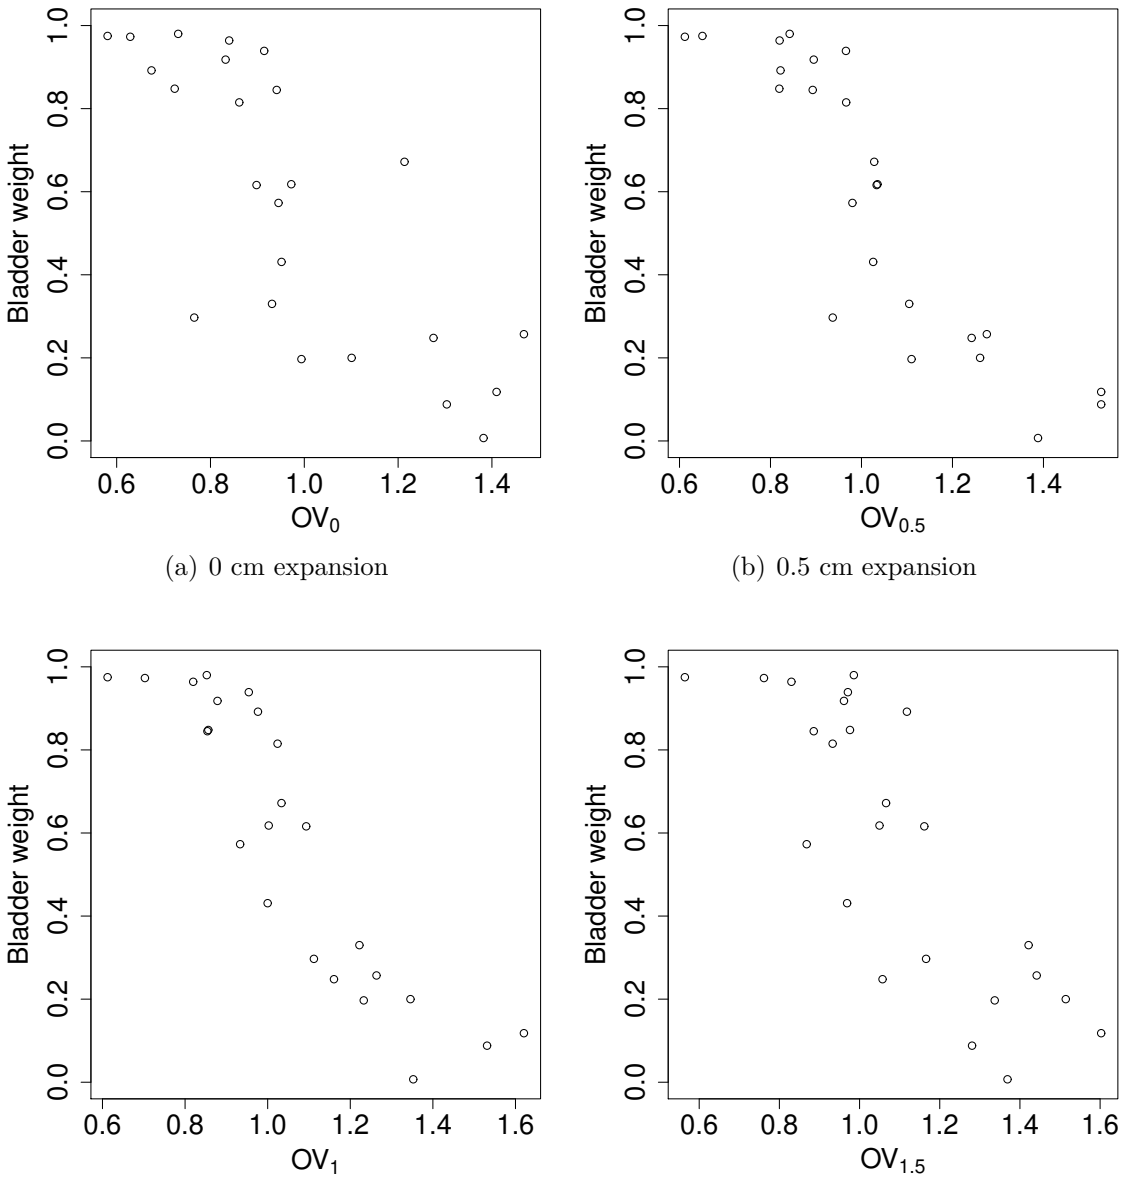

(a) 0 cm expansion

(b) 0.5 cm expansion

(c) 1 cm expansion

(d) 1.5 cm expansion

Figure 5.4: Scatterplots of the bladder weight and rectum-bladder OVH point ratios at different PTV expansion distances.

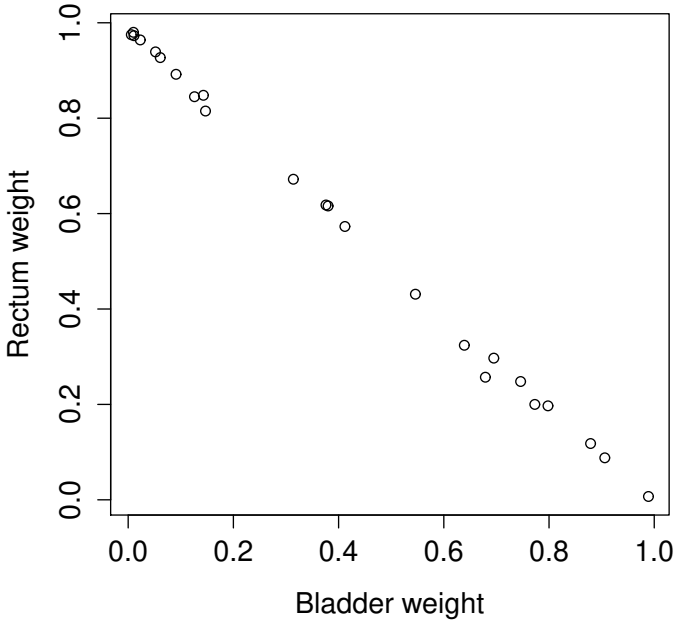

Figure 5.5: Scatterplots of the bladder weight and rectum weight.

for model (5.1). We also tried many models with multiple OV_X values as predictors, but the results were similar to the case with only OV_1.

Once the structure of the model was determined (using OV_1 as the independent variable), the predictive ability of the model was evaluated by comparing treatment plans generated from the IOM weights and treatment plans generated from the predicted weights. Again, we used a leave-one-out approach to generate weights for an out-of-sample patient. For example, for patient #1, we trained the model using the other 23 patients to determine values of θ_1 and θ_2. Then, we used the OV_1 of patient #1 to determine the "predicted weights". For each patient, dose distributions were generated through inverse planning using the treatment plan optimization formulation (3.1) with the individual's IOM weights ("IOM plan"), the individual's predicted weights ("predicted plan") and an average of all 24 IOM weight vectors ("average plan"). We measured objective values, V54.3 Gy and V70.0 Gy for both the bladder and rectum, and V54.3 Gy for the femoral heads. The clinical acceptability criteria were V54.3 Gy $\leq 50\%$ and V70.0 Gy $\leq 30\%$ for

the bladder and rectum, and V54.3 Gy $\leq 5\%$ for the femoral heads.

5.3 Results

5.3.1 Prediction of weights

Because of the leave-one-out method employed, model (5.1) has different parameter values θ_1 and θ_2 for each patient, as shown in Table 5.2. Similar θ_1 and θ_2 values across the patients suggest that model (5.1) using OV_1 is robust across different training sets. Figure 5.6 shows the regression line with patient #1 left out, as an example.

Table 5.2 shows the ℓ_2 distance between the IOM weight vector and the predicted weight vector for all 24 patients, as well as the distance between the IOM weight vector and the average weight vector. The difference between the IOM and predicted weights was smaller for 20 out of the 24 patients, and the average difference was smaller by roughly a factor of six.

5.3.2 Plan comparison: objective values

The IOM, predicted and average treatment plans were compared with respect to the attained objective values. Table 5.3 shows the absolute differences in the rectum and bladder objective values between the IOM and predicted treatment plans as well as between the IOM and average treatment plans for all patients. For both the bladder and rectum, the median difference between the IOM and predicted objective values was smaller than the median difference between the IOM and average objective values. These differences were statistically significant at the 95% level (p-values of 0.003 for both the bladder and rectum) according to a one-sided sign test (Walpole et al., 2006).

Table 5.2: Summary of predicted weights.

Pat.	Regress. params.		OV_1	Predicted weights				Weight comparison	
	θ_1	θ_2		Bladder	Rectum	L.Fem	R.Fem	$\|\|\boldsymbol{\alpha}^I - \boldsymbol{\alpha}^*\|\|_2$	$\|\|\boldsymbol{\alpha}^I - \bar{\boldsymbol{\alpha}}\|\|_2$
1	-10.600	9.720	1.263	0.163	0.817	0.010	0.010	0.029	0.165
2	-10.112	9.261	1.093	0.487	0.493	0.010	0.010	0.030	0.004
3	-10.060	9.166	1.620	0.018	0.962	0.010	0.010	0.017	0.413
4	-9.991	9.100	0.612	0.958	0.022	0.010	0.010	0.001	0.337
5	-10.119	9.214	0.856	0.877	0.103	0.010	0.010	0.003	0.158
6	-9.775	8.880	1.232	0.238	0.742	0.010	0.010	0.005	0.280
7	-10.212	9.276	1.002	0.695	0.285	0.010	0.010	0.014	0.005
8	-10.404	9.505	1.345	0.091	0.889	0.010	0.010	0.025	0.260
9	-10.071	9.172	0.854	0.877	0.103	0.010	0.010	0.002	0.164
10	-9.500	8.685	0.954	0.770	0.230	0.010	0.010	0.067	0.273
11	-10.525	9.639	1.222	0.223	0.757	0.010	0.010	0.025	0.105
12	-10.062	9.168	1.513	0.028	0.952	0.010	0.010	0.006	0.466
13	-9.609	8.685	1.161	0.379	0.601	0.010	0.010	0.038	0.209
14	-9.595	8.710	1.353	0.107	0.873	0.010	0.010	0.023	0.638
15	-10.824	9.823	0.933	0.816	0.164	0.010	0.010	0.121	0.000
16	-9.846	8.970	0.820	0.897	0.083	0.010	0.010	0.008	0.314
17	-9.967	9.078	0.703	0.944	0.036	0.010	0.010	0.002	0.332
18	-9.774	8.815	1.112	0.483	0.497	0.010	0.010	0.074	0.149
19	-9.947	9.075	1.033	0.623	0.357	0.010	0.010	0.004	0.022
20	-9.718	8.856	0.852	0.871	0.109	0.010	0.010	0.022	0.338
21	-9.707	8.907	1.024	0.627	0.353	0.010	0.010	0.078	0.135
22	-10.737	9.702	1.000	0.719	0.261	0.010	0.010	0.164	0.035
23	-9.795	8.926	0.878	0.851	0.129	0.010	0.010	0.010	0.257
24	-9.564	8.750	0.976	0.716	0.264	0.010	0.010	0.061	0.213
							Average	0.035	0.219

$\boldsymbol{\alpha}^I$, $\boldsymbol{\alpha}^*$, and $\bar{\boldsymbol{\alpha}}$ denote the IOM weights, predicted weights, and average IOM weights, respectively.

Table 5.3: The absolute differences in objective values between the IOM, predicted, and average treatment plans.

	Bladder		Rectum	
Pat.	|IOM-Pred|	|IOM-Avg|	|IOM-Pred|	|IOM-Avg|
1	0.27	0.50	0.43	0.20
2	0.15	0.11	0.02	0.01
3	0.85	0.34	0.01	0.31
4	0.03	0.06	0.23	0.64
5	0.09	0.48	0.16	0.24
6	0.05	0.23	0.09	0.33
7	0.07	0.08	0.13	0.04
8	0.18	0.35	0.16	0.29
9	0.15	0.19	0.05	0.61
10	0.09	0.20	0.18	0.29
11	0.01	0.32	0.23	0.05
12	0.58	0.26	0.03	0.13
13	0.11	0.24	0.05	0.33
14	0.44	0.78	0.09	0.27
15	0.06	0.02	0.19	0.01
16	0.06	0.29	0.30	0.50
17	0.03	0.40	0.37	1.28
18	0.09	0.09	0.10	0.24
19	0.05	0.32	0.01	0.16
20	0.08	0.33	0.37	0.68
21	0.13	0.07	0.44	0.61
22	0.01	0.04	0.05	0.06
23	0.09	0.37	0.11	0.47
24	0.07	0.26	0.60	1.04
Median	0.09	0.26	0.14	0.29
Mean	0.16	0.26	0.18	0.37
SD	0.19	0.18	0.16	0.32

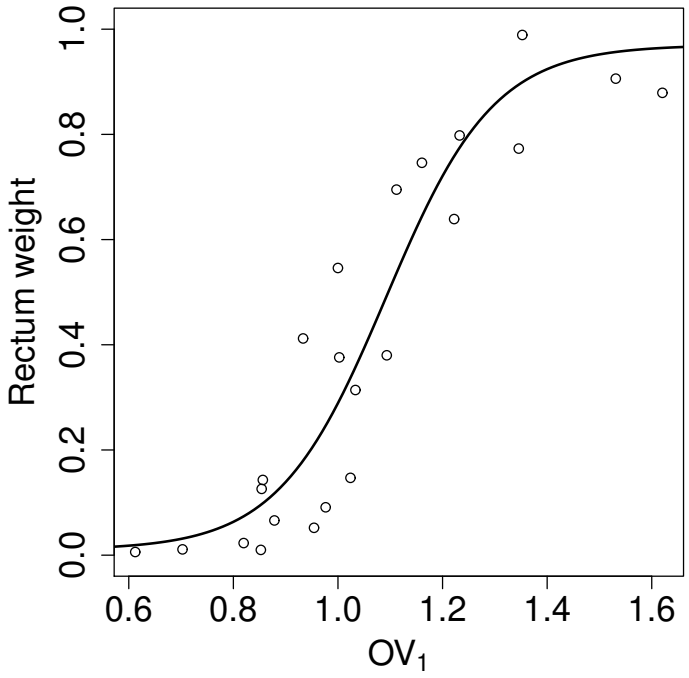

Figure 5.6: Regression line with patient #1 left out.

5.3.3 Plan comparison: OAR sparing

The actual dose distributions resulting from the weights provide a more clinically relevant interpretation of the weights. Figures 5.7 and 5.8 compare Vx Gy values achieved in the IOM, predicted, and average treatment plans. For all patients, the difference between the metrics achieved in the IOM and predicted plans was less than 5 percentage points for both the bladder and rectum (and less than 4 percentage points for all but one patient). The predicted V54.3 Gy values for the femoral heads were also within 5 percentage points of the IOM V54.3 Gy values, except for one patient. V78.0 Gy and V74.1 Gy for the CTV and PTV, respectively, were below 1% for all patients. It can be seen in Figures 5.7 and 5.8 that for most of the patients, the bladder and rectum Vx Gy values for the predicted plan are closer than the average plan to the IOM plan. For both bladder criteria, the absolute difference between the IOM and predicted plans was smaller than the absolute difference between the IOM and average plan for 18 patients. For both rectum criteria, 19 patients exhibited a smaller absolute difference between the IOM and

predicted plans.

5.3.4 Plan comparison: DVHs

To examine the similarity of the overall dose distributions, DVHs for the IOM and predicted treatment plans for each individual patient were compared. Figures 5.9, 5.10, 5.11, and 5.12 show the DVH comparison for patients #1, 3, 4, and 22, respectively. These four patients are representative of various anatomical geometries measured by OVH_1. DVHs for the IOM and predicted treatment plans were also compared to those for other Pareto optimal treatment plans in order to verify their relative proximity with respect to other achievable treatment plans. A thin solid line in Figures 5.9, 5.10, 5.11, and 5.12 corresponds to one Pareto optimal treatment plan. The DVH cloud generated by the thin lines represents a set of Pareto optimal dose distributions for each OAR. We generated the Pareto optimal treatments by solving the forward optimization problem (3.1) with 20 different normalized weight vectors, with components between 0.005 and 0.985. Figures 5.9, 5.10, 5.11, and 5.12 show that the predicted treatment plan is more similar to the IOM plan than most of the other Pareto optimal treatment plans. The agreement of the femoral head DVHs between the predicted and IOM treatment plans was not as strong, but still acceptable in almost all patients. Because the vast majority of the weight was placed on the bladder and rectum objectives for each patient, some fine tuning of the femoral head weights might improve the femoral head DVHs without substantially affecting the bladder and rectum DVHs. For almost all patients, femoral head weights of 0.01 seemed to suffice. Note that the DVH clouds do not represent a complete set of Pareto optimal set treatments; our purpose is simply to visualize the similarity between the IOM and predicted treatment plans compared to other optimized treatment plans.

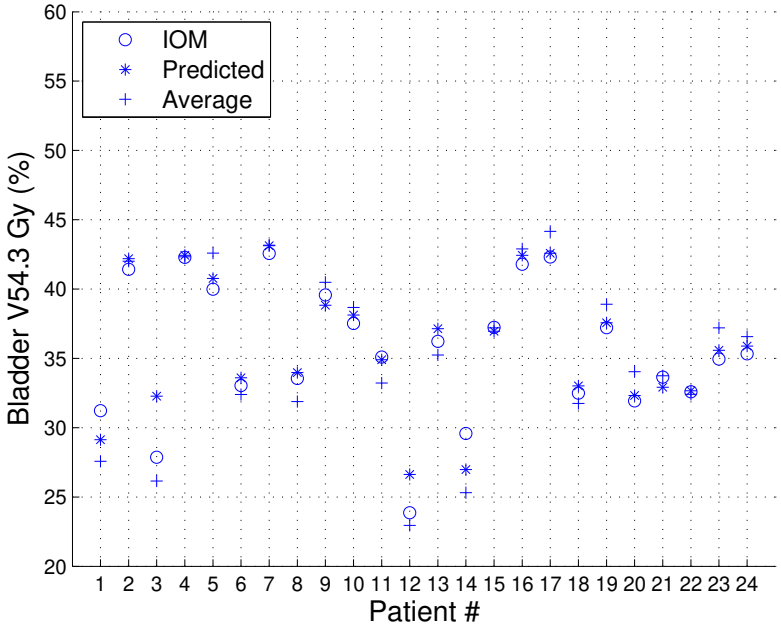

(a) V54.3 Gy for the rectum

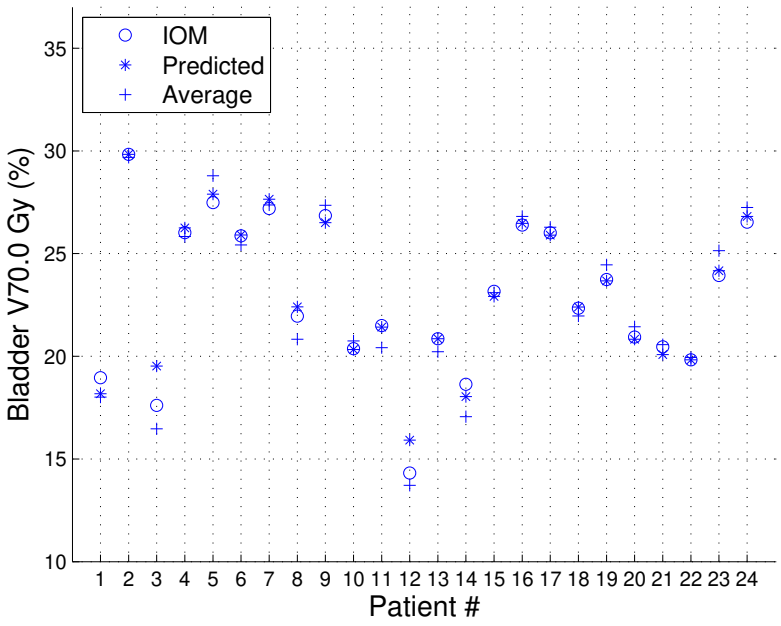

(b) V70.0 Gy for the rectum

Figure 5.7: Comparison of Vx Gy values for the IOM, predicted, and average treatment plans for the bladder.

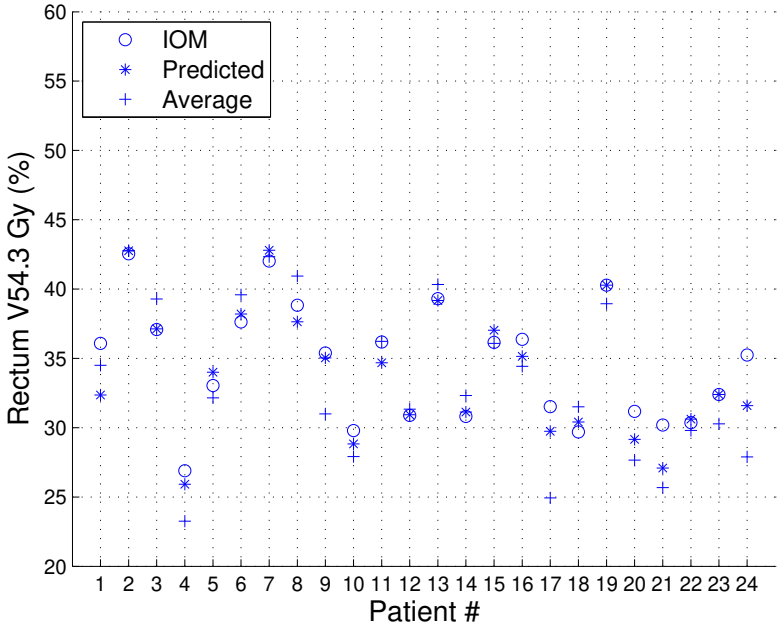

(a) V54.3 Gy for the rectum

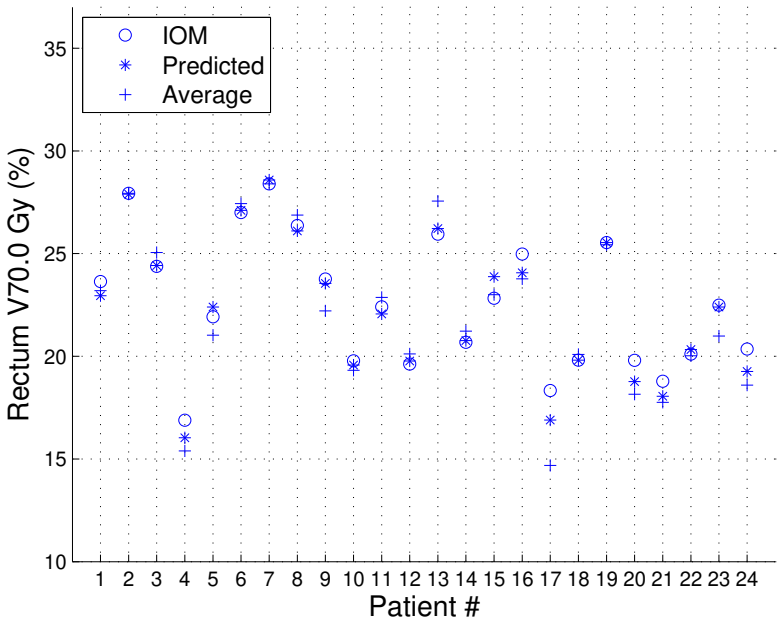

(b) V70.0 Gy for the rectum

Figure 5.8: Comparison of Vx Gy values for the IOM, predicted, and average treatment plans for the rectum.

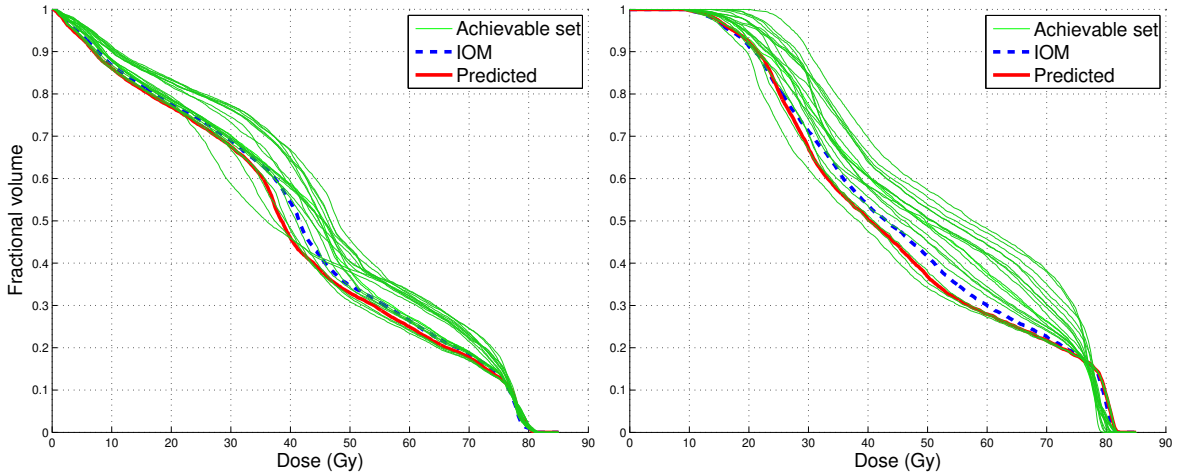

(a) Patient #1: Bladder

(b) Patient #1: Rectum

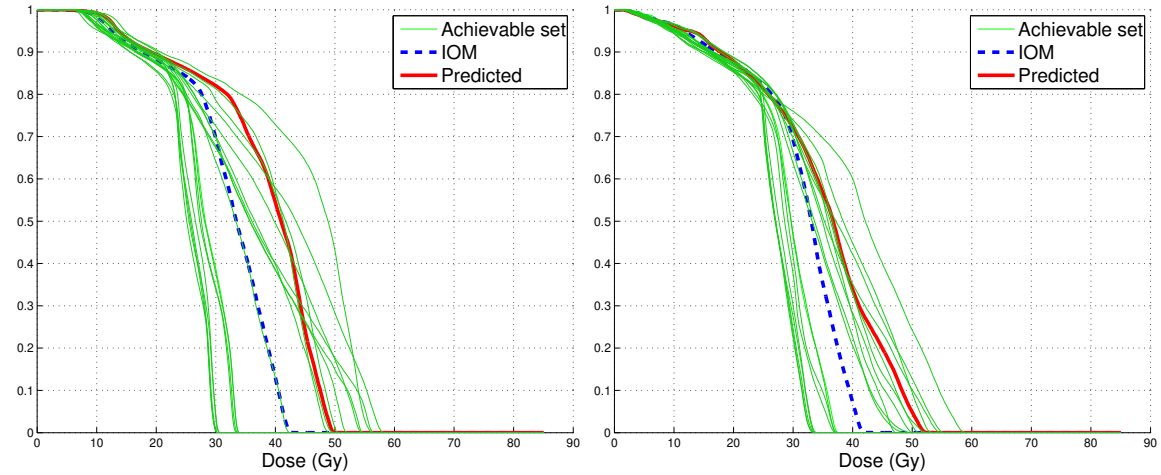

(c) Patient #1: Left femoral head

(d) Patient #1: Right femoral head

Figure 5.9: Comparison of DVHs from the IOM and the predicted weights within an approximate Pareto optimal set: patient #1

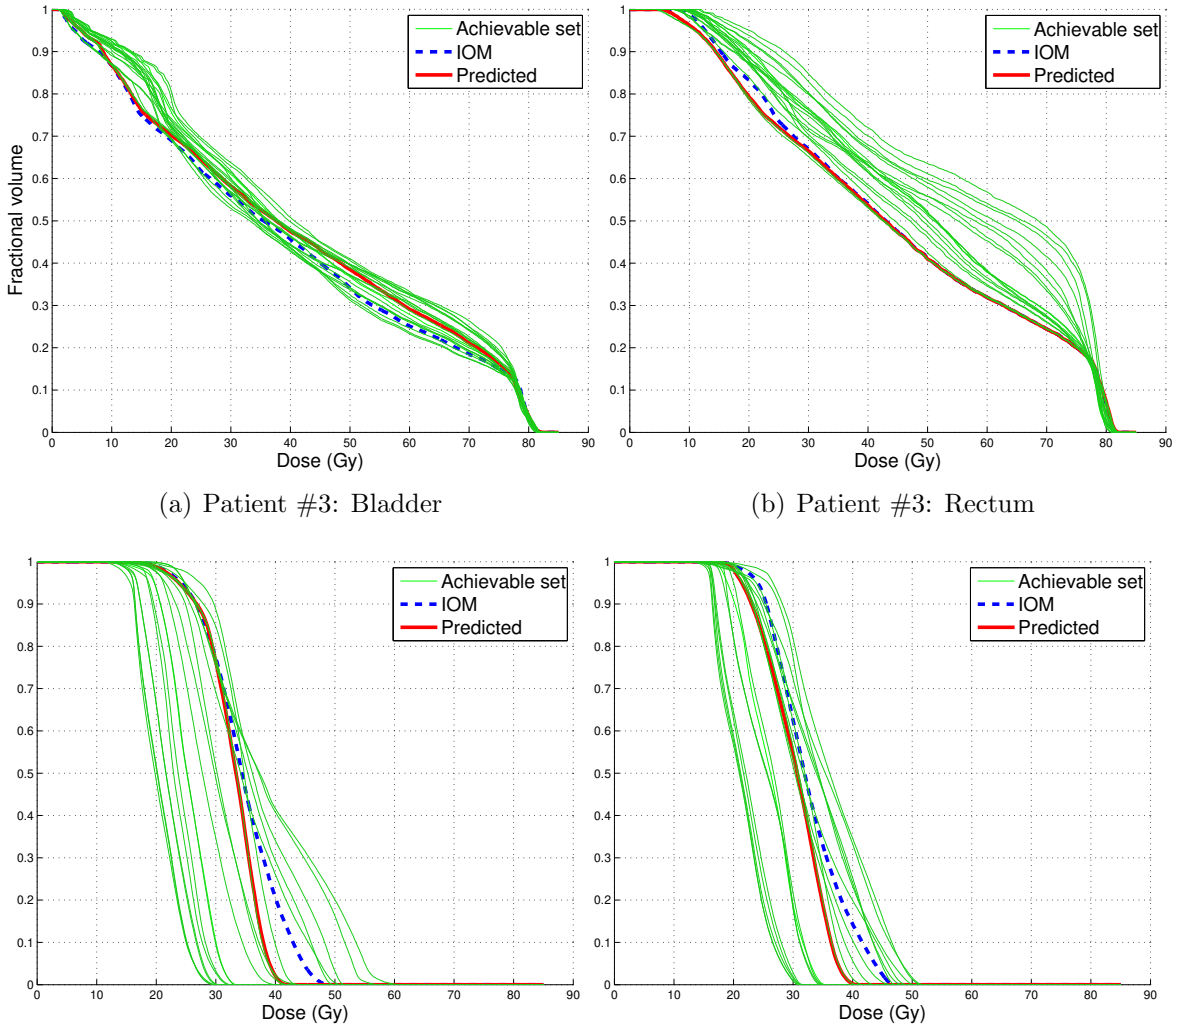

(a) Patient #3: Bladder

(b) Patient #3: Rectum

(c) Patient #3: Left femoral head

(d) Patient #3: Right femoral head

Figure 5.10: Comparison of DVHs from the IOM and the predicted weights within an approximate Pareto optimal set: patient #3

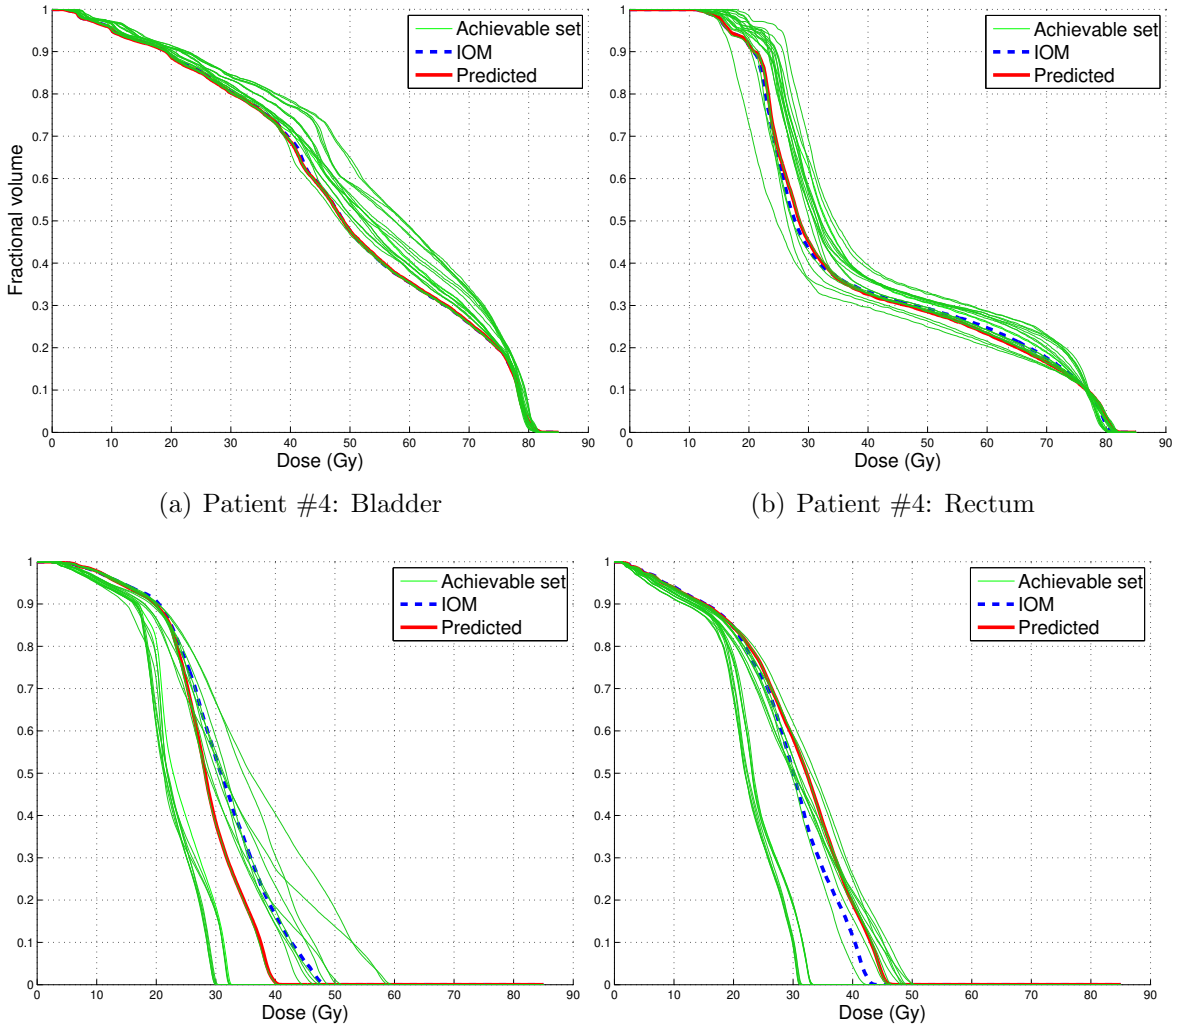

(a) Patient #4: Bladder

(b) Patient #4: Rectum

(c) Patient #4: Left femoral head

(d) Patient #4: Right femoral head

Figure 5.11: Comparison of DVHs from the IOM and the predicted weights within an approximate Pareto optimal set: patient #4

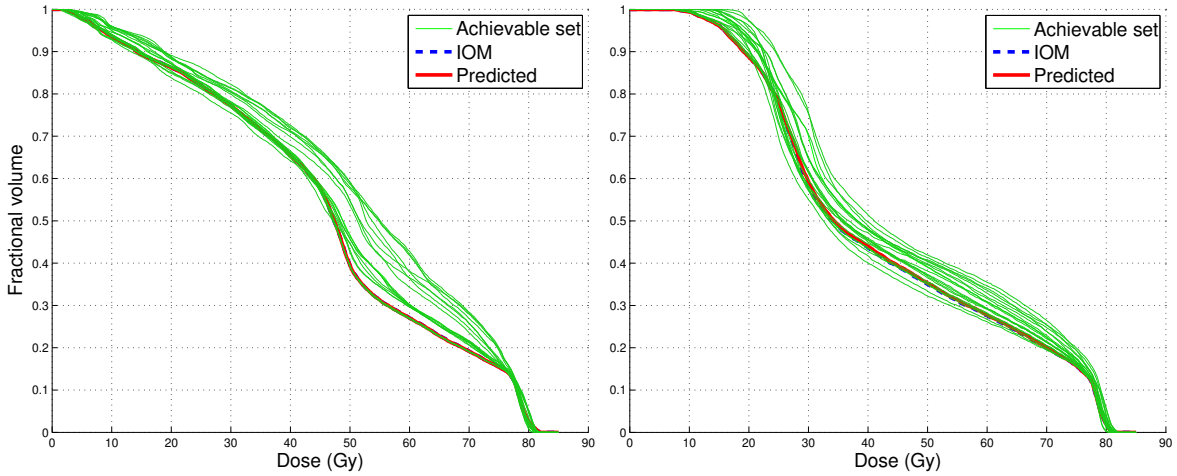

(a) Patient #22: Bladder

(b) Patient #22: Rectum

(c) Patient #22: Left femoral head

(d) Patient #22: Right femoral head

Figure 5.12: Comparison of DVHs from the IOM and the predicted weights within an approximate Pareto optimal set: patient #22

5.4 Discussion

Since our model "learns" from historical treatments plans, the model should acquire more expertise over time as more patient data is used to train the model. Furthermore, treatment plans obtained from the predicted weights will exhibit a trade-off preference that is consistent with historical treatments. Thus, such a prediction model may have value as a starting point for treatment planning by reducing variability in treatment quality and improve standardization of the planning process without sacrificing personalization. Personalization is still highly relevant because only the initial weights are derived from the statistical model, which would be followed by the clinical inverse planning process. Because the objective values and dose metrics achieved by the predicted weights are similar to those achieved by the IOM weights, predicted weights may provide a starting point for *a priori* treatment planning approaches. The similarity in the objective values suggests that predicted weights may also be useful in *a posteriori* approaches that explore the surface of Pareto optimal treatment plans (i.e., the set of non-dominated treatment plans in the objective space) (Craft et al., 2006).

One interpretation of the predicted weights is that they focus on the OAR that is harder to spare. Consider patient #3, who had the highest value of OV_1. One might expect that it would be harder to generate a dose distribution that satisfies the rectum criteria for this patient. Despite such a high predicted weight for the rectum objective, the rectum metrics achieved by the predicted weight for this patient are much higher (i.e., more likely to violate the acceptability limits) than the achieved bladder metrics, whose corresponding weight was very low. Similar observations can be made for other patients with high OV_1 values (e.g., #8, 12, and 14). On the other hand, patient #4 had the smallest OV_1 value and the achieved Vx Gy values for the rectum were very low (similar for other patients with low OV_1 values such as #16 and 17). Thus, one can interpret the prediction model as simultaneously determining how difficult it is to spare a certain OAR based on the patient anatomy and determining the appropriate weight to

accomplish this sparing.

Figure 5.13 shows the highest (i.e., worst) values of Vx Gy between the bladder and rectum for the IOM, predicted, and average treatment plans. For most of the patients (22 patients for V54.3 Gy and 19 patients for V70.0 Gy), the highest value comes from the average treatment plan. In other words, using average weights puts more patients at risk of violating a clinical acceptability criterion. This observation again suggests that the treatment planners who created the original plans had put more emphasis on sparing the "harder to spare" of the two primary OARs, and the predicted weights captured this clinical intention based on each patient's geometry.

If the predicted weights are similar to the IOM weights, then the resulting objective values and dose distributions will of course be similar. When the weights are not similar, however, there are some patients who still have similar objective values and dose distributions. The four patients who had the largest (absolute) deviation between the predicted and IOM rectum weights were patients #10, 15, 21, and 22. However, for these patients, the dose distributions generated from the predicted weights were still similar to those from the IOM weights (see Figures 5.7 and 5.8; also see Figures 5.12(a) – 5.12(d) for patient #22). Notice in Table 5.2 that these four patients all had OV_1 values close to 1 (0.95, 0.93, 1.02, and 1.00, respectively). That is, about the same amount of bladder and rectum overlapped the PTV at a 1 cm expansion. For such cases, it may be that the final dose distribution is less sensitive to the exact weight values. Sensitivity of the dose distribution to weights and patient anatomy is an important topic for future study and likely depends on the tumour site.

The approach described in this chapter has some limitations. The primary limitation is that the regression model currently only applies to the case where one (or two via normalization) weight is to be predicted. This limitation may be acceptable for a site like prostate with minimal anatomic variability, but a site that involves more complex trade-offs like head-and-neck may require a multinominal regression model where multiple

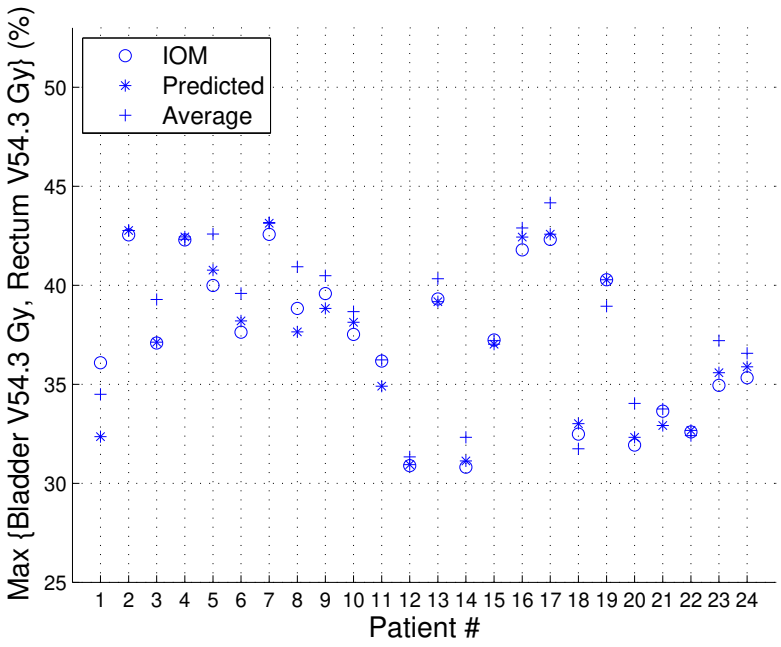

(a) Worst V54.3 Gy

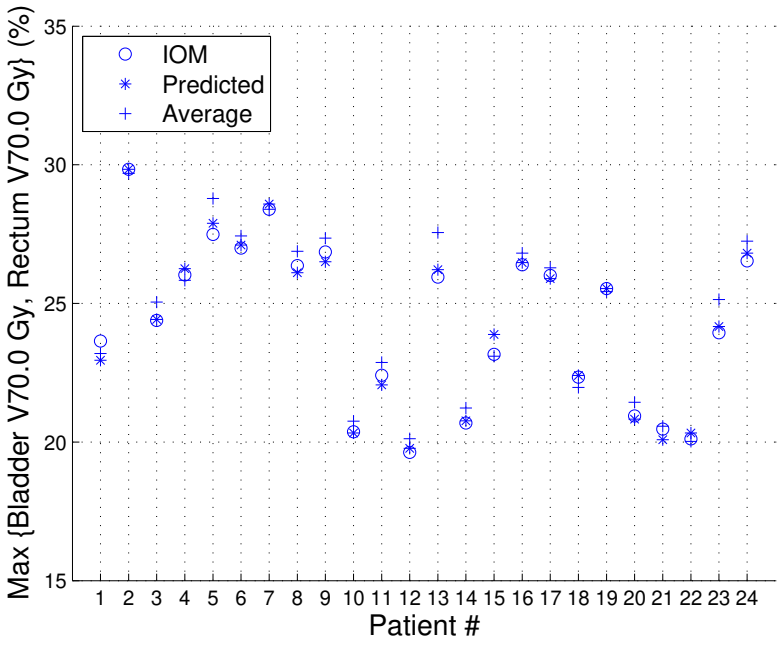

(b) Worst V70.0 Gy

Figure 5.13: Comparison of the worst Vx Gy values for the IOM, predicted, and average treatment plans.

weights are simultaneously predicted. Of course, such a model would require a much larger database of patients to train the model. We also note that the IOM weights would not be the same for different types of objective functions. That is, if we change the functional forms of the objective functions used in the inverse planning model, we may need to re-examine the geometrical factors that relate to the weight values and re-train a new prediction model. Similarly, the weights would also change given different treatment plan geometries (i.e., the number of beams, beam orientations, etc.). Lastly, the overlap volume of OARs with the PTV that we used as a predictor may not be able to detect complicated spatial characteristics of the anatomy. For example, two OARs that have the same overlap volume may be exposed to significantly different amount of radiation due to beam angles, which may lead to significantly different OAR weight values.

Overall, our goal in this chapter is to demonstrate a proof of concept, which we believe this simple model achieves for the prostate site. The results presented should be taken as a starting point for knowledge-base, geometry-driven weight determination in treatment planning, but further study with a larger patient database is needed to address complex sites with many more objectives.

5.5 Conclusion

In this chapter, we developed a statistical model to relate patient anatomy to weight values used in IMRT treatment planning for prostate cancer. In particular, we developed and validated an S-shaped relationship between the overlap volume ratio of the rectum to the bladder with a 1 cm PTV expansion and an optimal weight value of a penalty-based objective function for the rectum. By applying our model to 24 patients, we found that dose distributions generated from weights predicted by the model were similar to those generated by the optimal weights determined from a previously developed inverse optimization method. We found that the predicted weights outperform the average weights

and we believe that the predicted weights better quantify the treatment planner's original intention.

This study provides a proof of concept of geometry-driven weight determination, which may find application in various MCO approaches to treatment planning. The weights could be used as an informed starting point for the treatment planning process, which would hopefully converge more quickly as a result. The weights may also provide information about the most clinically relevant region of the Pareto surface to explore, again possibly reducing the search time to determine a final treatment plan. Overall, our model demonstrates that patient anatomy can be used to predict appropriate weights for treatment planning, which has implications for both standardization and personalization. From the standardization perspective, the model ensures that patients with similar anatomies will start off with similar weight values in the treatment planning process. However, because a full inverse planning process is still carried out with these weight values, the final treatment plan is still personalized for the patient at hand.

Chapter 6

Preference preservation in inverse multiobjective convex optimization

6.1 Introduction

To overcome noisy observations that cause the standard inverse optimization models to be ill-posed, Keshavarz et al. (2011) and Chapter 4 of this thesis proposed generalizations of existing inverse models of Iyengar and Kang (2005) and Ahuja and Orlin (2001), respectively. Keshavarz et al. (2011) imputes the parameters of a convex objective function to make an input solution approximately optimal. Solutions that are candidates to be optimal will be optimal under the imputed parameters, while solutions that are not candidates to be optimal will be suboptimal under the imputed parameters. Instead of enforcing the exact Karush-Kuhn-Tucker (KKT) conditions like Iyengar and Kang (2005), Keshavarz et al. (2011) introduce a vector of "residuals" in the KKT conditions. By minimizing these residuals, the inverse problem determines parameters that make a given solution "minimally" suboptimal. The objective function was parameterized as a weighted linear combination of pre-specified convex functions, with the weights being the parameters being imputed. Thus, this problem can be seen as an inverse multiobjective

optimization problem.

In Chapter 4, we considered inverse multiobjective linear optimization. In the linear case, the KKT residuals amount to a scalar that measures the duality gap induced by a given solution: a zero duality gap for a feasible solution on the boundary of the criterion-space polyhedron and a nonzero duality gap for interior or exterior points. It was shown that solving the residual-based inverse problem with a non-boundary point is equivalent to projecting the non-boundary point on to the boundary and solving the resulting standard inverse problem. The induced projection maintained the relative importance of each objective as encoded by the given solution in the overall weighted objective function, which we refer to as the "preference ordering" in this chapter. In other words, the objective values achieved by the inversely optimized weights were adjusted proportionally (component-wise) from the objective values corresponding to the given solution.

In multiobjective optimization, the decision maker's preferences are often represented by weights in the objective function. A solution that is generated by solving a multiobjective (forward) optimization problem implicitly encodes these preferences. Therefore, when inverse optimization is applied to a multiobjective problem and there are multiple sets of weight vectors that are optimal for the inverse problem, it may be desirable to choose a weight vector that maintains the initial preference ordering as closely as possible. Consider the following example from radiation therapy treatment planning for cancer, which is a multiobjective problem. The variables are the intensities of the beams that deliver radiation while the parameters include the weights of the conflicting objectives (i.e., spare healthy organs, irradiate the tumor, etc.). A historical treatment that was clinically acceptable and delivered to a patient encodes the final trade-offs between the objectives (e.g., organs) that the oncologist is willing to make. Thus, inversely optimized weights derived from such historical treatments should maintain those preferences if they are used to treat similar patients in the future (see Chapter 5).

In this chapter, we bring together the ideas of Keshavarz et al. (2011) and Chap-

ter 4 of this thesis to develop a general inverse optimization model for multiobjective convex optimization that (1) is convex, (2) accommodates any input solution, including one that cannot be optimal for the forward problem, and (3) determines a nonzero weight vector that maintains the preference ordering of the input solution. We propose a linear approximation to the convex model that can approximately preserve the preference ordering, and demonstrate that a successive linear programming algorithm can bridge the exact and approximate methods. We elucidate a connection between our approach and existing inverse models, and propose a general inverse convex programming framework that encompasses many of the inverse models discussed. Lastly, we present computational experiments that demonstrate our models using data from prostate cancer radiation therapy.

6.2 Inverse optimization methodology

We first define a canonical multiobjective convex optimization problem as the forward problem. Then, we briefly review the inverse optimization models from Iyengar and Kang (2005) and Keshavarz et al. (2011).

6.2.1 Preliminaries

Let $f_k : \mathbb{R}^n \to \mathbb{R}, k = 1, \ldots, K$ and $g_l : \mathbb{R}^n \to \mathbb{R}, l = 1, \ldots, L$ be convex functions. Let $\mathbf{x} \in \mathbb{R}^n, \mathbf{A} \in \mathbb{R}^{m \times n}$, and $\mathbf{b} \in \mathbb{R}^m$. We define the forward optimization problem as

$$\begin{aligned} \text{FOP}(\boldsymbol{\alpha}): \quad \underset{\mathbf{x}}{\text{minimize}} \quad & \sum_{k=1}^{K} \alpha_k f_k(\mathbf{x}) \\ \text{subject to} \quad & g_l(\mathbf{x}) \leq 0, \quad l = 1, \ldots, L, \\ & \mathbf{A}\mathbf{x} = \mathbf{b}, \end{aligned} \tag{6.1}$$

where α_k is the weight for the k-th objective function. Let $\mathbf{X}$ be the feasible region of problem (6.1). We assume $\boldsymbol{\alpha} \in \mathbb{R}^K_+ \backslash \{\mathbf{0}\}$, $f_k(\mathbf{x}) > 0, k = 1, \ldots, K$ for $\mathbf{x} \in \mathbf{X}$, and $\mathbf{A}$ has full rank. We also assume that Slater's condition holds (Boyd and Vandenberghe, 2004).

We define $\Omega(\boldsymbol{\alpha})$ to be the set of optimal solutions to FOP($\boldsymbol{\alpha}$) and $\Omega := \bigcup_{\boldsymbol{\alpha} \in \mathbb{R}^K_+ \backslash \{\mathbf{0}\}} \Omega(\boldsymbol{\alpha})$. A solution $\mathbf{x} \in \mathbf{X}$ is *weakly Pareto optimal* if there is no other $\mathbf{y} \in \mathbf{X}$ such that $f_k(\mathbf{y}) < f_k(\mathbf{x})$, for all $k = 1, \ldots, K$. On the other hand, a solution $\mathbf{x} \in \mathbf{X}$ is *strictly Pareto optimal* if there is no other $\mathbf{y} \in \mathbf{X}$ such that $f_k(\mathbf{y}) \leq f_k(\mathbf{x})$, for all $k = 1, \ldots, K$. It is known that for a convex multiobjective optimization problem, the set Ω consists of all (both weakly and strictly) Pareto optimal solutions (Ehrgott, 2005). Throughout this chapter references to "Pareto optimal solutions" include both weak and strict types, unless a specific type is specified. For any $S \subseteq \mathbf{X}$, we write $\mathbf{f}(S) = \{(f_1(\mathbf{x}), \ldots, f_K(\mathbf{x})) \,|\, \mathbf{x} \in S\}$. We denote $\mathbf{f}(\mathbf{X})$ as the feasible region in criterion space and the set $\mathbf{f}(\Omega)$ as the Pareto surface.

6.2.2 Inverse conic optimization by Iyengar and Kang (2005)

We begin by illustrating the approach of Iyengar and Kang (2005) using our forward problem (6.1). Given a solution $\hat{\mathbf{x}} \in \mathbf{X}$, a weight vector that makes $\hat{\mathbf{x}}$ optimal can be found by solving the following problem:

$$\begin{aligned} \underset{\boldsymbol{\alpha},\boldsymbol{\sigma},\boldsymbol{\pi}}{\text{minimize}} \quad & 0 \\ \text{subject to} \quad & \sum_{k=1}^{K} \alpha_k \nabla_{\mathbf{x}} f_k(\hat{\mathbf{x}}) + \sum_{l=1}^{L} \sigma_l \nabla_{\mathbf{x}} g_l(\hat{\mathbf{x}}) - \mathbf{A}'\boldsymbol{\pi} = \mathbf{0}, \\ & \sigma_l g_l(\hat{\mathbf{x}}) = 0, \quad l = 1, \ldots, L, \\ & \boldsymbol{\alpha} \geq \mathbf{0},\ \boldsymbol{\sigma} \geq \mathbf{0}. \end{aligned} \tag{6.2}$$

Constraints in problem (6.2) correspond to the KKT conditions for the forward problem (6.1) with Lagrange multipliers $\boldsymbol{\sigma}$ and $\boldsymbol{\pi}$. If $\boldsymbol{\alpha}^*$ is an optimal solution that arises from solving (6.2), then $\hat{\mathbf{x}} \in \Omega(\boldsymbol{\alpha}^*)$. Although an objective function $||\boldsymbol{\alpha} - \hat{\boldsymbol{\alpha}}||$ was included

in Iyengar and Kang (2005), where $\hat{\boldsymbol{\alpha}}$ is some given weight vector, we omit it here and in subsequent inverse formulations to allow for an objective of minimizing "residuals" to be introduced. Note that an arbitrary $\hat{\mathbf{x}}$ need not be in Ω, in which case formulation (6.2) returns $\boldsymbol{\alpha}^* = \mathbf{0}$ as the unique solution (see Example 6.2.1).

Example 6.2.1 Consider the following bi-objective convex optimization problem:

$$\begin{aligned} \underset{\mathbf{x}}{\text{minimize}} \quad & \alpha_1 f_1(x_1, x_2) + \alpha_2 f_2(x_1, x_2) \\ \text{subject to} \quad & (x_1 - 2)^2 + (x_2 - 2)^2 - 1 \leq 0, \end{aligned} \tag{6.3}$$

where $f_1(x_1, x_2) = 4x_1^2 + x_2^2$ and $f_2(x_1, x_2) = x_1^2 + 4x_2^2$. Constraints of the corresponding inverse problem given an input solution $\hat{\mathbf{x}}$ are

$$\begin{aligned} & (4\alpha_1 + \alpha_2)\hat{x}_1 + (\hat{x}_1 - 2)\sigma = 0, \\ & (\alpha_1 + 4\alpha_2)\hat{x}_2 + (\hat{x}_2 - 2)\sigma = 0, \\ & ((\hat{x}_1 - 2)^2 + (\hat{x}_2 - 2)^2 - 1)\sigma = 0, \\ & \boldsymbol{\alpha} \geq \mathbf{0}, \sigma \geq 0. \end{aligned} \tag{6.4}$$

Consider the following three solutions: $\hat{\mathbf{x}}^a = ((4 - \sqrt{2})/2, (4 - \sqrt{2})/2) \in \Omega$, $\hat{\mathbf{x}}^b = (1.7, 1.3) \in \mathbf{X} \setminus \Omega$, and $\hat{\mathbf{x}}^c = (1, 1) \notin \mathbf{X}$ (see Figure 6.1). For $\hat{\mathbf{x}}^a$, constraints (6.4) become

$$\begin{aligned} & (4 - \sqrt{2})(4\alpha_1 + \alpha_2) - \sqrt{2}\sigma = 0, \\ & (4 - \sqrt{2})(\alpha_1 + 4\alpha_2) - \sqrt{2}\sigma = 0, \\ & \boldsymbol{\alpha} \geq \mathbf{0}, \sigma \geq 0, \end{aligned} \tag{6.5}$$

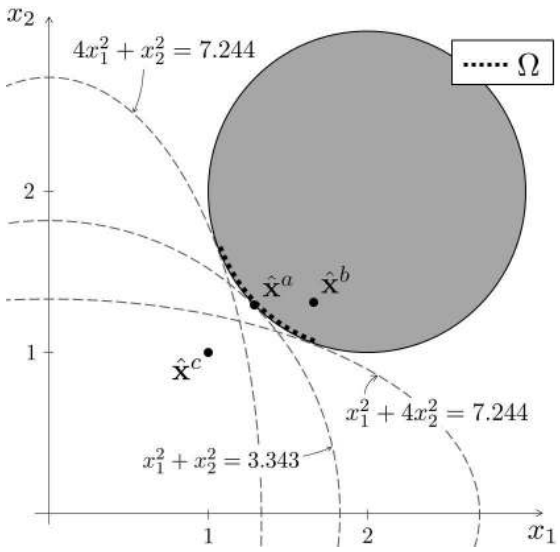

Figure 6.1: Illustration of Example 6.2.1

which hold when $\alpha_1 = \alpha_2 \geq 0$. For $\hat{\mathbf{x}}^b$, constraints (6.4) become

$$\begin{aligned} &1.7(4\alpha_1 + \alpha_2) - 0.3\sigma = 0, \\ &1.3(\alpha_1 + 4\alpha_2) - 0.7\sigma = 0, \\ &\boldsymbol{\alpha} \geq \mathbf{0}, \sigma = 0, \end{aligned} \tag{6.6}$$

which are satisfied only when $\alpha_1 = \alpha_2 = 0$. For $\hat{\mathbf{x}}^c$, constraints (6.4) become

$$\begin{aligned} &(4\alpha_1 + \alpha_2) - \sigma = 0, \\ &(\alpha_1 + 4\alpha_2) - \sigma = 0, \\ &\boldsymbol{\alpha} \geq \mathbf{0}, \sigma = 0, \end{aligned} \tag{6.7}$$

which again are satisfied only when $\alpha_1 = \alpha_2 = 0$.

6.2.3 Objective function imputation by Keshavarz et al. (2011)

Keshavarz et al. (2011) explicitly considered noise or modeling errors in imputing an objective function from a convex model so that the observed data from an unknown, complex system is as consistent as possible with the proposed model. They relaxed the

KKT conditions by allowing positive "residuals," which are then minimized by solving the following problem:

$$
\begin{aligned}
\text{KES}(\hat{\mathbf{x}}): \quad \underset{\boldsymbol{\alpha},\boldsymbol{\sigma},\boldsymbol{\pi},\boldsymbol{\delta},\boldsymbol{\gamma},\boldsymbol{\rho}}{\text{minimize}} \quad & \phi(\boldsymbol{\delta},\boldsymbol{\gamma},\boldsymbol{\rho}) \\
\text{subject to} \quad & \sum_{k=1}^{K} \alpha_k \nabla_{\mathbf{x}} f_k(\hat{\mathbf{x}}) + \sum_{l=1}^{L} \sigma_l \nabla_{\mathbf{x}} g_l(\hat{\mathbf{x}}) - \mathbf{A}'\boldsymbol{\pi} = \boldsymbol{\delta}, \\
& \sigma_l g_l(\hat{\mathbf{x}}) = \gamma_l, \quad l = 1,\ldots,L, \\
& \pi_j(\mathbf{a}_j'\hat{\mathbf{x}} - b_j) = \rho_j, \quad j = 1,\ldots,m, \\
& \alpha_1 = 1, \\
& \boldsymbol{\alpha} \geq \mathbf{0},\ \boldsymbol{\sigma} \geq \mathbf{0},
\end{aligned}
\tag{6.8}
$$

where $\mathbf{a}_j'$ denotes the j-th row of $\mathbf{A}$ and $\phi(\boldsymbol{\delta},\boldsymbol{\gamma},\boldsymbol{\rho}) = 0$ if and only if $\boldsymbol{\delta} = \mathbf{0}$, $\boldsymbol{\gamma} = \mathbf{0}$, and $\boldsymbol{\rho} = \mathbf{0}$ (e.g., $\phi(\boldsymbol{\delta},\boldsymbol{\gamma},\boldsymbol{\rho}) = ||\boldsymbol{\delta}||_2^2 + ||\boldsymbol{\gamma}||_2^2 + ||\boldsymbol{\rho}||_2^2$). The fourth constraint ensures that $\boldsymbol{\alpha} = \mathbf{0}$ is not a feasible solution and serves to implicitly normalize the resulting weight vector. The original version of formulation (6.8) in Keshavarz et al. (2011) accommodated multiple input data points each with their own residuals, but we simply illustrate their method with a single data point $\hat{\mathbf{x}}$. The extension to multiple input data points is straightforward. The original model also excluded the $\boldsymbol{\rho}$ residual as the focus was on $\hat{\mathbf{x}} \in \mathbf{X}$, but we include it here to emphasize the applicability of the model to points $\hat{\mathbf{x}} \notin \mathbf{X}$. As noted in their paper, the choice of ϕ impacts the solution. However, how the weights are normalized can also have a large impact, as we demonstrate later.

In multiobjective optimization applications, a solution $\hat{\mathbf{x}} \in \Omega$ implicitly encodes a preference ordering among the objectives, namely the weight vector that generated it. Thus, for $\hat{\mathbf{x}} \notin \Omega$ it may be desirable for the inverse formulation to derive weights that not only are as consistent as possible with the observation $\hat{\mathbf{x}}$ (i.e., minimized residuals), but also preserve the preference ordering as much as possible (i.e., a particular projection on to Ω).

Example 6.2.2 We revisit Example 6.2.1 and apply formulation (6.8) with $\hat{\mathbf{x}}^b$:

$$\begin{aligned} \underset{\boldsymbol{\alpha},\sigma,\boldsymbol{\delta},\gamma}{\text{minimize}} \quad & \delta_1^2 + \delta_2^2 + \gamma^2 \\ \text{subject to} \quad & 1.7(4\alpha_1 + \alpha_2) - 0.3\sigma = \delta_1, \\ & 1.3(\alpha_1 + 4\alpha_2) - 0.7\sigma = \delta_2, \\ & -0.42\sigma = \gamma, \\ & \alpha_1 = 1, \\ & \boldsymbol{\alpha} \geq \mathbf{0},\ \sigma \geq 0. \end{aligned} \tag{6.9}$$

The optimal solution is $(\alpha_1^*, \alpha_2^*, \sigma^*, \delta_1^*, \delta_2^*, \gamma^*) = (1, 0, 0, 6.8, 1.3, 0)$. The corresponding FOP($\boldsymbol{\alpha}^*$) generates the optimal solution $\mathbf{x}^* = (1.067, 1.641)$. Comparing $\mathbf{f}(\hat{\mathbf{x}}^b) = (13.250, 9.650)$ to $\mathbf{f}(\mathbf{x}^*) = (7.244, 11.910)$ we see that preferences are not maintained: one objective value decreases while the other increases. Furthermore, if the normalization constraint is changed to $\alpha_2 = 1$, then the optimal weight vector changes to $\boldsymbol{\alpha}^* = (0, 1)$ and the new optimal solution is $\mathbf{x}^* = (1.641, 1.067)$ with objective vector $\mathbf{f}(\mathbf{x}^*) = (11.910, 7.244)$, which is completely opposite to the previous case.

6.2.4 A preference-preserving inverse optimization model

Our goal is to formulate a convex inverse model that determines a weight vector $\boldsymbol{\alpha}^*$ and corresponding optimal solution $\mathbf{x}^*$ to FOP($\boldsymbol{\alpha}^*$) that preserve the preference ordering encoded by a given $\hat{\mathbf{x}}$. We begin with a definition to formalize our idea of preference preservation.

Definition 6.2.1 (Perfect relative preference preservation) *A solution* $\mathbf{x}$ *perfectly preserves the relative preference ordering encoded by* $\hat{\mathbf{x}}$ *if for some* $\epsilon > 0$, $f_k(\mathbf{x}) = \epsilon f_k(\hat{\mathbf{x}})$ *for all* $k = 1, \ldots, K$. *We will also use this definition to describe* $\boldsymbol{\alpha}$, *if a preference-preserving solution* $\mathbf{x}$ *is in the set* $\Omega(\boldsymbol{\alpha})$.

Definition 6.2.1 says that the objective values achieved by a preference-preserving $\mathbf{x}$ are adjusted component-wise by the same relative amount from those achieved by $\hat{\mathbf{x}}$. Geometrically, this definition means that $\mathbf{f}(\mathbf{x})$ lies on the line joining $\mathbf{f}(\hat{\mathbf{x}})$ and the origin. Since our focus is on identifying preference-preserving Pareto optimal solutions, we are interested in identifying where the line joining $\mathbf{f}(\hat{\mathbf{x}})$ and the origin intersects the Pareto surface, if at all.

The following formulation, which is called the inverse optimization problem (IOP), provides a necessary and sufficient condition for whether there exists a preference-preserving $\mathbf{x}^* \in \Omega$:

$$\begin{aligned} \text{IOP}(\hat{\mathbf{x}}): \quad \underset{\epsilon,\mathbf{x}}{\text{minimize}} \quad & \epsilon \\ \text{subject to} \quad & \epsilon f_k(\hat{\mathbf{x}}) \geq f_k(\mathbf{x}), \quad k = 1, \ldots, K, \\ & g_l(\mathbf{x}) \leq 0, \quad l = 1, \ldots, L, \\ & \mathbf{A}\mathbf{x} = \mathbf{b}. \end{aligned} \tag{6.10}$$

Theorem 6.2.1 *There exists a solution in Ω that perfectly preserves the preference ordering encoded by $\hat{\mathbf{x}}$ if and only if there exists a solution $(\epsilon^*, \mathbf{x}^*)$ that is optimal for* IOP($\hat{\mathbf{x}}$) *with the first K constraints being tight.*

Proof: The part of the result related to preference preservation follows directly from the definition. So the proof focuses on the optimality portion.

($\Leftarrow$) Let $(\epsilon^*, \mathbf{x}^*)$ be an optimal solution to IOP($\hat{\mathbf{x}}$), which implies $\mathbf{x}^* \in \mathbf{X}$. All that remains is to prove $\mathbf{x}^* \in \Omega$. Let $\boldsymbol{\alpha} \geq \mathbf{0}$, $\boldsymbol{\sigma} \geq \mathbf{0}$, and $\boldsymbol{\pi}$ be the Lagrange multipliers associated with the first, second, and third sets of constraints of problem (6.10), respectively. Consider the Lagrangian associated with problem (6.10):

$$L(\boldsymbol{\alpha}, \boldsymbol{\sigma}, \boldsymbol{\pi}, \epsilon, \mathbf{x}) = \epsilon + \sum_{k=1}^{K} \alpha_k (f_k(\mathbf{x}) - \epsilon f_k(\hat{\mathbf{x}})) + \sum_{l=1}^{L} \sigma_l g_l(\mathbf{x}) - \boldsymbol{\pi}'(\mathbf{A}\mathbf{x} - \mathbf{b}). \tag{6.11}$$

A solution $(\epsilon^*, \mathbf{x}^*)$ is optimal for problem (6.10) if and only if there exists $(\boldsymbol{\alpha}, \boldsymbol{\sigma}, \boldsymbol{\pi}) \in$

$\mathbb{R}^K_+ \times \mathbb{R}^L_+ \times \mathbb{R}^m$ that satisfies the following system of equations:

$$\begin{aligned}
&1 - \sum_{k=1}^{K} \alpha_k f_k(\hat{\mathbf{x}}) = 0, \quad (\nabla_\epsilon L(\boldsymbol{\alpha}, \boldsymbol{\sigma}, \boldsymbol{\pi}, \epsilon^*, \mathbf{x}^*) = 0) \\
&\sum_{k=1}^{K} \alpha_k \nabla_{\mathbf{x}} f_k(\mathbf{x}^*) + \sum_{l=1}^{L} \sigma_l \nabla_{\mathbf{x}} g_l(\mathbf{x}^*) - \mathbf{A}'\boldsymbol{\pi} = \mathbf{0}, \quad (\nabla_{\mathbf{x}} L(\boldsymbol{\alpha}, \boldsymbol{\sigma}, \boldsymbol{\pi}, \epsilon^*, \mathbf{x}^*) = 0) \\
&\alpha_k(\epsilon^* f_k(\hat{\mathbf{x}}) - f_k(\mathbf{x}^*)) = 0, \quad k = 1, \ldots, K, \\
&\sigma_l g_l(\mathbf{x}^*) = 0, \quad l = 1, \ldots, L.
\end{aligned} \tag{6.12}$$

The second and fourth sets of equations in (6.12) are equivalent to the KKT conditions for the original forward optimization problem (6.1). Thus, if a solution $\mathbf{x}^*$ is optimal for problem (6.10), $\mathbf{x}^* \in \Omega$.

($\Rightarrow$) Let $(\epsilon^*, \mathbf{x}^*)$ satisfy $\epsilon^* \mathbf{f}(\hat{\mathbf{x}}) = \mathbf{f}(\mathbf{x}^*)$ and $\mathbf{x}^* \in \Omega$. The KKT conditions for the forward problem imply there exists $(\boldsymbol{\alpha}^*, \boldsymbol{\sigma}^*, \boldsymbol{\pi}^*)$ that satisfies the second and fourth set of equations in (6.12). What remains is to show that the first and third set of equations in (6.12) are also satisfied. The third set of equations is satisfied since $\mathbf{x}^*$ is preference-preserving. The first equation can be satisfied through a re-scaling of $\boldsymbol{\alpha}^*$, which is possible since $\boldsymbol{\alpha}^*$ cannot be identically zero (by complementary slackness). Thus, $(\epsilon^*, \mathbf{x}^*)$ is optimal for (6.10). □

The previous result implies that an optimal $\mathbf{x}^*$ to (6.10) is also optimal for the forward problem. The next result identifies the optimal weight vector that $\mathbf{x}^*$ is optimal with respect to. The proof is omitted because it follows directly from the proof of Theorem 6.2.1.

Corollary 6.2.1 *Let $\boldsymbol{\alpha}^*$ be a vector of optimal Lagrange multipliers associated with the first set of constraints in formulation (6.10). An optimal solution $\mathbf{x}^*$ to (6.10) satisfies $\mathbf{x}^* \in \Omega(\boldsymbol{\alpha}^*)$.*

Thus, solving IOP($\hat{\mathbf{x}}$) can simultaneously identify a preference-preserving $\mathbf{x}^* \in \Omega$ as well as a corresponding optimal weight vector $\boldsymbol{\alpha}^*$. Note that the complementary slackness

conditions associated with problem (6.10) provide us with a simple and intuitive condition for when preferences are perfectly preserved.

Corollary 6.2.2 *Let* $\mathbf{x}^*$ *be an optimal solution to* IOP($\hat{\mathbf{x}}$) *and* $\boldsymbol{\alpha}^*$ *be a vector of optimal Lagrange multipliers associated with the first set of constraints in formulation* (6.10)*. If* $\alpha_k^* > 0$ *for all* $k = 1, \ldots, K$*, then* $\mathbf{x}^*$ *perfectly preserves the relative preference ordering encoded by* $\hat{\mathbf{x}}$*.*

It is also well known from multiobjective optimization that if $\boldsymbol{\alpha}^* > \mathbf{0}$, then solving FOP($\boldsymbol{\alpha}^*$) generates a strictly Pareto optimal solution (Ehrgott, 2005). Thus, when $\boldsymbol{\alpha}^* > \mathbf{0}$, we are guaranteed a preference-preserving, strictly Pareto optimal solution. The following example illustrates that if there is at least one k such that $\alpha_k^* = 0$, then there may or may not be a solution in Ω that can achieve perfect preference preservation.

Example 6.2.3 Consider Example 6.2.1 again and two new solutions $\hat{\mathbf{x}}^d = (1.725, 1.121)$ and $\hat{\mathbf{x}}^e = (1.789, 1.096)$. Figure 6.2 shows the feasible region and set of Pareto optimal solutions in criterion space (i.e., $\mathbf{f}(\mathbf{X})$ and $\mathbf{f}(\Omega)$, respectively), along with the two points $\mathbf{f}(\hat{\mathbf{x}}^d) = (13.160, 8.004)$ and $\mathbf{f}(\hat{\mathbf{x}}^e) = (14.000, 8.004)$. Solving IOP($\hat{\mathbf{x}}^d$) and IOP($\hat{\mathbf{x}}^e$) both return the same optimal weight vector $\boldsymbol{\alpha}^* = (0, 1)$ and the same optimal solution $\mathbf{x}^* = \bar{\mathbf{x}}^d$. It is clear from Figure 6.2 that $\bar{\mathbf{x}}^d$ perfectly preserves the preferences encoded by $\hat{\mathbf{x}}^d$ but not $\hat{\mathbf{x}}^e$. Because $\mathbf{f}(\bar{\mathbf{x}}^e) \notin \mathbf{f}(\Omega)$, there is no non-zero weight vector that makes $\bar{\mathbf{x}}^e$ optimal to the forward problem.

Example 6.2.3 illustrates the intuitive fact that when a given $\hat{\mathbf{x}}$ is sufficiently "inferior" with respect to a particular objective k, the optimal weight vector determined from inverse optimization will return $\alpha_k^* = 0$. In this case preferences will not be perfectly preserved. However, preferences will still be preserved for the objectives with positive weight.

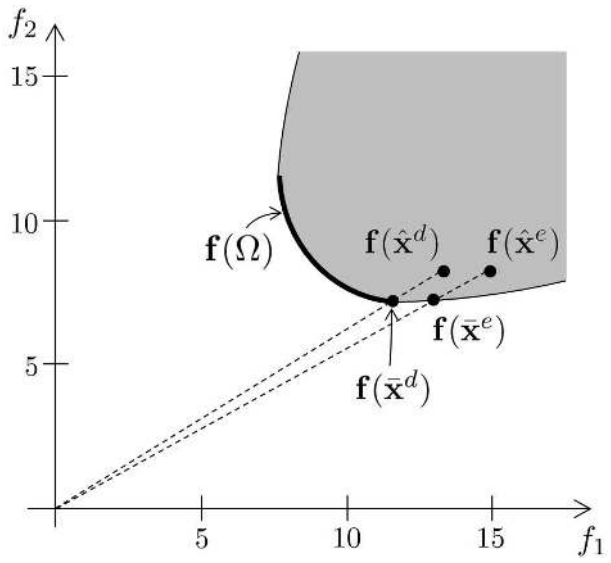

Figure 6.2: Illustration of Example 6.2.3

Instead of maintaining a relative preference ordering among the objective values, an absolute preference ordering can be maintained through the following formulation:

$$\begin{aligned} \underset{\epsilon,\mathbf{x}}{\text{minimize}} \quad & \epsilon \\ \text{subject to} \quad & \epsilon \geq f_k(\mathbf{x}) - f_k(\hat{\mathbf{x}}), \quad k = 1, \ldots, K, \\ & g_l(\mathbf{x}) \leq 0, \quad l = 1, \ldots, L, \\ & \mathbf{A}\mathbf{x} = \mathbf{b}. \end{aligned} \tag{6.13}$$

Formulation (6.13) aims to maintain the absolute difference between $f_k(\hat{\mathbf{x}})$ and $f_k(\mathbf{x}^*)$ for all $k = 1, \ldots, K$. Analogous results to Theorem 6.2.1 and Corollaries 6.2.1 and 6.2.2 hold for formulation (6.13). We focus on the relative case (i.e., formulation (6.10)) in the remainder of this chapter.

Next, we provide a linear approximation to formulation (6.10) that serves as the foundation for a specialized solution methodology. The linear approximation also helps elucidate connections with other models as we will discuss in the next section. The linearized inverse optimization problem (LIOP) can be formulated by linearizing the first

two sets of constraints of (6.10) around some point $\tilde{\mathbf{x}}$:

$$\begin{aligned} \text{LIOP}(\hat{\mathbf{x}}, \tilde{\mathbf{x}}) : \underset{\epsilon, \mathbf{x}}{\text{minimize}} \quad & \epsilon \\ \text{subject to} \quad & \epsilon f_k(\hat{\mathbf{x}}) \geq f_k(\tilde{\mathbf{x}}) + \nabla_{\mathbf{x}} f_k(\tilde{\mathbf{x}})'(\mathbf{x} - \tilde{\mathbf{x}}), \quad k = 1, \ldots, K, \\ & g_l(\tilde{\mathbf{x}}) + \nabla_{\mathbf{x}} g_l(\tilde{\mathbf{x}})'(\mathbf{x} - \tilde{\mathbf{x}}) \leq 0, \quad l = 1, \ldots, L, \\ & \mathbf{A}\mathbf{x} = \mathbf{b}. \end{aligned} \tag{6.14}$$

Since f_k and g_l are convex, it is straightforward to see that formulation (6.14) forms an outer approximation – and thus provides a lower bound – to formulation (6.10). If $\tilde{\mathbf{x}} = \mathbf{x}^*$, where $\mathbf{x}^*$ is an optimal solution to problem (6.10), then an optimal dual solution $\boldsymbol{\alpha}^*$ to (6.14) is also an optimal dual solution to (6.10). Unfortunately, $\mathbf{x}^*$ is not known *a priori*. One option is simply to solve (6.14) with $\tilde{\mathbf{x}} = \hat{\mathbf{x}}$ (i.e., solve LIOP($\hat{\mathbf{x}}$,$\hat{\mathbf{x}}$); from here on out we refer to (6.14) with $\tilde{\mathbf{x}} = \hat{\mathbf{x}}$ as LIOP($\hat{\mathbf{x}}$)). Note that linearization may render problem (6.14) unbounded. In this case, a trust-region $\mathbf{x} \in [\hat{\mathbf{x}} - \kappa\mathbf{e}, \hat{\mathbf{x}} + \kappa\mathbf{e}]$ for some κ can be added (Bazaraa et al., 2006).

Building on this idea, since (6.14) is a first-order approximation of (6.10), we can employ a successive linear programming (SLP) algorithm (Zhang et al., 1985; Bazaraa et al., 2006) to solve formulation (6.10). This algorithm repeatedly solves the linear problem (6.14) using a trust region approach, generating an optimal solution $\mathbf{x}_i^*$ in iteration i that is used as the input vector $\tilde{\mathbf{x}}$ in iteration $i + 1$. In Section 6.4, we implement the SLP algorithm from Zhang et al. (1985). We refer the reader to that paper for details of the algorithm and a proof of convergence.

We focus on the linearization of problem (6.10) because of the connection to the model of Keshavarz et al. (2011) it provides, and also because the linearized problem and SLP algorithm work well for our quadratic numerical examples in Section 6.4. However, we note that for general convex problems, successive quadratic programming may be a better choice (Bazaraa et al., 2006).

6.3 Relationships between inverse optimization models

The linearization described in the previous subsection provides a nice bridge between formulation (6.10) and the model of Keshavarz et al. (2011). First, note that if $\hat{\mathbf{x}} \in \Omega$, then an optimal $\boldsymbol{\alpha}^*$ derived from one of IOP($\hat{\mathbf{x}}$), LIOP($\hat{\mathbf{x}}$), or KES($\hat{\mathbf{x}}$) is optimal for all of them. If $f_k,\ k = 1, \ldots, K$ and $g_l,\ l = 1, \ldots, L$ are linear, then IOP($\hat{\mathbf{x}}$) and LIOP($\hat{\mathbf{x}}$) are equivalent. The next result draws an equivalence between LIOP($\hat{\mathbf{x}}$) and KES($\hat{\mathbf{x}}$), thus illustrating how the KES model can be modified to retain a given preference ordering.

Proposition 6.3.1 KES($\hat{\mathbf{x}}$) *is equivalent to* LIOP($\hat{\mathbf{x}}$) *if* $\boldsymbol{\delta} = \mathbf{0}, \phi(\boldsymbol{\delta}, \boldsymbol{\gamma}, \boldsymbol{\rho}) = -\boldsymbol{\gamma}'\mathbf{e} + \boldsymbol{\rho}'\mathbf{e}$, *and* $\boldsymbol{\alpha}$ *is normalized by the constraint* $\sum_{k=1}^{K} \alpha_k f_k(\hat{\mathbf{x}}) = 1$ *instead of* $\alpha_1 = 1$.

Proof: Consider the dual of LIOP($\hat{\mathbf{x}}$):

$$\begin{aligned} \underset{\boldsymbol{\alpha},\boldsymbol{\sigma},\boldsymbol{\pi}}{\text{maximize}} \quad & \mathbf{b}'\boldsymbol{\pi} + \sum_{k=1}^{K} \alpha_k (f_k(\hat{\mathbf{x}}) - \hat{\mathbf{x}}'\nabla_{\mathbf{x}} f_k(\hat{\mathbf{x}})) + \sum_{l=1}^{L} \sigma_l (g_l(\hat{\mathbf{x}}) - \hat{\mathbf{x}}'\nabla_{\mathbf{x}} g_l(\hat{\mathbf{x}})) \\ \text{subject to} \quad & \sum_{k=1}^{K} \alpha_k f_k(\hat{\mathbf{x}}) = 1, \\ & \sum_{k=1}^{K} \alpha_k \nabla_{\mathbf{x}} f_k(\hat{\mathbf{x}}) + \sum_{l=1}^{L} \sigma_l \nabla_{\mathbf{x}} g_l(\hat{\mathbf{x}}) - \mathbf{A}'\boldsymbol{\pi} = \mathbf{0}, \\ & \boldsymbol{\alpha} \geq \mathbf{0},\ \boldsymbol{\sigma} \geq \mathbf{0}. \end{aligned} \tag{6.15}$$

We multiply the first constraint in (6.15) by $\hat{\epsilon}$ to get

$$\hat{\epsilon} \sum_{k=1}^{K} \alpha_k f_k(\hat{\mathbf{x}}) - \hat{\epsilon} = 0, \tag{6.16}$$

and take the inner product of the second constraint in (6.15) with $\hat{\mathbf{x}}$ to get

$$\sum_{k=1}^{K} \alpha_k \hat{\mathbf{x}}'\nabla_{\mathbf{x}} f_k(\hat{\mathbf{x}}) + \sum_{l=1}^{L} \sigma_l \hat{\mathbf{x}}'\nabla_{\mathbf{x}} g_l(\hat{\mathbf{x}}) - \boldsymbol{\pi}'\mathbf{A}\hat{\mathbf{x}} = 0. \tag{6.17}$$

If we add the left-hand sides of equations (6.16) and (6.17) to the objective function of problem (6.15), we obtain the following problem:

$$\begin{aligned}
\underset{\boldsymbol{\alpha},\boldsymbol{\sigma},\boldsymbol{\pi}}{\text{maximize}} \quad & \hat{\epsilon} + \sum_{k=1}^{K} \alpha_k (f_k(\hat{\mathbf{x}}) - \hat{\epsilon} f_k(\hat{\mathbf{x}})) + \sum_{l=1}^{L} \sigma_l g_l(\hat{\mathbf{x}}) - \boldsymbol{\pi}'(\mathbf{A}\hat{\mathbf{x}} - \mathbf{b}), \\
\text{subject to} \quad & \sum_{k=1}^{K} \alpha_k f_k(\hat{\mathbf{x}}) = 1, \\
& \sum_{k=1}^{K} \alpha_k \nabla_{\mathbf{x}} f_k(\hat{\mathbf{x}}) + \sum_{l=1}^{L} \sigma_l \nabla_{\mathbf{x}} g_l(\hat{\mathbf{x}}) - \mathbf{A}'\boldsymbol{\pi} = \mathbf{0}, \\
& \boldsymbol{\alpha} \geq \mathbf{0}, \ \boldsymbol{\sigma} \geq \mathbf{0}.
\end{aligned} \tag{6.18}$$

Problem (6.18) is equivalent to the Lagrangian dual problem (6.11) with a fixed solution $(\hat{\epsilon}, \hat{\mathbf{x}})$ and the associated stationary conditions satisfied. That is, problem (6.18) minimizes the sum of the residuals only in the complementary slackness conditions with respect to $(\hat{\epsilon}, \hat{\mathbf{x}})$. Note that $\hat{\epsilon} + \sum_{k=1}^{K} \alpha_k (f_k(\hat{\mathbf{x}}) - \hat{\epsilon} f_k(\hat{\mathbf{x}})) = 1$, so solving (6.18) is equivalent to solving

$$\begin{aligned}
\underset{\boldsymbol{\alpha},\boldsymbol{\sigma},\boldsymbol{\pi}}{\text{maximize}} \quad & \sum_{l=1}^{L} \sigma_l g_l(\hat{\mathbf{x}}) - \boldsymbol{\pi}'(\mathbf{A}\hat{\mathbf{x}} - \mathbf{b}), \\
\text{subject to} \quad & \sum_{k=1}^{K} \alpha_k f_k(\hat{\mathbf{x}}) = 1, \\
& \sum_{k=1}^{K} \alpha_k \nabla_{\mathbf{x}} f_k(\hat{\mathbf{x}}) + \sum_{l=1}^{L} \sigma_l \nabla_{\mathbf{x}} g_l(\hat{\mathbf{x}}) - \mathbf{A}'\boldsymbol{\pi} = \mathbf{0}, \\
& \boldsymbol{\alpha} \geq \mathbf{0}, \ \boldsymbol{\sigma} \geq \mathbf{0}.
\end{aligned} \tag{6.19}$$

KES($\hat{\mathbf{x}}$) with $\boldsymbol{\delta} = \mathbf{0}, \phi(\boldsymbol{\delta}, \boldsymbol{\gamma}, \boldsymbol{\rho}) = -\boldsymbol{\gamma}'\mathbf{e} + \boldsymbol{\rho}'\mathbf{e}$, and $\boldsymbol{\alpha}$ normalized by the constraint

$\sum_{k=1}^{K} \alpha_k f_k(\hat{\mathbf{x}}) = 1$ is

$$
\begin{aligned}
\underset{\boldsymbol{\alpha},\boldsymbol{\sigma},\boldsymbol{\pi},\boldsymbol{\gamma},\boldsymbol{\rho}}{\text{minimize}} \quad & -\boldsymbol{\gamma}'\mathbf{e} + \boldsymbol{\rho}'\mathbf{e} \\
\text{subject to} \quad & \sum_{k=1}^{K} \alpha_k \nabla_{\mathbf{x}} f_k(\hat{\mathbf{x}}) + \sum_{l=1}^{L} \sigma_l \nabla_{\mathbf{x}} g_l(\hat{\mathbf{x}}) - \mathbf{A}'\boldsymbol{\pi} = \mathbf{0}, \\
& \sigma_l g_l(\hat{\mathbf{x}}) = \gamma_l, \quad l = 1, \ldots, L, \\
& \pi_j(\mathbf{a}_j'\hat{\mathbf{x}} - b_j) = \rho_j, \quad j = 1, \ldots, m, \\
& \sum_{k=1}^{K} \alpha_k f_k(\hat{\mathbf{x}}) = 1, \\
& \boldsymbol{\alpha} \geq \mathbf{0}, \ \boldsymbol{\sigma} \geq \mathbf{0},
\end{aligned}
\tag{6.20}
$$

which is equivalent to solving problem (6.19). □

Note that if $\hat{\mathbf{x}} \in \mathbf{X}$, the form of $\phi(\boldsymbol{\delta}, \boldsymbol{\gamma}, \boldsymbol{\rho})$ in the statement of Proposition 6.3.1 can be replaced with $\phi(\boldsymbol{\delta}, \boldsymbol{\gamma}, \boldsymbol{\rho}) = ||\boldsymbol{\gamma}||_1 + ||\boldsymbol{\rho}||_1$. The next result is analogous to Proposition 6.3.1 for the case of absolute preference orderings and its proof is omitted.

Proposition 6.3.2 KES($\hat{\mathbf{x}}$) *is equivalent to the linearization of formulation* (6.13) *at* $\hat{\mathbf{x}}$*,*

$$
\begin{aligned}
\underset{\epsilon,\mathbf{x}}{\textit{minimize}} \quad & \epsilon \\
\textit{subject to} \quad & \epsilon \geq \nabla_{\mathbf{x}} f_k(\hat{\mathbf{x}})'(\mathbf{x} - \hat{\mathbf{x}}), \quad k = 1, \ldots, K, \\
& g_l(\hat{\mathbf{x}}) + \nabla_{\mathbf{x}} g_l(\hat{\mathbf{x}})'(\mathbf{x} - \hat{\mathbf{x}}) \leq 0, \quad l = 1, \ldots, L, \\
& \mathbf{A}\mathbf{x} = \mathbf{b},
\end{aligned}
\tag{6.21}
$$

if $\boldsymbol{\delta} = \mathbf{0}, \phi(\boldsymbol{\delta}, \boldsymbol{\gamma}, \boldsymbol{\rho}) = -\boldsymbol{\gamma}'\mathbf{e} + \boldsymbol{\rho}'\mathbf{e}$*, and* $\boldsymbol{\alpha}$ *is normalized by the constraint* $\sum_{k=1}^{K} \alpha_k = 1$ *instead of* $\alpha_1 = 1$*.*

Propositions 6.3.1 and 6.3.2 state that by specifying $\boldsymbol{\delta}$, the function $\phi(\boldsymbol{\delta}, \boldsymbol{\gamma}, \boldsymbol{\rho})$ and the normalization constraint in KES($\hat{\mathbf{x}}$), a decision maker can enforce a particular preference ordering in the inverse optimization process. These observations suggest that a simple

general formulation can encompass several of the previous models presented:

$$\begin{aligned}
\underset{\boldsymbol{\alpha},\boldsymbol{\sigma},\boldsymbol{\pi},\boldsymbol{\delta},\boldsymbol{\gamma},\boldsymbol{\rho}}{\text{minimize}} \quad & \phi(\boldsymbol{\delta},\boldsymbol{\gamma},\boldsymbol{\rho}) \\
\text{subject to} \quad & \sum_{k=1}^{K} \alpha_k \nabla_{\mathbf{x}} f_k(\hat{\mathbf{x}}) + \sum_{l=1}^{L} \sigma_l \nabla_{\mathbf{x}} g_l(\hat{\mathbf{x}}) + \mathbf{A}'\boldsymbol{\pi} = \boldsymbol{\delta}, \\
& \sigma_l g_l(\hat{\mathbf{x}}) = \gamma_l, \quad l = 1,\ldots,L, \\
& \boldsymbol{\pi}'(\mathbf{A}\hat{\mathbf{x}} - \mathbf{b}) = \boldsymbol{\rho}, \\
& \boldsymbol{\alpha}'\mathbf{h} = 1, \\
& \boldsymbol{\alpha} \geq \mathbf{0},\ \boldsymbol{\sigma} \geq \mathbf{0}.
\end{aligned} \tag{6.22}$$

Formulation (6.22) can be specialized to the model of Keshavarz et al. (2011) without explicit consideration of preferences if $\mathbf{h} = \mathbf{e}_k$, where $\mathbf{e}_k$ is a vector of zeros except for a one in the k-th component. It can be specialized to the relative LIOP model if $\boldsymbol{\delta} = \mathbf{0}, \phi(\boldsymbol{\delta},\boldsymbol{\gamma},\boldsymbol{\rho}) = -\boldsymbol{\gamma}'\mathbf{e} + \boldsymbol{\rho}'\mathbf{e}, \mathbf{h} = \mathbf{f}(\hat{\mathbf{x}})$ or the absolute LIOP model if $\boldsymbol{\delta} = \mathbf{0}, \phi(\boldsymbol{\delta},\boldsymbol{\gamma}) = -\boldsymbol{\gamma}'\mathbf{e} + \boldsymbol{\rho}'\mathbf{e}, \mathbf{h} = \mathbf{e}$. If f_k and g_l are linear functions and $\mathbf{h} = \mathbf{f}(\hat{\mathbf{x}})$ or $\mathbf{h} = \mathbf{e}$, then formulation (6.22) is equivalent to the generalized inverse linear optimization models (4.3) or (4.5), respectively. If $\hat{\mathbf{x}} \in \Omega$, then formulation (6.22) specializes to the standard inverse optimization models of Ahuja and Orlin (2001) and Iyengar and Kang (2005) for linear programming and convex programming, respectively.

6.4 Computational results

In IMRT, historical treatments implicitly encode preferences or accepted trade-offs between the conflicting objectives. Thus, any knowledge-based treatment planning framework should consider maintaining those preferences. In fact, we believe maintaining the relative values of the objective function weights gives the optimization engine more flexibility to generate a personalized treatment respecting the patient's anatomy compared to simply porting over achievable dose metrics. In this section, we use data from

prostate cancer radiation therapy to demonstrate the models presented in this chapter. We demonstrate to what extent the initial preference ordering is preserved by the different inverse models as they attempt to learn weights for a convex quadratic treatment planning problem.

6.4.1 A multiobjective forward formulation

We again consider the clinical forward multiobjective optimization formulation (3.1). For each OAR k, we associate an objective $f_k(\mathbf{w}) = ||(\mathbf{D}_k\mathbf{w}-\theta_k\mathbf{e})_+||_2^2$, where θ_k denotes a dose threshold for structure k, and $(\cdot)_+$ denotes a vector with the operator $\max\{0,\cdot\}$ applied to each component. In the experiments below, we consider five OARs (i.e., $|\mathcal{K}| = 5$): the bladder, rectum, left and right femoral heads, and POR, i.e., a ring of healthy tissue around the PTV which is used to encourage conformity of the dose around the PTV. We let $\theta_k = 30$ for the left and right femoral heads, and $\theta_k = 50$ for the other three OARs, based on a protocol at Princess Margaret Cancer Centre.

6.4.2 Impact of different inverse optimization models on preferences

We used 24 historical treatments delivered at Princess Margaret Cancer Centre for our input vectors. Similar to Section 4.4, our discussion in this section will be based on the compact inverse formulations presented in Section 6.2 instead of the clinical specific formulations, that is, we use $\hat{\mathbf{x}}$ as a historical treatment input to the inverse models. For each of the historical treatments $\hat{\mathbf{x}}$, we derived weights using KES($\hat{\mathbf{x}}$), IOP($\hat{\mathbf{x}}$), LIOP($\hat{\mathbf{x}}$), and the SLP algorithm. For KES($\hat{\mathbf{x}}$), we solved five different instances, corresponding to five different ways of normalizing the weight vector (i.e., $\alpha_k = 1, k = 1,\ldots,5$). Parameters and termination criteria for the SLP algorithm were chosen based on Zhang et al. (1985) (e.g., the algorithm terminates when the ℓ_2 norm difference between solutions in

two consecutive iterations is less than 0.001). All the models were solved using CPLEX 12.3 on a computer with a 3.07 GHz 12-core CPU and 32 GB of RAM.

Tables 6.1 and 6.2 summarize the results of applying the five variations of the KES($\hat{\mathbf{x}}$) model to all patients. It can be seen that the imputed weight values are heavily dependent on which normalizing constraint is used – note that weights were re-normalized so they sum to one after solving the inverse model. For all patients, the weights in Tables 6.1 and 6.2 suggest that the KES model gives the vast majority of the weight to the objective whose weight is used in the normalization constraint. For each instance, $f_k(\mathbf{x}^*)/f_k(\hat{\mathbf{x}})$ denotes the ratio of the objective value achieved by $\boldsymbol{\alpha}^*$ to the objective value associated with $\hat{\mathbf{x}}$ for objective k. In Tables 6.1 and 6.2, we see that there is no well-defined pattern in the ratios; generally they are less than one, which implies that $\mathbf{x}^*$ is an improvement on $\hat{\mathbf{x}}$. However, the amount of improvement varies from objective to objective and from model to model. Furthermore, KES1 provides an example where the dose on the left femur actually increases. Overall, these results suggest that if preference preservation in the inverse process is important (e.g., in multiobjective problems), it requires explicit consideration in the inverse model.

Tables 6.3 and 6.4 show the results from applying the IOP model, LIOP model, and SLP algorithm to all patients. For the SLP algorithm, the solution time is the total running time through all iterations. As expected, the component-wise ratio obtained by the IOP model is equivalent to ϵ^* unless the corresponding weight is zero in which case the ratio is at most ϵ^* (see Corollary 6.2.2). The practical interpretation of this result is that although the ratio $f_k(\mathbf{x}^*)/f_k(\hat{\mathbf{x}}) = \epsilon^*$ is achievable, the objective value can be further decreased without any sacrifice. Note that preserving the preference ordering using the IOP model comes at a higher computational cost compared to the LIOP model, which only approximately maintains the preference ordering. Also note that ϵ^* from LIOP is less than ϵ^* from IOP as expected, because LIOP is an outer approximation to IOP. Finally, the SLP algorithm strikes a middle ground between the LIOP and IOP models

Table 6.1: Comparing the model of Keshavarz et al. (2011) with different weights fixed to one: patients #1 – 12.

		KES1		KES2		KES3		KES4		KES5	
Pat.	OAR	$\boldsymbol{\alpha}^*$	$\frac{f_k(\mathbf{x}^*)}{f_k(\hat{\mathbf{x}})}$	$\boldsymbol{\alpha}^*$	$\frac{f_k(\mathbf{x}^*)}{f_k(\hat{\mathbf{x}})}$	$\boldsymbol{\alpha}^*$	$\frac{f_k(\mathbf{x}^*)}{f_k(\hat{\mathbf{x}})}$	$\boldsymbol{\alpha}^*$	$\frac{f_k(\mathbf{x}^*)}{f_k(\hat{\mathbf{x}})}$	$\boldsymbol{\alpha}^*$	$\frac{f_k(\mathbf{x}^*)}{f_k(\hat{\mathbf{x}})}$
	Blad	0.973	0.653	0.003	0.837	0.000	0.782	0.000	0.811	0.001	0.823
	Rect	0.004	1.000	0.865	0.823	0.000	0.949	0.000	0.963	0.002	0.972
1	LFem	0.000	1.143	0.005	0.232	0.992	0.000	0.004	0.001	0.000	0.023
	RFem	0.000	0.671	0.004	0.177	0.004	0.002	0.992	0.000	0.000	0.050
	POR	0.023	0.929	0.123	0.614	0.004	0.412	0.004	0.411	0.997	0.404
	Blad	0.906	0.805	0.002	1.178	0.001	0.850	0.001	0.861	0.007	0.996
	Rect	0.000	0.783	0.987	0.555	0.000	0.697	0.000	0.702	0.000	0.718
2	LFem	0.015	0.636	0.001	2.853	0.998	0.000	0.001	0.143	0.001	0.968
	RFem	0.003	2.053	0.001	3.139	0.000	0.249	0.997	0.000	0.000	1.181
	POR	0.076	0.687	0.009	1.282	0.001	0.493	0.001	0.467	0.992	0.393
	Blad	0.941	0.688	0.006	0.963	0.000	0.893	0.000	0.833	0.000	0.938
	Rect	0.006	0.872	0.917	0.683	0.000	0.779	0.000	0.743	0.001	0.921
3	LFem	0.015	0.637	0.008	0.156	0.996	0.000	0.000	0.018	0.001	0.438
	RFem	0.009	1.828	0.007	1.601	0.001	0.013	0.999	0.000	0.001	0.098
	POR	0.029	0.832	0.062	0.785	0.003	0.374	0.001	0.397	0.997	0.342
	Blad	0.647	0.582	0.029	0.743	0.027	0.689	0.027	0.690	0.032	0.746
	Rect	0.071	0.878	0.658	0.806	0.070	0.863	0.070	0.861	0.085	0.882
4	LFem	0.034	0.165	0.042	0.162	0.659	0.001	0.035	0.069	0.042	0.149
	RFem	0.029	0.262	0.037	0.163	0.036	0.066	0.661	0.002	0.041	0.130
	POR	0.219	0.439	0.234	0.445	0.208	0.363	0.207	0.365	0.800	0.337
	Blad	0.805	0.765	0.032	0.895	0.001	0.803	0.001	0.806	0.009	0.969
	Rect	0.015	0.938	0.835	0.809	0.000	0.901	0.000	0.914	0.003	1.012
5	LFem	0.015	0.146	0.006	0.423	0.998	0.000	0.001	0.007	0.001	0.008
	RFem	0.033	0.115	0.014	0.319	0.000	0.005	0.997	0.000	0.002	0.021
	POR	0.132	0.705	0.113	0.665	0.001	0.511	0.001	0.502	0.985	0.407
	Blad	0.844	0.769	0.027	0.920	0.002	0.811	0.003	0.811	0.008	0.925
	Rect	0.041	0.914	0.848	0.756	0.002	0.875	0.001	0.863	0.003	0.959
6	LFem	0.020	0.135	0.004	1.096	0.991	0.000	0.003	0.004	0.000	0.022
	RFem	0.009	0.174	0.008	0.343	0.002	0.003	0.989	0.000	0.000	0.423
	POR	0.086	0.605	0.113	0.601	0.003	0.443	0.004	0.449	0.989	0.383
	Blad	0.976	0.764	0.016	0.945	0.000	0.876	0.000	0.884	0.001	0.976
	Rect	0.007	0.991	0.921	0.838	0.000	0.892	0.000	0.888	0.004	1.021
7	LFem	0.001	0.356	0.003	0.289	0.999	0.000	0.001	0.004	0.000	0.038
	RFem	0.002	1.157	0.003	0.756	0.001	0.007	0.999	0.000	0.000	0.223
	POR	0.014	1.168	0.057	0.741	0.000	0.438	0.000	0.443	0.995	0.383
	Blad	0.802	0.801	0.016	0.991	0.002	0.927	0.002	0.929	0.003	0.979
	Rect	0.021	0.922	0.795	0.776	0.001	0.944	0.001	0.924	0.004	0.956
8	LFem	0.007	0.138	0.005	0.223	0.975	0.000	0.003	0.026	0.000	0.141
	RFem	0.008	0.096	0.009	0.123	0.003	0.035	0.980	0.000	0.000	0.100
	POR	0.162	0.702	0.175	0.702	0.019	0.563	0.014	0.578	0.993	0.509
	Blad	0.918	0.686	0.044	0.823	0.001	0.754	0.001	0.788	0.003	0.871
	Rect	0.006	0.973	0.821	0.742	0.000	0.885	0.001	0.834	0.001	0.945
9	LFem	0.008	0.583	0.005	3.975	0.996	0.000	0.000	0.171	0.000	2.352
	RFem	0.012	0.485	0.005	1.830	0.001	0.030	0.996	0.000	0.000	0.526
	POR	0.056	0.765	0.125	0.646	0.002	0.437	0.002	0.407	0.996	0.354
	Blad	0.885	0.662	0.002	0.934	0.000	0.782	0.001	0.785	0.001	0.860
	Rect	0.000	0.972	0.955	0.689	0.000	0.856	0.000	0.912	0.000	0.976
10	LFem	0.015	0.175	0.005	0.131	0.998	0.000	0.004	0.006	0.001	0.021
	RFem	0.003	0.742	0.000	1.403	0.001	0.024	0.991	0.000	0.000	0.193
	POR	0.097	0.689	0.038	1.034	0.001	0.434	0.004	0.437	0.998	0.414
	Blad	0.950	0.667	0.026	0.854	0.001	0.757	0.001	0.745	0.000	0.813
	Rect	0.013	0.976	0.922	0.677	0.000	0.869	0.000	0.888	0.001	0.912
11	LFem	0.011	0.305	0.007	0.785	0.996	0.000	0.002	0.020	0.000	0.325
	RFem	0.001	1.653	0.008	0.372	0.001	0.033	0.995	0.000	0.000	0.570
	POR	0.025	0.873	0.037	1.059	0.002	0.425	0.002	0.439	0.999	0.386
	Blad	0.996	0.504	0.000	1.709	0.000	0.603	0.000	0.601	0.000	0.724
	Rect	0.000	1.239	0.996	0.625	0.000	0.696	0.000	0.699	0.000	0.843
12	LFem	0.000	2.917	0.000	8.791	1.000	0.000	0.000	0.006	0.000	0.254
	RFem	0.000	2.052	0.000	6.575	0.000	0.005	1.000	0.000	0.000	0.056
	POR	0.004	1.163	0.004	1.693	0.000	0.353	0.000	0.354	1.000	0.278

KES#: KES with #-th weight fixed to one. Weights were renormalized so they add up to one.

Table 6.2: Comparing the model of Keshavarz et al. (2011) with different weights fixed to one: patients #13 – 24.

		KES1		KES2		KES3		KES4		KES5	
Pat.	OAR	$\boldsymbol{\alpha}^*$	$\frac{f_k(\mathbf{x}^*)}{f_k(\hat{\mathbf{x}})}$	$\boldsymbol{\alpha}^*$	$\frac{f_k(\mathbf{x}^*)}{f_k(\hat{\mathbf{x}})}$	$\boldsymbol{\alpha}^*$	$\frac{f_k(\mathbf{x}^*)}{f_k(\hat{\mathbf{x}})}$	$\boldsymbol{\alpha}^*$	$\frac{f_k(\mathbf{x}^*)}{f_k(\hat{\mathbf{x}})}$	$\boldsymbol{\alpha}^*$	$\frac{f_k(\mathbf{x}^*)}{f_k(\hat{\mathbf{x}})}$
	Blad	0.971	0.656	0.007	0.989	0.000	0.720	0.000	0.747	0.000	0.892
	Rect	0.008	0.952	0.937	0.656	0.000	0.743	0.000	0.751	0.003	0.818
13	LFem	0.003	1.547	0.008	0.641	1.000	0.000	0.000	0.022	0.000	0.071
	RFem	0.004	1.257	0.002	1.514	0.000	0.033	1.000	0.000	0.000	0.372
	POR	0.014	0.904	0.046	0.737	0.000	0.430	0.000	0.373	0.997	0.317
	Blad	0.944	0.687	0.007	0.812	0.000	0.765	0.000	0.766	0.000	0.881
	Rect	0.016	0.925	0.900	0.768	0.000	0.838	0.000	0.866	0.006	1.017
14	LFem	0.001	0.122	0.001	0.383	0.999	0.000	0.000	0.003	0.000	0.559
	RFem	0.004	0.698	0.001	1.837	0.001	0.001	0.999	0.000	0.000	0.140
	POR	0.035	0.720	0.091	0.588	0.000	0.384	0.001	0.372	0.994	0.337
	Blad	0.915	0.738	0.005	0.983	0.000	0.855	0.000	0.843	0.006	0.944
	Rect	0.001	1.005	0.953	0.768	0.000	0.915	0.000	0.924	0.000	0.946
15	LFem	0.002	0.574	0.001	1.021	0.998	0.000	0.000	0.036	0.000	0.205
	RFem	0.001	0.313	0.001	0.642	0.001	0.034	0.999	0.000	0.000	0.264
	POR	0.081	0.647	0.040	0.816	0.001	0.482	0.001	0.495	0.994	0.441
	Blad	0.887	0.787	0.004	1.021	0.002	0.849	0.001	0.846	0.007	1.033
	Rect	0.000	0.928	0.972	0.759	0.000	0.908	0.000	0.903	0.000	0.944
16	LFem	0.003	0.494	0.001	0.821	0.997	0.000	0.000	0.051	0.000	0.083
	RFem	0.004	0.282	0.000	1.665	0.000	0.040	0.997	0.000	0.000	0.121
	POR	0.106	0.649	0.023	0.918	0.001	0.419	0.002	0.430	0.993	0.334
	Blad	0.953	0.723	0.001	1.230	0.000	0.765	0.001	0.745	0.001	0.910
	Rect	0.000	0.847	0.998	0.586	0.000	0.659	0.000	0.644	0.000	0.718
17	LFem	0.004	1.550	0.000	3.589	0.999	0.000	0.000	0.073	0.000	0.166
	RFem	0.004	0.311	0.000	1.192	0.001	0.035	0.999	0.000	0.000	0.072
	POR	0.039	0.677	0.001	1.372	0.000	0.395	0.000	0.488	0.999	0.330
	Blad	0.806	0.790	0.040	0.888	0.003	0.820	0.004	0.833	0.012	0.942
	Rect	0.006	1.068	0.783	0.823	0.000	0.996	0.000	1.003	0.003	1.039
18	LFem	0.005	0.150	0.002	0.359	0.995	0.000	0.000	0.020	0.000	0.052
	RFem	0.009	0.363	0.006	0.184	0.000	0.033	0.992	0.000	0.001	0.041
	POR	0.174	0.625	0.169	0.680	0.002	0.510	0.004	0.501	0.984	0.432
	Blad	0.956	0.719	0.004	0.952	0.000	0.789	0.000	0.801	0.000	0.885
	Rect	0.003	1.086	0.888	0.785	0.000	0.855	0.000	0.844	0.001	0.943
19	LFem	0.012	3.755	0.012	1.625	0.999	0.000	0.000	0.012	0.000	0.979
	RFem	0.004	2.471	0.011	1.105	0.001	0.023	1.000	0.000	0.000	0.052
	POR	0.025	0.908	0.085	0.698	0.000	0.411	0.000	0.417	0.999	0.370
	Blad	0.973	0.711	0.003	0.993	0.000	0.781	0.000	0.773	0.000	0.895
	Rect	0.002	1.006	0.960	0.640	0.000	0.726	0.000	0.702	0.000	0.838
20	LFem	0.005	0.605	0.002	4.023	1.000	0.000	0.002	0.001	0.000	0.580
	RFem	0.002	0.658	0.004	0.348	0.000	0.002	0.998	0.000	0.000	0.031
	POR	0.018	0.812	0.031	0.886	0.000	0.350	0.000	0.374	1.000	0.322
	Blad	0.992	0.650	0.007	0.936	0.000	0.729	0.000	0.704	0.000	0.859
	Rect	0.001	1.191	0.986	0.619	0.000	0.708	0.000	0.693	0.000	0.780
21	LFem	0.002	0.756	0.002	1.968	0.999	0.000	0.000	0.044	0.000	1.804
	RFem	0.001	2.653	0.002	2.902	0.001	0.009	1.000	0.000	0.000	0.297
	POR	0.004	1.254	0.003	1.086	0.000	0.343	0.000	0.374	1.000	0.274
	Blad	0.923	0.713	0.007	0.906	0.001	0.762	0.001	0.761	0.003	0.917
	Rect	0.000	1.012	0.979	0.752	0.000	0.909	0.000	0.908	0.000	0.974
22	LFem	0.004	1.928	0.001	5.600	0.997	0.000	0.000	0.045	0.000	2.053
	RFem	0.021	0.260	0.003	2.284	0.001	0.018	0.998	0.000	0.000	2.362
	POR	0.052	0.705	0.010	1.355	0.001	0.412	0.001	0.412	0.997	0.325
	Blad	0.924	0.787	0.011	1.073	0.000	0.857	0.038	0.854	0.003	1.015
	Rect	0.003	0.940	0.966	0.756	0.000	0.842	0.043	0.838	0.000	0.934
23	LFem	0.007	1.010	0.001	1.832	1.000	0.000	0.417	0.025	0.000	0.542
	RFem	0.018	0.330	0.000	3.162	0.000	0.028	0.000	0.000	0.000	0.137
	POR	0.048	0.995	0.022	1.125	0.001	0.482	0.502	0.487	0.997	0.414
	Blad	0.836	0.814	0.037	0.949	0.001	0.932	0.001	0.939	0.006	1.077
	Rect	0.020	0.919	0.866	0.771	0.000	0.909	0.001	0.879	0.003	0.950
24	LFem	0.004	0.866	0.005	0.876	0.994	0.000	0.000	0.497	0.000	1.317
	RFem	0.004	0.888	0.006	0.649	0.002	0.013	0.991	0.000	0.000	0.520
	POR	0.136	0.679	0.086	0.713	0.003	0.516	0.006	0.486	0.991	0.422

KES#: KES with #-th weight fixed to one. Weights were renormalized so they add up to one.

in terms of maintaining preferences and computational efficiency.

Next, we compared clinical performance of the weights obtained from the IOP, LIOP, SLP, and the model of Keshavarz et al. (2011). We solved the forward treatment planning formulation with the convex objective functions with different weights obtained from the different models, and compared the resulting treatment plans. For the KES results for each patient, we report the performance of the model with the normalization constraint $\alpha_{k'} = 1$, where k' is the structure with the highest inverse weight determined by the IOP model. For example, in the case of patient #1, the IOP model says that the rectum weight is the highest, so we use the inverse weights derived from the KES model that employs the normalization constraint $\alpha_2 = 1$ (the rectum is the second objective). Figure 6.3 shows the comparison of the DVHs for the clinical plan and a plan generated by the weights from our IOP model for patient #1. The two sets of DVHs are quite similar, with the DVHs from the IOP weights nearly equally distanced from the clinical DVHs (i.e., improved) for each organ, which implies that the treatment planner's preference on the five structures has been maintained.

Figure 6.4 compares DVHs for the clinical plan and a plan generated by the weights obtained from the model of Keshavarz et al. (2011) for patient #1. The DVH for the rectum is less shifted below than those for the other organs, which means the rectum gets less emphasis than what it should have received in order to have had the overall preference ordering maintained. Compared to Figure 6.3, the weights from our IOP model achieve slightly more sparing in the rectum (especially in a higher dose region) while sacrificing quite a lot in the other organs, which is apparently more consistent with the original intention of the treatment planner who designed the clinical plan. Figures 6.5 and 6.6, and Figures 6.7 and 6.8 show similar results for patients #2 and #3, respectively.

The difference in the clinical performance between the IOP weights and the weights from the model of Keshavarz et al. (2011) is more noticeable in patient #3's case. Note from Table 6.3 that patient #3 had IOP weights that are almost equally distributed

Table 6.3: Comparison of the results from IOP, LIOP, and SLP algorithm: patients #1 – 12.

		IOP				LIOP				SLP			
Pat	Organ	ϵ^*	$\boldsymbol{\alpha}^*$	$\frac{f_k(\mathbf{x}^*)}{f_k(\bar{\mathbf{x}})}$	Time (s)	ϵ^*	$\boldsymbol{\alpha}^*$	$\frac{f_k(\mathbf{x}^*)}{f_k(\bar{\mathbf{x}})}$	Time (s)	ϵ^*	$\boldsymbol{\alpha}^*$	$\frac{f_k(\mathbf{x}^*)}{f_k(\bar{\mathbf{x}})}$	Time (s)
	Blad		0.014	0.812			0.013	0.813			0.014	0.815	
	Rect		0.935	0.812			0.936	0.812			0.935	0.812	
1	LFem	0.812	0.002	0.812	1582	0.769	0.001	0.905	168	0.813	0.002	0.746	663 [31]
	RFem		0.000	0.809			0.002	0.695			0.000	0.932	
	POR		0.049	0.812			0.048	0.816			0.049	0.813	
	Blad		0.932	0.803			0.937	0.802			0.933	0.803	
	Rect		0.002	0.802			0.007	0.754			0.002	0.802	
2	LFem	0.803	0.012	0.803	1384	0.748	0.013	0.771	82	0.803	0.011	0.846	1001 [57]
	RFem		0.018	0.803			0.015	0.965			0.018	0.808	
	POR		0.036	0.803			0.028	0.825			0.036	0.803	
	Blad		0.471	0.712			0.456	0.713			0.474	0.712	
	Rect		0.412	0.712			0.438	0.709			0.409	0.712	
3	LFem	0.712	0.008	0.712	764	0.617	0.012	0.525	33	0.712	0.007	0.751	361 [31]
	RFem		0.033	0.712			0.027	0.889			0.033	0.696	
	POR		0.076	0.712			0.067	0.746			0.077	0.709	
	Blad		0.011	0.780			0.005	0.814			0.012	0.773	
	Rect		0.929	0.780			0.938	0.778			0.925	0.780	
4	LFem	0.780	0.013	0.780	1052	0.720	0.012	0.797	69	0.781	0.013	0.755	714 [81]
	RFem		0.011	0.780			0.010	0.850			0.012	0.744	
	POR		0.036	0.780			0.035	0.820			0.038	0.771	
	Blad		0.138	0.809			0.202	0.798			0.142	0.809	
	Rect		0.777	0.809			0.701	0.787			0.773	0.809	
5	LFem	0.809	0.001	0.800	1171	0.729	0.000	1.508	63	0.809	0.000	0.940	418 [30]
	RFem		0.005	0.809			0.006	1.559			0.004	0.883	
	POR		0.079	0.809			0.092	0.871			0.081	0.804	
	Blad		0.479	0.787			0.503	0.783			0.491	0.785	
	Rect		0.480	0.787			0.467	0.788			0.471	0.787	
6	LFem	0.787	0.001	0.785	2725	0.679	0.002	0.737	150	0.788	0.000	1.207	1081 [35]
	RFem		0.004	0.787			0.005	0.687			0.004	0.842	
	POR		0.036	0.787			0.023	0.878			0.034	0.807	
	Blad		0.175	0.842			0.099	0.866			0.173	0.842	
	Rect		0.782	0.842			0.862	0.839			0.784	0.842	
7	LFem	0.842	0.000	0.575	1192	0.807	0.000	0.767	52	0.842	0.000	0.574	324 [17]
	RFem		0.000	0.709			0.001	0.791			0.000	0.728	
	POR		0.043	0.841			0.038	0.855			0.043	0.847	
	Blad		0.537	0.805			0.476	0.805			0.531	0.804	
	Rect		0.343	0.805			0.459	0.793			0.363	0.803	
8	LFem	0.805	0.000	0.804	3242	0.764	0.000	0.808	182	0.805	0.000	0.559	774 [20]
	RFem		0.000	0.804			0.000	0.660			0.000	0.483	
	POR		0.120	0.805			0.065	0.899			0.106	0.823	
	Blad		0.175	0.750			0.105	0.777			0.172	0.750	
	Rect		0.701	0.750			0.728	0.752			0.709	0.749	
9	LFem	0.750	0.028	0.750	822	0.683	0.042	0.520	49	0.751	0.026	0.804	626 [53]
	RFem		0.022	0.750			0.023	0.629			0.022	0.780	
	POR		0.074	0.750			0.102	0.665			0.071	0.764	
	Blad		0.168	0.709			0.173	0.715			0.169	0.708	
	Rect		0.694	0.709			0.672	0.712			0.693	0.709	
10	LFem	0.709	0.000	0.708	1316	0.590	0.000	0.335	75	0.709	0.000	0.372	591 [24]
	RFem		0.001	0.709			0.000	0.885			0.001	0.753	
	POR		0.137	0.709			0.155	0.677			0.137	0.709	
	Blad		0.424	0.713			0.472	0.709			0.414	0.714	
	Rect		0.478	0.713			0.429	0.718			0.487	0.712	
11	LFem	0.713	0.005	0.713	825	0.589	0.003	0.890	46	0.714	0.004	0.797	795 [75]
	RFem		0.004	0.713			0.003	0.850			0.004	0.788	
	POR		0.089	0.713			0.093	0.698			0.091	0.709	
	Blad		0.058	0.646			0.049	0.661			0.058	0.646	
	Rect		0.891	0.646			0.882	0.649			0.891	0.646	
12	LFem	0.646	0.003	0.646	1017	0.477	0.004	0.461	96	0.646	0.003	0.692	601 [42]
	RFem		0.004	0.646			0.004	0.505			0.004	0.602	
	POR		0.044	0.646			0.061	0.551			0.044	0.643	

[·] denotes the number of iterations.

Table 6.4: Comparison of the results from IOP, LIOP, and SLP algorithm: patients #13 – 24.

		IOP				LIOP				SLP			
Pat	Organ	ϵ^*	$\boldsymbol{\alpha}^*$	$\frac{f_k(\mathbf{x}^*)}{f_k(\bar{\mathbf{x}})}$	Time (s)	ϵ^*	$\boldsymbol{\alpha}^*$	$\frac{f_k(\mathbf{x}^*)}{f_k(\bar{\mathbf{x}})}$	Time (s)	ϵ^*	$\boldsymbol{\alpha}^*$	$\frac{f_k(\mathbf{x}^*)}{f_k(\bar{\mathbf{x}})}$	Time (s)
	Blad		0.460	0.691			0.736	0.657			0.456	0.671	
	Rect		0.476	0.691			0.220	0.693			0.482	0.670	
13	LFem	0.671	0.018	0.691	1246	0.521	0.029	0.475	42	0.671	0.016	0.756	295 [31]
	RFem		0.005	0.691			0.004	0.793			0.005	0.642	
	POR		0.041	0.691			0.011	0.798			0.041	0.675	
	Blad		0.037	0.762			0.016	0.787			0.034	0.764	
	Rect		0.920	0.762			0.962	0.757			0.925	0.761	
14	LFem	0.762	0.001	0.762	904	0.667	0.000	2.289	13	0.763	0.001	0.542	189 [32]
	RFem		0.008	0.762			0.000	3.229			0.008	0.892	
	POR		0.034	0.762			0.022	0.896			0.032	0.773	
	Blad		0.174	0.779			0.178	0.778			0.167	0.780	
	Rect		0.775	0.779			0.771	0.779			0.782	0.778	
15	LFem	0.779	0.002	0.778	2140	0.702	0.003	0.662	127	0.779	0.002	0.792	684 [28]
	RFem		0.001	0.778			0.001	0.860			0.001	0.789	
	POR		0.048	0.778			0.047	0.784			0.048	0.782	
	Blad		0.611	0.789			0.819	0.780			0.626	0.788	
	Rect		0.329	0.789			0.151	0.800			0.314	0.791	
16	LFem	0.789	0.002	0.789	2987	0.732	0.001	1.212	117	0.789	0.003	0.754	694 [24]
	RFem		0.000	0.788			0.000	1.196			0.000	0.968	
	POR		0.058	0.789			0.029	0.971			0.057	0.788	
	Blad		0.939	0.724			0.923	0.725			0.938	0.724	
	Rect		0.021	0.724			0.028	0.703			0.021	0.723	
17	LFem	0.724	0.015	0.724	2013	0.616	0.008	1.046	64	0.724	0.015	0.742	395 [25]
	RFem		0.001	0.721			0.000	0.604			0.002	0.643	
	POR		0.024	0.724			0.041	0.638			0.024	0.725	
	Blad		0.187	0.820			0.148	0.827			0.185	0.819	
	Rect		0.716	0.820			0.762	0.816			0.721	0.819	
18	LFem	0.820	0.001	0.820	1098	0.772	0.000	1.231	66	0.820	0.000	1.220	490 [37]
	RFem		0.001	0.819			0.000	1.098			0.000	1.091	
	POR		0.095	0.820			0.090	0.838			0.094	0.829	
	Blad		0.130	0.788			0.083	0.809			0.129	0.788	
	Rect		0.760	0.788			0.814	0.784			0.761	0.788	
19	LFem	0.788	0.035	0.787	801	0.743	0.035	0.826	24	0.788	0.036	0.780	244 [23]
	RFem		0.022	0.788			0.018	0.914			0.021	0.800	
	POR		0.053	0.788			0.050	0.812			0.053	0.788	
	Blad		0.904	0.717			0.921	0.717			0.903	0.717	
	Rect		0.064	0.717			0.045	0.738			0.065	0.714	
20	LFem	0.717	0.003	0.713	9937	0.616	0.000	0.948	151	0.717	0.003	0.693	1150 [17]
	RFem		0.004	0.716			0.004	0.602			0.004	0.781	
	POR		0.026	0.717			0.030	0.668			0.025	0.727	
	Blad		0.708	0.671			0.744	0.670			0.710	0.671	
	Rect		0.256	0.671			0.192	0.684			0.253	0.672	
21	LFem	0.671	0.006	0.671	631	0.502	0.020	0.262	33	0.672	0.006	0.650	471 [52]
	RFem		0.020	0.671			0.031	0.412			0.020	0.674	
	POR		0.010	0.671			0.013	0.646			0.011	0.656	
	Blad		0.082	0.773			0.072	0.782			0.081	0.774	
	Rect		0.850	0.773			0.857	0.774			0.856	0.772	
22	LFem	0.773	0.012	0.773	1365	0.657	0.010	1.031	171	0.777	0.011	0.803	746 [51]
	RFem		0.012	0.773			0.011	0.771			0.011	0.819	
	POR		0.044	0.773			0.050	0.733			0.041	0.800	
	Blad		0.678	0.803			0.747	0.795			0.680	0.803	
	Rect		0.227	0.803			0.180	0.809			0.224	0.804	
23	LFem	0.803	0.010	0.804	1073	0.746	0.010	0.735	54	0.804	0.010	0.754	279 [23]
	RFem		0.010	0.803			0.027	0.382			0.011	0.744	
	POR		0.075	0.803			0.036	0.994			0.075	0.804	
	Blad		0.694	0.816			0.765	0.810			0.667	0.817	
	Rect		0.224	0.816			0.189	0.817			0.248	0.813	
24	LFem	0.816	0.006	0.816	2078	0.764	0.005	0.953	103	0.820	0.005	0.910	998 [49]
	RFem		0.005	0.816			0.006	0.667			0.005	0.860	
	POR		0.071	0.816			0.035	0.988			0.075	0.804	

[·] denotes the number of iterations.

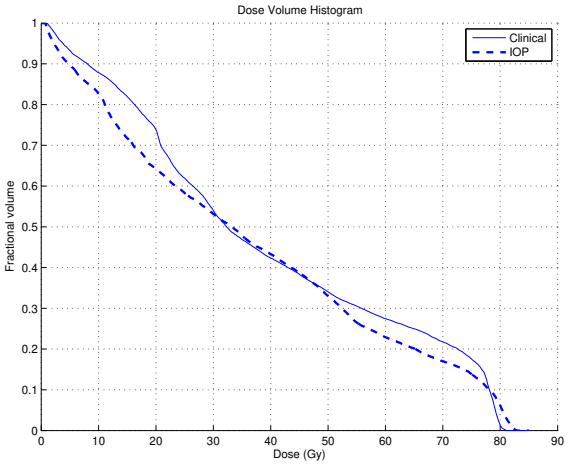

(a) Bladder

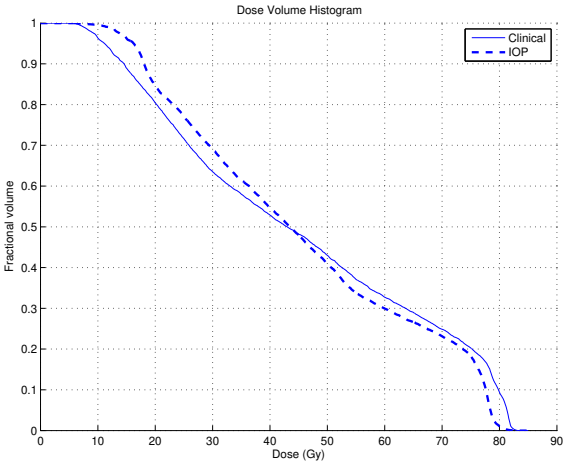

(b) Rectum

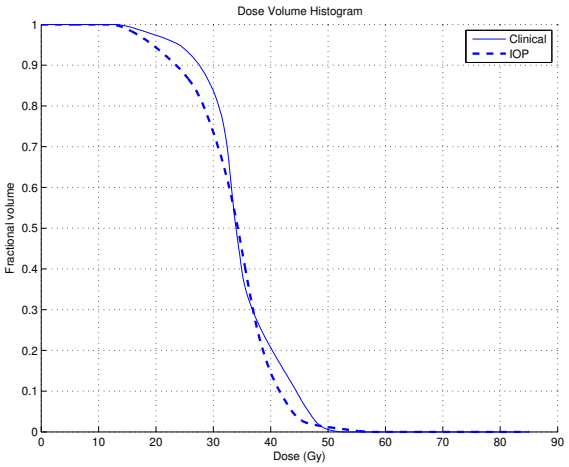

(c) Left femoral head

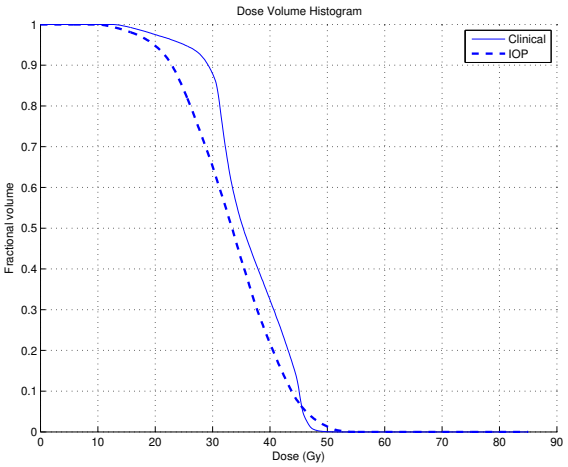

(d) Right femoral head

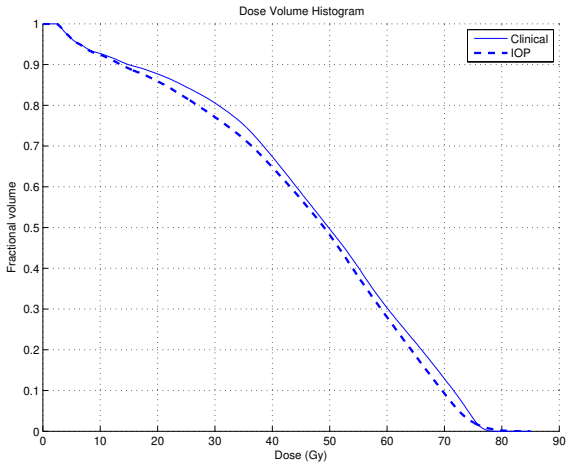

(e) POR

Figure 6.3: Comparison of DVHs from the clinical plan and a plan generated by IOP weights for patient #1

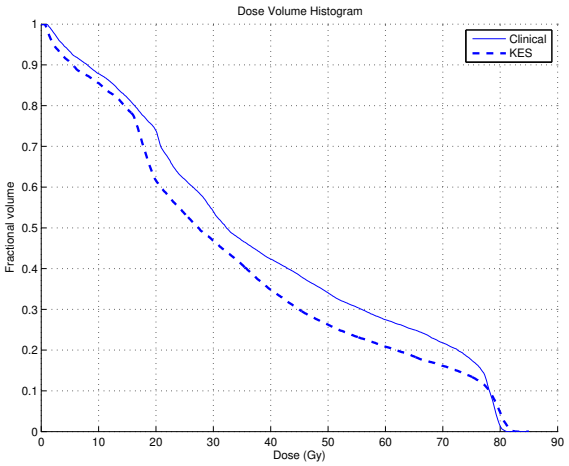

(a) Bladder

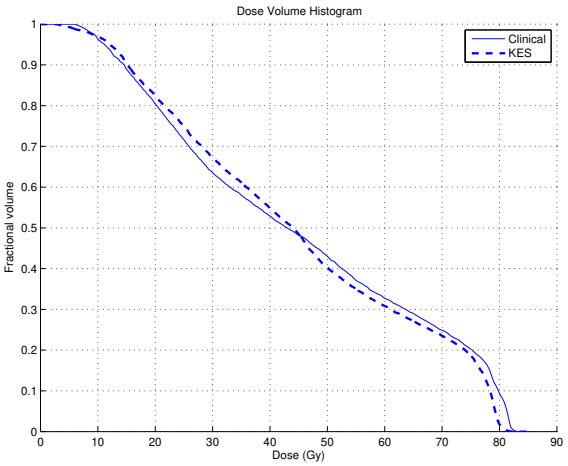

(b) Rectum

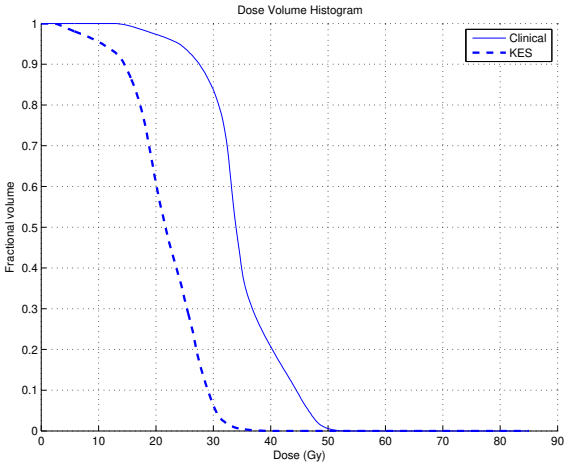

(c) Left femoral head

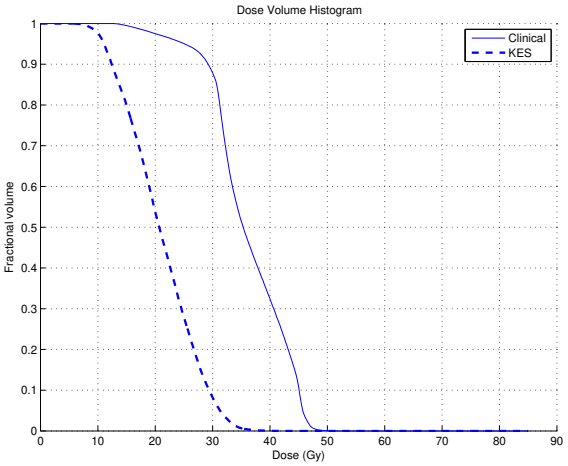

(d) Right femoral head

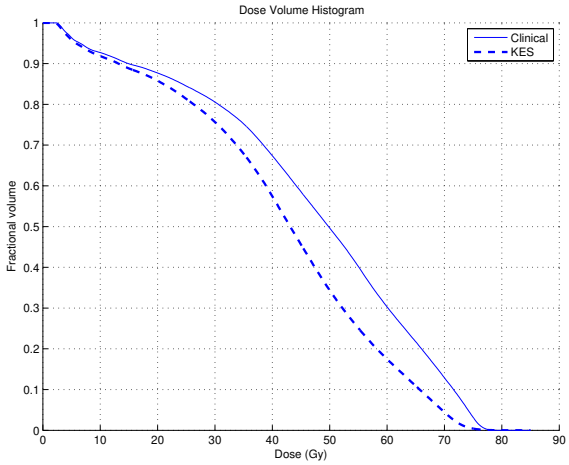

(e) POR

Figure 6.4: Comparison of DVHs from the clinical plan and a plan generated by weights from the model of Keshavarz et al. (2011) for patient #1

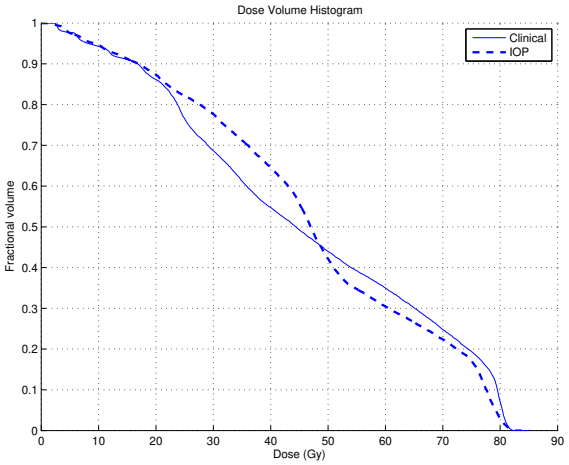

(a) Bladder

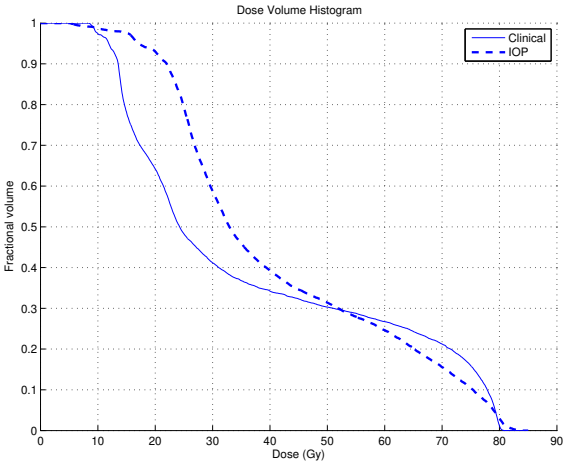

(b) Rectum

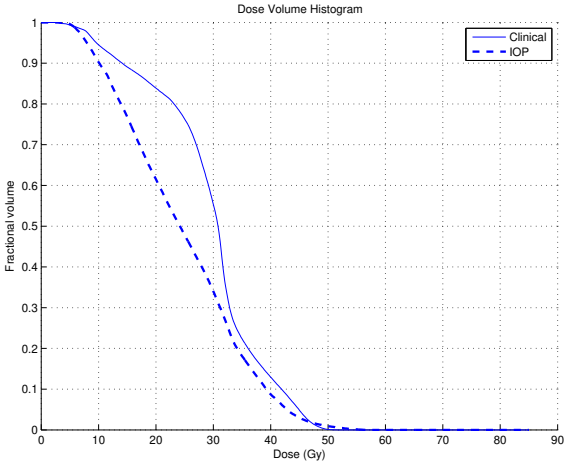

(c) Left femoral head

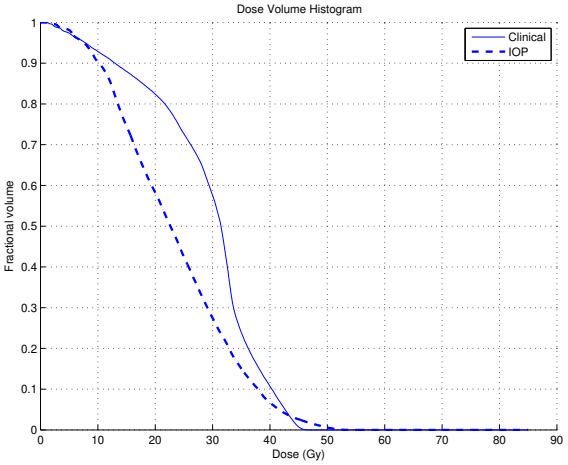

(d) Right femoral head

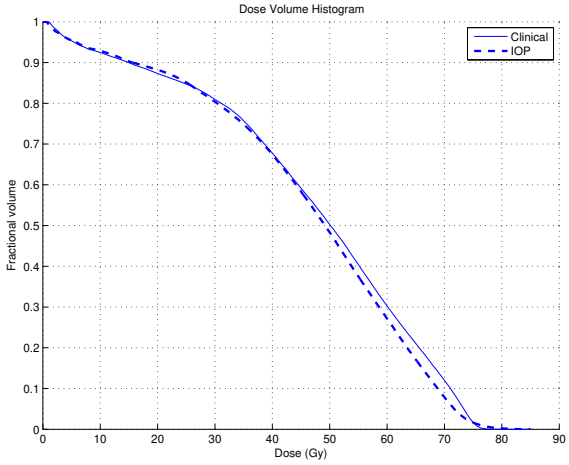

(e) POR

Figure 6.5: Comparison of DVHs from the clinical plan and a plan generated by IOP weights for patient #2

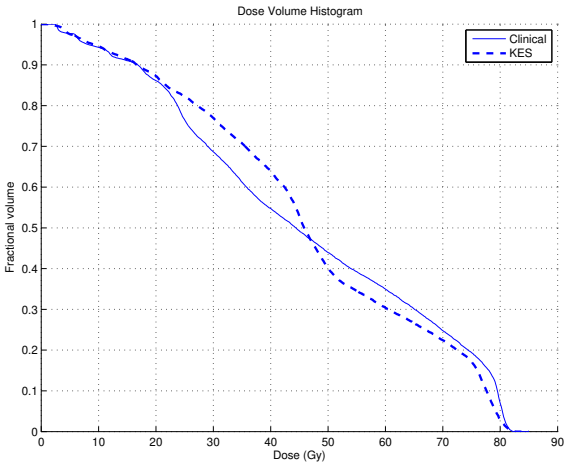

(a) Bladder

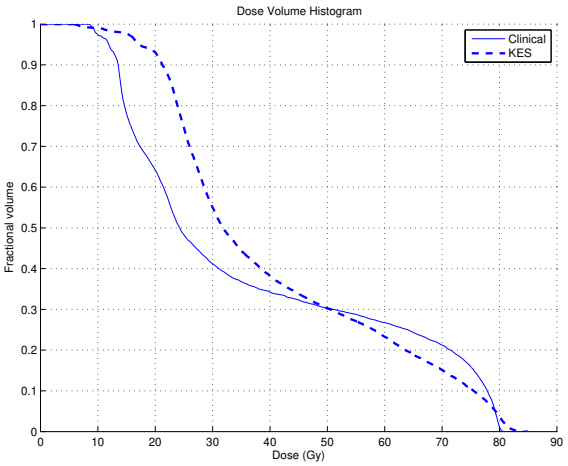

(b) Rectum

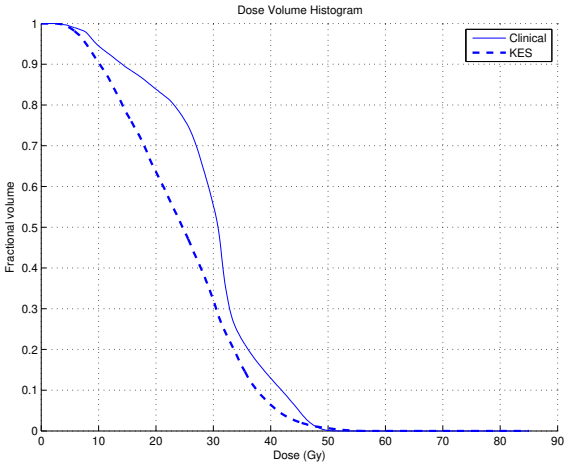

(c) Left femoral head

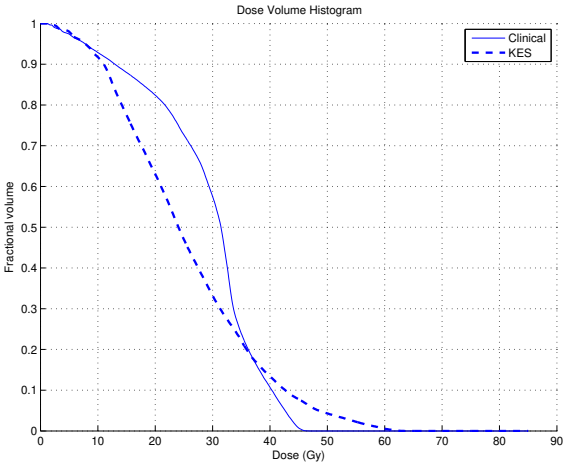

(d) Right femoral head

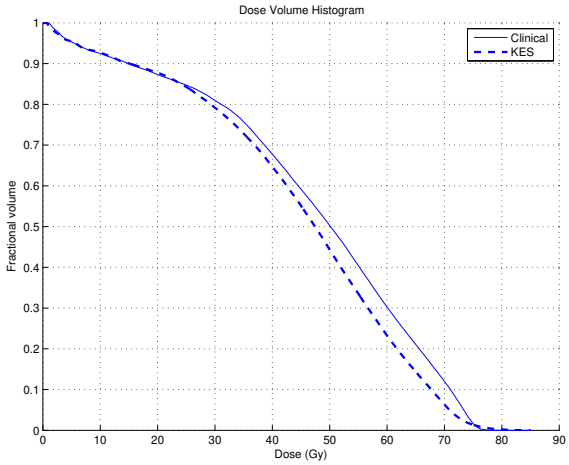

(e) POR

Figure 6.6: Comparison of DVHs from the clinical plan and a plan generated by weights from the model of Keshavarz et al. (2011) for patient #2

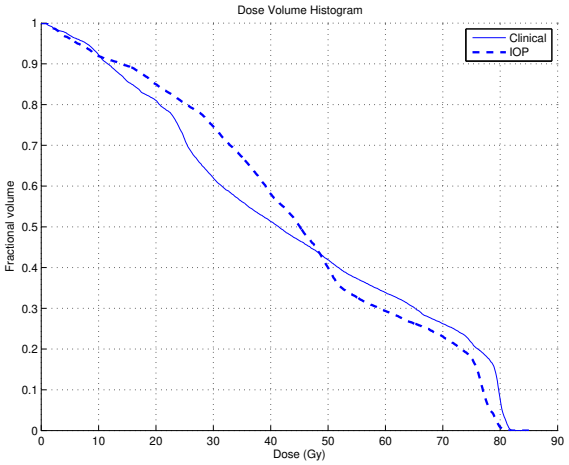

(a) Bladder

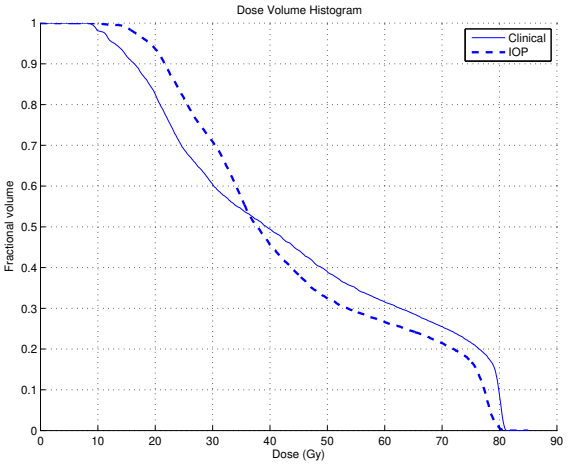

(b) Rectum

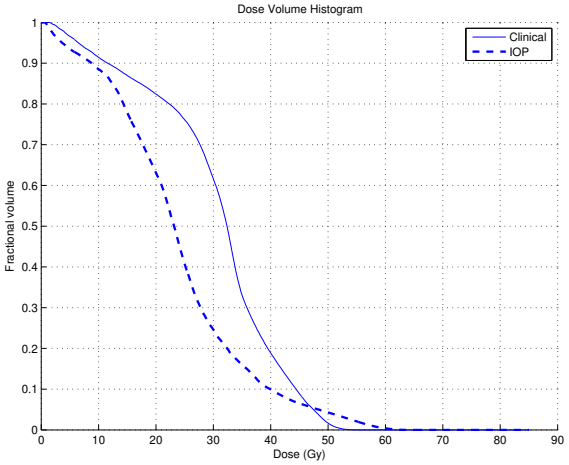

(c) Left femoral head

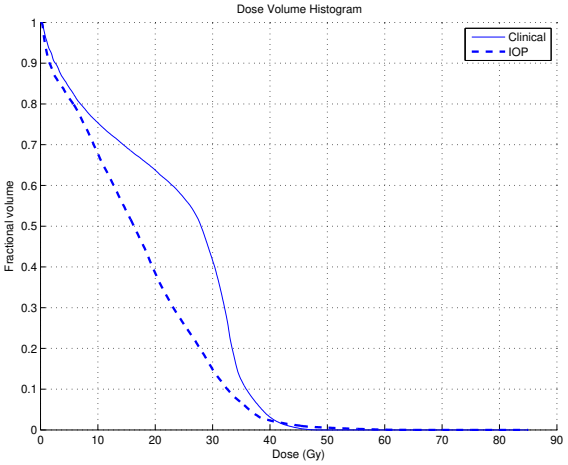

(d) Right femoral head

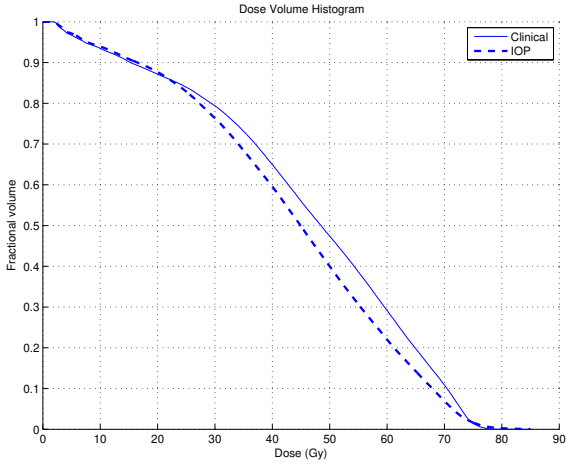

(e) POR

Figure 6.7: Comparison of DVHs from the clinical plan and a plan generated by IOP weights for patient #3

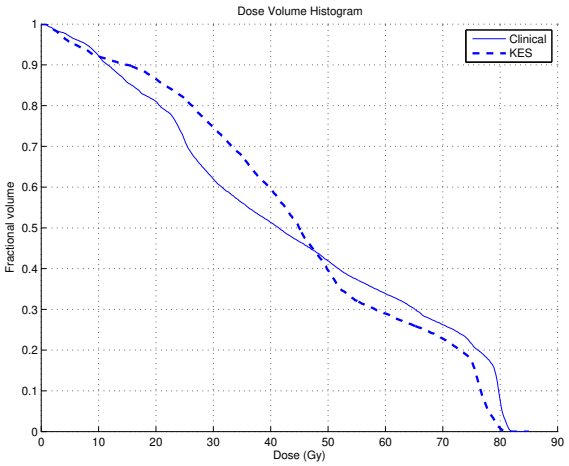

(a) Bladder

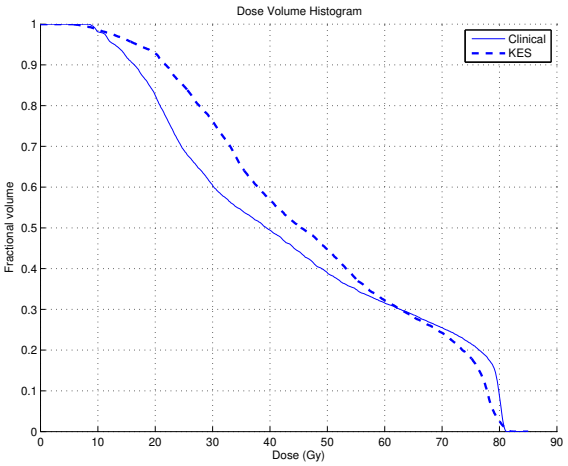

(b) Rectum

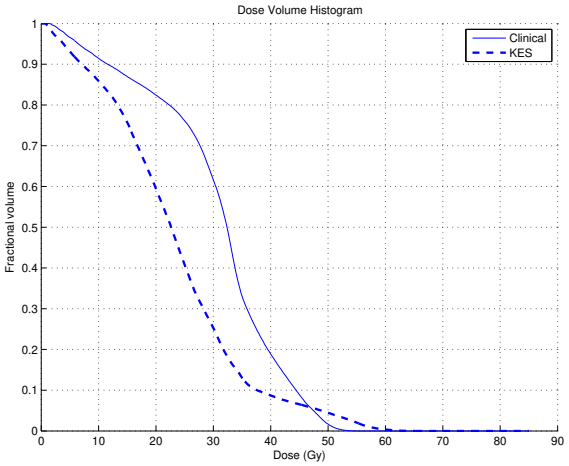

(c) Left femoral head

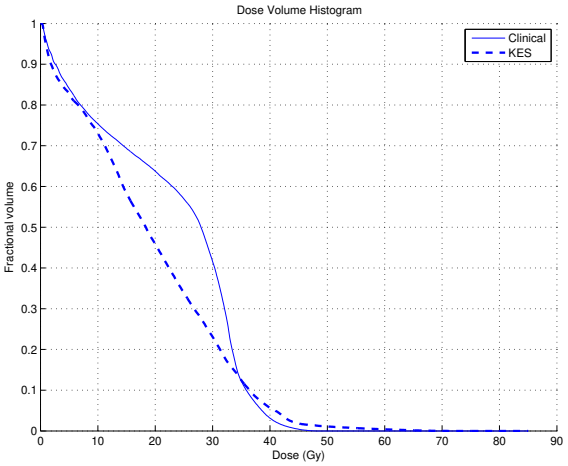

(d) Right femoral head

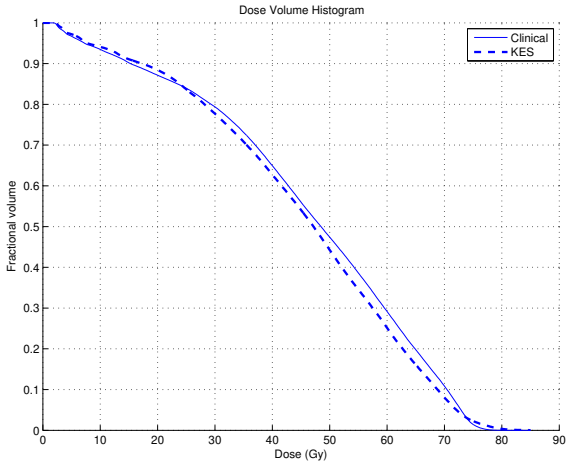

(e) POR

Figure 6.8: Comparison of DVHs from the clinical plan and a plan generated by weights from the model of Keshavarz et al. (2011) for patient #3

between the bladder and rectum objectives. However, if we fix the bladder weight, which is the highest weight for the IOP, to one in the model of Keshavarz et al. (2011), their model returns weights that are very biased towards the bladder: 0.941 (see Table 6.1). Comparing Figures 6.7 and 6.8, as a result, the DVHs from the model of Keshavarz et al. (2011) clearly lose significant amount of the rectum sparing which is not consistent with the original intention of the treatment planner revealed in the clinical DVHs.

We note that the "informed" version of Keshavarz et al. (2011), where the weight determined to be highest by the IOP model is fixed to one, also retains the preference ordering encoded by a given treatment plan to a reasonable extent. As shown in Figures 6.3–6.8, the difference between the DVHs for the IOP model and Keshavarz et al. (2011) often looks not so substantial. However, one should keep in mind that which weight is to be fixed to one in the model of Keshavarz et al. (2011) can only be determined after the IOP model is solved. There seems to be no particular way to identify which weight to fix to one without guidance of our IOP model.

Figure 6.9 shows the DVH comparison between the IOP and LIOP models, which reinforces the similarity of the results from the two models. The DVH comparison between the IOP model and the SLP algorithm is omitted because from the weights in Tables 6.3 and 6.4 it is obvious that the two models will lead to very similar treatment plans. The DVH comparison results discussed in this section were similar throughout all the remaining patients.

Figure 6.10 summarizes the performance of KES, IOP, LIOP and SLP across all 24 patients in terms of the solution time and preference preservation, the latter of which is quantified using the variance of the component-wise ratios $f_k(\mathbf{x}^*)/f_k(\hat{\mathbf{x}})$ across k. We see that the KES model solves quickly but exhibits high variance in the objective function ratios. On the other hand, the IOP model generally takes the longest to solve, but closely preserves the initial preferences. The LIOP model and the SLP algorithm by extension strike a balance between these two conflicting goals. Figure 6.11 shows the performance

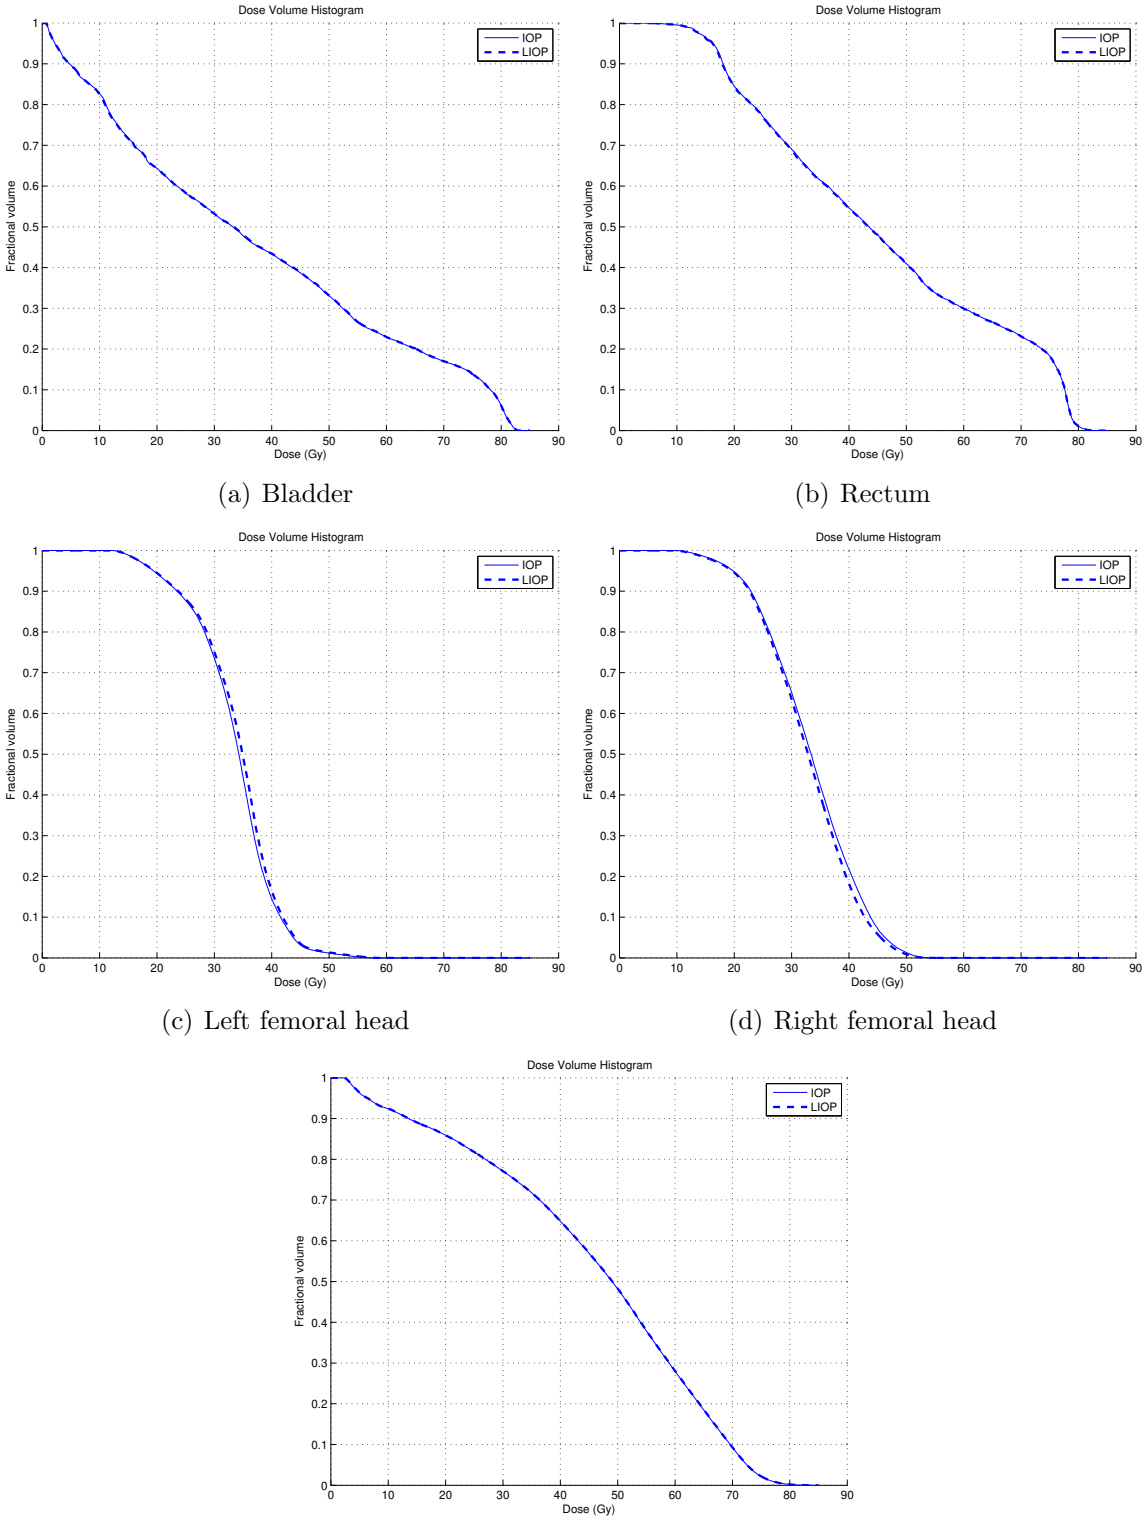

(a) Bladder

(b) Rectum

(c) Left femoral head

(d) Right femoral head

(e) POR

Figure 6.9: Comparison of DVHs from plans generated by IOP and LIOP weights for Patient #1

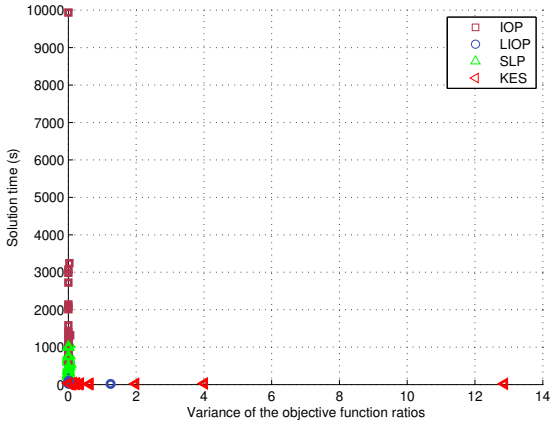

(a) Full

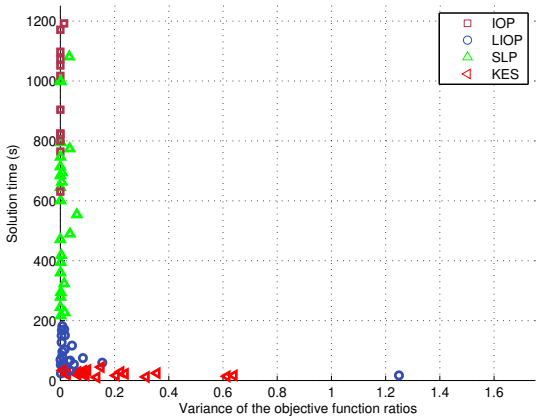

(b) Zoomed-in

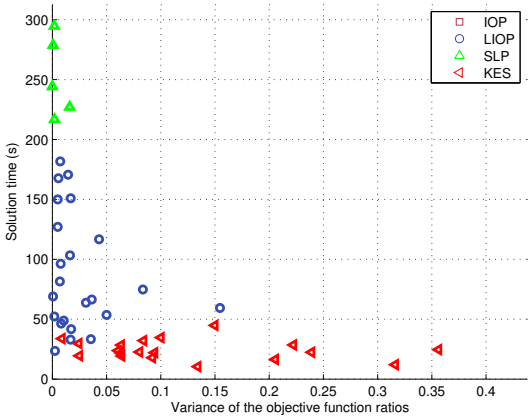

(c) Zoomed-in further

Figure 6.10: The trade-off between preference preservation (variance of the objective function ratios) and solution time

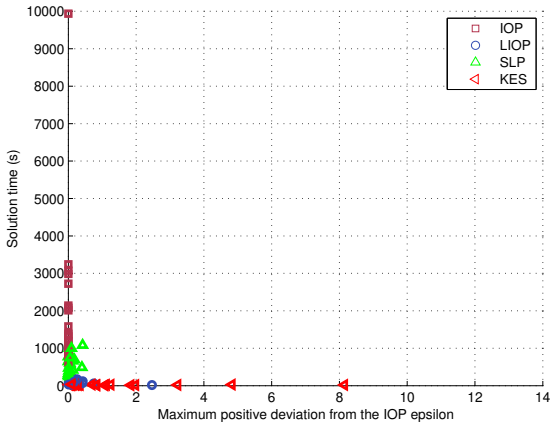

(a) Full

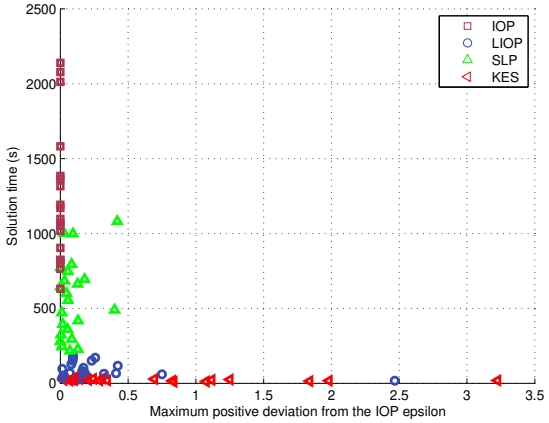

(b) Zoomed-in

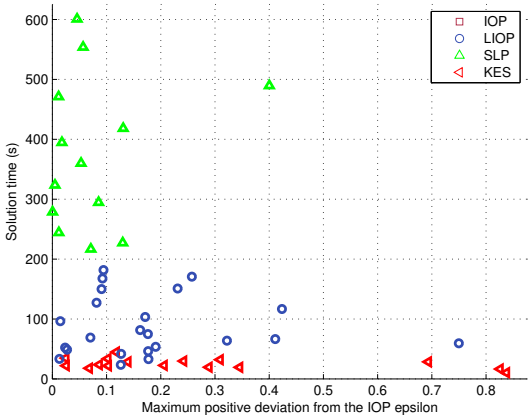

(c) Zoomed-in further

Figure 6.11: The trade-off between the maximum objective function ratio and solution time

Table 6.5: Comparison of average results from IOP, SLP, LIOP, and KES

	IOP	SLP	LIOP	KES
$\text{Var}(\frac{\mathbf{f}(\mathbf{x}^*)}{\mathbf{f}(\hat{\mathbf{x}})})$	0.004	0.009	0.076	0.938
$\text{Max}(\frac{\mathbf{f}(\mathbf{x}^*)}{\mathbf{f}(\hat{\mathbf{x}})} - \frac{\mathbf{f}(\mathbf{x}^{IOP})}{\mathbf{f}(\hat{\mathbf{x}})})$	0	0.031	0.223	1.107
$\lvert\epsilon^{\text{IOP}} - \epsilon^*\rvert$	0	0.001	0.086	N/A
$\lvert\lvert\boldsymbol{\alpha}^{\text{IOP}} - \boldsymbol{\alpha}^*\rvert\rvert_2$	0	0.007	0.075	0.279
Solution time (s)	1,883	569	85	23

KES fixes the weight determined highest by IOP to one. $\mathbf{x}^{\text{IOP}}$, ϵ^{IOP}, and $\boldsymbol{\alpha}^{\text{IOP}}$ denote the optimal $\mathbf{x}$, ϵ and $\boldsymbol{\alpha}$ obtained by IOP, respectively.

of the different models in terms of the solution time and the maximum objective function ratio, which shows similar results to Figure 6.10.

Table 6.5 summarizes the numerical results across all patients, reinforcing the trade-offs between the models. Note that the trade-offs exhibited by the SLP solutions can be adjusted so that the performance sweeps out a frontier between IOP and LIOP. For example, setting weaker termination criteria would bring the performance of the SLP algorithm towards the LIOP model.

Although a search for the best algorithm that solves the IOP quickest is out of the scope of this chapter, it is widely accepted that the successive quadratic programming (SQP) algorithm generally works better than the SLP algorithm. However, our computational experiments show that (see Table 6.6) the convergence rate of the SQP algorithm is slower than that of the SLP algorithm for our quadratic problem. The solution times could be improved by relaxing the stopping criterion, but in general there seems to be no benefit from using the SQP algorithm over the SLP algorithm.

6.5 Conclusion

Trade-offs and preferences are critical in multiobjective optimization. In this chapter, we developed a new approach to inverse convex multiobjective optimization that explicitly considers the preferences encoded by a given initial solution. Our approach is general and

Table 6.6: Results from the SQP algorithm for all patients

Pat	Organ	ϵ^*	$\boldsymbol{\alpha}^*$	$\frac{f_k(\mathbf{x}^*)}{f_k(\bar{\mathbf{x}})}$	Time (s)	Pat	Organ	ϵ^*	$\boldsymbol{\alpha}^*$	$\frac{f_k(\mathbf{x}^*)}{f_k(\bar{\mathbf{x}})}$	Time (s)
	Blad		0.014	0.812			Blad		0.426	0.692	
	Rect		0.935	0.812	2263		Rect		0.520	0.692	835
1	LFem	0.812	0.002	0.814	[4]	13	LFem	0.692	0.013	0.695	[4]
	RFem		0.000	0.814			RFem		0.005	0.692	
	POR		0.049	0.812			POR		0.036	0.693	
	Blad		0.950	0.801			Blad		0.018	0.785	
	Rect		0.003	0.797	509		Rect		0.955	0.779	267
2	LFem	0.794	0.009	1.006	[2]	14	LFem	0.771	0.001	0.896	[2]
	RFem		0.016	0.972			RFem		0.000	0.216	
	POR		0.022	0.873			POR		0.027	0.825	
	Blad		0.472	0.712			Blad		0.174	0.779	
	Rect		0.411	0.712	952		Rect		0.775	0.779	2309
3	LFem	0.712	0.008	0.712	[4]	15	LFem	0.779	0.002	0.779	[4]
	RFem		0.033	0.713			RFem		0.001	0.779	
	POR		0.076	0.712			POR		0.048	0.779	
	Blad		0.005	0.721			Blad		0.611	0.789	
	Rect		0.950	0.500	446		Rect		0.329	0.789	3115
4	LFem	0.763	0.009	0.687	[2]	16	LFem	0.789	0.002	0.789	[5]
	RFem		0.007	0.816			RFem		0.000	0.789	
	POR		0.029	0.718			POR		0.058	0.789	
	Blad		0.138	0.809			Blad		0.943	0.724	
	Rect		0.777	0.809	1580		Rect		0.020	0.725	892
5	LFem	0.809	0.001	0.814	[5]	17	LFem	0.724	0.014	0.770	[3]
	RFem		0.005	0.812			RFem		0.000	0.777	
	POR		0.079	0.809			POR		0.023	0.734	
	Blad		0.517	0.801			Blad		0.152	0.825	
	Rect		0.449	0.800	900		Rect		0.772	0.815	353
6	LFem	0.788	0.002	0.861	[2]	18	LFem	0.761	0.000	1.443	[2]
	RFem		0.005	0.819			RFem		0.000	1.272	
	POR		0.027	0.836			POR		0.076	0.879	
	Blad		0.175	0.842			Blad		0.129	0.787	
	Rect		0.782	0.842	1480		Rect		0.768	0.787	401
7	LFem	0.842	0.000	0.578	[4]	19	LFem	0.787	0.032	0.926	[3]
	RFem		0.000	0.722			RFem		0.019	0.908	
	POR		0.043	0.842			POR		0.052	0.799	
	Blad		0.477	0.805			Blad		0.908	0.717	
	Rect		0.453	0.794	1412		Rect		0.062	0.717	946
8	LFem	0.781	0.000	0.773	[2]	20	LFem	0.717	0.000	1.340	[5]
	RFem		0.000	0.641			RFem		0.004	0.717	
	POR		0.070	0.885			POR		0.026	0.717	
	Blad		0.175	0.750			Blad		0.709	0.671	
	Rect		0.701	0.750	885		Rect		0.256	0.671	637
9	LFem	0.750	0.028	0.752	[4]	21	LFem	0.671	0.006	0.677	[4]
	RFem		0.022	0.753			RFem		0.019	0.686	
	POR		0.074	0.750			POR		0.010	0.671	
	Blad		0.166	0.709			Blad		0.074	0.778	
	Rect		0.695	0.709	2331		Rect		0.872	0.770	723
10	LFem	0.709	0.000	0.186	[5]	22	LFem	0.761	0.007	1.476	[2]
	RFem		0.000	1.064			RFem		0.007	1.223	
	POR		0.139	0.709			POR		0.040	0.810	
	Blad		0.424	0.713			Blad		0.680	0.802	
	Rect		0.478	0.713	1037		Rect		0.233	0.802	803
11	LFem	0.713	0.005	0.713	[5]	23	LFem	0.802	0.009	0.841	[3]
	RFem		0.004	0.713			RFem		0.009	0.903	
	POR		0.089	0.713			POR		0.069	0.828	
	Blad		0.047	0.664			Blad		0.694	0.816	
	Rect		0.906	0.644	411		Rect		0.224	0.816	2627
12	LFem	0.642	0.003	0.708	[2]	24	LFem	0.816	0.006	0.816	[5]
	RFem		0.004	0.687			RFem		0.005	0.816	
	POR		0.040	0.672			POR		0.071	0.816	

$[\cdot]$ denotes the number of iterations.

maintains the complexity of the underlying forward problem. We formulated a general model that encompasses many of the inverse models in the literature and demonstrated how an inverse model can be adjusted to make it preference-preserving. Practical applications of this work include knowledge-based treatment planning in radiation therapy, where treatment planning preferences encoded in historical treatments can be maintained in the generation of plans for new patients.

Chapter 7

Conclusions

The traditional goal of inverse optimization is to find a cost vector that makes a given solution optimal for the underlying forward optimization formulation. However, often the forward formulation used to derive the inverse problem does not exactly match the system that was used to generate the given solution. Or, the given solution itself is often not an optimal solution to the system. The result is an ill-posed inverse optimization problem where there is no nonzero cost vector that can make the given solution optimal. This thesis addresses this situation by generalizing the traditional approach to solving the inverse problem. For any given forward optimization formulation, our generalized inverse optimization approach finds a cost vector that renders a given solution minimally suboptimal while preserving the original intention of a decision maker who generated the solution.

Due to its ability to produce a non-trivial inverse solution for any pair of a given solution and a forward formulation and measure how well the original preference is preserved through the forward formulation, our generalized inverse optimization approach identifies a compact forward formulation that can replace the complex formulation used in the actual system. Furthermore, the statistical analysis we propose for analyzing the inverse solutions can suggest how to use the identified forward formulation for the future

optimization process in an effective and standardized way. In prostate radiation therapy treatment planning, our generalized inverse method was used to determine clinical objectives that are fewer than used in practice as well as the values of objective function weights that allow the identified objectives to reproduce a clinical treatment plan. Regression analysis to relate the weights to a patient's anatomical characteristics showed the potential for inverse optimization to support automated and knowledge-based treatment planning.

Although this thesis focuses on prostate radiation therapy treatment planning as the application, our methodology can be applied to any optimization problem where data-driven parameter estimation is relevant. For example, our generalized models can replace the standard model used in many existing applications, since they specialize to the standard model when a given solution can be optimal for the underlying forward problem. Replacing the standard model with our generalized model in general inverse optimization problems will add the feature that allows them to accommodate more data with no additional cost of complexity.

Considering growing interest in big data in the operations research community, we believe that there is increasing opportunity for applications of our inverse optimization approach. The more data is available, the higher chance there is that the inverse optimization technique will more accurately and rigorously reconstruct the forward system. Our generalization to the standard inverse optimization approach amplifies this potential by robustly accommodating and learning from data regardless of its source and structure. Our approach can identify how inconsistent given data is with the underlying forward problem, which would allow us to filter out irrelevant data. Ultimately, we believe the new, generalized inverse optimization methodology we propose will open the doors for inverse optimization to better adapt and harness the power of big data.

Inverse optimization as a tool for retrieving the modeling parameters is a retrospective study. Once parameter values are inferred from historical or desirable solutions via inverse

optimization, study of how to use the inferred parameters for future decision making should follow. One should keep in mind that data should be collected from a wide range so that it can account for any changes or uncertainties that may be involved in the future decision making process. Also, it should be carefully reviewed if the system itself has improved and the existing data has become too out of date to capture new features of the system. Overall, the use of the inverse optimization technology should always be accompanied with critically selecting data and extracting meaningful value from the results.

In a broader context, we view inverse optimization also as a means of strengthening other optimization techniques by helping them actively make use of historical data and learn a decision maker's preferences. For example, by incorporating the idea of inverse optimization into Pareto surface approximation techniques for multiobjective optimization, one can explicitly incorporate the decision maker's preference into determining a search direction through the Pareto surface and identifying the most relevant area of the surface. The challenge in determining prioritization of multiple criteria in preemptive goal programming also relates to the challenge in determining parameter values. In preemptive goal programming, multiple criteria are prioritized and solved sequentially according to their relative importance, which can be quantified by inverse optimization. Lastly, inverse optimization may provide insight for how to derive uncertainty sets for a cost vector in robust optimization. Using historical solutions as input uncertainty sets may be efficiently and objectively constructed like the recent study that derives uncertainty sets from adaptive questionnaires (Bertsimas and O'Hair, 2013).

There are some specific future research directions this thesis offers. First, a new measure posed by generalizing the traditional concept of inverse optimization is how well given data "fits" the underlying forward problem. This measure is similar in nature to goodness-of-fit measures used for a statistical model. For example, given multiple input solutions and some forward problem, one may evaluate the goodness-of-fit between

the solutions and the forward problem via a measure that resembles some well-known statistic such as R^2. This measure will provide a rigorous way of identifying the forward problem that best fits the given data. Establishing such a measure is currently our ongoing research.

In radiation therapy, the weight prediction model presented in Chapter 5 can only predict one weight at a time, assuming other weights are fixed or linearly dependent on the one that is to be predicted. An extended prediction model that can predict multiple weights simultaneously will be an enabler for the proposed concept to be applied to more complex cancer sites such as head and neck, where there are typically more than 10 critical organs close to each other around the tumour. A multinomial regression approach to predicting multiple weights using patient anatomy as a predictor is currently our ongoing work. Ultimately, whether the weight prediction approach can be successfully embedded in the current knowledge-based, automated treatment planning framework should be clinically tested.

Bibliography

Abbeel, P. and A. Y. Ng (2004). Apprenticeship learning via inverse reinforcement learning. In *Proc. ICML*, pp. 1–8.

Ahuja, R. K. and J. B. Orlin (2001). Inverse optimization. *Oper. Res. 49*(5), 771–783.

Ahuja, R. K. and J. B. Orlin (2002). Combinatorial algorithms for inverse network flow problems. *Networks 40*(4), 181–187.

Alizadeh, B. and R. E. Burkard (2011). Combinatorial algorithms for inverse absolute and vertex 1-center location. *Networks 58*(3), 190–200.

Alizadeh, B., R. E. Burkard, and U. Pferschy (2009). Inverse 1-center location problems with edge length augmentation on trees. *Computing 86*(4), 331–343.

American Cancer Society (2014). *Cancer Facts and Figures 2014*. Atlanta, GA: American Cancer Society.

Arridge, S. R. (1999). Optical tomography in medical imaging. *Inverse Problems 15*(2), R41–R93.

Bazaraa, M. S., H. D. Sherali, and C. M. Shetty (2006). *Nonlinear Programming: Theory and Algorithms* (3rd ed.). John Wiley and Sons.

Benson, H. P. (1978). Existence of efficient solutions for vector maximization problems. *J. Optim. Theory Appl. 26*(4), 569–580.

Benson, H. P. (1998). An outer approximation algorithm for generating all efficient extreme points in the outcome set of a multiple objective linear programming problem. *J. Glob. Optim. 13*(1), 1–24.

Bertsimas, D. and R. Mazumder (2014). Least quantile regression via modern optimization. *Ann. Statist. 42*(6), 2494–2525.

Bertsimas, D. and A. O'Hair (2013). Learning preferences under noise and loss aversion: An optimization approach. *Oper. Res. 61*(5), 1190–1199.

Bertsimas, D. and J. Tsitsiklis (1997). *Introduction to Linear Optimization* (1st ed.). Athena Scientific.

Bortfeld, T. (2006). IMRT: A review and preview. *Phys. Med. Biol. 51*(13), R363–R379.

Boyd, S. and L. Vandenberghe (2004). *Convex Optimization*. New York, NY: Cambridge University Press.

Burton, D. and P. L. Toint (1992). On an instance of the inverse shortest paths problem. *Math. Programming 53*(1), 45–61.

Cai, M., X. Yang, and J. Zhang (1999). The complexity analysis of the inverse center location problem. *J. Glob. Optim. 15*(2), 213–218.

Canadian Cancer Society's Advisory Committee on Cancer Statistics (2014). *Canadian Cancer Statistics 2014*. Toronto, ON: Canadian Cancer Society.

Chan, T. C. Y., T. Craig, T. Lee, and M. B. Sharpe (2014). Generalized inverse multiobjective optimization with application to cancer therapy. *Oper. Res. 62*(3), 680–695.

Chanyavanich, V., S. K. Das, W. R. Lee, and J. Y. Lo (2011). Knowledge-based IMRT treatment planning for prostate cancer. *Med. Phys. 38*(5), 2515–2522.

Chung, Y. and M. Damange (2008). Some inverse traveling salesman problems. *Electronic Notes in Discrete Mathematics 30*, 9–14.

Cotrutz, C., M. Lahanas, C. Kappas, and D. Baltas (2001). A multiobjective gradient-based dose optimization algorithm for external beam conformal radiotherapy. *Phys. Med. Biol. 46*(8), 2161–2175.

Cotrutz, C. and L. Xing (2002). Using voxel-dependent importance factors for interactive DVH-based dose optimization. *Phys. Med. Biol. 47*(10), 1659–1669.

Craft, D. (2011). A guide to using multi-criteria optimization (MCO) for IMRT planning in RayStation. Technical report, Department of Radiation Oncology, Massachusetts General Hospital.

Craft, D., T. Halabi, H. A. Shih, and T. Bortfeld (2007). An approach for practical multiobjective IMRT treatment planning. *Int. J. Radiat. Oncol. Biol. Phys. 69*(5), 1600–1607.

Craft, D., P. Suss, and T. Bortfeld (2007). The tradeoff between treatment plan quality and required number of monitor units in inensity-modulated radioherapy. *Int. J. Radiat. Oncol. Biol. Phys. 67*(5), 1596–1605.

Craft, D. L., T. F. Halabi, H. A. Shih, and T. R. Bortfeld (2006). Approximating convex pareto surfaces in multiobjective radiotherapy planning. *Med. Phys. 33*(9), 3399–3407.

Dale, E. and D. R. Olsen (1997). Specification of the dose to organs at risk in external beam radiotherapy. *Acta. Oncol. 36*(2), 129–135.

Darboux, G. (1894). *Leçons sur la théorie générale des surfaces et les applications géométriques du calcul infinitésimal*. Gauthier-Villars.

Das, I. and J. E. Dennis (1998). Normal-boundary intersection: A new method for

generating the Pareto surface in nonlinear multicriteria optimization problems. *SIAM J. Optim. 8*(3), 631–657.

Deasy, J. O., A. I. Blanco, and V. H. Clark (2003). CERR: A computational environment for radiotherapy research. *Med. Phys. 30*(5), 979–985.

Duan, Z. and L. Wang (2011). Heuristic algorithms for the inverse mixed integer linear programming problem. *J. Glob. Optim. 51*(3), 463–471.

Duin, C. W. and A. Volgenant (2006). Some inverse optimization problems under the Hamming distance. *Eur. J. Oper. Res. 170*(3), 887–899.

Ehrgott, M. (2005). *Multicriteria Optimization* (2nd ed.). Springer.

Ehrgott, M., L. Shao, and A. Schöbel (2011). An approximation algorithm for convex multi-objective programming problems. *J. Glob. Optim. 50*(3), 397–416.

Erkin, Z., M. D. Bailey, L. M. Maillart, A. J. Schaefer, and M. S. Roberts (2010). Eliciting patients' revealed preferences: An inverse Markov decision process approach. *Decision Analysis 7*(4), 358–365.

Foroudi, F., S. Tyldesley, L. Barbera, J. Huang, and W. J. Mackillop (2003). Evidence-based estimate of appropriate radiotherapy utilization rate for prostate cancer. *Int. J. Radiat. Oncol. Biol. Phys. 55*(1), 51–63.

Hamacher, H. W. and K. H. Küfer (2002). Inverse radiation therapy planning – a multiple objective optimization approach. *Discrete Appl. Math. 118*(1–2), 145–161.

He, Y., B. Zhang, and E. Yao (2005). Weighted inverse minimum spanning tree problems under Hamming distance. *J. Comb. Optim. 9*, 91–100.

Heuberger, C. (2004). Inverse combinatorial optimization: A survey on problems, methods, and results. *J. Comb. Optim. 8*(3), 329–361.

Hochbaum, D. S. (2003). Efficient algorithms for the inverse spanning-tree problem. *Oper. Res. 51*(5), 785–797.

Huang, S. (2005). Inverse problems of some NP-complete problems. In *Lecture Notes in Computer Science*, Volume 3521, pp. 422–426. Heidelberg: Springer.

Iyengar, G. and W. Kang (2005). Inverse conic programming with applications. *Oper. Res. Lett. 33*(3), 319–330.

Jeraj, R., P. J. Keall, and J. V. Siebers (2002). The effect of dose calculation accuracy on inverse treatment planning. *Phys. Med. Biol. 47*(3), 391–407.

Jiang, Y., L. Liu, and E. Yao (2010). Inverse minimum cost flow problems under the weighted Hamming distance. *Eur. J. Oper. Res. 207*(1), 50–54.

Keshavarz, A., Y. Wang, and S. Boyd (2011). Imputing a convex objective function. In *2011 IEEE International Symposium on Intelligent Control (ISIC)*, pp. 613–619.

Lee, E. K., T. Fox, and I. Crocker (2003). Integer programming applied to intensity-modulated radiation therapy treatment planning. *Ann. Oper. Res. 119*, 165–181.

Liu, L. and J. Zhang (2006). Inverse maximum flow problems under the weighted Hamming distance. *J. Comb. Optim. 12*(4), 395–408.

Lu, R., R. J. Radke, L. Hong, C.-S. Chui, J. Xiong, E. Yorke, and A. Jackson (2006). Learning the relationship between patient geometry and beam intensity in breast intensity-modulated radiotherapy. *IEEE Trans. Biomed. Eng. 53*(5), 908–920.

Marks, L. B. (1996). The impact of organ structure on radiation response. *Int. J. Radiat. Oncol. Biol. Phys. 34*(5), 1165–1171.

Moore, K. L., R. S. Brame, D. A. Low, and S. Mutic (2011). Experience-based quality control of clinical intensity-modulated radiotherapy planning. *Int. J. Radiat. Oncol. Biol. Phys. 81*(2), 545–551.

Niemierko, A. (1997). Reporting and analyzing dose distributions: A concept of equivalent uniform dose. *Med. Phys. 24*(1), 103–110.

Preciado-Walters, F., R. Rardin, M. Langer, and V. Thai (2004). A coupled column generation, mixed integer approach to optimal planning of intensity modulated radiation therapy for cancer. *Math. Programming 101*(2), 319–338.

R Core Team (2013). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.

Ramachandran, D. and E. Amir (2007). Bayesian inverse reinforcement learning. In *Proc. IJCAI*, pp. 2586–2591.

Ratliff, N. D., J. A. Bagnell, and M. A. Zinkevich (2006). Maximum margin planning. In *Proc. ICML*, pp. 729–736.

Romeijn, H. E., R. K. Ahuja, J. F. Dempsey, and A. Kumar (2006). A new linear programming approach to radiation therapy treatment planning problems. *Oper. Res. 54*(2), 201–216.

Romeijn, H. E., J. F. Dempsey, and J. G. Li (2004). A unifying framework for multicriteria fluence map optimization models. *Phys. Med. Biol. 49*(10), 1991–2013.

Schaefer, A. J. (2009). Inverse integer programming. *Optimization Letters 3*(4), 483–489.

Shao, L. and M. Ehrgott (2008). Approximately solving multiobjective linear programmes in objective space and an application in radiotherapy treatment planning. *Math. Method. Oper. Res. 68*(2), 257–276.

Shepard, D. M., M. C. Ferris, G. H. Olivera, and T. R. Mackie (1999). Optimizing the delivery of radiation therapy to cancer patients. *SIAM Review 41*(4), 721–744.

Sokkalingam, P. T., R. K. Ahuja, and J. B. Orlin (1999). Solving inverse spanning tree problems through network flow techniques. *Oper. Res. 47*(2), 291–298.

Tarantola, A. (1987). *Inverse Problem Theory: Methods for Data Fitting and Model Parameter Estimation*. Elsevier.

Thieke, C., T. Bortfeld, and K. Küfer (2002). Characterization of dose distributions through the max and mean dose concept. *Acta. Oncol. 41*(2), 158–161.

Troutt, M. D., A. A. Brandyberry, C. Sohn, and S. K. Tadisina (2008). Linear programming system identification: The general nonnegative parameters case. *Eur. J. Oper. Res. 185*(1), 63–75.

Troutt, M. D., W. Pang, and S. Hou (2006). Behavioral estimation of mathematical programming objective function coefficient. *Management Sci. 52*(3), 422–434.

Walpole, R. E., R. H. Myers, S. L. Myers, and K. Ye (2006). *Probability and Statistics for Engineers and Scientists* (8 ed.). Prentice Hall.

Wang, L. (2009). Cutting plane algorithms for the inverse mixed integer linear programming problem. *Oper. Res. Lett. 37*(2), 114–116.

Webb, S. (1994). Optimizing the planning of intensity-modulated radiotherapy. *Phys. Med. Biol. 39*, 2229–2246.

Wei, Q., J. Zhang, and X. Zhang (2000). An inverse DEA model for inputs/outputs estimate. *Eur. J. Oper. Res. 121*(1), 151–163.

Wu, B., F. Ricchetti, G. Sanguineti, M. Kazhdan, P. Simari, M. Chuang, R. Taylor, R. Jacques, and T. McNutt (2009). Patient geometry-driven information retrieval for IMRT treatment plan quality control. *Med. Phys. 36*(12), 5497–5505.

Wu, B., F. Ricchetti, G. Sanguineti, M. Kazhdan, P. Simari, R. Jacques, R. Taylor, and T. McNutt (2011). Data-driven approach to generating achievable dose-volume histogram objectives in intensity-modulated radiotherapy planning. *Int. J. Radiat. Oncol. Biol. Phys. 79*(4), 1241–1247.

Wu, C., G. H. Olivera, R. Jeraj, H. Keller, and T. R. Mackie (2003). Treatment plan modification using voxel-based weighting factors/dose prescription. *Phys. Med. Biol. 48*(15), 2479–2491.

Xiao, X., L. Zhang, and J. Zhang (2009). A smoothing newton method for a type of inverse semi-definite quadratic programming problem. *J. Comput. Appl. Math. 223*(1), 485–498.

Xing, L., J. G. Li, S. Donaldson, Q. T. Le, and A. L. Boyer (1999). Optimization of importance factors in inverse planning. *Phys. Med. Biol. 44*(10), 2525–2536.

Yan, D., F. Vicini, J. Wong, and A. Martinez (1997). Adaptive radiation therapy. *Phys. Med. Biol. 42*(1), 123–132.

Yang, C. and J. Zhang (1999). Two general methods for inverse optimization problems. *Appl. Math. Lett. 12*, 69–72.

Yang, C., J. Zhang, and Z. Ma (1997). Inverse maximum flow and minimum cut problems. *Optimization 40*(2), 147–170.

Yang, X. (1997). Complexity of partial inverse assignment problem and partial inverse cut problem. *RAIRO Oper. Res. 35*(1), 117–126.

Yang, Y., E. C. Ford, B. Wu, M. Pinkawa, B. van Triest, P. Campbell, D. Y. Song, and T. R. McNutt (2013). An overlap-volume-histogram based method for rectal dose prediction and automated treatment planning in the external beam prostate radiotherapy following hydrogel injection. *Med. Phys. 40*(1), 011709–1–011709–10.

Yu, Y. (1997). Multiobjective decision theory for computational optimization in radiation therapy. *Med. Phys. 24*(9), 1445–1454.

Zelefsky, M. J., H. Chan, M. Hunt, Y. Yamada, A. M. Shippy, and H. Amols (2006).

Long-term outcome of high dose intensity modulated radiation therapy for patients with clinically localized prostate cancer. *J. Urol. 176*(4), 1415–1419.

Zeleny, M. (1974). *Linear Multiobjective Programming*. New York, NY: Springer Verlag.

Zeleny, M. (1986). Optimal system design with multiple criteria: De Novo programming approach. *Engineering Costs and Production Economics 10*, 89–94.

Zeleny, M. (2010). Multiobjective optimization, systems design and De Novo programming. In C. Zopounidis and P. M. Pardalos (Eds.), *Handbook of Multicriteria Analysis*.

Zhang, B., J. Zhang, and Y. He (2005). The center location improvement problem under the Hamming distance. *J. Comb. Optim. 9*(2), 187–198.

Zhang, B., J. Zhang, and Y. He (2006). Constrained inverse minimum spanning tree problems under the bottleneck-type Hamming distance. *J. Glob. Optim. 34*(3), 467–474.

Zhang, J. and M. Cai (1998). Inverse problem of minimum cuts. *Math. Method. Oper. Res. 47*(1), 51–58.

Zhang, J., N.-H. Kim, and L. Lasdon (1985). An improved successive linear programming algorithm. *Management Sci. 31*(10), 1312–1331.

Zhang, J. and Z. Liu (1996). Calculating some inverse linear programming problems. *J. Comput. Appl. Math. 72*(2), 261–273.

Zhang, J. and Z. Liu (1999). A further study on inverse linear programming problems. *J. Comput. Appl. Math. 106*, 345–359.

Zhang, J. and Z. Liu (2002). A general model of some inverse optimization problems and its solution method under L_∞ norm. *J. Comb. Optim. 6*, 207–227.

Zhang, J., Z. Ma, and C. Yang (1995). A column generation method for inverse shortest path problems. *Math. Method. Oper. Res. 41*(3), 347–358.

Zhang, J. and C. Xu (2010). Inverse optimization for linearly constrained convex separable programming problems. *Eur. J. Oper. Res. 200*(3), 671–679.

Zhang, J., L. Zhang, and X. Xiao (2010). A perturbation approach for an inverse quadratic programming problem. *Math. Method. Oper. Res. 72*(3), 379–404.

Zhu, L., L. Lee, Y. Ma, Y. Ye, R. Mazzeo, and L. Xing (2008). Using total-variation regularization for intensity modulated radiation therapy inverse planning with field-specific numbers of segments. *Phys. Med. Biol. 53*(23), 6653–6672.

Zhu, X., Y. Ge, T. Li, D. Thongphiew, F.-F. Yin, and Q. J. Wu (2011). A planning quality evaluation tool for prostate adaptive IMRT based on machine learning. *Med. Phys. 38*(2), 719–726.

Ziebart, B. D., A. Maas, J. A. Bagnell, and A. K. Dey (2008). Maximum entropy inverse reinforcement learning. In *Proc. AAAI*, pp. 1433–1438.