Generalized inverse multi-objective optimization with application to cancer therapy

Timothy C. Y. Chan, Tim Craig, Taewoo Lee, Michael B. Sharpe

Submitted to *Operations Research*
manuscript (Please, provide the mansucript number!)

Generalized Inverse Multi-Objective Optimization with Application to Cancer Therapy

Timothy C. Y. Chan
Department of Mechanical and Industrial Engineering, University of Toronto, 5 King's College Road, Toronto, Canada, ON M5S 3G8, tcychan@mie.utoronto.ca

Tim Craig
Radiation Medicine Program, Princess Margaret Cancer Centre, University of Toronto, 610 University of Avenue, Toronto, Canada, ON M5G 2M9, Tim.Craig@rmp.uhn.on.ca

Taewoo Lee
Department of Mechanical and Industrial Engineering, University of Toronto, 5 King's College Road, Toronto, Canada, ON M5S 3G8, taewoo.lee@utoronto.ca

Michael B. Sharpe
Radiation Medicine Program, Princess Margaret Cancer Centre, University of Toronto, 610 University of Avenue, Toronto, Canada, ON M5G 2M9, Michael.Sharpe@rmp.uhn.on.ca

We generalize the standard method of solving inverse optimization problems to allow for the solution of inverse problems that would otherwise be ill-posed or infeasible. In multi-objective linear optimization, given a solution that is not a weakly efficient solution to the forward problem, our method generates objective function weights that make the given solution a near-weakly efficient solution. Our generalized inverse optimization model specializes to the standard model when the given solution is weakly efficient and retains the complexity of the underlying forward problem. We provide a novel interpretation of our inverse formulation as the dual of the well-known Benson's method, and by doing so, develop a new connection between inverse optimization and Pareto surface approximation techniques. We apply our method to prostate cancer data obtained from Princess Margaret Cancer Centre in Toronto, Canada. We demonstrate that clinically acceptable treatments can be generated using a small number of objective functions and inversely optimized weights – current treatments are designed using a complex formulation with a large parameter space in a trial-and-error re-optimization process. We also show that our method can identify objective functions that are most influential in treatment plan optimization.

Key words: inverse optimization; multi-objective linear optimization; radiation therapy treatment planning

1. Introduction

An inverse problem takes as input the actual observations from a system and determines parameter values for a model that describe the system dynamics and are consistent with the original observa-

tions. Analogously, an inverse optimization problem determines the values of the cost coefficients given a feasible solution to the original "forward" optimization problem. For example, given a feasible solution $\hat{\mathbf{x}}$ to the forward linear program $\text{LP}(\hat{\mathbf{c}}) = \min_{\mathbf{x}}\{\hat{\mathbf{c}}'\mathbf{x} \,|\, \mathbf{A}\mathbf{x} \geq \mathbf{b}\}$, the inverse optimization problem seeks to find a cost vector $\mathbf{c}$ such that $\hat{\mathbf{x}}$ is optimal for $\text{LP}(\mathbf{c})$ (Ahuja and Orlin 2001). Determining such a vector $\mathbf{c}$ amounts to finding a feasible solution $(\mathbf{p}, \mathbf{c})$ to the system $\mathbf{p}'\mathbf{b} = \mathbf{c}'\hat{\mathbf{x}}$ (strong duality), $\mathbf{p}'\mathbf{A} = \mathbf{c}'$, $\mathbf{p} \geq \mathbf{0}$ (dual feasibility). Alternatively, one can use the complementary slackness conditions in place of strong duality, as was done in Ahuja and Orlin (2001). To choose between potentially multiple $\mathbf{c}$ satisfying the system, one may add an objective function that minimizes $||\mathbf{c} - \hat{\mathbf{c}}||_1$ or $||\mathbf{c} - \hat{\mathbf{c}}||_\infty$. Using either objective function, the above inverse optimization problem is a linear program.

We consider an inverse optimization problem where the given solution $\hat{\mathbf{x}}$ is not a candidate to be an optimal solution to the forward problem. One possibility is that $\hat{\mathbf{x}}$ is an interior point of a full-dimensional feasible region. Such an input would cause the standard inverse model to return the (uninteresting) zero vector, or would render the problem infeasible should the zero vector be excluded from the inverse feasible region. A second possibility is that $\hat{\mathbf{x}}$ is infeasible for the forward problem. In this paper, we develop a generalized inverse 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, 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 multi-objective linear optimization context, which has not been considered previously, and develop connections with established multi-objective optimization literature, especially Pareto surface approximation techniques. Both the multi-objective and generalization aspects of this research are motivated by a practical problem in radiation therapy treatment planning.

Intensity-modulated radiation therapy (IMRT) is an advanced cancer treatment technology that uses beams of high energy x-rays to deliver radiation to a tumor. In IMRT, each radiation beam is modeled as a set of many small beamlets, whose intensities are optimized to deliver a dose of radiation that conforms to the shape of the tumor. The resulting beamlet intensity profile is converted into a collection of deliverable aperture shapes using a multileaf collimator (MLC). Since treatments need to balance conflicting goals such as delivering a high dose to the tumor while keeping healthy tissue dose low, IMRT treatment planning is typically modeled as a multi-objective optimization problem.

The standard approach to solving a multi-objective IMRT optimization problem is to construct a single composite objective function by taking a weighted sum of the individual objectives. 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.

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 (multi-objective) optimization problems.

1.1. Relevant literature and contributions

Inverse optimization. Early research in inverse problems was driven by geophysical scientists who wanted to estimate model parameters such as transmission times of seismic waves, using actual observations of earthquake data (Tarantola 2005). Since Burton and Toint (1992) introduced the inverse shortest path problem for seismic tomography, an increasing amount of attention in the optimization community has been devoted to inverse optimization. Early work focused on inverse network and combinatorial optimization problems (e.g., Burton and Toint (1992), Zhang and Liu (1996), Hochbaum (2003)) – a comprehensive survey is provided by Heuberger (2004). Recently, inverse optimization methods for general integer programming (Schaefer 2009) and mixed integer programming (Wang 2009) have been developed. Ahuja and Orlin (2001) presented a general framework for inverse linear programming, which has been extended to other convex optimization problems (Iyengar and Kang 2005, Zhang and Xu 2010). Keshavarz et al. (2011) propose a similar idea to ours for inverse convex optimization by allowing the KKT conditions to be satisfied with some positive residuals, and focus on computational examples to demonstrate their approach.

Inverse optimization methods have been developed for a variety of applications including demand management (Carr and Lovejoy 2000), auctions (Beil and Wein 2003), production planning (Troutt et al. 2006), health care decision making (Erkin et al. 2010) and finance (Bertsimas et al. 2012).

Iterative approaches to weight determination. 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. 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 multi-objective 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.

Pareto surface approximation. Many methods exist to approximate the Pareto surface of a multi-objective optimization problem. For example, Das and Dennis (1998) developed a method to generate a "well-distributed" set of points that provide a good approximation to the Pareto surface. Benson (1998) and Shao and Ehrgott (2008) devised algorithms to find supporting hyperplanes of the Pareto surface in multi-objective linear programming. "Sandwich algorithms" to approximate the Pareto surface have also been widely studied (e.g., Klamroth et al. (2002), Craft et al. (2006), Karasakal and Köksalan (2009), Rennen et al. (2011), Bokrantz and Forsgren (2013)). Some of these methods end up having close connections with our inverse optimization method.

Our specific contributions in this paper 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 multi-objective linear optimization. We elucidate connections between inverse optimization and established multi-objective 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.

Unless otherwise indicated, omitted proofs can be found in the Appendix.

2. A Motivational Clinical Example

2.1. Prostate cancer and radiation therapy

Prostate cancer is the most commonly diagnosed cancer among men in North America (American Cancer Society 2012, Canadian Cancer Society 2012). 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 tumor, 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 tumor dose levels, which is the primary clinical challenge in treating prostate cancer currently.

2.2. Radiation therapy treatment plan optimization

Figure 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

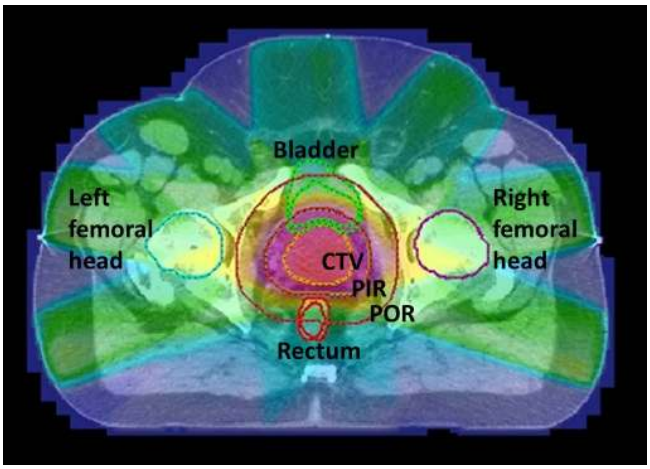

Figure 1 Anatomy of prostate site.

Table 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

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 tumor 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. 2004). In this paper, we address such criteria using a more tractable

alternative, as outlined in Section 5. Table 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 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 tumor 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.

2.3. The 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 1 is solved once. Then 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 paper is on the determination of the weights.

3. Inverse optimization methodology

The derivation of our inverse optimization methodology is based on the canonical multi-objective 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 multi-objective optimization problem can be written as

$$\begin{aligned} \text{FOP}(\boldsymbol{\alpha}): \quad \underset{\mathbf{x}}{\text{minimize}} \quad & \boldsymbol{\alpha}'\mathbf{C}\mathbf{x} \\ \text{subject to} \quad & \mathbf{A}\mathbf{x} = \mathbf{b}, \\ & \mathbf{x} \geq \mathbf{0}, \end{aligned} \tag{1}$$

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).

3.1. Preliminaries

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 (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 2). 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 (1). The following is a fundamental result of multi-objective linear programming (see Ehrgott (2005)).

PROPOSITION 1. *In a multi-objective 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 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.

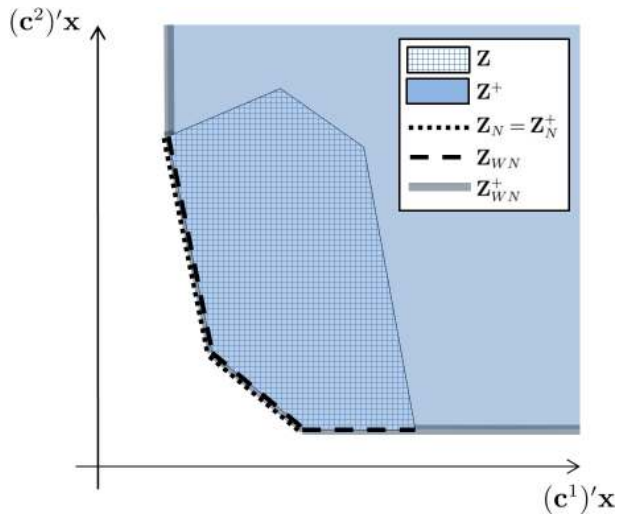

Figure 2 Geometry of criterion space.

3.2. An inverse formulation

Given $\hat{\mathbf{x}} \in \mathbf{X}$, we formulate the inverse problem associated with the forward problem (1) as

$$\begin{aligned} \text{IOP}(\mathbf{C}, \hat{\mathbf{x}}): \quad \underset{\boldsymbol{\alpha}, \mathbf{p}}{\text{minimize}} \quad & 0 \\ \text{subject to} \quad & \mathbf{A}'\mathbf{p} \leq \mathbf{C}'\boldsymbol{\alpha}, \\ & \boldsymbol{\alpha}'\mathbf{C}\hat{\mathbf{x}} = \mathbf{b}'\mathbf{p}, \\ & \boldsymbol{\alpha}'\mathbf{e} = 1, \\ & \boldsymbol{\alpha} \geq \mathbf{0}. \end{aligned} \tag{2}$$

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 (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 2, it is unlikely that a specific objective function is needed.

3.3. A generalized inverse formulation

By Proposition 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 (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 (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

$$\begin{aligned} \text{GIOP}_r(\mathbf{C}, \hat{\mathbf{x}}, \epsilon_r): \quad \underset{\boldsymbol{\alpha}, \mathbf{p}}{\text{minimize}} \quad & 0 \\ \text{subject to} \quad & \mathbf{A}'\mathbf{p} \leq \mathbf{C}'\boldsymbol{\alpha}, \\ & \boldsymbol{\alpha}'\mathbf{C}\hat{\mathbf{x}} = \epsilon_r\, \mathbf{b}'\mathbf{p}, \\ & \boldsymbol{\alpha}'\mathbf{e} = 1, \\ & \boldsymbol{\alpha} \geq \mathbf{0}. \end{aligned} \tag{3}$$

In formulation (3), $\epsilon_r > 0$ is a parameter. Formulation (3) can be solved repeatedly with varying ϵ_r until the smallest ϵ_r, denoted ϵ_r^*, is found such that (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 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* $\mathbf{C}\hat{\mathbf{x}} \notin \mathbf{Z}^+, \epsilon_r^* \leq 1$.

Note that instead of solving formulation (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^*:

$$\begin{aligned} \underset{\boldsymbol{\alpha},\mathbf{p}}{\text{minimize}} \quad & \boldsymbol{\alpha}'\mathbf{C}\hat{\mathbf{x}} \\ \text{subject to} \quad & \mathbf{A}'\mathbf{p} \leq \mathbf{C}'\boldsymbol{\alpha}, \\ & \mathbf{b}'\mathbf{p} = 1, \\ & \boldsymbol{\alpha} \geq \mathbf{0}. \end{aligned} \tag{4}$$

In formulation (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) must be normalized post-hoc. If we let $\bar{\boldsymbol{\alpha}}$ be an optimal solution to formulation (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 (3) and (4) interchangeably, depending on which is more convenient for the context. In particular, the theoretical analysis will focus on formulation (3), while the computational results will be based on formulation (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

$$\begin{aligned} \text{GIOP}_a(\mathbf{C}, \hat{\mathbf{x}}): \quad \underset{\boldsymbol{\alpha},\mathbf{p},\epsilon_a}{\text{minimize}} \quad & \epsilon_a \\ \text{subject to} \quad & \mathbf{A}'\mathbf{p} \leq \mathbf{C}'\boldsymbol{\alpha}, \\ & \boldsymbol{\alpha}'\mathbf{C}\hat{\mathbf{x}} = \mathbf{b}'\mathbf{p} + \epsilon_a, \\ & \boldsymbol{\alpha}'\mathbf{e} = 1, \\ & \boldsymbol{\alpha} \geq \mathbf{0}. \end{aligned} \tag{5}$$

Let ϵ_a^* be the optimal value of problem (5). If $\epsilon_r^* = 1$ or $\epsilon_a^* = 0$, then the two GIOP formulations become the standard inverse formulation (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.

3.4. 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 3 shows that even if $\hat{\mathbf{x}}$ is infeasible, the inverse problem can be feasible.

PROPOSITION 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($\mathbf{C}, \tilde{\mathbf{x}}$):

$$\begin{aligned} -\underset{\mathbf{y},\pi,\sigma}{\text{minimize}} \quad & \sigma \\ \text{subject to} \quad & \mathbf{A}\mathbf{y} = \pi\mathbf{b}, \\ & \mathbf{C}\mathbf{y} \leq \pi\mathbf{C}\tilde{\mathbf{x}} + \sigma\mathbf{e}, \\ & \mathbf{y} \geq \mathbf{0}. \end{aligned} \tag{6}$$

To show that IOP($\mathbf{C}, \tilde{\mathbf{x}}$) is feasible, it suffices to show that formulation (6) has a finite optimal solution. Moreover, if formulation (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 (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 (6) becomes

$$\sigma \geq \pi \max_{k \in \mathcal{K}} ((\mathbf{c}^k)'\mathbf{x} - (\mathbf{c}^k)'\tilde{\mathbf{x}}). \tag{7}$$

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 (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 (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. $\square$

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

PROPOSITION 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* $\mathbf{Cx}^* = \mathbf{C}\hat{\mathbf{x}} - \epsilon_a\mathbf{e}$.

COROLLARY 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.*

Corollary 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 $\boldsymbol{\alpha}'\mathbf{C}\hat{\mathbf{x}}/\epsilon_r \leq \boldsymbol{\alpha}'\mathbf{Cx}$ for all $\mathbf{Cx} \in \mathbf{Z}$. 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.

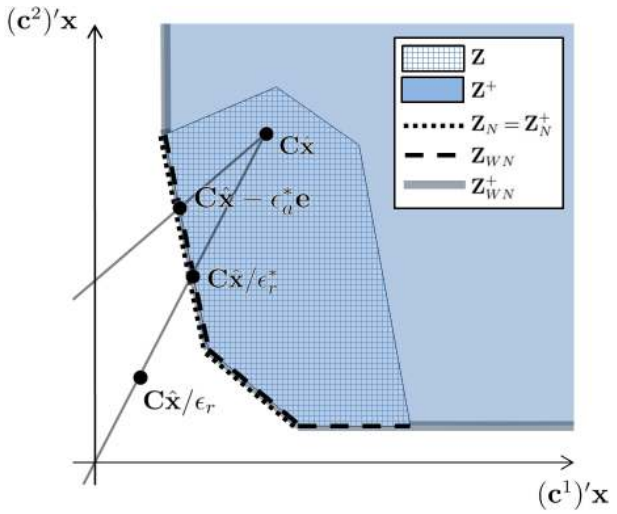

Figure 3 Adjustments of a point $\mathbf{C}\hat{\mathbf{x}}$ as an input for GIOP to a non-dominated point $\mathbf{C}\mathbf{x}^*$.

THEOREM 1. *(a) If* $(\boldsymbol{\alpha}^*, \mathbf{p}^*)$ *is an optimal solution to* $\text{GIOP}_r(\mathbf{C}, \hat{\mathbf{x}}, \epsilon_r^*)$*, then it is an optimal solution to* $\text{IOP}(\mathbf{C}, \mathbf{x}^*)$ *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* $\text{GIOP}_a(\mathbf{C}, \hat{\mathbf{x}})$*, then it is an optimal solution to* $\text{IOP}(\mathbf{C}, \mathbf{x}^*)$ *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 $\text{GIOP}_r(\mathbf{C}, \hat{\mathbf{x}}, \epsilon_r)$ and $\text{IOP}(\mathbf{C}, \mathbf{x}^*)$, by Proposition 4. Thus, if $(\boldsymbol{\alpha}^*, \mathbf{p}^*)$ is optimal for $\text{GIOP}_r(\mathbf{C}, \hat{\mathbf{x}}, \epsilon_r^*)$, it is optimal for $\text{IOP}(\mathbf{C}, \mathbf{x}^*)$. Lastly, recall that ϵ_r^* is the smallest value of ϵ_r such that $\text{GIOP}_r(\mathbf{C}, \hat{\mathbf{x}}, \epsilon_r)$ 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, $\text{IOP}(\mathbf{C}, \mathbf{x}^*)$ cannot be feasible by Proposition 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 $\mathbf{C}\bar{\mathbf{x}} = \mathbf{C}\hat{\mathbf{x}}/\bar{\epsilon}_r \in \mathbb{R}_+^{|\mathcal{K}|} \setminus \mathbf{Z}^+$ and $\text{IOP}(\mathbf{C}, \bar{\mathbf{x}})$ is feasible by Proposition 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 3). 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 $\text{GIOP}_r(\mathbf{C}, \hat{\mathbf{x}}, \epsilon_r^*)$ is also an optimal solution to $\text{IOP}(\mathbf{C}, \mathbf{x}^*)$ 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 $\text{GIOP}_a(\mathbf{C}, \hat{\mathbf{x}})$ is also optimal for $\text{IOP}(\mathbf{C}, \mathbf{x}^*)$ 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 multi-objective 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 multi-objective linear optimization, both $\mathbf{Z}$ and $\mathbf{Z}^+$ are polyhedra (Zeleny 1974). If a point $\mathbf{C}\mathbf{x} \in \mathbf{Z}^+_{WN}$ is a vertex of $\mathbf{Z}^+$, then many weight vectors $\boldsymbol{\alpha}$ can generate $\mathbf{C}\mathbf{x}$. The same is true if two or more constraints defining $\mathbf{Z}^+$ are active at $\mathbf{C}\mathbf{x}$. This observation leads to the following result, whose proof is omitted.

COROLLARY 2. *An optimal solution $\boldsymbol{\alpha}^*$ to $\text{IOP}(\mathbf{C}, \mathbf{x}^*)$ is unique if and only if $\mathbf{C}\mathbf{x}^* \in \mathbf{Z}^+_{WN}$ 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.

3.5. Validating GIOP solutions

Let $\boldsymbol{\alpha}^*$ be an optimal solution to $\text{GIOP}_r(\mathbf{C}, \hat{\mathbf{x}}, \epsilon_r^*)$. 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 $\text{FOP}(\boldsymbol{\alpha}^*)$ 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 $(\boldsymbol{\alpha}^*)'\mathbf{C}\mathbf{x}^* = (\boldsymbol{\alpha}^*)'\mathbf{C}\hat{\mathbf{x}}/\epsilon_r^*$. The same is true for

$\text{GIOP}_a(\mathbf{C}, \hat{\mathbf{x}})$. Thus, solving $\text{FOP}(\boldsymbol{\alpha}^*)$ 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 5. *(a) Let $\boldsymbol{\alpha}^*$ be an optimal solution to GIOP$_r$ formulation (4), $\epsilon_r^* = \boldsymbol{\alpha}^{*\prime}\mathbf{C}\hat{\mathbf{x}}$, and $\mathbf{y}_r^{dual}$ 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 (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 $\alpha_k^* = 0$.

Proposition 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. 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 multi-objective optimization problem. Given $\mathbf{C}\hat{\mathbf{x}} \in \mathbf{Z}$, Benson's formulation is

$$\begin{aligned} \underset{\mathbf{x},\boldsymbol{\epsilon}}{\text{maximize}} \quad & \boldsymbol{\epsilon}'\mathbf{e} \\ \text{subject to} \quad & \mathbf{C}\mathbf{x} \leq \mathbf{C}\hat{\mathbf{x}} - \boldsymbol{\epsilon}, \\ & \mathbf{A}\mathbf{x} = \mathbf{b}, \\ & \mathbf{x} \geq \mathbf{0}, \quad \boldsymbol{\epsilon} \geq \mathbf{0}. \end{aligned} \tag{8}$$

Formulation (8) 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 (8). 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 (8) is unbounded, then no efficient solutions exist.

Variants of formulation (8) 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 (8) was solved to find a supporting hyperplane of $\mathbf{Z}^+$ at $\mathbf{C}\mathbf{x}^*$. Shao and Ehrgott (2008) improved the method in Benson (1998) by using a variant of formulation (8) to find $\mathbf{C}\mathbf{x}^* \in \mathbf{Z}^+_{WN}$, instead of using univariate search. A second LP was still used to find the supporting hyperplane of $\mathbf{Z}^+$ at $\mathbf{C}\mathbf{x}^*$. Next, we draw a connection between our GIOP formulations and formulation (8), which allows us to find both $\mathbf{C}\mathbf{x}^*$ and the associated supporting hyperplane by solving a single LP.

The dual of $\text{GIOP}_a(\mathbf{C}, \hat{\mathbf{x}})$ formulation (5) can be written as

$$\begin{aligned} \underset{\mathbf{x},\sigma}{\text{maximize}} \quad & \sigma \\ \text{subject to} \quad & \mathbf{C}\mathbf{x} \leq \mathbf{C}\hat{\mathbf{x}} - \sigma\mathbf{e}, \\ & \mathbf{A}\mathbf{x} = \mathbf{b}, \\ & \mathbf{x} \geq \mathbf{0}. \end{aligned} \tag{9}$$

Notice that formulations (8) and (9) 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 $\mathbf{C}\mathbf{x}^*$ 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 2, whose proof is straightforward and omitted, formalizes these ideas.

THEOREM 2. *Let* $\hat{\mathbf{x}} \in \mathbf{X}$.

(*a*) *Let* $\boldsymbol{\epsilon} = \sigma \mathbf{C}\hat{\mathbf{x}}$ *in formulation* (8) *and let* $\boldsymbol{\lambda}_r^*$ *be an optimal dual vector associated with the first constraint in formulation* (8). *Let* $\boldsymbol{\alpha}_r^*$ *be an optimal solution to GIOP*$_r$ *formulation* (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* (8) *and let* $\boldsymbol{\lambda}_a^*$ *be an optimal dual vector associated with the first constraint in in formulation* (8). *Let* $\boldsymbol{\alpha}_a^*$ *be an optimal solution to GIOP*$_a$ *formulation* (5). *Then a solution is optimal for FOP*$(\boldsymbol{\lambda}_a^*)$ *if and only if it is optimal for FOP*$(\boldsymbol{\alpha}_a^*)$.

Theorem 2 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 (8) 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.

5. Computational Results

In this section, we demonstrate the use of our GIOP formulations in the context of radiation therapy treatment planning.

5.1. A multi-objective forward formulation

We formulate a multi-objective forward problem that emulates the clinical IMRT treatment planning formulation presented in Table 1. Since the tumor 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. We consider 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 $\mathcal{B}$ be the set of beamlets and w_b be the intensity delivered by beamlet $b \in \mathcal{B}$. 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 $\mathcal{I}$ 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. 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

$$\begin{aligned}
\underset{\mathbf{w}}{\text{minimize}} \quad & \sum_{i\in\mathcal{I}} \frac{\alpha_i}{|\mathcal{O}_i|} \sum_{v\in\mathcal{O}_i} \max\left\{0, \sum_{b\in\mathcal{B}} D_{v,b}w_b - \theta_v^i\right\} + \sum_{j\in\mathcal{J}} \alpha_j \max_{v\in\mathcal{O}_j} \sum_{b\in\mathcal{B}} D_{v,b}w_b, \\
\text{subject to} \quad & \sum_{b\in\mathcal{B}} D_{v,b}w_b \geq \ell_v, \quad \forall v \in \mathcal{T}, \\
& \sum_{b\in\mathcal{B}} D_{v,b}w_b \leq u_v, \quad \forall v \in \mathcal{V}, \\
& \frac{\beta_1}{|\mathcal{B}|}\sum_{b'\in\mathcal{B}} w_{b'} \leq \ w_b \leq \frac{\beta_2}{|\mathcal{B}|}\sum_{b'\in\mathcal{B}} w_{b'}, \quad \forall b \in \mathcal{B}, \\
& w_b \geq 0, \quad \forall b \in \mathcal{B},
\end{aligned} \tag{10}$$

where ℓ_v is a lower bound on the dose to voxel $v \in \mathcal{T}$, u_v is an upper bound on the dose to voxel $v \in \mathcal{V}$, and 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 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. 2007b). More sophisticated CVaR constraints (Romeijn et al. 2006), general piece-wise linear objective functions (Craft et al. 2007a), and linear EUD objectives (Thieke et al. 2002) can also be incorporated. However, our computational results suggest that this simple formulation is sufficient to replicate clinical-quality plans. Formulation (10) can be linearized in a standard fashion and represented as FOP($\boldsymbol{\alpha}$) formulation (1) by introducing auxiliary variables and constraints. The values of these auxiliary variables need to be determined in order to be used as input in our GIOP models. These values are uniquely determined by a given beamlet intensity profile $\hat{\mathbf{w}}$ or a dose distribution that specifies the dose to every voxel v, $d_v = \sum_{b\in\mathcal{B}} D_{v,b}\hat{w}_b$, through a direct application of the two max expressions in formulation (10).

We consider three different instances of formulation (10), 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).

In all instances, Table 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 $v \in \mathcal{V} \setminus$ POR,

Table 2 Problem sizes and solution times for the FOP and GIOP formulations (model (A)).

	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 3 Results from GIOP_r for model (A).

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

Table 4 Results from GIOP_a for model (A).

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

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.

5.2. Reproducing a clinical plan with fewer objectives

We obtained treatment plans for 12 prostate cancer patients who had previously received radiation therapy at Princess Margaret Cancer Centre. Details about these patients and their treatments are provided in Appendix B. An initialization step to process the clinical data is described in Appendix C. 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. We solved formulation (4) for GIOP_r, and (5) for GIOP_a. The sizes of the FOP, GIOP_r, and GIOP_a formulations for patient #1 are shown in Table 2 along with solution (CPU) times.

Tables 3 and 4 show the objective function weights, $\boldsymbol{\alpha}_r^*$ and $\boldsymbol{\alpha}_a^*$, that result from solving GIOP_r and GIOP_a, respectively, using patient #1's treatment as input. As discussed in Section 3.5, using

dual information to validate the weights is the most accurate method. In Table 3, using the definition of $\mathbf{x}_r^{dual}$ from Proposition 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 and Theorem 1). Similar results are seen in Table 4 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 $\mathbf{C}\mathbf{x}_a^{dual}$ are 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 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. 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 4 plots the DVHs of the CTV, PTV, and four OARs for both the clinically achievable and inversely optimized plans using weights from GIOP$_r$.

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

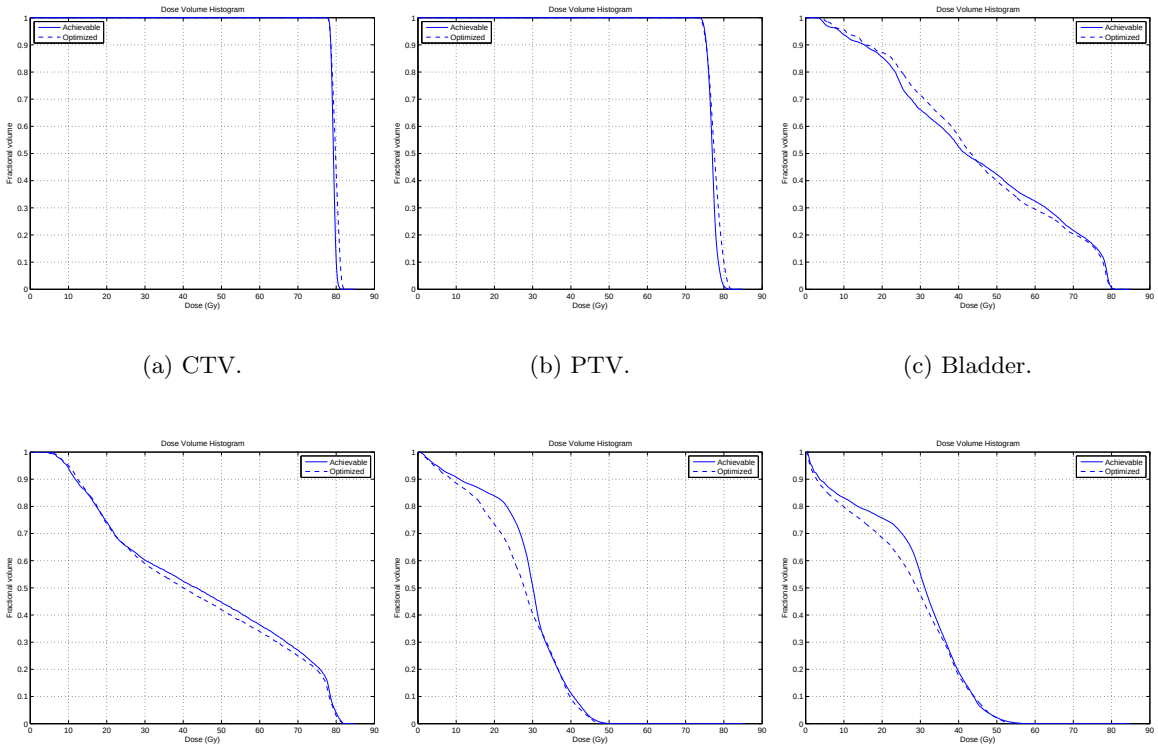

(a) CTV. (b) PTV. (c) Bladder.

(d) Rectum. (e) Left femoral head. (f) Right femoral head.

Figure 4 DVHs of the clinically achievable and inversely optimized plans for patient #1 using GIOP$_r$.

clinical plan is feasible but not on the efficient frontier of the forward problem. Figure 5 shows the similar-looking clinical and inversely optimized dose distributions on a CT image.

Lastly, we present a quantitative comparison in Table 5. 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 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 Table 3. Similar results to those shown in this section were obtained for the remaining 11 patients. A summary of the GIOP$_r$ results for all 12 patients is provided in Table 6. The GIOP$_a$ results were similar and omitted.

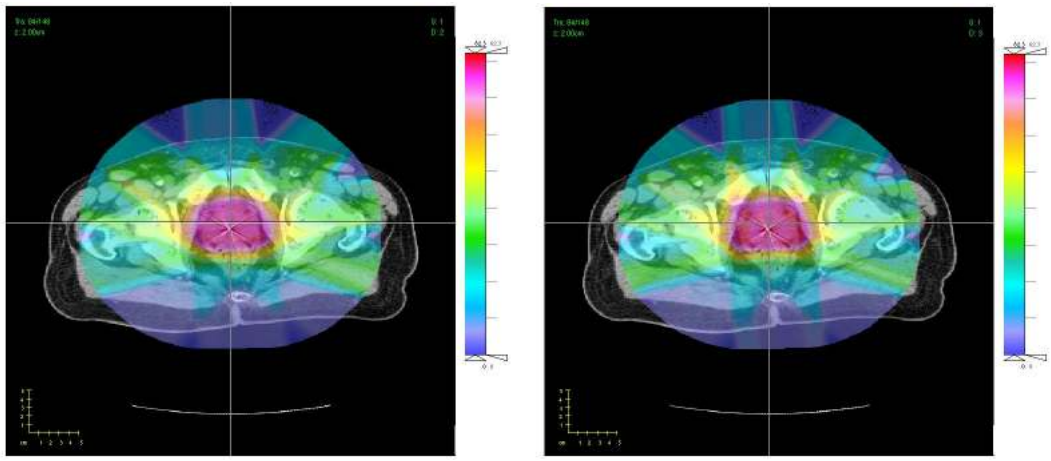

(a) Clinically achievable dose distribution. (b) Inversely optimized dose distribution.

Figure 5 Dose distributions of the clinically achievable and inversely optimized plans for patient #1 using GIOP$_r$.

Table 5 Dose metrics computed from inversely optimized plans.

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 6 Summary of results from GIOP$_r$ model (A) for all 12 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 7 Summary of results from GIOP$_r$ model (C) for all 12 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.

5.3. 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. This leads to clinical treatment planning formulations that have many more objectives than needed and 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 7. Notice that the values of ϵ_r^* in Table 7 are less than their corresponding values in Table 6. This 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 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 5 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). We believe this 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.

Table 6 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 7 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 $\theta^i = 50, 60,$ 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 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 8. Note that the values of ϵ_r^* from Table 8 are very similar to those in Table 7, 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 patient #1 are shown in Figure 6. 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 5).

Table 8 Summary of results from GIOP$_r$ model (B) for all 12 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.

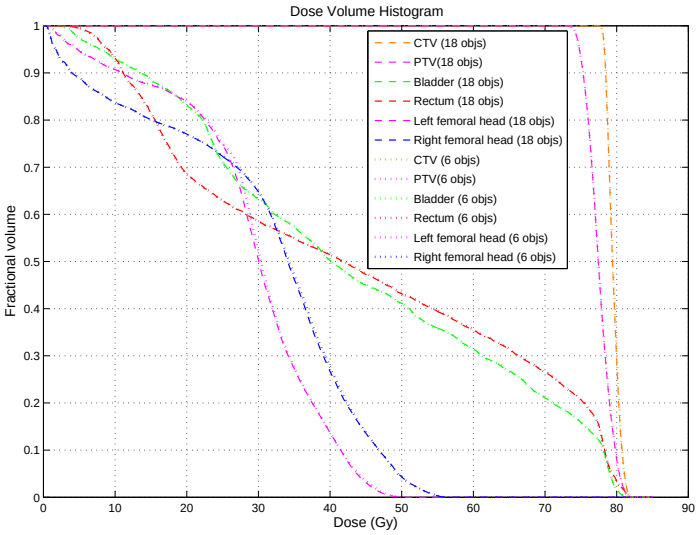

Figure 6 DVHs of inversely optimized plans from model (B) (six objectives) and model (C) (18 objectives).

5.4. 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 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 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. A useful measure of geometry that could be used is the overlap volume histogram (OVH), which measures the cumulative overlap volume between an OAR and the PTV as a function of an expanding PTV (Wu et al. 2009).

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 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 conducted a sensitivity analysis (Appendix D) and showed that having more objectives could actually increase the sensitivity of the solution. 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 results in this paper 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 5.3 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

$$\text{EUD}(\mathbf{d};a) = \left(\frac{1}{|\mathcal{O}|}\sum_{v\in\mathcal{O}} {d_v}^a\right)^{\frac{1}{a}}, \tag{11}$$

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 tumor 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 (11) for OARs (Thieke et al. 2002):

$$\text{EUD}_L(\mathbf{d};\lambda) = \lambda d^{mean} + (1-\lambda)d^{max}. \tag{12}$$

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 (11), comparable treatments can be created (Thieke et al. 2002). Note that for a tumor or target structure, d^{max} would be replaced by d^{min}, the minimum dose to the

structure. The results in Section 5.3 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.

6. Conclusion

This paper 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 multi-objective 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 multi-objective 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. For example, the solution of an inverse multi-objective 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, extending our method to general convex optimization opens the door for much broader application.

Appendix

A. Proofs

Proof of Proposition 2 Let $\mathbf{C}\hat{\mathbf{x}} \in \mathbf{Z}^+$ and $\mathbf{C}\mathbf{x}^* = \mathbf{C}\hat{\mathbf{x}}/\epsilon_r^*$. By Theorem 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 1 $\mathbf{C}\mathbf{x}^* \leq \mathbf{C}\hat{\mathbf{x}}$, and therefore $\epsilon_r^* \leq 1$. □

Proof of Proposition 4 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 $\boldsymbol{\alpha}'\mathbf{C}\mathbf{x} = \mathbf{b}'\mathbf{p}$ for $\mathbf{C}\mathbf{x} = \mathbf{C}\hat{\mathbf{x}}/\epsilon_r$. Since all other constraints are the same between the formulations $\text{GIOP}_r(\mathbf{C}, \hat{\mathbf{x}}, \epsilon_r)$ and $\text{IOP}(\mathbf{C}, \mathbf{x})$, the statement holds. The proof for part (b) is similar and omitted. □

Proof of Corollary 1 Consider $\text{GIOP}_r(\mathbf{C}, \hat{\mathbf{x}}, \epsilon_r)$. For ϵ_r sufficiently large, $\mathbf{0} \leq \mathbf{C}\hat{\mathbf{x}}/\epsilon_r \leq \mathbf{C}\mathbf{x}$ for some $\mathbf{C}\mathbf{x} \in \mathbf{Z}_{WN}^+$. 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 3, $\text{IOP}(\mathbf{C}, \tilde{\mathbf{x}})$ is feasible and then by Proposition 4, $\text{GIOP}_r(\mathbf{C}, \hat{\mathbf{x}}, \epsilon_r)$ is feasible, as required. The proof for part (b) is similar and omitted. □

Proof of Proposition 5 For part (a), let γ be the dual variable associated with the first set of constraints in formulation (4). Then the dual of formulation (4) can be written as:

$$\begin{aligned} \underset{\mathbf{y},\gamma}{\text{maximize}} \quad & \gamma \\ \text{subject to} \quad & \mathbf{A}\mathbf{y} = \gamma\mathbf{b}, \\ & \mathbf{C}\mathbf{y} \leq \mathbf{C}\hat{\mathbf{x}}, \\ & \mathbf{y} \geq \mathbf{0}. \end{aligned} \tag{13}$$

Let $(\mathbf{y}_r^{dual}, \gamma^*)$ be an optimal solution to formulation (13). By strong duality, $\gamma^* = \epsilon_r^*$. Then from formulation (13), the vector $\mathbf{x}_r^{dual} = \mathbf{y}_r^{dual}/\epsilon_r^*$ must satisfy

$$\epsilon_r^* \leq \min_{k \in \mathcal{K}} \frac{(\mathbf{c}^k)'\hat{\mathbf{x}}}{(\mathbf{c}^k)'\mathbf{x}_r^{dual}}. \tag{14}$$

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 1) and $\mathbf{x}_r^{dual} \in \mathbf{X}$ (from formulation (13)), $\mathbf{C}\mathbf{x}_r^{dual} \in \mathbf{Z}_{WN}$.

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

B. Patient Data

The 12 patient datasets obtained from Princess Margaret Cancer Centre each had a CTV comprising the prostate gland (± 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.

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.

C. Data Initialization

Two pieces of data are needed to solve the GIOPs: the delivered beamlet intensities, and the patient-dependent clinical dose influence matrix, $\mathbf{D}^{clin}$ (see formulation (10)). 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 $\mathbf{D}^{CERR}$:

$$\begin{aligned} \underset{\mathbf{w}}{\text{minimize}} \quad & \sum_{v\in\mathcal{V}} \left(\sum_{b\in\mathcal{B}} D^{CERR}_{v,b} w_b - d^{clin}_v \right)^2 \\ \text{subject to} \quad & \text{constraints from problem (10).} \end{aligned} \tag{15}$$

In Section 5, we treat the achievable dose distribution and $\mathbf{D}^{CERR}$ as being the clinical dose distribution and dose influence matrix, respectively. Note that because our GIOP formulations can handle input vectors that are infeasible for the forward problem, the second part of the data initialization (solving formulation (15)) can be omitted.

D. Sensitivity of the formulation to weights

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 did this 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 Figure 7.

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 from Section 4.3 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 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.

The sensitivity of the bladder and rectum, especially in the higher dose region, seems to be much lower than the femoral heads. This is likely due to the proximity of the bladder and rectum to the PTV – 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). This is the subject of ongoing research.

Acknowledgments

The authors gratefully acknowledge helpful feedback and comments from the Associate Editor and three anonymous reviewers. This research was supported in part by the Natural Sciences and Engineering Research Council of Canada under grant #385972-2010.

References

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

American Cancer Society. 2012. *American Cancer Statistics 2012*.

Beil, D. R., L. M. Wein. 2003. An inverse-optimization-based mechanism to support a multiattribute RFQ process. *Management Sci.* **49**(11) 1529–1545.

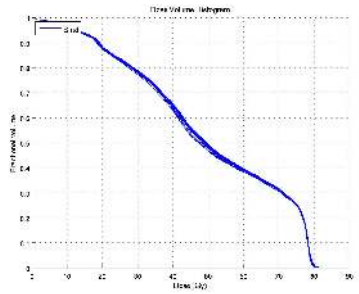

(a) Bladder, model (A) (4-obj)

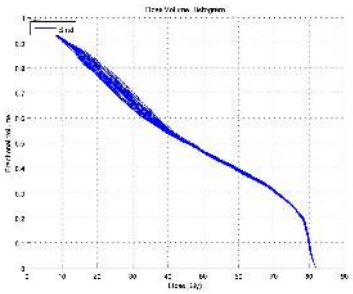

(b) Bladder, model (B) (6-obj)

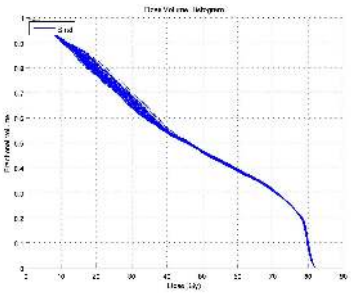

(c) Bladder, model (C) (18-obj)

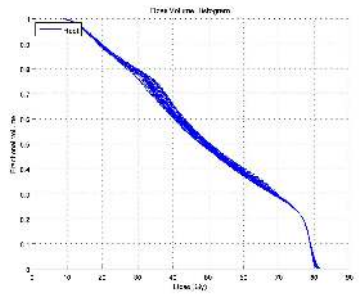

(d) Rectum, model (A)

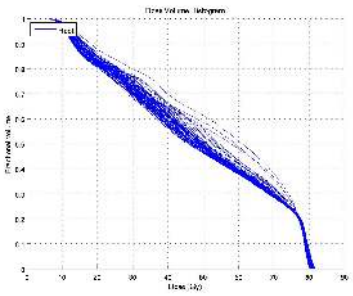

(e) Rectum, model (B)

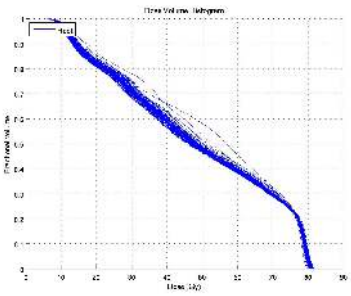

(f) Rectum, model (C)

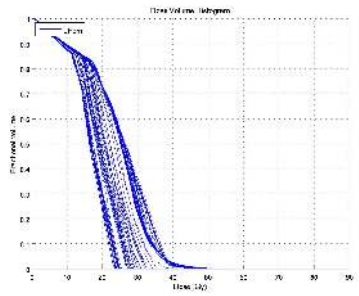

(g) Left femur, model (A)

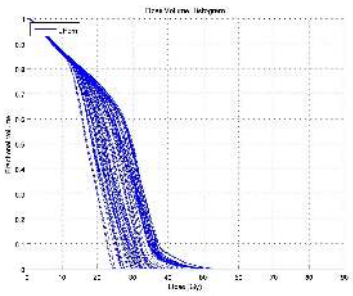

(h) Left femur, model (B)

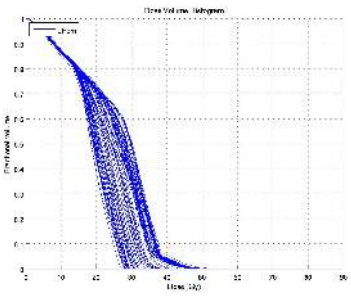

(i) Left femur, model (C)

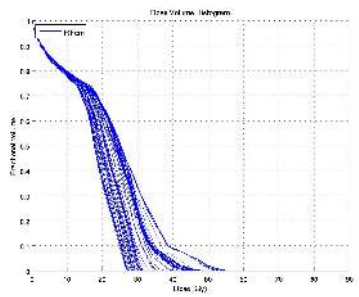

(j) Right femur, model (A)

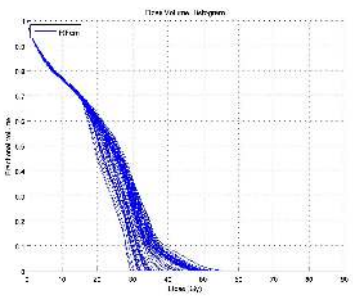

(k) Right femur, model (B)

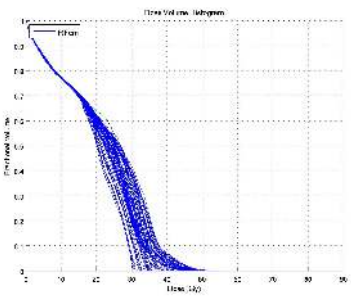

(l) Right femur, model (C)

Figure 7 DVH clouds from 100 sets of weights in the neighborhood of the inversely optimized weights.

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., C. Gupta, I. Ch. Paschalidis. 2012. Inverse optimization: A new perspective on the black-litterman model. *Oper. Res.* **60**(6) 1389–1403.

Bokrantz, R., A. Forsgren. 2013. An algorithm for approximating convex Pareto surfaces based on dual techniques. *INFORMS J. Comput.* **25**(2) 377–393.

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

Canadian Cancer Society. 2012. *Canadian Cancer Statistics 2012*.

Carr, S., W. Lovejoy. 2000. The inverse newsvendor problem: Choosing an optimal demand portfolio for capacitated resources. *Management Sci.* **46**(7) 912–927.

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

Cotrutz, C., 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. Tech. rep., Department of Radiation Oncology, Massachusetts General Hospital.

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

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

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

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

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

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

Ehrgott, M., L. Shao, A. Schöbel. 2011. An approximation algorithm for convex multi-objective programming problems. *J. Glob. Optim.* **50**(1) 397–416.

Erkin, Z., M. D. Bailey, L. M. Maillart, A. J. Schaefer, 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, 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., K. H. Küfer. 2002. Inverse radiation therapy planning – a multiple objective optimization approach. *Discrete Appl. Math.* **118**(1–2) 145–161.

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.

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

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

Karasakal, E., M. Köksalan. 2009. Generating a representative subset of the nondominated frontier in multiple criteria decision making. *Oper. Res.* **57**(1) 187–199.

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

Klamroth, K., J. Tind, M. M. Wiecek. 2002. Unbiased approximation in multicriteria optimization. *Math. Method. Oper. Res.* **56**(3) 413–437.

Lee, E. K., T. Fox, I. Crocker. 2003. Integer programming applied to intensity-modulated radiation therapy treatment planning. *Ann. Oper. Res.* **119**(1–4) 165–181.

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

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, 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.

Rennen, G., E. R. van Dam, D. den Hertog. 2011. Enhancement of sandwich algorithms for approximating higher-dimensional convex Pareto sets. *INFORMS J. Comput.* **23**(4) 493–517.

Romeijn, H. E., R. K. Ahuja, J. F. Dempsey, 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, J. G. Li. 2004. A unifying framework for multi-criteria 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., 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, T. R. Mackie. 1999. Optimizing the delivery of radiation therapy to cancer patients. *Siam Review* **41**(4) 721–744.

Tarantola, A. 2005. *Inverse Problem Theory and Methods for Model Parameter Estimation*. SIAM, Philadelphia.

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

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

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**(12) 2229–2246.

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

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

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

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

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, 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*. Springer Verlag, Berlin.

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

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

Zhu, L., L. Lee, Y. Ma, Y. Ye, R. Mazzeo, 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.