Objective Selection for Cancer Treatment: An Inverse Optimization Approach

Temitayo Ajayi
Department of Computational and Applied Mathematics, Rice University, 77005, tayo.ajayi25@gmail.com
Department of Radiation Oncology, The University of Texas MD Anderson Cancer Center, 77030

Taewoo Lee
Department of Industrial Engineering, University of Houston, 77004, tlee20@central.uh.edu

Andrew J. Schaefer
Department of Computational and Applied Mathematics, Rice University, 77005, andrew.schaefer@rice.edu

In radiation therapy treatment plan optimization, selecting a set of clinical objectives that are tractable and parsimonious yet effective is a challenging task. In clinical practice, this is typically done by trial and error based on the treatment planner's subjective assessment, which often makes the planning process inefficient and inconsistent. We develop the objective selection problem that infers a sparse set of objectives for prostate cancer treatment planning based on historical treatment data. We formulate the problem as a non-convex bilevel mixed-integer program using inverse optimization and highlight its connection with feature selection to propose multiple solution approaches, including greedy heuristics and regularized problems as well as application-specific methods that utilize anatomical information of the patients. Our results show that the proposed heuristics find objectives that are near optimal. Via curve analysis on dose-volume histograms, we show that the learned objectives closely represent latent clinical preferences.

Key words: Inverse optimization; objective selection; feature selection; greedy algorithm; multi-objective optimization; cancer therapy

1. Introduction

Cancer is the second-leading cause of death in the United States, with more than 1.7 million new cases and 600,000 deaths estimated in 2018 (American Cancer Society 2018). Approximately 40

percent of Americans will develop cancer in their lifetimes (American Cancer Society 2018). In this decade, cancer-associated expenditure will be between $124.4 and $157.7 billion (Mariotto et al. 2011).

Intensity-modulated radiation therapy (IMRT) is one of the primary methods to treat cancer. Because many clinical parameters are involved in creating a clinically desirable treatment, such as the number, angles, and intensities of the beams, radiation therapy is typically designed using mathematical optimization (Shepard et al. 1999, Bortfeld 1999). One substantial challenge in radiation therapy is the conflict between multiple clinical goals, such as delivering a sufficient amount of radiation dose to the tumor while sparing nearby healthy organs. Moreover, each healthy organ responds to radiation differently; thus, various clinical objectives are required to address issues with different organs. Treatment planners typically employ multi-objective optimization techniques to find optimal treatment parameter settings that generate desirable trade-offs among the set of objectives (e.g., Romeijn et al. (2004), Craft et al. (2007), Shao and Ehrgott (2008)).

Though many commercial treatment planning systems offer dedicated tools to deal with the forward optimization of the planning process, planning objectives are typically chosen subjectively based on the planner's experience (Craft 2011). Treatment planners adjust the selection of the objectives and associated parameters by trial and error as they repeatedly solve the optimization problem and evaluate the resulting treatment plan (Zhang et al. 2011). Consequently, the entire planning process often becomes slow and leads to inconsistent, suboptimal treatment quality (Lee et al. 2013, Boutilier et al. 2015).

The quality of the treatment plan and the efficiency of the planning process are heavily affected by the choice of objective functions. The chosen objectives affect computational complexity as well as the realism of the optimization problem; for example, linear objective functions can make the problem computationally more efficient, but the resulting formulation may not adequately reflect clinical reality. In addition, it is beneficial if the objectives are widely applicable so the constructed model with the chosen objectives can be used repeatedly for many similar patients. Thus, it is

crucial to find a sparse set of objectives that adequately reflect reality and efficiently generate high-quality treatment plans.

In this paper, we propose a novel, data-driven approach to select a sparse set of effective clinical objectives for radiation therapy treatment planning. Finding a good set of objectives that best represents previously exhibited clinical preferences can be viewed as a feature selection problem. Given an input set of historical treatment plans, we aim to learn objectives that can render the given treatments near optimal for the underlying multi-objective optimization problem. By doing so, we reproduce preferences that were implicit in the treatments, which can guide the generation of clinically desirable treatments for new patients efficiently. The task of inferring optimization parameters from data has been widely studied in the inverse optimization literature. Hence, we employ an inverse optimization approach integrated with feature selection to develop the objective selection problem. Our framework not only selects a subset of objectives with a specified sparsity, but also simultaneously weights the selected subset of objectives. Specifically, our contributions in this work are as follows:

1. We propose a novel, data-driven approach to select the best subset of objectives in multi-objective optimization problems. We formulate the objective selection problem using inverse optimization and show that an extensive formulation can be formulated as a bilevel, non-convex mixed-integer program.
2. We establish a connection between objective selection and feature selection and adopt the solution techniques for feature selection to approximately solve the objective selection problem, including greedy algorithm and regularization approaches. We also propose the application-specific, anatomy-based greedy algorithm that exploits the problem structure of cancer treatment planning.

2. Literature Review

Many different objectives have been proposed in IMRT planning, such as dose-volume objectives (Halabi et al. 2006, Wu et al. 2011), equivalent uniform dose (EUD) objectives (Wu et al. 2002,

Choi and Deasy 2002), quadratic penalty objectives (Breedveld et al. 2006, Romeijn et al. 2006), and minimum, maximum, and mean dose objectives and combinations of them (Thieke et al. 2002, Craft et al. 2012). Currently, there is no consensus on which objectives should be used for different cancers and patients; final specification is often a task for treatment planners within the trial-and-error process (Xing et al. 1999, Cotrutz et al. 2001).

There are many studies in the literature that relate to the computation of weights for a given set of objectives in radiation therapy. Xing et al. (1999) propose an algorithm that automates the iterative weight adjustment process guided by a scoring function that measures clinical acceptability. Zhang et al. (2006) use a "sensitivity-guided" approach to balance the importance of multiple objectives in the treatment planning process. Similarly, Zhang et al. (2011) include an automated parameter adjustment to iteratively alter the objective function within the treatment planning framework. Wilkens et al. (2007) prioritize objectives to construct a treatment plan progressively, where the order of the objectives indicates the relative importance instead of the weights. Lee et al. (2013) and Boutilier et al. (2015) employ inverse optimization and machine learning techniques to predict weights from the anatomical features of patients. However, these studies weight the entire pool of objectives; they do not focus on creating the optimization model by selecting objectives. The number of objectives and the complexity of the model are an important issue, especially with recent studies on automated- and knowledge-based treatment planning (Chanyavanich et al. 2011, McIntosh and Purdie 2016, Babier et al. 2018, Kearney et al. 2018, Mahmood et al. 2018).

Inverse optimization has received growing attention as a tool for using data to determine modeling parameters that lead to an efficient and effective optimization formulation. Inverse optimization generally seeks parameters of an objective function that best explain the system behavior or a decision-maker's preferences for various types of optimization problems such as linear programs and network optimization problems (e.g., Burton and Toint (1992), Zhang and Liu (1996) and Ahuja and Orlin (2001)), convex optimization (Iyengar and Kang 2005, Keshavarz et al. 2011), and integer (Schaefer 2009) and mixed-integer programs (Wang (2009), Duan and Wang (2011)

and Lamperski and Schaefer (2015)). Over the last decade, the scope of inverse optimization has expanded to accommodate noise, measurement errors, and uncertainty in data; the goal in such settings is to find the objective function parameters that make the input data as optimal as possible (Bertsimas et al. 2015, Aswani et al. 2018, Chan et al. 2019). In the presence of such imperfect data, the inverse problem is typically non-convex—Chan et al. (2014) and Chan et al. (2019) propose efficient exact algorithms when the underlying problem is linear; for general inverse convex programming, various approximations and heuristics have been proposed (Bertsimas et al. 2015, Aswani et al. 2018, Esfahani et al. 2018).

Inverse optimization has also been recently studied in the multi-objective optimization setting (Keshavarz et al. 2011, Chan and Lee 2018, Naghavi et al. 2019, Gebken and Peitz 2021). However, the existing inverse multi-objective optimization approaches focus only on finding optimal weighting factors for a pre-specified set of objectives. For example, Chan and Lee (2018) apply inverse optimization to radiation therapy treatment planning, yet their focus is to find weighting factors for a given set of objectives. Gebken and Peitz (2021) use a singular-value decomposition approach, but they focus only on finding objective weights for optimal solutions of unconstrained problems.

We observe that objective selection is a special case of feature selection in the sense that its goal is to find objectives that best fit the preferences exhibited in the given treatment data. Feature selection is a well-known problem in which one selects "a subset of variables from the input which can efficiently describe the input data while reducing effects from noise or irrelevant variables and still provide good prediction results" (Chandrashekar and Sahin 2014). Feature selection is NP-hard and computationally intractable in general (Amaldi and Kann 1998, Guyon and Elisseeff 2003), and numerous solution methods have been proposed including filter, regularization, and wrapper methods. Filter methods rely on ranking techniques and select features that have a rank above a pre-determined threshold (Chandrashekar and Sahin 2014). In a regularized method, the extensive form of the feature selection problem is convexified and a penalty term is added to

the objective function of the problem to encourage the selection of a sparse subset of features. Sparsity regularization requires an appropriate regularization parameter, which can be a difficult problem itself (Galatsanos and Katsaggelos 1992, Bazán and Francisco 2009). Wrapper methods use some black-box solver (see Martí (2015) for details) to evaluate the performance of a subset of features and employ a heuristic to select a good subset of features (Chandrashekar and Sahin 2014). Although some of the methods are known to produce an optimal subset of features theoretically (e.g., branch-and-bound), they often need to examine an exponential number of subsets. Thus, it is typically desirable to develop more computationally efficient methods with potentially some loss of optimality (Pudil et al. 1994).

As an alternative, sequential wrapper methods allow for the subset to dynamically change size, and the feature subset is evaluated at various iterations to determine when the subset is good enough. The *greedy algorithm*, known as forward selection in some data science contexts, is a sequential wrapper method prevalent in combinatorial optimization. In a greedy algorithm, features are selected iteratively based on which feature improves a given selection criterion. Although greedy algorithms can be sensitive to local optima, they are simple and the selection process interacts directly with the model produced by the selected features (Saeys et al. 2007).

Certain structural properties of the problem can impact the effectiveness of the greedy algorithm (Nemhauser et al. 1978). In particular, submodularity plays an important role in deriving performance guarantees (i.e., bounds) for the greedy algorithm in various machine learning problems (Krause et al. 2008, Shamaiah et al. 2010). In addition, the greedy algorithm has also been used for feature selection with non-submodular functions (e.g., Das and Kempe (2011)) as well as other optimization problems without an examination of submodularity (Gottlieb et al. 2003). Indeed, the greedy algorithm can efficiently obtain a good solution if the subproblems solved within the iterations have good structure, such as convexity. We show in Section 3 that the objective selection problem contains a convex subproblem, which strengthens the case to use the greedy algorithm.

3. Methodology

In this section, we develop a mathematical framework to select objectives for radiation therapy treatment planning. Although our focus is on cancer treatment planning, we present the ensuing formulations in a general mathematical setting.

3.1. Forward Multi-objective Optimization

Let $\mathcal{K} = \{1, \ldots, K\}$ be the index set for objectives $f_1(\mathbf{x}), \ldots, f_K(\mathbf{x})$, where $\mathbf{x} \in \mathbb{R}^n$. Given a subset $\mathcal{S} \subseteq \mathcal{K}$, we define the forward multi-objective optimization (MO) problem as follows:

$$\begin{aligned} \text{MO}(\mathcal{S}) = \quad & \min_{\mathbf{x}} \quad \mathbf{f}(\mathbf{x}) = [f_k(\mathbf{x})]_{k \in \mathcal{S}} \\ & \text{subject to} \quad \mathbf{A}\mathbf{x} = \mathbf{b}, \\ & \qquad\qquad \mathbf{g}(\mathbf{x}) \leq \mathbf{0}, \end{aligned} \tag{MO}$$

where we use $[f_k(\mathbf{x})]_{k \in \mathcal{S}}$ to denote a vector-valued function. The inequality constraints are defined by the vector-valued function $\mathbf{g} : \mathbb{R}^n \to \mathbb{R}^L$, with $\mathcal{L} = \{1, ..., L\}$, and the equality constraints are defined by $\mathbf{A} \in \mathbb{R}^{m \times n}$ and $\mathbf{b} \in \mathbb{R}^m$. In the context of radiation therapy, the objectives f_k may correspond to penalties for excessive radiation exposure to different organs at risk, $\mathbf{g}, \mathbf{A}$, and $\mathbf{b}$ represent dose-delivery constraints for various voxels in patient anatomy, and the decision variable $\mathbf{x}$ represents a treatment plan expressed as the amount of radiation to be delivered from each "beamlet." More details for radiation therapy-specific formulations are given in Section 4. We emphasize, however, that the methodology presented in this work applies to general multi-objective optimization problems and thus, we focus on the general MO framework above.

In general, there is no solution that can simultaneously minimize f_k for all $k \in \mathcal{K}$. Instead, the goal of MO is to find a set of solutions that are not dominated by other feasible solutions in any of the objectives. Formally, a solution to an MO problem is called *weakly Pareto optimal* if there does not exist $\mathbf{y}$ that is feasible for MO($\mathcal{S}$) such that $f_k(\mathbf{x}) > f_k(\mathbf{y})$ for all $k \in \mathcal{S}$. We let $\Omega(\mathcal{S})$ be the set of weakly Pareto optimal solutions for MO($\mathcal{S}$), also referred to as the Pareto surface.

Assumption 1. *The functions f_k and g_ℓ are convex and differentiable for all $k \in \mathcal{K}$ and $\ell \in \mathcal{L}$, $\mathbf{f}(\mathbf{x}) > \mathbf{0}$ for all $\mathbf{x}$ feasible for* MO($\mathcal{S}$) *for all $\mathcal{S} \subseteq \mathcal{K}$, and Slater's interior feasibility condition holds.*

Assumption 1 provides some basic structure to our objective selection framework and also permits the inclusion of nonlinear objectives and constraints.

3.2. Objective Selection Using Inverse Multi-objective Optimization

Given an input solution as data (e.g., a solution chosen by a decision maker), objective selection focuses on selecting a set of no more than θ objectives (with weights) from a pool of candidate objectives that capture the preferences implicitly exhibited in the input. In other words, the objective selection problem can be seen as a form of inverse optimization to find objectives that render the given solution as close as possible to the Pareto surface associated with the objectives, in terms of some specified (optimality) distance function. Given a feasible input solution $\widehat{\mathbf{x}}$, a set of objectives $\mathcal{S}$, and $\mathbf{x} \in \Omega(\mathcal{S})$, consider the following distance functions:

- $d^{[r]}(\mathbf{x}, \widehat{\mathbf{x}}, \mathcal{S}) = \max\limits_{k \in \mathcal{S}} \frac{f_k(\mathbf{x})}{f_k(\widehat{\mathbf{x}})}$,
- $d^{[q]}(\mathbf{x}, \widehat{\mathbf{x}}, \mathcal{S}) = \|\mathbf{f}(\mathbf{x}) - \mathbf{f}(\hat{\mathbf{x}})\|_q, q \in \{1, 2, \ldots, \infty\}$.

Although there are many other distance functions we can consider, the above functions are considered particularly relevant in the multi-objective optimization context in that a (weakly) Pareto optimal solution $\mathbf{x}$ that minimizes such distances from $\widehat{\mathbf{x}}$ is considered to preserve the original trade-offs in the objective values (Eskelinen and Miettinen 2012, van Haveren et al. 2017). Though one might consider a more straightforward distance such as $\|\widehat{\mathbf{x}} - \mathbf{x}\|_q$, recent studies show that this can cause the resulting objective values to be inconsistent with the input values $\mathbf{f}(\widehat{\mathbf{x}})$ and is more appropriate for single-objective optimization (Keshavarz et al. 2011, Chan and Lee 2018). Minimizing $d^{[r]}$ above results in minimizing the (relative) duality gap of a modified version of MO where the objective function is a weighted sum of the components of $\mathbf{f}$. Moreover, using $d^{[r]}$ in inverse multi-objective optimization leads to weight vectors that preserve relative tradeoff preferences (Chan and Lee 2018). We remark that $d^{[r]}$ is not a metric strictly speaking; for instance, it does not obey the triangle inequality. However, it does lead to valuable insight about multi-objective trade-offs and provides structure to the objective selection problem that we exploit. Lin

(2005) discusses uses of $d^{[q]}$ for a specific case in which the input data correspond to a set of ideal solutions (optimal for a single-objective problem) for the same multi-objective problem.

Our objective selection framework accepts multiple inputs. Given P inputs, the goal is to find a set of objectives such that the resulting Pareto surfaces are as close to the input solutions as possible. Let $\mathcal{P} = \{1, ..., P\}$ be the index set for the inputs and $\widehat{\chi} = \{\widehat{\mathbf{x}}_{(p)}\}_{p\in\mathcal{P}}$. Note that each input $\widehat{\mathbf{x}}_{(p)}$ may originate from a different forward optimization problem, in which case objectives and constraint parameters from each optimization problem are also indexed by (p). In the treatment planning context, each input is a previous clinical treatment plan for each patient.

We remark that the objectives across different input solutions share the same functional format so they can be considered the same "type." For example, $f_{(1),1} : \mathbb{R}^{n_{(1)}} \to \mathbb{R}$, the first objective for input 1, and $f_{(2),1} : \mathbb{R}^{n_{(2)}} \to \mathbb{R}$, the first objective for input 2, can both represent the L_2 norm of the solution for the respective problems even though the dimensions of the domains differ in size. Thus, each objective index $k \in \mathcal{K}$ represents the common type of the objectives consistent across the inputs. Each forward problem also leads to a different Pareto surface $\Omega_{(p)}(\mathcal{S})$.

We consider two distances between a set of input solutions and the respective Pareto surfaces:

- $\delta(\chi, \widehat{\chi}, \mathcal{S}) = \sum_{p\in\mathcal{P}} d_{(p)}(\mathbf{x}_{(p)}, \widehat{\mathbf{x}}_{(p)}, \mathcal{S})$
- $\gamma(\chi, \widehat{\chi}, \mathcal{S}) = \min_{\boldsymbol{\alpha}\in\mathcal{A}(\mathcal{S})} \left\{ \boldsymbol{\alpha}^\top \sum_{p\in\mathcal{P}} \mathbf{f}_{(p)}(\widehat{\mathbf{x}}_{(p)}) \big/ \boldsymbol{\alpha}^\top \sum_{p\in\mathcal{P}} \mathbf{f}_{(p)}(\mathbf{x}_{(p)}) \right\}$,

where $\mathcal{A}(\mathcal{S}) = \{\boldsymbol{\alpha} \in \mathbb{R}^K_+ \backslash \{\mathbf{0}\} \mid \alpha_k = 0 \text{ for all } k \notin \mathcal{S}\}$ is the set of admissible weight vectors for a set of objectives $\mathcal{S}$. Different single input distance metrics may be used for δ, for instance, $d^{[r]}$. Notice that δ is a separable measure, unlike γ. We refer to γ as the *batch duality gap*. For most of this section, we focus on the separable case, but we also present results with the batch duality gap.

We write the objective selection problem (OS) with input set $\mathcal{P}$ as follows:

$$\begin{aligned} \text{OS}(\widehat{\chi}) = \quad & \min_{\chi,\mathcal{S}} \quad \delta(\chi, \widehat{\chi}, \mathcal{S}) \\ & \text{subject to} \quad \mathbf{x}_{(p)} \in \Omega_{(p)}(\mathcal{S}),\ \forall\ p \in \mathcal{P}, \\ & \qquad\qquad 1 \le |\mathcal{S}| \le \theta, \\ & \qquad\qquad \mathcal{S} \subseteq \mathcal{K}. \end{aligned} \tag{OS}$$

For all input data points, the above problem finds a common set of objectives $\mathcal{S}$, i.e., the same types of objectives, although the underlying forward problem formulation is different across the inputs. If the input data points are all collected from the same forward optimization problem, we can rewrite the above formulation simply by replacing the first constraint with $\mathcal{X} \subset \Omega(\mathcal{S})$. We remark that the objective selection problem reduces to the normal multi-objective inverse problem of Chan et al. (2014) and Chan and Lee (2018) if $\theta = |\mathcal{K}|$ and $\mathcal{P}$ is a singleton.

3.2.1. Extensive Formulation We first present a general extensive formulation for OS($\hat{\mathcal{X}}$). Given that $\mathbf{f}$ and $\mathbf{g}$ are convex in the forward problem MO (Assumption 1), all weakly Pareto optimal solutions for MO($\mathcal{S}$) can be obtained by the following formulation with the weighted combination of the objectives:

$$\text{WMO}(\boldsymbol{\alpha}) = \min_{\mathbf{x}} \{\boldsymbol{\alpha}^\top \mathbf{f}(\mathbf{x}) \mid \mathbf{A}\mathbf{x} = \mathbf{b}, \mathbf{g}(\mathbf{x}) \leq \mathbf{0}\}. \tag{WMO}$$

Let $\mathcal{O}(\boldsymbol{\alpha})$ denote the set of optimal solutions to WMO($\boldsymbol{\alpha}$), and given input index p, let $\mathcal{O}_{(p)}(\boldsymbol{\alpha})$ denote the optimal solutions of WMO($\boldsymbol{\alpha}$) with data $\mathbf{f}_{(p)}, \mathbf{A}_{(p)}, \mathbf{b}_{(p)}$, and $\mathbf{g}_{(p)}$.

Lemma 1. *For any $\mathcal{S} \subseteq \mathcal{K}$, a solution $\mathbf{x}$ is weakly Pareto optimal for* MO($\mathcal{S}$) *if and only if there exists a non-zero weight vector $\boldsymbol{\alpha} \in \mathbb{R}^K_+$ with $\alpha_k = 0$ for all $k \notin \mathcal{S}$ such that $\mathbf{x}$ is an optimal solution to* WMO($\boldsymbol{\alpha}$). *That is, $\mathbf{x} \in \Omega(\mathcal{S}) \iff \mathbf{x} \in \bigcup_{\boldsymbol{\alpha}} \{\mathcal{O}(\boldsymbol{\alpha}) \,|\, \boldsymbol{\alpha} \in \mathbb{R}^K_+ \backslash \{\mathbf{0}\}, \alpha_k = 0$ for all $k \notin \mathcal{S}\}$.*

Proof of Lemma 1 is omitted as it follows directly from Theorem 3.15 of Ehrgott (2005).

Let $\mathcal{S}(\boldsymbol{\xi}) = \{k \in \mathcal{K} \mid \xi_k = 1\}$ where $\boldsymbol{\xi} \in \mathbb{B}^K$ is a binary variable. From Lemma 1, an extensive formulation of OS for the set of P input data points, $\widehat{\chi}$, can be written as follows:

$$\min_{\boldsymbol{\alpha}, \boldsymbol{\xi}, \chi} \quad \delta(\chi, \widehat{\chi}, \mathcal{S}(\boldsymbol{\xi})) \tag{1a}$$

$$\text{subject to} \quad \sum_{k \in \mathcal{K}} \xi_k \leq \theta, \tag{1b}$$

$$\boldsymbol{\alpha}_{(p)} \leq \boldsymbol{\xi}, \ \forall\, p \in \mathcal{P}, \tag{1c}$$

$$\boldsymbol{\alpha}_{(p)} \in \mathcal{A}(\mathcal{S}(\boldsymbol{\xi})), \ \forall\, p \in \mathcal{P}, \tag{1d}$$

$$\boldsymbol{\xi} \in \mathbb{B}^K, \tag{1e}$$

$$\mathbf{x}_{(p)} \in \mathcal{O}_{(p)}(\boldsymbol{\alpha}_{(p)}),\ \forall\ p \in \mathcal{P}. \tag{1f}$$

The binary variable $\boldsymbol{\xi}$ represents whether or not each candidate objective is chosen—if $\xi_k = 0$, objective k is not chosen as the corresponding weight α_k is forced to be 0. We assume $\alpha_k \leq 1$ for each objective k; this is without loss of generality, as any weight vector for the forward problem MO can be scaled by a positive constant without affecting the set of optimal solutions. Note that the above formulation assumes each data point $\widehat{\mathbf{x}}_{(p)}$ is associated with a different weight vector, i.e., collected from a decision maker with different preferences. If we assume that all data points represent the same preferences and thus should be assigned the same weight values, we modify the formulation by fixing $\boldsymbol{\alpha}_{(p)} = \boldsymbol{\alpha}$, for all $p \in \mathcal{P}$. Assuming each of the objectives $f_{(p),k}$ for all $k \in \mathcal{K}$ and constraint vectors $\mathbf{g}_{(p)}$ are differentiable for all data points $p \in \mathcal{P}$, formulation (1) can be further rewritten as the following non-convex mixed-integer program:

$$\min_{\boldsymbol{\alpha},\boldsymbol{\xi},\chi,\sigma,\pi} \quad \delta(\chi, \widehat{\chi}, \mathcal{S}(\boldsymbol{\xi})) \tag{2a}$$

$$\text{subject to} \quad \sum_{k\in\mathcal{K}} \xi_k \leq \theta, \tag{2b}$$

$$\boldsymbol{\alpha}_{(p)} \leq \boldsymbol{\xi}, \tag{2c}$$

$$\boldsymbol{\alpha}_{(p)}^{\top}\mathbf{e} = 1,\ \forall\ p \in \mathcal{P}, \tag{2d}$$

$$\mathbf{A}_{(p)}\mathbf{x}_{(p)} = \mathbf{b}_{(p)},\ \forall\ p \in \mathcal{P}, \tag{2e}$$

$$\mathbf{g}_{(p)}(\mathbf{x}_{(p)}) \leq \mathbf{0},\ \forall\ p \in \mathcal{P}, \tag{2f}$$

$$\sum_{k\in\mathcal{K}} \alpha_{(p),k}\nabla f_{(p),k}(\mathbf{x}_{(p)}) + \sum_{\ell\in\mathcal{L}} \sigma_{(p),\ell}\nabla g_{(p),\ell}(\mathbf{x}_{(p)}) + \mathbf{A}_{(p)}^{\top}\boldsymbol{\pi}_{(p)} = \mathbf{0},\ \forall\ p \in \mathcal{P}, \tag{2g}$$

$$\boldsymbol{\sigma}_{(p)} \circ \mathbf{g}_{(p)}(\mathbf{x}_{(p)}) = \mathbf{0},\ \forall\ p \in \mathcal{P}, \tag{2h}$$

$$\boldsymbol{\sigma}_{(p)} \geq \mathbf{0},\ \forall\ p \in \mathcal{P}, \tag{2i}$$

$$\boldsymbol{\alpha}_{(p)} \geq \mathbf{0},\ \forall\ p \in \mathcal{P}, \tag{2j}$$

$$\boldsymbol{\xi} \in \mathbb{B}^K, \tag{2k}$$

where constraints (2f)–(2j) represent the KKT conditions for data point p, replacing (1f), ($\circ$) is the Hadamard product, and $\mathbf{e}$ is a vector of ones of the appropriate dimension.

PROPOSITION 1. *$\mathcal{F}_1$, the feasible region of* (1)*, and $\mathcal{F}_2$, the feasible region of* (2) *are related as follows:*

$$\{(\boldsymbol{\alpha},\boldsymbol{\xi},\chi) \mid (\boldsymbol{\alpha},\boldsymbol{\xi},\chi,\boldsymbol{\sigma},\boldsymbol{\pi}) \in \mathcal{F}_2 \text{ for some } (\boldsymbol{\sigma},\boldsymbol{\pi})\} = \{(\boldsymbol{\alpha},\boldsymbol{\xi},\chi) \in \mathcal{F}_1 \mid \boldsymbol{\alpha}_{(p)}^\top \mathbf{e} = 1, \text{ for all } p \in \mathcal{P}\}.$$

Thus, the optimal objective values of (1) *and* (2) *are equal.*

Note that (2) is a mixed-integer non-convex problem, which is intractable in general. We show that formulation (2) is tractable if the underlying forward problem is linear, i.e., all of the candidate objectives and the constraints are linear. Suppose $\mathbf{f}_{(p)}$ and $\mathbf{g}_{(p)}$ are both affine functions for each $p \in \mathcal{P}$ so that the weighted forward MO for input p can be written as follows:

$$\min_{\mathbf{x}} \{\boldsymbol{\alpha}_{(p)}^\top \mathbf{C}_{(p)}\mathbf{x} \mid \mathbf{A}_{(p)}\mathbf{x} = \mathbf{b}, \mathbf{G}_{(p)}\mathbf{x} \leq \mathbf{h}_{(p)}\}. \tag{3}$$

Then the corresponding OS problem with input set $\mathcal{P}$ can be written as a mixed-integer program:

$$\min_{\mathbf{x},\boldsymbol{\alpha},\boldsymbol{\pi},\boldsymbol{\sigma},\mathbf{s},\boldsymbol{\xi}} \quad \sum_{p\in\mathcal{P}} d_{(p)}(\mathbf{x}_{(p)}, \widehat{\mathbf{x}}_{(p)}, \mathcal{S}) \tag{4a}$$

$$\text{subject to} \quad \sum_{k\in\mathcal{K}} \boldsymbol{\xi}_k \leq \theta, \tag{4b}$$

$$\boldsymbol{\alpha}_{(p)} \leq \boldsymbol{\xi}, \ \forall\ p \in \mathcal{P}, \tag{4c}$$

$$\boldsymbol{\alpha}_{(p)}^\top \mathbf{e} = 1, \ \forall\ p \in \mathcal{P}, \tag{4d}$$

$$\mathbf{A}_{(p)}\mathbf{x}_{(p)} = \mathbf{b}_{(p)}, \ \forall\ p \in \mathcal{P}, \tag{4e}$$

$$\mathbf{G}_{(p)}\mathbf{x}_{(p)} \leq \mathbf{h}_{(p)}, \ \forall\ p \in \mathcal{P}, \tag{4f}$$

$$\boldsymbol{\alpha}_{(p)}^\top \mathbf{C}_{(p)} + \boldsymbol{\sigma}_{(p)}^\top \mathbf{G}_{(p)} + \boldsymbol{\pi}_{(p)}^\top \mathbf{A}_{(p)} = \mathbf{0}, \ \forall\ p \in \mathcal{P}, \tag{4g}$$

$$\mathbf{h}_{(p)} - \mathbf{G}_{(p)}\mathbf{x}_{(p)} \leq M\mathbf{s}_{(p)}, \ \forall\ p \in \mathcal{P}, \tag{4h}$$

$$\boldsymbol{\sigma}_{(p)} \leq M(1 - \mathbf{s}_{(p)}), \ \forall\ p \in \mathcal{P}, \tag{4i}$$

$$\boldsymbol{\sigma}_{(p)} \geq \mathbf{0}, \ \forall\ p \in \mathcal{P}, \tag{4j}$$

$$\boldsymbol{\alpha}_{(p)} \geq \mathbf{0}, \ \forall\ p \in \mathcal{P}, \tag{4k}$$

$$\boldsymbol{\xi} \in \mathbb{B}^K, \tag{4l}$$

$$\mathbf{s}_{(p)} \in \mathbb{B}^L, \ \forall\ p \in \mathcal{P}, \tag{4m}$$

where M is a suitably large constant. Observe that (4e)–(4k) again represent the KKT conditions in a slightly different way than in (2); we introduce auxiliary variables $\mathbf{s}$ and big-M constraints to handle the complementary slackness. Again, if all P input solutions are assumed to represent the same preferences over the common set of objectives (i.e., common weights), we let $\boldsymbol{\alpha}_{(p)} = \boldsymbol{\alpha}, \ \forall\ p \in \mathcal{P}$.

When $\mathbf{f}_{(p)}$ and $\mathbf{g}_{(p)}$ are nonlinear functions, the OS formulation is generally non-convex. Furthermore, we observe that the general OS problem can be seen as a bilevel optimization problem, where the upper-level problem chooses the set of objectives first, followed by the lower-level problem that chooses the corresponding weight values. The challenge is that even the lower-level problem itself is not tractable due to the non-convexity in the KKT constraints (multiplication of $\boldsymbol{\alpha}$ and $\mathbf{f}(\mathbf{x})$ both of which are variables). In the sequel, we analyze the structure of the lower-level problem and show that it becomes tractable under some distance function d, enabling efficient solution approaches for the OS problem.

3.2.2. Restricted Inverse Problem Consider OS($\widehat{\chi}$) as the following bilevel framework: the upper-level problem first picks a set of objectives; the lower-level problem follows and finds an optimal weight vector for each data point $\widehat{\mathbf{x}}_{(p)}$ such that the associated distance function is minimized. We call the lower-level problem the restricted inverse problem (RP). Using reformulations, this bilevel problem can be expressed as a zero-sum problem. Given a pre-specified choice of objectives, $\mathcal{S}$, we present the following formulation to find a weight vector that minimizes $d_{(p)}(\mathbf{x}_{(p)}, \widehat{\mathbf{x}}_{(p)}, \mathcal{S}) = \max_{k \in \mathcal{S}} \{f_{(p),k}(\mathbf{x}_{(p)}) / f_{(p),k}(\widehat{\mathbf{x}}_{(p)})\}$ for each data point p:

$$\text{RP}(\widehat{\chi}, \mathcal{S}) = \quad \min_{\mathcal{X}} \quad \sum_{p \in \mathcal{P}} \max_{k \in \mathcal{K}} \left\{ f_{(p),k}(\mathbf{x}_{(p)}) / f_{(p),k}(\widehat{\mathbf{x}}_{(p)}) \right\} \tag{5a}$$

$$\text{subject to} \quad \mathbf{A}_{(p)} \mathbf{x}_{(p)} = \mathbf{b}_{(p)}, \ \forall\ p \in \mathcal{P}, \tag{5b}$$

$$\mathbf{g}_{(p)}(\mathbf{x}_{(p)}) \leq \mathbf{0}, \ \forall\ p \in \mathcal{P}. \tag{5c}$$

When $\mathcal{S} = \mathcal{K}$, we call the restricted inverse problem the *normal inverse problem*. The objective function of the above problem can be rewritten as a convex function after including auxiliary variables and convex constraints.

PROPOSITION 2. *Let $\chi^* = \{\mathbf{x}^*_{(p)}\}_{p\in\mathcal{P}}$ be an optimal solution for $\mathrm{RP}(\widehat{\chi}, \mathcal{S})$. Then, for all $p \in \mathcal{P}$:*

(i) $\mathbf{x}^*_{(p)} \in \Omega_{(p)}(\mathcal{S})$.

(ii) $1/\max\limits_{k\in\mathcal{K}}\{f_{(p),k}(\mathbf{x}^*_{(p)})/f_{(p),k}(\widehat{\mathbf{x}}_{(p)})\}$ *is the minimum relative duality gap with respect to* $\widehat{\mathbf{x}}_{(p)}$.

In addition, a convex reformulation of $\mathrm{RP}(\widehat{\chi}, \mathcal{S})$ *exists where optimal weight vectors* $\boldsymbol{\alpha}^*_{(p)}$, *such that* $\mathbf{x}^*_{(p)} \in \mathcal{O}_{(p)}(\boldsymbol{\alpha}^*_{(p)})$, *can be computed as Lagrange multipliers, for all* $p \in \mathcal{P}$.

Proposition 2 states that given a fixed subset of candidate objectives, one can solve the restricted inverse problem with multiple inputs. Moreover, the weight vectors can be recovered using KKT conditions to compute the Lagrange multipliers. Thus, $\mathrm{RP}(\widehat{\chi}, \mathcal{S})$ can serve as a black-box solver in sequential wrapper algorithms common in feature selection (Theorem 1). Returning to the bilevel view of the objective selection problem, the lower-level problem can be reformulated as formulation (EC.12), where the objective of the follower is to minimize $\sum\limits_{p\in\mathcal{P}} \epsilon_{(p)}$, the sum of the reciprocals of the relative duality gaps; in other words, the follower maximizes the sum of the duality gaps. The goal of the leader, however, is to choose the subset that minimizes the sum of the duality gap. This illustrates that the objective selection problem in this case is a zero-sum bilevel problem (Theorem 1).

THEOREM 1. *Given the distance functions* $d_{(p)}(\mathbf{x}, \widehat{\mathbf{x}}, \mathcal{S}) = \max\limits_{k\in\mathcal{S}} \left\{f_{(p),k}(\mathbf{x})/f_{(p),k}(\widehat{\mathbf{x}})\right\}$, *for all* $p \in \mathcal{P}$, *we have* $\mathrm{OS}(\widehat{\chi}) = \min\limits_{\mathcal{S}\subseteq\mathcal{K}, 1\leq|\mathcal{S}|\leq\theta} \mathrm{RP}(\widehat{\chi}, \mathcal{S})$.

Given an appropriate restricted inverse problem for other distance functions, one can produce an analog of Theorem 1. However, in general, the corresponding restricted problem is non-convex (due to the KKT conditions that ensure Pareto optimality), which may be difficult to solve. Thus, the existence of a convex formulation for $d^{[r]}$ is a special case.

An additional special case for which the restricted problem is tractable is when all of the objectives and constraints are affine, which leads to the mixed-integer linear program (4). However, many objectives in radiation therapy treatment planning are nonlinear. When $\mathcal{P}$ is a singleton, another tractable restricted problem results from the distance function $d^{[q,\mathrm{I}]}(\mathbf{x}, \widehat{\chi}_{\mathcal{K}}, \mathcal{S}) = \left(\sum_{k \in \mathcal{S}} |f_k(\mathbf{x}) - f_k(\widehat{\mathbf{x}}_k)|^q \right)^{1/q}$, a variant of $d^{[q]}$. The distance function $d^{[q,\mathrm{I}]}$ relies on multiple "ideal" solutions that are optimized for a single objective: $\widehat{\chi}_{\mathcal{K}} = \{\widehat{\mathbf{x}}_k\}_{k \in \mathcal{K}}$, where $\widehat{\mathbf{x}}_k \in \arg\min_{\mathbf{x}} \{f_k(\mathbf{x}) \mid \mathbf{A}\mathbf{x} = \mathbf{b}, \mathbf{g}(\mathbf{x}) \leq \mathbf{0}\}$, for all $k \in \mathcal{K}$ (Lin 2005); see the e-companion for details. In the radiation therapy context, however, this implies that the input data consists of $|\mathcal{K}|$ treatment plans for the patient each arising from a distinct single-objective optimization problem, which is more restrictive. Hence, for increased clinical relevance and flexibility in the inputs, we focus on using $d^{[r]}$ in the subsequent analysis and numerical results.

If the data points are assumed to come from forward problems with the same preferences, we consider the batch duality gap distance. This objective selection problem, which we call OS-W to indicate that the problem selects the same objectives and weights for all patients in the batch, is as follows:

$$\begin{aligned} \text{OS-W}(\widehat{\chi}) = \quad & \min_{\chi, \mathcal{S}} \quad \gamma(\chi, \widehat{\chi}, \mathcal{S}) \\ & \text{subject to} \quad \chi \in \widetilde{\Omega}(\mathcal{S}), \\ & \qquad\qquad\ \ 1 \leq |\mathcal{S}| \leq \theta, \\ & \qquad\qquad\ \ \mathcal{S} \subseteq \mathcal{K}, \end{aligned} \tag{OS-W}$$

where $\widetilde{\Omega}(\mathcal{S}) = \bigcup_{\boldsymbol{\alpha} \in \mathcal{A}(\mathcal{S})} \arg\min_{\chi} \{\boldsymbol{\alpha}^\top \sum_{p \in \mathcal{P}} \mathbf{f}_{(p)}(\mathbf{x}_{(p)}) \big| \mathbf{A}_{(p)}\mathbf{x}_{(p)} = \mathbf{b}_{(p)}, \mathbf{g}_{(p)}(\mathbf{x}_{(p)}) \leq 0, \ \forall\ p \in \mathcal{P}\}$. We modify the RP formulation to find the the same weights for the chosen objectives for all data points (RP-W).

$$\text{RP-W}(\widehat{\chi}, \mathcal{S}) = \quad \min_{\mathbf{x}, \boldsymbol{\alpha}} \quad \left(\boldsymbol{\alpha}^\top \sum_{p \in \mathcal{P}} \mathbf{f}_{(p)}(\mathbf{x}_{(p)}) \right) \Big/ \left(\boldsymbol{\alpha}^\top \sum_{p \in \mathcal{P}} \mathbf{f}_{(p)}(\widehat{\mathbf{x}}_{(p)}) \right) \tag{6a}$$

$$\text{subject to} \quad \mathbf{A}_{(p)}\mathbf{x}_{(p)} = \mathbf{b}_{(p)}, \ \forall\ p \in \mathcal{P}, \tag{6b}$$

$$\mathbf{g}_{(p)}(\mathbf{x}_{(p)}) \leq \mathbf{0}, \ \forall\ p \in \mathcal{P}, \tag{6c}$$

$$\boldsymbol{\alpha} \in \mathcal{A}(\mathcal{S}). \tag{6d}$$

THEOREM 2. *There exists a convex reformulation of* RP-W$(\widehat{\chi}, \mathcal{S})$ *whose Lagrange multipliers are optimal weights that minimize the batch duality gap. Also,* OS-W$(\widehat{\chi}) = \min_{\mathcal{S} \subseteq \mathcal{K}, 1 \leq |\mathcal{S}| \leq \theta}$ RP-W$(\widehat{\chi}, \mathcal{S})$.

Next we establish a relationship between the distance metrics $d^{[r]}$ and $d^{[q]}$. For the single-input case, RP-S and RP-W are equivalent; thus, we simply denote this restricted problem by RP and refrain from labeling data with the subscript (p). Given a subset $\mathcal{T} \subseteq \mathcal{K}$, let $||\cdot||_{q,\mathcal{T}} : \mathbb{R}^{|\mathcal{T}|} \to \mathbb{R}$ be the q-norm in $\mathbb{R}^{|\mathcal{T}|}$. Also, let $\mathbf{f}_{\mathcal{T}} = [f_k]_{k \in \mathcal{T}}$.

PROPOSITION 3. *Consider a single-input restricted inverse problem* RP$(\hat{\mathbf{x}}, \mathcal{S})$ *and an optimal solution* $(\mathbf{x}^*, \boldsymbol{\alpha}^*)$. *Let* $\mathcal{S}^*$ *be the support of* $\boldsymbol{\alpha}^*$, *e.g.,* $\mathcal{S}^* = \{k \in \mathcal{S} \mid \alpha_k^* > 0\}$. *Then the relative gap of* WMO$(\boldsymbol{\alpha}^*)$ *at* $\hat{x}$ *is*

$$\frac{||\mathbf{f}_{\mathcal{S}^*}(\hat{\mathbf{x}})||_{q,\mathcal{S}^*}}{||\mathbf{f}_{\mathcal{S}^*}(\hat{\mathbf{x}})||_{q,\mathcal{S}^*} - ||\mathbf{f}_{\mathcal{S}^*}(\hat{\mathbf{x}}) - \mathbf{f}_{\mathcal{S}^*}(\mathbf{x}^*)||_{q,\mathcal{S}^*}}.$$

Proposition 3 highlights a relationship between the proposed distance metrics as follows. Let $\mathcal{S}$ and $\mathcal{T}$ (with $|\mathcal{S}| = |\mathcal{T}|$) be the support of two weight vectors obtained from optimal solutions of restricted problems, and assume that all of the candidate objectives f_k have been normalized such that $f_k(\widehat{\mathbf{x}}) = 1$, for all $k \in \mathcal{S} \cup \mathcal{T}$, thus $||\mathbf{f}_{\mathcal{S}}(\widehat{\mathbf{x}})||_{q,\mathcal{S}} = ||\mathbf{f}_{\mathcal{T}}(\widehat{\mathbf{x}})||_{q,\mathcal{T}}$. Then if the minimum relative gap over $\mathcal{S}$ at $\widehat{\mathbf{x}}$ is $r_{\mathcal{S}}$ (similarly for $\mathcal{T}$) and $r_{\mathcal{S}} > r_{\mathcal{T}}$, we have $||\mathbf{f}_{\mathcal{S}}(\widehat{\mathbf{x}}) - \mathbf{f}_{\mathcal{S}}(\mathbf{x}^*_{\mathcal{S}})||_{q,\mathcal{S}} > ||\mathbf{f}_{\mathcal{T}}(\widehat{\mathbf{x}}) - \mathbf{f}_{\mathcal{T}}(\mathbf{x}^*_{\mathcal{T}})||_{q,\mathcal{T}}$; that is, in this case, a larger relative gap corresponds to a larger norm of the difference of function values. From a computational perspective, in this specific setting, optimizing with respect to the relative gap results in the same optimal objective sets as when using the q-norm. Despite this particular relationship between the relative gap and q-norm in Proposition 3, we opt to focus primarily on our formulations for optimizing the relative gap and leave q-norm-based objective selection for future research.

Recall from Section 2 that one can view the objective selection problem as a feature selection problem because the goal is to find a sparse set of objectives (no more than θ of them) of an optimization problem that fits the given input $\widehat{\mathbf{x}}$ or set of inputs $\widehat{\chi}$. Theorems 1 and 2 reinforce this connection and demonstrate that the difficult objective selection problem exhibits structure

that makes certain feature selection algorithms suitable. In particular, greedy approaches, widely used in feature selection problems, require solving a limited number of convex programs to find good sets of objectives and weights and thus may be a reasonable choice for the objective selection problems in treatment planning and other applications.

3.2.3. Regularization Because the objective selection problem is closely related to feature selection, another natural solution approach is to use regularization. Let $\boldsymbol{\lambda} \in \mathbb{R}_+^{|\mathcal{K}|}$, then a regularized objective selection problem is

$$\begin{aligned} \mathrm{RE}(\widehat{\mathbf{x}}) = \quad & \max_{\mathbf{x},\boldsymbol{\alpha}} \quad -\boldsymbol{\alpha}^\top \boldsymbol{\lambda} + \big(\boldsymbol{\alpha}^\top \mathbf{f}(\mathbf{x})\big) \,/\, \big(\boldsymbol{\alpha}^\top \mathbf{f}(\widehat{\mathbf{x}})\big) && \text{(7a)} \\ \text{subject to} \quad & \mathbf{x} \in \mathcal{O}(\boldsymbol{\alpha}), && \text{(7b)} \\ & \boldsymbol{\alpha}^\top \mathbf{f}(\widehat{\mathbf{x}}) = 1, && \text{(7c)} \\ & \boldsymbol{\alpha} \in \mathbb{R}_+^{|\mathcal{K}|}. && \text{(7d)} \end{aligned}$$

This regularized model maximizes the reciprocal of the relative gap minus the regularization term. If the regularization parameter is $\boldsymbol{\lambda} = \mathbf{0}$, then (7) solves the normal inverse problem (not the objective selection problem) over all candidate objectives because minimizing the relative gap is equivalent to maximizing its reciprocal, assuming both are well defined.

Consider the following model (8) for a single input:

$$\begin{aligned} \mathrm{REGP}(\widehat{\mathbf{x}}) = \quad & \min_{\mathbf{x},\epsilon} \quad \epsilon && \text{(8a)} \\ \text{subject to} \quad & f_k(\mathbf{x}) \le \epsilon f_k(\widehat{\mathbf{x}}) + \lambda_k, \ \forall\ k \in \mathcal{K}, && \text{(8b)} \\ & \mathbf{A}\mathbf{x} = \mathbf{b}, && \text{(8c)} \\ & \mathbf{g}(\mathbf{x}) \le \mathbf{0}. && \text{(8d)} \end{aligned}$$

Theorem 3 reveals that $\boldsymbol{\lambda}$ also serves as a regularization parameter in $\mathrm{REGP}(\widehat{\mathbf{x}})$ for the objective weights $\boldsymbol{\alpha}$.

THEOREM 3. *If an optimal primal-dual pair for* $\mathrm{REGP}(\widehat{\mathbf{x}})$ *exists with zero absolute duality gap, then the dual multipliers of* (8b) *yield optimal weights to the regularized objective selection problem* $\mathrm{RE}(\widehat{\mathbf{x}})$.

A common issue in regularization methods is determining regularization parameters such that the regularized problem mimics the original non-convex problem. In the context of objective selection, the regularization parameter $\boldsymbol{\lambda}$ must be tuned to select an objective set of the appropriate sparsity. Finally, we remark that this regularization approach can extend to accept multiple inputs if one adapts the problem RP-W$(\widehat{\mathbf{x}}_{(p)}, \mathcal{K})$ in a similar way by adding $\boldsymbol{\lambda}$ to constraints (EC.13a).

In our application of the objective selection approach to radiation therapy treatment planning, we adapt the feature selection perspectives and employ greedy algorithms to solve the computationally challenging problem. The specific algorithms that we propose further capitalize on application-specific knowledge as well as unique problem structure that enables anatomy-based heuristics. Thus, we first discuss structural details of the radiation therapy treatment planning problems in the next section, and provide detailed illustrations of our solution approaches in Section 5.

4. Objective Selection in Radiation Therapy Treatment Planning

4.1. Radiation Therapy Context

We apply objective selection to IMRT treatment planning and focus on prostate cancer. Prostate cancer is one of the most common cancer types in American men, accounting for nearly 20% of new cancer diagnoses in men in 2018 (American Cancer Society 2018). Furthermore, prostate cancer will be among the highest in cost increase of medical care from 2010 to 2020 (Mariotto et al. 2011). Although prostate cancer has a relatively high survival rate, complications due to radiation (e.g, radiation exposure to healthy organs) are still one of the biggest concerns about prostate cancer treatment (American Cancer Society 2018). The reasons above suggest it is critical to design high-quality treatment plans efficiently and consistently. Because there exist multiple clinical criteria which are generally not achievable simultaneously, clinical trade-offs and associated preferences, reflected in the treatment planning objectives (and their weights), are critical features that describe the administered treatment.

4.2. Model Formulation

We specify the models presented in Section 3 with details specific to IMRT treatment planning for prostate cancer. General goals in treatment planning are to spare organs at risk (OARs) while

delivering sufficient radiation dose to the tumor. In prostate cancer, generally four OARs are considered: the bladder, rectum, left femoral head, and right femoral head (Chanyavanich et al. 2011).

Let $\mathscr{B}$ be the set of beamlets, where each beamlet $b \in \mathscr{B}$ is associated with a decision variable w_b that determines the beamlet's radiation intensity. The patient's body is discretized into a set $\mathscr{V}$ of volume elements called voxels. The amount of dose delivered to voxel $v \in \mathscr{V}$ by beamlet $b \in \mathscr{B}$ is denoted by $D_{v,b}$ and the entire matrix by $\mathbf{D} \in \mathbb{R}^{|\mathscr{V}| \times |\mathscr{B}|}$. The anatomical regions of interest (ROIs) for prostate cancer treatment include the four OARs, the clinical target volume (CTV), the planning target volume (PTV), and the set of remaining voxels near the tumor, often referred to as normal tissue (Normal). The subsets of voxels that comprise each of the structures are denoted by $\mathscr{V}_{\text{CTV}}, \mathscr{V}_{\text{PTV}}, \mathscr{V}_{\text{Blad}}, \mathscr{V}_{\text{Rect}}, \mathscr{V}_{\text{LFem}}, \mathscr{V}_{\text{RFem}}$, and $\mathscr{V}_{\text{Normal}}$, respectively.

The objective selection problem considers multiple forms of candidate objectives. To permit a wide range of objectives to be considered, we use both maximum dose and threshold penalty objectives (both linear and quadratic) for these four organs at risk. In addition, we propose objectives for the target volumes as well. For both the CTV and PTV, we include a quadratic target dose error objective and a heterogeneity objective that measures the variance in dose to the structure. We summarize the candidate objectives below and note that all of them map from $\mathbb{R}^{|\mathscr{B}|}$ to $\mathbb{R}$.

1. Piecewise linear dose threshold functions (L1): Consider an OAR $\rho \in$ {Blad, Rect, LFem, RFem} with threshold value τ Gy. We define this objective function as $f^{L1}_{\tau,\rho}(\mathbf{w}) = \frac{1}{|\mathscr{V}_\rho|} \sum_{v \in \mathscr{V}_\rho} \max\{\sum_{b \in \mathscr{B}} D_{v,b} w_b - \tau, 0\}$. We consider threshold values $\tau \in \{0, 20, 40, 60\}$ Gy for the bladder and rectum and $\tau \in \{0, 20\}$ Gy for the femoral heads.
2. Piecewise quadratic dose threshold functions (L2): For each $\rho \in$ {Blad, Rect, LFem, RFem} with threshold value τ Gy, we define the L2 objective function as $f^{L2}_{\tau,\rho}(\mathbf{w}) = \frac{1}{|\mathscr{V}_\rho|} \sum_{v \in \mathscr{V}_\rho} \max\{\sum_{b \in \mathscr{B}} D_{v,b} w_b - \tau, 0\}^2$. We consider threshold values $\tau \in \{0, 20, 40, 60\}$ Gy for the bladder and rectum and $\tau \in \{0, 20\}$ Gy for the femoral heads.
3. Maximum dose functions (Max): For each $\rho \in$ {Blad, Rect, LFem, RFem}, the maximum dose objective function is $f^{\text{Max}}_{\rho}(\mathbf{w}) = \max_{v \in \mathscr{V}_\rho} \sum_{b \in \mathscr{B}} D_{v,b} w_b$.

Structure \ Function Type	L1	Max	L2	DE	HD
Bladder	1-4	5	6-9		
CTV				10	11
Left Femoral Head	12-13	14	15-16		
PTV				17	18
Right Femoral Head	19-20	21	22-23		
Rectum	24-27	28	29-32		

Table 1 A summary of the 32 functions used in objective selection. L1 and L2 function labels increase with the threshold magnitude.

4. Dose error functions (DE): Consider a target volume $\rho \in \{\text{CTV, PTV}\}$ with target dose ϕ_ρ. Define the target dose function as $f_\rho^{\text{DE}}(\mathbf{w}) = \frac{1}{|\mathscr{V}_\rho|} \sum_{v \in \mathscr{V}_\rho} \left(\phi_\rho - \sum_{b \in \mathscr{B}} D_{v,b} w_b \right)^2$. We consider a target dose of 80Gy for the CTV and 77 Gy for the PTV.
5. Heterogeneous dose functions (HD): For each target volume $\rho \in \{\text{CTV, PTV}\}$, define the function as $f_\rho^{HD}(\mathbf{w}) = \frac{1}{|\mathscr{V}_\rho|} \sum_{v \in \mathscr{V}_\rho} \left(\sum_{b \in \mathscr{B}} D_{v,b} w_b - \sum_{v' \in \mathscr{V}_\rho} \sum_{b \in \mathscr{B}} D_{v,b} w_b \right)^2$.

The L1 and L2 functions penalize radiation doses to OARs that exceed thresholds (i.e., L_1-norm and squared L_2-norm, respectively). We note that we do not include equivalent uniform dose functions (Niemierko 1997) or other dose-volume objectives; they do not result in a convex optimization problem in general. However, such functions have been estimated by weighted combinations of mean and maximum dose functions in the literature (Nimerierko 1999, Thieke et al. 2002), which can be represented through the 0-threshold versions of L1 and L2 objectives and the maximum dose objectives. The DE objectives can be appropriate if a specific dose is thought to be preferable for a target volume, and HD objectives can promote homogeneity in the radiation delivered to a target volume (without a specific target dose). We note that all of the objectives are nonnegative; to make them strictly positive, we add a small constant term (.01) to each of them. Table 1 summarizes the labels for the functions in the candidate objective pool.

The radiation dose required for the CTV and PTV for prostate cancer treatment is generally hard-constrained within upper and lower bounds, denoted by u_ρ and l_ρ, where $\rho \in \{\text{CTV,PTV}\}$. As the bladder, rectum, and normal tissue are close to the CTV and PTV yet radiation dose to

these structures is supposed to not exceed dose delivered to the CTV and PTV, dose upper bounds are also introduced to them, also denoted by $u_\rho, \rho \in \{\text{Blad, Rect, Normal}\}$. Additional constraints are generally introduced to discourage heterogeneous intensity maps; we require the intensity of each beamlet to be within some fixed ratio (lower and higher) of the average beamlet intensity. Denote these ratios by η_L and η_U. Based on the above notation, the forward optimization model for IMRT treatment planning is included in the e-companion (see Section EC.1).

5. Solution Approaches

As mentioned, there is a close connection between feature selection and objective selection problems. This motivates the use of sequential forward selection approaches to approximately solve the objective selection problem. We consider the problem of selecting $\theta = 6$ objectives, due to the six specific structures (bladder, rectum, left femoral head, right femoral head, CTV, and PTV). We propose three forward selection approaches to select objectives: one approach is a classical greedy algorithm that optimizes the distance function of the objective selection problem; the second approach iteratively searches through each ROI for the best objective in a greedy fashion; the third approach is a variant of the greedy algorithm that finds a solution even more efficiently by exploiting the unique structure of the problem that reflects patients' anatomical characteristics. In addition to the forward selection approaches, we implement the regularization method. We also propose extensions of these approaches to select objectives for a group of patients.

5.1. Greedy Algorithms and Regularization

In the generic greedy algorithm, given a single data input (e.g., one patient), for each currently unselected objective, we solve the restricted problem $\text{RP}(\widehat{\mathbf{x}}, \mathcal{S})$ by setting $\mathcal{S}$ to the union of the unselected objective and the current selected objective set. The objective that decreases the relative duality gap the most is added to the selected set, and the process repeats in the next iteration. We refer the reader to Nemhauser et al. (1978) for more details about the greedy algorithm. We denote this method *G-Solo* when it is applied to a single patient.

In the ROI-restricted greedy algorithm, given a prespecified ordering of the structures, at iteration k, we greedily select an objective from the k-th ROI. For instance, if the k-th ROI is the left femoral head, we solve the restricted problem considering a left femoral head objective (see Table 1) along with the currently selected objectives of the previous $k-1$ ROIs. The left femoral head objective that decreases the relative duality gap the most is added to the selected set, and we repeat the procedure for the next ROI. This method ensures that exactly one objective per structure is selected. We denote the ROI-restricted greedy algorithm by *GR-Solo*.

For the regularization approach, we solve (8) for the specified patient. We set $\boldsymbol{\lambda}$ equal to six times the vector of ones in $\mathbb{R}^{|\mathcal{K}|}$ so that the regularization term is $6||\boldsymbol{\alpha}||_1$, where $\boldsymbol{\alpha}$ is the weight vector. We note that further optimization and parameter tuning may find a regularization penalty that works robustly. For a single patient, we refer to this approach as *R-Solo*.

5.2. Anatomy-Based Approach

Our specific application in radiation therapy offers some structure that can allow us to search for objectives in a different way beyond the relative duality gap. Each objective function f_k is characterized by its ROI's matrix, $\mathbf{D}_{\rho_k}$, the type of objective (L1, L2, Max, DE, HD), and the threshold τ_k (we interpret $\tau_k = 0$ for objectives without thresholds). In particular, the $\mathbf{D}_\rho$ matrices quantify the influence of each beamlet on each voxel and thus provide geometric information on the structures. Hence, we can define a metric that indicates the similarity of two objectives from the same patient in order to select dissimilar objectives, thus reducing redundancy, that describe past planning decisions. We let $\hat{\gamma}(i,j)$ be the distance between the Gramian matrices $\mathbf{D}_{\rho_i}$ and $\mathbf{D}_{\rho_j}$ defined by Lim et al. (2019); additional details can be found in Section EC.2. Then, we define the dissimilarity between f_i and f_j by $\gamma(i,j) = 5\hat{\gamma}(i,j) + 5\mathbb{1}_{\text{type(i)}\neq\text{type(j)}} + |\tau_i - \tau_j|$, where $\mathbb{1}_{\text{type(i)}\neq\text{type(j)}}$ captures if objectives i and j are not of the same type (e.g., both L1). Thus, γ indicates the anatomical and function type similarity between two objectives from the same patient. Define the function $T : \mathcal{K}^2 \to \mathbb{R}$ by $T(\mathcal{S}) = \sum_{i,j\in\mathcal{S}} \gamma(i,j)$; we call T the total edge function because it is the sum of edge weights of a weighted complete graph where each node corresponds to each objective

$i \in \mathcal{S}$ and $\gamma(i,j)$ represents the edge weight. We refer to the set of edge weights $\{\gamma(i,j)\}_{i,j\in\mathcal{K}}$ as an *E-vector*.

In the anatomy-based greedy algorithm for a single patient, which we call *A-Solo*, the first objective function is selected by minimizing the duality gap (i.e., solving problem (5) repeatedly with each candidate objective and choosing the objective that gives the minimum duality gap). Each subsequent objective is then selected by maximizing the increase in the total edge function T (see Algorithm 1 in Section EC.2). As there is no optimization problem involved to add subsequent objectives and evaluating the total edge function is relatively easy and can be done *a priori*, this anatomy-based approach is faster than the generic greedy algorithm, G-Solo.

5.3. Batch-Input Objective Selection

Given a group of patients, we consider two different variants of the objective selection problem: (i) finding common objective sets for all patients in the same cluster and (ii) common objectives *and* weights for all patients in the same cluster.

(i) Finding common objectives: Objectives are greedily selected based on the sum of individual duality gaps, which is obtained by solving patients' individual restricted inverse problems (5).

(ii) Finding common objectives and weights: The regularization approach is generalized to apply to multiple patients. Objectives for the batch are computed as $\tilde{\mathbf{f}}(\chi) = \sum_{p\in\mathcal{P}} \mathbf{f}(\mathbf{x}_p)$, where $\mathcal{P}$ is the batch of patients. We solve the regularization problem (8) using $\tilde{\mathbf{f}}$ and input solution χ, again using the regularization penalty of $6||\boldsymbol{\alpha}||_1$.

We refer to the first approach (i) as *G-Batch-S* because it uses the greedy algorithm on the group of patients to select common objectives (not weights). In the same way, we label (ii) as *R-Batch-W* because it uses the regularization approach to find common objectives and weights for the group of patients. A summary of solution approaches developed for objective selection is displayed in Table 2. Overall, the presented solution approaches can be used to develop a treatment planning procedure that reduces the time-intensive burden of manual objective selection.

Method	Single Patient	Multiple Patients: Set (S), Weight (W)
Anatomy-based	A-Solo	
Greedy Alg.	G-Solo	G-Batch-S
Greedy Alg. by ROI	GR-Solo	
Regularization	R-Solo	R-Batch-W

Table 2 A summary of solution approaches for objective selection

6. Computational Results

In this section, we compare the performance of our proposed solution methods for the objective selection problem for prostate cancer therapy. We consider a total of 32 candidate objectives (see Table 1); because no efficient method is known to produce an optimal set objectives (outside of a prohibitively expensive exhaustive search), we do not provide results for optimal sets of objectives. Instead we focus on the performance of our solution approaches and use the normal inverse problem (without regularization or objective selection) as a benchmark, where appropriate. The stopping criterion for iterative approaches was six or seven iterations. We use 24 synthetic treatment planning data sets generated in Chan et al. (2014) and Chan and Lee (2018). Chan et al. (2014) used CERR (Computational Environment for Radiotherapy Research; (Deasy et al. 2003, 2006)) and an auxiliary optimization problem to generate the dose delivery constraint parameters (i.e., the $\mathbf{D}$ matrix in the forward optimization problem in Section EC.1) as well as input treatment plans for the objective selection problem to mimic what would have been used in the clinical setting. In this process, the authors used dose distribution data obtained from the Princess Margaret Cancer Centre as a reference by matching CTV, PTV, and OAR information and radiation beam environment parameters to the clinical data, which are available from Chan et al. (2014). All planning-related parameter data (e.g., upper and lower dose limits for the CTV, PTV, and OARs) were obtained from Chan et al. (2014) and are available in Section EC.1.

Figure 1 displays the distribution of problem sizes for instances of the restricted problem for each patient while running G-Solo. Problem sizes for instances of the restricted problem in other

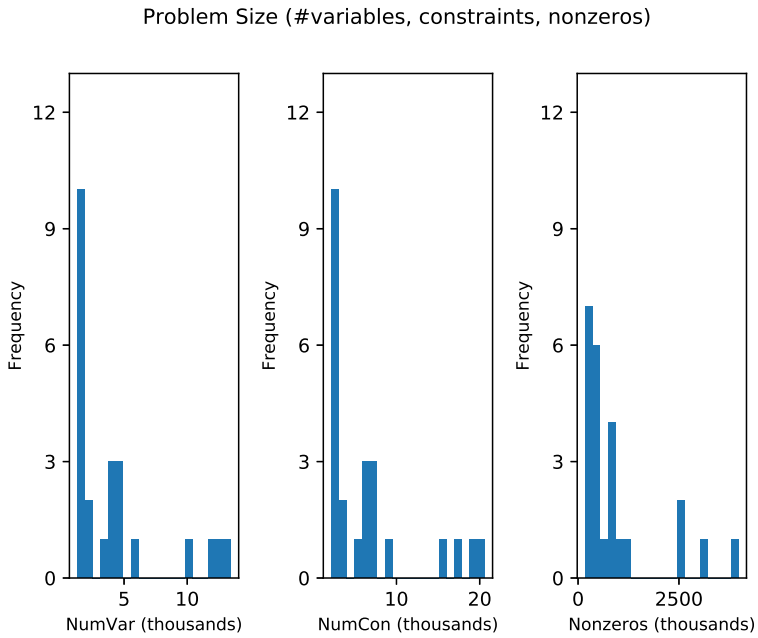

Figure 1 The problem sizes for each patient's restricted problem. The problem size data was taken from one of the G-Solo instances, although the problem size is roughly the same for each instance in each iteration of a solution method as well as the regularization problem.

iterative methods and regularization problem are roughly the same as that of G-Solo. The number or variables ranges from 1,000 to 14,000 and the number of constraints from 2,000 to 21,000.

Table 3 shows the diversity and trends of the objective sets selected by G-Solo, GR-Solo, A-Solo, and R-Solo for four patients. Results for other patients are included in the e-companion (Tables EC.1–EC.5). Where applicable, the last number in the objective indicates the dose threshold; for instance, RFem.L2.20 refers to the 20 Gy piecewise-quadratic dose threshold objective for the right femoral head. For the iterative approaches, the "Obj k" column in Table 3 indicates the k-th objective selected; for regularization, this column shows the objective with the k-th largest weight (there may be additional objectives that have positive weight).

For all of the solution approaches, no ROI dominates the objective selection; an assortment of ROIs are represented in each case, including the target volumes (CTV and PTV). For example, consider Patient 1. For all solution approaches, the first objective is either the maximum dose objective for the bladder or the analogous objective for the rectum. The maximum dose objective for

Px	Method	M.Time (s)	I.Gap	M.Gap	Obj 1	Obj 2	Obj 3	Obj 4	Obj 5	Obj 6
1	G-Solo	1065.61	1.102	1.108	Rect.Max	PTV.HD	Rect.L1.0	LFem.Max	Blad.Max	RFem.Max
1	GR-Solo	368.89	1.102	1.11	Blad.Max	Rect.Max	LFem.Max	RFem.Max	CTV.HD	PTV.DE
1	A-Solo	229.09	1.102	1.121	Rect.Max	Blad.L1.60	RFem.L2.0	LFem.L1.20	CTV.DE	RFem.L2.20
1	R-Solo	296.9	1.102	1.111	Rect.Max	Blad.Max	Rect.L1.0	Blad.L1.0	LFem.L1.0	RFem.L1.0
2	G-Solo	1013.8	1.084	1.086	Blad.Max	Blad.L1.0	RFem.Max	PTV.DE	Rect.Max	Rect.L1.0
2	GR-Solo	388.58	1.084	1.101	Blad.Max	Rect.L1.0	LFem.Max	RFem.Max	CTV.HD	PTV.HD
2	A-Solo	242.12	1.084	1.107	Blad.Max	RFem.L1.20	PTV.DE	LFem.L2.20	PTV.HD	RFem.L2.20
2	R-Solo	516.31	1.084	1.091	Blad.Max	Rect.Max	RFem.Max	Blad.L1.0	Blad.L1.20	Blad.L1.40
3	G-Solo	1406.8	1.078	1.08	Blad.Max	Blad.L1.0	Rect.L1.0	Rect.Max	PTV.DE	RFem.Max
3	GR-Solo	502.34	1.078	1.088	Blad.Max	Rect.L1.0	LFem.L1.0	RFem.L1.0	CTV.HD	PTV.DE
3	A-Solo	325.36	1.078	1.095	Blad.Max	LFem.L1.20	RFem.L2.0	CTV.DE	LFem.L2.20	RFem.L1.20
3	R-Solo	528.2	1.078	1.088	Blad.L1.0	Rect.L1.0	Blad.Max	Rect.Max	LFem.L1.0	RFem.L1.0
4	G-Solo	831.73	1.086	1.088	Rect.Max	PTV.HD	Blad.L1.0	Blad.Max	CTV.HD	PTV.DE
4	GR-Solo	314.04	1.086	1.09	Blad.Max	Rect.Max	LFem.Max	RFem.Max	CTV.HD	PTV.DE
4	A-Solo	191.12	1.086	1.093	Rect.Max	Blad.L1.60	RFem.L2.0	LFem.L1.20	PTV.DE	LFem.L2.20
4	R-Solo	480.41	1.086	1.094	Rect.Max	Blad.Max	Blad.L1.0	Rect.L1.0	RFem.L1.0	LFem.Max

Table 3 Objectives chosen by various objective selection methods and corresponding solution times and relative gaps (labeled M.Time and M.Gap, respectively) and relative gaps from the normal inverse problems (labeled I.Gap) for Patients 1–4. Where applicable, the last number in the objective indicates the threshold.

the rectum is selected by all methods for Patient 1 and the maximum dose objective for the bladder is selected by all but one method (A-Solo). Beyond these two objectives, there is an assortment of other objectives, including target volume and femoral head objectives. Similar observations can be made for other patients.

Table 3 and Tables EC.1–EC.5 show that maximum dose objectives are commonly selected amongst all methods across all patients, and this is especially true for the bladder and rectum (see highlighted entries). In all cases (patients and solution methods), either the maximum bladder dose objective or the maximum rectum dose objective is selected, and for all but two cases, one of these objectives is the first objective. These are, in fact, the most commonly selected objectives. Different threshold values are also selected for threshold penalty objectives. In addition, the dose error and heterogeneous dose objectives for the target volumes are selected at roughly the same frequency. From Tables 3 and EC.1–EC.5, we observe that typically, either G-Solo or R-Solo performs the best among our methods, in terms of the relative duality gap. We show in other results (including Figure 4) that these are also frequently the most computationally expensive methods.

While Table 3 and Tables EC.1–EC.5 show the different sets of objectives selected, Figures 2–3 and Figures EC.1–EC.2 display the progressive decrease in relative gap during the iterations of the solution approaches. For example, Figures 2–3 compare the performance of the objectives found by the solution approaches and 20 randomly selected objective sets (with optimized weights) for

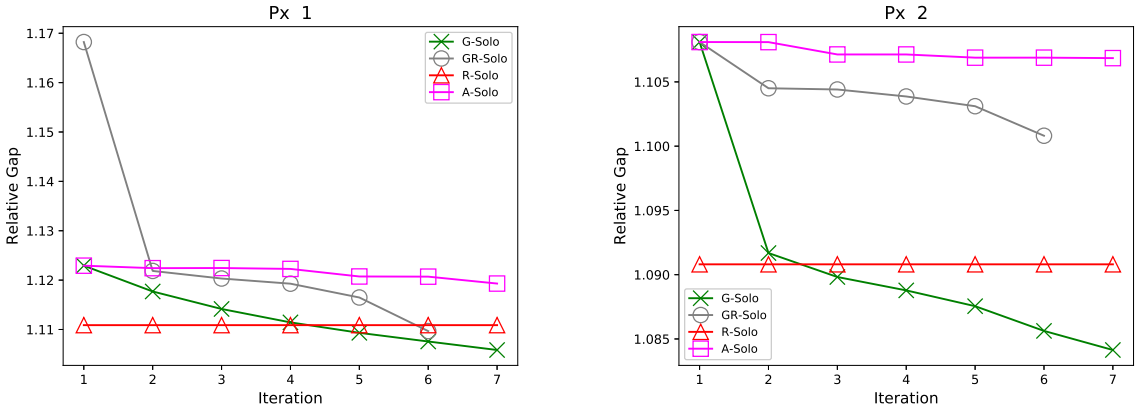

(a) Patient 1 (b) Patient 2

Figure 2 Relative duality gaps of objective selection solution approaches. The iterative methods continue to (weakly) improve at each iteration, and the relative duality gaps are bounded below by 1.

Patients 1 and 2. Figure 2 displays the relative gaps for the solution methods. For R-Solo we display the same relative gap for each iteration, as it is not an iterative method, and we show six iterations of GR-Solo because there are six ROIs with candidate objectives. We include a seventh iteration for A-Solo and G-Solo, and note that there is little improvement for the last iteration. Because some of the randomly selected objective sets produced very large duality gaps, the y-axis of Figure 3 represents the reciprocal of the duality gap. Iterations for the random objectives are determined by the order in which the objectives are randomly selected. All of the proposed solution methods perform substantially better than arbitrary objective selection in terms of relative gap minimization.

Figure 4 shows the run-times of solution approaches for four patients. Although the cumulative time of the first two or three iterations of G-Solo is cheaper than R-Solo (solved only once), the total cost of G-Solo after six greedy iterations is more expensive than any other single solution method. Hence, depending on the number of necessary iterations, either G-Solo or R-Solo could be more efficient. Recall that for A-Solo, the first objective is chosen greedily, hence the overlap with G-Solo. Each subsequent iteration (up to the penultimate) is determined from the E-vector, which leads to little subsequent cost. For R-Solo, the time to compute the weight vectors from the KKT conditions is included starting from the first iteration, while it is only included in the last

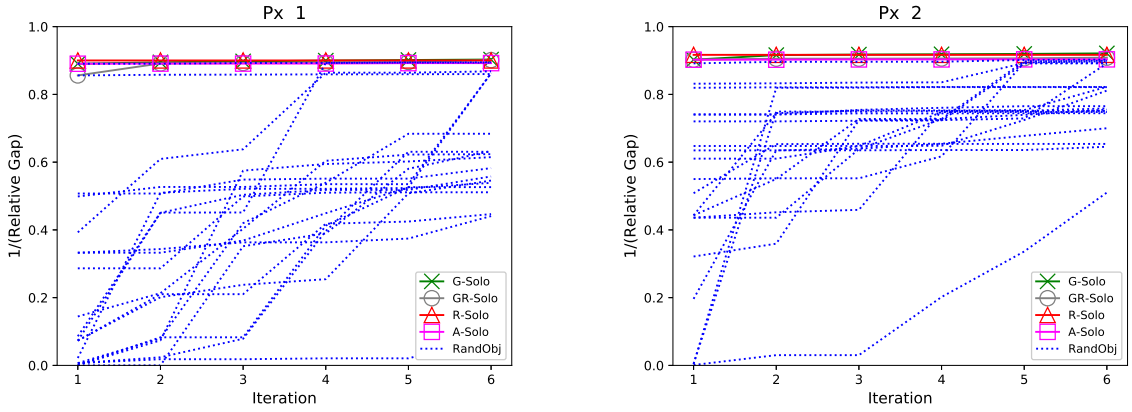

(a) Patient 1 (b) Patient 2

Figure 3 Relative duality gaps of objective selection solution approaches and random objective sets. Due to how large some random objective relative duality gaps are, we display the gap's reciprocal.

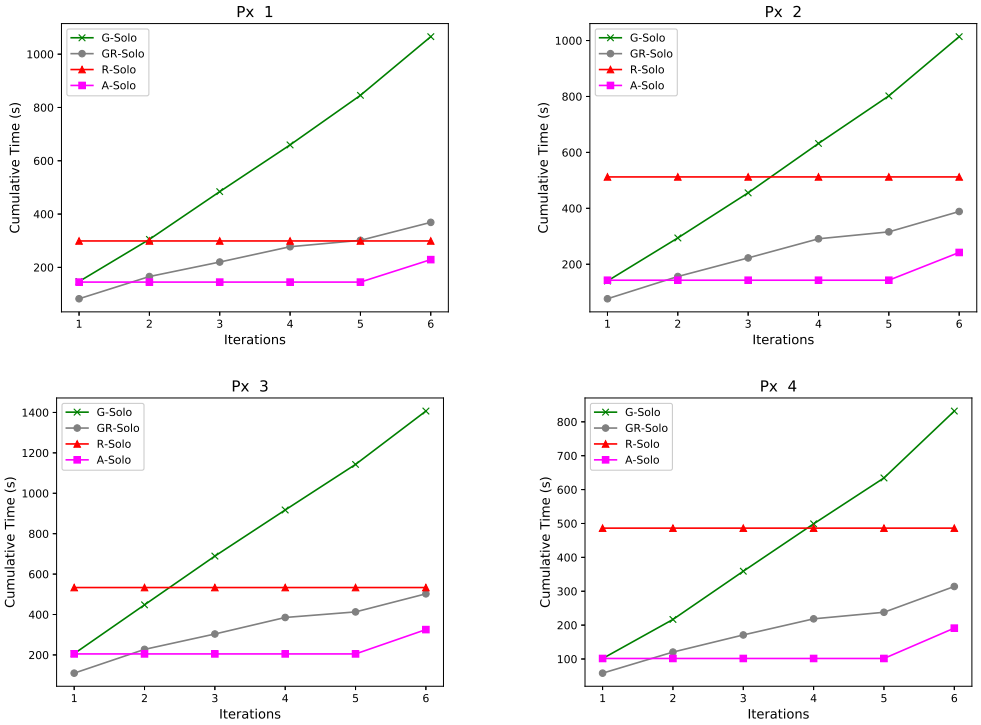

Figure 4 Run-times for various solution approaches. For the anatomy-based heuristic, the time to compute the weights is added to the last iteration.

iteration for iterative methods. Figure 5 summarizes the distribution of run-times for the solution approaches for all patients. R-Solo is typically more expensive than the other two methods, A-Solo

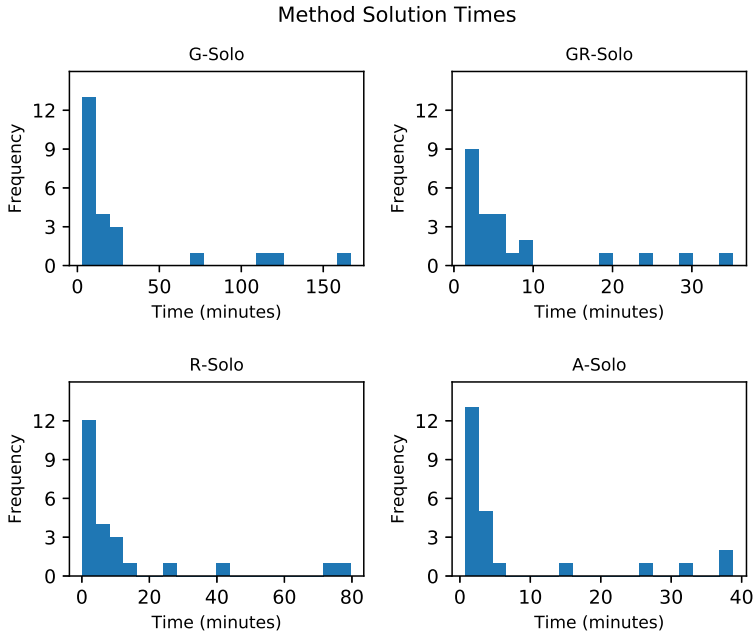

Figure 5 Solution method run-time distributions for objective selection. The time includes computation of the weight vectors after the final iteration (or regularized inverse problem). Note that the x-axes have different ranges for each solution method. The mean solve times in minutes across all patients are: 28.9 (G-Solo), 8.0 (GR-Solo), 12.2 (R-Solo), and 8.2 (A-Solo).

and GR-Solo. Most of the patients complete within a small range of time, but there are consistent outliers which take longer lengths of time, which matches the presence of a few patients who have large problem sizes (see Figure 1).

Among the proposed solution approaches, R-Solo is the only one for which the number of selected objectives cannot be directly set, and this can lead to more than θ chosen objectives. Figure 6 displays the objective weights for Patients 1-4 from R-Solo. R-Solo clearly selects a small subset of objectives for Patients 2-4, but more objectives are selected for Patient 1. In fact, for several cases R-Solo selects more than θ objectives; examples include Patients 8, 14, 18, and 22 in Figure 7. Moreover, it is not clear how to set the regularization parameter for each patient for different cardinality parameters θ; for instance, too few objectives may have been selected for Patient 2 if the regularization penalty was too high in this case. Thus, in general, using R-Solo to select a sparse set of objectives may involve an additional parameter tuning phase.

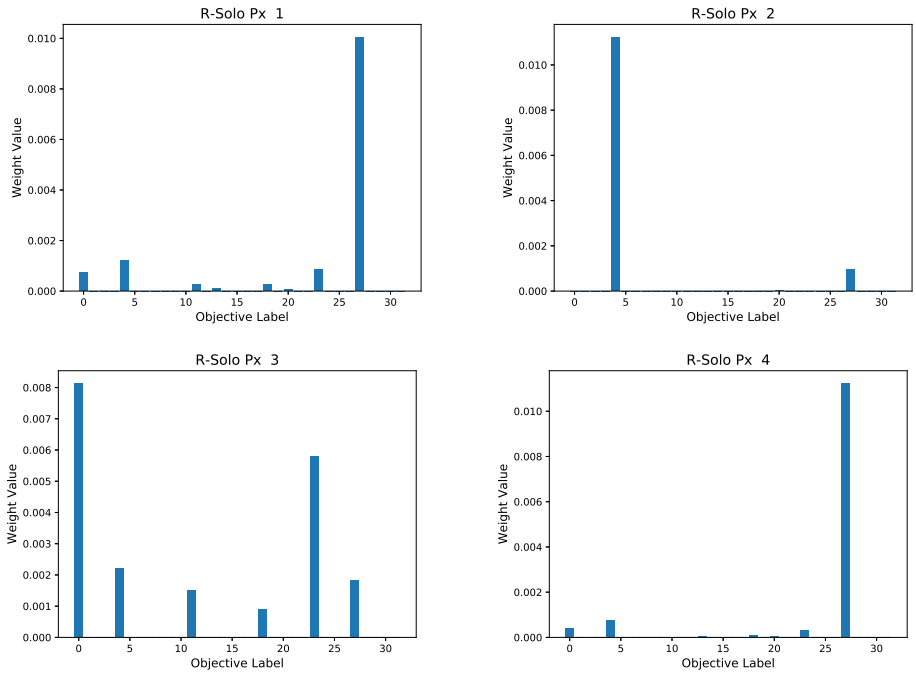

Figure 6 Objective weights for R-Solo (Patients 1–4). The number of selected objectives in the regularization method is not explicitly controlled.

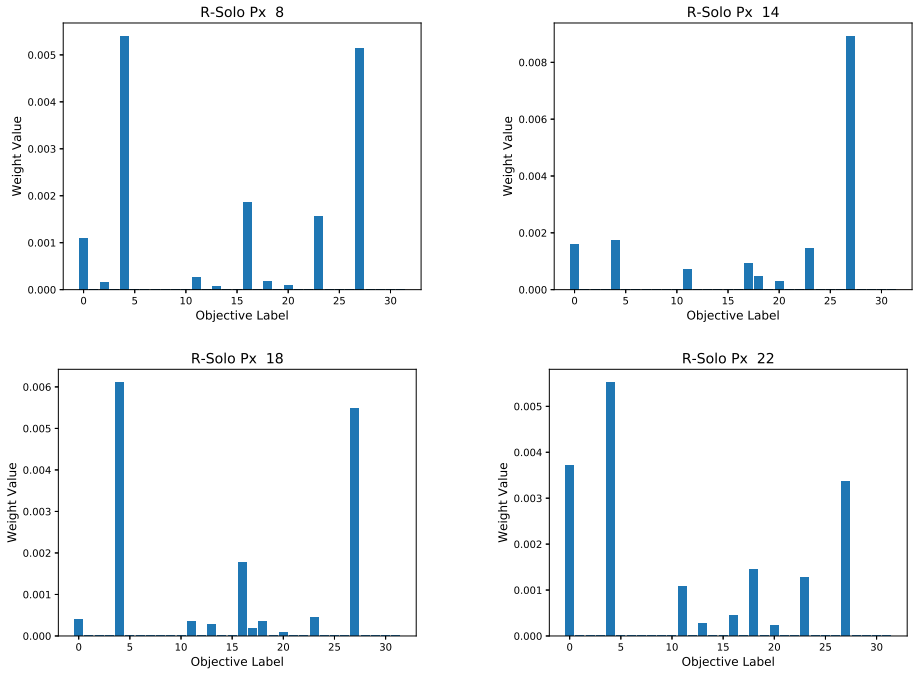

Figure 7 Objective weights for R-Solo (Patients 8,14,18,22). The number of selected objectives in the regularization method is not explicitly controlled.

Next we consider the problem of finding objectives and weights from a set of patients, as opposed to the single-patient setting. This problem is particularly relevant when there are multiple patients available in the data that can be used for objective selection for a new patient. We selected 18 out of 24 patients as the training cohort, and the rest formed the testing cohort. Recall that we consider two different batch methods: G-Batch-S and R-Batch-W. G-Batch-S was used to select an objective set $\mathcal{S}_{\text{batch}}$ for the entire group of patients. Then, the relative gap was computed for each patient (both in the training cohort and testing cohort) by solving the restricted inverse problem RP($\mathcal{S}_{\text{batch}}$). The objectives selected by G-Batch-S were: Blad.Max, Rect.Max, PTV.DE, CTV.HD, Blad.L1.0, and LFem.Max. The size of the solved optimization problems is single-patient based, as the restricted inverse problems are solved for each patient independently. The time per iteration was approximately 28 minutes with a total run-time of 2.8 hours. The second batch approach, R-Batch-W, selected weights for the entire patient set based on the eighteen patients in the training cohort. After computing weights, the forward problem for each patient was solved using these weights to compute the relative gap. The total run-time was approximately 3.6 hours (approximately 45,000 variables and 700,000 constraints).

Table 4 shows the relative duality gaps when using the single-patient and batch methods. The "Inverse" column of Table 4 displays the gaps using the normal inverse problem (neither objective selection nor regularization, see Section 3.2.2), which provides a lower bound on the relative duality gaps achievable by the other solution approaches. We observe that our approaches are often close to this lower bound. There is no apparent difference in the batch methods' performance between the training and test cohorts.

Next, we give clinical interpretations of our objective selection approaches. We use the dose-volume histogram (DVH) to compare treatment plans generated by objectives from the solution approaches and the original treatment plans. DVHs are a clinical standard for visualizing and evaluating radiation therapy treatment quality; they display the percentage volume of each structure receiving a certain level of dose or higher. Figures 8 and 9 show that, for Patients 1 and 2, respectively, the DVHs produced by the solutions approaches are similar to the clinical DVHs compared to the randomly selected objectives for the various ROIs.

Patient	In Training Cohort	G-Solo	G-Batch-S	R-Solo	R-Batch-W	A-Solo	GR-Solo	Inverse
5	No	1.06	1.06	1.08	1.15	1.08	1.06	1.06
7	No	1.1	1.1	1.11	1.1	1.12	1.11	1.1
9	No	1.06	1.06	1.06	1.09	1.07	1.06	1.05
10	No	1.08	1.08	1.08	1.08	1.09	1.08	1.07
11	No	1.08	1.08	1.08	1.1	1.09	1.09	1.07
12	No	1.06	1.06	1.07	1.13	1.08	1.07	1.06
1	Yes	1.11	1.11	1.11	1.18	1.12	1.11	1.1
2	Yes	1.09	1.09	1.09	1.1	1.11	1.1	1.08
3	Yes	1.08	1.09	1.09	1.12	1.09	1.09	1.08
4	Yes	1.09	1.09	1.09	1.13	1.09	1.09	1.09
6	Yes	1.05	1.06	1.06	1.12	1.07	1.06	1.05
8	Yes	1.1	1.11	1.11	1.16	1.17	1.11	1.1
13	Yes	1.23	1.24	1.26	1.61	1.28	1.24	1.23
14	Yes	1.29	1.32	1.35	1.64	1.4	1.3	1.28
15	Yes	1.16	1.17	1.19	1.33	1.21	1.17	1.16
16	Yes	1.21	1.21	1.25	1.36	1.28	1.21	1.2
17	Yes	1.14	1.14	1.18	1.26	1.23	1.14	1.13
18	Yes	1.31	1.31	1.33	1.41	1.43	1.31	1.3
19	Yes	1.26	1.28	1.28	1.46	1.32	1.28	1.26
20	Yes	1.26	1.26	1.28	1.37	1.34	1.26	1.25
21	Yes	1.12	1.13	1.13	1.24	1.13	1.12	1.11
22	Yes	1.24	1.24	1.26	1.46	1.35	1.25	1.22
23	Yes	1.19	1.19	1.21	1.31	1.26	1.19	1.18
24	Yes	1.25	1.25	1.27	1.77	1.3	1.25	1.24

Table 4 Relative duality gaps of patients using single patient and batch methods. The training set included 18 patients, based on which the objectives and weights were determined, which were then applied to all patients individually to compute the duality gaps.

We further compare the DVHs using shape and curve analysis via three different distance metrics: Euclidean norm, discrete Fréchet distance, and Procrustes distance. We refer the reader to Jekel et al. (2018) and references in Table 5 for additional details on these metrics. Using each metric, patient, and OAR, we compute the distance between the G-Solo DVHs and clinical DVHs as well as the distance between the DVHs generated by random, arbitrarily selected objectives and clinical DVHs for all patients. For a given patient, OAR, and metric, we use the average of the metric values over all random trials. We use the average of the random objectives trials in order to compare our results to naive, arbitrary objective selection and demonstrate that such an arbitrary approach leads to high variability in the treatment quality. The "G-Solo" columns of Table 6 show each of the metric values (for the bladder) for each patient using the solution from forward problem parametrized by the G-Solo objectives and weights; that is, it displays the distance between the G-Solo DVH and the clinical DVH. The "Random" column shows the average metric value over the twenty random trials for each patient, i.e., the average over trials of distances between the Random (trial t) DVH and the clinical DVH. For each metric and each OAR, the "G-Solo" column

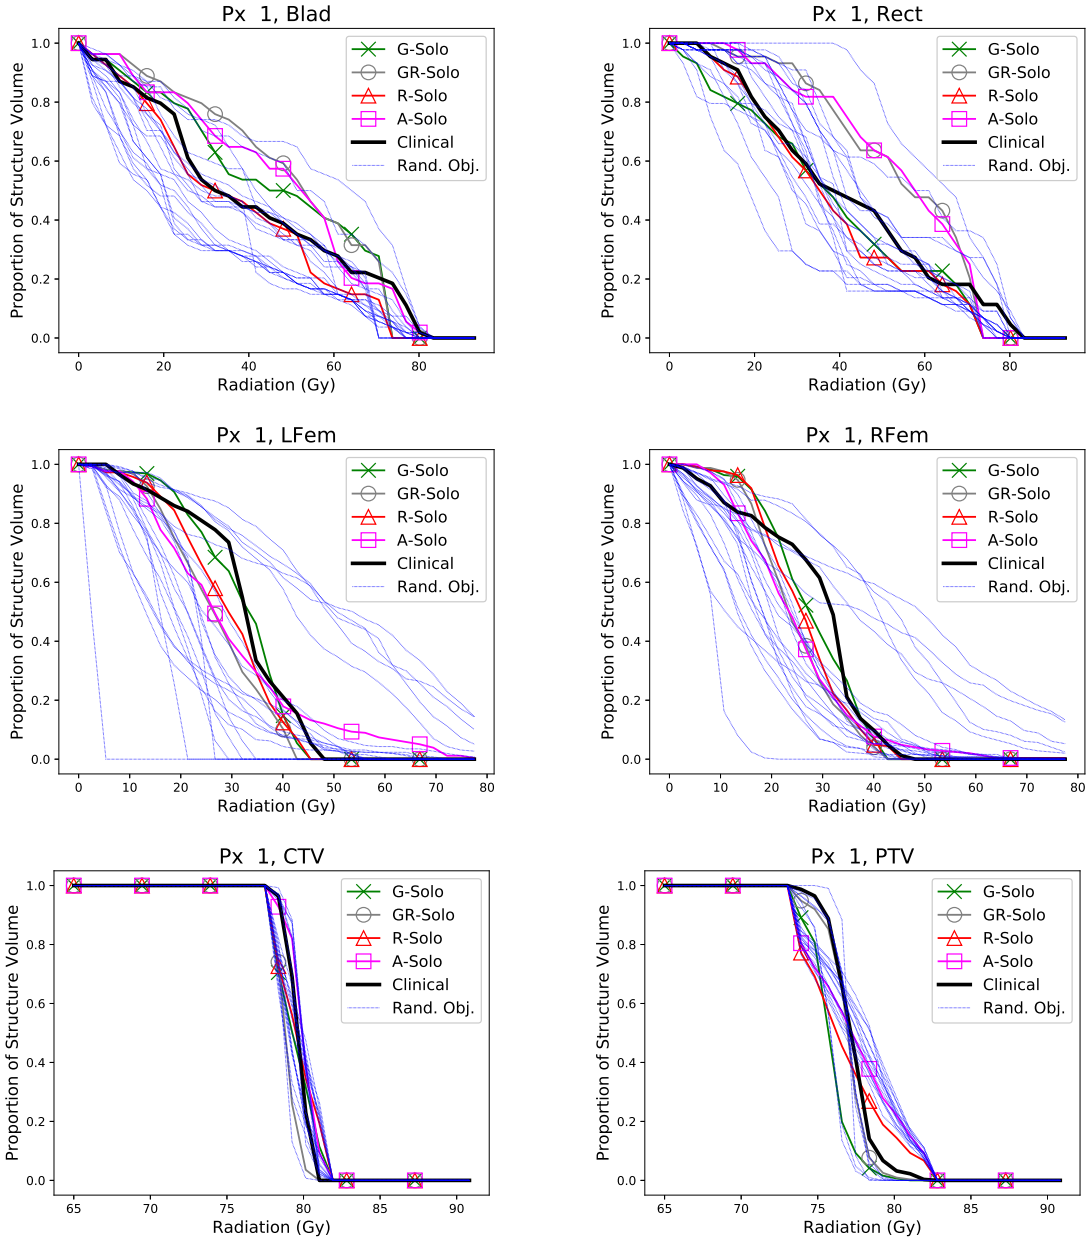

Figure 8 DVHs for various solution approaches (Patient 1).

was compared to the "Random" column using a two-sided Wilcoxon signed-rank test with a null hypothesis that there is no difference in the distribution producing each column and a significance level of .05. Table 7 shows that the p-values from the tests are less than 0.05 for all of the metrics and OARs, which indicates that the differences between the G-Solo DVHs and those generated by random objectives are statistically significant.

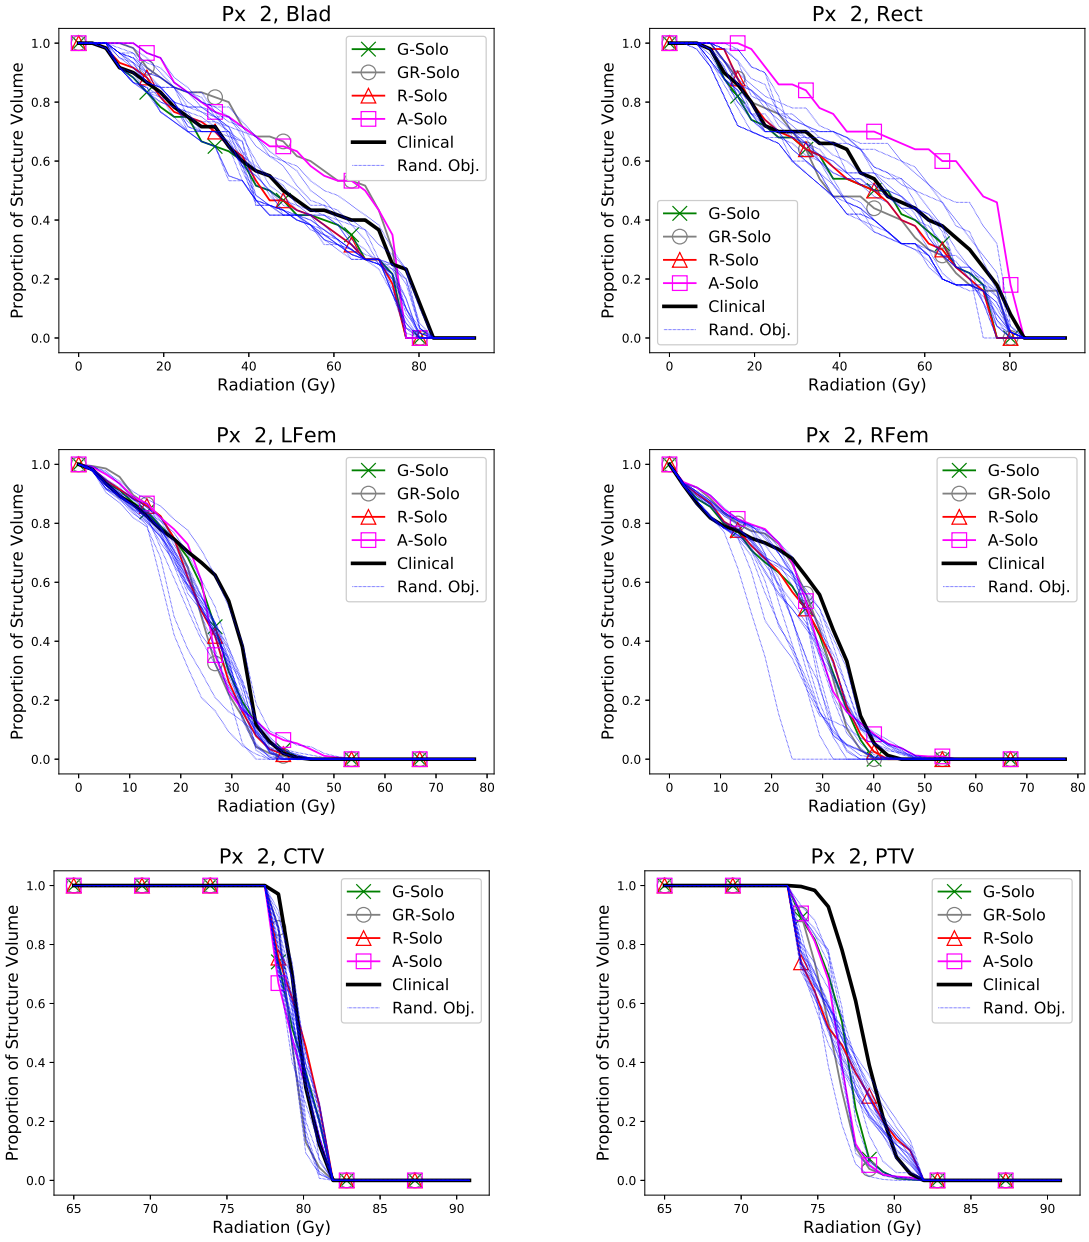

Figure 9 DVHs for various solution approaches (Patient 2).

Comparison Metric	G-Solo	Random Objectives	Referenecs
Duality Gap	$\frac{\boldsymbol{\alpha}^\top \mathbf{f}(\mathbf{x}_{\text{Clin}})}{\boldsymbol{\alpha}^\top \mathbf{f}(\mathbf{x}_G)}$	$\frac{1}{T}\sum_{t=1}^{T} \frac{\boldsymbol{\alpha}^\top \mathbf{f}(\mathbf{x}_{\text{Clin}})}{\boldsymbol{\alpha}^\top \mathbf{f}(\mathbf{x}_{R,t})}$	Chan et al. (2014)
Euclidean Norm	$\|\|\text{DVH}_{\text{Clin}} - \text{DVH}_G\|\|_2$	$\frac{1}{T}\sum_{t=1}^{T} \|\|\text{DVH}_{\text{Clin}} - \text{DVH}_{R,t}\|\|_2$	
Fréchet Distance	$\text{FD}(\text{DVH}_{\text{Clin}} - \text{DVH}_G)$	$\frac{1}{T}\sum_{t=1}^{T} \text{FD}(\text{DVH}_{\text{Clin}} - \text{DVH}_{R,t})$	Fréchet (1906)
Procrustes Distance	$\text{PD}(\text{DVH}_{\text{Clin}} - \text{DVH}_G)$	$\frac{1}{T}\sum_{t=1}^{T} \text{PD}(\text{DVH}_{\text{Clin}} - \text{DVH}_{R,t})$	Goodall (1991)

Table 5 Metrics for comparing treatment plan DVHs.

Px.	G-Solo (Euc)	Rand. (Euc)	G-Solo (Fréc)	Rand. (Fréc)	G-Solo (Proc)	Rand. (Proc)
1	2.13	2.9	0.185	0.269	7.18e-06	6.18e-06
2	1.2	1.47	0.233	0.153	2.91e-06	3.04e-06
3	0.924	1.22	0.127	0.117	1.23e-06	1.73e-06
4	1.52	2.68	0.138	0.251	3.47e-06	7.59e-06
5	1.0	1.99	0.185	0.178	2.74e-06	3.6e-06
6	1.13	1.68	0.219	0.193	3.33e-06	4.01e-06
7	1.63	2.46	0.254	0.217	5.72e-06	5.52e-06
8	1.16	2.92	0.26	0.253	4.9e-06	7.42e-06
9	1.52	2.47	0.157	0.193	3.66e-06	7e-06
10	1.09	2.01	0.154	0.166	1.71e-06	3.83e-06
11	1.13	2.18	0.21	0.205	3.94e-06	5.18e-06
12	0.772	2.0	0.106	0.181	6.96e-07	4.42e-06
13	3.62	4.8	0.462	0.435	3.51e-05	1.74e-05
14	1.18	2.84	0.176	0.315	5.28e-06	1.19e-05
15	3.2	3.26	0.28	0.336	2.35e-05	1.72e-05
16	3.5	5.35	0.3	0.46	1.41e-05	1.54e-05
17	2.15	3.62	0.3	0.333	1.06e-05	9.39e-06
18	2.96	4.75	0.24	0.394	1.12e-05	1.64e-05
19	2.09	4.12	0.182	0.364	6.04e-06	1.15e-05
20	3.21	4.84	0.357	0.407	7.06e-06	1.91e-05
21	1.0	3.06	0.217	0.285	5.15e-06	9.45e-06
22	2.38	3.61	0.238	0.371	8.38e-06	1.57e-05
23	2.52	2.45	0.238	0.233	9.33e-06	8.3e-06
24	2.76	3.48	0.4	0.434	1.7e-05	1.89e-05

Table 6 DVH comparison metrics for the bladder. The "G-Solo" columns have one entry per patient. The "Random" objective columns are the average values over 20 trials for each patient.

Figure 10 shows that there appears to be a linear relationship between the duality gap and the three distance metrics used for the bladder and rectum DVH comparisons, associated with the objectives selected by the greedy algorithm (filled circles). Similarly, for the average of the arbitrarily chosen objectives, there also appears to be a (possibly weaker) linear relationship between the corresponding duality gap and the DVH distance metrics. Similar results for the femoral heads' DVHs can be found in Figure EC.7 in the e-companion. These results provide an important clinical implication that the duality gap as a result of the restricted inverse problem associated with a certain set of objectives can be used as a metric for how similar the resulting dose distributions are to the benchmark (clincial) DVHs. The results for the DVH comparison are further supported by an individual voxel dose analysis, in which we compare the dose delivered by the clinical and G-Solo- and random-objective-based treatment plans; details are in Tables EC.9 and EC.10. These tables show that the doses from the G-Solo treatment plans are more similar to the clinical treatment plan doses than those of the random objective treatment plans.

ROI	Euclidean	Fréchet	Procrustes
Bladder	2.4e-05	0.0093	0.018
Rectum	1.8e-05	0.0021	0.00036
L. Femur	0.028	0.0086	0.003
R. Femur	0.0018	0.00044	0.00044

Table 7 The p-values from a comparison of DVH curves (G-Solo vs. Random) using Wilcoxon signed-rank tests.

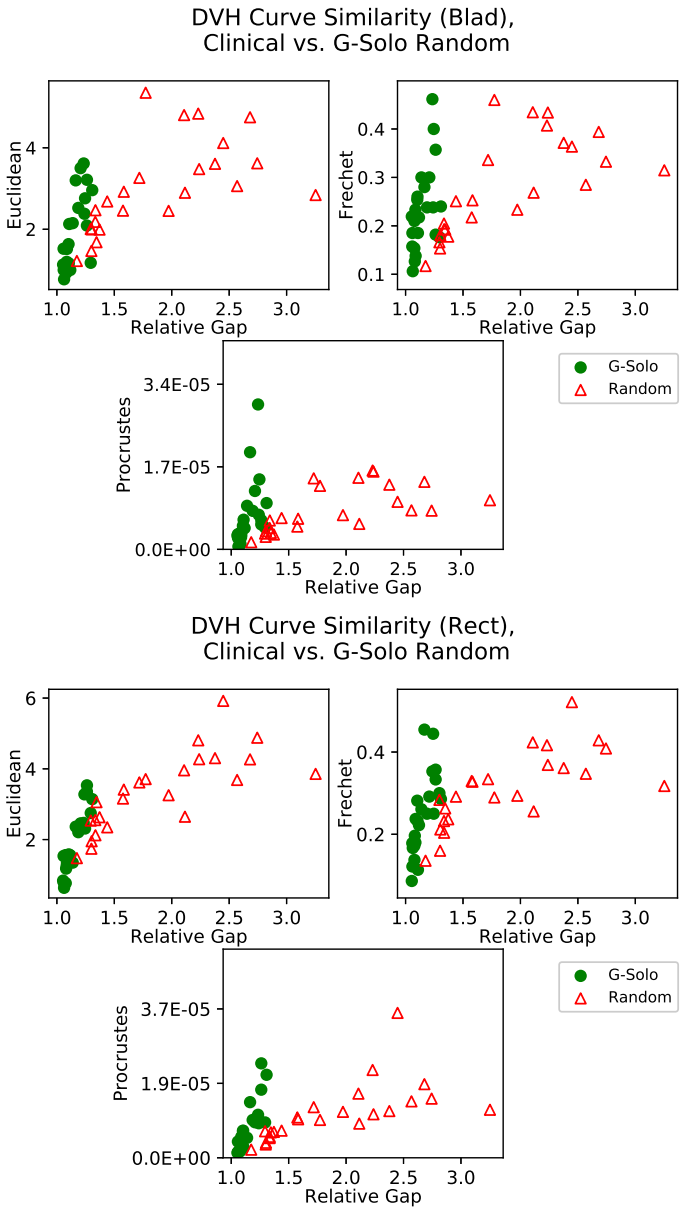

Figure 10 Relationship between distance metrics for DVH similarity and the duality gap (Bladder and Rectum).

In summary, we observe that individualized treatment plans produced by recovering clinical preferences in the form of planning objectives and weights closely replicate the original clinical plans. Moreover, empirically, the relative gaps produced by G-Solo and R-Solo are similar. We observe some trade-offs between the proposed approaches. The iterative solution approaches can be beneficial when an appropriate θ is unknown because they allow the user to continue to run the greedy algorithm until the relative gap decreases below a tolerance and thus dynamically select

a small, but representative set of objectives. On the other hand, the regularization approach (R-Solo) has no direct control over the size of the set of selected objectives, and the regularization penalty parameter may require tuning. However, the iterative approaches (e.g., G-Solo) typically become computationally more expensive than R-Solo when the number of objectives to be selected increases.

We also observe that the objectives selected from the training cohort also resulted in small relative gaps for the testing cohort. Hence, determining a set of objectives and weights based on an existing batch of patients may lead to useful IMRT planning parameters for future similar patients, which can minimize the manual, time-consuming objective selection process in current practice. Thus, our proposed methodology can lead to more efficient and consistent treatment planning. Overall, our results support the following potential treatment planning procedure:

1. Given an initial database of patients, cluster patients according to appropriate anatomical and clinical characteristics. For each cluster, find objectives (using G-Batch-S) or objectives and weights (using R-Batch-W).
2. For each new patient p:
   - 2.a Match patient p to some cluster c.
   - 2.b If using G-Batch-S: find a similar patient p' within cluster c. Solve the restricted inverse problem for patient p' using the objectives from G-Batch-S, and obtain the weights. Use these weights for the forward planning problem of patient p.
   - 2.c If using R-Batch-W: use the objectives and weights from R-Batch-W applied to cluster c for the forward planning problem of patient p.

Clustering patients is beyond the scope of this paper; however, finding patients with similar characteristics can be done by analyzing the patients' anatomical geometry or CT images (Wu et al. 2009, Chanyavanich et al. 2011, Moore et al. 2011). Our results, as a proof of concept, demonstrate the ability of our solution approaches to select objectives for a single patient and groups of patients, which potentially can improve future planning procedures as described above.

7. Conclusion

In this paper, we introduce the objective selection problem that finds a set of convex objectives that capture the preferences in input data and produces weakly Pareto optimal solutions with respect to these preferences. We formulate the problem as a non-convex bilevel mixed-integer program using inverse optimization and propose iterative and regularized approaches to approximately solve it by connecting objective selection with feature selection. Moreover, we demonstrate that we can extend our objective selection framework to multiple patients. The proposed solution approaches for the objective selection problem can lead to efficient and streamlined treatment planning that avoids time-consuming and often unguided trial and error. Future work includes exploration into other distance metrics for the restricted inverse problem as well as further analysis of the anatomy-based greedy algorithm.

Acknowledgements

The authors thank the referees, the associate editor, David Mildebrath, Silviya Valeva, and Saumya Sinha of Rice University for helpful comments. This project was funded in part by NSF grants CMMI-1826297 and CMMI-1933373.

References

Ahuja R, Orlin J (2001) Inverse optimization. *Operations Research* 49(5):771–783.

Amaldi E, Kann V (1998) On the approximability of minimizing nonzero variables or unsatisfied relations in linear systems. *Theoretical Computer Science* 209(1):237 – 260.

American Cancer Society (2018) Cancer Facts & Figures 2018.

Aswani A, Shen ZJM, Siddiq A (2018) Inverse optimization with noisy data. *Operations Research* 66(3):870–892.

Babier A, Boutilier JJ, McNiven AL, Chan TCY (2018) Knowledge-based automated planning for oropharyngeal cancer. *Medical Physics* 45(7):2875–2833.

Bazán FSV, Francisco JB (2009) An improved fixed-point algorithm for determining a Tikhonov regularization parameter. *Inverse Problems* 25(4).

Bertsimas D, Gupta V, Paschalidis IC (2015) Data-driven estimation in equilibrium using inverse optimization. *Mathematical Programming* 153(2):595–633.

Bortfeld T (1999) Optimized planning using physical objectives and constraints. *Seminars in Radiation Oncology* 9(1):20–34.

Boutilier JJ, Lee T, Craig T, Sharpe MB, Chan TCY (2015) Models for predicting objective function weights in prostate cancer IMRT. *Medical Physics* 42(4):1586–1595.

Breedveld S, Storchi PRM, Keijzer M, Heijmen BJM (2006) Fast, multiple optimizations of quadratic dose objective functions in IMRT. *Physics in Medicine & Biology* 51:3569–3579.

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

Chan T, Craig T, Lee T, Sharpe M (2014) Generalized inverse multiobjective optimization with application to cancer therapy. *Operations Research* 62(3):680–695.

Chan TCY, Lee T (2018) Trade-off preservation in inverse multi-objective convex optimization. *European Journal of Operational Research* 270(1):25–39.

Chan TCY, Lee T, Terekhov D (2019) Inverse optimization: Closed-form solutions, geometry, and goodness of fit. *Management Science* 65(3):1115–1135.

Chandrashekar G, Sahin F (2014) A survey on feature selection methods. *Computers and Electrical Engineering* 40:16–28.

Chanyavanich V, Das SK, Lee WR, Lo JY (2011) Knowledge-based IMRT treatment for prostate cancer. *Medical Physics* 38:2515–2522.

Choi B, Deasy JO (2002) The generalized equivalent uniform dose function as a basis for intensity-modulated treatment planning. *Physics in Medicine & Biology* 47(20):3579–3589.

Cotrutz C, Lahanas M, Kappas C, Baltas D (2001) A multiobjective gradient-based dose optimization algorithim for external beam conformal radiotherapy. *Physics in Medicine & Biology* 46:2161–2175.

Craft D (2011) A guide to using multi-criteria optimization (MCO) for IMRT planning in Raystation. Technical report, Harvard Medical School, Department of Radiation Oncology.

Craft D, Halabi T, Shih HA, Bortfeld T (2007) An approach for practical multiobjective IMRT treatment planning. *International Journal of Radiation Oncology, Biology, Physics* 69(5):1600–1607.

Craft DL, Hong TS, Shih HA, Bortfeld TR (2012) Improved planning time and plan quality through multicriteria optimization for intensity-modulated radiotherapy. *International Journal of Radiation Oncology, Biology, Physics* 82(1):e83–e90.

Das A, Kempe D (2011) Submodular meets spectral: Greedy algorithms for subset selection, sparse approximation and dictionary selection. *Proceedings of the 28th International Conference on International Conference on Machine Learning*, 1057–1064.

Deasy J, Lee E, Bortfeld T, Langer M, Zakarian K, Alaly J, Zhang Y, Liu H, Mohan R, Ahuja R, Pollack A, Purdy J, Rardin R (2006) A collaboratory for radiation therapy treatment planning optimization research. *Annals of Operations Research* 148(1):55–63.

Deasy JO, Blanco AI, Clark VH (2003) CERR: A computational environment for radiotherapy research. *Medical Physics* 30(5):979–985.

Duan Z, Wang L (2011) Heuristic algorithms for the inverse mixed integer linear programming problem. *Journal of Global Optimization* 51:463–471.

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

Esfahani PM, Shafieezadeh-Abadeh S, Hanasusanto GA, Kuhn D (2018) Data-driven inverse optimization with imperfect information. *Mathematical Programming* 167(1):191–234.

Eskelinen P, Miettinen K (2012) Trade-off analysis approach for interactive nonlinear multiobjective optimization. *OR Spectrum* 34(4):803–816.

Fréchet MM (1906) Sur quelques points du calcul fonctionnel. *Rendiconti del Circolo Matematico di Palermo (1884-1940)* 22(1):1–72.

Galatsanos NP, Katsaggelos AK (1992) Methods for choosing the regularization parameter and estimating the noise variance in image restoration and their relation. *IEEE Transactions on Image Processing* 1(3):322–336.

Gebken B, Peitz S (2021) Inverse multiobjective optimization: Inferring decision criteria from data. *Journal of Global Optimization* 80:3–29.

Goodall C (1991) Procrustes methods in the statistical analysis of shape. *Journal of the Royal Statistical Society: Series B (Methodological)* 53(2):285–321.

Gottlieb J, Puchta M, Solnon C (2003) A study of greedy, local search, and ant colony optimization approaches for car sequencing problems. Cagnoni S, et al., eds., *Applications of Evolutionary Computing*, 246–257 (Berlin, Heidelberg: Springer Berlin Heidelberg).

Guyon I, Elisseeff A (2003) An introduction to variable and feature selection. *Journal of Machine Learning Research* 3:1157–1182.

Halabi T, Craft D, Bortfeld T (2006) Dose-volume objectives in multi-criteria optimization. *Physics in Medicine & Biology* 51:3809–3818.

Iyengar G, Kang W (2005) Inverse conic programming with applications. *Operations Research Letters* 33:319–330.

Jekel CF, Venter G, Venter MP, Stander N, Haftka RT (2018) Similarity measures for identifying material parameters from hysteresis loops using inverse analysis. *International Journal of Material Forming* 12(3):355–378.

Kearney V, Chan JW, Haaf S, Descovich M, Solberg TD (2018) DoseNet: A volumetric dose prediction algorithm using 3d fully-convolutional neural networks. *Physics in Medicine & Biology* 63(23):235022.

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

Krause A, Singh A, Guestrin C (2008) Near-optimal sensor placement in Gaussian processes: Theory, efficient algorithms and empirical studies. *Journal of Machine Learning* 9:235–284.

Lamperski JB, Schaefer AJ (2015) A polyhedral characterization of the inverse-feasible region of a mixed-integer program. *Operations Research Letters* 43(6):575–578.

Lee T, Hammad M, Chan TCY, Craig T, Sharpe MB (2013) Predicting objective function weights from patient anatomy in prostate IMRT treatment planning. *Medical Physics* 40(12):121706.

Lim L, Sepulchre R, Ye K (2019) Geometric distance between positive definite matrices of different dimensions. *IEEE Transactions on Information Theory* 65(9):5401–5405.

Lin JG (2005) On min-norm and min-max methods of multi-objective optimization. *Mathematical Programming* 103:1–33.

Mahmood R, Babier A, McNiven A, Diamant A, Chan TCY (2018) Automated treatment planning in radiation therapy using generative adversarial networks. Doshi-Velez F, Fackler J, Jung K, Kale D, Ranganath R, Wallace B, Wiens J, eds., *Proceedings of the 3rd Machine Learning for Healthcare Conference*, volume 85 of *Proceedings of Machine Learning Research*, 484–499 (PMLR).

Mariotto AB, Yabroff KR, Shao Y, Feuer EJ, Brown ML (2011) Projections of the cost of cancer care in the United States: 2010-2010. *Journal of the National Cancer Institute* 103(2):117–128.

Martí R (2015) Black-box solvers in combinatorial optimization. *2015 International Conference on Industrial Engineering and Systems Management (IESM)*, 2–2.

McIntosh C, Purdie TG (2016) Voxel-based dose prediction with multi-patient atlas selection for automated radiotherapy treatment planning. *Physics in Medicine & Biology* 62:415.

Moore KL, Brame RS, Low DA, Mutic S (2011) Experience-based quality control of clinical intensity-modulated radiotherapy planning. *International Journal of Radiation Oncology, Biology, Physics* 81(2):545–551.

Naghavi M, Foroughi AA, Zarepisheh M (2019) Inverse optimization for multi-objective linear programming. *Optimization Letters* 13(2):281–294.

Nemhauser G, Wolsey L, Fisher M (1978) An analysis of approximations for maximizing submodular set functions - I. *Mathematical Programming* 14(1):265–294.

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

Nimerierko A (1999) A generalized concept of equivalent uniform dose. *Medical Physics* 26:1100.

Pudil P, Novovičová J, Kittler J (1994) Floating search methods in feature selection. *Pattern Recognition Letters* 15(11):1119–1125.

Romeijn HE, Ahuja RK, Dempsey JF, Kumar A (2006) A new linear programming approach to radiation therapy treatment planning problems. *Operations Research* 54(2):201–216.

Romeijn HE, Dempsey JF, Li JG (2004) A unifying framework for multi-criteria fluence map optimization models. *Physics in Medicine & Biology* 49(10):1991–2013.

Saeys Y, Inza In, Larrañaga P (2007) A review of feature selection techniques in bioinformatics. *Bioinformatics* 23(19):2507–2517.

Schaefer AJ (2009) Inverse integer programming. *Optimization Letters* 3(4):483–489.

Shamaiah M, Banerjee S, Vikalo H (2010) Greedy sensor selection: Leveraging submodularity. *49th IEEE Conference on Decision and Control (CDC)*, 2572–2577.

Shao L, Ehrgott M (2008) Approximately solving multiobjective linear programmes in objective space and an application in radiotherapy treatment planning. *Mathematical Methods of Operations Research* 68(2):257–276.

Shepard DM, Ferris MC, Olivera GH, Mackie TR (1999) Optimizing the delivery of radiation therapy to cancer patients. *SIAM Review* 41(4):721–744.

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

van Haveren R, Breedveld S, Keijzer M, Voet P, Heijmen B, Ogryczak W (2017) Lexicographic extension of the reference point method applied in radiation therapy treatment planning. *European Journal of Operational Research* 263(1):247–257.

Wang L (2009) Cutting plane algorithms for the inverse mixed integer linear programming problem. *Operations Research Letters* 37(2):114–116.

Wilkens JJ, Alaly JR, Zakarian K, Thorstad WL, Deasy JO (2007) IMRT treatment planning based on prioritizing prescription goals. *Physics in Medicine & Biology* 52(6):1675–1692.

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

Wu B, Ricchetti F, Sanguineti G, Kazhdan M, Simari P, Jacques R, Taylor R, McNutt T (2011) Data-driven approach to generating achievable dose-volume histogram objectives in intensity-modulated radiotherapy planning. *International Journal of Radiation Oncology, Biology, Physics* 79(4):1241–1247.

Wu Q, Mohan R, Niemierko A, Schmidt-Ullrich R (2002) Optimization of intensity-modulated radiotherapy plans based on the equivalent uniform dose. *International Journal of Radiation Oncology, Biology, Physics* 52(1):224–235.

Xing L, Li JG, Donaldson S, Le QT, Boyer AL (1999) Optimization of importance factors in inverse planning. *Physics in Medicine & Biology* 44(10):2525–2536.

Zhang J, Liu Z (1996) Calculating some inverse linear programming problems. *Journal of Computational and Applied Mathematics* 72(2):261–273.

Zhang X, Li X, Quan EM, Pan X, Li Y (2011) A methodology for automatic intensity-modulated radiation treatment planning for lung cancer. *Physics in Medicine & Biology* 56:3873–3893.

Zhang X, Wang X, Dong L, Liu H, Mohan R (2006) A sensitivity-guided algorithm for automated determination of IMRT objective function parameters. *Medical Physics* 33(8):2935–2944.

Temitayo Ajayi is an Operations Research Analyst at Nature Source Improved Plants, where he uses optimization to improve genomic selection for agricultural crops. His additional research interests include stochastic optimization and healthcare operations research applications, including various aspects of cancer treatment decision-making.

Taewoo Lee is an assistant professor of Industrial Engineering at the University of Houston. His primary research interests include inverse optimization and optimization under uncertainty and their applications to healthcare operations and medical decision-making problems.

Andrew J. Schaefer is the Noah Harding Chair and Professor of Computational and Applied Mathematics at Rice University. His research interests include stochastic optimization methodology and its applications to healthcare problems. In particular, he is interested in optimizing decisions arising in the treatment of a variety of diseases, including end-stage liver disease, HIV/AIDS, influenza, and cancer.