Knowledge-Based Automated Planning for Oropharyngeal Cancer

by

Aaron Babier

A thesis submitted in conformity with the requirements
for the degree of Master of Applied Science
Graduate Department of Mechanical and Industrial Engineering
University of Toronto

Abstract

Knowledge-Based Automated Planning for Oropharyngeal Cancer

Aaron Babier

Master of Applied Science

Graduate Department of Mechanical and Industrial Engineering

University of Toronto

2017

We automatically generate intensity-modulated radiation therapy plans for oropharyngeal cancer by combining knowledge-based planning (KBP) predictions with an inverse optimization (IO) pipeline into a single automated treatment planning pipeline. We extended two existing KBP methods, which use patients' anatomical geometry to predict achievable dose volume histograms (DVHs), and developed the first IO method that takes DVHs as direct inputs. The DVH predictions from KBP are put into the IO pipeline to automatically generate treatment plans via an intermediate step using objective function weights and an inverse planning problem. This step enables our automated planning pipeline to seamlessly fuse with the current treatment planning paradigm to increase its efficiency. Our automated pipeline can replicate, and often improve upon the clinical treatment plans by reducing the dose to healthy tissue and increasing primary target coverage. These results have been validated using a large cohort of 217 oropharyngeal cancer patients.

Dedicated to Leslie.

Acknowledgements

I would like to thank the late Michael Sharpe for his guidance and encouragement throughout the early stages of this thesis. Mike was a tremendous teacher and a truly inspirational mentor.

Of course, thank you Professor Timothy Chan for being an outstanding advisor. I am incredibly grateful for all of your support both inside and out of academics. Working with you has been the most rewarding part of my life as a student, and I could not be more excited to continue working under you in the coming years.

Thank you to my committee members Andrea McNiven and Professor Scott Sanner for providing valuable insights on this work. I would like to extend an extra special thank you to Andrea for taking the time over the past two years to help me better understand medical physics.

I also need to thank my fellow labmates – Philip, Justin, Chris, Neal, Islay, Rafid, Ian, and Ben – who made our lab an fun place to learn. It was always comforting to come into the supportive and stimulating environment.

Thank you Minha for all of your help and encouragement over the past year. I have learned so much from you, and I cannot imagine finishing this without you.

Lastly, thank you to my parents for supporting me since day one, and always taking an interest in my development.

Contents

Chapter 1

Introduction

Cancer is the leading cause of death in Canada, and nearly half of all Canadians are expected to be diagnosed with this disease in their lifetime.[5] Treatment methods for cancer include surgery, chemotherapy, and radiation therapy. In radiation therapy, the tissues identified as cancerous, or potentially cancerous, are designated as targets. During treatment, the ionizing energy is directed toward one or more targets, and healthy structures are inevitably damaged in the process. The most important healthy structures are referred to as organs at risk (OARs), and minimizing the damage to them is critical for a successful treatment.

External beam radiation therapy is a branch of treatments that includes intensity-modulated radiation therapy (IMRT) and volumetric arc therapy (VMAT). Both modalities deliver dose to a patient with a beam projected from a linear accelerator (LINAC). The key difference is the angles at which the treatment is delivered. IMRT delivers radiation from a small set of angles selected from a 360° arc around the patient, but VMAT delivers radiation in a continuous 360° arc. Although we focus exclusively on using IMRT, it is straightforward to apply our methodology to VMAT. Our methods are developed specifically for the treatment of oropharyngeal cancer, for its proximity to OARs renders it one of the most difficult sites to treat. Throughout this thesis, we validate our techniques using a dataset of 217 oropharyngeal cancer treatment plans that were delivered at Princess Margaret Cancer Centre located in Toronto, Canada.

1.1 Intensity-Modulated Radiation Therapy

The IMRT treatment process can be decomposed into three main steps. First, the patient's anatomy is imaged using computed tomography (CT) and magnetic resonance imaging. Second, a treatment plan, which dictates how the treatment is delivered, is devised. Third, the treatment is delivered to the patient. Our contributions are to the optimization portion of devising the IMRT plan, but we outline the complete process below to provide context for the reader.

The IMRT treatment process begins with imaging the anatomy of a patient via a CT scan. The CT scan produces a series of two-dimensional cross-section images, and those images, called slices, can be combined to create a three-dimensional representation the patient's anatomy. Next, the relevant OARs and targets for the treatment are contoured on all slices; collectively, these OARs and targets are referred to as regions of interest (ROIs) . The CT scans also provide information on how radiation travels through the patient. This can be used to calculate the dose that will be deposited in various tissues from a radiation beam of known intensity.

Next, the treatment plan is devised with three primary components: the angles from which radiation beams are delivered, the desired intensity modulation of the beam at each angle, and the shape that the beam must take in order to achieve the desired modulation. The purpose of this plan is to direct the LINAC, which delivers the plan, as it rotates around the patient's bed during the treatment. The LINAC stops at several angles in its rotation to irradiate the patient. At each angle, the LINAC produces a radiation beam of constant intensity; while the window through which the beam exits and travels towards the patient takes a series of irregular shapes called apertures. The window is created by a multi-leaf collimator (MLC), which consists of numerous rows of paired tungsten strips. Each pair consists of tungsten strips on opposing sides of the collimator that can slide toward the centre within its row. The tungsten strips block radiation from exiting the LINAC and the movement allows the MLC to form apertures. As different areas of the MLC permit emission of radiation for different durations of time, the dose delivered at each angle becomes non-uniform across the irradiated region. The cumulative dose that is delivered to each structure should correspond to the dose distribution of the treatment plan. Ideally, the apertures and beam angles are selected such that the cancer receives sufficient radiation while the dose delivered to OARs is minimized. However, making favourable selections

is a difficult non-convex optimization problem. Across a population of clinical treatment plans, beam angles vary only slightly, and are often approximately, if not exactly, equidistant. In contrast, the apertures used in treatment plans are more variable.

Approximations are generally made to render the aperture selection problem tractable. Specifically, fluence map optimization (FMO) problems use linear approximations to represent the aperture selection, with the assumption that the beam angles are already known. To construct an FMO problem, the beam at each angle is decomposed into a grid of small beamlets, which become the decision variables in a multi-objective optimization problem. In a treatment plan generated by FMO, each beamlet is assigned an intensity level, and the collective grid of all beamlet intensities is called a fluence map. A good fluence map delivers sufficient dose to the target without damaging OARs, but finding the best patient-specific tradeoffs remains a challenging part of the planning process. Additionally, the fluence map should have low heterogeneity; which means that that the intensities of neighbouring beamlets should be relatively similar. The sum of positive gradients (SPG) is a metric that quantifies heterogeneity.[9] It is reflective of the plan's complexity, and it is directly correlated to the delivery time of the treatment plan. A high SPG value indicates an unrealistic treatment plan due to its long delivery time. As well, the prolonged treatment time results in greater accumulation of minor dose leakage, which ultimately exposes the patient to more radiation. In this thesis, we generate treatment plans of realistic complexity by solving FMO problems that include an SPG objective function.

Treatment planning is driven largely by trial-and-error. The process involves human treatment planners who iteratively tune objective function weights in an inverse planning problem (IPP), which is a multi-objective optimization problem, until they generate a plan that is deemed acceptable by a physician. When the plan is unacceptable, planners may also change the form of the IPP by introducing new objectives and constraints. They may also resort to contouring artificial structures, which are structures without anatomical meaning, around undesirable dosimetric features such as dose hot spots. As a result, every plan produced by formulating and solving an IPP is highly dependent on the skill and experience of the treatment planner. In response, researchers have attempted to revise the current planning process by developing methods for knowledge-based planning (KBP) and automated planning. KBP

methods leverage historical treatment plans to help planners make informed decisions for new patients, and automated planning make the planning process efficient by minimizing human intervention.

1.2 Knowledge-Based Planning

Knowledge-based planning refers to a collection of methods that leverage a library of previous treatment plans, and predict features of an acceptable plan for a new patient. The most prominent KBP methods use patients' anatomical geometry to predict acceptable dose-volume histograms (DVHs).[29;30;31;36;41] DVHs are high-level snapshots of a treatment plan's dose distribution within each ROI; they are cumulative histograms that plot fractional volume of a structure that receives a minimum of the given dose. However, they are inadmissible as input to the IPPs that generate treatment plans, and are often positioned as tools to guide a treatment planner towards an acceptable plan.

Current KBP methods have several limitations. They are generally designed for disease sites with a single target, yet several complex sites have multiple targets. Additionally, they only predict DVHs for OARs (i.e., not targets), and those predictions may be physically impossible. Without a target DVH prediction, these KBP methods fail to provide a complete picture of a patient-specific treatment plan.

1.3 Automated Planning

Automated planning methods generate treatment plans for new patients without human intervention. Breast cancer treatment planning remains the only site where treatment planning is fully automated,[20] but other methods require some level of human intervention, particularly with contouring.[13;28] For complicated sites like head and neck automated methods use KBP to leverage information from previous plans. These methods are typically blackboxes that generate treatment plans from patient geometry, and they cannot be adjusted by tuning and IPP that treatment planners are already familiar with.[13;18;28;29]

Advocates for fully automated planning will argue human intervention defeats the purpose of the automation. Realistically, however, patients and physicians need the option to make more

personal plans. The ability for humans to intervene is especially important for complicated cases where there are several competing clinical priorities.[22] This inspired us to construct an automated method that initializes IPPs, which are familiar to treatment planners, with weights that produce high quality plans. To achieve this we turn to inverse optimization (IO).

Inverse optimization methods quantify the weights embedded in a treatment plan. An IO method can reverse engineer objective function weights from clinical dose objectives, which is effectively the inverse of inverse planning. Historically this technique has been used to map a population of clinical treatment plans to the same IPP (i.e., same set of objectives, constraints, and structures).[6] This provided the means to predict objective function weights from patient anatomy, and give planners a "warm-start".[4;16] The application of IO to IMRT has been limited to treatment plans for prostate cancer, which is a relatively simple site in comparison to head and neck.

1.4 Contributions

The goals of this thesis are to combine KBP with IO, and build a pipeline that generates treatment plans (i.e., optimized fluence maps) using patient anatomy with an IPP that is familiar to human treatment planners. We are the first to input KBP predictions into an IO model (Figure 1.1). Using the IO model eliminates the task of manually tuning objective weights to achieve the predicted DVH, which is a process that is susceptible to human error. The KBP methods that predict objective function weights output a similar product. However, our framework can be easily generalized to other disease sites and KBP methods. Each of our methods was validated with 217 oropharyngeal cancer treatment plans that were delivered at Princess Margaret Cancer Centre located in Toronto, Canada.

We made three contributions to the field:

1. We developed a novel method to account for SPG using inverse optimization. Our IO method infers the minimum SPG required to deliver a set of DVHs. In contrast, previous IO methods require a complete dose distribution as input, and need constraints to limit fluence map heterogeneity. The method was validated with a model that generalized to over 200 treatment plans for a complex disease site, and the DVHs created using IO and

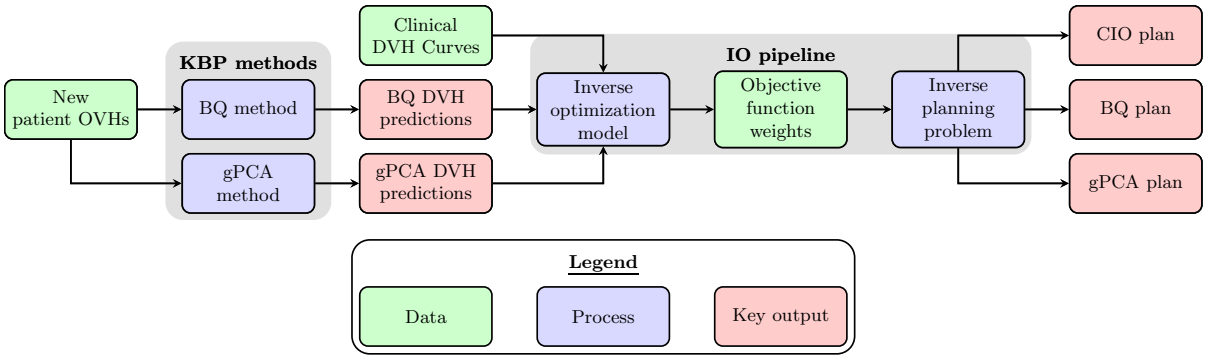

Figure 1.1: Overview of computational framework.

IPP closely reproduced the clinical DVHs

2. We extended two prominent KBP methods from the literature to output realistic DVHs for OARs and targets, and to generalize their application to patients with multiple targets. Previous KBP methods are limited to predicting the DVHs of OARs in disease sites with a single target, and are prone to predicting DVHs that are physically impossible. Both methods were tested on a large cohort of clinical treatment plans with leave-one-out cross validation. We compared the predicted DVHs to their clinical counterparts, and found that each extended KBP method predicted significantly different DVHs.

3. We combined each of the two KBP methods with IO to create a pipeline that automatically generates treatment plans. We compared both sets of these plans from to clinical plans using several metrics, and found that the plans generated using each KBP method made drastically different tradeoffs in terms of OAR sparing, target coverage, and delivery complexity. Nearly half of the plans generated by our pipeline from the most promising KBP method were non-inferior to clinical plans in terms of clinical planning criteria satisfactio.

We have also made the following contributions to the literature:

Publications

1. **A. Babier**, J.J. Boutilier, A. McNiven, M.B. Sharpe, T.C.Y. Chan. "Coupling inverse methods to produce IMRT treatment plans from DVH curves." *Under review at Medical Physics.*

2. **A. Babier**, J.J. Boutilier, A. McNiven, T.C.Y. Chan. "Knowledge-based automated planning for oropharyngeal cancer." *In preparation.*

Presentations

1. **A. Babier**, J.J. Boutilier, A.L. McNiven, T.C.Y. Chan. "Using inverse optimization to evaluate knowledge based planning methods for oropharyngeal cancer." *2017 INFORMS Annual Meeting*, Houston, TX, October 2017.

2. **A. Babier**, J.J. Boutilier, A.L. McNiven, T.C.Y. Chan."Knowledge-based automated planning for oropharyngeal cancer." *2017 AAPM Annual Meeting*, Denver, CO, August 2017.

3. **A. Babier**, J.J. Boutilier, A.L. McNiven, T.C.Y. Chan. "Reverse engineering fluence maps from dose volume histograms using inverse optimization and inverse planning." *2017 AAPM Annual Meeting*, Denver, CO, July 2017.

4. **A. Babier**, J.J. Boutilier, A.L. McNiven, T.C.Y. Chan. "Coupling knowledge-based planning and inverse optimization." *2017 IFORS Annual Meeting*, Quebec City, QC, July 2017.

5. **A. Babier**, J.J. Boutilier, A.L. McNiven, M.B. Sharpe, T.C.Y. Chan. "Automating cancer treatment planning with inverse optimization." *Centre in Computational Science and Engineering Seminar*, Toronto, ON, December 2016.

6. **A. Babier**, J.J. Boutilier, A.L. McNiven, M.B. Sharpe, T.C.Y. Chan. "Automating cancer treatment planning with inverse optimization." *UTORG Lunch Talk*, Toronto, ON, November 2016.

7. **A. Babier**, J.J. Boutilier, A.L. McNiven, M.B. Sharpe, T.C.Y. Chan. "Using inverse optimization to produce IMRT treatment plans from DVH curves." *2016 INFORMS Annual Meeting*, Nashville, TN, November 2016.

8. **A. Babier**, J.J. Boutilier, A.L. McNiven, M.B. Sharpe, T.C.Y. Chan. "Using inverse optimization to produce IMRT treatment plans from DVH curves." *MIE Symposium*, Toronto, ON, June 2016.

1.5 Organization

This thesis is organized as follows. In Chapter 2, we develop an inverse optimization model that uses DVHs as input (Contribution 1). In Chapter 3, we extend two existing KBP methods, and combine them with our IO method to build an automated planning pipeline (Contributions 2 and 3). In Chapter 4, we close with a summary and provide suggestions regarding directions for future work.

Chapter 2

Inverse Optimization

2.1 Introduction

At the heart of clinical treatment planning is an inverse planning problem (IPP), which is an optimization problem that generates fluence maps or deliverable segments given other parameters, such as weights on different objective functions. IPP-generated treatment plans are traditionally arrived at through manual, iterative parameter tuning of the underlying optimization model. Advances in the traditional approach include methods that automatically tune objective function weights to produce or improve a treatment plan.[8;15;17;32;33;35;37;39] Recently, inverse optimization (IO) methods[6] were developed to improve the treatment plan generation process by recovering objective function weights from previously delivered clinical plans – in effect, the inverse of inverse planning. These parameters can be predicted from an individual patient's geometry[4;16] and used to initialize the planning process in a more personalized manner, with the goal of having fewer replanning iterations. Previous IO methods assumed that an entire three-dimensional dose distribution was available as input.

The question of whether IO methods can generate acceptable objective weights from only a dose volume histogram (DVH), rather than a full dose distribution, has not been explored yet. This question is important to resolve given the rising interest in knowledge-based planning (KBP) methods, many of which focus on predicting achievable DVHs for de novo patients.[26;29;30;31;36] Such "target" DVHs can be used as a guide for treatment planners to know when to terminate the planning process. A related stream of research aims to automatically de-

rive fluence-based treatment plans given desired DVHs or DVH features, with motivation from adaptive radiation therapy (ART).[17;37;38] These methods are appealing because they eliminate manual weight tuning from the planning process, often relying on some norm minimization between desired and delivered dose instead. However, when the produced solution is unacceptable, it is unclear how a planner can improve the plan within this framework. In contrast, IO can fit into an existing clinical planning pipeline that planners are familiar with, so the process of refining a plan is intuitive.

Additionally, previous fluence-based optimization models that generate treatment plans from DVHs do not explicitly consider deliverability. A method to achieve a given DVH with minimal heterogeneity would be especially desirable because less heterogeneous fluence maps typically translate into plans with fewer monitor units.[27] The sum of positive gradients (SPG) method was developed as a linear approximation for the number of monitor units required to deliver a fluence map, and can be minimized tractably in fluence-based optimization models.[9] By integrating SPG into an IO model, we can infer the minimum delivery complexity – a critical factor in evaluating plan quality[11;19] – required to generate given DVHs.

In this paper, we propose a novel IO model to recover objective function weights and infer fluence map heterogeneity from a set of DVH curves. We then use these objective function weights in an IPP to generate a fluence map and corresponding DVHs that are similar to the input DVHs. A similar method applied to three-dimensional dose distributions was previously developed for prostate cancer,[6] which has relatively simple geometry. In this paper, we generalize that approach to use DVHs as input and accommodate more complex treatment sites, focusing on head and neck cancer. We formulate an IPP with linearizable objectives suitable for head and neck cancer treatment, including an SPG objective to control fluence map heterogeneity, and then we develop the corresponding IO model. This approach represents the first time SPG has been used in an IO model. We apply our IO framework to 217 sets of head and neck DVHs derived from clinical treatment plans delivered at the Princess Margaret Cancer Centre in Toronto, Canada. This dataset represents the largest set of patients used in an optimization-based study for head and neck cancer treatment planning in the literature.

2.2 Methods and Materials

2.2.1 Data

We obtained 217 oropharynx step and shoot intensity-modulated radiation therapy (IMRT) treatment plans from the Princess Margaret Cancer Centre in Toronto, Canada. Data access was granted under UHN REB 15-9076-CE. All plans were designed by the head and neck group consisting of multiple oncologists and treatment planners specializing in head and neck cancer. The plans were all generated using Pinnacle^3 (Philips, Madison, WI) to deliver 9 approximately equispaced coplanar fields with 6 MV IMRT, to satisfy a prescription of 70 Gy to the highest-risk target over 35 fractions. All treatment plans were extracted from the treatment planning system and imported into MATLAB via A Computational Environment for Radiotherapy Research (CERR).[10] For each treatment plan, we calculated DVHs using a bin size of 0.1 Gy. All optimization problems were solved with Gurobi 6.0 with default settings (Gurobi Optimization, Houston, TX).

2.2.2 Inverse Planning Problem (IPP)

In this section, we propose a series of linearizable objective functions that can be combined to construct a linear IPP that is suitable for head and neck treatment planning in general. The framework for an IO model is also outlined, and approaches to validate the resulting plans are presented.

Preliminaries and notation

We began by identifying a set of targets $\mathcal{T}$ and healthy structures $\mathcal{I}$ for each patient. Each target t was a planning target volume (PTV) with a prescribed dose θ^t, denoted as PTVθ^t. All treatment plans included in this study had both a PTV56 and PTV70; 130 plans also had a PTV63. The healthy structures contained in $\mathcal{I}$ were the brain stem, spinal cord, right parotid, left parotid, larynx, esophagus, and mandible. A dose-shaping structure called the limPostNeck was also included in $\mathcal{I}$, which was contoured around the posterior neck with the intent to limit fibrosis in all clinical plans. Each target structure t was divided into a set of voxels $\mathcal{O}^t$ and each healthy structure i was divided into a set of voxels $\mathcal{O}^i$.

Patient anatomy was discretized into voxels of dimension 1mm × 1mm × 2mm, which we call the *full resolution*. We also defined a *sampled resolution* by selecting voxels at the vertices of a 4mm × 4mm × 2mm grid overlaid on the original grid. To differentiate between the two resolutions, parameters from the sampled resolution are denoted with a bar (e.g., $\bar{\mathbf{A}}$ vs. $\mathbf{A}$).

Next, a beam was simulated from nine equidistant gantry angles 0°, 40°, ..., 320°, denoted by set $\mathcal{K}$. At each angle $k \in \mathcal{K}$, the beam was divided into a grid of beamlets, which were partitioned into a set of rows $\mathcal{R}^k$. Each row $r \in \mathcal{R}^k$ contained a set of beamlets $\mathcal{B}^r$. The relationship between the intensity w_b of beamlet b and dose d_v deposited to voxel v was determined using the influence matrix $D_{v,b}$ generated by the IMRTP library from CERR and given by equation (2.1):

$$d_v = \sum_{k \in \mathcal{K}} \sum_{r \in \mathcal{R}^k} \sum_{b \in \mathcal{B}^r} D_{v,b} w_b. \tag{2.1}$$

OAR objectives

We chose to limit the average dose and the maximum point dose to healthy tissues by including an objective for mean dose z^i and max dose y^i objectives for each OAR $i \in \mathcal{I}$ as

$$z^i = \underset{v \in \mathcal{O}^i}{\text{mean}} \{d_v\}, \quad \forall i \in \mathcal{I}, \tag{2.2}$$

$$y^i = \max_{v \in \mathcal{O}^i} \{d_v\}, \quad \forall i \in \mathcal{I}. \tag{2.3}$$

These two types of objective functions had corresponding weights γ^i and β^i, respectively.

Mean and max dose are the two extreme limits of equivalent uniform dose (EUD),[25] so they provide limited control over the full DVH curve. To control the DVH between the mean and max dose, we used a series of penalty-based objectives h^f to penalize dose exceeding a fixed threshold $f \in \mathcal{F}^i$ for each OAR $i \in \mathcal{I}$, with corresponding objective weights κ^f. The penalty objective h^f was defined as

$$h^f = \underset{v \in \mathcal{O}^i}{\text{mean}} \{\max\{0, d_v - f\}\}, \quad \forall f \in \mathcal{F}^i, \ \forall i \in \mathcal{I}. \tag{2.4}$$

Each OAR i was assigned six penalty thresholds with magnitudes of 0.25, 0.50, 0.75, 0.85, 0.95, and 0.975 times the maximum dose to that OAR from the clinical plan. Similar objectives have

been used in previous linear IPPs.[6;12;14;21]

Target objectives

For the targets, we minimized deviation from a prescribed dose θ^t. The average deviation below θ^t is the average *underdose* to target t, which was represented by l^t and assigned a corresponding weight ϕ^t. Similarly, the average deviation above θ^t is the average *overdose* to target t, which was represented by u^t and assigned a weight ψ^t. The average underdose and overdose objectives are defined below:

$$l^t = \underset{v\in\mathcal{O}^t}{\text{mean}} \left\{\max\{0, \theta^t - d_v\}\right\}, \quad \forall t \in \mathcal{T}, \tag{2.5}$$

$$u^t = \underset{v\in\mathcal{O}^t}{\text{mean}} \left\{\max\{0, d_v - \theta^t\}\right\}, \quad \forall t \in \mathcal{T}. \tag{2.6}$$

Similar to the OARs, we defined a maximum dose objective y^t to each target $t \in \mathcal{T}$, with corresponding weight β^t:

$$y^t = \max_{v\in\mathcal{O}^t}\{d_v\}, \quad \forall t \in \mathcal{T}. \tag{2.7}$$

Plan complexity

The overall SPG across all angles in $\mathcal{K}$ was defined as $\sum_{k\in\mathcal{K}} m^k$, where

$$m^k = \max_{r\in\mathcal{R}^k} \left\{\sum_{b\in\mathcal{B}^r} \max\{0, w_b - w_{b'}\}\right\}, \quad \forall k \in \mathcal{K}, \tag{2.8}$$

is the SPG for angle k, defined as the maximum *sum of positive intensity differences* over all rows $r \in \mathcal{R}^k$. The sum of positive intensity differences in row r was defined as the positive differences between the intensity w_b and the intensity $w_{b'}$, where b' is the beamlet to the immediate right of beamlet b, both in row r. When b' did not exist (i.e., at the end of the row), we set $w_{b'}$ to zero.

IPP model

To form the IPP, we linearized equations (2.2)-(2.8) by introducing appropriate auxiliary variables and constraints, and summed those terms in a weighted objective function. The conceptual

IPP model is given in (2.9):

$$\begin{aligned} \underset{z,y,x,l,u,c,g,w}{\text{minimize}} \quad & \sum_{i\in\mathcal{I}}(\gamma^i z^i + \beta^i y^i + \sum_{f\in\mathcal{F}^i}\kappa^f h^f) + \sum_{t\in\mathcal{T}}(\beta^t y^t + \phi^t l^t + \psi^t u^t) + \sum_{k\in\mathcal{K}} m^k \\ \text{subject to} \quad & (2.2) - (2.8). \end{aligned} \tag{2.9}$$

The complete model is given in the Appendix, including details of the linearizations of equations (2.2)–(2.8) required to transform (2.9) into a linear optimization problem.

2.2.3 Inverse optimization model

To streamline the presentation and development of the IO model, we first write the linearized IPP model (2.9) in matrix notation (see Appendix for complete expansion). The linearized versions of objective functions (2.2)-(2.7) are expressed as rows of a matrix $\mathbf{C}$, and their objective weights are captured in the vector $\boldsymbol{\alpha}$. The SPG objective from formulation (2.8) was represented by vector $\mathbf{g}$. The constraints that linearize the objectives in formulation (2.9) are represented by matrix $\mathbf{A}$ and vector $\mathbf{b}$. The vector $\mathbf{x}$ is the decision vector, including auxiliary variables and the vector of beamlet intensities $\mathbf{w}$. Given this notation, the equivalent linear programming formulation of (2.9) is given by (2.10):

$$\begin{aligned} \underset{\mathbf{x}}{\text{minimize}} \quad & \boldsymbol{\alpha}'\mathbf{C}\mathbf{x} + \mathbf{g}'\mathbf{x} \\ \text{subject to} \quad & \mathbf{A}\mathbf{x} \geq \mathbf{b}, \\ & \mathbf{x} \geq \mathbf{0}. \end{aligned} \tag{2.10}$$

The dual of formulation (2.10) is straightforward to derive[2] and shown in formulation (2.11), where $\mathbf{p}$ is the vector of dual variables corresponding to the constraint $\mathbf{A}\mathbf{x} \geq \mathbf{b}$:

$$\begin{aligned} \underset{\mathbf{p}}{\text{maximize}} \quad & \mathbf{b}'\mathbf{p} \\ \text{subject to} \quad & \mathbf{C}'\boldsymbol{\alpha} + \mathbf{g} \geq \mathbf{A}'\mathbf{p}, \\ & \mathbf{p} \geq \mathbf{0}. \end{aligned} \tag{2.11}$$

The IO model (2.12) takes the form of absolute duality gap minimization,[6] subject to the dual feasibility constraints:

$$\begin{aligned} \underset{\boldsymbol{\alpha},\mathbf{p}}{\text{minimize}} \quad & \boldsymbol{\alpha}'\mathbf{C}\hat{\mathbf{x}} + \mathbf{g}'\hat{\mathbf{x}} - \mathbf{b}'\mathbf{p} \\ \text{subject to} \quad & \mathbf{C}'\boldsymbol{\alpha} + \mathbf{g} \geq \mathbf{A}'\mathbf{p}, \\ & \boldsymbol{\alpha} \geq \mathbf{0}, \mathbf{p} \geq \mathbf{0}. \end{aligned} \tag{2.12}$$

In formulation (2.12), the vector $\hat{\mathbf{x}}$ represents the variables associated with a given treatment plan (i.e., the fluence map and auxiliary variable values). However, this vector is not directly observed since we assume we only have access to DVHs from the historical plans. Having the DVHs though, is sufficient to calculate the vector $\mathbf{C}\hat{\mathbf{x}}$, which represents dose objectives (e.g., mean dose to an organ). The scalar $\mathbf{g}'\hat{\mathbf{x}}$ is another unknown quantity since the SPG is not directly observable from DVH curves. Fortunately, this quantity is a constant in the objective function of (2.12), so we can omit it from the inverse optimization model. The final IO model is

$$\begin{aligned} \underset{\boldsymbol{\alpha},\mathbf{p}}{\text{minimize}} \quad & \boldsymbol{\alpha}'\mathbf{C}\hat{\mathbf{x}} - \mathbf{b}'\mathbf{p} \\ \text{subject to} \quad & \mathbf{C}'\boldsymbol{\alpha} + \mathbf{g} \geq \mathbf{A}'\mathbf{p}, \\ & \boldsymbol{\alpha} \geq \mathbf{0}, \mathbf{p} \geq \mathbf{0}. \end{aligned} \tag{2.13}$$

We used model (2.13) in the experiments to turn DVHs into objective function weights. In total, there were eight objectives for each OAR and three objectives for each target.

2.2.4 Analyses

Figure 3.1 outlines the computational pipeline that was applied to each patient's clinical treatment plan. We began by extracting a vector of dosimetric features (i.e., objective function values based on dose delivered) $\mathbf{C}\hat{\mathbf{x}}$ from the full resolution, clinical DVH curves. We constructed $\bar{\mathbf{A}}$ (constraint matrix) from the sampled influence matrix $\bar{\mathbf{D}}$. We constructed the sampled $\bar{\mathbf{C}}$, $\bar{\mathbf{g}}$, and $\bar{\mathbf{b}}$ for formulation (2.13) similarly. Given $\mathbf{C}\hat{\mathbf{x}}$, formulation (2.13) generated a weight vector $\boldsymbol{\alpha}$, which was input into the IPP (i.e., formulation (2.9)) to produce a set of optimized beamlets $\mathbf{w}$. Finally, we calculated the *inverse* dose distribution using equation (2.1), the beamlets $\mathbf{w}$, and the influence matrix $\mathbf{D}$. The corresponding plan and DVHs will be referred to as the *inverse plan* and *inverse DVHs*.

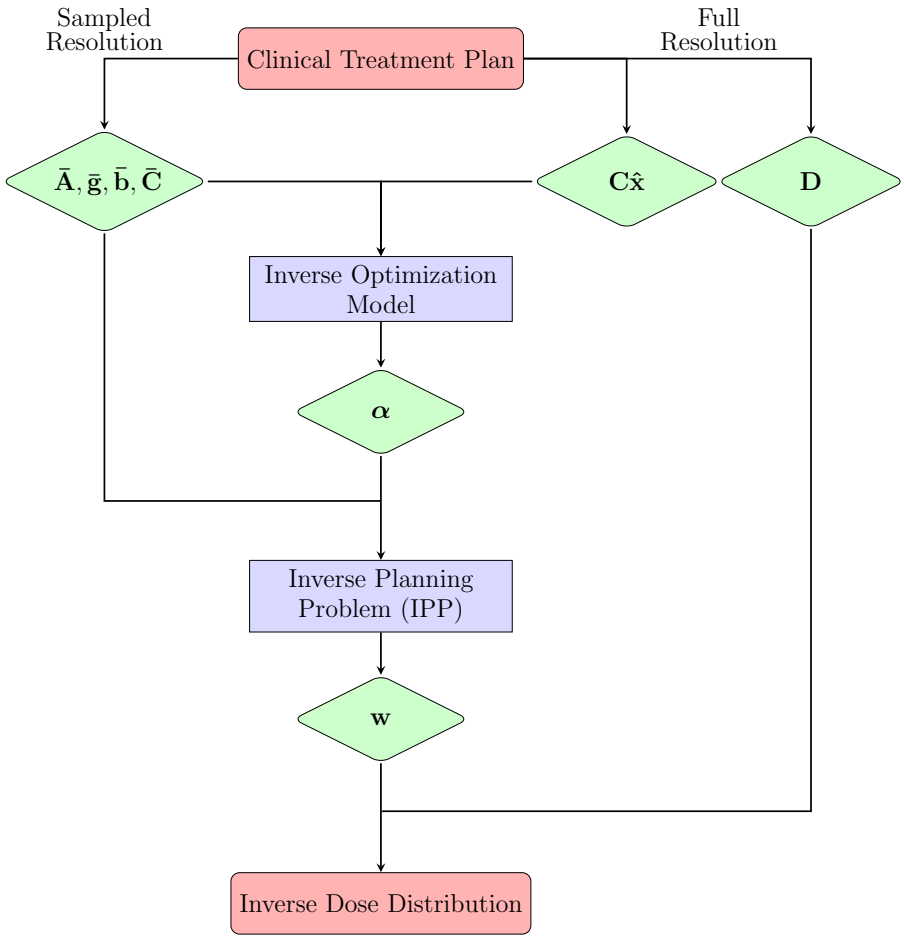

Figure 2.1: Overview of computational framework.

Analysis 1: Impact of voxel sampling

For each objective function, we determined its value from the DVH of both the sampled and full resolution clinical dose distributions. We calculated the correlation coefficient for this objective over all treatment plans in our dataset. For each objective, we also calculated the difference in values between the sampled and full resolutions. Then, we determined the median, 5^{th} percentile, and 95^{th} percentile of these differences over all treatment plans.

Analysis 2: Difference between clinical and inverse plan objective values

For each objective, we computed the difference in value between the clinical and inverse plan (i.e., clinical minus inverse). The resulting distribution over all treatment plans for each objective was shown in a box plot, such that points on the positive vertical axis represented cases

where the objective function value from the inverse plan was less than the clinical plan. For each objective function, we also calculated the frequency of nonzero objective weights generated by the IO method.

Analysis 3: Fractional volume difference between clinical and inverse DVHs

We computed both the positive and negative median fractional volume difference between each clinical DVH curve and the corresponding inverse DVH curve. Positive differences correspond to instances where the inverse plans delivered less dose to a fractional volume than the clinical plan, and negative differences correspond to the reverse case. This evaluation resulted in two median values for each OAR for each treatment plan, one value for the median of positive differences and another for the median of negative differences. We again used box plots to visualize the distribution over the population of treatment plans for each structure.

Analysis 4: Planning criteria satisfaction

We compared the clinical and inverse treatment plans in terms of satisfaction of the clinical planning criteria outlined in Table 3.1. We determined the frequency and degree to which both the clinical and inverse treatment plans violated the planning criteria.

Table 2.1: The primary planning criteria used for treating oropharynx cancer at Princess Margaret Cancer Centre with a prescription of 70 Gy to the high-risk targets over 35 fractions. Scalar $\mathcal{D}_v$ is the dose to a fractional volume v, $\mathcal{D}_{mean}$ is the mean dose to a structure, and $\mathcal{D}_{max}$ is the max dose to a structure.

Structure	Criterion
Brain Stem	$\mathcal{D}_{max} \leq 54$ Gy
Spinal Cord	$\mathcal{D}_{max} \leq 48$ Gy
Right Parotid	$\mathcal{D}_{mean} \leq 26$ Gy
Left Parotid	$\mathcal{D}_{mean} \leq 26$ Gy
Larynx	$\mathcal{D}_{mean} \leq 45$ Gy
Esophagus	$\mathcal{D}_{mean} \leq 45$ Gy
Mandible	$\mathcal{D}_{max} \leq 73.5$ Gy
PTV70	$\mathcal{D}_{99} \geq 65.1$ Gy
PTV63	$\mathcal{D}_{99} \geq 58.6$ Gy
PTV56	$\mathcal{D}_{99} \geq 53.2$ Gy

Analysis 5: Fluence map heterogeneity

For each DVH, we compared the fluence map heterogeneity inferred by IO model (13) and an IO model that did not incorporate SPG. Instead, this other IO model used a hard constraint to keep all beamlets within 50% of the mean beamlet intensity, which is taken from a prostate IO model from the literature.[6] To differentiate between the two inverse methods, any feature from the non-SPG IO model is denoted with an asterisk (e.g., IO* vs. IO, IPP* vs. IPP, etc.). We calculated the mean and standard deviation of SPG* over all inverse* plans, and compared that to the SPG from all inverse plans derived from (2.13). We also evaluated the frequency that inverse* plans satisfied the clinical criteria from Table 3.1.

2.3 Results

2.3.1 Impact of voxel sampling

Table 2.2 shows the correlation between the objective function values constructed from the full and sampled resolution DVH curves. The values of each objective function were highly correlated, indicating that sampling did not degrade the dose calculation. For the mean and threshold dose objectives, the variability was very small, at most 0.50 Gy difference between the 5^{th} and 95^{th} percentiles of the distribution. The two max objectives had slightly more variability, up to a 2.86 Gy difference at the 95^{th} percentile.

Table 2.2: Correlation coefficients for the objective function values extracted from the full and sampled resolution DVH curves, and percentiles for their differences.

Objective	Correlation	Objective Value Difference (Gy)		
		5^{th} Percentile	Median	95^{th} Percentile
Mean OAR Dose (z^i)	0.999	-0.27	-0.01	0.22
Max OAR Dose (y^i)	0.998	0.00	0.51	2.86
Threshold Dose Penalties (x^f)	1.000	-0.50	-0.03	0.01
Mean Target Underdose (l^t)	0.999	-0.01	0.00	0.01
Mean Target Overdose (u^t)	0.999	-0.02	0.00	0.02
Max Target Dose (y^t)	0.992	0.00	0.28	1.04

2.3.2 Difference between clinical and inverse plan objective values

Figure 2.2 presents the difference between the inverse and clinical plan objective function values on the horizontal-axis for the structures on the vertical axis. In cases where the inverse objectives were better than clinical objectives (i.e., positive x-axis values), the median relative improvement was 1.6%. The median relative difference was similarly small, 0.8%, when the clinical plan outperformed the inverse plan. Over all objective function value differences, only 6.4% of them were outside 1.5 times the interquartile range. Generally the inverse plans performed slightly better on the OAR objectives while the clinical plans performed slightly better on the target objectives. Overall though, the objective values between clinical and inverse plans

were very similar across the board, indicating that the IO method was able to closely reproduce the clinical doses.

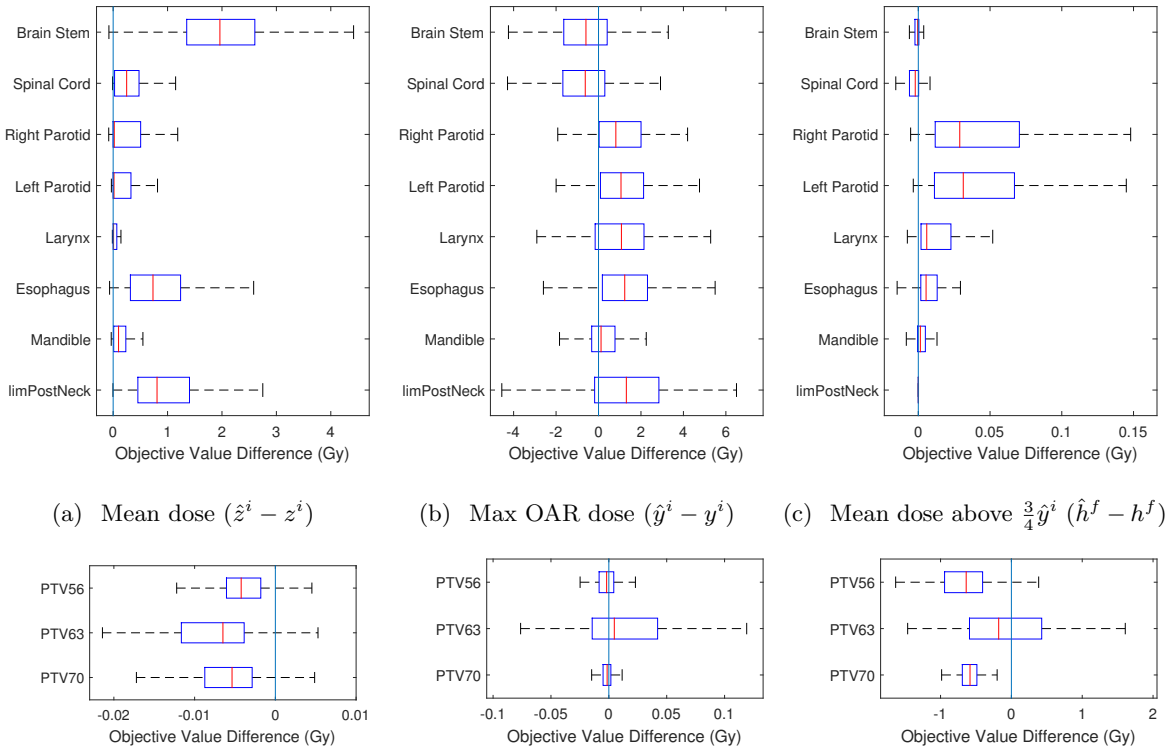

(a) Mean dose $(\hat{z}^i - z^i)$ (b) Max OAR dose $(\hat{y}^i - y^i)$ (c) Mean dose above $\frac{3}{4}\hat{y}^i$ $(\hat{h}^f - h^f)$

(d) Mean target underdose $(\hat{l}^t - l^t)$ (e) Mean target overdose $(\hat{u}^t - u^t)$ (f) Max target dose $(\hat{y}^t - y^t)$

Figure 2.2: The distribution of objective value differences between inverse and clinical plans. A positive difference implies the inverse plan achieved a lower (i.e., better) objective function value than the clinical plan. The upper and lower boundaries of each box represent the 75^{th} and 25^{th} percentiles respectively, and the horizontal line in the box depicts the median. Whiskers extend to 1.5 times the interquartile range. A line across each plot provides a reference point for zero difference.

The IPPs had between 51-73 objective functions, depending on the whether certain OARs or the PTV63 was contoured for a particular patient. The vector of objective function weights produced by the IO model was sparse in general (Table 2.3). Across all patients and objective functions, 38.5% of the objective weights were nonzero. The frequency of nonzero weights depended on the type of objective function. Imputed weights for the target underdose objective were always positive. Similarly, weights for the target overdose objective were positive in over

90% of patients. On the other hand, weights for the max OAR and max target dose objectives were zero over 90% of the time.

Table 2.3: Frequency of nonzero objective function weights produced by the IO model.

Objective Weight	Nonzero Frequency (%)
Mean Dose (γ^i)	28.5
Max Dose (β^i)	8.6
Threshold Dose Penalties (κ^f)	36.4
Target Underdose (ψ^t)	100.0
Target Overdose (ϕ^t)	91.4
Max Target Dose (β^t)	0.8

2.3.3 Fractional volume difference between clinical and inverse DVHs

Figure 2.3 shows the median positive and negative differences between the inverse and clinical DVH curves on the vertical axis against all structures of interest on the horizontal axis. Only 5% of all DVH differences were outside 1.5 time the interquartile range for each structure, but even those fractional volume differences were small (Figure 2.4). The largest median positive difference (i.e., inverse had better DVH) was 4% fractional volume, which occurred for the brain stem objective. The largest median negative difference was -3% fractional volume, which occurred for the larynx objective. Overall, the DVH differences were small, with the median across all structures being $\pm 1.4\%$ fractional volume.

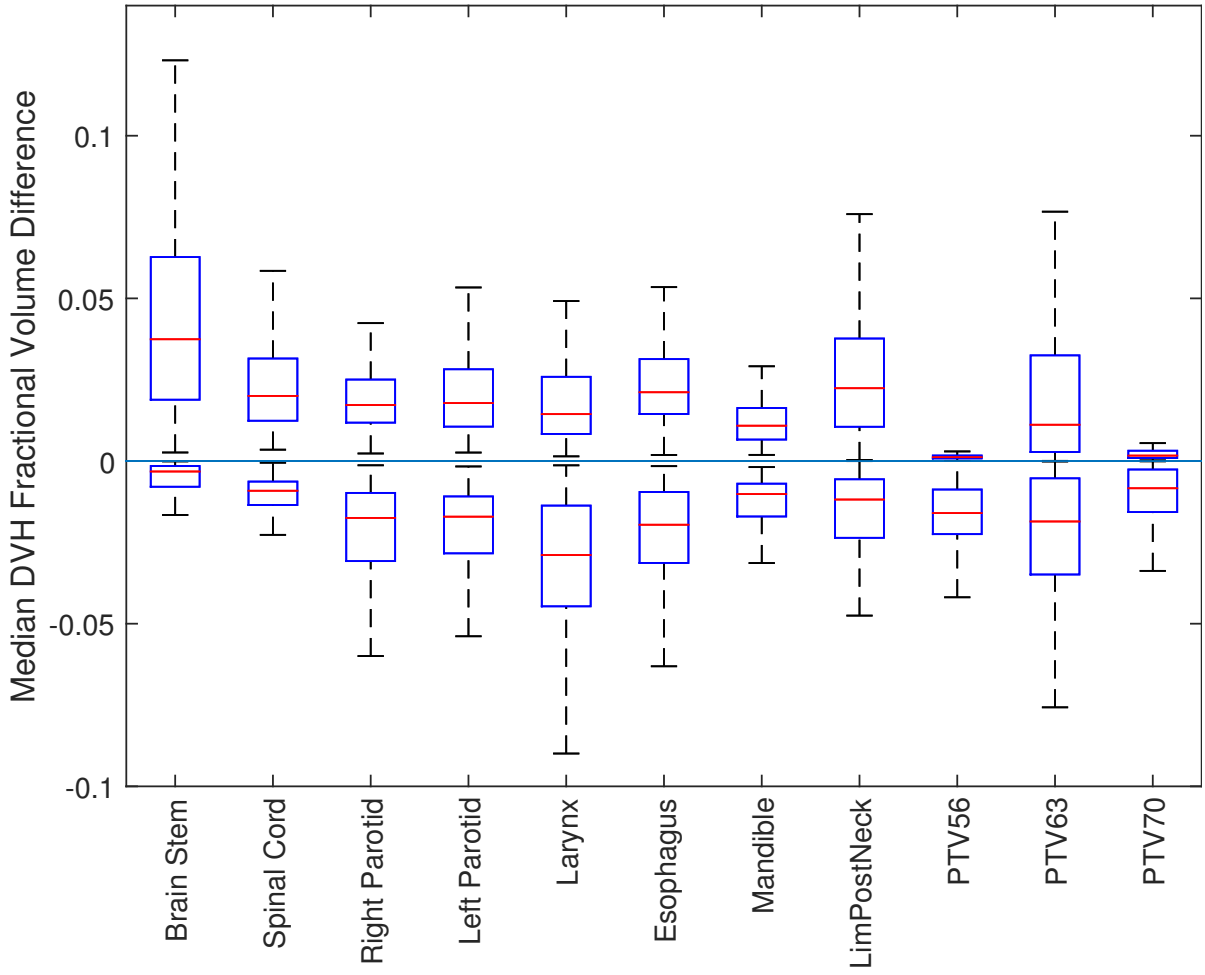

Figure 2.3: The distribution of positive and negative median DVH differences between inverse and clinical plans. Positive differences correspond the inverse plan achieving lower dose to a fractional volume compared to the clinical plan. The upper and lower boundaries of each box represent the 75^{th} and 25^{th} percentiles respectively, and the horizontal line in the box depicts the median. Whiskers extend to 1.5 times the interquartile range. A line across each plot provides a reference point for zero difference.

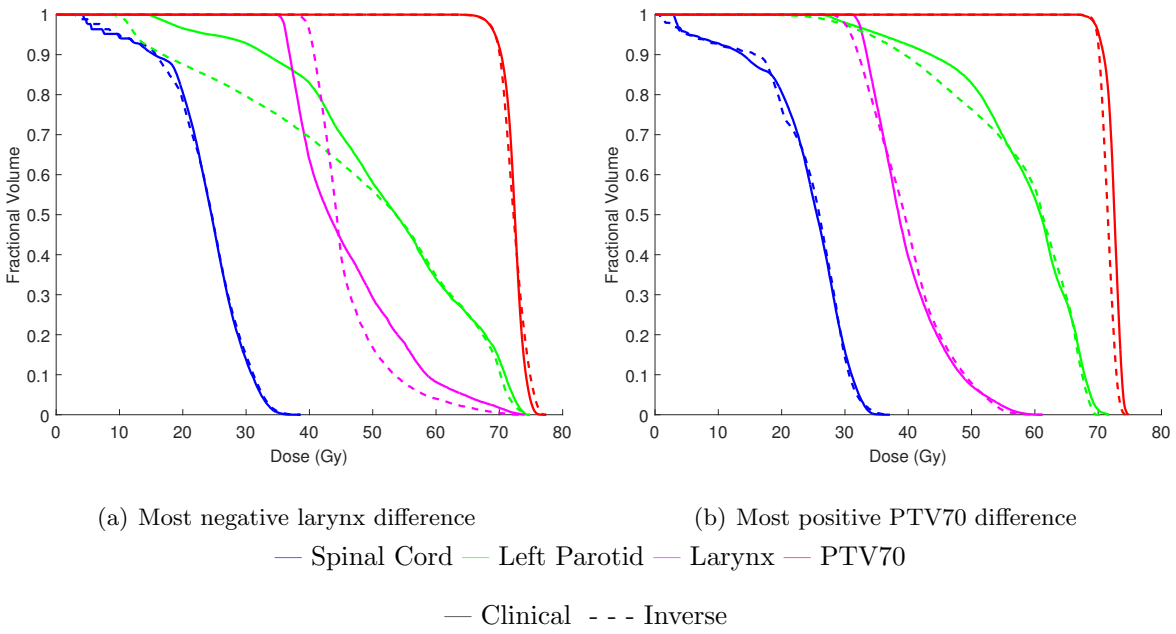

(a) Most negative larynx difference (b) Most positive PTV70 difference

Figure 2.4: Sample DVHs from: (a) The plan with the most negative larynx DVH difference (-0.182 fractional volume). (b) The plan with the most positive PTV70 DVH differences (0.113 fractional volume).

2.3.4 Planning criteria satisfaction

Table 2.4 summarizes how often the inverse and clinical plans achieved the planning criteria from Table 3.1. Across all structures and patients, the inverse plans achieved 68.3% of the planning criteria compared to 68.9% for the clinical plans. Planning criteria satisfaction varied across structures for both the inverse and clinical plans. For example, both clinical and inverse plans were able to satisfy the brain stem, spinal cord, and PTV63 criteria in at least 99.5% of the patients. On the other hand, both inverse and clinical plans satisfied at most 20% of the criteria for the left and right parotid.

Across all structures, the performance of the inverse and clinical plans with respect to the stated planning criteria was very similar. For instance, the inverse and clinical plans met the planning criteria for the larynx and esophagus at exactly the same frequency. For most structures, the difference in planning criteria satisfaction between the two plans was at most one percentage point. The exceptions were the mandible, where the inverse plan satisfied the criteria more frequently (by 7.9 percentage points) and the PTV56 and PTV70, where the clinical plan

satisfied the criteria more frequently (by 9.3 and 5.5 percentage points, respectively). Note that in these three structures, the amount of planning criteria violation was similar (difference of at most 0.5 Gy on average) for both the inverse and clinical plans.

Table 2.4: The frequency that clinical and inverse plans satisfied the planning criteria. Criteria violation is the mean ± standard deviation taken only over the cases that violated the criteria.

Structure	Frequency			Violation (Gy)	
	Clinical	Inverse	Abs. diff.	Clinical	Inverse
Brain Stem	100%	99.5%	−0.5%	0.0 ± 0.0	0.1 ± 0.0
Spinal Cord	100%	99.5%	−0.5%	0.0 ± 0.0	0.6 ± 0.0
Right Parotid	19.1%	20.0%	0.9%	13.4 ± 9.9	13.2 ± 9.7
Left Parotid	17.5%	18.5%	1.0%	13.8 ± 10.2	13.6 ± 9.9
Larynx	61.2%	61.2%	0.0%	7.3 ± 6.7	7.3 ± 6.7
Esophagus	96.2%	96.2%	0.0%	5.1 ± 5.0	4.4 ± 4.3
Mandible	64.5%	72.4%	7.9%	0.7 ± 0.6	0.6 ± 0.5
PTV56	51.2%	41.9%	−9.3%	1.1 ± 0.9	1.5 ± 1.2
PTV63	100%	100%	0.0%	0.0 ± 0.0	0.0 ± 0.0
PTV70	89.4%	83.9%	−5.5%	1.2 ± 1.7	1.5 ± 1.2

2.3.5 Fluence map heterogeneity

For every patient in our cohort the inverse plan satisfied as many or strictly more clinical criteria than the inverse* plans, and always with less heterogeneous fluence maps. The average SPG (53 ± 3) was 40.4% lower than the average SPG* (89 ± 15). Over all inverse plans, many more clinical criteria were satisfied (68.3%) compared to inverse* plans (39.9%). These results demonstrate that including SPG in the objective function produced superior plans. Slices from the inverse dose distribution also appeared more similar to clinical distributions (Figure 2.5).

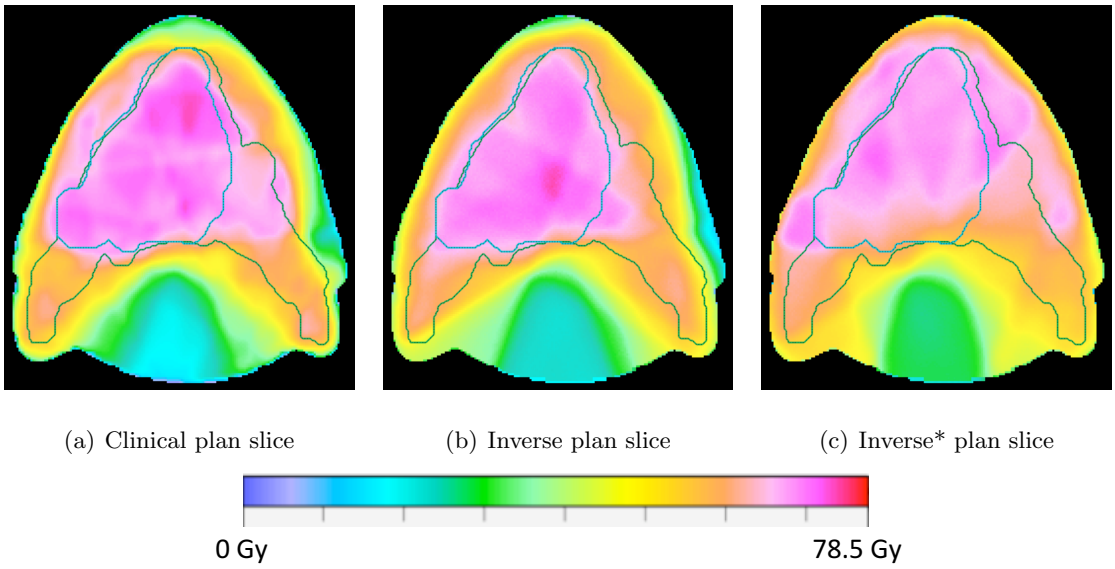

(a) Clinical plan slice (b) Inverse plan slice (c) Inverse* plan slice

Figure 2.5: The dose distribution over the same CT slice from the (a) clinical, (b) inverse, and (c) inverse* plans for patient 25. The following PTV56 and PTV70 are contoured in green and blue, respectively.

2.4 Discussion

In this paper, we developed a method that can reverse engineer – using inverse optimization – objective function weights given DVHs corresponding to a clinical treatment plan. We then used these weights to generate a fluence map via the corresponding inverse planning problem. By computing objective function values, DVH curve differences, and planning criteria satisfaction, we found that the inverse and clinical plans were very similar. Moreover, we demonstrated that our method produces plans that satisfy more clinical criteria than the plans from previous IO methods, and with consistently less heterogeneous fluence maps.

Learning optimization model parameters for the purpose of automatically generating a treatment plan may be useful in a clinical pipeline that requires frequent replanning (e.g., adaptive radiation therapy) or atlas-based approaches (e.g., knowledge-based planning). The process of solving the IO model to get weights and then using those weights in an IPP is easily linkable so that the process of generating a plan from an input DVH is automated. Previous studies have shown that the output of IO can be used to train machine learning models to predict

objective function weights for de novo patients in a hypothetical KBP pipeline.[4;16] At the very least, such an IO approach may be an efficient way to initialize the treatment planning (or replanning) process. Because the IO approach builds on the existing clinical pipeline of using an IPP, a treatment planner will be able to easily replan if additional iterations are needed after the initial plan is generated.

Consistent with other IO models for IMRT, we observed sparsity in the resulting inversely optimized weight vector.[6;7] In particular, across all patients and objective functions, over 60% of the objective weights were zero. This result suggests that many fewer objectives than the ones we included are needed to re-create a given DVH. Target objectives almost always had positive weights, while max dose objectives almost always did not. Some objectives were correlated (e.g., penalizing dose above 95% of the maximum OAR dose and penalizing dose above 97.5% of the maximum OAR dose) and they all worked together to shape the overall dose distribution. These correlations likely led to some objectives having zero weight, because giving weight to one objective effectively accounted for other correlated objectives. Mathematically, sparsity in the weight vector may be due to the implicit regularization of the weight vector that is performed by duality-gap-based IO models.[7] Other studies have aimed to produce fluence maps from DVH curve features.[37;38;39;40] However, they limited their analysis to a small number of treatment plans (no more than eight). In contrast, we applied our approach to a cohort of over 200 clinical cases, the largest in the literature to date on this topic, and demonstrated consistency of results across several key metrics over this large cohort. Another key difference in our approach is the use of IO, which is an intermediate step in the process of generating a treatment plan from DVHs. Our method also differs from previous studies in the explicit inclusion of a secondary objective to minimize fluence map heterogeneity. This feature also distinguishes our model from previous IO papers in radiation therapy,[6] where a heuristic approach to control beamlet heterogeneity was used. By using the sum of positive gradients,[9] our approach constructs a fluence map that matches input DVHs closely with minimal beamlet heterogeneity. Also note that IO models in general require a normalization constraint on the objective weights to prevent a trivial solution from being realized (i.e., $\boldsymbol{\alpha} = \mathbf{0}$). Such a normalization is unnecessary in our model because of the presence of the SPG objective, which provides an "anchor" to the weight vector and prevents $\boldsymbol{\alpha} = \mathbf{0}$ from being a feasible solution.

A limitation of the method presented in this paper is that it relies on the IPP to be a linear optimization problem. An extension of this approach to convex optimization would be straightforward, but of course, not all clinical objectives of interest are linear or even convex. Another limitation is the use of a fluence-based model; the development of a similar IO approach for direct aperture optimization may be a fruitful direction for future research. Finally, our approach is not as straightforward to implement as an approach that generates fluence maps from DVHs by taking, for example, a normed difference between a desired and delivered dose distribution. Development of general libraries to implement IO, without knowledge of duality theory, may help lower the barrier to using our approach.

2.5 Conclusion

In this paper, we developed a general inverse optimization method that, when coupled with an inverse planning problem, can automatically generate fluence maps from input DVHs. We tested our approach on over 200 head and neck cancer cases, the largest dataset and most complex site for this task in the literature to date. We demonstrated that our approach can generate DVHs that closely match the input DVHs. Our framework may be useful in adaptive radiation therapy or knowledge-based planning pipelines, and has the potential to provide a "warm-start" for the inverse planning process when some attributes of a desirable treatment plan are known.

Chapter 3

Automated Planning

3.1 Introduction

Automated planning promises to make treatment planning more efficient. Aside from breast cancer treatment planning,[20] commercialized methods that automate treatment planning planning are in their infancy and must be initialized with human input. One move toward automation is knowledge-based planning (KBP), which comprises a group of methods that learn from historical treatment plans and predict attributes of desirable plans for new patients.[28;29;30;31;36;41] Unfortunately, it is often difficult to incorporate those predictions into the inverse planning problem (IPP) that generates a treatment plan. As a consequence, these KBP methods are typically positioned as tools that guide treatment planners toward an acceptable plan. Alternatively, KBP predictions can be input to an automated planning engine, but those methods introduce a new planning paradigm; where planners are unable to adjust the final plan.[13;18;23;24;28;34]

Inverse optimization (IO) is a method that can quantify the relative importance of objective functions with objective function weights, which can be input directly to an IPP. Thus far, the input to IO methods has been limited to objective functions from clinical plans,[1;6;16] yet many KBP predictions also take a form that is nearly admissible to IO. In this paper, we are the first to unite KBP predictions directly with an IO model, and eliminate the burden of manually tuning weights to achieve a predicted dose volume histogram (DVH); a process that is time-consuming process and susceptible to human error. A similar problem has also been examined by a branch of KBP methods that generate objective function weights directly from patient

anatomy, but they are focused exclusively on a prostate cancer treatment planning planning with invariable objective functions.[4;16] In contrast, our method leverages IO to generate the weights for an IPP with objective functions that are tailored specifically to achieve the predicted DVH.

There are several candidate KBP methods that could provide input for an IO model. Query methods are one branch of KBP that predict single DVH points.[28;29;30;31] A second branch combines principal component analysis (PCA) and linear regression to predict full DVHs.[36;41] These methods leverage geometric information from overlap volume histograms (OVHs),[30] and most also consider structure volumes. Unfortunately, these methods are specialized for sites with a single target[28] and can predict the DVHs for OARs but not targets. There is also evidence that those KBP predictions are prone to over exaggerated OAR sparing.[3] Additionally, very little is known about the explicit tradeoffs required to achieve predicted DVHs. Fortunately, IO is an efficient tool that can reverse engineer objective function weights from DVH curves, and applying it to KBP predictions is less labour intensive than the iterative treatment planning approach.

The purpose of this paper is to develop a pipeline that generates treatment plans from patient anatomy. We expand two KBP methods from the literature, a query-based and a PCA-based method, to make them output realistic DVHs for both OARs and targets for patients with multiple targets. Next, the predictions are input to an IO pipeline[1] that generates the objective function weights for an IPP, which we solve to produce a plan that is similar to the input DVHs. We validated our approach with over two hundred oropharyngeal cancer treatment plans by predicting achievable DVHs, and engineering parameters to make plans that actually achieve the predictions.

3.2 Methods and Materials

In this section, we outline our automated planning pipeline (Figure 3.1). We began with one of our two KBP methods, which predict achievable DVHs of targets and OARs. Then those DVHs were input to an IO pipeline to produce a fluence-based treatment plan. The clinical DVHs were input to the same IO pipeline to generate benchmark plans. All optimization problems

were solved with Gurobi 6.0 (Gurobi Optimization, Houston, TX) with default settings.

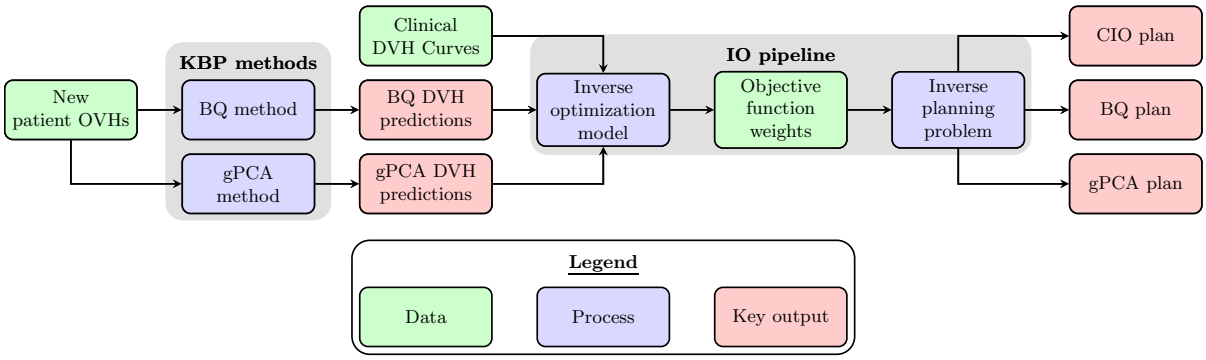

Figure 3.1: Overview of computational framework.

3.2.1 Data

We obtained 217 oropharynx step and shoot intensity-modulated radiation therapy (IMRT) treatment plans from the Princess Margaret Cancer Centre in Toronto, Canada. Data access was granted under UHN REB 15-9076-CE. We imported the treatment plans to MATLAB via A Computational Environment for Radiotherapy Research (CERR).[10] For the regions of interest (ROIs) in each treatment plan, we calculated DVHs and OVHs with 1001 equally sized bins from 0.0% to 100.0% fractional volume.

3.2.2 KBP Methods

For each patient we identified the targets $\mathcal{T}$ and OARs $\mathcal{I}$. All targets $t \in \mathcal{T}$ were planning target volumes (PTVs) with a dose prescribed by an oncologist. All 217 plans in this study had a PTV56 and PTV70; 130 of the 217 plans also had a PTV63. The OARs included the brain stem, spinal cord, right parotid, left parotid, larynx, esophagus, and mandible. The *limPostNeck*, a dose-shaping structure, was also included in $\mathcal{I}$ to limit fibrosis around the posterior neck. The anatomy of all patients was discretized into voxels of dimension 3mm $\times$ 3mm $\times$ 2mm.

We present two prominent KBP methods from the literature and extend them to predict DVHs for a set of OARs $i \in \mathcal{I}$ and targets $t \in \mathcal{T}$ for an out-of-sample patient j. The methods require a training set of previous treatment plans $p \in \mathcal{P}$. Both methods consider discrete fractional volumes $f \in \mathcal{F}$ and predict the dose $D_f^{j,i}$ or $D_f^{j,t}$ to be delivered to an OAR i or a

target t, respectively. They rely on a relationship, specifically OVHs, between the fraction of one structure contained by the isotropic expansion of another; to differentiate them, we denote the isotropically expanded structure with an asterisk (e.g., t^*). All OVHs were constructed from the minimum isotropic expansion $r^{p,i}_{f,t^*}$ of a target $t^* \in \mathcal{T}$ required to contain a fractional volume f of OAR i in a patient p. The PCA based method also used the OVHs constructed from the minimum isotropic expansion $r^{p,i}_{f,t^*}$ of each target $t^* \in \mathcal{T}$ required to contain a fractional volume f of all other targets $t \in \mathcal{T}$ in a patient p. Similarly, we also calculated the OVHs for the out-of-sample patient j.

Bagging query method

The traditional query method predicts a series of DVH points for an OAR i based on its proximity to a single target t^*.[30] It can be written as

$$D^{j,i}_f = \min_{p \in \mathcal{P}} \left\{ D^{p,i}_f \mid r^{j,i}_{f,t^*} \geq r^{p,i}_{f,t^*} \right\}, \quad \forall f \in \mathcal{F},\ \forall i \in \mathcal{I}, \tag{3.1}$$

which means that for an OAR i the dose delivered to the fractional volume f in a new patient j is equal to the lowest dose delivered to that same f over a subset of plans from $\mathcal{P}$; where the plans in that subset are those with a fraction volume f that is closer to the target t^* than that same f in the new patient j. This query method assumes that that the dose delivered to OARs is inversely proportional to the distance from a target, but ignores other factors. There are also variations of the this method that limit the plans contained in $\mathcal{P}$.[28;29;31] Unfortunately, there is evidence that query methods are susceptible to overfitting,[3] and the absence of target DVH predictions makes them poor candidates for inverse optimization. To address this issue, we constructed the *bagging query* (BQ) method, which limits overfitting, and accounts for sites with multiple targets. Our BQ method predicts the DVHs for an out-of-sample patient j as a weighted average that is proportional to how often the original query method selects a plan over a series of fractional volumes. To formulate this method we began by constructing a set $\mathcal{H}^{j,i}_{f,t^*}$, which indexed any patient h that would be selected from $\mathcal{P}$ by the original query method

for a given fractional volume f, OAR i, and target t^*:

$$\mathcal{H}^{j,i}_{f,t^*} = \underset{p\in\mathcal{P}}{\operatorname{argmin}} \left\{ D^{p,i}_f \mid r^{j,i}_{f,t^*} \geq r^{p,i}_{f,t^*} \right\}, \quad \forall f \in \mathcal{F},\ \forall t^* \in \mathcal{T},\ \forall i \in \mathcal{I}. \tag{3.2}$$

We refer to the plans $h \in \mathcal{H}^{j,i}_{f,t^*}$ as the *indexed plans*. In order to evaluate the weighted average, we counted the number of indexed plans for each OAR i and target t^* as $n^{j,i}$ and n^{j,t^*}, respectively:

$$n^{j,i} = \sum_{t^*\in\mathcal{T}} \sum_{f\in\mathcal{F}} |\mathcal{H}^{j,i}_{f,t^*}|, \quad \forall i \in \mathcal{I}, \tag{3.3}$$

$$n^{j,t^*} = \sum_{i\in\mathcal{I}} \sum_{f\in\mathcal{F}} |\mathcal{H}^{j,i}_{f,t^*}|, \quad \forall t^* \in \mathcal{T}. \tag{3.4}$$

Next, we predicted DVHs as the weighted average of indexed plans. The dose to all fractional volumes $f^* \in \mathcal{F}$ were predicted for each OAR i and target t^* as $\tilde{D}^{j,i}_{f^*}$ and $\tilde{D}^{j,t^*}_{f^*}$, respectively:

$$\tilde{D}^{j,i}_{f^*} = \sum_{t^*\in\mathcal{T}} \sum_{f\in\mathcal{F}} \sum_{h\in\mathcal{H}^{j,i}_{f,t^*}} \frac{D^{h,i}_{f^*}}{n^{j,i}}, \quad \forall i \in \mathcal{I},\ \forall f^* \in \mathcal{F}, \tag{3.5}$$

$$\tilde{D}^{j,t^*}_{f^*} = \sum_{i\in\mathcal{I}} \sum_{f\in\mathcal{F}} \sum_{h\in\mathcal{H}^{j,i}_{f,t^*}} \frac{D^{h,i}_{f^*}}{n^{j,t^*}}, \quad \forall t^* \in \mathcal{T},\ \forall f^* \in \mathcal{F}. \tag{3.6}$$

Generalized PCA method

Previous PCA methods are limited to predicting dose for OARs in patients with one target.[36;41] Thus, we formulated the *generalized PCA* (gPCA) method to predict DVHs for both OARs and targets in plans with multiple targets. In this section, we decompose the gPCA method into three steps outlined in Figure 3.2. Similar to previous PCA methods, we sampled 50 points in increments of 2% between 98% and 0% fractional volume across each DVH; the dose delivered to 100% fractional volume was excluded because it will receive at least 0 Gy. Similarly, 50 points were sampled from the OVH between 100% and 2% fractional volume. To differentiate between the two resolutions all sets with elements from the sampled curves are denoted with a hat (e.g., $\hat{\mathcal{F}}$ vs. $\mathcal{F}$). We applied PCA to the sampled DVHs from the training set $\mathcal{P}$, and generated the DVH principal components (PCs) $d^{p,i}_k$ and $d^{p,t}_k$ for each OAR i and target t, respectively. Each k

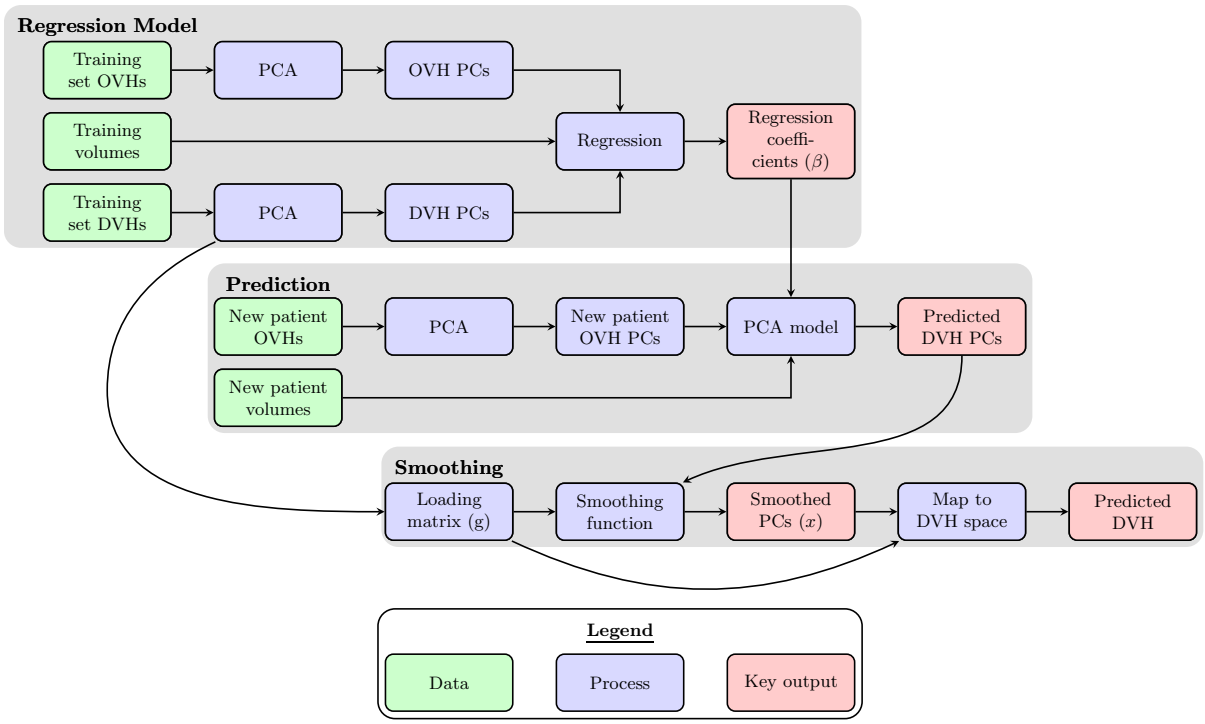

Figure 3.2: An overview of the gPCA method with the three main steps emphasized in grey boxes.

comes from an index set $\mathcal{K}$ that corresponds to the PCs, and similar to previous PCA methods we only considered the first three PCs of the sampled DVHs (i.e., $|\hat{\mathcal{K}}| = 3$). We also generated the elements $g^s_{f,k}$ for the loading matrix to project between PC space and DVH space, so that

$$D^{p,s}_f = \sum_{k \in \mathcal{K}} g^s_{f,k} d^{p,s}_k, \quad \forall f \in \hat{\mathcal{F}},\ \forall s \in \{\mathcal{I}, \mathcal{T}\},\ \forall p \in \mathcal{P}. \tag{3.7}$$

Similarly, we applied PCA to the sampled OVHs from the training set $\mathcal{P}$, but also for the out-of-sample patient j. Then the OVH PCs $o^{p,i}_{c,t^*}$ and $o^{p,t}_{c,t^*}$ were generated for the OVH t^* shares with each OAR i and target t, respectively. Each c comes from an index set $\mathcal{C}$ of the PCs, which we again truncate to the first three PCs of the sampled OVH (i.e., $|\hat{\mathcal{C}}| = 3$). The volumes of each OAR v^p_i and target v^p_t were also extracted from each patient $p \in \mathcal{P}$, and a series of regression coefficients β were introduced to reproduce the original model[36;41] with a single target t^* as

$$d^{j,i}_k = \beta^k_0 + \sum_{c \in \hat{\mathcal{C}}} \beta^k_c o^{p,i}_{c,t^*} + \sum_{i \in \mathcal{I}} \beta^k_i v^p_i + \beta^k_{t^*} v^p_{t^*}, \quad \forall k \in \hat{\mathcal{K}},\ \forall i \in \mathcal{I}. \tag{3.8}$$

We modified that model to predict the DVH PCs $\bar{d}_k^{j,s}$ for all ROIs $s \in \{\mathcal{I}, \mathcal{T}\}$ in plans with multiple targets as

$$\bar{d}_k^{j,s} = \beta_0^k + \sum_{t^* \in \mathcal{T}} \sum_{c \in \hat{\mathcal{C}}} \beta_c^k o_{c,t^*}^{j,s} + \sum_{i \in \mathcal{I}} \beta_i^k v_i^p + \sum_{t \in \mathcal{T}} \beta_t^k v_t^p, \quad \forall k \in \hat{\mathcal{K}},\ \forall s \in \{\mathcal{I}, \mathcal{T}\}. \tag{3.9}$$

Like all PCA methods, equation (3.9) is prone to predicting PCs that map to insensible DVHs; meaning a fractional volume f', which is immediately smaller than f, is assigned a lower dose than f even though that is physically impossible. To address this issue, we developed a *smoothing procedure* to produce PCs that map to a non-increasing DVH:

$$\begin{aligned} \underset{x}{\text{minimize}} \quad & \sum_{k \in \hat{\mathcal{K}}} \left(x_k^{j,s} - \bar{d}_k^{j,s} \right)^2 \\ \text{subject to} \quad & \sum_{k \in \hat{\mathcal{K}}} g_{f,k}^s x_k^{j,s} \leq \sum_{k \in \hat{\mathcal{K}}} g_{f',k}^s x_k^{j,s}, \quad \forall f \in \hat{\mathcal{F}},\ s \in \{\mathcal{I}, \mathcal{T}\}. \end{aligned} \tag{3.10}$$

We solved equation (3.10) to generate *smooth* PCs $x_k^{j,s}$, which were converted into DVH predictions via

$$\bar{D}_f^{j,s} = \sum_{k \in \hat{\mathcal{K}}} g_{f,k}^s x_k^{j,s}, \quad \forall f \in \hat{\mathcal{F}},\ \forall s \in \{\mathcal{I}, \mathcal{T}\}. \tag{3.11}$$

3.2.3 Inverse optimization

A previously developed IO model[1] was used to reverse engineer objective function weights from the DVH curves. This IO method also inferred the minimum sum of positive gradients (SPG), a measure of plan complexity,[9] necessary to deliver the input DVHs. Next, those weights were input to an inverse planning problem, which was solved to generate an *inverse plan*. We use the name of the method that generated the input to the IO to identify the type of inverse plans (i.e., *clinical IO* (CIO) *plans*, *BQ plans*, or *gPCA plans*).

3.2.4 Analysis

We began by training both KBP methods using leave-one-out cross-validation, which resulted in two sets of predicted DVHs for each patient. Next, the predicted and clinical DVHs were input into the IO pipeline to generate inverse plans. We also generated another set of inverse

plans by re-solving the IPP with the same weights as the KBP plans, but subject to a constraint that kept their SPG less than or equal to the CIO SPG. This group of plans will be denoted with a prime (e.g., BQ′ or gPCA′), and represent plans with the same level of complexity as the CIO plans. After all this we have a clinical, CIO, BQ, gPCA, BQ′, and gPCA′ plan for each patient.

Analysis 1: DVH prediction error

To assess DVH prediction error, we computed the median fractional volume difference between each clinical DVH curve and the predicted DVH curve from both KBP methods. Positive differences correspond to instances where the clinical plans delivered less dose to a fractional volume than the prediction, and negative differences correspond to the reverse case. For each KBP method, the distribution of differences across the population was visualized with a box plot, and we used the two-sided Wilcoxon rank-sum test to determine if the predictions were significantly different from the clinical DVHs. A one-tailed Wilcoxon rank-sum test was also used to determine if the predictions were significantly greater than or less than clinical DVHs. For each fractional volume f of every ROI s over the population of predicted plans $\mathcal{J}$ (i.e., $|\mathcal{J}|$ = 217), we calculated the average dose $\tilde{D}^s_{f,\text{mean}}$ and $\bar{D}^s_{f,\text{mean}}$ of the BQ and gPCA predictions, respectively:

$$\tilde{D}^s_{f,\text{mean}} = \underset{j\in\mathcal{J}}{\text{mean}} \left\{ \tilde{D}^{j,t}_f \right\}, \quad \forall f \in \mathcal{F},\ \forall s \in \{\mathcal{I}, \mathcal{T}\}, \tag{3.12}$$

$$\bar{D}^s_{f,\text{mean}} = \underset{j\in\mathcal{J}}{\text{mean}} \left\{ \tilde{D}^{j,t}_f \right\}, \quad \forall f \in \hat{\mathcal{F}},\ \forall s \in \{\mathcal{I}, \mathcal{T}\}. \tag{3.13}$$

These averages were merged to generate the population average DVHs for both of the KBP methods, and they were compared to the clinical population average DVHs to illustrate the bias of each KBP method

Analysis 2: DVH reproduction error

We computed the median difference between DVHs input to the IO and their corresponding inverse plans. Over the each type of generated plan, the distribution of differences is displayed with a box plot; negative differences were when the inverse plan delivered less dose to a fractional

volume than the prediction, and positive differences correspond to the reverse case.

Analysis 3: Planning criteria satisfaction

We compared the clinical plans to the inverse plans in terms of satisfying the clinical planning criteria outlined in Table 3.1. Then at the individual criterion level, we evaluated the frequency with which the plan populations satisfied each criterion. For all inverse plans we also evaluated the SPG relative to the corresponding CIO plan, and recorded the average and standard deviation (SD) of the relative SPG for each type of plan. We also recognize that not all clinical criteria are weighted equally. Thus, we evaluated the proportion of inverse plans that satisfied the same clinical criteria as their respective clinical plan, and those that fell short.

Table 3.1: The planning criteria used to evaluate our plan. We define $\mathcal{D}_v$ as the dose to a fractional volume v, $\mathcal{D}_{mean}$ as the mean dose to a structure, and $\mathcal{D}_{max}$ as the max dose to a structure.

Structure	Criteria
Brain Stem	$\mathcal{D}_{max} \leq 54$ Gy
Spinal Cord	$\mathcal{D}_{max} \leq 48$ Gy
Right Parotid	$\mathcal{D}_{mean} \leq 26$ Gy
Left Parotid	$\mathcal{D}_{mean} \leq 26$ Gy
Larynx	$\mathcal{D}_{mean} \leq 45$ Gy
Esophagus	$\mathcal{D}_{mean} \leq 45$ Gy
Mandible	$\mathcal{D}_{max} \leq 73.5$ Gy
PTV70	$\mathcal{D}_{99} \geq 65.1$ Gy
PTV63	$\mathcal{D}_{99} \geq 58.6$ Gy
PTV56	$\mathcal{D}_{99} \geq 53.2$ Gy

Analysis 4: Planning criteria error

For all clinical planning criteria we evaluated the difference between the clinical plans and the inverse plans we generated. Over each of our plan populations, the distribution of differences was visualized with a box plot. Points along the negative horizontal axis are the inverse plans

that performed better than the clinical plans with respect to the clinical planning criteria; positive differences correspond to the reverse case.

3.3 Results

3.3.1 DVH prediction error

For each structure on the vertical axis, Figure 3.3 displays the difference between the clinical and predicted plan DVHs on the horizontal axis. Over all OAR fractional volumes, we found that DVHs predicted by the BQ method were significantly lower than the clinical DVHs ($p < 0.05$), but the predictions from the gPCA method were not significantly different from the clinical DVHs ($p > 0.05$). Similarly, the target DVHs produced by the BQ method were significantly lower than the clinical DVHs ($p < 0.05$); while the differences between the gPCA predictions and the clinical DVHs were not significant ($p > 0.05$). On average, the KBP methods produce DVHs that are noticeably different (Figure 3.4).

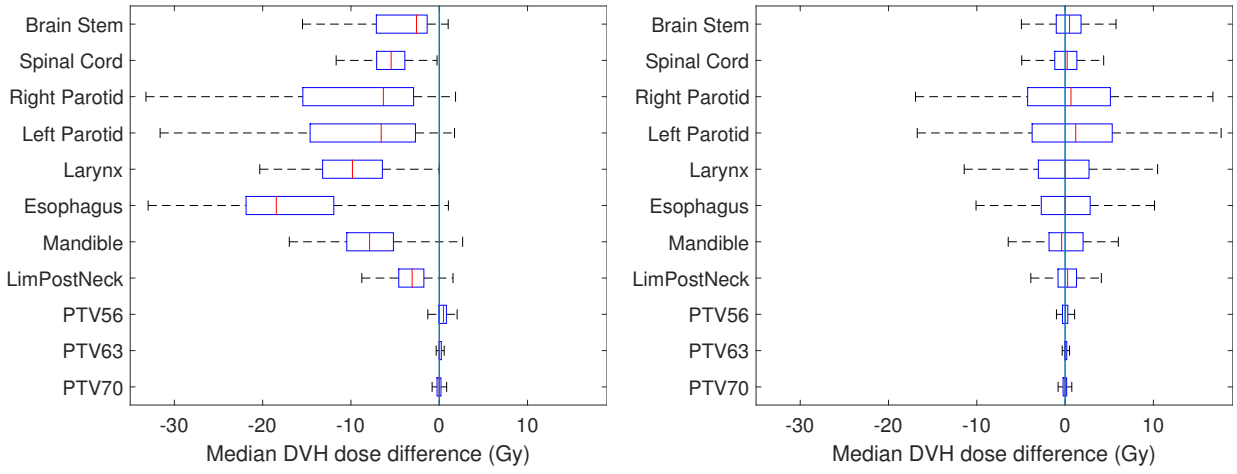

(a) BQ prediction minus clinical

(b) gPCA prediction minus clinical

Figure 3.3: The distribution of median DVH differences between clinical plans and KBP predictions. A positive difference implies the clinical plan achieved lower dose than the predicted plan. The upper and lower boundaries of each box represent the 75th and 25th percentiles respectively, and the horizontal line in the box depicts the median. Whiskers extend to 1.5 times the interquartile range. A line across each plot provides a reference point for zero difference.

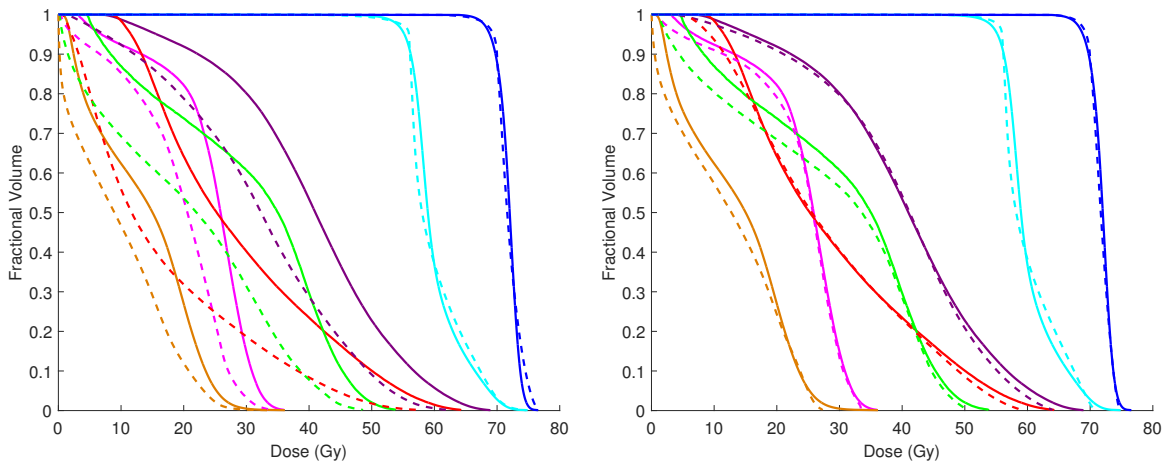

(a) BQ predictions and clinical

(b) gPCA prediction and clinical

Spinal Cord Mandible Left Parotid Esophagus LimPostNeck PTV56 PTV70
Clinical Plan KBP Prediction

Figure 3.4: The average DVH produced by BQ and gPCA predictions is benchmarked against the average clinical DVHs.

3.3.2 DVH reproduction error

Most of the inverse plans were similar to their respective input DVHs, as shown by Figure 3.5. Median DVH differences across all structures ranged from -2.78 Gy to 0.28 Gy, -1.21 Gy to 0.20 Gy, and -2.99 Gy to 0.18 Gy for the CIO, BQ, and gPCA plans, respectively. The most negative median difference occurred for the brain stem in the CIO and gPCA plans, and for the esophagus in the BQ plans. For the BQ′ and gPCA′ plans, however, the median DVH difference ranged from -0.73 Gy to 9.65 Gy and -2.87 Gy to 2.48 Gy, respectively, which were much wider ranges than the original KBP plans.

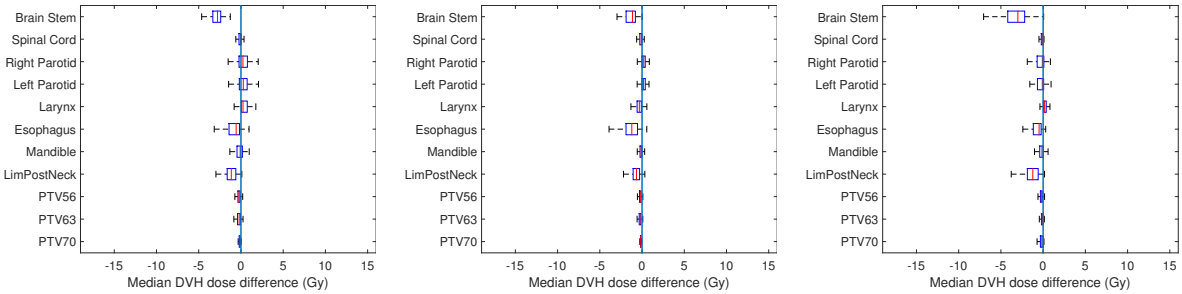

(a) CIO plan minus clinical (b) BQ plans (c) gPCA plans

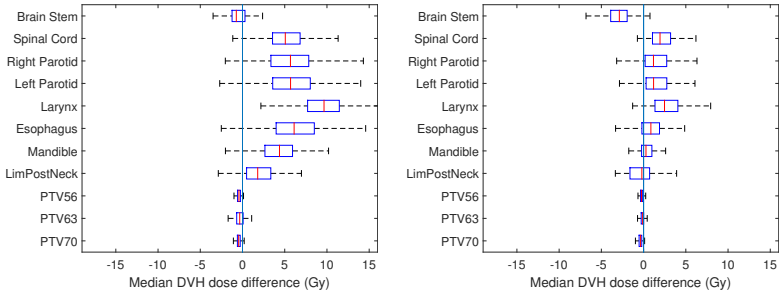

(d) BQ plans (e) gPCA plans

Figure 3.5: The distribution of median DVH differences between inverse plans and their respective input DVHs. A positive difference implies the inverse plan achieved lower dose than the input DVHs. The upper and lower boundaries of each box represent the 75th and 25th percentiles respectively, and the horizontal line in the box depicts the median. Whiskers extend to 1.5 times the interquartile range. A line across each plot provides a reference point for zero difference.

3.3.3 Clinical criteria satisfaction

Table 3.2 summarizes how often each population of plans satisfied the clinical planning criteria (Table 3.1). Across all criteria, the CIO plans were better than clinical plans at sparing OARs (by 1.9 percentage points) but worse at achieving target coverage (by 4.6 percentage points). The BQ plans also achieved superior OAR sparing compared to clinical plans (by 5.8 percentage points) but with less target criteria achieved (by 6.9 percentage points). In contrast, the gPCA plans achieved fewer OAR criteria than the clinical plans (by 4.0 percentage points) but achieved target coverage much more frequently (by 17.8 percentage points).

The BQ plans required a fluence map that was on average 33.6% more complex than the

CIO plans. While the gPCA plans were only 8.3% more complex than the CIO plans on average. With the SPG constrained, the typical BQ′ plan only satisfied 57.0% of the clinical criteria, which was 16.4 percentage points lower than the BQ plans. In contrast, the gPCA′ plans satisfied 71.0% of the clinical criteria, which was only 1.7 percentage points lower than the gPCA plans, but still better than the clinical plans by 0.3 percentage points.

Table 3.2: The frequency with which each type of plan satisfied the clinical planning criteria, and the relative SPG they required.

	Clinical	CIO	BQ	gPCA	BQ′	gPCA′
Brain Stem	100.0%	100.0%	99.5%	99.5%	99.1%	99.5%
Spinal Cord	100.0%	100.0%	100.0%	100.0%	73.3%	98.6%
Right Parotid	20.5%	21.9%	44.6%	19.4%	25.71%	16.9%
Left Parotid	17.1%	18.0%	48.6%	15.1%	19.9%	12.8%
Larynx	61.7%	61.7%	89.6%	61.7%	69.2%	44.3%
Esophagus	96.7%	97.6%	100.0%	99.1%	99.5%	98.6%
Mandible	78.8%	88.9%	39.5%	52.3%	47.9%	64.1%
All OARs	68.2%	70.1%	74.4%	64.2%	62.2%	62.4%
PTV56	52.5%	43.8%	27.1%	90.4%	9.7%	87.6%
PTV63	100.0%	100.0%	100.0%	100.0%	100.0%	99.2%
PTV70	88.9%	85.7%	96.3%	97.3%	43.8%	96.3%
All Targets	77.4%	72.8%	70.5%	95.2%	43.4%	93.6%
All Structures	70.7%	70.9%	73.4%	72.7%	57.0 %	71.0%
Average SPG	—	1.000	1.336	1.083	1.000	0.999
SPG SD	—	0.000	0.231	0.076	0.000	0.004

Table 3.3 summarizes the proportion of treatment plans that satisfy the same clinical criteria as the clinical plans, and the proportion with some unsatisfied. Across all patients, only 39.17% of the BQ plans achieved the same clinical criteria as the clinical plans compared to 52.07% for the gPCA plans. When the SPG was constrained the differences were much larger. For example, the clinical criteria satisfied in the clinical plans was achieved in 85.71% of CIO plans, but only 14.29% and 49.77% for the BQ′ and gPCA′ plans, respectively.

Table 3.3: The frequency with which each type of plan achieved the clinical criteria that the clinical plans satisfied.

Achieved criteria	CIO	BQ	gPCA	BQ′	gPCA′
0	85.71%	39.17%	52.07%	14.29%	49.77%
1 or fewer	100.00%	82.03%	92.17%	43.32%	87.10%
2 or fewer	100.00%	99.54%	99.54%	80.18%	98.16%
3 or fewer	100.00%	99.54	99.54%	97.70%	99.54%
4 or fewer	100.00%	100.00%	100.00%	100.00%	100.00%

3.3.4 Clinical criteria error

Figure 3.6 illustrates how our generated plans performed in terms the dose delivered at the clinical planning criteria, relative to the clinical plans. The CIO plans were the most similar to the clinical plans; with the median difference in each criteria ranging from -1.31 Gy to 0.24 Gy. The next most similar plans were the gPCA′ plans, which had median differences that ranged from -0.72 Gy to 1.56 Gy. Overall the BQ plans had the most negative clinical criteria difference of -10.51 Gy, and the BQ′ had the most positive median difference of 8.92 Gy.

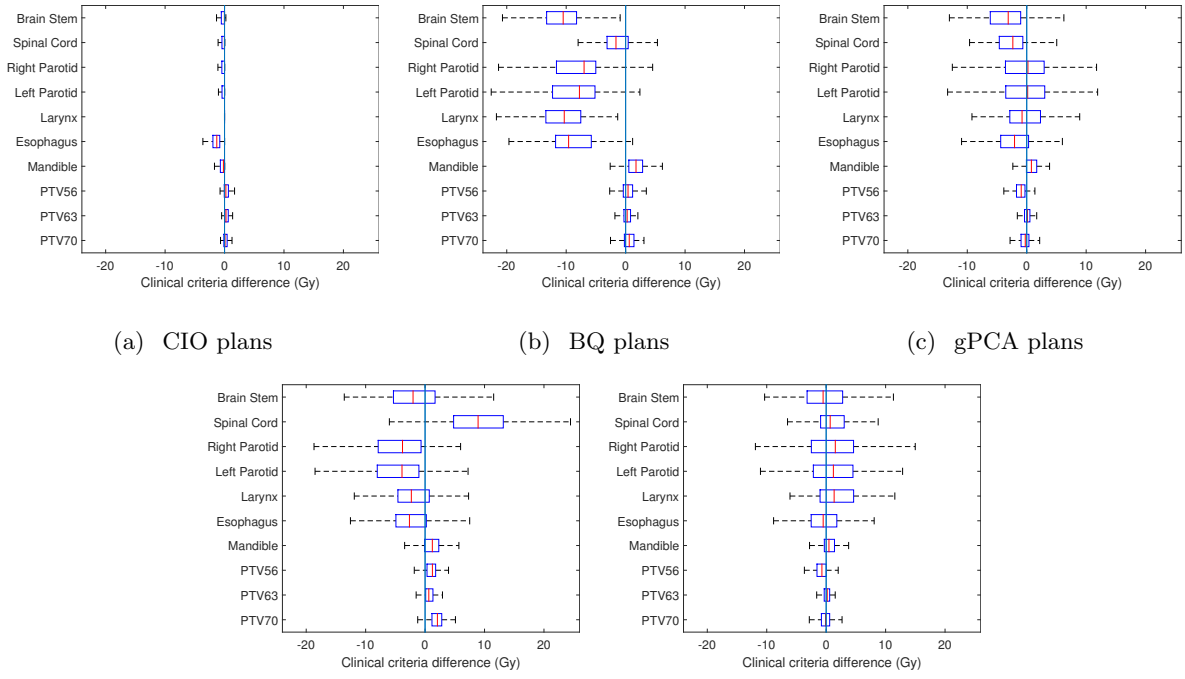

(a) CIO plans (b) BQ plans (c) gPCA plans

(d) BQ' plans (e) gPCA' plans

Figure 3.6: The difference between the inverse and clinical plans, in terms of dose delivered at the clinical planning criteria. A positive difference implies that the clinical plan was better than the inverse plan in that criterion. The upper and lower boundaries of each box represent the 75th and 25th percentiles respectively, and the horizontal line in the box depicts the median. Whiskers extend to 1.5 times the interquartile range. A line across each plot provides a reference point for zero difference.

3.4 Discussion

In this paper, we developed a methodology to unite knowledge-based planning with inverse optimization to automatically generate treatment plans from patient anatomy. Previous KBP methods were inadmissible to IO, so we the extended methods from the literature to address that shortcoming. Our methods typically predicted very different DVHs. Nevertheless, all predictions were input into an IO pipeline, and the DVHs were accurately reproduced, but with a clear discrepancy in fluence complexity. We evaluated each component of the pipeline, and the pipeline as a whole. We examined how close the KBP predictions were to clinical plans, and how well the IO method recreated those plans. The final plans were evaluated based on

how they met a selection of clinical planning criteria (Table 3.1).

Two existing KBP methods were extended, so they predicted achievable DVHs for all ROIs. The BQ method is an extension of previous query methods[28;29;30;31] that protects against overfitting with a variation of *bagging*, and predicts full DVHs as opposed to single points. We also generalized the PCA method,[36;41] so that it predicts the DVHs for all ROIs in sites with multiple targets. Both of these new KBP methods also ensure that the predicted DVHs are physically possible. We evaluated the difference between each clinical and predicted DVH, and found that the two methods had different bias. The BQ method generally predicted that OARs could be spared much more than the clinical plans achieved. In contrast, the gPCA method generally predicted more conservative DVHs that were similar to the clinical curves. This is similar to the findings in previous work;[3] where a query-based method overfit the training set, and the more conservative PCA-based approach did not. Beyond speculation, however, any tradeoffs made to achieve an overfit prediction were unclear.

The stage where a planner manually tunes objective function weights to recreate a KBP DVHs should automated, which would require a method to convert predicted DVHs into objective function weights. We demonstrated that a IO designed specifically to recreate DVHs[1] can be that automated method. The IO method leverages the DVHs to construct objectives functions that are tailored specifically for the patient-specific DVHs. In contrast, previous methods involving IO and machine learning generated weights directly from patient geometry in a way that is restricted to one site and with a fixed set of objective functions.[4;16] Integrating other KBP methods or applying our framework to other sites is also more straightforward than the previous approaches.

Our pipeline reproduced the DVHs from two dissimilar KBP methods within a small median DVH difference, which demonstrates that both KBP methods produce achievable DVHs. Each population of KBP plans, however, required vastly different SPGs to deliver; suggesting that information like fluence heterogeneity is encoded in predicted DVHs. In previous automated planning methods the plans were limited to a fixed measure of complexity, so the idea of an automated methods inferring the required complexity is novel.[13;18;28] One of these techniques[28] also leveraged a query prediction, however, they recorded much better target coverage with poor OAR sparing, which was a stark contrast to our BQ plans. When we limited the KBP

plans to be within the same SPG as the CIO plans, we found that they performed much worse than the original KBP plans in terms of clinical criteria; although the degradation was much more severe in the BQ′ plans than the gPCA′ plans. Nevertheless, nearly half of all gPCA′ plans were non-inferior to their clinical counterparts in terms of clinical criteria satisfaction, and over 85% violated one or fewer criteria that the clinical plans achieved. This finding illustrates that our pipeline can provide a "warm-start" for the IPP.

A limitation of this work is its dependence on the quality and quantity of the plan library used to train the KBP methods. Limiting the dependence that our pipeline has on KBP could help generate more consistently high quality plans. Another limitation of this approach is the omission of objectives that favour favourable features in a full three dimensional dose distribution. Extensions of this work could examine three dimensional KBP methods, and incorporate geometry into the dose objectives. Moreover, the IO method is a linear program, so the non-convex clinical planning criteria for targets are only included implicitly. Nevertheless, a linear IPP can generate plans that achieve sufficient target coverage, which was demonstrated over the population of gPCA plans.

3.5 Conclusion

In this paper we developed a framework the unites knowledge-based planing with inverse optimization, and created an knowledge-based automated treatment planning pipeline. The approach was tested on over 200 oropharyngeal patients, which represents the largest dataset used for an automated planning method for head and neck. We demonstrated that nearly 50% of plans generated by our pipeline can achieve the same criteria as the clinical plans, and over 85% violated one or fewer criteria that the clinical plans achieved. Once implemented into a clinical planning pipeline, tuning weights to achieve predicted DVHs will be fully automated, and planners will only need to make small adjustment to tailor the plan for a new patient.

Chapter 4

Conclusion and Future Directions

The success of an IMRT treatment relies heavily on the quality of its treatment plan, yet it is highly variable within and across clinics. An abundance of KBP methods attempt to address this issue by providing treatment planners with patient-specific guideline DVH curves to make planning more consistent. However, these KBP methods still require planners to do manual parameter tuning. Optimization methods can translate these DVH guidelines into objective function weights, which can be used as direct input for an IPP. As such, new treatment planning tools that use these methods can be integrated seamlessly into the current planning workflow. In this thesis, we developed the foundation for a pipeline that automatically generates treatment plans from patient anatomical geometry. The pipeline consists of two main components. The first component is one of two KBP methods that use patient geometry to predict DVH curves. The predictions are inputs to the second component, which is an IO method that reverse engineers objective function weights from DVH curves. Together, these components form a novel knowledge-based automated planning method.

At the core of our automated pipeline is the IO method that we developed. We tested our formulation with 217 clinical treatment plans, which is the largest patient cohort used to validate a method that reconstructs treatment plans from DVHs. The method retrieves the objective function weights and infers the complexity of a fluence map that is required to deliver a set of clinical DVHs. We showed that inferring the fluence map SPG produces less complex plans than the previous IO method applied to radiation therapy, which constrained the intensity of all beamlets in a fluence map. Moreover, our IO method produced objective function weights that

create plans with greater clinical criteria satisfaction than the plans generated by the previous IO method.

We adapted two KBP methods from the literature to predict the DVHs for all regions of interest in a disease site with multiple targets. Each KBP was tested on a large cohort of oropharyngeal treatment plans using leave-one-out cross validation. These predicted DVHs were inputs to our inverse optimization method, which produced the objective function weights for the IPP used to generate high-quality treatment plans. These plans can then be adjusted using the familiar IPP tuning approach to meet the personal preferences of the patient. With our most promising KBP method, over 85% of the generated plans satisfied the same, or one less, clinical criteria as the original clinical plans

There are several future directions for this work. A natural extension would be aperture optimization, which takes our fluence maps as inputs or direct aperture optimization, which takes objective function weights as inputs, to generate the series of MLC aperture formations necessary to deliver the treatment plan. Additionally, our IPP problem treats all voxels of a structure equally, but this is not reflective of reality. Generally there are particular subregions of OARs that play a more critical role in preserving a patient's quality of life. Although we limited our approach to a structure-based multi-objective optimization problem, generalizing our methodology to a voxel-based optimization problem could address this issue. Lastly, our KBP and IO methods relied heavily on DVH curves, and did not consider the spatial dose distribution within the ROIs, which is also an indicator of plan quality. This means that two plans with the same DVHs may have different spatial dose distributions that result in very different patient outcomes. It is straightforward to apply our IO model to a three-dimensional dose distribution. However, this would require to accompanying changes to other components of our pipeline: the inclusion of three-dimensional spatial information in the objective or constraints of the IPP to leverage the additional information, and a KBP method that predicts a full three dimensional dose distribution to use as input for the IO model.

Bibliography

[1] A. Babier, J. J. Boutilier, M. B. Sharpe, A. L. McNiven, and T. C. Y. Chan. Reverse engineering fluence maps from dose volume histograms using inverse optimization and inverse planning.

[2] D. Bertsimas and J. N. Tsitsiklis. Introduction to linear optimization. *Athena Scientific*, 1997.

[3] J. J. Boutilier, T. Craig, M. B. Sharpe, and T. C. Y Chan. Sample size requirements for knowledge-based treatment planning. *Med. Phys.*, 43(3):1212–21, 2016.

[4] J. J. Boutilier, T. Lee, T. Craig, M. B. Sharpe, and T. C. Y. Chan. Models for predicting objective function weights in prostate cancer IMRT. *Med. Phys.*, 42(4):1586–95, 2015.

[5] Canadian Cancer Society. *Canadian Cancer Society's Advisory Committee on Cancer Statistics. Canadian Cancer Statistics 2016.* 2016.

[6] T. C. Y. Chan, T. Craig, T. Lee, and M. B. Sharpe. Generalized inverse multiobjective optimization with application to cancer therapy. *Oper. Res.*, 62(3):680–95, 2014.

[7] T. C. Y. Chan, T. Lee, and D. Terekhov. Goodness of fit in inverse optimization. 2015.

[8] C. Cotrutz and L. Xing. Using voxel-dependent importance factors for interactive DVH-based dose optimization. *Phys. Med. Biol.*, 47(10):1659–69, 2002.

[9] D. Craft, P. Suss, and T. Bortfeld. The tradeoff between treatment plan quality and required number of monitor units in intensity-modulated radiotherapy. *Int. J. Radiat. Oncol. Biol. Phys.*, 67(5):1596–605, 2007.

[10] J. O. Deasy, A. I. Blanco, and V. H. Clark. CERR: a computational environment for radiotherapy research. *Med. Phys.*, 30(5):979–85, 2003.

[11] G. A. Ezzell, J. M. Galvin, D. Low, J. R. Palta, I. Rosen, M. B. Sharpe, P. Xia, Y. Xiao, L. XIng, and C. X. Yu. Guidance document on delivery, treatment planning, and clinical implementation of IMRT: report of the IMRT subcommittee of the AAPM radiation therapy committee. *Med. Phys.*, 30(8):2089–115, 2003.

[12] H. W. Hamacher and K.-H. Küfer. Inverse radiation therapy planning—a multiple objective optimization approach. *Discrete Appl. Math.*, 118:145– 161, 2002.

[13] I. Hazell, K. Bzdusek, P. Kumar, C. R. Hansen, A. Bertelsen, J. G. Eriksen, J. Johansen, and C. Brink. Automatic planning of head and neck treatment plans. *J. Appl. Clin. Med. Phys.*, 17(1):272–282, 2016.

[14] A. Holder. Designing radiotherapy plans with elastic constraints and interior point methods. *Health Care Manag. Sci.*, 6(1):5–16, Feb 2003.

[15] C. Holdsworth, M. Kim, J. Liao, and M. Phillips. The use of a multiobjective evolutionary algorithm to increase flexibility in the search for better IMRT plans. *Med. Phys.*, 39(4):2261 – 74, 2012.

[16] T. Lee, M. Hammad, T. C. Y. Chan, T. Craig, and M. B. Sharpe. Predicting objective function weights from patient anatomy in prostate IMRT treatment planning. *Med. Phys.*, 40:121706, 2013.

[17] N. Li, M. Zarepisheh, A. Uribe-Sanchez, K. Moore, Z. Tian, X. Zhen, Y. J. Graves, Q. Gautier, L. Mell, L. Zhou, X. Jia, and S. Jiang. Automatic treatment plan re-optimization for adaptive radiotherapy guided with the initial plan DVHs. *Phys. Med. Biol.*, 58(24):8725–38, 2013.

[18] C. McIntosh, M. Welch, A. McNiven, D. A. Jaffray, and T. G. Purdie. Fully automated treatment planning for head and neck radiotherapy using a voxel-based dose prediction and dose mimicking method. *Phys. Med. Biol.*, 62(15):5926–5944, 2017.

[19] A. L. McNiven, M. B. Sharpe, and T. G. Purdie. A new metric for assessing IMRT modulation complexity and plan deliverability. *Med. Phys.*, 37(2):505–15, 2010.

[20] T. G. Purdie, R. E. Dinniwell, A. Fyles, and M. B. Sharpe. Automation and intensity modulated radiation therapy for individualized high-quality tangent breast treatment plans. *Int. J. Radiat. Oncol. Biol. Phys.*, 90(3):688–95, 2014.

[21] H. E. Romeijn, R. K. Ahuja, J. F. Dempsey, and A. Kumar. A new linear programming approach to radiation therapy treatment planning problems. *Oper. Res.*, 54(2):201 – 16, 2006.

[22] M. B. Sharpe, K. L. Moore, and C. G. Orton. Point/counterpoint: Within the next ten years treatment planning will become fully automated without the need for human intervention. *Med. Phys.*, 41(12):120601, 2014.

[23] S. Shiraishi and K. L. Moore. Knowledge-based prediction of three-dimensional dose distributions for external beam radiotherapy. *Med. Phys.*, 43(1):378, 2016.

[24] S. Shiraishi, J. Tan, L. A. Olsen, and K. L. Moore. Knowledge-based prediction of plan quality metrics in intracranial stereotactic radiosurgery. *Med. Phys.*, 42(2):908, 2015.

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

[26] J. P. Tol, A. R. Delaney, M. Dahele, B. J. Slotman, and W. F. A. R. Verbakel. Evaluation of a knowledge-based planning solution for head and neck cancer. *Int. J. Radiat. Oncol. Biol. Phys.*, 91(3):612–20, 2015.

[27] S. Webb, D. J. Convery, and P. M. Evans. Inverse planning with constraints to generate smoothed intensity-modulated beams. *Phys. Med. Biol.*, 43(10):2785–94, 1998.

[28] B. Wu, M. Kusters, M. Kunze-Busch, T. Dijkema, T. McNutt, G. Sanguineti, K. Bzdusek, A. Dritschilo, and D. Pang. Cross-institutional knowledge-based planning (KBP) implementation and its performance comparison to auto-planning engine (APE). *Radiother. Oncol.*, 123(1):57–62, 2017.

[29] B. Wu, T. McNutt, M. Zahurak, P. Simari, D. Pang, R. Taylor, and G. Sanguineti. Fully automated simultaneous integrated boosted-intensity modulated radiation therapy treatment planning is feasible for head-and-neck cancer: a prospective clinical study. *Int. J. Radiat. Oncol. Biol. Phys.*, 84(5):e647–53, 2012.

[30] B. Wu, F. Ricchetti, G. Sanguineti, M. Kazhdan, P. Simari, M. Chuang, R. Taylor, R. Jacques, and T. McNutt. Patient geometry-driven information retrieval for IMRT treatment plan quality control. *Med. Phys.*, 36(12):5497–505, 2009.

[31] B. Wu, F. Ricchetti, G. Sanguineti, M. Kazhdan, P. Simari, R. Jacques, R. Taylor, and T. McNutt. Data-driven approach to generating achievable dose-volume histogram objectives in intensity-modulated radiotherapy planning. *Int. J. Radiat. Oncol. Biol. Phys.*, 79(4):1241–7, 2011.

[32] C. Wu, G. H. Olivera, R. Jeraj, H. Keller, and T. R. Mackie. Treatment plan modification using voxel-based weighting factors/dose prescription. *Phys. Med. Biol.*, 48(15):2479–91, 2003.

[33] L. Xing, J. G. Li, S. Donaldson, Q. T. Le, and A. L. Boyer. Optimization of importance factors in inverse planning. *Phys. Med. Biol.*, 44(10):2525–36, 1999.

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

[35] Y. Yu. Multiobjective decision theory for computational optimization in radiation therapy. *Med. Phys.*, 24(9):1445–54, 1997.

[36] L. Yuan, Y. Ge, W. R. Lee, F. F. Yin, J. P. Kirkpatrick, and Q. J. Wu. Quantitative analysis of the factors which affect the interpatient organ-at-risk dose sparing variation in IMRT plans. *Med. Phys.*, 39(11):6868–78, 2012.

[37] M. Zarepisheh, T. Long, N. Li, Z. Tian, H. E. Romeijn, X. Jia, and S.B. Jiang. A DVH-

guided IMRT optimization algorithm for automatic treatment planning and adaptive radiotherapy replanning. *Med. Phys.*, 41(6):061711, 2014.

[38] M. Zarepisheh, M. Shakourifar, G. Trigila, P. S. Ghomi, S. Couzens, A. Abebe, L. Noreña, S. B. Jiang W. Shang, and Y. Zinchenko. A moment-based approach for DVH-guided radiotherapy treatment plan optimization. *Phys. Med. Biol.*, 58:1869–1887, 2013.

[39] M. Zarepisheh, A. F. Uribe-Sanchez, N. Li, X. Jia, and S. B. Jiang. A multicriteria framework with voxel-dependent parameters for radiotherapy treatment plan optimization. *Med. Phys.*, 41(4):041705, 2014.

[40] Y. Zhang and M. Merritt. Dose-volume-based IMRT fluence optimization: A fast least-squares approach with differentiability. *Linear Algebra and its Applications*, 428:1365 – 87, 2008.

[41] X. Zhu, Y. Ge, T. Li, D. Thongphiew, F. Yin, and Q. J. Wu. A planning quality evaluation tool for prostate adaptive IMRT based on machine learning. *Med. Phys.*, 38(2):719–26, 2011.

Appendix A

Appendix

A.1 Matrix Notation

In this section, we provide the explicit details of the matrix notation presented in the Methods. Let $\mathbf{e}_\zeta$ be a vector of all ones and $\mathbf{0}_\zeta$ be a vector of all zeros, both of length ζ. A matrix of zeros is denoted as $\mathbf{0}_{\zeta\times\omega}$ with ζ rows and ω columns. The identity matrix with ζ rows and columns is denoted $\mathbf{I}_\zeta$. The matrix $\mathbf{E}_{\zeta\times\omega}$ is block diagonal with ω blocks, each consisting of the column vector $\mathbf{e}_\zeta$.

A.1.1 OAR objective matrix notation

The variables for each OAR $i \in \mathcal{I}$ are represented by $\mathbf{x}^i$, which contains the mean dose z^i, max dose y^i, and a vector $\mathbf{h}$ with elements h^f for each penalty $f \in \mathcal{F}^i$. A vector of $\boldsymbol{\rho}_v^f$ is also included to represent the dose exceeding each penalty $f \in \mathcal{F}^i$ for all voxels $v \in \mathcal{O}^i$:

$$\mathbf{x}^i = \begin{bmatrix} z^i \\ y^i \\ \mathbf{h}^f \\ \boldsymbol{\rho}_v^f \end{bmatrix}. \tag{A.1}$$

These $\mathbf{x}^i$ vectors are manipulated by an objective function matrix $\mathbf{C}^i$ to linearize the OAR objectives. The matrix $\mathbf{C}^i$ has a column for each variable in $\mathbf{x}^i$, and a row for each objective

function:

$$\mathbf{C}^i = \begin{bmatrix} 1 & & & \\ & 1 & & \\ & & \mathbf{I}_{|\mathcal{F}^\mathbf{i}|} & \mathbf{0}_{|\mathcal{F}|\times|\mathcal{F}^i|\cdot|\mathcal{O}^i|} \end{bmatrix}. \tag{A.2}$$

The weights for each objective are contained in $\boldsymbol{\alpha}^i$. The elements in vector $\boldsymbol{\kappa}$ correspond to each objective function value κ^f for each $f \in \mathcal{F}$:

$$\boldsymbol{\alpha}^i = \begin{bmatrix} \gamma^i \\ \beta^i \\ \boldsymbol{\kappa} \end{bmatrix}. \tag{A.3}$$

Each variable in $\mathbf{x}^i$ has a column in $\mathbf{A}^i$ corresponding to the constraint matrix:

$$\mathbf{A}^i = \begin{bmatrix} 1 & & & \\ & \mathbf{e}_{|\mathcal{O}^\mathbf{i}|} & & \\ & & \mathbf{I}_{|\mathcal{F}^\mathbf{i}|} & -\frac{1}{|\mathcal{O}^i|}\mathbf{E}_{|\mathcal{F}|\times|\mathcal{F}^\mathbf{i}|\cdot|\mathcal{O}^\mathbf{i}|} \\ & & & \mathbf{I}_{|\mathcal{F}^\mathbf{i}|\cdot|\mathcal{O}^\mathbf{i}|} \\ & & & \mathbf{I}_{|\mathcal{F}^\mathbf{i}|\cdot|\mathcal{O}^\mathbf{i}|} \end{bmatrix}. \tag{A.4}$$

The elements of the vector $\mathbf{A}^i\mathbf{x}^i$ must be less than or equal to each element in the vector $\mathbf{b}^i$. Vector $\mathbf{f}$ represents a vector of the linear penalties in $\mathcal{F}$, where each $f \in \mathcal{F}$ is repeated $|\mathcal{O}^i|$ times:

$$\mathbf{b}^i = \begin{bmatrix} 0 \\ \mathbf{0}_{|\mathcal{O}^\mathbf{i}|} \\ \mathbf{0}_{|\mathcal{F}^\mathbf{i}|} \\ -\mathbf{f} \\ \mathbf{0}_{|\mathcal{F}^\mathbf{i}|\cdot|\mathcal{O}^\mathbf{i}|} \end{bmatrix}. \tag{A.5}$$

These matrices are linked to a vector of beamlet intensities $\mathbf{w}$ via a matrix $\boldsymbol{\Lambda}^i$. The dose influence matrix for OAR i is $\mathbf{D}^i$, and $\mathbf{D}^i_{|\mathcal{F}^i|}$ corresponds to an influence matrix that is vertically

concatenated $|\mathcal{F}^i|$ times (i.e., repeated):

$$\mathbf{\Lambda}^i = \begin{bmatrix} -\frac{\mathbf{1}}{|\mathcal{O}^\mathbf{i}|}\mathbf{e}'_{|\mathcal{O}^\mathbf{i}|}\mathbf{D} \\ -\mathbf{D} \\ \mathbf{0}_{|\mathcal{F}^\mathbf{i}|} \\ -\mathbf{D}_{|\mathcal{F}^\mathbf{i}|} \\ \mathbf{0}_{|\mathcal{F}^\mathbf{i}|\cdot|\mathcal{O}^\mathbf{i}|\times|\mathcal{B}|} \end{bmatrix}. \tag{A.6}$$

A.1.2 Target objective matrix notation

The matrix form of objectives, constraints, and variables for target structures $t \in \mathcal{T}$ are simiilar to those for the OAR objectives:

$$\mathbf{A}^t = \begin{bmatrix} \mathbf{e}_{|\mathcal{O}^\mathbf{t}|} & & & \\ & 1 & -\frac{1}{|\mathcal{O}^t|}\mathbf{e}'_{|\mathcal{O}^\mathbf{t}|} & \\ & & 1 & -\frac{1}{|\mathcal{O}^t|}\mathbf{e}'_{|\mathcal{O}^\mathbf{t}|} \\ & & \mathbf{I}_{|\mathcal{O}^\mathbf{t}|} & \\ & & & \mathbf{I}_{|\mathcal{O}^\mathbf{t}|} \\ & & \mathbf{I}_{|\mathcal{O}^\mathbf{t}|} & \\ & & & \mathbf{I}_{|\mathcal{O}^\mathbf{t}|} \end{bmatrix} \quad \mathbf{\Lambda}^t = \begin{bmatrix} -\mathbf{D}_{|\mathcal{O}^\mathbf{t}|\times|\mathcal{B}|} \\ 0 \\ 0 \\ -\mathbf{D}_{|\mathcal{O}^\mathbf{t}|\times|\mathcal{B}|} \\ \mathbf{0}_{|\mathcal{O}^\mathbf{t}|\times|\mathcal{B}|} \\ \mathbf{D}_{|\mathcal{O}^\mathbf{t}|\times|\mathcal{B}|} \\ \mathbf{0}_{|\mathcal{O}^\mathbf{t}|\times|\mathcal{B}|} \end{bmatrix} \quad \mathbf{b}^t = \begin{bmatrix} \mathbf{0}_{|\mathcal{O}^\mathbf{t}|} \\ 0 \\ 0 \\ -\theta^\mathbf{t} \cdot \mathbf{e}_{|\mathcal{O}^\mathbf{t}|} \\ \mathbf{0}_{|\mathcal{O}^\mathbf{t}|} \\ \theta^\mathbf{t} \cdot \mathbf{e}_{|\mathcal{O}^\mathbf{t}|} \\ \mathbf{0}_{|\mathcal{O}^\mathbf{t}|} \end{bmatrix}$$

$$\boldsymbol{\alpha}^t = \begin{bmatrix} \beta^t \\ \phi^t \\ \psi^t \end{bmatrix} \quad \mathbf{C}^t = \begin{bmatrix} 1 & & & \\ & 1 & & \\ & & 1 & \mathbf{0}'_{2\cdot|\mathcal{O}^t|} \end{bmatrix} \quad \mathbf{x}^t = \begin{bmatrix} y^t \\ u^t \\ l^t \\ \boldsymbol{\mu}_v^t \\ \boldsymbol{\epsilon}_v^t \end{bmatrix}$$

A.1.3 Beamlet objective matrix notation

The variables for the SPG are represented by $\mathbf{x}^{\mathcal{K}}$. A vector $\mathbf{m}$ is defined with elements equal to m^k for each $k \in \mathcal{K}$. Another vector $\boldsymbol{\lambda}$ is defined with an element λ_b for each $b \in \mathcal{B}$

$$\mathbf{x}^{\mathcal{K}} = \begin{bmatrix} \mathbf{m} \\ \boldsymbol{\lambda} \end{bmatrix}. \tag{A.7}$$

These vectors are manipulated by an objective function $\mathbf{g}^{\mathcal{K}}$ to linearize the SPG objectives

$$\mathbf{g}^{\mathcal{K}} = \begin{bmatrix} \mathbf{e}_{|\mathcal{K}|} \\ \mathbf{0}_{|\mathcal{B}|} \end{bmatrix}. \tag{A.8}$$

Each variable in $\mathbf{x}^{\mathcal{K}}$ has a column in $\mathbf{A}^{\mathcal{K}}$ corresponding to the constraint matrix:

$$\mathbf{A}^{\mathcal{K}} = \begin{bmatrix} \mathbf{0}_{|\mathcal{B}|\times|\mathcal{K}|} & \mathbf{I}_{|\mathcal{B}|\times|\mathcal{B}|} \\ \mathbf{0}_{|\mathcal{B}|\times|\mathcal{K}|} & \mathbf{I}_{|\mathcal{B}|\times|\mathcal{B}|} \\ \mathbf{E}_{|\mathcal{R}|\times|\mathcal{K}|} & -\mathbf{E}_{|\mathcal{R}|\times|\mathcal{B}|} \end{bmatrix}. \tag{A.9}$$

The elements of the vector $\mathbf{A}^{\mathcal{K}}\mathbf{x}^{\mathcal{K}}$ must be less than or equal to each element in the vector $\mathbf{b}^{\mathcal{K}}$:

$$\mathbf{b}^{\mathcal{K}} = \begin{bmatrix} \mathbf{0}_{|\mathcal{B}|} \\ \mathbf{0}_{|\mathcal{B}|} \\ \mathbf{0}_{|\mathcal{R}|} \end{bmatrix}. \tag{A.10}$$

These matrices are linked to a vector of beamlet intensities $\mathbf{w}$ via a matrix $\boldsymbol{\Lambda}^{\mathcal{K}}$. The matrix $\mathbf{G}_{|\mathcal{B}|\times|\mathcal{B}|}$ is constructed such that $\mathbf{G}_{|\mathcal{B}|\times|\mathcal{B}|}\mathbf{w}$ computes $w_b - w_{b'}$ for each $b \in \mathcal{B}$:

$$\boldsymbol{\Lambda}^{\mathcal{K}} = \begin{bmatrix} -\mathbf{G}_{|\mathcal{B}|\times|\mathcal{B}|} \\ \mathbf{0}_{|\mathcal{B}|\times|\mathcal{B}|} \\ \mathbf{0}_{|\mathcal{R}|\times|\mathcal{B}|} \end{bmatrix}. \tag{A.11}$$

A.1.4 Complete matrix notation

This subsection combines the matrices and vectors into those from the IPP in formulation (2.9)

$$\mathbf{A} = \begin{bmatrix} \mathbf{A}^i & & & & \boldsymbol{\Lambda}^i \\ & \ddots & & & \vdots \\ & & \mathbf{A}^t & & \boldsymbol{\Lambda}^t \\ & & & \ddots & \vdots \\ & & & \mathbf{A}^w & \boldsymbol{\Lambda}^{\mathcal{K}} \end{bmatrix} \quad \mathbf{b} = \begin{bmatrix} \mathbf{b}^i \\ \vdots \\ \mathbf{b}^t \\ \vdots \\ \mathbf{b}^{\mathcal{K}} \end{bmatrix} \quad \mathbf{x} = \begin{bmatrix} \mathbf{x}^i \\ \vdots \\ \mathbf{x}^t \\ \vdots \\ \mathbf{x}^{\mathcal{K}} \\ \mathbf{w} \end{bmatrix}$$

$$\boldsymbol{\alpha} = \begin{bmatrix} \boldsymbol{\alpha}^i \\ \vdots \\ \boldsymbol{\alpha}^t \\ \vdots \end{bmatrix} \quad \mathbf{C} = \begin{bmatrix} \mathbf{C}^i & & & \\ & \ddots & & \\ & & \mathbf{C}^t & \\ & & & \ddots \end{bmatrix} \quad \mathbf{g} = \begin{bmatrix} 0 \\ \vdots \\ \mathbf{g}^{\mathcal{K}} \\ \mathbf{0}_{|\mathcal{B}|} \end{bmatrix}$$

A.2 Detailed optimization models

A.2.1 Inverse planning problem

The following model is the complete linear program corresponding to the IPP (2.9).

$$
\begin{aligned}
\underset{z,y,x,l,u,m,g,w}{\text{minimize}} \quad & \sum_{i\in\mathcal{I}}\left(\gamma^i z^i + \beta^i y^i + \sum_{f\in\mathcal{F}^i}\kappa^f \sum_{v\in\mathcal{O}^i} x_v^f\right) + \sum_{t\in\mathcal{T}}\left(\beta^t y^t + \phi^t l^t + \psi^t u^t\right) + \sum_{k\in\mathcal{K}} m^k \\
\text{subject to} \quad & z^i = \frac{1}{|\mathcal{O}^i|}\sum_{v\in\mathcal{O}^i}\sum_{k\in\mathcal{K}}\sum_{r\in\mathcal{R}^k}\sum_{b\in\mathcal{B}^r} D_{v,b}w_b, \quad \forall i\in\mathcal{I}, \\
& y^i \geq \sum_{k\in\mathcal{K}}\sum_{r\in\mathcal{R}^k}\sum_{b\in\mathcal{B}^r} D_{v,b}w_b, \quad \forall v\in\mathcal{O}^i,\ \forall i\in\mathcal{I}, \\
& y^t \geq \sum_{k\in\mathcal{K}}\sum_{r\in\mathcal{R}^k}\sum_{b\in\mathcal{B}^r} D_{v,b}w_b, \quad \forall v\in\mathcal{O}^t, \forall t\in\mathcal{T}, \\
& h^f \geq \sum_{v\in\mathcal{O}^i}\rho_v^f, \quad \forall f\in\mathcal{F}^i,\ \forall i\in\mathcal{I}, \\
& \rho_v^f \geq \frac{1}{|\mathcal{O}^i|}\left(\sum_{k\in\mathcal{K}}\sum_{r\in\mathcal{R}^k}\sum_{b\in\mathcal{B}^r} D_{v,b}w_b - f\right), \quad \forall v\in\mathcal{O}^i,\ \forall f\in\mathcal{F}^i,\ \forall i\in\mathcal{I}, \\
& \rho_v^f \geq 0, \quad \forall v\in\mathcal{O}^i,\ \forall f\in\mathcal{F}^i,\ \forall i\in\mathcal{I}, \\
& l^t = \sum_{v\in\mathcal{O}^t}\epsilon_v^t, \quad \forall t\in\mathcal{T}, \\
& \epsilon_v^t \geq \frac{1}{|\mathcal{O}^t|}\left(\theta^t - \sum_{k\in\mathcal{K}}\sum_{r\in\mathcal{R}^k}\sum_{b\in\mathcal{B}^r} D_{v,b}w_b\right), \quad \forall v\in\mathcal{O}^t,\ \forall t\in\mathcal{T}, \\
& \epsilon_v^t \geq 0, \quad \forall v\in\mathcal{O}^t, \forall t\in\mathcal{T}, \\
& u^t = \sum_{v\in\mathcal{O}^t}\mu_v^t, \quad \forall t\in\mathcal{T}, \\
& \mu_v^t \geq \frac{1}{|\mathcal{O}^t|}\left(\sum_{k\in\mathcal{K}}\sum_{r\in\mathcal{R}^k}\sum_{b\in\mathcal{B}^r} D_{v,b}w_b - \theta^t\right), \quad \forall v\in\mathcal{O}^t,\ \forall t\in\mathcal{T}, \\
& \mu_v^t \geq 0, \quad \forall v\in\mathcal{O}^t,\ \forall t\in\mathcal{T}, \\
& \lambda_b \geq w_b - w_{b'}, \quad \forall b\in\mathcal{B}^r,\ \forall r\in\mathcal{R}^k,\ \forall k\in\mathcal{K}, \\
& m^k \geq \sum_{b\in\mathcal{B}^r}\lambda_b, \quad \forall r\in\mathcal{R}^k,\ \forall k\in\mathcal{K}, \\
& w_b \geq 0,\ \lambda_b \geq 0, \quad \forall b\in\mathcal{B}^r,\ \forall r\in\mathcal{R}^k,\ \forall k\in\mathcal{K}.
\end{aligned}
\tag{A.12}
$$

A.2.2 Dual of inverse planning problem

The following model is the complete linear program corresponding to formulation (2.11) and is the dual of the IPP (A.12):

$$
\begin{aligned}
\underset{q,\pi,n,\xi,s,n,\sigma,\delta}{\text{maximize}} \quad & \sum_{t\in\mathcal{T}}\sum_{v\in\mathcal{O}^t}\theta^t\frac{q_v^t-\pi_v^t}{|\mathcal{O}^t|}-\sum_{i\in\mathcal{I}}\sum_{f\in\mathcal{F}^i}\sum_{v\in\mathcal{O}^i}\frac{n_v^f f}{|\mathcal{O}^i|} \\
\text{subject to} \quad & \sum_{i\in\mathcal{I}}\sum_{v\in\mathcal{O}^i}D_{v,b}\left(\xi^i+s_v^i+\sum_{f\in\mathcal{F}^i}\frac{n_v^f}{|\mathcal{O}^i|}\right)+\sum_{t\in\mathcal{T}}\sum_{v\in\mathcal{O}^t}D_{v,b}\left(\frac{\pi_v^t-q_v^t}{|\mathcal{O}^t|}+s_v^t\right) \\
& \qquad +\sigma_b-\sigma_{b'}\geq 0 \quad \forall b\in\mathcal{B}^r,\ \forall r\in\mathcal{R}^k,\ \forall k\in\mathcal{K}, \\
& \gamma^i=\xi^i, \quad \forall i\in\mathcal{I}, \\
& \beta^i=\sum_{v\in\mathcal{O}^i}s_v^i, \quad \forall i\in\mathcal{I}, \\
& \kappa^f\geq n_v^f, \quad \forall v\in\mathcal{O}^i,\ \forall f\in\mathcal{F}^i,\ \forall i\in\mathcal{I}, \\
& \beta^t=\sum_{v\in\mathcal{O}^t}s_v^t, \quad \forall v\in\mathcal{O}^t,\ \forall t\in\mathcal{T}, \\
& \psi^t\geq\pi_v^t, \quad \forall v\in\mathcal{O}^t,\ \forall t\in\mathcal{T}, \\
& \phi^t\geq q_v^t, \quad \forall v\in\mathcal{O}^t,\ \forall t\in\mathcal{T}, \\
& \delta^r\geq\sigma_b, \quad \forall b\in\mathcal{B}^r,\ \forall r\in\mathcal{R}^k,\ \forall k\in\mathcal{K}, \\
& 1=\sum_{r\in\mathcal{R}^k}\delta^r, \quad \forall\ k\in\mathcal{K}, \\
& \xi^i\geq 0, \quad \forall i\in\mathcal{I}, \\
& s_v^i\geq 0, \quad \forall v\in\mathcal{O}^i,\forall i\in\mathcal{I}, \\
& n_v^f\geq 0, \quad \forall v\in\mathcal{O}^i,\forall f\in\mathcal{F}^i,\forall i\in\mathcal{I}, \\
& s_v^t\geq 0, \quad \forall v\in\mathcal{O}^t,\forall t\in\mathcal{T}, \\
& \pi_v^t\geq 0, \quad \forall v\in\mathcal{O}^t,\forall t\in\mathcal{T}, \\
& q_v^t\geq 0, \quad \forall v\in\mathcal{O}^t,\forall t\in\mathcal{T}, \\
& \sigma_b\geq 0, \quad \forall b\in\mathcal{B}^r,\ \forall r\in\mathcal{R}^k,\ \forall k\in\mathcal{K}, \\
& \delta^r\geq 0, \quad \forall r\in\mathcal{R}^k, \quad \forall k\in\mathcal{K}.
\end{aligned}
\tag{A.13}
$$

A.2.3 Inverse optimization model

The following model is the complete linear program corresponding to formulation (2.13) and is the IO model corresponding to the IPP (A.12):

$$
\begin{aligned}
\underset{\substack{\gamma,\beta,\kappa,\phi,\psi,\\ q,\pi,n,\xi,s,n,\sigma,\delta}}{\text{minimize}} \quad & \left(\sum_{i\in\mathcal{I}}\left(\gamma^i\hat{z}^i+\beta^i\hat{y}^i+\sum_{f\in\mathcal{F}^i}\kappa^f\hat{h}^f\right)+\sum_{t\in\mathcal{T}}(\beta^t\hat{y}^t+\phi^t\hat{l}^t+\psi^t\hat{u}^t)\right)\\
& -\left(\sum_{t\in\mathcal{T}}\sum_{v\in\mathcal{O}^t}\theta^t\frac{q_v^t-r_v^t}{|\mathcal{O}^t|}-\sum_{i\in\mathcal{I}}\sum_{f\in\mathcal{F}^i}\sum_{v\in\mathcal{O}^i}\frac{n_v^f f}{|\mathcal{O}^i|}\right)\\
\text{subject to} \quad & \sum_{i\in\mathcal{I}}\sum_{v\in\mathcal{O}^i}D_{v,b}\left(\xi^i+s_v^i+\sum_{f\in\mathcal{F}^i}\frac{n_v^f}{|\mathcal{O}^i|}\right)+\sum_{t\in\mathcal{T}}\sum_{v\in\mathcal{O}^t}D_{v,b}\left(\frac{\pi_v^t-q_v^t}{|\mathcal{O}^t|}+s_v^t\right)\\
& \qquad +\sigma_b-\sigma_{b'}\geq 0 \quad \forall b\in\mathcal{B}^r,\ \forall r\in\mathcal{R}^k,\ \forall k\in\mathcal{K},\\
& \gamma^i=\xi^i, \quad \forall i\in\mathcal{I},\\
& \beta^i=\sum_{v\in\mathcal{O}^i}s_v^i, \quad \forall i\in\mathcal{I},\\
& \kappa^f\geq n_v^f, \quad \forall v\in\mathcal{O}^i,\ \forall f\in\mathcal{F}^i,\ \forall i\in\mathcal{I},\\
& \beta^t=\sum_{v\in\mathcal{O}^t}s_v^t, \quad \forall v\in\mathcal{O}^t,\ \forall t\in\mathcal{T},\\
& \psi^t\geq\pi_v^t, \quad \forall v\in\mathcal{O}^t,\ \forall t\in\mathcal{T},\\
& \phi^t\geq q_v^t, \quad \forall v\in\mathcal{O}^t,\ \forall t\in\mathcal{T},\\
& \delta^r\geq\sigma_b, \quad \forall b\in\mathcal{B}^r,\ \forall r\in\mathcal{R}^k,\ \forall k\in\mathcal{K},\\
& 1=\sum_{r\in\mathcal{R}^k}\delta^r, \quad \forall\ k\in\mathcal{K},\\
& \xi^i\geq 0, \quad \forall i\in\mathcal{I},\\
& s_v^i\geq 0, \quad \forall v\in\mathcal{O}^i,\forall i\in\mathcal{I},\\
& n_v^f\geq 0, \quad \forall v\in\mathcal{O}^i,\forall f\in\mathcal{F}^i,\forall i\in\mathcal{I},\\
& s_v^t\geq 0, \quad \forall v\in\mathcal{O}^t,\forall t\in\mathcal{T},\\
& \pi_v^t\geq 0,\ q_v^t\geq 0, \quad \forall v\in\mathcal{O}^t,\forall t\in\mathcal{T},\\
& \sigma_b\geq 0, \quad \forall b\in\mathcal{B}^r,\ \forall r\in\mathcal{R}^k,\ \forall k\in\mathcal{K},\\
& \delta^r\geq 0, \quad \forall r\in\mathcal{R}^k, \quad \forall k\in\mathcal{K}.
\end{aligned}
\tag{A.14}
$$