Operations Research for Health Care ▮ (▮▮▮▮) ▮▮▮–▮▮▮

Contents lists available at ScienceDirect

Operations Research for Health Care

journal homepage: www.elsevier.com/locate/orhc

Constraint generation methods for robust optimization in radiation therapy

Houra Mahmoudzadeh [a,*], Thomas G. Purdie [b,c,d], Timothy C.Y. Chan [a]

[a] Department of Mechanical and Industrial Engineering, University of Toronto, Toronto, Canada
[b] Radiation Medicine Program, UHN Princess Margaret Cancer Centre, Toronto, Canada
[c] Department of Radiation Oncology, University of Toronto, Toronto, Canada
[d] Techna Institute for the Advancement of Technology for Health, Toronto, Canada

ARTICLE INFO

Article history:
Received 31 October 2014
Accepted 24 March 2015
Available online xxxx

Keywords:
Decomposition
Constraint generation
Robust optimization
Radiation therapy
Breast cancer

ABSTRACT

We develop a constraint generation solution method for robust optimization problems in radiation therapy in which the problems include a large number of robust constraints. Each robust constraint must hold for any realization of an uncertain parameter within a given uncertainty set. Because the problems are large scale, the robust counterpart is computationally challenging to solve. To address this challenge, we explore different strategies of adding constraints in a constraint generation solution approach. We motivate and demonstrate our approach using robust intensity-modulated radiation therapy treatment planning for breast cancer. We use clinical data to compare the computational efficiency of our constraint generation strategies with that of directly solving the robust counterpart.

1. Introduction

Robust Optimization (RO) deals with optimization problems in which some problem parameters are uncertain and modeled as belonging to an "uncertainty set" [1,2]. One of the areas in which robust optimization has been applied is radiation therapy (RT) treatment planning. In RT, the goal is to deliver radiation beams from different angles to a cancer patient so that the beams intersect at the cancerous target (i.e., a tumor), while sparing as much of the surrounding healthy tissue as possible. Robust optimization has been used to manage uncertainties in RT treatment planning problems including uncertainties in patient geometry [3], dose calculations [4], breathing motion [5–10], and range and setup errors in proton therapy [11–14].

Much effort in RO is placed on deriving tractable robust counterparts, which are finite-sized deterministic equivalents to the original RO problem. However, the resulting robust counterpart can still be quite large and computationally challenging to solve for real-world problem instances. In radiation therapy treatment planning for example, the original problem is often of a very large scale and the robust counterpart is even larger. Therefore, there is a need for specialized solution methods to solve these problems.

Decomposition methods, which have a long history [15,16], represent a wide range of methods that can be used to solve large-scale optimization problems. Constraint generation is one type of decomposition method that has been used extensively to solve large-scale optimization problems in applications such as timetable scheduling [17], network reliability [18], network design [19], facility location [20], and network interdiction [21,22]. Oskoorouchi et al. [23] developed an interior point constraint generation algorithm for semi-infinite problems that was applied to radiation therapy.

In this paper, we develop a family of constraint generation strategies to solve large-scale robust optimization problems in radiation therapy. We focus on problems with multiple sets of robust constraints, which necessitates exploring different strategies for choosing constraints to be added at each iteration. We test several strategies for finding and adding constraints efficiently. We also compare the computational efficiency of the constraint generation methods with that of directly solving the robust counterpart. Our solution approach is motivated by the robust intensity-modulated radiation therapy (IMRT) treatment planning problem for breast cancer, in which there exists a large number of robust constraints [8].

2. Breast cancer IMRT treatment planning

Previously, a robust optimization model that incorporated conditional value-at-risk (CVaR) [24,25] was developed for breast

* Corresponding author. Tel.: +1 416 978 4585; fax: +1 416 978 7753.
E-mail address: houra@mie.utoronto.ca (H. Mahmoudzadeh).

http://dx.doi.org/10.1016/j.orhc.2015.03.003
2211-6923/

cancer IMRT to control the tail dose to the tumor and organs-at-risk under breathing motion uncertainty [8]. In this section, we first introduce this problem and then briefly describe the optimization model in order to motivate the development of our constraint generation strategies.

2.1. Background

In IMRT, the radiation beams can be modeled as a collection of small *beamlets* whose intensities are optimized. An optimization problem is to find the intensity of each beamlet such that sufficient dose is delivered to the tumor while minimizing the dose to the surrounding healthy tissue. Our approach to breast cancer IMRT follows the clinical protocol at the Princess Margaret Cancer Centre [26,27]. There are two opposed beams that are tangent to the body and deliver radiation to a target volume in the breast tissue. In left-sided breast cancer, parts of the left lung and the heart are usually inside the treatment field and are considered to be organs-at-risk (OAR). Since the radiation is delivered to the patient while the patient is breathing, the organs move and deform throughout the course of the treatment. In particular, the heart may move inside the treatment field and become exposed to excessive radiation. Four-dimensional computed tomography (4D-CT) images are used to obtain geometrical information about the organs over the phases of the patient's breathing cycle from inhale to exhale. The uncertain parameter in this problem is the patient's breathing pattern, which is modeled as a probability mass function (PMF) that captures the fraction of time that the patient spends in each breathing phase. We construct the uncertainty set by including upper and lower error bounds on a nominal breathing pattern[8].

2.2. A robust-CVaR optimization model for breast cancer IMRT

Here, we briefly introduce the optimization model from Chan et al. [8], which proposes the robust optimization model that we consider. Let w_b be the intensity of beamlet $b \in \mathcal{B}$, where $\mathcal{B}$ is the set of all beamlets. The body is discretized into small volumetric pixels called "voxels". Let $\mathcal{V}^T$ and $\mathcal{V}^H$ be the set of all voxels in the clinical target volume (inside the breast) and the heart, respectively. A breathing PMF is defined over the set of breathing phases $\mathcal{I}$.

The total dose to each voxel is the sum of the dose accumulated over all breathing phases and depends on the uncertain fraction of the time spent at each phase. Let $\Delta_{v,i,b}$ be the *influence matrix*, which quantifies the amount of dose that voxel v receives when the patient is in phase i per unit intensity of beamlet b. A robust upper β-CVaR constraint on the target can be formulated as:

$$\overline{\zeta}^T_\beta + \frac{1}{1-\beta}\sum_{v\in\mathcal{V}^T}\max\left\{0, \sum_{i\in\mathcal{I}}\sum_{b\in\mathcal{B}}\Delta_{v,i,b}\tilde{p}(i)w_b - \overline{\zeta}^T_\beta\right\} \leq U^T_\beta, \quad \forall \tilde{\mathbf{p}} \in \mathcal{P}. \tag{1}$$

The variable ζ is the value-at-risk (VaR) of the dose distribution, which captures the β percent of an organ that is receiving the highest amount of dose. Parameter U^T_β is the upper bound on the conditional-value-at-risk (CVaR) which is the average β% of the tail of the dose distribution. Constraint (1) is a β-CVaR constraint and must be met for all breathing patterns $\tilde{\mathbf{p}}$ in the given uncertainty set $\mathcal{P}$, which is a polyhedral set. Lower CVaR constraints can be formulated similarly. Model (2) shows the robust-CVaR IMRT model with one set of upper and lower β-CVaR constraints for limiting overdose and underdose to the target.

$$\text{minimize} \quad \frac{1}{|\mathcal{V}^H|}\sum_{v\in\mathcal{V}^H}\sum_{i\in\mathcal{I}}\sum_{b\in\mathcal{B}}\Delta_{v,i,b}p(i)w_b \tag{2a}$$

$$\text{subject to} \quad \overline{\zeta}^T_\beta + \frac{1}{(1-\beta)|\mathcal{V}^T|}\sum_{v\in\mathcal{V}^T}\overline{d}^T_{v,\beta} \leq U^T_\beta, \quad \forall\beta\in\overline{A}^T, \tag{2b}$$

$$\overline{d}^T_{v,\beta} \geq \sum_{i\in\mathcal{I}}\sum_{b\in\mathcal{B}}\Delta_{v,i,b}\tilde{p}(i)w_b - \overline{\zeta}^T_\beta, \quad \forall v\in\mathcal{V}^T,\ \beta\in\overline{A}^T,\ \tilde{\mathbf{p}}\in\mathcal{P}, \tag{2c}$$

$$\underline{\zeta}^T_\beta - \frac{1}{(1-\beta)|\mathcal{V}^T|}\sum_{v\in\mathcal{V}^T}\underline{d}^T_{v,\beta} \geq L^T_\beta, \quad \forall\beta\in\underline{A}^T, \tag{2d}$$

$$\underline{d}^T_{v,\beta} \geq \underline{\zeta}^T_\beta - \sum_{i\in\mathcal{I}}\sum_{b\in\mathcal{B}}\Delta_{v,i,b}\tilde{p}(i)w_b, \quad \forall v\in\mathcal{V}^T,\ \beta\in\underline{A}^T,\ \tilde{\mathbf{p}}\in\mathcal{P}, \tag{2e}$$

$$\overline{\zeta}^T_\beta \geq 0, \quad \forall\beta\in\overline{A}^T, \tag{2f}$$

$$\underline{\zeta}^T_\beta \geq 0, \quad \forall\beta\in\underline{A}^T, \tag{2g}$$

$$\overline{d}^T_{v,\beta} \geq 0, \quad \forall v\in\mathcal{V}^T,\ \beta\in\overline{A}^T, \tag{2h}$$

$$\underline{d}^T_{v,\beta} \geq 0, \quad \forall v\in\mathcal{V}^T,\ \beta\in\underline{A}^T, \tag{2i}$$

$$w_b \geq 0, \quad \forall b\in\mathcal{B}. \tag{2j}$$

In a robust IMRT problem, there may exist multiple CVaR constraints on an organ for different values of β. Notice that constraint (2c) and (2e) must hold for all voxels in the tumor. We define a *type* of robust constraint as one that must hold for the same set of voxels for the same β value and in the same inequality direction (i.e., upper or lower CVaR constraints) for all $\tilde{\mathbf{p}} \in \mathcal{P}$. For example, all the upper β-CVaR constraints in (2c) are of the same type, although they are separate constraints for each voxel on the tumor. On the other hand, the sets of robust constraints (2c) and (2e) are of different types. Similarly, constraints for different β values or for different organs would be of different types.

Because the original problem is linear and the uncertainty set is polyhedral, the robust counterpart of this problem is linear. However, the large number of robust constraints make the robust counterpart very large. Alternatively, because the uncertainty set is polyhedral, an equivalent reformulation exists by simply enumerating the vertices of $\mathcal{P}$. In general though, enumerating all the vertices of a polyhedron is NP-hard [28] and could lead to an exponential number of constraints. Thus, we consider constraint generation as an alternative solution method.

3. A constraint generation solution method

In this section, we first formulate a general form of the previous RO problem with multiple uncertain constraints. Then, we define the steps of the constraint generation algorithm and develop several constraint addition strategies.

3.1. A robust optimization problem with uncertain constraints

Consider a robust optimization problem with uncertain constraints of different types $k \in \mathcal{K}$ that must hold for every $v \in \mathcal{V}(k)$. Let $\mathbf{w}$ be the decision vector. The uncertainty is in the vector $\tilde{\mathbf{p}} \in \mathcal{P}$, which is a parameter that affects all robust constraints. In other words, all robust constraints must hold for all values of $\tilde{\mathbf{p}}$ in the uncertainty set $\mathcal{P}$. Let $\mathbf{c}$ be the vector of objective function coefficients and $\mathbf{A}_{v,k}$ be the constraint coefficient matrix for constraint type $k \in \mathcal{K}$ for every $v \in \mathcal{V}(k)$. Given that $|\mathcal{I}|$ and $|\mathcal{B}|$ are the sizes of the vectors $\tilde{\mathbf{p}}$ and $\mathbf{w}$, respectively, the size of the matrix $\mathbf{A}_{v,k}$ is $|\mathcal{I}| \times |\mathcal{B}|$ for each v, k. The parameter $a_{v,k}$ is the right hand side

and corresponds to the upper or lower bound on the dose to be delivered. The general robust optimization problem is

$$\underset{\mathbf{w}}{\text{minimize}} \quad \mathbf{c}'\mathbf{w} \tag{3a}$$

$$\text{subject to} \quad \tilde{\mathbf{p}}'\mathbf{A}_{v,k}\mathbf{w} \geq a_{v,k}, \quad \forall v \in \mathcal{V}(k),\ k \in \mathcal{K},\ \tilde{\mathbf{p}} \in \mathcal{P}, \tag{3b}$$

$$\mathbf{w} \geq \mathbf{0}. \tag{3c}$$

For simplicity, we omit constraints that do not depend on $\mathcal{P}$, which may also be present in the problem.

3.2. Problem decomposition

We decompose (3) into a master problem and a subproblem. The master problem is a linear program. An optimal solution to the master problem is passed to the subproblem, which identifies constraints to be added to the master problem or provides a certificate of optimality for the original problem.

3.2.1. The master problem

Formulation (4) shows a general form of the master problem.

$$\underset{\mathbf{w}}{\text{minimize}} \quad \mathbf{c}'\mathbf{w} \tag{4a}$$

$$\text{subject to} \quad \tilde{\mathbf{p}}'_{v,k}\mathbf{A}_{v,k}\mathbf{w} \geq a_{v,k}, \quad \tilde{\mathbf{p}}_{v,k} \in \mathcal{P}^n_{v,k}, \quad \forall v \in \mathcal{V}(k),\ k \in \mathcal{K}, \tag{4b}$$

$$\mathbf{w} \geq \mathbf{0}.$$

The initial set $\mathcal{P}^1_{v,k}$ consists of a single $\mathbf{p} \in \mathcal{P}$ for all $v \in \mathcal{V}(k), k \in \mathcal{K}$. At every iteration n, a finite number of constraints are added to the master problem, updating the set $\mathcal{P}^n_{v,k}$. Thus, the master problem is a linear program. An optimal solution $\mathbf{w}^*$ is passed to the subproblem.

3.2.2. The subproblem

The subproblem (5) finds a distinct vector $\mathbf{p}_{v,k} \in \mathcal{P}$ for each v, k for the corresponding constraint in (4b) to maximize the violation of the constraints.

$$\underset{\{\mathbf{p}_{v,k}: v \in \mathcal{V}(k),\ k \in \mathcal{K}\}}{\text{maximize}} \quad \sum_{v \in \mathcal{V}(k)} \sum_{k \in \mathcal{K}} \text{viol}_{v,k} \tag{5a}$$

$$\text{subject to} \quad \text{viol}_{v,k} = \max\left\{0, a_{v,k} - \mathbf{p}'_{v,k}\mathbf{A}_{v,k}\mathbf{w}^*\right\}, \quad \forall v \in \mathcal{V}(k),\ k \in \mathcal{K}, \tag{5b}$$

$$\mathbf{p}_{v,k} \in \mathcal{P}, \quad \forall v \in \mathcal{V}(k),\ k \in \mathcal{K}. \tag{5c}$$

If the optimal value of the subproblem is 0, then the current solution to the master problem is optimal. The optimal solution $\mathbf{p}^*_{v,k}$ generates the largest violation for the constraint index v of type k. In other words, $p^*_{v,k}$ is the breathing pattern that generates the worst-case violation of dose metric type k for voxel v. Note that the subproblem is separable in v, k. In general, the complexity of the subproblem depends on the uncertainty set $\mathcal{P}$. Because we consider a polyhedral uncertainty set in this paper, the subproblem is linear.

3.2.3. Constraint addition strategies

Our robust model (3) has $|\mathcal{K}|$ sets of robust constraints, each constituting a set of constraints for all $v \in \mathcal{V}(k)$. An optimal solution to the subproblem may identify numerous $\mathbf{p}^*_{v,k}$ that result in a positive violation of its corresponding constraint. This leads to several natural questions about adding constraints to the master problem:

- How many constraints should be added to the master problem at every iteration?
- Since different $\mathbf{p}^*_{v,k}$ vectors may generate worst case violations for different v and k, which one(s) should be added to the master problem?
- Is it efficient to add a PMF, $\mathbf{p}$, for all v of the same constraint type k? In other words, is there a PMF that generates a large aggregate violation over many $v \in \mathcal{V}(k)$?

Let $\mathbf{p}^*_{v,k}$ generate the maximum violation for constraint index v in constraint type k. Let $\mathbf{p}^*$ be the PMF that generates the highest violation among all constraints, which corresponds to a particular index v^* and constraint type k^*. In other words, $\mathbf{p}^* = \mathbf{p}^*_{v^*,k^*}$ where $(v^*, k^*) = \underset{v,k}{\text{argmax}}(\text{viol}^*_{v,k})$. Some of the constraint addition strategies we explore will add the same PMF, $\mathbf{p}^*$, for all voxels and other strategies add different PMFs, $\mathbf{p}^*_{v,k}$, for each voxel. First assume we have only one type of robust constraint ($|\mathcal{K}| = 1$) for all $v \in \mathcal{V}$, and we need to choose constraint index v for which we need to add constraints to the master problem. We call these *index-based* strategies. These strategies are as follows:

S1. Add $\mathbf{p}^*$ for all $v \in \mathcal{V}(k)$.
S2. Add $\mathbf{p}^*$ for all $v \in \mathcal{V}(k)$ for which the maximum violation is greater than zero.
S3. Add $\mathbf{p}^*_{v,k}$ for all $v \in \mathcal{V}(k)$ for which the maximum violation is greater than zero.
S4. Add $\mathbf{p}^*$ for all $v \in \mathcal{V}(k)$ for which the maximum violation is greater than some threshold $\delta > 0$.
S5. Add $\mathbf{p}^*_{v,k}$ for all $v \in \mathcal{V}(k)$ for which the maximum violation is greater than some threshold $\delta > 0$.
S6. Add $\mathbf{p}^*$ for only v^*.

The strategies are ordered based on the number of constraints that they add at each iteration. Strategies S2 and S3 add the same number of constraints but for different $\mathbf{p}$ vectors in the first iteration. The same is true for strategies S4 and S5.

Now consider the case where we have more than one type of constraint. All the index-based strategies can be implemented either for all constraint types or only for constraint type k^*. We call this second decision the *type-based* strategy. Thus, in total, we have twelve strategies to test, which are summarized in Table 1.

3.2.4. Solution update

After updating the master problem with new constraints, standard dual simplex iterations are used to find a new optimal solution. The solution update time depends on the number of constraints added at each iteration, which depends on the constraint addition strategy employed. The last column in Table 1 shows the number of constraints that are added in each iteration for each strategy.

3.2.5. Stopping criterion

The algorithm terminates when the maximum violation of all constraints is less than a given tolerance $\epsilon > 0$. In Section 4, we explore the sensitivity of the results to different values of ϵ.

3.2.6. Overview of the solution method

Algorithm 1 shows an overview of the solution method.

4. Results

We applied the constraint generation approach with all the different addition strategies on the breast cancer IMRT treatment planning model (2). A clinical patient dataset was provided by the Princess Margaret Cancer Centre, Toronto, Canada. There were 6824 voxels in the target volume and 13 249 voxels in the heart. The beam included 901 beamlets. The problem included upper and lower robust-CVaR constraints on the target, with $\beta = 0.5\%$,

Please cite this article in press as: H. Mahmoudzadeh, et al., Constraint generation methods for robust optimization in radiation therapy, Operations Research for Health Care (2015), http://dx.doi.org/10.1016/j.orhc.2015.03.003

Table 1
Strategies for adding constraints.

Strategy name	Added $\mathbf{p}$	Index-based	Type-based	Constraints per it.
S1-1	$\mathbf{p}^*$	$\forall v \in \mathcal{V}(k)$	$k = k^*$	$\lvert\mathcal{V}(k^*)\rvert$
S1-2			$\forall k$	$\sum_{k\in\mathcal{K}} \lvert\mathcal{V}(k)\rvert$
S2-1	$\mathbf{p}^*$	$\forall v \in \mathcal{V}(k) : \lvert\text{viol}_{v,k}\rvert > 0$	$k = k^*$	$\leq \lvert\mathcal{V}(k^*)\rvert$
S2-2			$\forall k$	$\leq \sum_{k\in\mathcal{K}} \lvert\mathcal{V}(k)\rvert$
S3-1	$\mathbf{p}^*_{v,k}$	$\forall v \in \mathcal{V}(k) : \lvert\text{viol}_{v,k}\rvert > 0$	$k = k^*$	$\leq \lvert\mathcal{V}(k^*)\rvert$
S3-2			$\forall k$	$\leq \sum_{k\in\mathcal{K}} \lvert\mathcal{V}(k)\rvert$
S4-1	$\mathbf{p}^*$	$\forall v \in \mathcal{V}(k) : \lvert\text{viol}_{v,k}\rvert > \delta$	$k = k^*$	$\leq \lvert\mathcal{V}(k^*)\rvert$
S4-2			$\forall k$	$\leq \sum_{k\in\mathcal{K}} \lvert\mathcal{V}(k)\rvert$
S5-1	$\mathbf{p}^*_{v,k}$	$\forall v \in \mathcal{V}(k) : \lvert\text{viol}_{v,k}\rvert > \delta$	$k = k^*$	$\leq \lvert\mathcal{V}(k^*)\rvert$
S5-2			$\forall k$	$\leq \sum_{k\in\mathcal{K}} \lvert\mathcal{V}(k)\rvert$
S6-1	$\mathbf{p}^*$	$v = v^*$	$k = k^*$	1
S6-2			$\forall k$	$\lvert\mathcal{K}\rvert$

Algorithm 1

1: $n = 1$. Let $\mathcal{P}^1_{v,k}$ be a single element from $\mathcal{P}$, for all $v \in \mathcal{V}(k)$, $k \in \mathcal{K}$.
2: Solve master problem (4) with $\mathcal{P}^n_{v,k}$ and pass optimal solution $\mathbf{w}^n$ to the subproblem.
3: Solve the subproblem (5) and find optimal $\mathbf{p}^n_{v,k}$ for all $v \in \mathcal{V}(k)$, $k \in \mathcal{K}$.
4: **if** the optimal value of the subproblem is positive **and** the stopping criterion is not met **then**
5: go to 9.
6: **else**
7: go to 10.
8: **end if**
9: Add new constraints to the set $\mathcal{P}^n_{v,k}$ to construct $\mathcal{P}^{n+1}_{v,k}$, increment n, and go to 2.
10: Output $\mathbf{w}^n$ and stop.

$U = 45.79$ and $\beta = 95\%$, $L = 39.01$, respectively. The parameter δ was set to 0.1 Gy which is considered a clinically negligible dose. The objective was to minimize the mean dose to the heart.

There were 13,648 ($\sum_{k\in\mathcal{K}} |\mathcal{V}(k)|$) robust constraints in total. The constraint generation algorithm was coded using C++ and CPLEX 12.1 was used to solve the optimization problems. The robust counterpart was solved using both C++/CPLEX and AMPL/CPLEX for comparison. All trials were run using a single Linux node of a Dell PowerEdge R410 computer with a 3.07 GHz 12-core CPU and 32 GB of RAM.

4.1. Computational efficiency

Table 2 shows the total run times for all constraint generation strategies and the robust counterpart. Strategy S3 tailors the added vector $\mathbf{p}$ to each violated constraint indexed by v, which results in a faster solution time than strategy S2, which adds the same vector $\mathbf{p}$ to all violated constraints. Both strategies S2 and S3 start with adding the same number of constraints in the first iteration, but strategy S3 takes fewer iterations than strategy S2, since the constraints it adds are more effective. The same trend is seen when comparing strategies S5 and S4. Overall, tailoring the added vector $\mathbf{p}$ to each violated constraint produces the fastest solution times. S6 did not find an optimal solution within the time limit provided, so we did not further examine this strategy. The robust counterparts were an order of magnitude slower than the fastest constraint generation approach. Table 1 also suggests tailoring the added constraint to the particular type of robust constraint that exhibited a violation can speed up the computation. This was true for S1 and S3. AMPL's pre-processing capability was the reason AMPL's robust counterpart solved more quickly than the C++ equivalent.

Table 3 shows the breakdown between master and subproblem computation times for each constraint generation strategy.

Table 2
Computation times (min).

Strategy	Constraint generation						Robust counterpart	
	S1	S2	S3	S4	S5	S6	AMPL	C++
$k = k^*$	170	35	12	38	13	>1000	180	322
$\forall k$	195	30	19	38	13	>1000		

Table 3
Detailed computation times (min) and the number of iterations for each strategy. The total time includes post-optimization processing between the iterations of the algorithm.

Strategy	Subproblem	Master problem	Total time	# of iterations
S1-1	0.29	156	170	40
S1-2	0.18	184	195	24
S2-1	0.37	19	35	51
S2-2	0.27	17	30	37
S3-1	0.05	6	12	7
S3-2	0.06	12	19	6
S4-1	0.54	17	38	73
S4-2	0.50	18	38	69
S5-1	0.07	7	13	10
S5-2	0.04	7	13	5

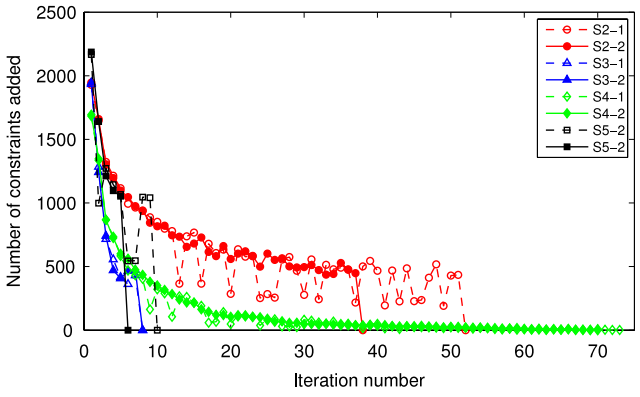

Fig. 1. The number of constraints added at each iteration.

For all strategies, the time to solve the initial master problem was the same (4.2 min). The subproblem column reports the total time spent solving the subproblem over all iterations. Similarly, the time spent for updating the master problem over all iterations is shown.

Fig. 1 shows the exact number of constraints that were added per iteration for strategies S2–S5. Strategies S1 and S6 add a fixed number of constraints and are not shown. The strategies that add $\mathbf{p}$ vectors for all types of constraints ($\forall k$) have a smoother trend compared to those that add them for only one type ($k = k^*$). The reason is that different types of constraints will most likely have different worst-case $\mathbf{p}$ vectors at each iteration.

Please cite this article in press as: H. Mahmoudzadeh, et al., Constraint generation methods for robust optimization in radiation therapy, Operations Research for Health Care (2015), http://dx.doi.org/10.1016/j.orhc.2015.03.003

Table 4
The effect of using different values of ϵ. The last two columns show the percentage of voxels that have a violation for some constraint and the mean of positive violations for all voxels, for the upper bound (U) and lower bound (L) constraints.

ϵ	# it.	Run time (min)	Maximum violation (Gy)	Percentage of voxels		Mean of viols >0	
				U	L	U	L
1	3	10.07	0.90	5.29	4.60	0.79	0.90
0.8	5	11.05	0.57	3.53	3.50	8.8×10^{-3}	0.17
0.5	6	11.50	0.37	2.29	3.09	3.5×10^{-3}	4.7×10^{-3}
0.1	7	11.83	0.031	3.33	3.33	4.4×10^{-5}	4.4×10^{-5}
0.03	8	12.14	0.026	2.62	2.78	1.5×10^{-4}	3.0×10^{-7}
0.01	9	12.39	1.8×10^{-11}	2.21	3.18	1.8×10^{-13}	2.6×10^{-13}
10^{-10}	9	12.39	1.8×10^{-11}	2.21	3.18	1.8×10^{-13}	2.6×10^{-13}

Fig. 2 compares the solution times (in min) of the strategies versus the average number of constraints that are added per iteration. It shows that adding a moderate number of constraints in each iteration generally results in the shortest computation time. Strategies S1 and S6 (omitted from the figure) exhibit the largest computation time and add the largest and the smallest number of constraints in each iteration, respectively.

4.2. Solution quality vs. computation time

The results presented in Section 4 are based on $\epsilon = 0.1$. In this section, we compare the solution quality and computation time for different values of ϵ using strategy S3-1, which is the fastest strategy. To quantify the solution quality, we calculate the maximum violation (in units of Gy, which is a measure of radiation dose) of all constraints in the final solution.

Table 4 compares the results using different values of ϵ. It shows the number of iterations, total run time and the maximum violation from any of the constraints in units of Gray (Gy). The last two columns show the percentage of voxels that have a violation for some constraint and the mean of positive violations for all voxels, for the upper bound (U) and lower bound (L) constraints. It can be seen that as ϵ is varied, the solution time is minimally affected with a changing ϵ and the number and magnitude of constraint violations is well-controlled. Fig. 3 shows a dose–volume histogram (DVH) which provides a more clinical view of the solution quality. A DVH shows the fraction of an organ that receives a certain dose or higher. As Fig. 3 illustrates, the solutions corresponding to $\epsilon = 1$ and $\epsilon = 10^{-10}$ result in treatments with essentially identical dosimetric properties.

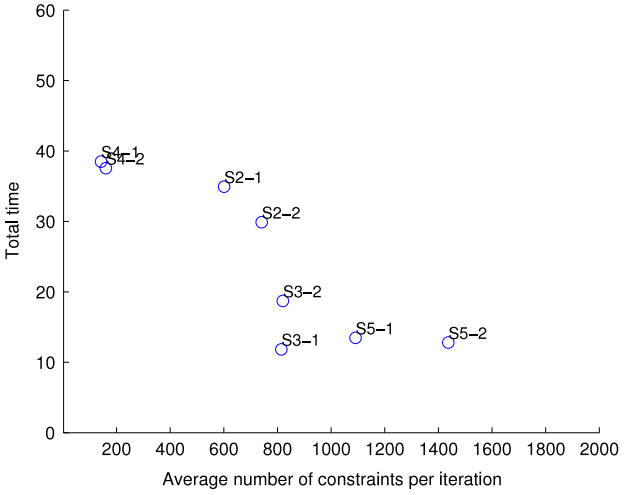

Fig. 2. Comparing the average number of constraints per iteration and total time.

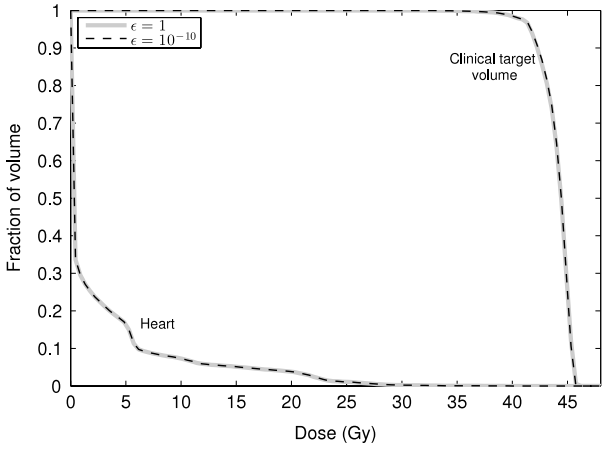

Fig. 3. The effect of ϵ on solution quality.

5. Discussion

The fact that strategy S3-1 was the fastest constraint addition strategy reinforces the need to tackle individual voxels separately as each may have different maximum violations for different PMFs. It seems that adding a moderate number of constraints at each iteration resulted in the best performance. The natural strategy of adding one constraint at a time (S6) turned out to be the most inefficient, due to the large number of iterations that were required to converge. At the other end of the spectrum, adding the same vector $\mathbf{p}$ for each k and v also resulted in long computation times.

For most strategies, adding the new constraint(s) for the same type $k = k^*$ results in a shorter computation time than adding them for all k. This confirms that identifying different types of constraints can help reduce the computation time. In the IMRT case, for example, inhale- and exhale-weighted breathing patterns can have different effects on upper and lower CVaR constraints on the dose to each organ.

Adding a violation threshold $\delta > 0$ (used in S4 and S5) resulted in a slightly higher computation time in most of the cases, due to an increase in the number of iterations and corresponding decrease in the number of constraints added per iteration. Our hypothesis was that some constraints added earlier in the algorithm might obviate the need for additional constraints later, because those corresponding infeasibilities would have been addressed by the earlier constraints. This turned out not to be the case.

Finally, we solved the robust problem by enumerating all vertices of the uncertainty set (which was possible here since it was only five dimensional) as another comparison. This version of the problem took about 30 min to solve, which is faster than the robust counterpart, but still much slower than the best constraint generation strategy.

6. Conclusions and future work

In this paper, we developed a decomposition-based solution method for robust optimization problems with a large number

Please cite this article in press as: H. Mahmoudzadeh, et al., Constraint generation methods for robust optimization in radiation therapy, Operations Research for Health Care (2015), http://dx.doi.org/10.1016/j.orhc.2015.03.003

of uncertain constraints. We developed strategies for finding and adding constraints at each iteration. We defined types of constraints and demonstrated the computational benefit of categorizing the constraints based on their type. We applied our method to a large-scale robust IMRT optimization model for breast cancer and compared the computation time of our solution method with that of solving the robust counterpart. Our results showed that constraint generation can typically save one order of magnitude in computation time.

We believe that using geometric information about the voxels (i.e., which voxels are in close proximity to each other) may result in further reductions in computation time for the breast cancer IMRT problem. Also, clustering the breathing patterns based on similarity and prioritizing the addition of different breathing patterns may also help improve computational efficiency. These are topics for future study.

Acknowledgments

The authors would like to thank the EURO Summer Institute on Operational Research Applied to Health in a Modern World 2014 (ESI XXXI), in which a first draft of this paper was presented and discussed. Computational infrastructure used in this paper was provided by the High Performance Computing Virtual Laboratory (HPCVL). This research was supported in part by an Early Researcher Award ER11-08-058 from the Government of Ontario, and the Canadian Breast Cancer Foundation - Ontario Region.

References

[1] A. Ben-Tal, L. El Ghaoui, A. Nemirovski, Robust Optimization, Princeton University Press, 2009.
[2] D. Bertsimas, D.B. Brown, C. Caramanis, Theory and applications of robust optimization, SIAM Rev. 53 (3) (2011) 464–501.
[3] M. Chu, Y. Zinchenko, S.G. Henderson, M.B. Sharpe, Robust optimization for intensity modulated radiation therapy treatment planning under uncertainty, Phys. Med. Biol. 50 (2005) 5463–5477.
[4] A. Ólafsson, S.J. Wright, Efficient schemes for robust IMRT treatment planning, Phys. Med. Biol. 51 (21) (2006) 5621.
[5] T. Bortfeld, T.C.Y. Chan, A. Trofimov, J.N. Tsitsiklis, Robust management of motion uncertainty in intensity-modulated radiation therapy, Oper. Res. 56 (6) (2008) 1461–1473.
[6] T.C.Y. Chan, T. Bortfeld, J.N. Tsitsiklis, A robust approach to IMRT optimization, Phys. Med. Biol. 51 (2006) 2567–2583.
[7] T.C.Y. Chan, V.V. Mišić, Adaptive and robust radiation therapy optimization for lung cancer, European J. Oper. Res. 231 (3) (2013) 745–756.
[8] T.C.Y. Chan, H. Mahmoudzadeh, T.G. Purdie, A robust-CVaR optimization approach with application to breast cancer therapy, European J. Oper. Res. (2014).
[9] V.V. Mišić, T.C.Y. Chan, The perils of adapting to dose errors in radiation therapy, PLoS One (2015) in press.
[10] P.A. Mar, T.C.Y. Chan, Adaptive and robust radiation therapy in the presence of drift, Phys. Med. Biol. 60 (2015) 3599–3615.
[11] J. Unkelbach, T.C. Chan, T. Bortfeld, Accounting for range uncertainties in the optimization of intensity modulated proton therapy, Phys. Med. Biol. 52 (10) (2007) 2755.
[12] W. Liu, X. Zhang, Y. Li, R. Mohan, Robust optimization of intensity modulated proton therapy, Med. Phys. 39 (2) (2012) 1079–1091.
[13] W. Cao, G.J. Lim, A. Lee, Y. Li, W. Liu, X. Ronald Zhu, X. Zhang, Uncertainty incorporated beam angle optimization for IMPT treatment planning, Med. Phys. 39 (8) (2012) 5248–5256.
[14] A. Fredriksson, R. Bokrantz, A critical evaluation of worst case optimization methods for robust intensity-modulated proton therapy planning, Med. Phys. 41 (8) (2014) 081701.
[15] J.F. Benders, Partitioning procedures for solving mixed-variables programming problems, Numer. Math. 4 (1) (1962) 238–252.
[16] G.B. Dantzig, Large-scale system optimization: a review. Tech. Rep., DTIC Document, 1965.
[17] M.A. Odijk, A constraint generation algorithm for the construction of periodic railway timetables, Transp. Res. B 30 (6) (1996) 455–464.
[18] J. Shaio, Constraint generation for network reliability problems, Ann. Oper. Res. 106 (1–4) (2001) 155–180.
[19] A.K. Andreas, J.C. Smith, Decomposition algorithms for the design of a nonsimultaneous capacitated evacuation tree network, Networks 53 (2) (2009) 91–103.
[20] A. Siddiq, Robust facility location under demand location uncertainty (Ph.D. thesis), University of Toronto, 2013.
[21] G. Brown, M. Carlyle, J. Salmerón, K. Wood, Defending critical infrastructure, Interfaces 36 (6) (2006) 530–544.
[22] G.G. Brown, W.M. Carlyle, R.C. Harney, E.M. Skroch, R.K. Wood, Interdicting a nuclear-weapons project, Oper. Res. 57 (4) (2009) 866–877.
[23] M.R. Oskoorouchi, H.R. Ghaffari, T. Terlaky, D.M. Aleman, An interior point constraint generation algorithm for semi-infinite optimization with health-care application, Oper. Res. 59 (5) (2011) 1184–1197.
[24] R. Rockafellar, S. Uryasev, Optimization of conditional value-at-risk, J. Risk 2 (2000) 21–42.
[25] R. Rockafellar, S. Uryasev, Conditional value-at-risk for general loss distributions, J. Banking Finance 26 (7) (2002) 1443–1471.
[26] T.G. Purdie, R.E. Dinniwell, D. Letourneau, C. Hill, M.B. Sharpe, Automated planning of tangential breast intensity modulated radiotherapy using heuristic optimization, Int. J. Radiat. Oncol. Biol. Phys. (2011).
[27] T.G. Purdie, R.E. Dinniwell, A. Fyles, M.B. Sharpe, Automation and intensity modulated radiation therapy for individualized high-quality tangent breast treatment plans, Int. J. Radiat. Oncol. Biol. Phys. (2014).
[28] L. Khachiyan, E. Boros, K. Borys, K. Elbassioni, V. Gurvich, Generating all vertices of a polyhedron is hard, Discrete Comput. Geom. 39 (1–3) (2008) 174–190.

Please cite this article in press as: H. Mahmoudzadeh, et al., Constraint generation methods for robust optimization in radiation therapy, Operations Research for Health Care (2015), http://dx.doi.org/10.1016/j.orhc.2015.03.003