Accelerated robust optimization algorithm for proton therapy treatment planning.

Gregory Buti[1], Kevin Souris[1], Ana M. Barragán Montero[1], Marie Cohilis[1], John A. Lee[1], Edmond Sterpin[1,2]

[1]Université Catholique de Louvain, Institut de Recherche Expérimentale et Clinique (IREC), Center of Molecular Imaging, Radiotherapy and Oncology (MIRO), Brussels, Belgium

[2]Katholieke Universiteit Leuven, Department of Oncology, Laboratory of Experimental Radiotherapy, Leuven, Belgium

Version typeset February 10, 2020

E-mail: gregory.buti@uclouvain.be

Abstract

Purpose: Robust optimization is a computational expensive process resulting in long plan computation times. This issue is especially critical for moving targets as these need a large number of uncertainty scenarios to robustly optimize their treatment plans. In this study, we propose a novel worst-case robust optimization algorithm, called dynamic minimax, that accelerates the conventional minimax optimization. Dynamic minimax optimization aims at speeding up the plan optimization process by decreasing the number of evaluated scenarios in the optimization.

Methods: For a given pool of scenarios (for instance 63 = 7 setup × 3 range × 3 breathing phases), the proposed dynamic minimax algorithm only considers a reduced number of candidate-worst scenarios, selected from the full 63 scenario set. These scenarios are updated throughout the optimization by randomly sampling new scenarios according to a hidden variable P, called the 'probability acceptance function', which associates with each scenario the probability of it being selected as the worst case. By doing so, the algorithm favors scenarios that are mostly "active", that is, frequently evaluated as the worst case. Additionally, unconsidered scenarios have the possibility to be re-considered, later on in the optimization, depending on the convergence towards a particular solution.

The proposed algorithm was implemented in the open-source robust optimizer MIROpt and tested for six 4D-IMPT lung tumor patients with various tumor sizes

and motions. Treatment plans were evaluated by performing comprehensive robustness tests (simulating range errors, systematic setup errors and breathing motion) using the open-source Monte-Carlo dose engine MCsquare.

Results: The dynamic minimax algorithm achieved an optimization time gain of 84%, on average. The dynamic minimax optimization results in a significantly noisier optimization process due to the fact that more scenarios are accessed in the optimization. However, the increased noise level does not harm the final quality of the plan. In fact, the plan quality is similar between dynamic and conventional minimax optimization with regards to target coverage and normal tissue sparing: on average, the difference in worst-case D95 is 0.2 Gy and the difference in mean lung dose and mean heart dose is 0.4 Gy and 0.1 Gy, respectively (evaluated in the nominal scenario).

Conclusions: The proposed worst-case 4D-robust optimization algorithm achieves a significant optimization time gain of 84%, without compromising target coverage or normal tissue sparing.

Keywords— proton therapy, robust optimization, minimax

I. Introduction

The superior dose distributions produced by intensity-modulated proton therapy (IMPT) indicate a potential for improved patient outcome as compared to conventional X-ray radiotherapy.[1,2,3] However, it is of critical importance that the IMPT treatment plan is made sufficiently robust in order to prevent an unacceptable deterioration of the treatment at the moment of delivery. Successful treatment planning strategies must therefore take into account treatment uncertainties such as tumor motion, setup and range errors.[4,5,6,7] In proton therapy treatment planning, the most effective way of handling these uncertainties is to simulate them during the plan optimization process. This approach has led to the development of robust optimization algorithms which provide an alternative to more conventional margin-based approaches.[8,9,10,11]

In general, the different robust optimization algorithms can be classified into two main groups: (1) probabilistic (or stochastic) optimization and (2) worst-case robust optimization.[12,13] Both groups aim at covering treatment uncertainties by simulating a discrete set of treatment uncertainty scenarios (i.e., realizations of specific combinations of treatment errors). However, the algorithms differ in the way in which the objective function is minimized. Probabilistic optimization algorithms minimize the expected value of the objective function. In contrast, in worst-case robust optimization, the worst-case scenario (the one with the highest objective function value) is chosen, at each iteration, to minimize the objective function.

In this study, we focus on worst-case robust optimization. Different approaches for worst-case robust optimization have been proposed, depending on the way the worst-case scenario is defined. For instance, in voxel-wise worst-case optimization, the worst-case scenario is defined by considering the worst-case value for each individual voxel, among all scenarios (i.e., high dose in organ-at-risk voxels and low dose in the target voxels).[8,9] However, this approach results in a non-physical and potentially overly conservative solution.[10,14] For this reason, Fredriksson *et. al* introduced the so-called '*minimax*' optimization where, for each uncertainty scenario, the objective function is computed for all voxels simultaneously.[10] *Minimax* optimization for IMPT treatment plans have shown to yield clinically acceptable target coverage, in the presence of treatment uncertainties, for a variety of tumor locations.[15,16] The main drawback of both *minimax* and voxel-wise worst-case optimization is their computationally expensive nature, both in terms of the plan computation time and memory consumption. This is due to the following two main issues: first, dose-influence

matrices must be computed and stored for each treatment uncertainty scenario and second, dose distributions must be re-evaluated, at each iteration, for all scenarios defined within the uncertainty set. Because the uncertainty sources (such as tumor motion, setup error and range errors) are usually handled in a mutually independent way, moving targets are especially resource demanding, as their increased number of uncertainty sources amount to a large number of scenarios. This limits the potential of *minimax* optimization as a standard clinical tool and prevents its applicability in online-adaptive workflows.[17]

An example of an approach that aims at reducing the plan computation time is to reduce the number of uncertainty scenarios, with the goal of limiting the number of scenario evaluations during optimization. To this end, in a previous study, a planning strategy was proposed that pre-selects a reduced set of relevant uncertainty scenarios, resulting in a significant plan computation time gain.[18] In contrast, in this study, the full pre-defined uncertainty set is maintained, but we propose an approximate *'dynamic' minimax* algorithm that deals with the inherently long optimization time of the conventional *minimax* optimization algorithm. We focus on accelerating *minimax* optimization by considering only a reduced set of scenarios, selected from the full uncertainty set. This reduced set is then dynamically updated throughout the optimization process, in order to retain only those scenarios that are mostly active in guiding the optimization solution. The present study aims to address the feasibility of this *dynamic minimax* optimization and analyses the time gain with respect to *conventional minimax*. In order to illustrate the proposed method, six lung cancer patients with various tumor sizes and motions are used.

II. Material and Methods

In this section, first, the *conventional minimax* optimization algorithm is formalized, followed by a detailed presentation of the proposed *dynamic minimax* optimization algorithm. Afterwards, an overview is given of the optimization software and patient data used for the testing and evaluation of the respective methods.

II.A. Conventional Minimax Optimization

By representing S as the pre-defined set of uncertainty scenarios s, *conventional minimax* optimization is typically formulated as:

$$\min_{w} \max_{s} \{f_{obj}(d(w,s))\} \quad (1)$$
$$\text{s.t.} \begin{cases} w \geq 0 \\ s \in S, \end{cases}$$

with f_{obj} as the objective function, d the dose distribution and w the optimization variables (i.e., the spot weights) which are constrained to allow only positive solutions. The *conventional minimax* algorithm is characterized by the following three steps performed at each iteration of optimization: (1) the dose distribution is computed for all scenarios s in S with the objective function f_{obj} evaluated in each of the scenarios, (2) the worst-case scenario is selected as the scenario in which the objective function attains its highest value and (3) the spot weights w are updated by minimizing the objective function of the current worst-case scenario.

II.B. Dynamic Minimax Optimization

The proposed algorithm differs from the *conventional minimax* optimization algorithm by decomposing the pre-defined uncertainty set S into two scenario pools: (1) an 'active pool' S_A of candidate-worst scenarios (the pool size of S_A is denoted as N_A) and (2) a 'dead pool' S_D containing the leftover scenarios (the number of dead pool scenarios is denoted as N_D). Hence, the union of both pools is equal to S ($S_A \cup S_D = S$). From this point onward, we denote the active pool scenarios and dead pool scenarios as 'active scenarios' and 'dead scenarios', respectively. The idea is to identify the scenarios that are mostly used in guiding the optimization solution and include these scenarios into the active pool S_A. Subsequently, at each iteration, only the active scenarios ($s \in S_A$) are considered. Hence, the *dynamic minimax* algorithm can be re-formulated as follows:

$$\min_{w} \max_{s} \{f_{obj}(d(w,s))\}$$
$$\text{s.t.} \begin{cases} w \geq 0 \\ s \in S_A. \end{cases}$$

The active scenarios ($s \in S_A$) are probabilistically selected, based on an auxiliary variable P, the so-called 'acceptance probability set' $P = \{P_s \mid s \in S\}$ which associates with each scenario the

probability that it might be evaluated as the worst case.[1] P serves a similar role to the acceptance probability function commonly found in simulated annealing optimization schedules.[19] Because P plays a key role in the *dynamic minimax* algorithm, we explain in the following two paragraphs (1) how P is updated over time and (2) how active scenarios are subsequently selected from P.

II.B.1. Acceptance probability set P

At each iteration, the acceptance probability P is updated by performing two steps. In the first step, the value P_s of the current worst-case scenario ($s = s_{worst}$) is incremented by a factor $\alpha(t)$:

$$P_s(t) = P_s(t-1) + \alpha(t) \text{ if } s = s_{worst}, \tag{2}$$

followed by a re-normalization of P:

$$P_s(t) = P_s(t) \times \frac{1}{1+\alpha(t)} \ \forall \ s \in S, \tag{3}$$

with t the iteration number and $\alpha(t)$ a global time-varying parameter. Following simulated annealing optimization, $\alpha(t)$ is chosen to decay over time and is defined as $\alpha(t) = 1/t$. In doing so, P gradually reduces its sensitivity to fluctuations in the optimization process (so-called optimization noise). In the second step, the values P_s of the current dead scenarios ($s \in S_D$) are incremented by a factor $\alpha(t)/N_D$:

$$P_s(t) = P_s(t-1) + \frac{\alpha(t)}{N_D} \ \forall \ s \in S_D(t), \tag{4}$$

again followed by a re-normalization of P:

$$P_s(t) = P_s(t) \times \frac{1}{1+\alpha(t)} \ \forall \ s \in S. \tag{5}$$

Step 2 is performed in order to add the possibility that yet unconsidered (i.e., dead) scenarios may become active at a later point in the optimization. In Eq. 4, $\alpha(t)$ is weighted by the size of the dead pool, ensuring that a worst-case evaluation (Eq. 2) weights more than its absence from the active pool. The re-normalization steps of P (Eqs. 3 and 5) are necessary to maintain at all times, a total probability mass of 1 (see Section II.B.2.). Additionally, they serve to effectively reduce the values of inactive scenarios (that is, scenarios present in the active pool but not contributing to the optimization) so that these can eventually be discarded.

[1] It must be noted that this scenario 'probability' P_s does not bear a resemblance with the uncertainty probability of the scenario, typically used in probabilistic optimization.

II.B.2. Active pool S_A

Throughout the optimization process, the active pool scenarios are selected by randomly sampling (without replacement), N_A number of scenarios according to their probabilities specified in P. In other words, each scenario can only be drawn once, with the probabilities in P normalized after each draw, in order to maintain a probability mass of 1.

In practice, the active pool is updated at discrete points during the optimization process (in our case at an iteration interval of $\Delta t = 10$). At the start, P is initialized by assigning a uniform probability distribution with no scenarios left unconsidered (i.e., all scenarios $s \in S$ are evaluated). After the first active pool update, the active pool size is set to its reduced size and active scenarios will be selected using the method described above. Furthermore, because the organ-at-risk (OARs) objectives are evaluated in the nominal scenario only, the nominal scenario is always included active pool throughout the entire optimization process.

In general, the *dynamic minimax* algorithm is characterized by the size of the active pool N_A, which is a user-defined parameter. In Section III., we will investigate how the choice of N_A influences the resulting optimization process.

II.C. Optimization Software

The proposed *dynamic minimax* algorithm was implemented in the open-source treatment planning system MIROpt, coded in Matlab (MathWorks, Natick, United States).[20,21] MIROpt uses the open-source Monte Carlo dose engine MCsquare for its dose calculations (MCsquare has been validated for clinical practice from commissioning measurements).[22,23] Dose calculations are performed with 10^5 ions per spot on a $2\times2\times2$ mm^3 dose grid and the spot weights are optimized using a gradient descent algorithm. Constraints on the optimization variable (spot weights w) are handled by a simple projection method, that is, negative values of w are projected to the admissible solution space by setting their values to zero. In order to compare the optimization times of the different optimization algorithms, the maximum number of iterations obtained from the *conventional minimax* optimization is subsequently used in the *dynamic minimax* optimizations.

A quadratic objective function is used to penalize deviations from the pre-defined treatment planning objectives. Target planning objectives were handled robustly (i.e., evaluated for all considered uncertainty scenarios) whilst the OAR objectives were evaluated in the nominal scenario

only. Plan optimization was performed on a 256GB RAM system with a 2x8 Core Intel Xeon processor (E5-2667 v3) @3.20 GHz.

For the *dynamic minimax* optimizations, both the worst-case objective function used to guide the optimization (i.e. evaluated only for the active pool scenarios) as well as the 'real' worst-case objective function (i.e. evaluated for all scenarios) will be reported in the results Section III. Generally, the latter is unavailable as the *dynamic minimax* optimization does not evaluate all uncertainty scenarios at each iteration. However, in order to compare the different methods, additional *dynamic minimax* optimizations are performed where all uncertainty scenarios are evaluated, storing the real worst-case scenarios as well.

II.D. Robustness Evaluation

The robustness of all resulting plans was evaluated with MCsquare, by using a comprehensive approach in which the dose distribution is recomputed on a set of 250 treatment error scenarios. The error scenarios include effects of systematic setup errors, range errors and respiratory motion.[24] Setup errors and range errors are sampled from normal distributions with a standard deviation of 2 mm and 1.6%,[25] respectively, whilst respiratory motion is modeled by recomputing the dose on each breathing phase CT and accumulating the dose on the reference (time-averaged mid-position (MidP)) CT.[26] A 90% confidence interval is generated in the dosimetric space by discarding the 10% worst scenarios (based on the target D_{95}) of the above-mentioned 250 error scenarios.[24] The number of protons is selected in order to reach a statistical uncertainty of 1%.

For the dosimetric plan evaluations, the target DVH metrics (CTV D_{95} and CTV D_5) are calculated in the worst-case *evaluation* scenario, i.e. the scenario where the lowest target coverage is obtained (based on CTV D_{95}) within the 90% confidence interval, generated using the method mentioned above. Similarly, the CTV bandwidths (BW) at the D_{95} and D_5 dose levels, are calculated within the same 90% confidence interval. The OAR DVH metrics are calculated in the nominal scenario only, meaning that the dose distribution is recomputed on the nominal planning CT with a statistical uncertainty of 1%.

II.E. Patient Cases

Six lung tumor patients were chosen to test the proposed optimization algorithm, as their treatment planning typically involves a large number of optimization scenarios, causing long plan optimization times. Patient data were characterized by a 4D-CT image set, binned in ten breathing phases, evenly spaced in time. All patients presented a single tumor volume, delineated on the MidP-CT. The main features of the patient cohort are summarized in Table 1. All patients had a dose prescription of 60 Gy to the clinical-target-volume (CTV) with target coverage considered acceptable if 95% of the CTV received more than 95% of the prescribed dose (= 57 Gy), whilst no more than 5% of the CTV received over 105% of the prescribed dose (= 63 Gy), for the worst-case scenario.

All treatment plans used the MidP-CT as the nominal planning CT which was created with the open-source platform OpenReggui.[26,27] Treatment plans were optimized using uncertainty scenarios that contain setup errors, range errors and respiratory motion. Similar to other studies, uncertainty parameters were chosen as combinations of 5 mm setup errors in the three directions (left-right, anterior-posterior and superior-inferior), ±3% range error and maximum inhale and exhale breathing phases, generating an uncertainty set of 63 scenarios (= 7 setup error scenarios × 3 range error scenarios × 3 breathing phases).[6,8,10,28] Setup and range errors are modeled by rigidly shifting the CT image and uniformly scaling the CT mass densities (obtained from the CT image), respectively. All treatment plans were designed using a configuration of three co-planar beams, delivered via IMPT with the pencil beam scanning (PBS) technique (see Table 1).

Table 1: Patient characteristics.

Patient	CTV size	Motion Amplitude			Tumor position	Beam angles
		LR	AP	SI		
	[cm^3]	[mm]	[mm]	[mm]		[°]
P1	152.6	4.2	2.1	3.1	RML	0, 270, 310
P2	107.7	3.1	2.9	3.7	LLL	90,135, 180
P3	41.3	1.4	2.9	0.8	RUL	180, 225, 270
P4	70.3	0.8	1.2	0.5	LUL	90, 135, 180
P5	109.6	2.2	1.8	6.6	RUL	180, 225, 270
P6	249.7	2.1	2.5	10.6	RLL	180, 225, 270

Tumor motion amplitude (in left-right (LR), anterior-posterior (AP) and superior-inferior (SI) directions). Tumor positions (right-middle lobe (RML), left-lower lobe (LLL), right-upper lobe (RUL), right-lower lobe (RLL) and left-upper lobe (LUL)).

III. Results

In this section, the performance of *dynamic minimax* optimization algorithm is compared to the *conventional minimax* optimization. As mentioned in Section II., the *conventional minimax* algorithm evaluates, at each iteration, all 63 scenarios in the uncertainty set. Because the *dynamic minimax* is characterized by the parameter N_A, we present the results for two different choices of N_A, that is, $N_A = 15$ and a more extreme case of $N_A = 5$. The quality of the optimizations will be measured, first, according to the value of the worst-case objective function value throughout the optimization process (so-called optimization curve) and second, from the dosimetric metrics (target coverage, robustness and OAR sparing) obtained after performing comprehensive robustness tests (see Section II.D.).

III.A. Optimization Data

Table 2 reports the plan optimization times, together with the final (worst-case) objective function value. For the *dynamic minimax* optimizations, both the final objective function f_{obj} as well as the 'real' final objective function f_{obj}^{real} (see Section II.C.) are reported.

Results show that the *dynamic minimax* algorithm achieved an average time gain of 84% and 67%, for the 5 and 15 active pool size optimizations, respectively. The final objective function

values of the different optimization methods are similar in magnitude for all test cases, with only a small difference between f_{obj}^{real} and f_{obj}.

In Fig. 1 (top and middle panels), the optimization curves of the three optimizations (*conventional minimax*, $N_A = 15$ and $N_A = 5$ *dynamic minimax*) are compared. All optimizations follow a similar trend but with the $N_A = 5$ optimization lying below the *conventional* throughout the entire optimization process. The $N_A = 5$ optimization does appear to be significantly the noisiest. Fig. 1 (middle) shows that the real worst-case optimization curve of the $N_A = 5$ optimization deviates slightly during an early stage but reaches similar values near the end of the optimization process.

Fig. 1 (bottom) shows the number of iterations that a scenario (ordered from 1 to 63) is selected as the worst case. Although mostly similar, the *conventional minimax* optimization accessed the least amount of scenarios, in order to reach its final solution. In contrast, the *dynamic minimax* optimizations use a larger number of rarely accessed scenarios with the bigger pool size matching closely the *conventional minimax* optimization.

Table 2: Plan optimization time, final worst-case objective function value f_{obj} (evaluated only for the active pool scenarios) and real final worst-case objective function value f_{obj}^{real} (evaluated for all scenarios). Plans of each patient (P1-6) were obtained using the *conventional minimax* optimization (Ref.) and *dynamic minimax* optimization algorithms with pool sizes of N_A=5 and N_A=15. The average time reductions (in %) are reported at the bottom.

	Optimization time [min]			**Final f_{obj} [Gy^2]**			**Final f_{obj}^{real} [Gy^2]**	
	Ref.	N_A=15	N_A=5	Ref.	N_A=15	N_A=5	N_A=15	N_A=5
P1	513	170	85	1.55	1.43	1.08	1.45	1.27
P2	396	142	72	0.74	0.61	0.55	0.63	0.69
P3	167	47	22	1.96	1.78	1.63	1.79	2.24
P4	219	79	32	2.97	2.50	1.71	2.57	3.02
P5	409	152	83	1.02	0.93	0.73	0.98	1.07
P6	758	213	107	6.0	5.3	4.4	5.3	4.8
ΔAvg.		-67%	-84%					

III.B. Dosimetric Results

Table 3 and Table 4 show the target and OAR DVH metrics for the obtained treatment plans. Target coverage metrics (D_{95} and D_5) are calculated in the worst-case *evaluation* scenario whilst the OAR metrics are calculated in the nominal scenario only (see Section II.D.). Furthermore, the

average difference between the value in the reference plan (obtained using *conventional minimax* optimization algorithm) with plans optimized using the *dynamic minimax* algorithms is shown for each metric.

On average, equal target coverage (worst-case CTV D_{95}) is obtained between the *conventional minimax* and $N_A = 15$ *dynamic minimax* optimization. The $N_A = 5$ *dynamic minimax* optimization improved worst-case CTV D_{95} slightly by 0.2 Gy, on average, with respect to the reference plans. OAR dose is similar between all studied plans (average difference of mean lung dose of only 0.2 Gy and 0.4 Gy between the *conventional minimax* and $N_A = 15$ and $N_A = 5$ *dynamic minimax* optimizations, respectively and difference in mean esophagus dose of -0.1 Gy and 0.1 Gy, respectively).

Fig. 2 displays the dose distribution together with the corresponding DVHs for each optimization method. Results indicate similar dose profiles between all plans with isodose lines that nearly coincide. This similarity translates to DVHs that have a similar sensitivity to the treatment errors (indicated by the CTV BWs in Table 3) and matching OAR DVH curves.

Table 3: Target coverage metrics (CTV D_{95} and D_5) and robustness metrics (CTV bandwidth (BW) at the D_{95} and D_5 dose level) for plans of all patients (P1-6), obtained using *conventional minimax* optimization (Ref.) and *dynamic minimax* optimization with pool sizes of N_A=5 and N_A=15. CTV D_{95} and D_5 are computed in the worst-case *evaluation* scenario.

	CTV					
	Worst-case D_{95} [Gy]			Worst-case D_5 [Gy]		
	Ref.	N_A=15	N_A=5	Ref.	N_A=15	N_A=5
P1	57.0	56.9	57.3	62.8	62.4	62.4
P2	57.6	57.6	57.3	61.8	61.8	61.7
P3	58.0	57.9	58.5	62.6	62.6	61.9
P4	58.2	58.3	58.6	62.1	62.1	62.4
P5	58.3	58.4	58.5	61.7	61.7	61.6
P6	57.2	57.4	57.2	64.2	63.6	63.6
ΔAvg.		0.0	+0.2		-0.2	-0.3
	BW at D_{95} [Gy]			BW at D_5 [Gy]		
	Ref.	N_A=15	N_A=5	Ref.	N_A=15	N_A=5
P1	1.9	1.8	1.5	1.1	0.9	1.0
P2	1.1	0.9	0.9	0.9	0.8	0.8
P3	0.6	0.8	0.3	1.1	1.2	1.4
P4	0.6	0.6	0.4	1.0	1.0	1.4
P5	0.6	0.5	0.4	0.6	0.7	0.5
P6	1.6	1.5	1.7	1.6	1.6	1.6
ΔAvg.		-0.1	-0.2		0.0	+0.1

Table 4: Organ-at-risk DVH metrics (lung, esophagus and heart) for plans of all patients (P1-6), obtained using *conventional minimax* optimization (Ref.) and *dynamic minimax* optimization with pool sizes of N_A=5 and N_A=15. Metrics have been computed in the nominal scenario.

	Lung						**Esophagus**			**Heart**		
	V_{20} [%]			D_{mean} [Gy]			D_{mean} [Gy]			V_{40} [%]		
	Ref.	N_A=15	N_A=5	Ref.	N_A=15	N_A=5	Ref.	N_A=15	N_A=5	Ref.	N_A=15	N_A=5
P1	26.4	26.5	28.3	13.5	13.6	14.1	2.0	2.0	2.1	3.2	3.3	3.4
P2	26.9	27.2	27.8	13.5	13.6	13.9	5.4	5.5	5.7	3.8	3.9	4.0
P3	13.4	13.4	13.6	7.0	7.0	7.2	4.8	4.8	5.0	0.0	0.0	0.0
P4	19.0	19.1	19.4	9.7	9.8	10.0	2.1	2.1	2.2	0.0	0.0	0.0
P5	21.9	22.1	22.4	10.6	10.8	10.9	7.9	8.0	8.3	1.1	1.2	1.2
P6	30.0	31.6	31.6	16.0	16.6	16.6	20.3	19.7	19.7	3.3	3.4	3.4
ΔAvg.		+0.4	+0.9		+0.2	+0.4		-0.1	+0.1		+0.1	+0.1

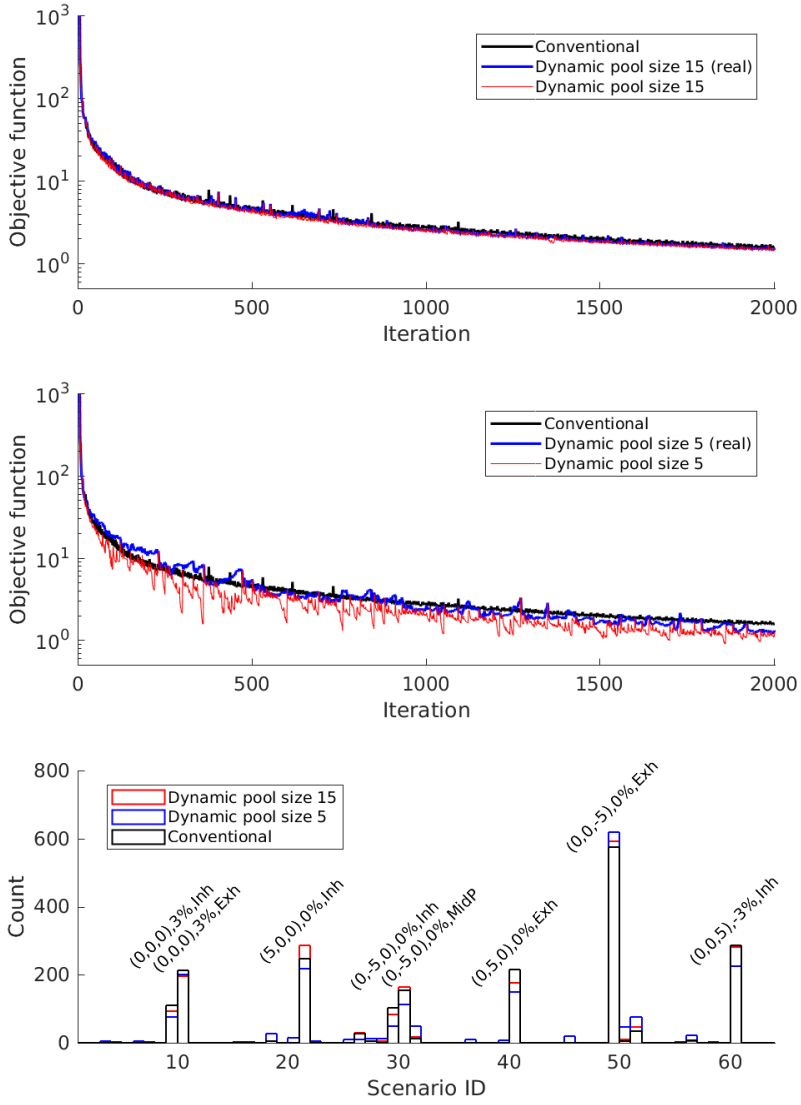

Figure 1: Comparison of *conventional minimax* and *dynamic minimax* optimizations (results of patient P1 are displayed). The top and middle panels show the progression of the (worst-case) objective function value throughout the optimization (top: pool size of N_A=15 and middle: pool size of N_A=5). For the *dynamic minimax* optimization, the worst-case objective function f_{obj} used to guide the optimization (i.e. evaluated only for the active pool scenarios) is displayed in red, whilst the real worst-case objective function f_{obj}^{real} (i.e. evaluated for all scenarios) is displayed in blue. The bottom panel shows the histogram displaying the number of iterations (= counts) that each scenario is evaluated as the worst case. The magnitude of the uncertainties is shown for the most counted scenarios. The uncertainties are displayed as follows: setup error (x,y,z) in mm in the left-right x, anterior-posterior y and superior-inferior z directions, range error and breathing phase (MidP, max inhale (Inh) or max exhale (Exh)).

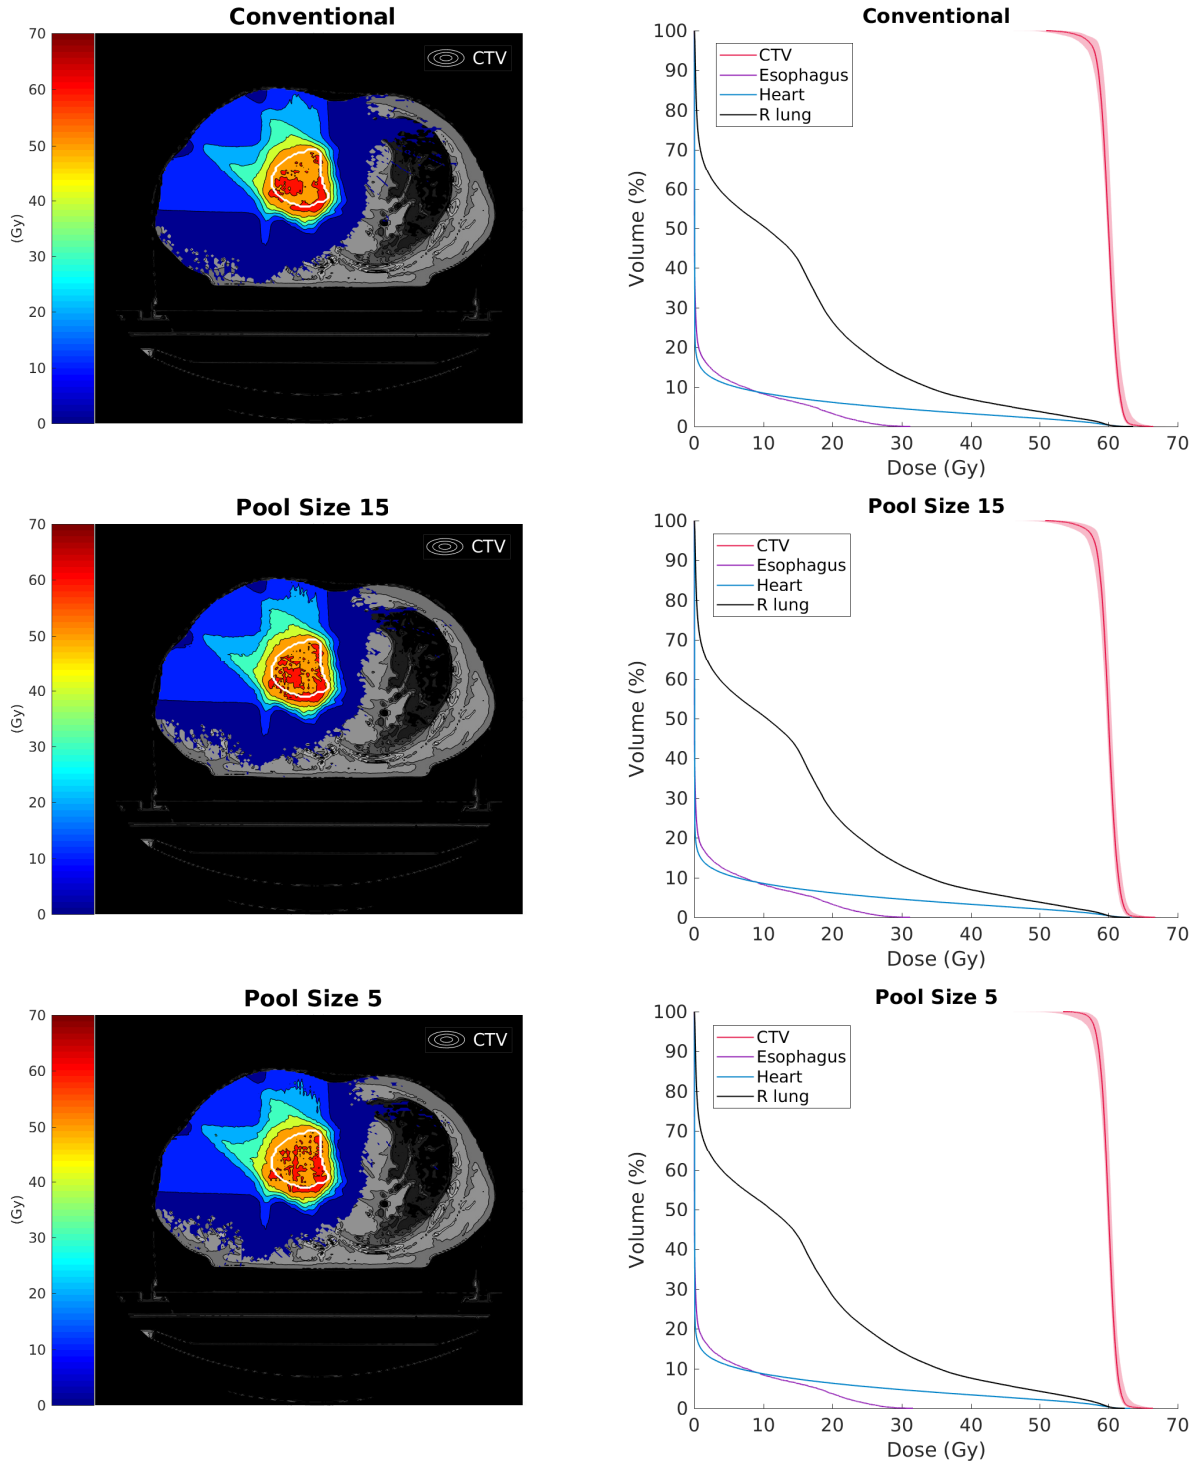

Figure 2: The left column shows the dose distributions for plans obtained using *conventional minimax* optimization and *dynamic minimax* optimization with pool sizes of N_A=5 and N_A=15 for patient P1. In each figure, the CTV is indicated in white. The right column shows the corresponding DVHs with the CTV-DVH band representing the evaluations in the considered error *evaluation* scenarios (see Section II.D.).

IV. Discussion

In *minimax* optimization, only the current worst-case scenario is used to guide the optimization solution. In the meantime, as noted by Fredriksson *et. al*,[14] *minimax* algorithms tend to neglect so-called 'easy' scenarios, that is, scenarios where there is little conflict between organ sparing and target coverage in the objective function. Hence, a substantial amount of computation time and resources are potentially wasted on scenario evaluations that are rarely the worst case. Fig. 1 (bottom) illustrates this feature of *minimax* optimization by showing that the optimizer only accesses a fraction of the full uncertainty set in order to reach its final solution. This suggests that the majority of scenarios produce either comparable dose distributions or produce dose distributions where the planning objectives are consistently well respected. Fredriksson argues that disregarding 'easy' scenarios is one of the main disadvantages of the *minimax* algorithm when comparing it to other classes of robust optimization algorithms.[14] In fact, it is exactly this drawback that the *dynamic minimax* algorithm attempts to address. By relying on the sparsity of active scenarios in the solution space, fewer scenarios are needed whilst still preserving most of the information of the full problem. In doing so, the computational cost of an iteration is significantly reduced (in other words, the number of scenario evaluations performed at each iteration is reduced), resulting in an accelerated optimization process (a time gain of up to 84% is obtained).

The optimization curves in Section III.A. show that by reducing the size of the active pool N_A, the optimization noise level increases. Fundamentally, worst-case robust optimization is inherently a noisy optimization process. This is explained by the fact that different optimization scenarios are used throughout the optimization as a result of the discontinuous *max* operator (see Eq. 1). Additionally, the projection method (see Section II.C.), to handle constraints on the optimization variables (the spot weights), also adds noise to the optimization. In addition to the above-mentioned noise sources, the *dynamic minimax* algorithm, will add optimization noise by regularly changing the possible optimization scenarios throughout the optimization process. This effect will be more pronounced for smaller active pool sizes, which change their composition more frequently. The additional noise level produced by the *dynamic minimax* algorithm is exemplified in Fig. 1 (bottom). As shown, optimizations with smaller pool sizes will explore an increased number of scenarios in the solution space. By increasing the pool size slightly (to $N_A = 15$), the noise is reduced to a level comparable in magnitude to the *conventional minimax* optimization. However, as the results of Section III.B. indicate, the increased optimization noise level does not harm the

final quality of the treatment plans. In fact, results indicate that a noisy optimization trajectory in the solution space might be advantageous in order to further explore and eventually find a better solution; this is an approach commonly employed in simulated annealing and stochastic gradient descent optimization schedules.

The *dynamic minimax* algorithm was tested for a patient population of six lung tumor cases. Therefore, in order to further validate the optimality of the proposed algorithm parameters (mainly the pool size), the algorithm should be tested for a wider set of patient cases. For instance, in highly complex cases (i.e. large tumor motion with considerable conflicts among the planning objectives), it might be advisable to employ a conservative approach and use a larger the pool size. This would guarantee that important scenarios are not missed throughout the optimization. Based on the results of the present study, by using a pool size of 15, almost equal results are obtained as for the *conventional minimax* whilst still achieving a significant plan optimization time gain of 67%.

It must be noted that this study only focuses on reducing the optimization time and does not deal with other computational aspects (such as the memory consumption) of *minimax* optimization. In particular, the computation of the beamlet dose-influence matrices gives a large contribution to the overall plan computation time (especially for Monte Carlo-based dose computations). The following solutions exist that can reduce the dose computation time and which could potentially be used in conjunction with the *dynamic minimax* optimization: first, the number of beamlet dose-influence matrices can be reduced by performing a pre-selection of relevant uncertainty scenarios,[18] and second, a hybrid Monte Carlo-pencil beam dose optimizer can be used to accelerate the plan computation time with Monte-Carlo like accuracy.[21]

V. Conclusions

In robust *minimax* optimization, the dose distributions must be evaluated for all uncertainty scenarios in order to evaluate their respective objective functions. As a result, the plan optimization time linearly scales with the number of pre-defined uncertainty scenarios. Especially for lung tumor patients, which need a large number of scenarios to robustly optimize their treatment plans, the associated computational burden may cause excessive plan computation times. This issue limits the use of robust optimization in the clinical environment.

In this study, we propose an approximate worst-case robust optimization algorithm that accelerates *minimax* optimization. The proposed *dynamic minimax* algorithm relies on the fact that *minimax* algorithms neglect so-called 'easy' scenarios where there is little conflict among the planning objectives. Therefore, instead of evaluating all scenarios in the pre-defined uncertainty set, only a reduced set of active pool scenarios is considered. Following stochastic annealing optimization schedules, these active scenarios are updated according to a variable called the 'acceptance probability set'. This variable expresses the probability that a scenario might be evaluated as the worst case. By doing so, only the scenarios that are contributing most to the optimization, at that moment, will be retained and accessible in order to guide the optimization solution. The proposed method was applied to 4D-robust *minimax* optimization and tested for six moving lung tumor cases. Results show that, on average, an optimization time gain of up to 84% is achieved without compromising either target robustness or normal tissue sparing.

Acknowledgements

Gregory Buti is supported by the Télévie Grant from the Belgian 'Fonds National pour la Recherche Scientifique' F.R.S-FNRS (Grant No. 7453918F). Computational resources have been provided by the supercomputing facilities of the Université Catholique de Louvain (CISM/UCL) and the Consortium des Équipements de Calcul Intensif en Fédération Wallonie Bruxelles (CÉCI) funded by the F.R.S.-FNRS under convention 2.5020.11. Kevin Souris is funded by the Walloon region (MECATECH/BIOWIN, Grant No. 8090). Ana M. Barragán Montero is funded by the Walloon region (PROTHERWAL/CHARP, Grant No. 7289). Marie Cohilis is supported by the Télévie Grant from the F.R.S-FNRS (Grant No. 7450517F). John A. Lee is a Senior Research Associate with the F.R.S.-FNRS.

References

[1] A. Elhammali, P. Blanchard, A. Yoder, Z. Liao, X. Zhang, X. R. Zhu, P. K. Allen, M. Jeter, J. Welsh, and Q.-N. Nguyen, Clinical outcomes after intensity-modulated proton therapy with concurrent chemotherapy for inoperable non-small cell lung cancer, Radiotherapy and Oncology **136**, 136–142 (2019).

[2] N. Nakamura, K. Hotta, S. Zenda, H. Baba, S. Kito, T. Akita, A. Motegi, H. Hojo, M. Nakamura, R. V. Parshuram, M. Okumura, and T. Akimoto, Hypofractionated proton beam therapy for centrally located lung cancer, Journal of Medical Imaging and Radiation Oncology **63**, 552–556 (2019).

[3] R. M. Hoshina, T. Matsuura, K. Umegaki, and S. Shimizu, A Literature Review of Proton Beam Therapy for Prostate Cancer in Japan, Journal of Clinical Medicine **8**, 48 (2019).

[4] P. C. Park, J. P. Cheung, X. R. Zhu, A. K. Lee, N. Sahoo, S. L. Tucker, W. Liu, H. Li, R. Mohan, L. E. Court, and L. Dong, Statistical Assessment of Proton Treatment Plans Under Setup and Range Uncertainties, International Journal of Radiation Oncology∗Biology∗Physics **86**, 1007–1013 (2013).

[5] S. Brousmiche, K. Souris, J. O. de Xivry, J. A. Lee, B. Macq, and J. Seco, Combined influence of CT random noise and HU-RSP calibration curve nonlinearities on proton range systematic errors, Physics in Medicine & Biology **62**, 8226–8245 (2017).

[6] A. J. Lomax, Intensity modulated proton therapy and its sensitivity to treatment uncertainties 1: the potential effects of calculational uncertainties, Physics in Medicine and Biology **53**, 1027–1042 (2008).

[7] A. J. Lomax, Intensity modulated proton therapy and its sensitivity to treatment uncertainties 2: the potential effects of inter-fraction and inter-field motions, Physics in Medicine and Biology **53**, 1043–1056 (2008).

[8] W. Liu, X. Zhang, Y. Li, and R. Mohan, Robust optimization of intensity modulated proton therapy, Medical Physics **39**, 1079–1091 (2012).

[9] D. Pflugfelder, J. J. Wilkens, and U. Oelfke, Worst case optimization: a method to account for uncertainties in the optimization of intensity modulated proton therapy, Physics in Medicine and Biology **53**, 1689–1700 (2008).

[10] A. Fredriksson, A. Forsgren, and B. Hårdemark, Minimax optimization for handling range and setup uncertainties in proton therapy, Medical Physics **38**, 1672–1684 (2011).

[11] J. Unkelbach, T. Bortfeld, B. C. Martin, and M. Soukup, Reducing the sensitivity of IMPT treatment plans to setup errors and range uncertainties via probabilistic treatment planning, Medical Physics **36**, 149–163 (2008).

[12] A. Fredriksson, A characterization of robust radiation therapy treatment planning methods-from expected value to worst case optimization, Medical Physics **39**, 5169–5181 (2012).

[13] J. Unkelbach, M. Alber, M. Bangert, R. Bokrantz, T. C. Y. Chan, J. O. Deasy, A. Fredriksson, B. L. Gorissen, M. van Herk, W. Liu, H. Mahmoudzadeh, O. Nohadani, J. V. Siebers, M. Witte, and H. Xu, Robust radiotherapy planning, Physics in Medicine & Biology **63**, 22TR02 (2018).

[14] A. Fredriksson and R. Bokrantz, A critical evaluation of worst case optimization methods for robust intensity-modulated proton therapy planning, Medical Physics **41**, 081701 (2014).

[15] L. V. van Dijk, R. J. H. M. Steenbakkers, B. ten Haken, H. P. van der Laan, A. A. van 't Veld, J. A. Langendijk, and E. W. Korevaar, Robust Intensity Modulated Proton Therapy (IMPT) Increases Estimated Clinical Benefit in Head and Neck Cancer Patients, PLOS ONE **11**, e0152477 (2016).

[16] D. Cummings, S. Tang, W. Ichter, P. Wang, J. D. Sturgeon, A. K. Lee, and C. Chang, Four-dimensional Plan Optimization for the Treatment of Lung Tumors Using Pencil-beam Scanning Proton Radiotherapy, Cureus (2018).

[17] K. Bernatowicz, X. Geets, A. Barragan, G. Janssens, K. Souris, and E. Sterpin, Feasibility of online IMPT adaptation using fast, automatic and robust dose restoration, Physics in Medicine & Biology **63**, 085018 (2018).

[18] G. Buti, K. Souris, A. M. B. Montero, J. A. Lee, and E. Sterpin, Towards fast and robust 4D optimization for moving tumors with scanned proton therapy, Medical Physics (2019).

[19] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, Optimization by Simulated Annealing, Science **220**, 671–680 (1983).

[20] A. M. B. Montero, Miropt - http://www.openmiropt.org/ Accessed November 2019.

[21] A. M. B. Montero, K. Souris, D. Sanchez-Parcerisa, E. Sterpin, and J. A. Lee, Performance of a hybrid Monte Carlo-Pencil Beam dose algorithm for proton therapy inverse planning, Medical Physics **45**, 846–862 (2017).

[22] K. Souris, MCsquare - http://www.openmcsquare.org/ Accessed June 2019.

[23] S. Huang, M. Kang, K. Souris, C. Ainsley, T. D. Solberg, J. E. McDonough, C. B. Simone, and L. Lin, Validation and clinical implementation of an accurate Monte Carlo code for pencil beam scanning proton therapy, Journal of Applied Clinical Medical Physics **19**, 558–572 (2018).

[24] K. Souris, A. B. Montero, G. Janssens, D. D. Perri, E. Sterpin, and J. A. Lee, Technical Note: Monte Carlo methods to comprehensively evaluate the robustness of 4D treatments in proton therapy, Medical Physics (2019).

[25] H. Paganetti, Range uncertainties in proton therapy and the role of Monte Carlo simulations, Physics in Medicine and Biology **57**, R99–R117 (2012).

[26] M. Wanet, E. Sterpin, G. Janssens, A. Delor, J. A. Lee, and X. Geets, Validation of the mid-position strategy for lung tumors in helical TomoTherapy, Radiotherapy and Oncology **110**, 529–537 (2014).

[27] G. Janssens, OpenReggui - https://www.openreggui.org/ Accessed June 2019.

[28] T. Inoue, J. Widder, L. V. van Dijk, H. Takegawa, M. Koizumi, M. Takashina, K. Usui, C. Kurokawa, S. Sugimoto, A. I. Saito, K. Sasai, A. A. van't Veld, J. A. Langendijk, and E. W. Korevaar, Limited Impact of Setup and Range Uncertainties, Breathing Motion, and Interplay Effects in Robustly Optimized Intensity Modulated Proton Therapy for Stage III Non-small Cell Lung Cancer, International Journal of Radiation Oncology∗Biology∗Physics **96**, 661–669 (2016).