Published in final edited form as:

Med Phys. 2020 July ; 47(7): 2779–2790. doi:10.1002/mp.14148.

Automated proton treatment planning with robust optimization using constrained hierarchical optimization

Vicki T. Taasti*, Linda Hong, Joseph O. Deasy, Masoud Zarepisheh

Department of Medical Physics, Memorial Sloan Kettering Cancer Center, New York, NY

Abstract

Purpose—We present a method for fully automated generation of high quality robust proton treatment plans using hierarchical optimization. To fill the gap between the two common extreme robust optimization approaches, i.e., stochastic and worst-case, a robust optimization approach based on the p -norm function is used whereby a single parameter, p , can be used to control the level of robustness in an intuitive way.

Methods—A fully automated approach to treatment planning using Expedited Constrained Hierarchical Optimization (ECHO) is implemented in our clinic for photon treatments. ECHO strictly enforces critical (inviolable) clinical criteria as hard constraints and improves the desirable clinical criteria sequentially, as much as is feasible. We extend our in-house developed ECHO codes for proton therapy and integrate it with a new approach for robust optimization. Multiple scenarios accounting for both setup and range uncertainties are included (thirteen scenarios), and the maximum/mean/dose-volume constraints on organs-at-risk (OARs) and target are fulfilled in all scenarios. We combine the objective functions of the individual scenarios using the p -norm function. The p -norm with a parameter $p=1$ or $p=\infty$ result in the stochastic or the worst-case approach, respectively; an intermediate robustness level is obtained by employing p -values in-between. While the worst-case approach only focuses on the worst-case scenario(s), the p -norm approach with a large p value ($p\approx 20$) resembles the worst-case approach without completely neglecting other scenarios. The proposed approach is evaluated on three head-and-neck (HN) patients and one water phantom with different parameters, $p\{1,2,5,10,20\}$. The results are compared against the stochastic approach (p -norm approach with $p=1$) and the worst-case approach as well as the non-robust approach (optimized solely on the nominal scenario).

Results—The proposed algorithm successfully generates automated robust proton plans on all cases. As opposed to the non-robust plans, the robust plans have narrower dose volume histogram (DVH) bands across all thirteen scenarios, and meet all hard constraints (i.e., maximum/mean/dose-volume constraints) on OARs and the target for all scenarios. The spread in the objective function values is largest for the stochastic approach ($p=1$) and decreases with increasing p towards the worst-case approach. Compared to the worst-case approach, the p -norm approach results in DVH bands for clinical target volume (CTV) which are closer to the prescription dose at a negligible cost in the DVH for the worst scenario, thereby improving the overall plan quality. On average, going from the worst-case approach to the p -norm approach with $p=20$, the median

*Corresponding author: vicki.taasti@maastro.nl; currently at Maastric Clinic, Maastricht, The Netherlands.

objective function value across all the scenarios is improved by 15% while the objective function value for the worst scenario is only degraded by 3%.

Conclusion—An automated treatment planning approach for proton therapy is developed, including robustness, dose-volume constraints, and the ability to control the robustness level using the p -norm parameter p , to fit the priorities deemed most important.

Keywords

Robust optimization; constrained optimization; proton therapy treatment planning; dose volume constraints; p -norm

1. Introduction

Proton therapy has the potential to spare healthy tissue to a higher extent than photon therapy due to the finite range of protons. However, the finite range also makes protons more sensitive to potential errors, such as setup or range estimation errors [1]. Robust optimization can take these errors into account during the treatment planning by simulating different errors and then optimizing the treatment plan with the objective of getting a good dose distribution even in the presence of error scenarios. Today, all major treatment planning systems (TPSs) provide robust optimization for proton treatment planning [1].

Several methods have been proposed to robustly optimize the treatment plan to ensure a high plan quality even under errors, but most suggested methods fall in one of two main categories [1]: 1) the stochastic programming approach where the expectation value of the objective functions over all the included scenarios is minimized (lower is better) [2], or 2) the worst-case approach (also known as minimax), where only the objective function in the worst scenario(s) is optimized [3]. While the worst-case approach ensures a good plan quality even for the most detrimental errors, it completely neglects other scenarios and does not try to improve them even if it is possible at no cost. The stochastic approach, on the other hand, treats all the scenarios equally and could potentially favor improvements in less severe scenarios and focus less on improving the more severe error scenarios which is also undesirable. It is therefore important to have an intermediate approach which considers all scenarios, similar to the stochastic approach, but weights the more severe scenarios higher, similar to the worst-case approach. Fredriksson proposed the *minimax stochastic programming* approach to address this issue where an uncertainty set for the scenario probabilities is defined and then the expected value of the objective function for the most unfavorable probability distribution is optimized. A special form of this approach, that has received more attention, is equivalent to the so-called conditional value at risk (CVaR) optimization which mainly concerns about a fraction of the most severe scenarios [4]. In this study, we employ the p -norm function over the objective functions for all scenarios to provide an intermediate robust optimization approach. The level of robustness and importance of focusing on more severe scenarios can also be easily and intuitively adjusted using the p -norm parameter.

Treatment planning is a labor-intensive and time-consuming process. Many techniques have been developed to automate or facilitate the treatment planning process including:

1) knowledge-based planning [6], 2) multi-criteria optimizations [7], and 3) hierarchical constrained optimization [8–10]. Different flavors of hierarchical constrained optimization have been developed in the past [9–11]. Our group has also recently developed a technique that we refer to as ECHO (Expedited Constrained Hierarchical Optimization) [7]. ECHO is being used in our daily clinical routinely, with more than 2000 patients treated to the end of 2019. Recently ECHO has been extended to also incorporate dose-volume constraints (DVCs) [12]. ECHO is designed to meet all the hard maximum/mean/dose-volume constraints on the OARs and target (note, DVCs are applied to OARs only, and not the target). The algorithm proceeds sequentially, in three steps: 1) maximizes the tumor coverage, 2) minimizes OARs doses, and 3) optimizes the plan delivery. In this study, we adapt the first two steps of ECHO for proton therapy and incorporate robustness into this framework to ensure all the hard maximum/mean/dose-volume constraints are fulfilled in all scenarios. To the best of our knowledge, this work is first to demonstrate integration of robust optimization with a method which strictly fulfills all the hard maximum/mean/dose-volume constraints. We employ the p -norm robust optimization approach to combine the objective functions in each of the two steps to optimize the clinical target volume (CTV) coverage, and OARs doses corresponding to all scenarios with special emphasis on more severe scenarios. As such, the proposed approach can automatically generate robust proton plans without any parameter tweaking. We demonstrate the optimization framework on three head and neck (HN) cancer patients and a phantom considering thirteen simulated error scenarios corresponding to systematic errors for setup and range estimation.

2. Materials and methods

2.1 p -norm approach for robust optimization

In robust optimization, error scenarios, e.g. setup and range estimation errors, are simulated and their corresponding dose distributions are included in the optimization in addition to the nominal scenario (the scenario, where no errors are included). Worst-case and stochastic optimization are two commonly used robustness approaches which optimize the maximum and expected value of the objective functions corresponding to dose distributions for all scenarios, respectively. The worst-case approach can be stated as follows:

$$\min_x \left(\max_{s \in S} f(d(x, s)) \right) \quad (1)$$

subject to

$$x \geq 0 \quad (2)$$

$$d(x, s) = A(s) \cdot x \quad (3)$$

where x is the optimization variable (the proton spot intensity vector), S is the set of error scenarios, $A(s)$ is the so-called dose influence matrix (corresponding to scenario s) representing the deposited dose at each voxel given a unit intensity for each spot. $d(x, s)$ is the resulting deposited dose, and $f(d(x, s))$ is the objective function evaluation of the deposited dose. This formulation is also known as the composite worst-case approach

where the maximum is taken over the objective functions, as opposed to over the voxels which would result into the voxel-wise worst-case approach [13]. Solving the above Minimax problem could be computationally challenging due to the non-differentiability of the objective function in equation (1). A common technique, which has been also followed in this study, is to replace $\max_{s \in S} f(d(x, s))$ with an auxiliary variable z in equation (1) and then adding the constraint $f(d(x, s)) \leq z$ for each scenario s . Fredriksson *et al.* [3] replaced $\max_{s \in S} f(d(x, s))$ with a to address the computational challenge. In this study, we also use p -norm function, although for a different purpose.

In the stochastic programming approach, the expected value of the objective functions is minimized and Eq. (1) is replaced with Eq. (4):

$$\min_x \left(\sum_{s \in S} w_s \cdot f(d(x, s)) \right) \quad (4)$$

Here the weights, w_s , can be interpreted as the probability for scenario s to occur [1]. The constraints given in Eq. (2) and (3) are retained in this equation, and they will be implied in the below formalism as well.

To control the robustness level in-between the worst-case and the stochastic approach, we propose to combine the error scenarios using the p -norm function. Our proposed robust optimization algorithm can then be stated as:

$$\min_x \left(\sum_{s \in S} (f(d(x, s)))^p \right)^{\frac{1}{p}} \quad (5)$$

where p is the p -norm parameter that controls the behavior of the function. For $p=\infty$ and $p=1$, this p -norm formalism is reduced to the worst-case and stochastic approach, respectively, and any intermediate p -value results in an intermediate robustness level between these two extreme approaches. When $p \rightarrow 1$ and $f \rightarrow 0$ (which is always the case for the objective functions used in this study, see Table 1), the p -norm is known to be convex that is critical to keep the algorithm tractable [5]. N_s is the number of scenarios, which in this study will be thirteen (section 2.4). Here we use equal weights for all scenarios, $w_s=1/N_s$, however, if there are probabilities associated with the scenarios, w_s , they can also be easily included in Eq. (5) similar to Eq. (4).

2.2 Expedited Constrained Hierarchical Optimization (ECHO) and Dose Volume Constraints (DVCs)

ECHO optimization proceeds in a hierarchical order and it solves three constrained optimization problems (Step-1 to 3). Step-1 optimizes the target coverage while respecting all the required clinical criteria (expressed as hard maximum/mean dose constraints on OARs and target). The results of Step-1 for the target dose distribution are turned into constraints for Step-2 to be preserved with a slight relaxation by introducing so-called slip

factors and then Step-2 minimizes the OAR doses beyond the required clinical criteria as much as possible. Step-3 smooths out the fluence map; this step we skip here for proton therapy. More details about ECHO can be found in the original publication [8]. DVCs have been later integrated into ECHO [12] by adding an extra Step-0 before Step-1 and 2. Hard DVCs can in principle be incorporated into optimization by introducing binary variables which would flag the voxels that exceed the DVC threshold and then restrict the number of voxels violating the threshold. However, the resulting mixed integer programming problem would be computationally intractable. Step-0 is essentially the same as Step-1 but it also includes the convex relaxation of DVCs obtained by lifting the binary restriction on the variables and allowing them to take any value between 0 and 1. Although Step-0 does not guarantee the satisfaction of DVCs due to the convex approximation, it provides a crucial information about voxels likely to receive doses less than the DVC threshold, hereafter referred to as sub-threshold voxels. Step-1 and 2 subsequently impose maximum dose constraints on these sub-threshold voxels which would in turn guarantee the DVCs to be satisfied. More details on the DVC method can be found in [12].

2.3 Robust automated proton treatment planning

To add robustness in ECHO, we first replace the PTV with the CTV and then enforce the hard maximum/mean dose constraints in all scenarios in Step-0 to 2. For p -norm approach, we need to modify the objective functions of all the steps according to Eq. (5) (and to Eq. (1) for the worst-case approach). In Step-2, the results of Step-1 corresponding to each scenario are turned into constraints to be preserved as follows:

$$f^{\text{Step1}}(d(x^{\text{Step2}}, s)) \leq f^{\text{Step1}}(d(x^{\text{Step1}}, s)) \cdot \tilde{s} \quad (6)$$

where f^{Step1} is the objective function in Step-1, $d(x^{\text{Step1}}, s)$ and $d(x^{\text{Step2}}, s)$ are the dose distribution for scenario s in Step-1 and 2, respectively, and $\tilde{s}$ is the slip factor ($\tilde{s} = 1.3$ in our implementation).

ECHO uses quadratic objective functions for the PTV/CTV and convex gEUD functions for the OARs. The gEUD parameters depend on the clinical criteria (e.g., gEUD(1) or gEUD(2) for parallel structures and gEUD(20) for serial structures). Table 1 lists the hard maximum/mean dose constraints as well as the objective functions used in this study for the HN patients. Since we do not have any clinical DVC criteria on these patients, we will later add an artificial DVC to test the performance of the proposed algorithm.

For DVCs, we need to add the convex relaxation version of the mixed integer programming problem into Step-0 for each scenario to encourage the optimizer to satisfy the DVCs. Let us assume we have a DVC of $V(d) \%$ that we would like to satisfy in all the scenarios. Then we have:

$$d_i(x, s) \leq d + (D_{\max} - d) \cdot b_i(s) \quad (7)$$

$$\frac{1}{n} \sum_i b_i(s) \leq v/100 \quad (8)$$

where D_{max} is the maximum dose allowed for the structure, and n is the number of voxels in the structure. If $b_i(s)$ is only allowed to be zero or one, then Eq. (7) and (8) represent the exact implementation of DVC using mixed integer programming. In that case, $b_i(s)$ will be equal to 1 if $d_i(x,s)$ is larger than the threshold d and then Eq. (8) restricts the number of voxels violating the threshold, leading to DVC satisfaction. Given the computational challenges of solving mixed integer programming problems, the convex relaxation of (7) and (8) is used by allowing $b_i(s)$ to take any value between zero and one which would in turn encourage but not necessarily satisfy the DVC.

After solving Step-0, the sub-threshold voxels (i.e., $(100-v)\%$ of the low dose voxels) corresponding to each scenario are identified and constrained in the next steps to satisfy the DVC:

$$d_i(x,s) \leq d \text{ for } i \in V_s$$

where V_s is the set of sub-threshold voxels for scenario s . It should be noted that the set of sub-threshold voxels could differ for different scenarios as illustrated in Figure 1 with a DVC of $V(10 \text{ Gy}) < 25\%$ on the mandible. After solving Step-0, the 75% of the low dose voxels (sub-threshold voxels) for each error scenario are identified to be constrained by the DVC threshold of $d=10 \text{ Gy}$ using the regular maximum dose constraints in Step-1 and 2. As depicted in Figure 1, while a lot of voxels are common among all the error scenarios (large pink points), there are some voxels that only belong to the sub-threshold voxels for one or a few error scenarios. The small red points do not belong to any sub-threshold voxel set and are therefore only constrained by the mandible maximum dose in the optimization.

2.4 Patient population and computational frameworks

All plans, robust and non-robust, are optimized with ECHO, which uses the commercial optimization engine KNITRO™ (Artelys Corp., Chicago, IL) that can handle general constrained non-linear optimization problems. Dose calculations are performed in MATLAB (The MathWorks, Inc., Natick, MA) with the standard version of the open-source dose calculation software matRad [15,16] with pencil beam dose calculation algorithm. The proton spots are placed on a rectangular grid and cover the PTV region plus 1 mm. In this study, we simulate range over- and under-shoots by rescaling the stopping power ratio (SPR) image by $\pm 3.5\%$, following typical range margin recipes used in proton treatment planning [17,18]. Setup errors are simulated by shifting the isocenter in either the x, y, or z direction by $\pm 3 \text{ mm}$. Following the recommendation by Kraan *et al.* [19], we combined the range and setup errors, resulting in thirteen error scenarios. The dose influence matrix for each of the thirteen scenarios is pre-computed and stored to accurately calculate the dose distribution for each scenario in the optimization process. Random and fractionation errors are not included in this study; however, they can be in principle incorporated [20].

Three HN cancer patients from The Cancer Imaging Archive (TCIA) [21,22] plus a water phantom are included in this study. Figure 2 shows a representative CT slice for each patient. For the HN patients, only the primary CTV is used in our optimizations. For the non-robust optimizations, a PTV is created by 5 mm expansion of CTV to account for 3 mm setup error plus range uncertainty. The water phantom has been commonly used in the previous studies [4,23,24] since the horse-shoe shaped target wraps closely around the OAR and can highlight the benefits of the robust optimization specially in the presence of the over- and under-shoot uncertainties (a 3D version of the phantom is used here). The objective functions and constraints for the three HN cases are described in Table 1. For the phantom, the prescription dose is also 70 Gy and the objective function and maximum dose constraints for the OAR are gEUD(20) and 54 Gy, respectively. Table 2 summarizes the beam configurations and the number of spots and voxels used in the optimization for each patient.

3. Results

3.1 Comparison of robust and non-robust plans

Figure 3 compares the non-robust plans generated using the original version of ECHO and the robust plans generated with the proposed robustified version of ECHO using the p -norm approach with $p=2$, for all four patients. The non-robust plans are optimized solely based on the nominal scenario and using the PTV as the target while the robust plans are optimized including all thirteen scenarios and using the CTV as the target. The dose volume histogram (DVH) bands represent the DVH corresponding to all scenarios and the solid DVH line represents the nominal scenario. To facilitate the visual comparisons, only one representative OAR for each patient is shown for which more noticeable differences between robust and non-robust plans are observed. It is obvious in Figure 3 that the robust plans have much tighter DVH bands (red DVH bands) as compared to the non-robust plans (blue DVH bands), as expected. The hard maximum/mean dose criteria are also always met in all scenarios for the robust plans while this may not be true for some scenarios of the non-robust plans, as can be seen in Figure 3H where the maximum dose criterion is marked. The robust plan could although compromise the DVH in the nominal scenario in order to improve other scenarios and that is usually referred to as the *price of robustness*. This can be easily seen in Figure 3C and 3G by comparing the DVHs for the nominal scenario for the plans with and without robustness represented by the solid red and blue lines, respectively. Given that ECHO optimizes the CTV first and OAR secondly, most of the robustness impacts can only be observed in the CTV.

3.2 Stochastic and worst-case approaches compared to the p -norm approach

Figure 4–7 compare the DVH of the plans generated by the worst-case approach against the plans generated using the p -norm approach with different $p \in \{1, 2, 10\}$ for all four patients. The p -norm approach with $p=1$ corresponds to the stochastic approach (equal weights for all scenarios); however, Eq. (1) is used for the worst-case approach. The solid and dashed lines represent the DVH for the nominal and worst scenario, respectively. Here the worst scenario is the scenario with the highest value of the objective function (see Table 1) as judged based on the CTV objective function. The differences are more pronounced on the

CTV than the OAR since ECHO optimizes the CTV first, whereby the search space of the feasible solution is smaller for the OAR optimization. The differences in the results for the OARs are therefore mainly determined by the results for the CTV in Step-1, rather than by the different robust optimization approaches. The results for all patients follow the same trend: The p -norm approach improves the overall DVH band on the CTV (i.e., closer to the prescription dose) and compromises the worst scenario (dashed line) compared to the worst-case approach. As the p -value increases, the results for the p -norm approach move toward the results for the worst-case approach; however, since the p -norm approach always optimizes all the scenarios, it is able to improve the other scenarios at almost no cost (patients 1, 3, and the phantom) or a slight cost (patient 2) for the worst scenario. It should be noted that the deterioration of the CTV coverage in Step-1 of ECHO using the worst-case approach could potentially result in an enlarged search space in the Step-2 which then provides more opportunity to improve the OAR doses and that is why the OAR dose sometimes can be better using the worst-case approach (patient 1, 3, and the phantom). This is not a trait of different robust optimization approaches, but rather a consequence of having a hierarchical optimization scheme with the target coverage optimization in the first step.

For patient 2, the CTV dose distribution in the worst-case approach (blue band in Figure 5A–C) is considerably degraded compared to the p -norm approaches. This is mainly due to the more complicated geometry of this patient, where the larynx is in the close vicinity of the CTV. When the dose distribution is shifted towards the larynx, the maximum dose constraint on larynx prevents the coverage for a small part of the CTV, turning this error scenario into the worst scenario which is the only scenario that the worst-case approach is optimizing. The p -norm approach, in contrast, includes all the scenarios and while it focuses more on the worst scenario, it still trades improvements in other scenarios at a small cost for the worst scenario.

While DVH bands used in Figures 4–7 provide useful information regarding plan robustness, they do not show the individual DVHs and therefore we cannot see how the plans are distributed in the bands. Depicting all individual DVHs, on the other hand, would clutter the figure and makes the visual comparison difficult. Figures 8 and 9 provide complementary information by depicting the CTV objective function (Figure 8) and PTV D98% metric (Figure 9) for all individual plans for each technique.

Figure 8 represents the CTV objective function values (quadratic dose deviation) for all the thirteen scenarios corresponding to the worst-case and p -norm approaches. The top of each boxplot corresponds to the worst scenario(s) and the nominal scenario is indicated by a red 'x'. The horizontal dashed line represents the CTV objective of the non-robust plan in the nominal scenario. As we move from the stochastic approach ($p=1$) to the worst-case approach, it can be obviously seen that the worst scenario is improving (the top of the boxplot is lowering), however, the median objective function of all scenarios (red horizontal line) is deteriorating (the red line is raising). For low p -values, there is a larger spread in the objective function values, but the median is closer to lower quantile, meaning that most scenarios have an objection function value close to the best objective function value, while only a single or a few scenarios have a high objective function value. On the contrary, the boxplot corresponding to the worst-case approach reveals that the median and the worst

objective functions are almost the same. This means that in the worst-case approach, many scenarios are almost as bad as the worst scenario(s) since no effort is made to improve them. Going from the worst-case approach to p -norm approach with $p=20$, the median objective function is improved by 15% at the negligible 3% cost in the worst scenario(s), on average for all four patients. The difference between the horizontal dashed line (non-robust plan objective value) and the 'x' mark (nominal scenario objective value in robust plan) demonstrates the price of robustness (i.e., the cost to the nominal scenario to improve other scenarios).

Figure 9 compares CTV D98% values for all the scenarios corresponding to different approaches. The bottom of boxplots and red horizontal lines correspond to the worst scenario and the median values, respectively. As we move from the stochastic approach ($p=1$) to the worst-case approach, it can be seen that, most of the time, the worst scenario(s) is improving (the bottom of the boxplot is raising), and the median value of all scenarios is deteriorating (the red line is lowering). The reason this is only true in most cases but not all in Figure 9, while it is true for all cases in Figure 8, is that Figure 8 shows the CTV quadratic objective function which is optimized, and although it correlates with the CTV D98%, it does not explicitly include that.

In terms of computational efficiency, all the tested robust optimization approaches performed similarly and we found no trend in the timing for the different approaches. For patient 1, the average optimization time was 36 minutes (27 to 44 minutes), for patient 2 it was 22 h (16.8 to 27.3 h), for patient 3 it was 28 h (22.5 to 32.5 h) and for the phantom 7 h (4.9 to 11.2 h).

3.3 Dose volume constraints (DVCs)

We do not have any DVCs in our clinic for HN patients and therefore we use some artificial DVCs on the mandible to test the proposed DVC algorithm. Figure 10 compares the robust plans for the p -norm approach ($p=2$) with (red) and without (blue) DVCs. The DVCs are denoted by 'x' and are chosen so that they are not already fulfilled in the original robust plans ($V(10\text{ Gy})\ 5\%$ for patient 1 and $V(10\text{ Gy})\ 25\%$ for patients 2, 3 and the phantom). As shown in Figure 9, the algorithm successfully meets the DVC for all scenarios and for all patients, but meeting the DVC constraints degrade the CTV coverage and homogeneity as expected.

It should be noted that we only consider DVCs on OARs (not on CTV coverage) so there is no conflict and infeasibility issue regarding DVCs. ECHO seeks for the best CTV coverage while meeting all DVCs on OARs. Very demanding DVCs could potentially result into a very bad CTV coverage (e.g., $V(10\text{ Gy})\ 0\%$ could result in no CTV coverage).

Discussion

In this study, we integrated robustness in our automated treatment planning system ECHO. The hard maximum/mean/DVC constraints are extended to include all the scenarios and the objective function based on the p -norm in our robust optimization approach encourages the CTV coverage and OAR doses to improve in all the scenarios. The p -norm parameter

p can be used to adjust the level of conservatism in robustness. A large p -value ($p \approx 20$) resembles the worst-case approach although without disregarding any scenario. We found the advantage of the p -norm approach over the worst-case approach is more pronounced on patients with more complicated geometries (patient 2 and the phantom) where the p -norm approach pushes all the scenarios to optimality at a negligible cost of the worst scenario(s).

The presented p -norm robust optimization approach can be implemented in any constrained or even unconstrained optimization framework since it does not introduce any constraint into the optimization problem. However, if one is only using linear objective functions and/or constraints, the p -norm approach destroys the linearity of the problem. Fredriksson has also introduced a robust optimization approach using the *minimax stochastic programming* technique to control the level of robustness in-between the stochastic and the worst-case approach [4]. This technique converts the objective functions into constraints and cannot be used in unconstrained optimization frameworks; however, it preserves the original form of the objective functions (e.g., linearity) and could be a method of choice if one is solving constrained linear optimization problems.

DVCs are inherently non-convex and could result in computationally intractable optimization problems. In this study, we initially solve an extra optimization problem including the convex relaxation of DVCs for all the scenarios. We subsequently identify the sub-threshold voxels for each scenario and impose maximum dose constraints on them in each of the following two steps. This approach ensures the DVCs satisfaction in all the scenarios and can also be integrated in any constrained optimization framework, including non-robust optimization [12].

It has previously been discussed how robustness can be incorporated in a prioritized optimization scheme similar to the ECHO framework [25]. In that study, the robustness is postponed to the last step of the hierarchical optimization, whereby it can be seen as a correction to the realized solution for the nominal scenario. While this approach could reduce the overall time as only one robust optimization is performed in contrast to our framework which includes robustness in both steps, it does not enforce the robustness and satisfaction of hard dose constraints for all the scenarios and rather encourages the robustness over the search space provided from the previous steps.

While hierarchical constrained optimization is a powerful tool to automate the treatment planning process, solving constrained optimization problems could be computationally expensive and incorporating robustness could further increase the computational burden. Future study is needed to optimize the computation time and also incorporate random and fractionation errors.

Conclusion

We have presented a new, fully automated treatment planning framework for robust proton treatment planning. We compared different levels of robustness settings, from the stochastic approach (represented by the p -norm approach with $p=1$) all the way to the worst-case approach. For all the approaches, the robustness is included in both the objective functions

and the constraints and the level of robustness can be easily controlled in the p -norm approach. All the hard maximum/mean/dose-volume constraints are fulfilled for all the scenarios without any parameter tweaking and the proposed technique maximizes the CTV coverage and minimizes the OAR doses in two steps. We conclude that this approach could fill an important gap in the need for automated proton planning.

Acknowledgements

This work was partially supported by MSK Cancer Center Support Grant/Core Grant from the NIH (P30 CA008748).

Bibliography

- [1]. Unkelbach J, Alber M, Bangert M, Bokrantz R, Chan TC, Deasy JO, Fredriksson A, Gorissens BL, van Herk M, Liu W, Mahmoudzadeh H, Nohadani O, Siebers JV, Witte M, Xu H. Robust radiotherapy planning. *Phys Med Biol* 2018;63:22TR02.
- [2]. Unkelbach J, Oelfke U. Incorporating organ movements in IMRT treatment planning for prostate cancer: minimizing uncertainties in the inverse planning process *Med Phys*. 2005;32:2471–83. [PubMed: 16193776]
- [3]. Fredriksson A, Forsgren A, Hårdemark B. Minimax optimization for handling range and setup uncertainties in proton therapy. *Med. Phys.* 2011;38:1672–1684. [PubMed: 21520880]
- [4]. Fredriksson A A characterization of robust radiation therapy treatment planning methods – from expected value to worst case optimization *Med Phys*. 2012;39:5169–81. [PubMed: 22894442]
- [5]. Choi B, Deasy JO. The generalized equivalent uniform dose function as a basis for intensity-modulated treatment planning. *Phys Med Biol*. 2002;47:3579–3589. [PubMed: 12433121]
- [6]. Chanyavanich V, Das SK, Lee WR, Lo JY. Knowledge-based IMRT treatment planning for prostate cancer. *Med Phys*. 2011;38:2515–2522. [PubMed: 21776786]
- [7]. Monz M, Küfer KH, Bortfeld TR, Thieke C. Pareto navigation—algorithmic foundation of interactive multi-criteria IMRT planning. *Phys Med Biol*. 2008;53:985–998. [PubMed: 18263953]
- [8]. Zarepisheh M, Hong L, Zhou Y, Hun Oh J, Mechalakos JG, Hunt MA, Mageras GS, Deasy JO. Automated intensity modulated treatment planning: The expedited constrained hierarchical optimization (ECHO) system. *Med Phys*. 2019;46:294–2954.
- [9]. Wilkens JJ, Alaly JR, Zakarian K, Thorstad WL, Deasy JO. IMRT treatment planning based on prioritizing prescription goals. *Phys Med Biol*. 2007;52:1675–1692. [PubMed: 17327656]
- [10]. Breedveld S, Storchi PRM, Heijmen BJM. The equivalence of multi-criteria methods for radiotherapy plan optimization. *Phys Med Biol*. 2009;54:7199–7209. [PubMed: 19920305]
- [11]. Jee KW, McShan DL, Fraass BA. Lexicographic ordering: intuitive multicriteria optimization for IMRT, *Phys Med Biol*. 2007;52:1845–1861. [PubMed: 17374915]
- [12]. Mukherjee S, Hong L, Deasy JO, Zarepisheh M. Integrating soft and hard dose-volume constraints into hierarchical constrained IMRT optimization. *Accepted Med Phys*. 2019.
- [13]. Fredriksson A, Bokrantz R. A critical evaluation of worst case optimization methods for robust intensity-modulated proton therapy planning *Med Phys*. 2014;41:081701. [PubMed: 25086511]
- [14]. Fredriksson and Bokrantz. The scenario-based generalization of radiation therapy margins. *Phys Med Biol*. 2016;61:2067–2082. [PubMed: 26895381]
- [15]. Wieser HP, Cisternas E, Wahl N, Ulrich S, Stadler A, Mescher H, Müller LR, Klinge T, Gabrys H, Burigo L, Mairani A, Ecker S, Ackermann B, Ellerbrock M, Parodi K, Jäkel O, Bangert M. Development of the open-source dose calculation and optimization toolkit matRad. *Med Phys*. 2017;44:2556–2568. [PubMed: 28370020]
- [16]. Wieser HP, Wahl N, Gabry HS, Müller LR, Pezzano G, Winter J, Ulrich S, Burigo L, Jäkel O, Bangert M. MatRad – An open-source treatment planning toolkit for educational purposes. *Med Phys Int* 2018;6:119–127.

- [17]. Paganetti H Range uncertainties in proton therapy and the role of Monte Carlo simulations. *Phys Med Biol* 2012;57:R99–R117. [PubMed: 22571913]
- [18]. Taasti VT, Bäumer C, Dahlgren CV, Deisher AJ, Ellerbrock M, Free J, Gora J, Kozera A, Lomax AJ, de Marzi L, Molinelli S, Teo BKK, Wohlfahrt P, Petersen JBB, Muren LP, Hansen DC, Richter C. Inter-centre variability of CT-based stopping-power prediction in particle therapy: survey-based evaluation. *Phys imag radiat oncol*. 2018;6:25–30.
- [19]. Kraan AC, van de Water S, Teguh DN, Al-Mamgani A, Madden T, Kooy HM, Heijmen BJM, Hoogeman MS. Dose uncertainties in IMPT for oropharyngeal cancer in the presence of anatomical, range, and setup errors. *Int J Radiat Oncol Biol Phys* 2013;87:888–896. [PubMed: 24351409]
- [20]. Tilly D, Ahnesjö A. Fast dose algorithm for generation of dose coverage probability for robustness analysis of fractionated radiotherapy. *Phys Med Biol*. 2015;60:5439–5454. [PubMed: 26118844]
- [21]. Clark K, Vendt B, Smith K, Freymann J, Kirby J, Koppel P, Moore S, Philips S, Maffitt D, Pringle M, Tarbox L, Prior F. The Cancer Imaging Archive (TCIA): Maintaining and Operating a Public Information Repository. *J Digit Imaging*. 2013;26:1045–1057. [PubMed: 23884657]
- [22]. Bejarano T, De Ornelas Couto M, Mihaylov IB. Head-and-neck squamous cell carcinoma patients with CT taken during pre-treatment, mid-treatment, and post-treatment Dataset. The Cancer Imaging Archive (2018). 10.7937/K9/TCIA.2018.13upr2xf
- [23]. Bangert M, Hennig P, Oelfke U. Analytical probabilistic modeling for radiation therapy treatment planning. *Phys Med Biol*. 2013;58:5401–5419. [PubMed: 23877218]
- [24]. Unkelbach J, Chan TCY, Bortfeld T. Accounting for range uncertainties in the optimization of intensity modulated proton therapy. *Phys Med Biol*. 2007;52:2755–2773. [PubMed: 17473350]
- [25]. Xie Y Applications of nonlinear optimization. *PhD Thesis* Graduate School of Arts and Sciences, Washington University, St. Louis, USA (2014). 10.7936/K76W987X

Figure 1:

Visualization of the sub-threshold voxels identified from Step-0 for different scenarios for patient 2 with a DVC of $V(10\text{ Gy}) = 25\%$ on the mandible. The small red points do not belong to any sub-threshold set and are only restricted by the mandible maximum dose constraint. The large points belong to the sub-threshold set of at least one error scenario and are restricted by a maximum dose constraint of 10 Gy in Step-1 and 2. The number of scenarios that each point belongs to is color coded. The CTV points (light blue) are depicted only to show the geometry.

Figure 2:

Representative CT slices for each of the four patients. The red arrows mark the beam directions. The phantom only consists of water; the different colors do therefore not correspond to different CT numbers but merely represent the different structures.

Figure 3:

DVH comparison for robust (red, using the p -norm approach with $p=2$) and non-robust (blue) plans for all four patients. The solid lines represent the DVHs for the nominal scenario and the shaded areas represent the DVH band for all thirteen scenarios. The black vertical line represents the prescription dose. The maximum dose constraint of 54 Gy on OAR for the phantom is indicated in H to highlight the fact that non-robust plan violates this constraint in some scenarios.

Figure 4:

(Patient 1): DVH comparison for the worst-case approach (blue) against the p -norm approach (red) with different values of the p -norm parameter p , note that $p=1$ corresponds to the stochastic approach. The solid lines represent the DVHs for the nominal scenario and the shaded areas represent the DVH band for all thirteen scenarios. The dashed lines represent the DVHs for the worst scenario, as judged based on the CTV function value.

Figure 6:

(Patient 3): DVH comparison for the worst-case approach (blue) against the p -norm approach (red). The solid lines represent the DVHs for the nominal scenario and the shaded areas represent the DVH band for all thirteen scenarios. The dashed lines represent the DVHs for worst scenario, as judged based on the CTV function value.

Figure 7:

(Phantom): DVH comparison for the worst-case approach (blue) against the p -norm approach (red). The solid lines represent the DVHs for the nominal scenario and the shaded areas represent the DVH band for all thirteen scenarios. The dashed lines represent the DVHs for the worst scenario, as judged based on the CTV objective function value.

Figure 8:

The CTV objective function comparisons for different robust optimization approaches. The boxplots represent the CTV objective functions associated to different scenarios and the nominal scenario is indicated by 'x'. The non-robust CTV objective function is indicated by the horizontal dashed line.

Figure 9:

Target coverage parameter, D98%. The box plots show the distribution of D98% over the thirteen scenarios for each of the robust optimization approaches. The red horizontal lines represent the median, and the red crosses represent the nominal scenario.

Figure 10:

Comparison of DVH curves with (red) and without (blue) DVC. For patient 2 (B), an insert is added to show the high dose region for the CTV, where the DVH curves for the optimization with DVC are degraded compared to the optimization without the DVC.

Table 1:

Objective functions and dose constraints for the target and OARs.

Organ	Objective function	Dose prescription / hard constraints
CTV (robust optimization)	Quadratic function	Prescription dose: 70 Gy Max dose constraint: 1.3×70 Gy
PTV (non-robust optimization)	Quadratic function	Prescription dose: 70 Gy Max dose constraint: 1.3×70 Gy
Mandible	gEUD $a = 20$ (approximating the maximum dose)	Max dose constraint: 70 Gy
Larynx	gEUD $a = 1$ (equals the mean dose)	Mean dose constraint: 20 Gy Max dose constraint: 1.05×70 Gy
Parotid (left and right)	gEUD $a = 1$	Mean dose constraint: 26 Gy Max dose constraint: 1.05×70 Gy
Constrictor muscles	gEUD $a = 1$	Mean dose constraint: 40 Gy Max dose constraint: 1.05×70 Gy
Brainstem	gEUD $a = 20$	Max dose constraint: 54 Gy
Spinal cord	gEUD $a = 2$	Max dose constraint: 45 Gy
CTV/PTV ring 1 (2 mm)	gEUD $a = 2$ + gEUD $a = 20$	Max dose constraint: 1.1×70 Gy
CTV/PTV ring 2 (20 mm)	gEUD $a = 1$ + gEUD $a = 2$	Max dose constraint: 70 Gy
Body minus ring 2	gEUD $a = 1$ + gEUD $a = 2$	Max dose constraint: 0.8×70 Gy

Table 2:

Patient information, including the beam configuration, the CTV and PTV volume, and the number of proton beam spots and voxels in the optimization.

	Patient 1	Patient 2	Patient 3	Water phantom
Beam configuration	220°, 320°	0°, 135°, 235°	0°, 135°, 235°	0°, 45°, 315°
CTV / PTV volume (cm ³)	2.7 / 12.4	59.1 / 107.5	72.4 / 129.9	40.1 / 72.6
Number of spots	707	5630	7034	3790
Total number of voxels for robust / non-robust optimization	53956 / 61024	93399 / 113167	118208 / 140661	35553 / 57439