Bi-level multi-criteria optimization to include linear energy transfer into proton treatment planning

Mara Schubert[1] and Katrin Teichert[1,*]

[1]Fraunhofer Institute for Industrial Mathematics (ITWM), Fraunhoferplatz 1, 67663 Kaiserslautern, Germany
[*]katrin.teichert@itwm.fraunhofer.de

September 30, 2023

Abstract

In proton therapy treatment planning, the aim is to ensure tumor control while sparing the various surrounding risk structures. The biological effect of the irradiation depends on both physical dose and linear energy transfer (LET). In order to include LET alongside physical dose in plan creation, we propose to formulate the proton treatment planning problem as a particularly structured multi-criteria bi-level optimization problem, which we call hierarchical.

We show that the hierarchical multi-criteria bi-level problem can be reduced to a standard multi-criteria optimization (MCO) problem employing a specific domination cone. As the unfavorable properties of this cone prohibit a direct application of standard MCO solution methods, we further illustrate how a more convenient approximate cone can be constructed.

Based on the found reduction to a standard MCO problem, we then describe a novel approach to calculate a Pareto front representation for the hierarchical multi-criteria bi-level problem. As a point of reference, we also discuss a second, more brute-force approach. We apply both approaches to a prostate and a head and neck case, and compare the results.

Short running title— Bi-level MCO for treatment planning
Keywords— bi-level optimization, multi-criteria optimization, decision making, nonlinear optimization, Pareto front approximation
Acknowledgements— This work was funded by DFG, grant number 441208898

1 Introduction

In external radiation therapy, a target area is irradiated from the outside in an attempt to destroy cancerous cells through radiation induced damage. At the same time, nearby risk structures must be spared as much as possible in order to avoid side effects. Finding the best balance between tumor control and risk sparing motivates the application of multi-criteria optimization (MCO) methods. Over the last two decades, MCO methodology has been successfully established for both photon and proton treatment planning [Küfer et al., 2002], [Küfer et al., 2003], [Monz et al., 2008], [Craft et al., 2012], [Breedveld et al., 2019]. The general idea behind these approaches is as follows. Firstly, a set of Pareto efficient plans is calculated that covers the range of clinically feasible treatment options (*plan creation*). Subsequently, the plan set can be interactively searched by the planner for the most preferable compromise (*decision making*).

Compared to photons, proton radiation is more complex, resulting in a less straightforward relationship between physical dose and biological effect. One way to more accurately reflect the biological effect of protons is to integrate linear energy transfer (LET) into the plan optimization, as higher LET leads to increased biological effectiveness [Paganetti and van Luijk, 2013], [Chaudhary et al., 2014]. A straightforward method to integrate LET into plan optimization - which we will apply as well - is to include objectives that evaluate the product LETxD of LET and physical dose [Unkelbach et al., 2016], [Bai et al., 2020], [Gu et al., 2021].

A complete integration of LET into multi-criteria proton treatment plan optimization has not yet been described. In [Giantsoudi et al., 2013], the authors depict a way to retroactively include LET into the decision making for an already optimized plan set, but do not consider LET in the plan creation.

In this work, we discuss how to properly include LET into multi-criteria proton treatment plan optimization. To this aim, we propose to use bi-level optimization as a mathematical framework. Let X denote the set of feasible treatment plans, let ϕ denote a vector-valued objective measuring the plan quality in terms of LETxD, and let F denote a vector-valued objective measuring the plan quality in terms of physical dose. Then the bi-level problem

$$\begin{aligned} \min_{x \in X} \quad & \phi(x) \\ s.t. \quad & F_i(x) \leq F_i(z^*), \quad i = 1, ..., n \cdot \\ & z^* \in \operatorname{argmin}\{F(z) : z \in X\} \end{aligned} \tag{1}$$

reflects our aim to improve LETxD without sacrificing aspects of the physical dose distribution. Note that if both ϕ and F were real-valued, solving (1) would be equivalent to a lexicographic approach. In the multi-criteria setting however, the problem has not been extensively studied.

In a previous work, Unkelbach et al. treated optimization of physical dose and LET as a two-step problem [Unkelbach et al., 2016]. After finding a desirable dose distribution, achieved objective values were incorporated in the constraints and - in a second step - LETxD in the organs at risk was minimized. However, a fixed scalarization was employed in both steps, resulting in a straightforward lexicographic procedure. In contrast, we integrate dose and LETxD objectives into a unitary multi-criterial problem formulation and then find a representation of the whole Pareto front.

One conceptual idea to solve (1) is the following: Find a domination cone and the corresponding induced order relation such that the solution set of a standard multi-criteria optimization problem under this ordering coincides with the solution set of the original bi-level problem. Then solve this standard MCO problem using established Pareto front approximation methods. In our case, the bi-level problem (1) is a multi-criteria optimization problem in both levels, but it has additional properties that are conductive to such a reformulation.

The first ones to follow such an approach were Fliege and Vincente in [Fliege and Vicente, 2006]. They reformulated specific bi-level problems - namely those with one objective in each level, and where the upper level constraints do not depend on the lower level variables - to bi-criteria problems, and they found an order relation such that the solutions of both problems were the same. To find a fitting cone, they approximated the original order relation with a weaker one such that the set of non-dominated points with regards to the approximate weaker relation was a subset of the original non-dominated points. Due to the properties of their cone, a linear scalarization was not necessarily sufficient to find the solutions of the multi-criteria problem. They proposed to use quadratic mappings instead.

In [Ivanenko and Plyasunov, 2008] the authors extended the ideas from [Fliege and Vicente, 2006] to problems where the upper level constraints depend on the lower level variables. They considered bi-level problems with a single objective in each level, thus the results are not directly applicable to our problem. They focused on relations and did not construct a cone, nor specified how the relation can be used for numerical optimization.

Ruuska, Miettinen and Wiecek [Ruuska et al., 2012] also generalized the approach from [Fliege and Vicente, 2006]. They considered multi-criteria problems in both levels, but again with the restriction that the feasible set of the upper variable does not depend on the lower level variable. They obtained a relation and a standard MCO problem whose efficient solutions coincide with those of the original bi-level problem. However, the obtained relation was not transitive. They proposed a weaker one instead and showed that there are cases where this weaker relation is transitive, but this does not hold for the general case. A generating cone was not found, and they did not elaborate on how to efficiently use the relation in numerical optimization routines.

In contrast to the above, problem (1) motivates us to investigate bi-level problems where the lower level problem is independent of the upper level variable, but the feasible set of the upper level problem depends on the lower level function value. We will show that this special structure enables us to find a particular order relation and a generating cone to identify the solutions of the bi-level problem with those of an MCO problem. Additionally, the problem structure allows us to define an approximate cone that is polyhedral and whose dual has non-empty interior. This two properties can be exploited in numerical optimization routines, as we will demonstrate.

Our paper is organized as follows. In the theoretical part of our paper (Section 2), we formally introduce this particular type of bi-level problem and investigate its theoretical properties. In particular, we show how the problem can be reformulated into a standard MCO problem by defining a specific domination cone. We then discuss two algorithmic approaches to solve the problem. The first algorithm draws on the approach from [Unkelbach et al., 2016] for proton treatment planning, generalizing it to the multi-criteria setting. The second algorithm exploits the aforementioned reformulation to a standard MCO problem. In the results part (Section 3) we apply both algorithms to two real-world proton planning problems and compare the results.

2 Materials and Methods

2.1 Terms and Definitions

First, we will introduce some general notations and definitions that will be used throughout the paper. The interior of a set S is denoted by $int(S)$, and the boundary by $bnd(S)$. For two sets S_1 and S_2, $S_1 + S_2$ refers to their Minkowski sum, and for a point s and a set S the sum $s + S$ is the Minkowski sum between $\{s\}$ and S. Further, the dual of a cone C is denoted by C^*.

The following Definitions 2.1 to 2.4 as well as Lemma 1 introduce fundamental concepts from the theory of multi-criteria optimization.

Definition 2.1. *A multi-criteria optimization problem is an optimization problem with a vector valued objective function. It can be written as*

$$\begin{aligned} &\min\ F(x) = (F_1(x), ..., F_n(x)) \\ &s.t.\ x \in X \end{aligned} \tag{MCO}$$

where $F : X \to \mathbb{R}^n$ and $X \subseteq R^m$ is the feasible set.

If the objective functions $F_1, ..., F_n$ and the feasible set X are convex, the MCO problem is convex. We denote the image of F as $F(X) = \mathcal{Y}$. Typically, not all objective functions can take their minimal possible value at the same time. Instead, the best possible compromises between the objectives are considered the solutions to the problem.

To define a minima with respect to the vector-valued evaluation F, an order relation on $\mathcal{Y}$ must be defined that determines if a vector is smaller than another one. An order relation is a relation that is reflexive, anti-symmetric and transitive. A convex, pointed and salient cone C is called proper and induces an order relation. This conic order relation is defined as $x \preccurlyeq_C y \Leftrightarrow y - x \in C$. The cone is referred to as domination cone. A prominent domination cone is $\mathbb{R}^n_{\geq 0}$, also called the standard domination cone. For this cone $x \preccurlyeq_{\mathbb{R}^n_{\geq 0}} y \Leftrightarrow x_i \leq y_i \quad \forall i = 1, ..., n$. If $\mathbb{R}^n_{\geq 0}$ is employed, we denote the induced order relation simply as $x \leq y$. The concept can also be applied to matrices, where, if not defined otherwise, the relation is understood element-wise as well.

The following definition explains when a solution is considered a best possible compromise.

Definition 2.2. *[Ehrgott, 2005][Def 2.1] A point $x \in X$ is called Pareto efficient or Pareto optimal if there is no other $\tilde{x} \in X$ such that $F(\tilde{x}) \preccurlyeq_C F(x)$. The point $F(x)$ in the image space is called non-dominated. The set of non-dominated points is*

$$P_C = \{F(x) | x \in X, \nexists \tilde{x} \in X : F(\tilde{x}) \preccurlyeq_C F(x))\} \tag{2}$$

and referred to as Pareto front.

We will write $C - \min$ to refer to the operator for the non-dominated points of a MCO problem with regards to cone C. For the standard domination cone we will use min. A non-dominated point can be understood as a point such that no other point is equally good or better in all objectives. A weakly non-dominated point is one where no other point is truly better in all objectives.

Definition 2.3. *[Ehrgott, 2005][Def 2.24] A point $F(x)$ is weakly non-dominated if there is no other point $F(\tilde{x})$ such that $F(\tilde{x}) \in F(x) - int(C)$.*

One way to find the Pareto efficient solutions of an MCO problem (MCO) is to scalarize it, i.e. to transform it into a parametrized standard optimization problem with a real-valued objective. The scalarized problem can then be solved with standard optimization methods such as gradient descent.

Definition 2.4. *The weighted sum scalarization of* (MCO) *is*

$$\begin{aligned} \min \; & \sum_{i=1}^{n} w_i F_i(x) \\ s.t. \; & x \in X \end{aligned} \tag{WS}$$

with $w \in \mathbb{R}_{\geq 0} \setminus \{0\}$.

If the weights for (WS) are chosen from the dual cone C^*, one can solve (WS) to obtain a (weakly) non-dominated point of (MCO). The proof for the following lemma from [Serna, 2012] can be found in the appendix.

Lemma 1. *[Serna, 2012][Lemma 1.23] If the weights w of* (WS) *are from $C^* \setminus \{0\}$ and $\sum_{i=1}^{n} w_i F_i(x^*) = \min\{\sum_{i=1}^{n} w_i F_i(x) \mid x \in X\}$, then $F(x^*)$ is a weakly non-dominated point of* (MCO) *with regards to C. If $w \in int(C^*)$, $F(x^*)$ is a non-dominated point.*

Since commonly the Pareto front cannot be calculated explicitly, some Pareto points are calculated and used to approximate the whole set. For convex MCO problems, the convex hull of the calculated points of the Pareto front is used as inner approximation and the intersection of the positive half spaces of the hyperplanes supporting the Pareto points $\bigcap\{z : w^T z \geq \sum_{i=1}^{n} w_i F_i(x^*)\}$ as outer approximation of $P_C + C$. Since the true Pareto front is in between the inner and outer approximation, their distance can be used as upper threshold of the approximation error. Algorithms that use an inner and outer approximation are called sandwiching algorithms. Their aim is to find an approximation with as few solves as possible. For this they calculate the next weight to be used for (WS) such that the newly calculated point reduces the maximal approximation error as far as possible. For a pseudo code formulation of the sandwiching algorithm, see [Bokrantz and Forsgren, 2012] Algorithm 3.1.

In the remainder of this section, we introduce the bi-level problem as a special kind of optimization problem. An optimization problem is a bi-level problem if the constraints of the so-called upper level problem include an optimization problem, the so-called lower level problem. In particular, the upper level optimization variable x_u is a parameter in the lower level problem.

Definition 2.5. *A bi-level problem can be formulated as*

$$\begin{aligned} \min_{x_u \in X_u, x_l \in X_l} \quad & \phi(x_u, x_l) \\ s.t. \quad & x_l \in \operatorname{argmin}\{F(x_u, x_l) : G_i(x_u, x_l) \leq 0, i = 1, ..., I\} \\ & \psi_j(x_u, x_l) \leq 0, j = 1, ..., J. \end{aligned} \tag{BLP}$$

A bi-level problem can be a multi-criteria problem in either level. This means ϕ or F can be vector-valued. Then finding a Pareto point of the upper level requires two decisions, one for the lower level and one for the upper level once the lower level is decided. Depending on the setting there are several scenarios possible. Stemming from game theory, two are common in literature. One is where the lower level will decide so the upper level can find the best points. This is called the optimistic approach. The opposite, where the lower level is decided such that the upper level will get the worst result, is the pessimistic approach.

As motivated in the introduction, our aim is to solve the particular multi-criteria bi-level problem (1). This problem has certain special properties. Firstly, the upper level objectives ϕ are only dependent on x_u and not on x_l. Secondly, the upper level constraints ψ are only dependent on the lower level objectives values, not on the variables themselves. We call bi-level problems with these properties hierarchical. We will refer to the domination cone for the lower level problem of (HBP) as C_L and for the one of the upper level as C_U. A hierarchical bi-level problem can be formalized as

$$\begin{aligned} C_U - \min_{x_u} \quad & \phi(x_u) \\ s.t. \quad & \psi_j(x_u, F(x_l)) \leq 0, j = 1, ..., J \\ & x_l \in C_L - \underset{z}{\operatorname{argmin}}\{F(z) : G_i(z) \leq 0, i = 1, ..., I\}. \end{aligned} \tag{HBP}$$

Since the function value of the lower level is deterministic for the feasible set of the upper one, not the chosen efficient point itself, a hierarchical bi-level problem will automatically lead to the optimistic approach.

2.2 Theoretical properties of the hierarchical bi-level problem

Following the concept originally introduced in [Fliege and Vicente, 2006] and discussed in the introduction, we may try to reduce the hierarchical bi-level problem (HBP) to a standard MCO problem by finding a suitable domination cone. In the following, we introduce such a problem reformulation, and we show that all Pareto efficient solutions of this reformulation can be identified with a Pareto efficient solution of (HBP) and the corresponding solution of the lower level problem. We call such a reformulation a *single level reduction* of the hierarchical bi-level problem. The single level reduction uses the objectives and constraints from both levels of (HBP) as congenial objectives and constraints, respectively.

Definition 2.6. *The single level reduction of* (HBP) *is*

$$\begin{aligned} C_{SLR} - \min_{x_u, x_l} \quad & \phi(x_u), F(x_l) \\ s.t. \quad & \psi_j(x_u, F(x_l)) \leq 0, \ j = 1, ..., J \\ & G_i(x_l) \leq 0, \ i = 1, ..., I. \end{aligned} \tag{SLR}$$

To achieve an identification of the Pareto front of (HBP) and (SLR) we define the cone for the single level reduction as follows.

Definition 2.7. *The single level reduction cone is*

$$C_{SLR}(C_U, C_L) = \{(u, l) \in \mathbb{R}^{n_1+n_2} : l \in C_L \setminus \{0\} \vee (l = 0, u \in C_U)\}. \tag{3}$$

If the used cones are clear, we will simply write C_{SLR}. The properties required for a domination cone to ensure that it induces an order relation mostly transfer from C_U and C_L to C_{SLR}.

Lemma 2. *If C_U and C_L are pointed/salient cones, then $C_{SLR}(C_U, C_L)$ is also pointed/salient. If C_U and $C_L \setminus \{0\}$ are convex cones, then $C_{SLR}(C_U, C_L)$ is also convex.*

The proof can be found in the appendix. Notably, the requirement that $C_L \setminus \{0\}$ is convex implies the convexity of C_L. Henceforth, it will be assumed that both C_L and C_U fulfill the perquisites of Lemma 2. The following theorem shows how the Pareto front of (HBP) with regards to C_L and C_U is connected to the Pareto front of (SLR) with regards to $C_{SLR}(C_U, C_L)$.

Theorem 2.1. *Consider a* (HBP) *with regards to pointed, salient, convex domination cones and for the cone of the lower level it additionally holds that $C_L \setminus \{0\}$ is convex.*

i Then, for any non-dominated point $(\phi, F)(x_u^*, x_l^*)$ of (SLR) *with regard to the single level reduction cone, $\phi(x_u^*)$ is a point of the Pareto front of* (HBP) *that is obtained if the chosen solution of the lower level is $F(x_l^*)$, as long as for any x_l^* the feasible region of the upper level of HBP is not empty.*

ii For any non-dominated point $\phi(x_u^*)$ of (HBP) *with the lower level solution $F(x_l^*)$, $(\phi, F)(x_u^*, x_l^*)$ is a point of the Pareto front of the corresponding single level reduction with regards to the single level reduction cone.*

Proof. Denote the Pareto front of (HBP) with regards to C_L and C_U as

$$P_{HBP} = \Big\{\phi(x_u) : \psi(x_u, F(x_l)) \le 0, \tag{4a}$$

$$x_l \in C_L - \underset{z}{\operatorname{argmin}}\{F(z) : G(z) \le 0\}, \tag{4b}$$

$$\nexists y_u : \phi(y_u) \preccurlyeq_{C_U} \phi(x_u), \phi_i(y_u) < \phi_i(x_u) \text{ for some i}, \tag{4c}$$

$$\psi(y_u, F(x_l)) \le 0\Big\}. \tag{4d}$$

The Pareto front of (SLR) with regards to C_{SLR} is

$$P_{SLR} = \left\{\begin{pmatrix} \phi(x_u) \\ F(x_l) \end{pmatrix} : G(x_l) \le 0, \psi(x_u, F(x_l)) \le 0, \right. \tag{5a}$$

$$\nexists (y_u, y_l) : G(y_l) \le 0, \psi(y_u, F(y_l)) \le 0, \tag{5b}$$

$$\Big[\Big(F(y_l) \preccurlyeq_{C_L \setminus \{0\}} F(x_l)\Big) \tag{5c}$$

$$\vee \Big(F(y_l) = F(x_l), \phi(y_u) \preccurlyeq_{C_U} \phi(x_u), \tag{5d}$$

$$\left. \exists j : \phi_j(y_u) < \phi_j(x_u)\Big)\Big]\right\}.$$

First, to show [i], let $\begin{pmatrix} \phi(x_u^*) \\ F(x_l^*) \end{pmatrix} \in P_{SLR}$. From (5a) one obtains that x_l^* fulfills the constraints of the lower level problem. (5a) to (5c) mean that for a fixed u, x_l is an efficient point of

$$C_L - \min_x\{F(x) : G(x) \le 0, \psi(u, F(x)) \le 0\}. \tag{6}$$

Contrary, for P_{HBP} x_l is part of a vector from the set

$$\{(u, x_l) : x_l \in C_L - \underset{z}{\operatorname{argmin}}\{F(z) : G(z) \leq 0\}\} \cap \{(u, x_l) : \psi(u, F(x_l)) \leq 0\}. \quad (7)$$

Clearly, the feasible region of (6) is a subset of the feasible region of the optimization problem in (7) and the minima of (7) is smaller or equal than the one of (6). Hence, if a point is in the set described by (7), it is also in (6). But the desired direction is the other way round. Assume that a point (u, x) fulfills (6) but not (7). Then there is a Pareto optimal point $F(y_l)$ of (7) that would dominate the optimal solution of (6) but does not fulfill the constraints. By assumption any solution x_l of the lower level problem leads to a non empty feasible region of the upper level. This means that there is a point (y_u, y_l) that fulfills (7). As reasoned above, this point is also in (6). This means there exists a (y_u, y_l) that was forbidden in (5b) and (5c).

Next, set (y_u, y_l) in (5b) to (y_u, x_l^*). Hence, there cannot be a y_u that fulfills $\psi(y_u, F(x_l^*)) \leq 0$ and (5d). This implies (4c) and (4d). Consequently, $\phi(x_u^*)$ is in P_{HBP} and the lower level solution chosen for this point is $F(x_l^*)$.

For the other direction assume that $\phi(x_u^*) \in P_{HBP}$ and the corresponding chosen lower level solution is x_l^*. Then (4a) and (4b) imply (5a). Next, assume that there is a (y_u, y_l) such that (5b) and (5c) hold. Then $G(y_l) \leq 0$ and $F(y_l) \preccurlyeq_{C_L \setminus \{0\}} F(x_l^*)$ contradict the assumption that x_l^* is an efficient point of the lower level problem, thus such a (y_u, y_l) cannot exists. Next, we have to exclude the existence of a point fulfilling (5b) and (5d). For $y_l = x_l^*$ this follows directly from (4c) and (4d). Assume there is an $y_l \neq x_l^*$ such that $F(y_l) = F(x_l^*)$. There cannot be an y_u such that (4c) and (4d) hold. If (5d) is violated so is (4c). Violation of $\psi(y_u, F(x_l^*)) \leq 0$ also means violation of $\psi(y_u, F(y_l)) \leq 0$ and thus also transfers directly to (5b). $\square$

This theorem relies on the specific structure of the hierarchical bi-level problem. Especially the last step requires that the lower level variable does only occur in the upper level constraints as evaluation of the lower level objective. Otherwise it could happen that there is a non-dominated point of the lower level problem that belongs to more than one efficient solution of the lower level problem. Then P_{HBP} will include the upper level Pareto point with regard to whichever point was set in the lower level, or if all lower level Pareto points are considered it will include the upper Pareto points belonging to all of them - even if in the upper level one dominates the other and the lower levels are the same. On the other hand, in P_{SLR} only the point that leads to the better upper level value would be included. The independence required in (HBP) prevents this case from happening. Efficient points of the lower level that belong to the same non-dominated point will result in the same non-dominated points in the upper level.

For Theorem 2.1[i] we required the set (7) to be nonempty for any efficient point x_l of the lower level problem. If this is not the case, P_{HBP} would not include a solution for these x_l, whereas P_{SLR} would include a point that is not Pareto optimal in the sense of the lower level problem, but the best that can be found with the additional constraints. As long as the constraints of the upper level only ensure that the lower level objectives evaluated at x_u are close enough to $F(x_l)$ the assumption will always hold since $x_u = x_l$ would be feasible.

To calculate a Pareto optimal point of a MCO problem a common method is to solve a corresponding weighted sum problem with weights from the dual cone of the domination cone. The obtained solution is weakly non-dominated (see Lemma 1). Unfortunately, due to the construction of C_{SLR} the set of weakly non-dominated points is far too large to be usable as solution of the optimization problem, since all points that are weakly non-dominated in the lower level are weakly non-dominated with regards to C_{SLR}, independent of their upper level value. To distinguish points that are the same in the lower level but different in the upper level is the conceptual idea behind the hierarchical bi-level problem. Hence, the weakly non-dominated solutions are insufficient for all use cases. The points of interest are the non-dominated ones. However, as the interior of C^*_{SLR} is empty, there are no weights in C^*_{SLR} that guarantee true Pareto optimality with regards to C_{SLR}. Consequently, the cone is unfit for calculations. C_{SLR} can be reformulated as

$$C_{SLR} = (C_L \setminus \{0\} \times \mathbb{R}^{n_2}) \cup (\{0\} \times C_U)\,. \tag{8}$$

If C_L and C_U are the standard domination cone, one obtains

$$\begin{aligned} C_{SLR}(\mathbb{R}^{n_1}_{\geq 0}, R^{n_2}_{\geq 0}) &= \left(\mathbb{R}^{n_1}_{\geq 0} \setminus \{0\} \times \mathbb{R}^{n_2}\right) \cup \left(\{0\} \times R^{n_2}_{\geq 0}\right) \\ &= \left(\mathbb{R}^{n_1}_{\geq 0} \times \mathbb{R}^{n_2}\right) \setminus \left(\{0\} \times (\mathbb{R}^{n_2}_{\geq 0})'\right), \end{aligned} \tag{9}$$

where $(\mathbb{R}^{n_2}_{\geq 0})'$ denotes the complement of $\mathbb{R}^{n_2}_{\geq 0}$ in $\mathbb{R}^{n_2}$. To get more non-dominated points we need a smaller approximation cone. The idea of the approximation cone is to exclude slightly more than $\left(\{0\} \times (\mathbb{R}^{n_2}_{\geq 0})'\right)$.

Definition 2.8. *Let I^n be a unit matrix of Dimension $n \times n$ and $\mathbf{0}^{n\times m}$ a matrix with only zeros of size $n \times m$ and let L describe the dimension of the lower level problem, i.e. $F(x_l) \in \mathbb{R}^L$, and U the dimension of the upper level problem. Then the approximation cone $C_{B(\mathcal{E})}$ of $C_{SLR}(\mathbb{R}^U_{\geq 0}, \mathbb{R}^L_{\geq 0})$ is the set*

$$C_{B(\mathcal{E})} = \{Q\lambda,\ \lambda \geq 0\}, \tag{10}$$

with

$$Q = \begin{pmatrix} I^U & \mathbf{0}^{U\times L} & -I^U \\ \mathbf{0}^{L\times U} & I^L & \mathcal{E} \end{pmatrix}, \tag{11}$$

where $\mathcal{E} \in \mathbb{R}^{L\times U}_{>0}$ is a matrix with entries $\epsilon_{i,j}$ that describe the required derivation in direction l_j for a negative step in direction u_i.

Notice how the blocks from Q stem from the lower and upper level domination cones, with the concept that a point where the upper level part is outside of C_U leads to costs in the lower level.

We will write λ as $(\lambda_u, \lambda_l, \lambda_B) \in \mathbb{R}^U_{\geq 0} \times \mathbb{R}^L_{\geq 0} \times \mathbb{R}^U_{\geq 0}$ to refer to the entries of λ that are multiplied with the respective blocks of Q. This means a point in $C_{B(\mathcal{E})}$ can be written as

$$\begin{pmatrix} \lambda_{u_1} - \lambda_{B_1} \\ \vdots \\ \lambda_{u_U} - \lambda_{B_U} \\ \lambda_{l_1} + \sum_{k=1}^{U} \epsilon_{k1}\lambda_{B_k} \\ \vdots \\ \lambda_{l_L} + \sum_{k=1}^{U} \epsilon_{kL}\lambda_{B_k} \end{pmatrix}. \tag{12}$$

The next step is to show that for $\mathcal{E} \to \mathbf{0}$ the Pareto front $P_{C_{B(\mathcal{E})}}$ converges towards $P_{SLR(\mathbb{R}^U_{\geq 0}, \mathbb{R}^L_{\geq 0})}$. To do so let us define a_{min} for a point $x = (x_u, x_l)$ as

$$a_{min} = \inf \left\{ a \in \mathbb{R}_{\geq 0} : \begin{pmatrix} x_u \\ 0 \end{pmatrix} + a \begin{pmatrix} 0 \\ x_l \end{pmatrix} \in C_{B(\mathcal{E})} \right\}. \tag{13}$$

With a_{min} it is possible to construct an approximation cone such that an arbitrary fixed point that is not in P_{SLR} is not in P_{C_B}.

Lemma 3. *Let s be the mapping $s : \mathcal{Y} \to \mathcal{Y}_l$, $s : (u, l) \mapsto l$. Assume that for any $y \in bnd(\mathcal{Y} + C_{B(\mathcal{E})})$ there is no $b = (b_u, b_l) \in \mathbb{R}^{U+L}$ with $\exists i, j, k : b_{u_i} < 0,\ b_{l_j} > 0,\ b_{l_k} = 0$ such that*

$$s(y + b) \in bnd(\mathcal{Y}_L + C_L) \text{ and} \tag{14}$$

$$y + b \in bnd(\mathcal{Y} + C_{B(\mathcal{E})}). \tag{15}$$

Let $z \in P_{SLR}$ and $v \notin P_{SLR}$, $v \in P_{C_{B(\mathcal{E})}}$ be two points such that z dominates v with regards to C_{SLR}. Then for $x = v - z$ equation (13) has a solution $a_{min} < \infty$.

Proof. It is to show that there is an $a \in \mathbb{R}_{\geq 0}$ such that $\begin{pmatrix} x_u \\ a \cdot x_l \end{pmatrix} \in C_{B(\mathcal{E})}$. This means there has to be a $\lambda \geq 0$ such that

$$\begin{pmatrix} x_u \\ a \cdot x_l \end{pmatrix} = \begin{pmatrix} \lambda_{u_1} - \lambda_{B_1} \\ \vdots \\ \lambda_{u_U} - \lambda_{B_U} \\ \lambda_{l_1} + \sum_{k=1}^{U} \epsilon_{k1} \lambda_{B_k} \\ \vdots \\ \lambda_{l_L} + \sum_{k=1}^{U} \epsilon_{kL} \lambda_{B_k} \end{pmatrix}. \tag{16}$$

Since z dominates v with regards to C_{SLR}, the point x is in C_{SLR} and consequently $x_l \in C_L = \mathbb{R}^L_{\geq 0}$. If $x_l = 0$, it follows that $x_u \geq 0$. But then z also dominates v with regards to $C_{B(\mathcal{E})}$ and $v \notin P_{C_{B(\mathcal{E})}}$. Hence, $x_l \geq 0$ and $\exists j$ such that $x_{l_j} > 0$. Clearly, if $x_u \geq 0$, v would not be in $P_{C_{B(\mathcal{E})}}$. Assume there is an i such that $x_{u_i} < 0$. For $x_{u_k} \geq 0$ choose $\lambda_{u_k} = x_{u_k}$ and $\lambda_{B_k} = 0$. For $x_{u_k} < 0$ choose $\lambda_{u_k} = 0$ and $\lambda_{B_k} = -x_{u_k}$. For (16) to have a solution, it has to hold $\exists\ a \in \mathbb{R}_{\geq 0} :\ a \cdot x_{l_j} \geq \sum_{i=1}^{U} \epsilon_{kj} \lambda_{B_k}\ \forall j = 1, ..., L$. The only case where this might not be the case is if there exists j such that $x_{l_j} = 0$. This would contradict the prerequisites. $\square$

From now on we will assume that the prerequisites of Lemma 3 hold. Notably, Lemma 3 does not necessarily require (14). This prerequisite is there to not exclude too many cases, since it ensures that (15) only has to hold for points that have the possibility of beeing non-dominated. Clearly, if one would solve (13) for a point $v \notin P_{C_{B(\mathcal{E})}}$ because it is dominated by z, one would obtain $a_{min} \leq 1$. This can be used to construct an approximation cone such that v is not on the corresponding Pareto front.

Lemma 4. *Let z and v be two points as described in Lemma 3. Further, let a_{min} be the solution of (13) for $v - z$ with regards to $C_{B(\mathcal{E})}$. Then z dominates v with regards to $C_{B(\frac{\mathcal{E}}{a_{min}})}$.*

Proof. Clearly, $a_{min} > 1$. Further, $a_{min} \cdot x_{l_j} \geq \sum_{i=1}^{U} \epsilon_{kj} \lambda_{B_k} \forall j = 1, ..., L$. Thus, $x_{l_j} \geq \sum_{i=1}^{U} \frac{\epsilon_{kj}}{a_{min}} \lambda_{B_k} \forall j = 1, ..., L$ and consequently $v - z \in C_{B(\frac{\mathcal{E}}{a_{min}})}$. $\square$

We have shown that for any point that is on the Pareto front of an approximation cone, but not on the actual Pareto front, a bigger approximation cone can be found such that the point is not on its front anymore. This can be used to show that for $\mathcal{E} \rightarrow \mathbf{0}$ the approximated Pareto front converges towards the true Pareto front.

Theorem 2.2. *If* (SLR) *is a convex problem, for $\mathcal{E} \rightarrow \mathbf{0}$ the set $P_{C_{B(\mathcal{E})}}$ converges towards P_{SLR}.*

Proof. For a series of $\epsilon-$matrices $(\mathcal{E}_n)_{n=1}^{\infty}$ with $\mathcal{E}_n \geq \mathcal{E}_{n+1}$ it follows that $C_{B(\mathcal{E}_n)} \subseteq C_{B(\mathcal{E}_{n+1})}$ and thus $P_{B(\mathcal{E}_n)} \supseteq P_{B(\mathcal{E}_{n+1})}$. Denote $P_{B(\mathcal{E}_n)}$ as P_n and notice that $(P_n)_{n=1}^{\infty}$ is a monotonous decreasing sequence of sets. Such sequences always converge. To show that the limit is P_{SLR}, let $B_\delta(P_{SLR})$ for some $\delta > 0$ be an open δ-ball around P_{SLR}. Consider the set V of points $v \in P_n$, $v \notin B_\delta(P_{SLR})$ with $d_{Haus}(v, P_{SLR}) = \delta$, where d_{Haus} is the Hausdorff metric. For each of these points $a_{min}(v) < \infty$. Consequently, $k = \max_{v \in V} a_{min}(v) < \infty$ and can be used to scale the $\mathcal{E}$ matrix for the next step in the sequence. Choose $\mathcal{E}_{n+1} = \frac{\mathcal{E}_n}{k}$. Then P_{n+1} and all subsequent elements of the set sequence have a Hausdorff distance from P_{SLR} smaller than δ. This construction works for every arbitrary small $\delta > 0$. Thus, $\lim_{n \rightarrow \infty} P_n = P_{SLR}$. $\square$

While the weakly non-dominated points with regards to $C_{B(\mathcal{E})}$ still include all points that are weakly non-dominated in the lower level independent of the upper level, the interior of the dual cone $C^*_{B(\mathcal{E})}$ is not empty. Thus, contrary to C_{SLR} it is possible to calculate truly non-dominated points with $C_{B(\mathcal{E})}$. In the next section, we introduce the reduction approach as one of two methods to calculate the Pareto front of HBP. The reduction approach makes use of all theoretic findings presented in this section.

2.3 Solving the hierarchical bi-level problem

In this section, we introduce two different approaches to solve the multi-criteria hierarchical bi-level problem (HBP). To work properly, both approaches require that all functions defining the bi-level problem are convex. For the remainder of this section, we therefore require the following assumption to hold.

Assumption 2.1. *For the hierarchical bi-level problem (HBP), we assume that the lower level objectives F_k $(k = 1, ..., n^l)$, the lower level constraints G_i $(i = 1, ..., I)$, the upper level objectives ϕ_k $(k = 1, ..., n^u)$ and the upper level constraints ψ_j $(j = 1, ..., J)$ are all convex.*

The performance of these two methods will later be compared in the results section when applied to specific proton treatment planning examples. As a subroutine, both methods use the sandwiching algorithm described in [Serna, 2012] and [Bokrantz and Forsgren, 2012]. The pseudo code for the sandwiching algorithm can be found in [Bokrantz and Forsgren, 2012] Algorithm 3.1. For our purposes, we define a stopping criterion for the sandwiching algorithm that is met if either the approximation error reaches a given target value δ, or the number of computed solutions reaches a given maximum number N. We employ the following signature

$$X^* = Sandwiching(\mathcal{M}, \delta, N, C := \mathbb{R}^n_{\geq 0}) \tag{17}$$

where X^* is the set of calculated Pareto efficient solutions and $\mathcal{M}$ is the problem instance. C denotes the domination cone, which can be any proper cone. In the standard case, C is the positive orthant $\mathbb{R}^n_{\geq 0}$.

2.3.1 Two stage approach

The first algorithm can be seen as a generalization of the approach of [Unkelbach et al., 2016]. It operates in two stages. In the *first stage*, a set of Pareto efficient solutions $Z = z^1, ...z^N$ to the (multi-criterial) lower level problem

$$\begin{aligned} \min_{z} \quad & F(z) \\ s.t. \quad & G_i(z) \leq 0, i = 1, ..., I. \end{aligned} . \tag{18}$$

is calculated using the sandwiching algorithm. In the *second stage*, for each solution z^k from the first stage, the Pareto front of the problem

$$\begin{aligned} \min_{x_u} \quad & \phi(x_u) \\ s.t. \quad & \psi_j(x_u, F(z^k)) \leq 0, j = 1, ..., J \end{aligned} . \tag{19}$$

is computed, again using the sandwiching algorithm. Problem (19) is the upper level problem of (HBP), where the lower level variable x^l is fixed to z^k. In Algorithm 1 the pseudo code for the two-stage approach is depicted.

Algorithm 1 Two stage approach

---

input:
- a multi-criteria hierarchical bi-level problem (HBP)
- a target approximation quality $\delta_1 \geq 0$ and a maximal number of solutions $N_1 \geq 0$ (for first stage)
- a target approximation quality $\delta_2 \geq 0$ and a maximal number of solutions $N_2 \geq 0$ (for second stage)

output:
- a set of Pareto efficient solutions X^* to (HBP)

Start
  $Z = z^1, ..., z^N \leftarrow Sandwiching((18), \delta_1, N_1)$
  for all $z^k \in Z$ **do**
    $X^{*,k} \leftarrow Sandwiching((19), \delta_2, N_2)$
  end for
  return $\bigcup_{k=1}^{N} X^{*,k}$
End

2.3.2 Reduction approach

The theoretical findings of Section 2.2 motivate a second, potentially more efficient method for solving (HBP). As shown in Theorem 2.1, it is possible to reformulate the multi-criteria hierarchical bi-level problem (HBP) to the standard multi-criteria problem (MCO) by defining a specific domination cone. As the sandwiching algorithm can be applied using any proper cone [Serna, 2012][Bokrantz and Forsgren, 2012], this allows the direct application of the Sandwiching algorithm to the reformulated problem.

Unfortunately, while the cone C_{SLR} in the reformulation described in 2.1 is indeed proper, the dual cone of C_{SLR} has empty interior. As a result, the Sandwiching algorithm can - in general - only produce weakly Pareto-efficient solutions. Creating only weakly Pareto-efficient solutions is however not a satisfying outcome when optimizing in any real-world scenario.

Fortunately, Definition 2.8 describes a proper cone $C_{B(\mathcal{E})}$ that approximates $C_{SLR}(\mathbb{R}^{n_1}_{\geq 0}, \mathbb{R}^{n_2}_{\geq 0})$ and whose dual has non-empty interior. We can, therefore, apply the Sandwiching algorithm to the reformulation using the approximate proper cone $C_{B(\mathcal{E})}$ instead of $C_{SLR}(\mathbb{R}^{n_1}_{\geq 0}, \mathbb{R}^{n_2}_{\geq 0})$. Moreover, because of the convergence result from Theorem 2.2, we can expect to obtain a set of solutions which approximates the Pareto front of HBP reasonably well provided that the entries of $\mathcal{E}$ are chosen suitably small. We call this the *reduction approach*. In Algorithm 2 the pseudo code for the reduction approach is depicted.

Algorithm 2 Reduction approach

input:
- a multi-criteria hierarchical bi-level problem (HBP)
- a target approximation quality $\delta \geq 0$ and a maximal number of solutions $N \geq 0$

output:
- a set of feasible solutions X to (HBP) whose images approximate the Pareto front of (HBP)

Start

Reformulate (HBP) to (MCO) according to Theorem 2.1

Choose suitably small $\epsilon > 0$

Define the domination cone $C_{B(\mathcal{E})}$ of 2.8 with $\epsilon_{i,j} := \epsilon$

return $Sandwiching((MCO), \delta, N, C_{B(\mathcal{E})})$

End

3 Results

In this section, we apply the two algorithms from 2.3 - the two-stage approach and the reduction approach - to calculate the Pareto front for two realistic proton treatment planning problems. The problems were taken from *matRad* [Cisternas et al., 2015], an open source software for radiation treatment developed by the German Cancer Research Center. For both approaches, we employ the commercially available nonlinear optimization solver *knitro* [Byrd et al., 2006] to solve the weighted sum problems within the Sandwiching subroutine.

As motivated in the introduction, we conceptualize the treatment planning problem as a multi-criteria bi-level problem of the form (1), which is a special case of of the

multi-criteria hierarchical bi-level problem (HBP). Each upper level objective ϕ_i, each lower level objective F_i, and each constraint defining the feasible set X - in addition to the physical constraint of positive irradiation - employs one of the following evaluation functions to evaluate the dose distribution (physical or LETxD) in a given volume V, which can be a target or risk structure.
The *underdose*, penalizing any voxel dose below the reference value d^{ref}:

$$\text{UD}(d) = \sum_{v \in V} (\max\left\{0, d^{\text{ref}} - d_v\right\})^2. \tag{20}$$

The *overdose*, penalizing any voxel dose above the reference value

$$\text{OD}(d) = \sum_{v \in V} (\max\left\{0, d_v - d^{\text{ref}}\right\})^2. \tag{21}$$

The *generalized equivalent uniform dose* for an exponent $p \geq 1$:

$$\text{gEUD}(d) = \sum_{v \in V} d_v^p. \tag{22}$$

3.1 Prostate case

Our first example was a prostate case. The target was irradiated from two proton beams from opposing sides (90° and 270°). There is a primary tumor (PTV 68) with prescription dose of 68 Gy and a larger surrounding target (PTV 56) with prescription dose of 56 Gy. The relevant risk structures were the rectum, the bladder, and the left and right femoral head, see Figure 1.

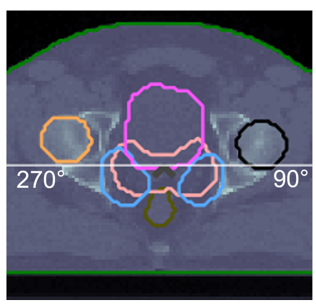

Figure 1: Geometry of the prostate case. PTV 56 (rose), PTV 68 (blue), bladder (pink), rectum (olive), left femoral head (black), right femoral head (orange).

We investigated an optimization model with two upper level objectives measuring the LETxD in bladder and rectum, and two lower level objectives measuring the physical dose in bladder and rectum. Target coverage, the sparing of the femoral heads and the unclassified tissue (Body) was ensured by constraints. The optimization model is depicted in Table 1.

We ran the two-stage approach with $N_1 = N_2 = 5$ and the reduction approach with $N = 25$. We also set all target approximation qualities to 0, such that for both approaches, the output set consisted of exactly 25 solutions. For the approximating cone in the reduction approach, we employed both $\epsilon = 0.1$ and $\epsilon = 0.05$.

volume	evaluation function	dose type
	upper level objectives	
Rectum	gEUD, $p = 2$	LETxD
Bladder	gEUD, $p = 2$	LETxD
	lower level objectives	
Rectum	gEUD, $p = 2$	physical
Bladder	gEUD, $p = 2$	physical
	lower level constraints	
PTV 68	UD, $d^{\text{ref}} = 68, \leq 10$	physical
PTV 68	OD, $d^{\text{ref}} = 73, \leq 5$	physical
PTV 56	UD, $d^{\text{ref}} = 56, \leq 10$	physical
PTV 56	OD, $d^{\text{ref}} = 61, \leq 5$	physical
Rt femoral head	gEUD, $p = 2, \leq 250$	physical
Lt femoral head	gEUD, $p = 2, \leq 250$	physical
Body	OD, $d^{\text{ref}} = 30, \leq 10$	physical

Table 1: The optimization problem for the prostate case.

The total computation time on a Lenovo T490s laptop was 9510 seconds for the two stage approach, while for the reduction approach the computation times were 5901 seconds ($\epsilon = 0.1$) and 5293 seconds ($\epsilon = 0.05$), respectively. One reason for the longer calculation time of the two stage approach is that in the first stage, 5 additional plans have to be optimized. Also, the individual optimization runs took longer on average for the two stage approach. One likely reason for this is that the tight constraints imposed on the second stage optimization problems make them more challenging to solve numerically. The calculation overhead on top of the individual optimization runs is negligible for both approaches and creates no meaningful difference in calculation time.

Figure 2 shows the objective space images of the calculated solution sets. On the left side, the points are projected on the lower level objectives evaluating the physical dose, while on the right side, their projections on the upper level objectives evaluating LETxD are displayed. Recall that the two-stage approach guarantees Pareto optimality with respect to the lower level objectives, while the points obtained from the reduction approach only create an approximation that improves with decreasing ϵ. This is confirmed in Figure 2, showing the lower level objective evaluations to be worse for the reduction approach. However, the difference between the point sets from the two approaches gets much smaller when employing $\epsilon = 0.05$ instead of $\epsilon = 0.1$.

On the other hand, the points obtained from the reduction approach reach better values for the upper level objectives, as can be seen in the right side plots of Figure 2. The reason for this is that for the reduction approach, but not the two-stage approach, some trade-off between lower and upper level objectives is allowed. Again, the difference between the point sets gets smaller when switching from $\epsilon = 0.1$ to $\epsilon = 0.05$.

Figure 3 shows the dose volume histogram and the dose averaged LET distribution

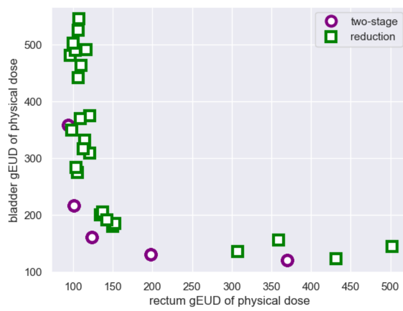

(a) $\epsilon = 0.1$: lower level

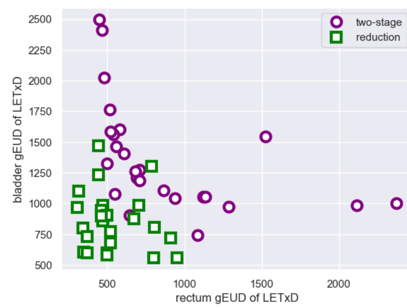

(b) $\epsilon = 0.1$: upper level

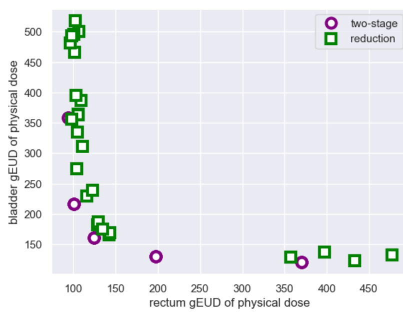

(c) $\epsilon = 0.05$: lower level

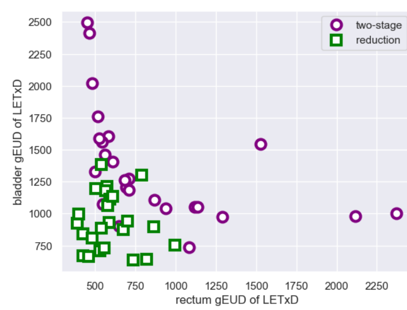

(d) $\epsilon = 0.05$: lower level

Figure 2: The objective space images of the solution sets obtained with the two-stage approach and the reduction approach for the prostate case, shown as projection views. Left: Projection on the lower level dose objectives for bladder and rectum. Right: Projection on the upper level LETxD objectives for bladder and rectum. The upper plots show the points obtained from the reduction approach for $\epsilon = 0.1$, the lower plots for $\epsilon = 0.05$. The points from the two-stage approach are the same in all plots.

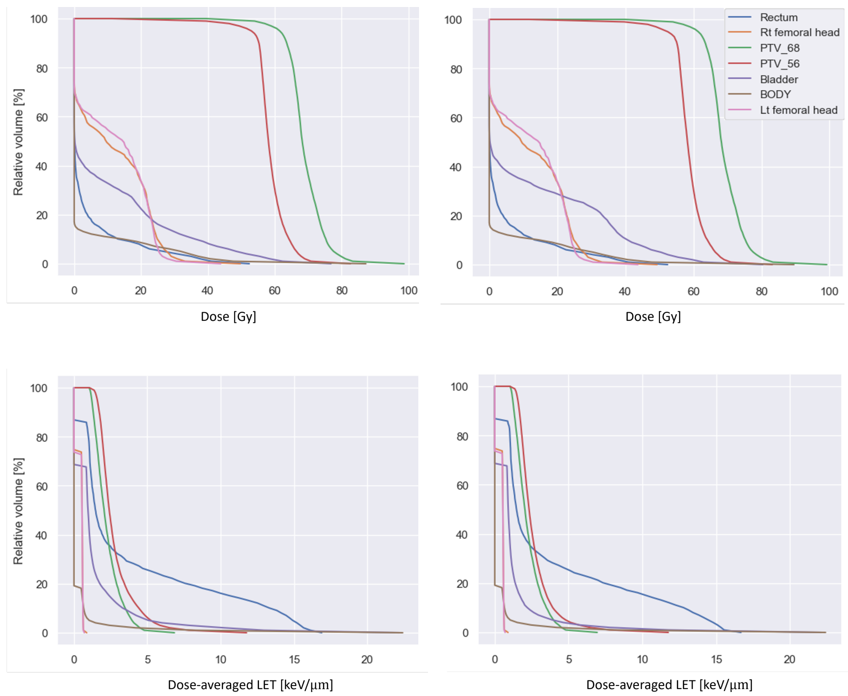

Figure 3: Two similar optimized plans for the prostate case, one obtained from the two-stage approach (left), one from the reduction approach (right). Top: Dose volume histogram showing, on the y-axis, the relative volume of a structure that receives at least the amount of dose on the x-axis. Bottom: Dose-averaged LET volume histogram, showing the relative volume of a structure that exhibits a dose averaged LET of at least the value on the x-axis.

for two comparable plans, one originating from the two-stage approach and one from the reduction approach.

3.2 Head and neck case

As a second example, we looked at a head and neck case with two beams from 45° and 315° respectively. There is a primary tumor (PTV 70) with prescription dose of 70 Gy and a larger surrounding target (PTV 63) with prescription dose of 63 Gy. The relevant risk structures are the left and right parotid, brain stem, cerebellum, spinal cord and larynx, see Figure 4.

Again, we investigated an optimization model with two upper level objectives, this time measuring LETxD in the left and right parotid, and two lower level objectives, measuring the physical dose in the left and right parotid. Target coverage and the sparing of brain stem, spinal cord, cerebellum and unclassified tissue (Skin) is ensured by constraints. The optimization model is given in Table 2.

Again, we ran the two-stage approach with $N_1 = N_2 = 5$ and the reduction

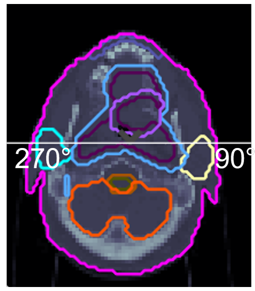

Figure 4: Geometry of the head and neck case: PTV 70 (pink), PTV 63 (blue), right parotid (light blue), left parotid (white), cerebellum (orange), brain stem (brown).

approach with $N = 25$ to obtain exactly 25 solutions with both approaches. For the approximating cone in the reduction approach, we employed both $\epsilon = 0.1$ and $\epsilon = 0.01$.

For the head and neck case, the total computation time was 14934 seconds for the two stage approach. Using the reduction approach, the calculation times were 8054 seconds ($\epsilon = 0.1$) and 5468 seconds ($\epsilon = 0.01$), respectively. Hence, as in the prostate case, the reduction approach was significantly faster than the two-stage approach.

Figure 5 shows the objective space images of the calculated solution sets as projections on the lower (left) and upper (right) level objectives. As in the prostate example, and corresponding to theory, the lower level objective evaluations are worse for the reduction approach while the upper level objective evaluations are better. As expected, the difference between the point sets shrinks significantly when switching from $\epsilon = 0.1$ to $\epsilon = 0.01$.

Figure 6 shows the dose volume histogram and the dose averaged LET distribution for two comparable plans obtained from the two different approaches.

4 Discussion

In this work, we described a concise and novel way of including LET in the plan optimization for proton treatment planning. To this aim, we defined a particular type of bi-level optimization problem which we called a hierarchical bi-level problem. We showed how this type of problem can be reduced to a standard multi-criteria problem, and how this reduction can be utilized to calculate a Pareto front representation. For two examplary proton planning problems, we observed a significant advantage in calculation time for this reduction approach compared to a second, more brute-force approach.

For hierarchical bi-level problems with convex objectives and constraints, the most significant advantage of the reduction approach lies in the direct application of the Sandwiching algorithm. By its use of inner and outer approximation, the Sandwiching algorithm is uniquely effective in creating a good representation of the Pareto front with very few weighted sum optimizations. In this paper, we did not investigate how effectively the calculated points represent the true Pareto front, as this would be have been a computationally intractable task for the rather large exemplary problems we discussed. However, as a future research topic, a computational study on a set of benchmark problems with analytically characterized Pareto fronts could provide this insight.

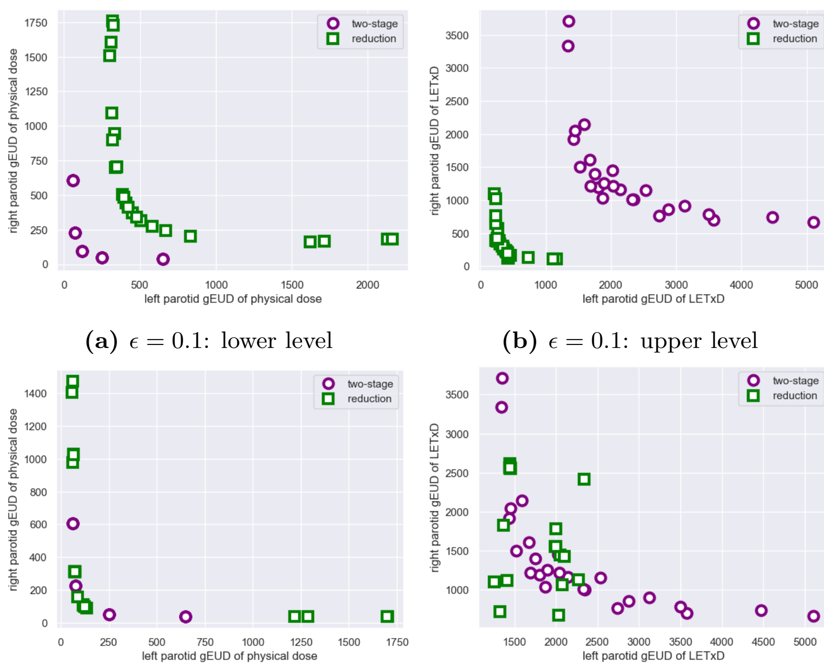

(a) $\epsilon = 0.1$: lower level

(b) $\epsilon = 0.1$: upper level

(c) $\epsilon = 0.01$: lower level

(d) $\epsilon = 0.01$: upper level

Figure 5: The objective space images of the solution sets obtained with the two-stage approach and the reduction approach for the head and neck case, shown as projection views. Left: Projection on the lower level dose objectives for left and right parotid. Right: Projection on the upper level LETxD objectives for left and right parotid. The upper plots show the points obtained from the reduction approach for $\epsilon = 0.1$, the lower plots for $\epsilon = 0.01$. The points from the two-stage approach are the same in all plots.

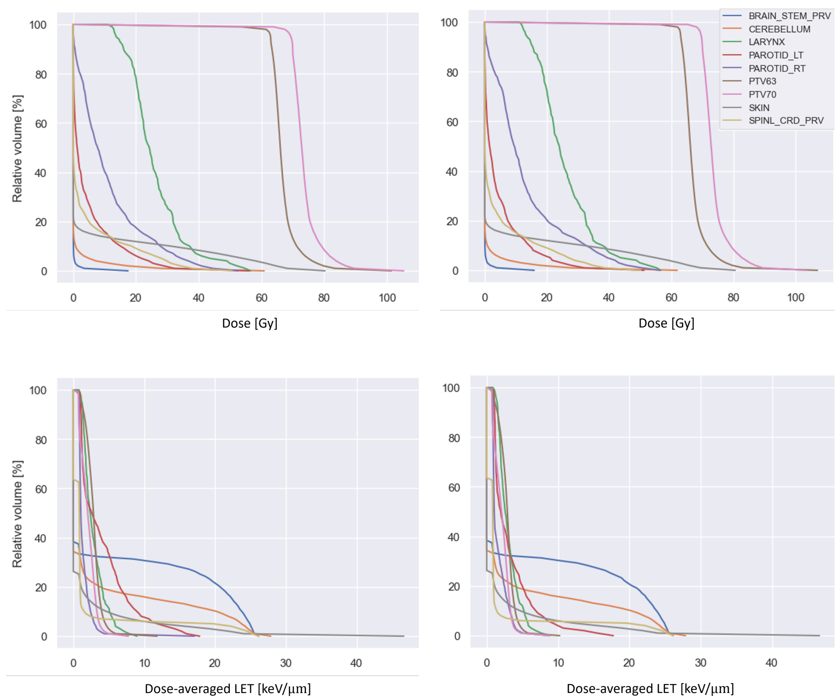

Figure 6: Two similar optimized plans for the head and neck case, one obtained from the two-stage approach (left), one from the reduction approach (right). Top: Dose volume histogram showing, on the y-axis, the relative volume of a structure that receives at least the amount of dose on the x-axis. Bottom: Dose-averaged LET volume histogram, showing the relative volume of a structure that exhibits a dose averaged LET of at least the value on the x-axis.

volume	evaluation function	dose type
	upper level objectives	
Left parotid	gEUD, $p = 2$	LETxD
Right parotid	gEUD, $p = 2$	LETxD
	lower level objectives	
Left parotid	gEUD, $p = 2$	physical
Right parotid	gEUD, $p = 2$	physical
	lower level constraints	
PTV 70	UD, $d^{\text{ref}} = 70, \leq 10$	physical
PTV 70	OD, $d^{\text{ref}} = 75, \leq 10$	physical
PTV 63	UD, $d^{\text{ref}} = 63, \leq 20$	physical
PTV 63	OD, $d^{\text{ref}} = 68, \leq 10$	physical
Skin	OD, $d^{\text{ref}} = 30, \leq 200$	physical
Skin	OD, $d^{\text{ref}} = 80, \leq 0$	physical
Brain stem	gEUD, $p = 2, \leq 10$	physical
Cerebellum	gEUD, $p = 2, \leq 25$	physical
Spinal cord	gEUD, $p = 2, \leq 100$	physical
Larynx	gEUD, $p = 2, \leq 750$	physical

Table 2: The optimization problem for the head and neck case.

An extension of the reduction approach to hierarchical problems with non-convex objectives and constraints is another interesting topic for research. Under these circumstances, the Sandwiching algorithm cannot be applied and would need to be replaced by an approximation algorithm suited for non-convex Pareto fronts, e.g. an epsilon constraint, hyperboxing or hypervolume method. The single level reduction itself, as presented in this work, does not require convex objectives nor constraints. However, convexity is indeed required to show the convergence of the Pareto fronts induced by a sequence of approximation cones towards the original front. Whether an approximation cone can be used in the case of non-convexity, and if so, how it can be defined, is a point of further investigation. For proton treatment plan optimization, an extension of the reduction approach to non-convex objectives and constraints would allow for a wider variety of dose evaluation functions, in particular the widely used dose-volume objective [Halabi et al., 2006].

Finally, we observe that the reduction approach can be extended to multi-level problems. Lemma 2 showed that the domination cone C_{SLR} needed for the single-level reduction retains many properties from the domination cone C_U of the upper level and the domination cone C_L of the lower level. If additionally $C_U \setminus \{0\}$ is convex, so is $C_{SLR} \setminus \{0\}$. This means that C_{SLR} has all the required properties of C_L. Hence, if all domination cones except the one of the highest level are convex even if $\{0\}$ is excluded, multi-level problems can recursively be reduced to a single-level problem with the presented approach. For this, repeatedly the two inner most levels are reduced and then form the new inner most level.

Competing interests

The authors have no competing interests to declare that are relevant to the content of this article. No funding was received to assist with the preparation of this manuscript.

References

[Bai et al., 2020] Bai, X., Lim, G., Grosshans, D., Mohan, R., and Cao, W. (2020). A biological effect-guided optimization approach using beam distal-edge avoidance for intensity-modulated proton therapy. *Medical physics*, 47(9):3816–3825.

[Bokrantz and Forsgren, 2012] Bokrantz, R. and Forsgren, A. (2012). An algorithm for approximating convex pareto surfaces based on dual techniques. *Informs Journal on Computing*, 25(2).

[Breedveld et al., 2019] Breedveld, S., Craft, D., Van Haveren, R., and Heijmen, B. (2019). Multi-criteria optimization and decision-making in radiotherapy. *European Journal of Operational Research*, 277(1):1–19.

[Byrd et al., 2006] Byrd, R. H., Nocedal, J., and Waltz, R. A. (2006). *Knitro: An Integrated Package for Nonlinear Optimization*, pages 35–59. Springer US.

[Chaudhary et al., 2014] Chaudhary, P., Marshall, T. I., Perozziello, F. M., Manti, L., Currell, F. J., Hanton, F., McMahon, S. J., Kavanagh, J. N., Cirrone, G. A. P., Romano, F., and Prise, K. M. (2014). Relative biological effectiveness variation along monoenergetic and modulated bragg peaks of a 62-mev therapeutic proton beam: a preclinical assessment. *International Journal of Radiation Oncology Biology Physics*, 90(1):27–35.

[Cisternas et al., 2015] Cisternas, E., Mairani, A., Ziegenhein, P., Jäkel, O., and Bangert, M. (2015). *matRad - a multi-modality open source 3D treatment planning toolkit*, pages 1608–1611. Springer International Publishing.

[Craft et al., 2012] Craft, D. L., Hong, T. S., Shih, H. A., and Bortfeld, T. R. (2012). Improved planning time and plan quality through multicriteria optimization for intensity-modulated radiotherapy. *Int. J. Rad. Onc. Biol. Phys.*, 82(1):e83–e90.

[Ehrgott, 2005] Ehrgott, M. (2005). *Multicriteria Optimization: With 88 figures and 12 tables*. Springer Berlin · Heidelberg, Berlin, Heidelberg, second edition edition.

[Fliege and Vicente, 2006] Fliege, J. and Vicente, L. N. (2006). Multicriteria approach to bilevel optimization. *Journal of Optimization Theory and Applications*, 131(2):209–225.

[Giantsoudi et al., 2013] Giantsoudi, D., Grassberger, C., Craft, D., Niemierko, A., Trofimov, A., and Paganetti, H. (2013). Linear energy transfer-guided optimization in intensity modulated proton therapy: feasibility study and clinical potential. *Int. J. Radiat. Oncol. Biol. Phys.*, 87:216–22.

[Gu et al., 2021] Gu, W., Ruan, D., Zou, W., Dong, L., and Sheng, K. (2021). Linear energy transfer weighted beam orientation optimization for intensity-modulated proton therapy. *Medical physics*, 48(1):57–70.

[Halabi et al., 2006] Halabi, T., Craft, D., and Bortfeld, T. (2006). Dose–volume objectives in multi-criteria optimization. *Physics in Medicine and Biology*, 51(15).

[Ivanenko and Plyasunov, 2008] Ivanenko, D. S. and Plyasunov, A. V. (2008). Reducibility of bilevel programming problems to vector optimization problems. *Journal of Applied and Industrial Mathematics*, 2(2):179–195.

[Küfer et al., 2002] Küfer, K. H., Hamacher, H., and Hans, W. (2002). Inverse radiation therapy planning – a multiple objective optimization approach. *Discrete Applied Mathematics*, 118.

[Küfer et al., 2003] Küfer, K.-H., Scherrer, A., Monz, M., Alonso, F., Trinkaus, H., Bortfeld, T., and Thieke, C. (2003). Intensity-modulated radiotherapy – a large scale multi-criteria programming problem. *OR Spectrum*, 25:223–249.

[Monz et al., 2008] Monz, M., Küfer, K. H., Bortfeld, T. R., and Thieke, C. (2008). Pareto navigation - algorithmic foundation of interactive multi-criteria IMRT planning. *Phys. Med. Biol.*, 53:985–998.

[Paganetti and van Luijk, 2013] Paganetti, H. and van Luijk, P. (2013). Biological considerations when comparing proton therapy with photon therapy. *Semin. Radiat. Oncol.*, 23:77–87.

[Ruuska et al., 2012] Ruuska, S., Miettinen, K., and Wiecek, M. M. (2012). Connections between single-level and bilevel multiobjective optimization. *Journal of Optimization Theory and Applications*, 153(1):60–74.

[Serna, 2012] Serna, J. I. (2012). *Multi-objective Optimization in Mixed Integer Problems with Application to the Beam Angle Optimization Problem in IMRT*. PhD thesis, Technical University of Kaiserslautern, Department of Mathematics.

[Unkelbach et al., 2016] Unkelbach, J., Botas, P., Giantsoudi, D., Gorissen, B. L., and Paganetti, H. (2016). Reoptimization of intensity modulated proton therapy plans based on linear energy transfer. *International journal of radiation oncology, biology, physics*, 96(5):1097–1106.

5 Appendix

5.1 Proof of lemma 1

Proof. [Serna, 2012][Lemma 1.23] Assume that for a $w \in C^* \setminus \{0\}$, z is an optimal solution of $\min\{\sum_{i=1}^n w_i F_i(x) \; s.t. \; x \in X\}$. Further, assume z is not weakly non-dominated. This means $\exists \tilde{z} \in F(X) : \tilde{z} \in z - int(C)$. Hence, there is a point $a \in int(C)$ such that $\tilde{z} = z - a$. With the properties of the dual cone it follows

$$w^T\tilde{z} = w^T(z-a) = w^Tz - \underbrace{w^Ta}_{>0} < w^Tz,$$

which contradicts the assumption. Likewise, for $\bar{w} \in int(C^*)$, assume that z is an optimal solution of the weighted sum but not non-dominated. Thus, there is an $b \in C$ and $\tilde{z} \in F(X)$ such that $\tilde{z} = z - b$. Again, this implies $\bar{w}^T\tilde{z} < \bar{w}^Tz$ and consequently contradicts the assumption. □

5.2 Proof of Lemma 2

Proof.
- C_{SLR} includes all points that are zero for the dimensions corresponding to the lower level and arbitrary values for $u \in C_U$. Thus, if C_U is pointed then C_{SLR} is pointed.

- Assume $d = (d_u, d_l) \in C_{SLR}$. Then $-d = (-d_u, -d_l)$. If there is an i such that $d_{l_i} \neq 0$ then $-d_l \notin C_L$ and consequently $-d \notin C_{SLR}$. If $d_l = 0$ and there is a j such that $d_{u_j} \neq 0$ then $-d_u \notin C_U$ and thus $d \notin C_{SLR}$.
- Assume $C_L \setminus \{0\}$ and C_U to be convex. Let $a, b \in C_{SLR}$. Clearly, if $a_l, b_l \in C_L \setminus \{0\}$, so is their sum and if $a_l = b_l = 0$ and $a_u, b_u \in C_U$, then $a_u + b_u \in C_U$. Without loss if generality assume $a_l \in C_L \setminus \{0\}$, $b_l = 0$ and $b_u \in C_U$. Then $a_l + b_l = a_l + 0 = a_l \in C_L \setminus \{0\}$. Hence, $a + b \in C_{SLR}$.

$\square$