An open source multi-criteria optimization framework for radiotherapy

Tobias Becher[1,2,3], Amit Ben Antony Bennan[1,2], Remo Cristoforetti[1,2,3], and Niklas Wahl[1,2]

[1]Department of Medical Physics in Radiation Oncology, German Cancer Research Center – DKFZ, Heidelberg, Germany
[2]Heidelberg Institute for Radiation Oncology (HIRO) and National Center for Radiation Research in Oncology (NCRO), Heidelberg, Germany
[3]Department of Physics and Astronomy, Heidelberg University, Heidelberg, Germany

Abstract We present an implementation of a Pareto plan navigation framework in the open source treatment planning toolkit `matRad`. It expands on the existing weighted sum approach and allows for calculation and navigation of Pareto surfaces. Furthermore we introduce a lexicographic approach as an alternative to the weighted sum method.

1 Introduction

A central concern of the treatment planning problem in radiotherapy is the handling of different trade-offs like target coverage against organ at risk (OAR) sparing. Since the different objectives often contradict each other there is no single optimal plan, but rather an entire set of *Pareto optimal* plans that cannot improve on one objective without worsening at least another.

Methods to solve this multi-criteria optimization (MCO) problem evolved from an iterative tuning of weights in a weighted sum objective function to automated approaches using Pareto front approximation or lexicographic optimization. Pareto front approximation techniques automatically compute a set of Pareto optimal plans and allow navigation of the resulting Pareto front in an interactive user interface. Lexicographic approaches, on the other hand, require definition of a *wish list*, e. g., based on a clinical protocol prioritizing treatment goals, to automatically reach an acceptable plan. These approaches are now commonly implemented and used in commercial treatment planning systems. The open-source community is, however, lacking comprehensive implementation of such methods for research and education.

Within the open-source radiotherapy treatment planning toolkit `matRad`, which up to its latest release relies on the weighted sum approach for non-linear treatment plan optimization, we developed a MCO module to allow for Pareto front approximation and navigation as well as lexicographic optimization. The implementations are based on well-established methods published in literature.

For Pareto front approximation and navigation, we employ a so-called *sandwich* algorithm [1] that iteratively chooses weights automatically to generate a plan database approximating the Pareto front. Afterwards, a navigation algorithm finds convex combinations of the pre-calculated plans to allow near continuous and interactive exploration of the solutions.

For lexicographic optimization, `matRad` was extended with a priority and goal system to enable definition of a wish list. The $2p\varepsilon c$-method [2] was chosen to consecutively optimize individual objectives constrained by previous solutions and goals. In the end, this method results in a plan optimally fulfilling all the defined goals while prioritizing minimization of the objective functions according to their priority in the wish list.

In this work, we describe the implementation of said MCO framework in `matRad` and demonstrate both methods by application to a prostate case.[1]

2 Materials and Methods

2.1 Software framework

The basis for our MCO framework is `matRad` (version 2.10.1 *Blaise*) [3, 4]. It is written in `MATLAB` and allows for IMRT and IMPT dose calculation and plan optimization using a weighted sum approach with user chosen penalties. Photons, protons and carbon ions are available for planning including biological optimization for the latter. Optimization relies on `IPOPT` [5] or Matlab's `fmincon` using an L-BFGS approximation for Hessian matrices.

To implement Pareto front approximation, linear programs and convex hull calculations are required. These are solved using the `Matlab` functions *linprog* and *convhulln*, respectively.

2.2 The multi-criteria optimization problem in radiotherapy

The general multi-criteria RT treatment planning problem can be expressed through a set of functions **f** and hard constraints **c**:

$$
\begin{aligned}
\underset{\mathbf{w}}{\text{minimize}} \quad & \mathbf{f}(\mathbf{d}) = [f_1(\mathbf{d}), f_2(\mathbf{d}), ..., f_n(\mathbf{d})] \\
\text{subject to} \quad & c_k^l \leq c_k(\mathbf{d}) \leq c_k^u \qquad \forall k, \\
& d_i = \sum_j D_{ij} w_j \qquad \forall i, \\
& w_i \geq 0 \qquad \forall i
\end{aligned}
\tag{1}
$$

d is a discretized vector of dose values with elements d_i related to the beamlet weights **w** via the precomputed dose influence matrix **D** such that $\mathbf{d} = \mathbf{D}\mathbf{w}$. One approach minimizes the weighted sum, $\sum_i p_i f_i$, where p_i are weights/penalties applied to each objective.Varying the penalties allows to obtain different plan options.

[1]The implementation is available in the `matRad` repository on `Github` (Commit ID at submission: f19635f).

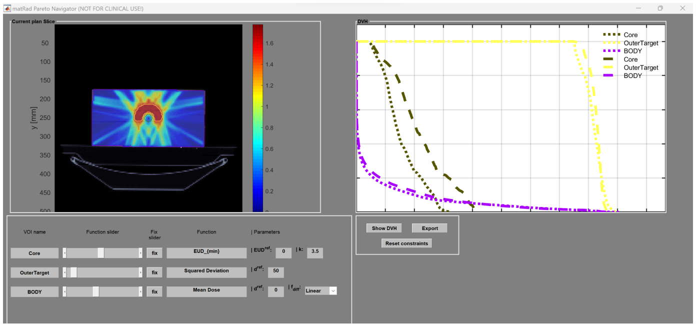

Figure 1: GUI for interactive plan navigation. Objective values can be adjusted by moving the sliders and limited through upper bounds. Additional buttons allow to show the current plan's DVH and to export it to the general `matRad` GUI.

2.3 Pareto front approximation

For Pareto front approximation, we rely on so-called *sandwich* algorithms that approximate the Pareto surface with polyhedral inner and outer approximations. Given an existing set of plans with corresponding objective function values $\mathbf{y}$ the inner approximation is formed by the corresponding convex hull while the outer approximation is bounded by the intersection of supporting hyperplanes. With the assistance of the maximum distance of inner and outer approximation a new weight vector is chosen to improve on the current approximation in the most improvable region.
The algorithm for our implementation was introduced by Rennen, Dam, and Hertog [1]. By optimizing each objective individually we calculate the anchor points of the Pareto front. The anchor point solutions are used to scale the objectives f_i to a range of 0 to 1, to better account for different dimensions of objective function values, using their highest values u_i and lowest value l_i. The scaled objectives f' are then given by $f_i' = \frac{f_i - l_i}{u_i - l_i}$. Afterwards the algorithm iterates over the bounding facets of the inner approximations and calculates new penalty vectors based on the normal vector of the facet with largest distance.
For optimization of individual weight computations found by the sandwich algorithm, we use `matRad's` existing weighted sum optimization framework, allowing us to use any optimizer interfaced by `matRad`. The sandwich algorithm continues until the maximal distance between inner and outer approximation falls below a predetermined threshold or a certain number of iterations is passed.
Finally, all solutions are stored in form of the resulting beamlet intensity vectors for direct dose calculation or plan navigation.

2.4 Plan navigation on the Pareto front

For visualization of the different plans after database calculation, a dedicated widget is available for `matRad's` graphical user interface (GUI). The widget is shown in Figure 1 with an example Pareto front approximation loaded.
Starting from an initial plan with objective function values $\hat{\mathbf{y}}$, a new plan can be investigated by moving across the Pareto surface using interactive sliders. Upon moving a slider, the corresponding objective function value is fixed and a new plan is found as a convex combination of pre-calculated plans using the algorithm introduced by Monz, Küfer, Bortfeld, et al. [6]. The core idea of the algorithm is to start from the point $\mathbf{y}^R$ obtained through the slider movement and then search along the line $\mathbf{y}^R + t\mathbf{1}$ for an intersection $\mathbf{y}'$ with the approximation of the Pareto surface, expressed as $\mathbf{y}' = \sum_k v_k \mathbf{y}^k$. $\mathbf{y}^k$ are the objective function values of the pre-calculated plans and v_k the corresponding convex coefficients, that are then used to calculate new weights $\mathbf{w}' = \sum_k v_k \mathbf{w}^k$ from the pre-calculated weights $\mathbf{w}^k$. While the new plans are not necessarily Pareto optimal,[2] they are still clinically relevant if the Pareto surface is sufficiently approximated.
To limit navigation to regions of interest, upper bounds may be added on specific objective functions. The range of the sliders is then adjusted to mimic the upper bound and since the change also limits the options for other objectives the respective slider ranges are changed as well.

2.5 Lexicographic optimization

Lexicographic optimization relies on a prioritized wish list formulating goals for each objective f_i. The chosen method for our implementation – the $2p\varepsilon c$-method [2, 7] – consists of two fully automated plan optimization phases of iterating to the wish list.
The first phase optimizes objectives individually one after the other (under the general constraints of Equation 1). In each step, moving to optimization of lower prioritized objectives adds additional constraints $f_j \leq \varepsilon_j$ to still respect higher priority objectives. ε_j is set to the corresponding aspired value g_j if it can be met and else to the achieved value f_j^* multiplied with a slack variable c (in our implementation chosen as $c = 1.03$).
The second phase then turns objectives that could not meet their goal value into hard constraints and they are not optimized again. For the remaining objectives it applies the ε-constraint method which can be formulated as

$$
\begin{aligned}
\underset{\mathbf{w}}{\text{minimize}} \quad & f_i(\mathbf{d}) & & \\
\text{subject to} \quad & c_k^l \leq c_k(\mathbf{d}) \leq c_k^u & & \forall k, \\
& f_j(\mathbf{d}) \leq \varepsilon_j & & \forall j \in \{1,..,n\} \backslash i, \\
& d_i = \sum_j D_{ij} w_j & & \forall i, \\
& w_i \geq 0 & & \forall i
\end{aligned}
\tag{2}
$$

Initially the ε values are set as achieved in the first phase and after optimization are changed to $c f_j^*$ for all following objectives.
This approach has the advantage that objectives that easily meet their goal value do not dominate the optimization in the second step, while ensuring that objectives which are not able to meet their goal value are still respected.

[2] for convex objectives it holds that $\mathbf{f}\left(\mathbf{D}\sum_k v_k \mathbf{w}^k\right) \leq \sum_k v_k \mathbf{y}^k$

	Objectives		Constraints	
Volume	Name	Parameters	Min[Gy]	Max[Gy]
Target	$f_{sq\ deviation}$	$\bar{d} = 50$	45	57
Core	f_{EUD}	$a = 3.5$	0	40
Body	f_{mean}	-	0	45

Table 1: Objectives and constraints used for the TG119 phantom.

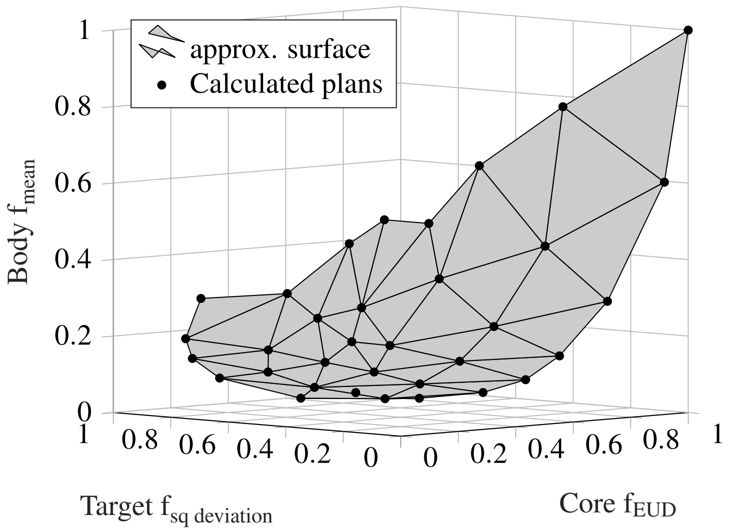

Figure 2: Example Pareto surface for the TG119 phantom. The (normalized) objective function values define the three axes. The plans obtained through optimizing the weight combinations found by the sandwich algorithm are shown as dots connected by the approximated Pareto surface.

3 Results

3.1 Pareto front approximation

The implementation of the Pareto front approximation algorithm was verified with the non-radiotherapy test case 1 from the original publication [1] in `MATLAB`.
It was then applied to two radiotherapy scenarios using the CORT dataset available within `matRad`. The first example uses the TG119 phantom shown in Figure 1. The three objectives set (see Table 1) allow to plot the Pareto surface approximation (see Figure 2).
The second example uses a prostate patient (shown in Table 2). In total, 105 plans were calculated for the database. In Figure 3, we show four plans: (1) anchor point plan focusing on the 68 Gy volume, (2) anchor point plan focusing on rectum, (3) a balanced plan from the precalculated database and (4) plan found with help of the navigation interface.

	Objectives		Constraints	
Volume	Name	Parameters	Min[Gy]	Max[Gy]
PTV 68	$f_{sq\ deviation}$	$\bar{d} = 68$	61.2	78.2
PTV 56	$f_{sq\ deviation}$	$\bar{d} = 56$	50.4	64.4
Bladder	f_{EUD}	$a = 3$	0	58
Rectum	f_{EUD}	$a = 8$	0	58
Body	f_{mean}	-	0	58

Table 2: Objectives and constraints used in the sandwich algorithm implementation for the prostate case.

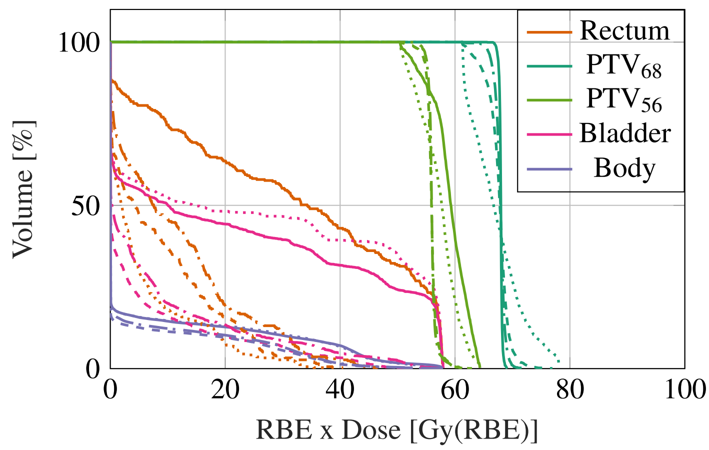

Figure 3: DVH for the four prostate plans. Anchor point plans (1) and (2) are shown as solid and dotted lines, respectively. The balanced database plans (3) and (4) are plotted dashed and dash-dotted.

3.2 Lexicographic optimization

To test the $2p\varepsilon c$-method a wish list on the prostate phantom presented in Table 3 was optimized.

Constraints				
Number	Volume	Name	Min[Gy]	Max[Gy]
1	Rectum	Min/Max Dose	0	58
2	Bladder	Min/Max Dose	0	58
3	Body	Min/Max Dose	0	58
Objectives				
Priority	Volume	Name	Param.	Goal
1	PTV 68	$f_{sq\ deviation}$	$\bar{d} = 68$	2
2	PTV 56	$f_{sq\ deviation}$	$\bar{d} = 56$	2
3	Rectum	EUD	$a = 8$	30
4	Bladder	EUD	$a = 3$	30
5	Body	Mean dose	-	10

Table 3: wish list for the prostate patient. Constraints for OAR are set as in the Pareto front approximation, while target constraints are omitted. The PTVs have highest priority and are are assigned a squared deviation objective each.

Table 4 shows achievable objective function values and chosen constraints in both phases of the $2p\varepsilon c$ approach. All objectives except for Rectum meet their aspired goal value in the first phase and are therefore optimized again. A DVH corresponding to the final plan can be found in Figure 4. The target coverage was achieved and while the Rectum goal is missed, the resulting plan was evaluated to be acceptable.

Priority	Volume	Goal	Achievable	New constraint
First step				
1	PTV 68	2	$5.9 \cdot 10^{-2}$	2
2	PTV 56	2	$2.6 \cdot 10^{-6}$	2
3	Rectum	20	22.6	23.3
4	Bladder	30	18.9	20
5	Body	10	5.1	10
Second step				
1	PTV 68	2	1.67	1.72
2	PTV 56	2	0.70	0.72
4	Bladder	30	23.8	24.5
5	Body	10	4.3	-

Table 4: Result table of the prostate patient. The goal on the rectum cannot be met in the first iteration and is therefore skipped in the second iteration.

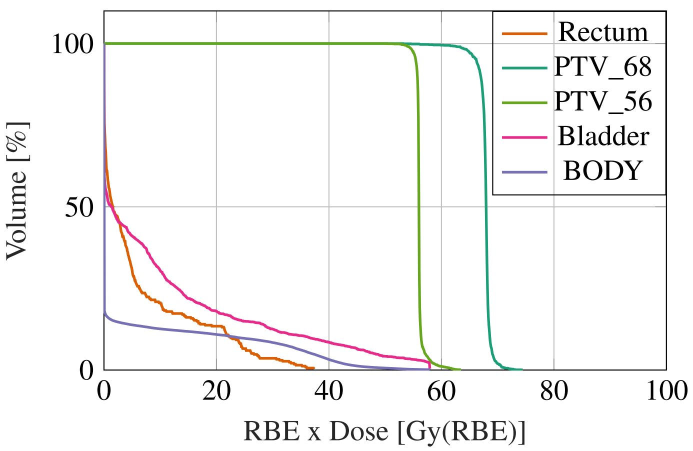

Figure 4: DVH for the prostate plan calculated with the $2p\varepsilon c$-method.

4 Discussion

Two established MCO treatment planning approaches were successfully implemented in the open-source treatment planning system `matRad`.

4.1 Pareto front approximation and navigation

The implemented sandwich algorithm was benchmarked against published results and the Pareto surface could be adequately approximated for two RT examples.
While navigation of the treatment plan is possible nearly in realtime, once the plan database is calculated using the sandwich algorithm, the computation time is limited by the constraint optimization method in `matRad`. Since `matRad` favors flexibility of computational performance optimization, the performance of the non-linear optimization approach is suboptimal for Pareto surface approximation. However, computational performance can be optimized in the future by users themselves or public community contributions.

Future improvements to the algorithm itself might include automatic clustering of correlated objectives to reduce the complexity of the calculation for high dimensional problems as well as the implementation of parallelizable sandwich algorithms.

4.2 Lexicographic optimization

Implementation of the $2p\varepsilon c$-method was demonstrated with a prostate plan and was shown to work as expected. It allows for a more automated and clinically oriented steering of the treatment plan than the weighted sum approach, however design of a suitable wish list depends on the planner's experience. Similar computational limits as in the Pareto front approximation approach apply, as the approach is limited in performance by `matRad's` underlying optimization algorithm.

5 Conclusion

We present our MCO framework in the open source treatment planning software `matRad` which includes: (1) A sandwich algorithm implementation to calculate Pareto optimal plans and a navigation interface to allow their exploration afterwards. (2) also introduced a lexicographic method for automated treatment plan generation. This package facilitates future research into advanced treatment plan optimization strategies.

References

[1] G. Rennen, E. van Dam, and D. den Hertog. *Enhancement of Sandwich Algorithms for Approximating Higher Dimensional Convex Pareto Sets*. Rochester, NY, June 30, 2009. DOI: 10.2139/ssrn.1427721.

[2] S. Breedveld, P. R. M. Storchi, and B. J. M. Heijmen. "The equivalence of multi-criteria methods for radiotherapy plan optimization". *Physics in Medicine & Biology* 54.23 (Nov. 2009), p. 7199. DOI: 10.1088/0031-9155/54/23/011.

[3] H.-P. Wieser, E. Cisternas, N. Wahl, et al. "Development of the open-source dose calculation and optimization toolkit matRad". *Medical Physics* 44.6 (2017). _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/mp.12251, pp. 2556–2568. DOI: 10.1002/mp.12251.

[4] B. Ackermann, M. Bangert, A. B. A. Bennan, et al. *e0404/matRad: Blaise v2.10.1*. Version 2.10.1. Nov. 20, 2020. DOI: 10.5281/zenodo.7107719.

[5] A. Wächter and L. T. Biegler. "On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming". *Mathematical Programming* 106.1 (Mar. 1, 2006), pp. 25–57. DOI: 10.1007/s10107-004-0559-y.

[6] M. Monz, K. H. Küfer, T. R. Bortfeld, et al. "Pareto navigation—algorithmic foundation of interactive multi-criteria IMRT planning". *Physics in Medicine and Biology* 53.4 (Feb. 21, 2008), pp. 985–998. DOI: 10.1088/0031-9155/53/4/011.

[7] S. Breedveld, P. R. M. Storchi, M. Keijzer, et al. "A novel approach to multi-criteria inverse planning for IMRT". *Physics in Medicine & Biology* 52.20 (Oct. 2007), p. 6339. DOI: 10.1088/0031-9155/52/20/016.