Volume 20 • Issue 4 • June 2016

Multi-criteria optimisation and decision-making in radiotherapy

Sebastiaan Breedveld & Rens van Haveren

When treating patients with radiotherapy, healthy organs may be damaged, resulting in loss of functionality. The distribution of the dose over the different organs is modelled as a multi-criteria optimisation problem. As there is not a single optimal (or "golden") solution, decision-making is used to find a clinically acceptable solution.

Radiotherapy (radiation therapy) is used in treating $\approx 50\%$ of cancer patients. Due to irradiation of healthy organs surrounding the tumour, severe side-effects may be induced, impacting the patient's quality-of-life. For example, reduced functionality of salivary glands may result in dry-mouth syndrome. This urges the patient to drink some water every half-an-hour (day and night). Damaged swallowing muscles impairs the consumption of solid food, and limits the patient to eating mashed food, or using a probe.

It goes without saying that each complication has a significant impact on the quality-of-life of the patient and it is thus of utmost importance to reduce the probability of developing treatment-induced complications. The dose levels for which complications occur are not fixed, but differ between patients.

Also, different organs have different radiosensitivity and it is therefore important to consider the dose to each organ separately. The dose to the different organs is correlated, resulting that sparing one completely may result in the inability to spare the others (see figure 1). This means that a trade-off has to be made between the doses to the different organs. A typical trade-off could be between the dose to the salivary glands and the dose to the swallowing muscles, where a balance has to be found between dry-mouth problems and swallowing problems. A real multi-criteria problem in radiotherapy typically consists of 20-30 criteria.

We use techniques from *multi-criteria optimisation* and *decision-making* (MCDM), subfields of Operational Research, to find clinically acceptable treatment plans, and to automate the decision-making process.

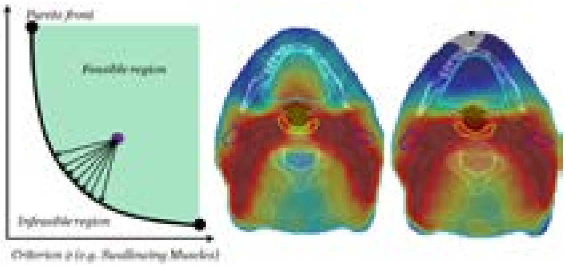

Figure 1: Left: Pareto front. A non-Pareto optimal solution (purple) can be improved in one or more criteria, until it is on the Pareto front. The black dots at the end of the curve are the extremities, i.e. full minimisation of one criterion. Middle and right: Pareto optimal treatment plans with different trade-offs. Cyan = mouth, yellow = a swallowing muscle, blue/green = salivary glands. Colourwash: More red (hotter) = higher dose (damage).

Multi-criteria optimisation

A multi-criteria optimisation problem is formulated as follows:

$$\begin{array}{ll} \text{minimise} & [f_1(x), f_2(x), \ldots, f_n(x)] \\ \text{subject to} & g(x) \leq 0 \end{array}$$

where $f_i(x)$ are nonlinear real-valued functions, $g(x)$ is a set of nonlinear constraints and $x \in \mathbb{R}^m$ the decision variables. The goal of multi-criteria optimisation is to find a balance between the criteria $[f_1, f_2, \ldots, f_n]$. An important concept is *Pareto-optimality*. A solution $[f_1^*, f_2^*, \ldots, f_n^*]$ is Pareto-optimal if an improvement for one criterion results in a deterioration for at least one of the others, i.e. there is no gain without loss. This concept is visualised in figure 1 for $n = 2$.

However, there is not a "golden" solution. What is an acceptable trade-off is determined by the long-term experience and insights of the physician (medical doctor). It is clear that there is not a sane trade-off near the extremities, as there is a huge gain for one with only a slight deterioration of the other. But what makes a clinically acceptable trade-off is less well defined. There are different approaches to handle this *decision-making* problem. One is to give full control to the physician, and let him/her determine the most favourable solution. Another approach is to formalise the decision-making process, allowing automated decision-making. The challenge of automated decision-making is to find a general configuration which results in clinically acceptable plans for different patients (within the same treatment group).

Pareto front navigation

By using Pareto front navigation, the physician (operator) can freely explore the trade-offs between different criteria (see [3] for an interactive example). This way, the different doses belonging to the different solutions are directly visualised. For a bicriteria problem, the Pareto front can be spanned by solving:

$$\begin{array}{ll} \text{minimise} & \alpha f_1(x) + (1-\alpha) f_2(x) \\ \text{subject to} & g(x) \leq 0 \end{array}$$

for weights $\alpha \in [0, 1]$. Each value of α gives a unique Pareto-optimal solution.

There are some practical problems, resulting from the *curse of dimensionality*. A uniform distribution of α does not necessarily result in a uniform distribution of points on the Pareto front, and doing this for $n > 3$ is not straightforward. Decision-making is also problematic: humans are in general capable of considering up to 5 criteria simultaneously, so it is questionable if the operator will always reach an optimal decision, or will reach the same solution a second time (reproducibility).

Sequential ϵ-constraint programming

Radiotherapy treatment plan optimisation requires consistent, reproducible and fast decision-making (for several reasons). The idea is to automate decision-making by mimicking the human decision-maker's decision pattern. This is possible by constructing a prioritised list with goal values to reach for each criterion. An example is given in table 1.

Table 1: Constraints and criteria

Constraints

	Name	Type	Limit
	Tumour	minimum	$43.7\ Gy$
	Nerves	maximum	$45\ Gy$

Criteria

Priority	Volume	Type	Goal
1	Salivary Glands	minimise	$40\ Gy$
2	Swallowing Muscles	minimise	$50\ Gy$
3	Salivary Glands	minimise	$20\ Gy$
4	Swallowing Muscles	minimise	$40\ Gy$
5	Salivary Glands	minimise	$0\ Gy$
6	Swallowing Muscles	minimise	$0\ Gy$

In this list, there are 2 hard constraints and 6 prioritised steps which are processed in sequential order. The idea is that in each step, the criterion is minimised up to, but not lower than the given *Goal*. This is to prevent that the solution is steered to one of the extremities, severely limiting the sparing of other criteria. This criterion is then added as a constraint to the problem. Formally, the steps taken are:

$$\begin{array}{ll} \text{minimise} & f_1(x) \\ \text{subject to} & g(x) \leq 0 \end{array}$$

Depending on the result $f_i(x^*)$, the new bound is chosen according to the following rule:

$$\epsilon_i = \begin{cases} b_i & f_i(x^*)\delta < b_i \\ f_i(x^*)\delta & f_i(x^*)\delta \geq b_i \end{cases}$$

where b_i is the Goal for f_i (table 1) and δ is a slight relaxation to create some space for the subsequent optimisations, usually set to 1.03 (3%). This step leads to more favourable trade-offs by steering the solution away from steep parts of the Pareto front.

The next optimisation optimises f_2, keeping f_1 constrained:

$$\begin{array}{ll} \text{minimise} & f_2(x) \\ \text{subject to} & g(x) \leq 0 \\ & f_1(x) \leq \epsilon_1 \end{array}$$

and so on.

The list in table 1 describes a balanced search between the 2 criteria, similar to a manual search. The advantage is that many criteria can be included, and that the search is deterministic. The challenge is to find a uniform configuration (list) which results in clinically acceptable plans for each patient.

Lexicographic Reference Point Method

A disadvantage of the previous method is that it requires many optimisations. Essentially, this method follows a path through the criterion-space (while obeying the imposed hard constraints, e.g. as in table 1). The idea of the Lexicographic Reference Point Method (LRPM) is to optimise along this path, resulting in a single optimisation problem. This implies shorter calculation times, in particular when the list contains many criteria and priorities. The idea of the LRPM is to express the preference structure through *reference points*. A reference point is a vector in which each coordinate represents a goal value for one of the criteria. These goal values are equally important to attain, e.g. the reference point $(f_1, f_2) = (20, 40)$ indicates that attaining 20 for f_1 is as important as attaining 40 for f_2. However, different reference points have different priorities: the first priority is to attain the worst reference point. If feasible, the focus is on attaining the second-worst reference point. This process continues until no improvements are possible. Intuitively, the worst-case scenario (first reference point) is steered in a prioritised manner towards the ideal scenario (last reference point). The flexibility of using multiple reference points allows us to prioritise the criteria, e.g. see figure 2: in between every pair of subsequent reference points, the focus alternates on decreasing either f_1 or f_2. This approach can be extended for an the general case (n criteria) in which one can specify a finite number of reference points that strictly decrease coordinatewise.

The next step is to design a *reference path*: a parametric curve through the criterion space that connects the reference points so that the parametric equations are strictly decreasing. Then, a single optimisation problem is formed [2] to process all objectives and their priorities. Intuitively, the reference path is followed until it intersects with the Pareto front. This intersection is then the returned Pareto-optimal solution.

To apply the LRPM with a uniform configuration for a group of patients, some modifications are needed. The intersection of the reference path and the Pareto front may occur on a steep part of the Pareto front (representing an undesired trade-off) for some patients, see figure 2. To overcome this issue, we need a mechanism that excludes the selection of solutions near the extremities. Hereto, *trade-off curves* are modelled into the LRPM: without the trade-off curves, the red plans are generated while the clinically favourable plans (green plans) are generated with the trade-off curves. In conclusion, even though the clinically favourable plans seem scattered and unstructured, there is an underlying preference structure that can be captured by the LRPM using a uniform configuration (reference points and trade-off curves). This allows us to apply the LRPM with fixed parameters per treatment site.

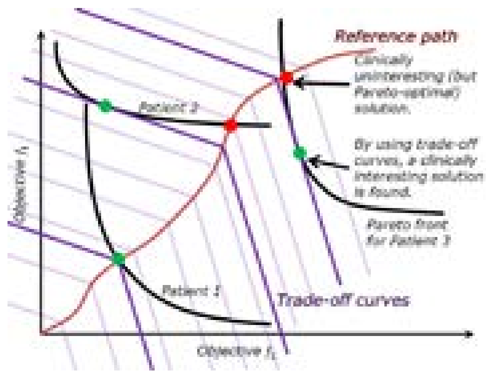

Figure 2: The effect of adding trade-off curves. The black lines are the Pareto fronts for three different patients and the green/red dots represent the generated plans with/without trade-off curves.

Automated Radiotherapy Treatment Planning

The ability to automate decision-making allows to cope with other complex parts of the problem, resulting in fully automated generation of radiotherapy treatment plans. The quality of these plans are in general equal, and often much better than existing clinical practise. Currently, these methods are in clinical use in the Erasmus University Medical Center in Rotterdam.

References

[1] Breedveld S., Storchi P. & Heijmen B. (2009) *The equivalence of multi-criteria methods for radiotherapy plan optimization* DOI:10.1088/0031-9155/54/23/011

[2] Van Haveren R., Breedveld S., Keijzer M., Voet P., Heijmen B. & Ogryczak W. *A Novel Lexicographic Reference Point Method with Bounded Trade-offs Applied in Radiation Therapy Treatment Planning* (to appear)

[3] Pareto front navigation example: http://sebastiaanbreedveld.nl/rt_paretonavigation.html

Mathematics