Expert Systems with Applications 42 (2015) 8337–8348

Contents lists available at ScienceDirect

Expert Systems with Applications

journal homepage: www.elsevier.com/locate/eswa

Interactive visual guidance for automated stereotactic radiosurgery treatment planning

Danielle A. Ripsman [a], Dionne M. Aleman [a,b,c], Kimia Ghobadi [a,*]

[a] Department of Mechanical and Industrial Engineering, University of Toronto, 5 King's College Road, Toronto, ON M5S 3G8, Canada
[b] Institute for Health Policy, Management and Evaluation, University of Toronto, 155 College Street, Suite 425, Toronto, ON M5T 3M6, Canada
[c] Techna Institute, University Health Network, 610 University Avenue, Toronto, ON M5G 2M9, Canada

ARTICLE INFO

Article history:
Available online 28 June 2015

Keywords:
Graphical user interface
Radiation therapy
Decision support
Expert systems
Treatment planning
Multi-objective parameter estimation

ABSTRACT

The growing technology industry has led to the steady enhancement of expert systems, often at the cost of increased complexity for the systems' end users. Efforts to improve the prescriptive elements of systems, however, often prove unsuccessful, since the nature of complex and high-dimensional decision problems is difficult to capture precisely by models and algorithms. To rectify this deficiency, complementary softwares may be used to accept decision-making input from users. In this paper, we introduce a graphical interface-based multi-criteria decision support system for designing radiation therapy treatment plans. While many automated strategies for treatment plan generation exist in the literature, they often require a large amount of iteration and *a priori* decision-making in practice, so much of the planning is done manually. Our interface, morDiRECT (the Medical Operations Research Laboratory's Display for Ranking and Evaluating Customized Treatments) uses the variability associated with the planning parameters to generate diverse plan sets automatically, creating a comprehensive and visible decision space for users. We demonstrate morDiRECT's generation process, built-in analytical tooling and graphical display using four clinical case studies. In three cases, we find plans that fully dominate the benchmark forward plans, as well as additional plans that possess potentially desirable tradeoffs for all cases. Our results demonstrate that with relatively little upfront effort, users can pre-generate and choose from a diverse set of clinically acceptable plans, leading to reliable treatments for head-and-neck patients.

1. Introduction

Medical technology has advanced considerably over the past few decades, paving the way for a rapid evolution in radiation therapy treatment tools (Schimpff, 2014). As these tools grow in sophistication, the analytical complexity and resultant cognitive loading placed on their operators increases dramatically (Ruotsalainen, Boman, Miettinen, & Tervo, 2009). In order to alleviate some of this demand, we introduce a graphical user interface (GUI) called morDiRECT (the Medical Operations Research Laboratory's Display for Ranking and Evaluating Customized Treatments) to support expert users through the non-trivial tasks of radiation therapy treatment planning and selection.

The radiation therapy delivery process can be broken down into five key stages, depicted in Fig. 1. The majority of the tools designed to facilitate this process aim to support the more mechanical Stages 1 and 5 (21st Century Oncology, 2013; Elekta AB, 2013; General Electric Company, 2013; IBA, 2013; Philips, 2013). In contrast, the intermediate planning done in Stages 2–4, tends to be less supported, relying heavily on human operators (National Cancer Institute, 2014). Due to the complex nature of these planning stages, automated strategies such as inverse planning have become a prevalent source of discussion in the radiation therapy literature (Romeijn & Dempsey, 2008; Webb, 2014). The high versatility of inverse planning has also led to extensions to similar problems within the field, such as Leksell Gamma Knife® radiosurgery treatment (Ferris & Shepard, 2000; Ferris, Lim, & Shepard, 2002; Ferris, Lim, & Shepard, 2003; Ghobadi, Ghaffari, Aleman, Jaffray, & Ruschin, 2012; Ghobadi, 2014; Shepard, Yu, Murphy, Bussière, & Bova,2015; Shepard et al., 2015; Wu et al., 2003; Wu et al., 2004).

Inverse planning methodologies fall into a class of algorithms that specifically target Stage 3 of the delivery process. Stage 3 is difficult by nature as it is associated with a detailed understanding of the technology, as well as the case at hand. While carrying out this stage (either manually or with the help of traditional inverse

* Corresponding author at: Sloan School of Management, MIT, MA, USA.
E-mail addresses: danielle.ripsman@mail.utoronto.ca (D.A. Ripsman), aleman@mie.utoronto.ca (D.M. Aleman), kimia@mie.utoronto.ca (K. Ghobadi).

http://dx.doi.org/10.1016/j.eswa.2015.06.028
0957-4174/

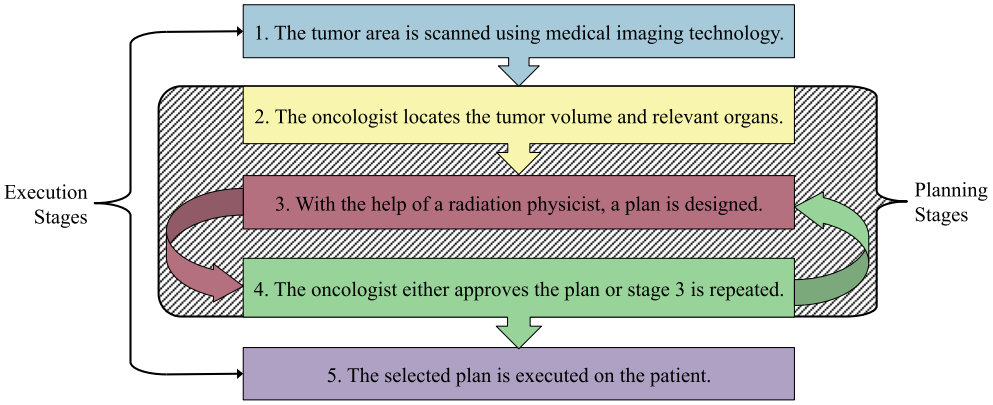

Fig. 1. Key stages in radiation therapy delivery.

planning software) the radiation physicist must use trial and error to balance a potentially large number of competing treatment objectives and complex machine specifications simultaneously. Failure to account for all the relevant objectives can be costly for both the patient and the hospital, as it may lead to the subsequent rejection of the proposed plan in Stage 4, and thus require further iterations. Proposals to remedy this issue typically employ either automatic or interface-based extensions to the existing inverse planning frameworks, in an effort to reduce the associated cognitive loading.

Among the automatic class of extensions are several multi-objective optimization techniques that have been broadened to incorporate expert preferences. These methods include hierarchical constraint tightening (or loosening) (Breedveld, Storchi, Keijzer, Heemink, & Heijmen, 2007), transformed statistical rankings (Lourenzutti & Krohling, 2014), decision theory-based rankings (Yu, 1997), stochastic analytical hierarchy processes (SAHP) (Cobuloglu & Büyüktahtakın, 2015) and lexicographic ordering (Long et al., 2012). A related form of automation is the case-based reasoning approach, which circumvents the need for concrete objectives in favour of choosing solutions based on similarities to past cases (Lolli, Ishizaka, Gamberini, Rimini, & Messori, 2015; Petrovic, Mishra, & Sundar, 2011). The major drawback of these algorithms is the rigidity that stems from an absence of human interference. While there is a potential gain in terms of speed and the reduction of human error, new errors are introduced by *a priori* decision-making that may not be universally acceptable. Additionally, each algorithm only outputs a single plan, meaning that in the case of a rejection, the burden for generating a subsequent plan is placed back on the radiation physicist.

In practice, treatment planning can involve a high degree of uncertainty (Romeijn & Dempsey, 2008) and even advanced algorithms tend to be quite specialized and are consequently no match for their human counterparts in terms of adaptability (Bonczek, Holsapple, & Whinston, 2014; Haight, 2010; Wickens, Lee, Liu, & Gordon-Becker, 2004). For this reason, successfully integrated decision support systems should provide a well-balanced allocation of tasks between experts and automation, motivating the interface-based class of extensions. Although interfaces are a well-established method to support human decision making (Grandjean & Kroemer, 1997; Korhonen & Wallenius, 1988; Lotov, Bushenkov, & Kamenev, 2004; Wickens et al., 2004), and have been more specifically addressed as useful in the field of medicine (Thyvalikakath et al., 2014; Aigner & Miksch, 2006; Gschwandtner, Aigner, Kaiser, Miksch, & Seyfang, 2011), they are still under-utilized in public health applications (Aigner & Miksch, 2006; Thyvalikakath et al., 2014; Yasnoff & Miller, 2014). Interfaces that are implemented frequently lack sophistication, making them less effective in their task of reducing the user's cognitive load (Wickens et al., 2004; Yasnoff & Miller, 2014).

Cotrutz and Xing (2002) present an interface concept for iterative radiotherapy plan improvement based on adjusting localized areas of a commonly used plan assessment plot called a dose volume histogram, though their methodology is only intended for fine tuning and their interface is not explicitly developed. Otto (2014) introduces a supervised approach for iterative dose design, along with a custom interface that uses a speedy approximation algorithm to ensure treatment plans are feasible. Since dosages must be designed before the computationally intensive optimization is run, plan characteristics such as the duration of the treatment will be unknown at the time of plan selection, and may consequently suffer in quality. Hence, while utilizing the planning system is simpler than unsupervised plan design, it is still a cognitively intensive task for clinicians.

Jain, Kahn, Drzymala, Emami, and Purdy (1993) also introduce a radiation therapy interface to support their plan ranking model, however, this fairly simple and tabular interface is only intended to support the selection process, not the plan generation. Hanne and Trinkaus (2003) present a fairly comprehensive spider plot interface called knowCube. While their interface does demonstrate a large range of functionality, their spider plot presentation modality makes it difficult to visualize multiple plan alternatives concurrently and their generation process is rigidly set to generate 1000 plan alternatives, rather than taking input from the planner. Lotov et al. (2004), Bortz et al. (2014) and Korhonen and Wallenius (1988) all discuss the design of Pareto front based interfaces, but do not deal with radiation therapy, while Craft, Halabi, Shih, and Bortfeld (2006) and Wang, Jin, Zhao, Peng, and Hu (2014) provide an analysis of Pareto tradeoffs in radiation therapy planning, but do not include an interface. Rosen, Liu, Childress, and Liao (2005), Ehrgott and Winz (2008) and Aubry, Beaulieu, Sévigny, Beaulieu, and Tremblay (2006), on the other hand, do use Pareto-optimality to generate radiation therapy interfaces. Rosen et al. (2005) introduce TPEx, a simplified dose volume histogram-based interface which allows experts to navigate through a number of allowable plans. The navigation, however, is performed strictly in terms of dose and volume properties and ultimately, the final plan is generated using a non-deterministic algorithm, leading to potential inconsistencies for the end user. Ehrgott and Winz (2008) and Aubry et al. (2006) both provide simpler interfaces, with basic filtering functionality for plan selection.

A prevailing issue with all the above-mentioned designs comes from the concept of choosing only Pareto optimal plans, while simultaneously limiting the number of objectives. By restricting the results to the Pareto front, plans with benefits that are unquantified in the objective function are discarded, obscuring potentially

high quality plans from the user. Navigating Pareto decision spaces graphically can also be quite hard on the user, due to high dimensionality associated with the multiple criteria.

We propose the morDiRECT interface as an alternative to the existing radiation therapy planning intervention techniques. Unlike its predecessors, morDiRECT provides a comprehensive surface plot-based decision support environment that limits dimensionality and places an emphasis on the decision space and the relationship between input parameters and output metrics, rather than outputs alone. It allows users to both generate and explore treatments from an unbiased range of potential plans, whereas many of its alternatives focus purely on generation or analysis and display Pareto optimal results only. While morDiRECT could work with a number of radiation therapy optimization algorithms, we illustrate its performance using a deterministic, non-Pareto radiosurgery optimization algorithm designed by Ghobadi (2014) and Ghobadi, Aleman, Jaffray, and Ruschin (2013). The interface's ability to support the radiation therapy planning and analysis processes is demonstrated on four clinical case studies.

2. Methods and materials

The morDiRECT planner, depicted in Fig. 2, is a graphical user interface designed using MATLAB® (The MathWorks, Inc.) for the purpose of selecting acceptable radiotherapy treatment plans. A typical morDiRECT planning session is comprised of three parts. The first is a supported information transfer between morDiRECT and its users, detailing how plans are to be run (e.g., planning algorithms and parameter selection) and stored (see Section 2.1). The second is an internal application of an inverse planning algorithm to generate the candidate plans (see Section 2.2). Finally, the third is the supported comparison and analysis of potential plans, using a set of built-in functions (see Section 2.3). The plans can be generated serially or in parallel, depending on available infrastructure; either way, plan generation is automated and does not require user supervision.

2.1. Inputs and outputs

Each treatment plan generated by the morDiRECT interface is the product of an independent optimization, incorporating a unique set of input parameters. These parameters are the algorithm-specific constant terms used to represent difficult-to-quantify features, such as mechanical specifications, dose regulations and relative weightings, used in many inverse planning techniques (Romeijn & Dempsey, 2008). Given perfect information, these values could be customized to each patient, yielding perfect plans. In reality, however, the parameter values must be estimated *a priori*. The quality of the estimate can only be determined after the fact, based on an analysis of the output summary statistics or goal metrics collected from the optimization. To assist in the navigation of this uncertain input–output relationship, morDiRECT allows users to run multiple potential plans per trial, which are then analyzed concurrently.

2.1.1. Input parameter selection

Due to their uncertain but substantial influence on the quality of the final plan, the unknown input parameters from the inverse planning algorithms are used as the drivers for morDiRECT's plan generation process. The values of these parameters are the inputs for the interface and they are selected manually through the Input Specification Menu (Fig. 3(a)). Using this menu, the value of any number of input parameters may be adjusted from the default as desired, while a set of up to three parameters are designated as the "driver" inputs, that is, parameters whose values will be explored. Driver inputs are chosen by selecting the checkbox to the left of the parameter's name. Drivers are denoted as $\mathcal{D} = \{D_1, D_2, D_3\}$, where D_i is the domain of values to be tested for driver i.

Once a parameter has been changed to a driver, the planner has the opportunity to assign it a range, rather than just a single value. This action will inform the interface that the optimization should be run multiple times, once at every value of the driving parameter. When multiple drivers are selected, plans are generated exhaustively for every combination of these input sets, such that the total number of plans generated is $\prod_{i=1}^{3}|D_i|$.

The number of drivers is restricted to three as a product of the interface's ultimate objective of conveying meaningful plan trade-offs and relationships to its users. Since outputs are communicated through the Evaluation Window (Fig. 2), where they are displayed with respect to their driving parameter values, additional drivers would force the display to either take on higher dimensional graphs, or divide the plans over larger numbers of subranges, both of which hinder an individual's ability to visualize the relationships between the inputs and outputs in memory. It should be noted, however, that users can still employ the interface to study the interactions of more than three parameters through the use of multiple runs, at an increased cost in terms of time and analytical complexity.

2.1.2. Output metric selection

Since different planners may prefer to evaluate plans based on differing criteria, morDiRECT allows users to choose their desired outputs through the Output Specification Menu (Fig. 3(b)). Within this menu, the planner may specify which outputs or features are to be stored and in what capacity.

A selected checkbox to the left of an output metric indicates that a particular metric is to be stored for later examination. Outputs with corresponding drop down menus provide flexible storage options. If the user chooses the "Plot and Display" option, the specified output will be visible in the final display, appearing both in its own self-titled plot, as well as in a summary panel. "Beam-on time", for example, has been selected for this option in Fig. 3(b) and it can be verified to have its own plot in the upper left corner of Fig. 2, as well as appearing in the Current Node panel of the display (Fig. 2, E). The user could alternatively select "Display Only", which would omit the plot, only displaying a summary metric, or "Plot Only", which does the opposite, revealing a plot but no summary data. Finally, the user can choose the "Store Only" option, if s/he does not wish to reveal a metric in the final display, but would like to save the values for future reference.

2.2. Plan generation algorithm

The selection of optimization algorithm, naturally, has a strong influence on the quality of the generated plan alternatives. Since some algorithms may be better suited to solving particular types of problems, users have the freedom to integrate any number of algorithms into the interface's toolbox, applying them to cases at their own discretion. For demonstration purposes, we will illustrate the continuous path (CP) algorithm for Elekta's Leksell Gamma Knife® Perfexion™ (PFX) (Elekta AB, 2013) stereotactic radiosurgery planning (Ghobadi, 2014; Ghobadi et al., 2013). The following section gives a brief overview of the CP algorithm, in relation to the interface.

The CP algorithm outputs treatment plans for PFX, a stereotactic radiosurgery delivery unit primarily designed to treat head-and-neck targets. The delivery component of PFX is made up of eight equidistant sources called sectors, which surround the patient's head, concentrating on a focal point within the gross

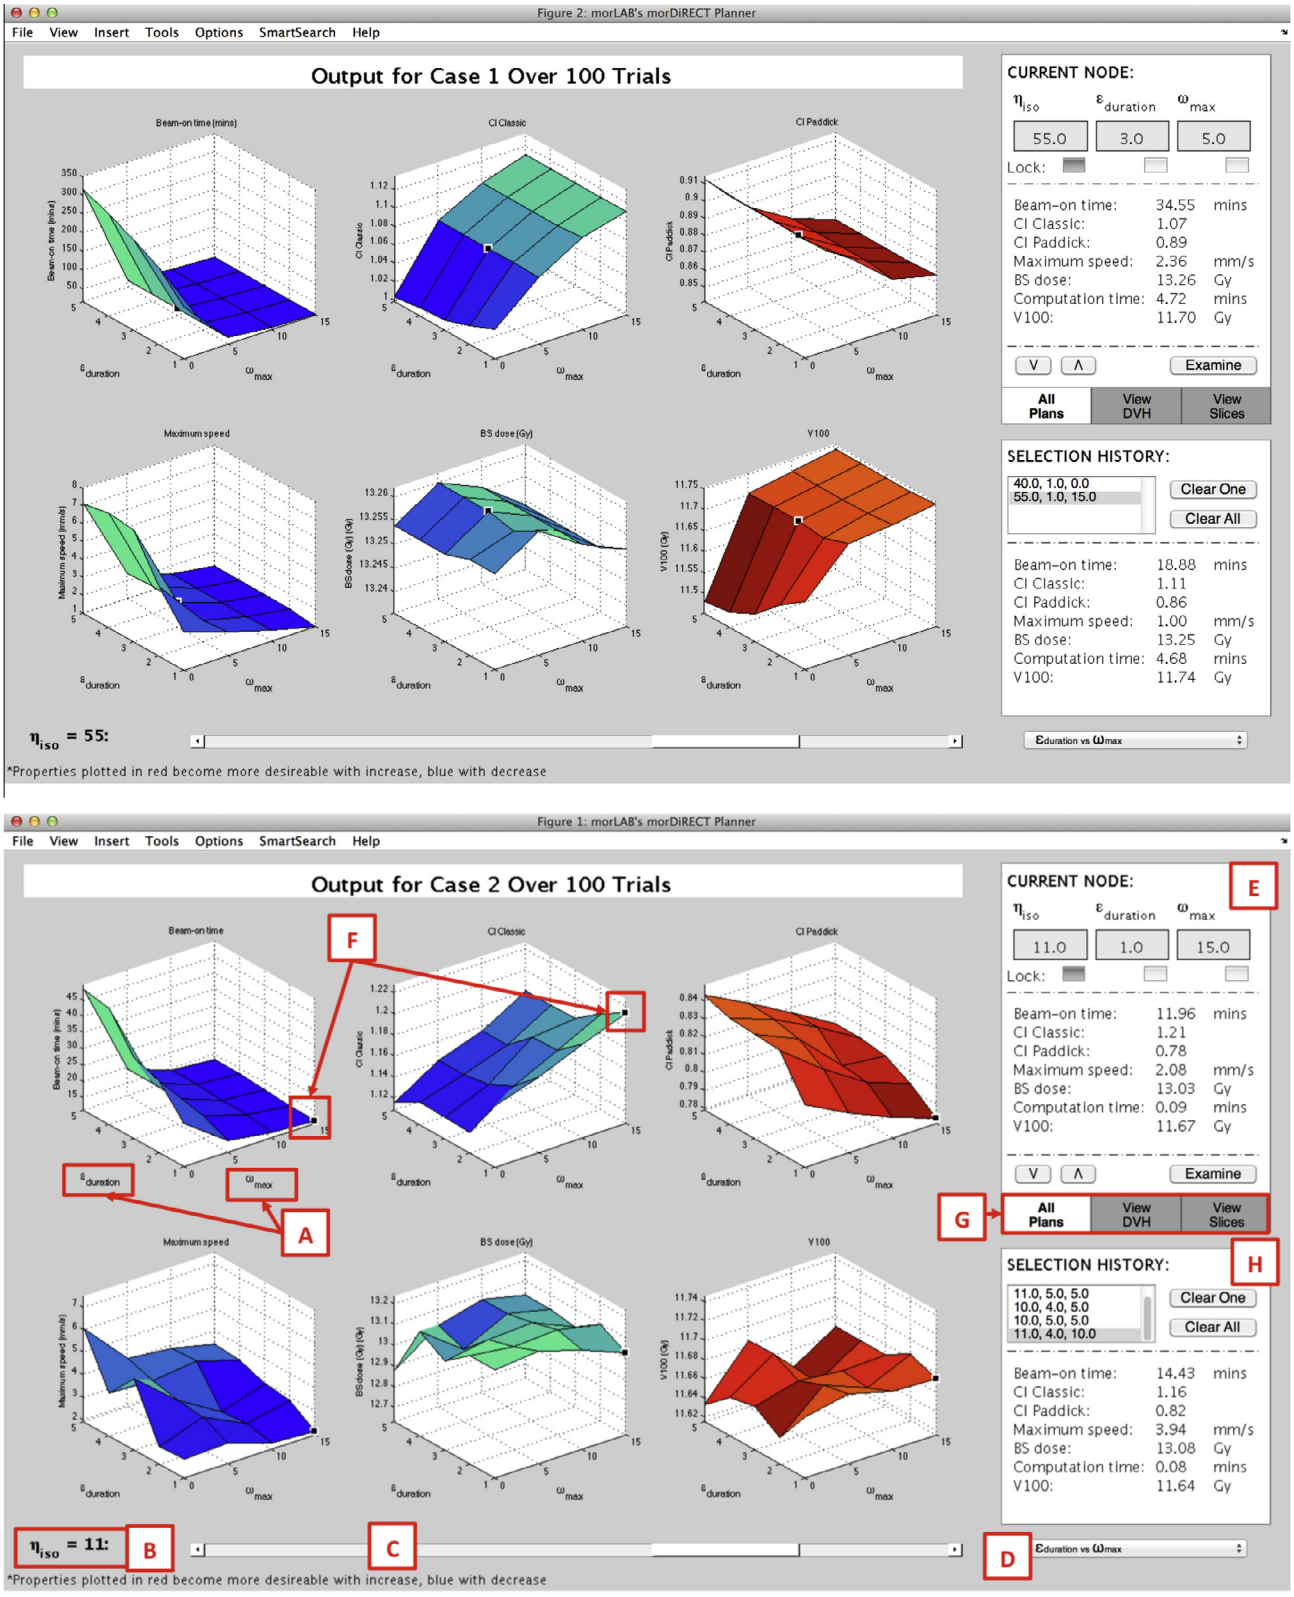

Fig. 2. Demonstration of morDiRECT's Evaluation Window, displaying analogous results from representative clinical case studies 1 (top) and 2 (bottom). A: Active driver inputs. B: Stationary projected value. C: Slider for changing value of projection. D: Pulldown menu to toggle active drivers. E: Current node summary panel. F: Plot markers. G: Plan examination tools. H: Selection history panel.

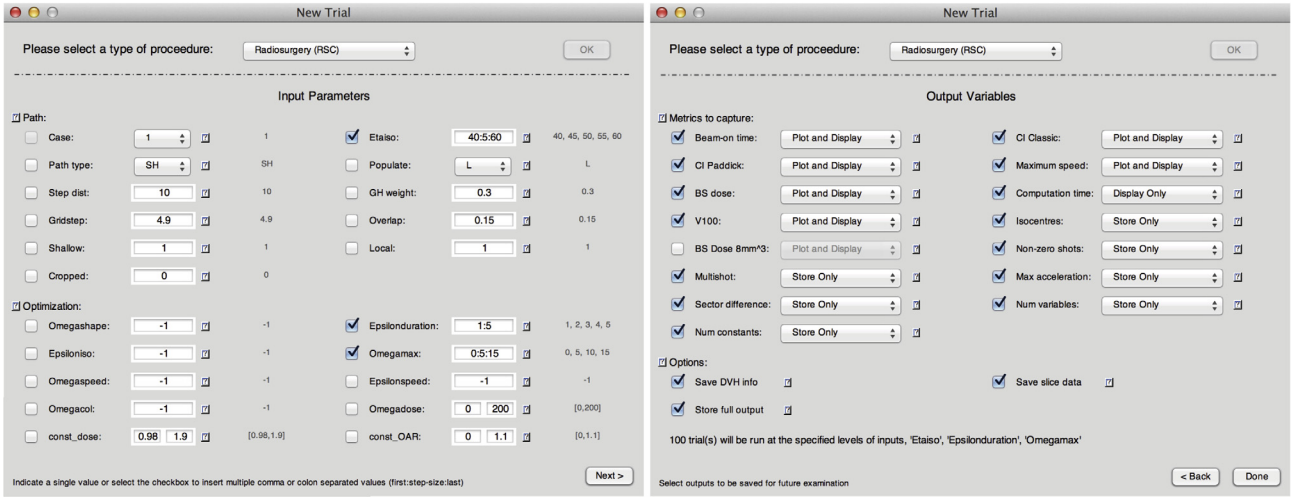

(a) Input Specification Menu (b) Output Specification Menu

Fig. 3. Input and output menus. (a) The user may adjust any input and select up to three driver inputs for which to specify a finite range of values to enumerate. (b) The user selects which output data to track, and whether to plot, display, or simply store data for future use. The selections made on both menus correspond to the upper Evaluation Window in Fig. 2.

tumour volume (GTV), called an isocentre. During the treatment delivery, the patient lies on a mechanically operated couch that shifts his/her entire body, including the GTV, aligning PFX with a series of predetermined isocentre locations. The magnitude of radiation delivered at a given isocentre is determined by the number of sectors selected for activation and their respective collimator diameters, or beam sizes. A total of three possible collimator diameter sizes are available in PFX, allowing for varying delivery shapes and conformities.

Inverse planning can be used to provide directions for PFX, indicating which isocentre to use, which sector collimator pairs to activate and the duration and magnitude of the activation. CP planning goes a step beyond the traditional inverse methodologies by enforcing extra restrictions that transform the selected isocentres into a connected path, providing a more fluid treatment delivery. This path connectivity introduces a quicker and more homogeneous treatment, but also the requirement for several new parameters to enforce mechanical limitations (e.g., couch speed).

The CP planning algorithm is a two-part process, where each part has its own set of adjustable parameters that act as input for the interface. The first part is a heuristic path selection, with parameter set $\mathcal{P} = \{P_1, P_2, \ldots, P_n\}$, and the second is a deterministic dose duration optimization, with parameter set $O = \{O_1, O_2, \ldots, O_m\}$. When the interface is used, the drivers are selected from these parameters sets such that $\mathcal{D} \subseteq \mathcal{P} \cup O$.

2.2.1. Heuristic path selection

The path selection heuristic is a grassfire and sphere-packing hybrid, which uses the volumetric features of the target to find a set of desirable isocentre locations, denoted Θ^{iso}. Conceptually, this greedy algorithm iteratively finds a point within the target that has the maximum tumour coverage. It then determines the largest allowable spherical shot that can be delivered to the area surrounding this point without exceeding the boundary of the tumour, and removes this area from the target. This process is repeated for a finite number of iterations. After all the selections have been completed, the path is made continuous using a greedy hamiltonian path and any gaps are filled with additional isocentres.

Since this algorithm is neither exhaustive nor optimal, but rather an empirically reliable method for finding a good path, many of the parameters may be adjusted according to user preference. These parameters make up set $\mathcal{P}$, which includes specification of the desired overlap between shots (“Overlap”), restriction of the maximum distance between two consecutive isocentre locations (“Step dist”) and specification of the number of isocentres selected $\eta_{\text{iso}} = |\Theta^{\text{iso}}|$ (“Etaiso”). These parameters are displayed in the “Path” section of the Input Specification Menu in Fig. 3(a), where η_{iso} is selected as one of the drivers for the CP algorithm.

2.2.2. Dose duration optimization

The results of the heuristic path selection are fed into an optimization model that determines both the magnitude and duration of radiation delivery at each isocentre location. This delivery plan is calculated through a linearized optimization model, solved using Gurobi Solver v5.6 (Gurobi Optimization, Inc.). The model is designed to penalize dose delivered to healthy tissue surrounding the tumour region, while enforcing a number of mechanical and clinical restrictions that ensure the target can be feasibly treated. A simplified representation of the objective function can be defined as follows:

$$\text{minimize} \quad f(z_r) = W_r(z_r - U_r)$$

where z_r is the dose delivered to a predetermined ring surrounding the target area, U_r is the maximum allowable dose to be delivered to this ring and W_r is the penalty weighting for overdosing.

The other major goals, such as restricting the allowable dose to the healthy and targeted tissues, are enforced in the constraint section of the model:

$$L_s \leqslant z_s \leqslant U_s$$

where U_s and L_s are the upper and lower limits on the dose z_s delivered to structure s, respectively. The set of weightings, thresholds and bounds from this model make up the set O of input parameters, displayed in the “Optimization” section of the Input Specification Menu (Fig. 3(a)). The majority of the remaining parameters in this set come from enforcing subjective restrictions. Examples include time differential limit, ϵ_{speed}, which prevents rapid accelerations and decelerations during delivery, but may detract from the overall quality of the plan; dose uniformity threshold $\epsilon_{\text{duration}}$, which encourages uniform treatment delivery, but may hinder plan flexibility; and weighting parameter $\omega_{\max}$, which restricts the largest

delivery time at a given isocentre, but may be detrimental for accuracy. Parameters $\epsilon_{\text{duration}}$ and $\omega_{\max}$ were selected as the second and third driver inputs for the CP algorithm.

2.3. Decision analysis features

The morDiRECT Evaluation Window (Fig. 2) displays the user-selected output metrics in a surface plot-based format. The plans are sorted onto the x- and y-axes based on the driver inputs used to generate the specified plan (Fig. 2, A). The z-axis on a given plot is used to represent how the plan is scored in terms of a single output metric. One complete plan description corresponds to a node taken from the same (x, y) point on the coordinates axes across every plot in a given display window (Fig. 2, F). Users may navigate this collection of plans using a combination of built-in navigation and analysis functions, with the ultimate goal of developing a better understanding of available treatment options.

2.3.1. Navigation tools

Plot navigation. Since there are only two axes available to indicate input parameters, but up to three driver inputs, any single view will only reveal a subset of the plans, corresponding to a specified value of the third driving parameter, or the projected value (Fig. 2, B). Using the slider (Fig. 2, C), the user can navigate to other values of projection and observe how the plot contours change as the projected value changes. The user may also toggle the drivers assigned to the axes and projected value using the pulldown menu (Fig. 2, D).

Plan summary tool. When the user clicks a node on any plot, the full information from the corresponding plan is automatically transferred to the Current Node summary panel (Fig. 2, E). From this panel, the user is given a number of alternatives, such as displaying detailed information about the plan either as dose volume histograms (DVHs) (Fig. 4(a)) or as isodose lines on slices, which are depictions of a given plan's dose distribution contours overlaid on 2D images of the target region (Fig. 4(b)), through the Plan Examination tabs (Fig. 2, G). The plan may also be sent to the History panel, discussed in Section 2.3.2.

Axis locking tool. The locking mechanism, located on the top of the Current Node panel allows the user to reduce the dimensionality of the graph by focusing on a smaller subset of plan alternatives. To utilize this mechanism, the user selects a point on one of the plots with a driver value of interest, followed by selecting the "Lock" control underneath that driver on the Current Node panel. The interface responds by temporarily clearing away all plans that do not correspond to the selected driver, reducing the plots to a single dimension, as in Fig. 5.

2.3.2. Analysis tools

History panel. Candidate plans can be sent to the History panel (Fig. 2, H) using the down arrow in the Current Node panel and later retrieved using the neighbouring up arrow. Any number of plans may be stored in the listbox and browsed at any time, allowing the user to aggregate and eliminate potential plan alternatives.

SmartSearch. The search tool (Fig. 6(a)) allows the user to instantly navigate to the plan with the best value of any one output metric. Clicking the name of the output automatically navigates the Evaluation Window view to the correct projection, highlighting the selected plan across all plots and displaying its metrics in the Current Node panel.

Custom search. The last option in the SmartSearch tool is a Custom Search through available plans. Selecting this option pulls up a window (Fig. 6(b)) that provides users with an opportunity to rank a specified number of treatment plans in terms of all key output metrics. The importance factors given to each output are used to create a normalized utility-based function that sorts all candidate plans according to the user's preference. After the user clicks "Ok", the requested number of top plans, sorted by the utility function, are sent to the History panel. This tool may be used repeatedly with adjusted values, allowing the user to find better compromises based on a given dataset.

Pairs plots. Using the Options menu (Fig. 7(a)), the user may select any number of outputs to compare in a pairwise comparison matrix (Fig. 7(b)).

3. Results and discussion

The application of the morDiRECT interface is demonstrated using four clinical head-and-neck case studies. For each case a decision-space of 100 plans is generated and searched for desirable plans, which are then compared to their respective clinical forward (manual) plans as an illustration of morDiRECT's performance. The following section takes users through the plan specification, decision space analysis, SmartSearch and Custom Search phases applied to creating and searching the decision space for each case, yielding the results in Table 1.

3.1. Run specifications

The cases were run with equivalent sets of drivers and identical output metrics. The following driving parameters were selected for

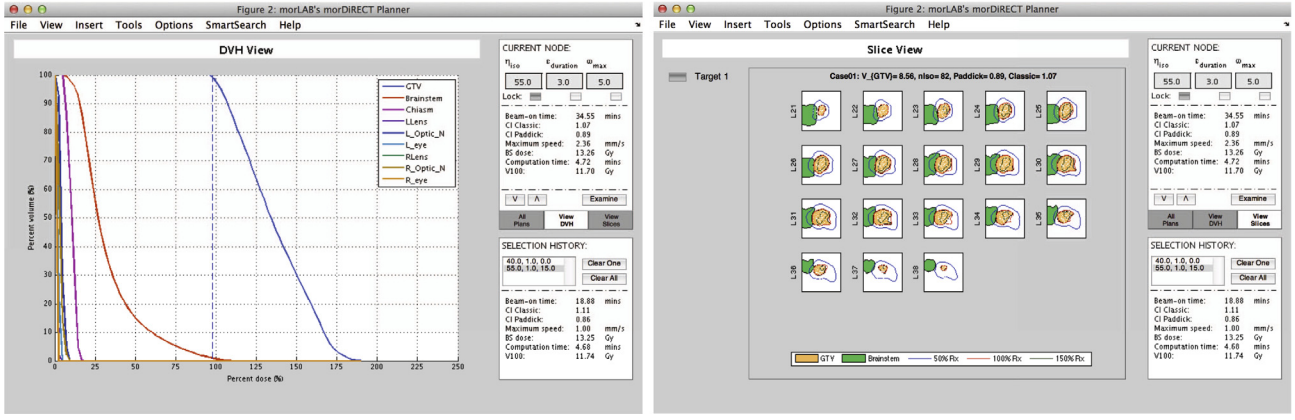

(a) Dose volume histogram (DVH) (b) Dose distribution slices

Fig. 4. The Plan Examination tabs allow users to view DVHs and isodose lines.

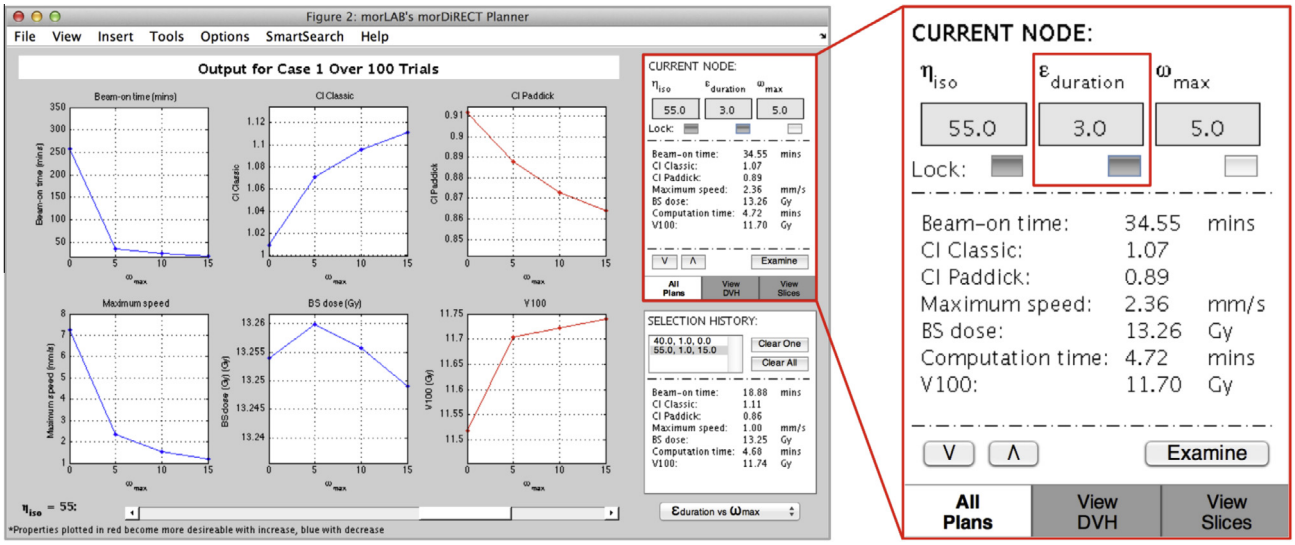

Fig. 5. The Current Node panel is used to lock the display at $\epsilon_{\text{duration}} = 1$, leaving $\omega_{\max}$ as the only active driver input, since the projected value ($\eta_{\text{iso}} = 55$) is already locked.

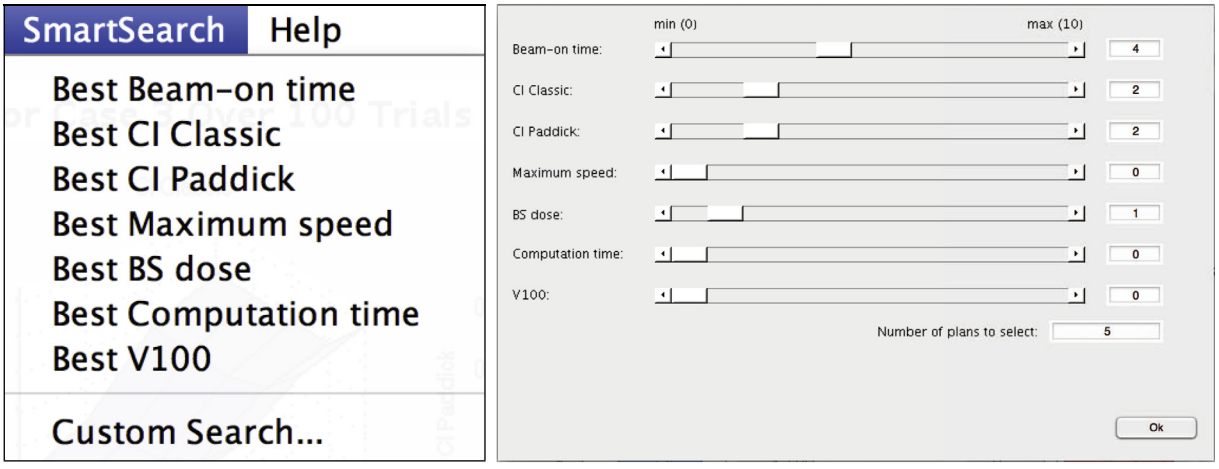

(a) SmartSearch tool (b) Custom Search tool

Fig. 6. Search tools. (a) The SmartSearch tool provides shortcuts to navigate to the best plan for the selected output metric; (b) The Custom Search tool allows the user to set a custom utility preference, optionally generating a ranked list of desirable outputs to be sent to the history panel for review.

running all four cases, due to their ability to produce empirically diverse results:

$$\mathcal{D} = \{\epsilon_{\text{duration}}, \omega_{\max}, \eta_{\text{iso}}\},$$

where $\epsilon_{\text{duration}} = \{1,2,3,4,5\}$, $\omega_{\max} = \{0,5,10,15\}$ and $\eta_{\text{iso}} = \{8s, 9s, 10s, 11s, 12s\}$. For our η_{iso} driver, the constant is an isocentre density approximation and s is a size multiplier assigned to account for the relative volume and number of voxels in a given GTV, denoted by ratios R_{vol} and R_{vox}, respectively. We used the following equation:

$$s = round((R_{vol} + R_{vox})/R_{vox})$$

to obtain values of 5, 1, 4 and 8 for cases 1, 2, 3 and 4, respectively.

The output metrics collected from each plan and a short description are provided in Table 2. The first five metrics are the focus of our analysis, since maximum speed was satisfied for all plans and computation time is not an indicator of treatment plan performance, but rather the optimization process itself. We use the information in the Desired Results column of Table 2 to guide our plan analysis, as a rough substitute for an oncologist's intuition that would be used to determine the suitability of a plan in practice. It should be noted that the V_{100} values utilized in the cases were not normalized, despite the common practice, in an effort to demonstrate the full extent of the tradeoffs available to the user.

3.2. Decision space analysis

The first stage of making an informed plan selection is understanding the range and diversity of the potential plans available, which in turn, sheds light on the nature of the tradeoffs at hand. Utilizing morDiRECT, a user can obtain this information at a glance. Each plot axis in the Evaluation Window is pinned to the upper and lower bounds of its respective output metric, helping the user visualize any deviation or clustering across the plots. A subset of this visual information is presented numerically as summary statistics for each of the four case studies in Table 3.

The output metric ranges will drive the focus of the user while selecting plans. For example, Cases 1 and 3 have a very narrow range for BS dose, within which, all plans are clinically acceptable according to our guidelines in Table 2. For this reason, the amount

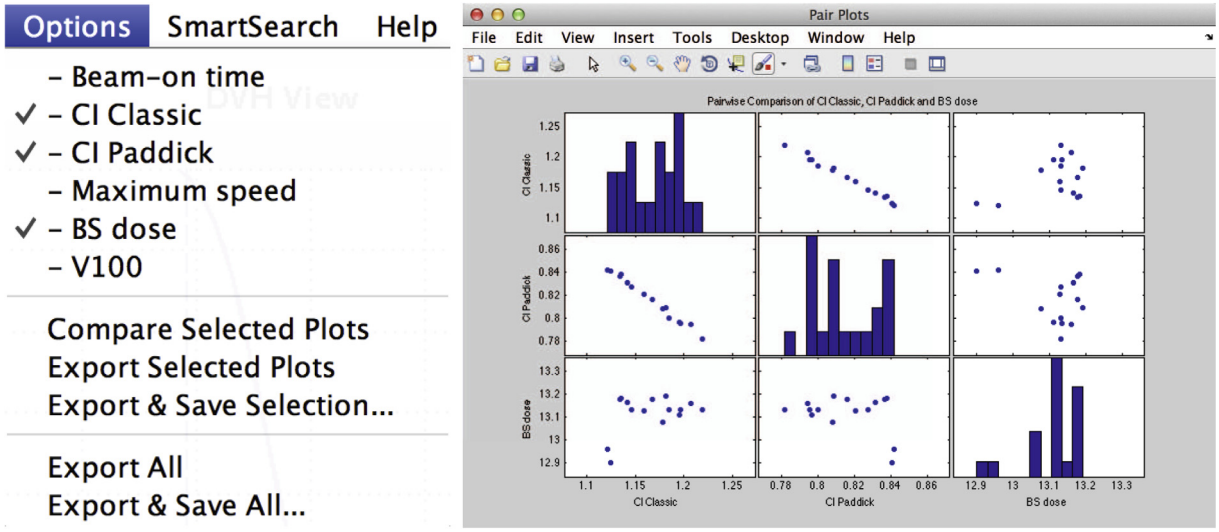

(a) Options menu (b) Pairwise comparisons

Fig. 7. Using the options menu (a), users can select the metrics to be compared via pairwise comparison plots (b).

Table 1
Summary of the findings from each of the four cases, with regards to their ability to dominate their respective manual forward plans.

Case	Dominating SmartSearch plans	Dominating Custom Search plan	Dominates forward
1	2	–	✓
2	0	4	✓
3	0	0	×
4	0	4	✓

of dosage delivered to the brainstem is unlikely to be a driving element, despite its overall importance to clinical success. Case 4 appears to contain a number of risky plans that will overdose the patient's brainstem, making BS dose a primary concern. Finally, for Case 2, the significance of BS dose will depend on tradeoff analysis and user intuition, since the values are all within the acceptable range, but some are significantly lower than others, which is favourable for patients who undergo multiple sessions or treatments.

In terms of finding good solutions for Case 1, the user will likely be concerned with the four remaining metrics, with a possible emphasis on the BOT metric, since it climbs nearly as high as six hours at the extremes. This value is around eleven times as large as the clinically developed forward plan and expert users would recognize this treatment time as prohibitively large. Similarly, Case 2 may benefit from an emphasis on conformity indices, Case 3 suffers in terms of BOT and conformity and Case 4 is a difficult case that requires vigilance across the board.

Table 2
Output metric guidelines.

Output metric	Description	Desired result
Beam-on time (BOT)	Total duration of the radiation delivery	Lower is better
Classic conformity index (CI)	Prescription dose (Rx) volume/target volume	Closer to 1.0 is better
Paddick CI	$(\text{Rx to target})^2/(\text{target volume} \times \text{Rx volume})$	Closer to 1.0 is better
Brainstem (BS) dose	Gy of radiation delivered to the brainstem	$\leqslant$15 Gy, lower is better
V_{100}	Volume of target receiving at least 100% of Rx	$\geqslant$ 98%, higher is better
Maximum speed	Fastest couch speed during treatment	$\leqslant$10 mm/s
Computation time	Computational time for the optimization	Lower is better

3.3. Application of SmartSearch

Using morDiRECT's decision support tools is far quicker than manually browsing and provides much more information than aggregates,which neglect the key interactions between metrics. Table 4 displays plans that were obtained by selecting the best plan for each output metric using SmartSearch compared to their manual forward plans as a loose indication of plan quality.

It is interesting to note, for example, that while Case 1's worst case BOT is prohibitively large, the resulting plan does have some merits, including a nearly perfect Classic CI, which for some users, may be a tradeoff worth considering. It is more likely, however, that due to the extreme nature of the tradeoff, the user would be interested in choosing a more well rounded plan from one of the 99 alternatives. We will proceed under this assumption in demonstrating the SmartSearch tool, which we recommend activating at early stages in the decision process due to its simple operation, requiring only two mouse clicks to obtain a potential treatment plan.

For Case 1, three of the five SmartSearch plans dominate the forward plan. On closer examination, however, the BOT and V_{100} SmartSearch plans have the same driver inputs, meaning they are actually the same plan. These two dominating Case 1 SmartSearch plans can then be compared in more detail, using the DVH and isodose slice viewers (Fig. 8). In an ideal DVH plot, 100% of the GTV would receive the prescribed dose while all other healthy tissue would receive none. Similarly, each ideal isodose slice would have the 100% Rx line following the exact contour of the GTV with no contours overlaying sensitive organs. Since ideal plans are not practically attainable, the generated plans do deliver some dose to sensitive organs, however, they are both clinically acceptable and comparable to one another, with very similar DVHs and isodose slices. The user would decide whether it is beneficial to spend an extra 8 min BOT in exchange for a slight improvement in conformity and brainstem avoidance, in which

Table 3
Summary statistics of output metrics.

	Figure	BOT	Classic CI	Paddick CI	BS dose	V_{100}
Case 1	Range	[17.7, 350]	[1.00, 1.13]	[0.84, 0.91]	[13.2, 13.3]	[95, 98]
	Mean	79.0	1.08	0.88	13.2	97
	Std dev	98.4	0.04	0.02	0.01	0.63
Case 2	Range	[11.2, 49.6]	[1.11, 1.23]	[0.78, 0.85]	[12.6, 13.2]	[97, 98]
	Mean	20.5	1.17	0.81	13.1	97
	Std dev	10.1	0.03	0.02	0.09	0.22
Case 3	Range	[21.2, 207]	[1.07, 1.24]	[0.80, 0.92]	[13.3, 13.3]	[98, 99]
	Mean	59.6	1.14	0.87	13.3	99
	Std dev	50.1	0.03	0.02	0.00	0.22
Case 4	Range	[28.1, 408]	[1.15, 1.43]	[0.68, 0.82]	[14.3, 15.1]	[96, 98]
	Mean	99.4	1.30	0.73	14.9	98
	Std dev	114	0.08	0.04	0.16	0.47

Table 4
SmartSearch plans for each output metric compared to the manual forward plan. Bolded values meet or exceed the forward plan quality for that metric.

	Plans	BOT (min)	Classic CI	Paddick CI	BS dose (Gy)	V_{100} (%)	Dominates forward	Driver inputs		
								η_{iso}	$\epsilon_{duration}$	ω_{max}
Case 1	Forward	32	1.14	0.85	14.4	98	–	–	–	–
	BOT	**17.7**	**1.12**	**0.86**	**13.3**	**98**	✓	50	1	15
	Classic	350.0	**1.00**	**0.91**	**13.3**	95	×	60	5	0
	Paddick	311.6	**1.02**	**0.91**	**13.3**	97	×	45	5	0
	BS dose	**25.5**	**1.11**	**0.86**	**13.2**	**98**	✓	40	1	10
	V_{100}	**17.7**	**1.12**	**0.86**	**13.3**	**98**	✓	50	1	15
Case 2	Forward	34	1.15	0.81	14.6	96	–	–	–	–
	BOT	**11.2**	1.22	0.78	**13.1**	**98**	×	10	1	15
	Classic	49.6	**1.11**	**0.85**	**12.6**	**97**	×	12	5	0
	Paddick	49.6	**1.11**	**0.85**	**12.6**	**97**	×	12	5	0
	BS dose	49.6	**1.11**	**0.85**	**12.6**	**97**	×	12	5	0
	V_{100}	**14.9**	1.19	0.80	**13.2**	**98**	×	8	4	15
Case 3	Forward	24	1.20	0.82	14.2	100	–	–	–	–
	BOT	**21.3**	**1.15**	**0.86**	**13.3**	99	×	48	4	15
	Classic	207.1	**1.07**	**0.91**	**13.3**	98	×	48	5	0
	Paddick	195.9	**1.08**	**0.92**	**13.3**	98	×	48	4	0
	BS dose	**21.3**	**1.15**	**0.86**	**13.3**	99	×	48	5	15
	V_{100}	43.0	**1.20**	**0.83**	**13.3**	99	×	32	2	15
Case 4	Forward	61	1.40	0.69	14.9	98	–	–	–	–
	BOT	**28.1**	**1.34**	**0.71**	15.1	**98**	×	96	4	15
	Classic	374.7	**1.15**	**0.82**	**14.6**	97	×	96	4	0
	Paddick	374.7	**1.15**	**0.82**	**14.6**	97	×	96	4	0
	BS dose	407.9	**1.15**	**0.81**	**14.3**	96	×	96	5	0
	V_{100}	**38.7**	**1.38**	**0.70**	15.1	**98**	×	64	2	10

case s/he would select the BS dose plan. Alternatively, the user could continue to search for different plan alternatives.

The plans selected for Cases 2 through 4 fail to dominate the forward plans using SmartSearch. Depending on the opinion of the user, the plan selected for Case 2 containing the best Classic and Paddick CI's as well as the lowest BS dose, may in fact be the preferred plan overall, despite its relatively high BOT and low V_{100}. If this were the case, the user would terminate the search, or at the very least, keep the plan in the selection history panel for future reference. For the purpose of this case study, however, we will continue to search Case 2's plan-space for a dominating plan.

Case 3, can never dominate its forward plan using the given set of driving parameters, since the best V_{100} is still less than the forward plan's perfect V_{100}. This situation is analogous to a real-world failure of the interface's inverse planning algorithm to meet clinical criteria, which can occur if the case is too difficult, criteria too ambitious or if, as in our case, the chosen driving parameters do not correspond with the desired region of the feasible treatment plan space. In such cases, although the entire algorithm may be rerun with different driving parameters in the hope of finding better solutions, it is also possible and time-saving to search the current plans for interesting and perhaps acceptable treatment plans that offer valuable tradeoffs. For instance, we note the tradeoffs made by the optimization for the best BOT and BS dose plans resulted in a far lower dosage to the patient's brainstem than the forward plan, while still maintaining a nearly perfect V_{100}. Since these plans both dominate the forward in all but the V_{100}, it is likely that they would be preferred over the forward plan by most users.

Finally Case 4 failed to dominate across the board, but unlike the previous three cases, none of the plans were even close to acceptable. The best BOT and V_{100} plans deliver far too much dosage to the brainstem, while the higher conformity and BS dose plans would take hours to run, making all five plans undesirable.

3.4. Application of Custom Search

Since Table 4 indicates that the more promising Case 2 and 4 SmartSearch plans fell short of the user's expectations, more balanced plans can be found using the Custom Search tool. The weightings chosen to yield the custom plan, shown in Table 5, are roughly based on the knowledge of the decision space gained from the previous two analysis stages. In selecting the custom weightings for Case 2, we placed a considerable importance

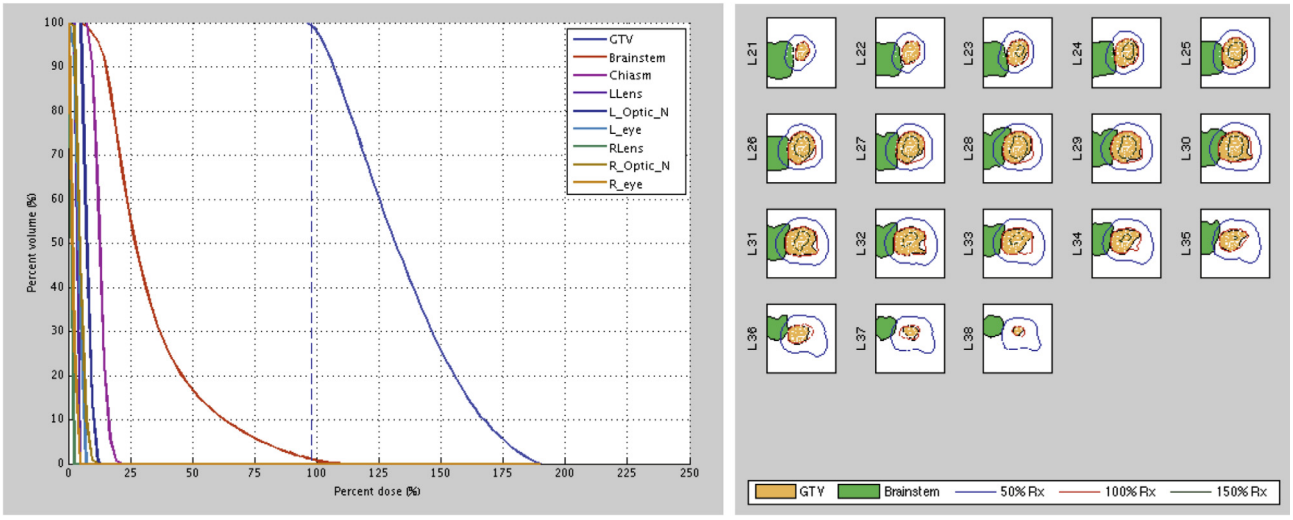

(a) BOT/V_{100} SmartSearch plan

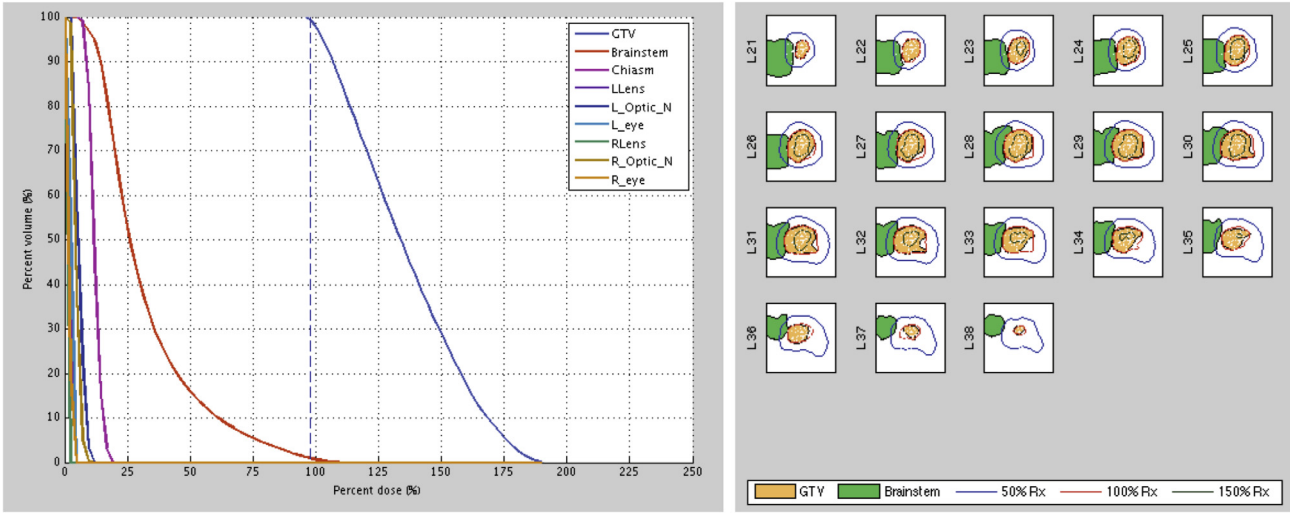

(b) BS dose SmartSearch plan

Fig. 8. Visualization of the Case 1 dominating SmartSearch plans.

Table 5
Custom plans from Cases 2 and 4, and the weightings applied, compared to the manual forward plans. Bolded values meet or exceed the forward plan quality for that metric.

	Plans	BOT (min)	Classic CI	Paddick CI	BS dose (Gy)	V_{100} (%)	Dominates forward	Driver inputs		
								η_{iso}	$\epsilon_{duration}$	ω_{max}
Case 2	Forward	34	1.15	0.81	14.6	96	–	–	–	–
	Custom 1	**19.2**	**1.14**	**0.83**	**13.0**	**97**	✓	11	4	5
	Custom 2	**19.2**	**1.14**	**0.83**	**13.0**	**97**	✓	11	5	5
	Custom 3	**19.2**	**1.14**	**0.83**	**13.2**	**97**	✓	10	4	5
	Custom 4	**19.2**	**1.14**	**0.83**	**13.2**	**97**	✓	10	5	5
	Custom 5	**14.4**	1.16	**0.82**	**13.1**	**97**	×	11	4	10
	Weights	4	2	2	1	0				
Case 4	Forward	61	1.40	0.69	14.9	98	–	–	–	–
	Custom 1	265.9	**1.17**	**0.80**	**14.4**	97	×	88	2	0
	Custom 2	**30.6**	**1.36**	**0.71**	**14.9**	**98**	✓	88	2	15
	Custom 3	**37.7**	**1.39**	**0.70**	**14.9**	**98**	✓	64	1	10
	Custom 4	**38.5**	**1.38**	**0.70**	**14.9**	**98**	✓	64	3	10
	Custom 5	**39.2**	**1.38**	**0.70**	**14.9**	**98**	✓	64	4	10
	Weights	4	1	1	6	3				

weighting on BOT, to circumvent the issue of excessively long treatment times, but without neglecting the conformity indices, which suffer when BOT is given exclusive preference. From our analysis of Table 3, we had observed that the Case 2 BS dose has a significantly wide range, so our function incorporates this metric into the selection process as well. For the purpose of dominating the forward plan, it is likely safe to neglect V_{100}, since its entire range dominates the forward plan's value of 96, although, in reality an expert may choose to value these features differently. Based on these conclusions, the Custom Search menu was filled in with the weightings, and the five best plans returned are presented in Table 5.

We follow a similar process for Case 4, noting that BS dose is a major issue, we give it the largest weighting, followed by BOT, which is found to have an alarmingly high upper bound, then V_{100}. Again, for the purpose of forward plan domination, we apply little importance to the CI's, since the forward plan is known to have weak conformity.

In both cases, four of the five custom plans dominate the forward plan, providing viable and balanced plan options and illustrating that with some intuition into the decision space, high-quality plans can quickly be found. It is also interesting to note that in the absence of a human expert, the weightings selected may lead to extreme tradeoffs that do not address the user's concerns. An example of this behaviour is the first custom plan for Case 4, which is unacceptable in terms of V_{100} and BOT, but is still chosen due to its excellent conformity and BS dose metrics.

4. Conclusions and future work

We presented the morDiRECT interface, a multi-criteria decision support system, which facilitates both the generation and decision making components of radiation therapy planning. Applying this tool to four clinical case studies, we demonstrated that high-quality plans can be easily generated, without the iterative process characteristic of radiation therapy planning. We used an intuitive process to quickly find excellent treatment plans from a selection of 100 potential plans for each case. We thus reaffirmed that supported decision making increases transparency and tradeoff recognition in an otherwise inaccessible black box function (i.e., inverse planning software).

Unlike many approaches presented in the literature, morDiRECT enables experts to control the outcome of the decision making process, rather than endeavouring to replace human input with automated class-based, statistical or plan ranking algorithms. Using real-time, human decision making provides a distinct advantage over purely automated methods in terms of flexibility in instances of difficult-to-quantify tradeoffs. It also provides benefits over other interface-based methods by removing the requirement for excessive *a priori* decision making and restrictive Pareto optimality constraints, that fail to capture potentially desirable tradeoffs.

There are, however, certain limitations associated with morDiRECT's approach. Users need to provide input ranges at the beginning of the process, which may not always be intuitive or understood by inexperienced operators. The general heuristic input ranges in this paper can be applied, but there is no guarantee of their performances for a different application. There are also repetitive calculations involved in the plan generation that are both memory and time intensive, requiring large amounts of computational power, such as a cluster, which may not be a universally viable option.

Future work will include the development of faster methods of plan generation, since computation time is a limiting factor for the tractability of the morDiRECT for larger scale cases. We would also like to investigate new ways of automating the user input component of the interface, to shift the emphasis further towards user decision making and away from any *a priori* input considerations.

Acknowledgements

This work was funded in part by Elekta Instrument, AB, Stockholm, Sweden, The Canada Foundation for Innovation, The Ontario Research Fund, and the Natural Sciences and Engineering Research Council of Canada (NSERC) PGS-D program.

References

21st Century Oncology. (2013). *Oncor digital medical linear accelerator*. <http://www.21stcenturyoncology.com/radiation-oncology/about-radiation-therapy/technology/oncor.aspx> (last accessed on February 12, 2014).

Aigner, W., & Miksch, S. (2006). Carevis: Integrated visualization of computerized protocols and temporal patient data. *Artificial Intelligence in Medicine, 37*(3), 203–218.

Aubry, J. F., Beaulieu, F., Sévigny, C., Beaulieu, L., & Tremblay, D. (2006). Multiobjective optimization with a modified simulated annealing algorithm for external beam radiotherapy treatment planning. *Medical Physics, 33*, 4718–4729.

Bonczek, R. H., Holsapple, C. W., & Whinston, A. B. (2014). *Foundations of decision support systems*. Academic Press.

Bortz, M., Burger, J., Asprion, N., Blagov, S., Böttcher, R., Nowak, U., et al. (2014). Multi-criteria optimization in chemical process design and decision support by navigation on Pareto sets. *Computers & Chemical Engineering, 60*, 354–363.

Breedveld, S., Storchi, P. R. M., Keijzer, M., Heemink, A. W., & Heijmen, B. J. M. (2007). A novel approach to multi-criteria inverse planning for IMRT. *Physics in Medicine and Biology, 52*(20), 6339–6353.

Cobuloglu, H. I., & Büyüktahtakın, İ. E. (2015). A stochastic multi-criteria decision analysis for sustainable biomass crop selection. *Expert Systems with Applications, 42*(15), 6065–6074 (last accessed on June 10, 2015).

Cotrutz, C., & Xing, L. (2002). Using voxel-dependent importance factors for interactive DVH-based dose optimization. *Physics in Medicine and Biology, 47*(10), 1659–1669.

Craft, D. L., Halabi, T. F., Shih, H. A., & Bortfeld, T. R. (2006). Approximating convex Pareto surfaces in multiobjective radiotherapy planning. *Medical Physics, 33*, 3399–3407.

Ehrgott, M., & Winz, I. (2008). Interactive decision support in radiation therapy treatment planning. *OR Spectrum, 30*(2), 311–329.

Elekta AB. (2013). *Gamma knife® surgery – stereotactic radiosurgery*. <http://www.elekta.com/healthcare-professionals/products/elekta-neuroscience/gamma-knife-surgery.html?utm_source=gamma-knife&utm_medium=redirect&utm_campaign=redirects> (last accessed on February 13, 2014).

Ferris, M. C., Lim, J., & Shepard, D. M. (2002). An optimization approach for radiosurgery treatment planning. *SIAM Journal on Optimization, 13*(3), 921–937.

Ferris, M. C., Lim, J., & Shepard, D. M. (2003). Radiosurgery treatment planning via nonlinear programming. *Annals of Operations Research, 119*(1–4), 247–260.

Ferris, M. C., & Shepard, D. M. (2000). Optimization of gamma knife radiosurgery. *Discrete Mathematical Problems with Medical Applications, 55*, 27–44.

General Electric Company. (2013). *Magnetic resonance imaging*. <http://www3.gehealthcare.com/en/Products/Categories/Magnetic_Resonance_Imaging> (last accessed on February 13, 2014).

Ghobadi, K. (2014). Optimization methods for patient positioning in Leksell Gamma Knife Perfexion, PhD Thesis, University of Toronto.

Ghobadi, K., Aleman, D. M., Jaffray, D. A., & Ruschin, M. (2013). Automated inverse planning over a continuous path using gamma knife perfexion. In *Proceedings of the ISERC annual conference, May 2013*.

Ghobadi, K., Ghaffari, H. R., Aleman, D. M., Jaffray, D. A., & Ruschin, M. (2012). Automated treatment planning for a dedicated multi-source intra-cranial radiosurgery treatment unit using projected gradient and grassfire algorithms. *Medical Physics, 39*(6), 3134–3141.

Grandjean, E., & Kroemer, K. H. E. (1997). *Fitting the task to the human. A textbook of occupational ergonomics* (5th ed.,). Taylor & Francis.

Gschwandtner, T., Aigner, W., Kaiser, K., Miksch, S., Seyfang, A., et al. (2011). Carecruiser: Exploring and visualizing plans, events, and effects interactively. In *Pacific visualization symposium* (pp. 43–50). IEEE.

Haight J. M. (2010). *Automation vs. Human intervention: What is the best fit for the best performance?* <http://www.asse.org/practicespecialties/tech/tech1.php> (last accessed on April 16, 2014).

Hanne, T., & Trinkaus, H. L. (2003). KnowCube for MCDM – visual and interactive support for multicriteria decision making. *Berichte des ITWM* (50).

IBA. (2013). *Proton therapy, changing the way we treat the cancer*. (2013). <http://www.iba-protontherapy.com/why-proton-therapy> (last accessed on February 13, 2014).

Jain, N. L., Kahn, M. G., Drzymala, R. E., Emami, B. E., & Purdy, J. A. (1993). Objective evaluation of 3-D radiation treatment plans: A decision-analytic tool incorporating treatment preferences of radiation oncologists. *International Journal of Radiation Oncology∗ Biology∗ Physics, 26*(2), 321–333.

Korhonen, P., & Wallenius, J. (1988). A Pareto race. *Naval Research Logistics, 35*(6), 615–623.

Lolli, F., Ishizaka, A., Gamberini, R., Rimini, B., & Messori, M. (2015). Flowsort-GDSS – a novel group multi-criteria decision support system for sorting problems with application to FMEA. *Expert Systems with Applications, 42*(17), 6342–6349.

Long, T., Matuszak, M., Feng, M., Fraass, B. A., Ten Haken, R. K., & Romeijn, H. E. (2012). Sensitivity analysis for lexicographic ordering in radiation therapy treatment planning. *Medical Physics, 39*, 3445–3455.

Lotov, A. V., Bushenkov, V. A., & Kamenev, G. K. (2004). *Interactive decision maps: Approximation and visualization of Pareto frontier*. *Applied optimization* (Vol. 89). Kluwer Academic Publishers.

Lourenzutti, R., & Krohling, R. A. (2014). The Hellinger distance in multicriteria decision making: An illustration to the topsis and todim methods. *Expert Systems with Applications, 41*(9), 4414–4421 (last accessed on June 10, 2015).

National Cancer Institute. (2010). *Radiation therapy for cancer*. <http://www.cancer.gov/cancertopics/factsheet/Therapy/radiation> (last accessed on February 13, 2014).

Otto, K. (2014). Real-time interactive treatment planning. *Physics in Medicine and Biology, 59*(17), 4845 (last accessed on June 10, 2015).

Petrovic, S., Mishra, N., & Sundar, S. (2011). A novel case based reasoning approach to radiotherapy planning. *Expert Systems with Applications, 38*(9), 10759–10769.

Philips. (2013). *CT scanners*. <http://www.healthcare.philips.com/main/products/ct/products/scanners/> (last accessed on February 13, 2014).

Romeijn, H. E., & Dempsey, J. F. (2008). Intensity modulated radiation therapy treatment plan optimization. *TOP, 16*(2), 215–243.

Rosen, I., Liu, H. H., Childress, N., & Liao, Z. (2005). Interactively exploring optimized treatment plans. *International Journal of Radiation Oncology∗ Biology∗ Physics, 61*(2), 570–582.

Ruotsalainen, H., Boman, E., Miettinen, K., & Tervo, J. (2009). Nonlinear interactive multiobjective optimization method for radiotherapy treatment planning with Boltzmann transport equation. *Contemporary Engineering Sciences, 2*(9), 391–422.

Schimpff, S. C. (2011). *Major advances in radiation therapy for cancer treatment*. <http://www.kevinmd.com/blog/2011/11/major-advances-radiation-therapy-cancer-treatment.html> (last accessed on February 13, 2014).

Shepard, D. M., Yu, C., Murphy, M. J., Bussière, M., & Bova, F. J. (2015). Treatment planning for stereotactic radiosurgery. In *Principles and practice of stereotactic radiosurgery* (pp. 73–94). Springer (last accessed on May 30, 2015).

Thyvalikakath, T. P., Dziabiak, M. P., Johnson, R., Torres-Urquidy, M. H., Acharya, A., Yabes, J., et al. (2014). Advancing cognitive engineering methods to support user interface design for electronic health records. *International Journal of Medical Informatics, 83*(4), 292–302 (last accessed on May 30, 2015).

Wang, J., Jin, X., Zhao, K., Peng, J., & Hu, W. (2014). Tu-c-17a-10: Patient features based dosimetric Pareto front prediction in esophagus cancer radiotherapy. *Medical Physics, 41*(6), 460 (last accessed on May 30, 2015).

Webb, S. (2014). The physical basis of IMRT and inverse planning. *British Institute of Radiology, 76*(910) (last accessed on May 30, 2015).

Wickens, C. D., Lee, J., Liu, Y. D., & Gordon-Becker, S. (2004). *An introduction to human factors engineering* (2nd ed.). Pearson.

Wu, Q. J., Chankong, V., Jitprapaikulsarn, S., Wessels, B. W., Einstein, D. B., Mathayomchan, B., et al. (2003). Real-time inverse planning for gamma knife radiosurgery. *Medical Physics, 30*(11), 2988–2995.

Wu, Q. J., Jitprapaikulsarn, S., Mathayomchan, B., Einstein, D., Maciunas, R. J., Pillai, K., et al. (2004). Clinical evaluation of a gamma knife inverse planning system. *Radiosurgery, 5*, 2988–2995.

Yasnoff, W. A., & Miller, P. L. (2014). *Public health informatics and information systems*. Springer.

Yu, Y. (1997). Multiobjective decision theory for computational optimization in radiation therapy. *Medical Physics, 24*(9), 1445–1454.