MAKING ON-LINE ADAPTIVE RADIOTHERAPY POSSIBLE USING ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING FOR EFFICIENT DAILY RE-PLANNING

Yves Archambault, Christopher Boylan, Drew Bullock, Tomasz Morgas, Jarkko Peltola, Emmi Ruokokoski, Angelo Genghi, Benjamin Haas, Pauli Suhonen, and Stephen Thompson

Varian Medical Systems, Inc. Palo Alto, California USA; Baden, Switzerland; and Helsinki, Finland

Abstract— **Adaptive therapy involves the ability to alter a radiotherapy treatment plan based on tumor and anatomical changes over a course of therapy. The goal is to better target the tumor, reduce dose to healthy tissue and potentially improve overall outcomes. To date, achieving this has typically required time-consuming re-planning between treatment sessions or monopolizing a linac for an extended period while a patient waits on the treatment couch for new plans to be generated. Neither of these alternatives has been deemed practical or affordable at scale, as very often clinics don't have the resources even if they have the tools.**

Consequently, Varian Medical Systems developed Ethos™ therapy, a radiotherapy treatment system that uses artificial intelligence (AI) and machine learning to accomplish adaptive radiotherapy. In this paper, we describe the technology that underlies the adaptive capabilities of the system.

Keywords— **Adaptive radiotherapy, artificial intelligence, machine learning, Ethos, RapidPlan, treatment planning, neural networks, .**

VII. Introduction

A number of challenges exists in the delivery of adaptive radiotherapy. Briefly, the challenges have been:

- The challenge of performing a full treatment planning workflow during a radiotherapy treatment session, while the patient is on the treatment couch, in treatment position. Treatment planning is a complicated task and requires time and significant attention and knowledge. The pressure imposed on clinicians performing on-couch adaptive therapy heightens opportunities for mistakes.
- The challenge of manually delineating influencers—anatomy that influences the shape of targets—during the detection of daily anatomy. Manual delineation of target structures for daily re-planning is time-consuming and technically challenging. The process can be particularly challenging in the presence of artifacts.
- The challenge of producing quality plans using inverse planning, which is typically performed by dedicated treatment planning staff and the manual or templated application of optimization cost-function based objectives or constraints. The skill and expertise of the treatment planner can have a major impact to the final plan quality. Additionally, plan generation in existing commercialized systems typically requires moderately complex user interactions which distract the focus during the on-couch session.

The Ethos system was designed to address these challenges. In this paper, we take a deeper look at the technologies within the Ethos system that address these challenges.

VIII. On-couch adaptive therapy workflow

The challenges involved in delivering on-couch adaptive therapy are addressed, in the Ethos system, through a re-planning workflow that has been reduced to well-defined and predictable clinical decision points in order to lower the cognitive load of the clinician.

Figure 1 (p. 3) illustrates the on-couch adaptive work-flow implemented in Ethos therapy. It guides the user by presenting focused information and asking for a single decision at the time.

- After the kV-CBCT image is acquired, it is presented to the clinician for evaluation. The clinician can either accept the image or decide to acquire a new one (Decision 1).
- Once the image is accepted, the system detects selected normal organ structures directly on the kV-CBCT. These structures are referred to as "influencer" structures. They are those structures that are in the closest proximity to the target(s) and have the biggest impact on their shape and position. The influencer structures are then presented to the clinician, who is asked to review and adjust them, and then to confirm that they are adequate (Decision 2).
- Once confirmed, the influencer structures are used to guide an algorithm that propagates the target structures from the planning, or reference, image to the kV-CBCT image. This ensures that features and relations between target and anatomy structures that were present on the reference image are preserved on the kV-CBCT image. The propagated targets and the detected normal anatomy structures comprise a new patient model—the session model. The user is asked to

review and accept the session patient model, with focus on the propagated targets (Decision 3).

- The session patient model is then used by the automated treatment planning to produce two plans. The scheduled plan is obtained by calculating dose from the reference plan on the session patient model while the adapted plan is obtained by running a new optimization and calculating dose using the session patient model. The two plans are then shown to the user so the appropriate one, scheduled or adapted, can be selected for treatment (Decision 4).

IX. Detection of daily anatomy

To address the challenges involved during detection of daily anatomy, an AI-based algorithm that is based on convolutional neural networks is used to contour the influencers on the session images.

Neural networks

A neural network is a collection of connected units or nodes called "artificial neurons" that behave much like biological neurons. They have an input layer, an output layer (prediction) and one or more hidden layers. The depth of the network depends on the number of hidden layers. Deep neural networks are neural networks with multiple hidden layers. Deep learning convolutional neural networks (CNNs) make the explicit assumption that the inputs are images, which allows for the incorporation of certain properties into their architecture. CNNs are best for solving problems related to image recognition, object detection, and other computer vision applications. A typical CNN can be viewed as a sequence of layers that transforms an image volume into an output volume.

Varian's in-house developed and trained deep learning model for Ethos therapy utilizes TensorFlow, CUDA and cuDNN libraries, and processes images on different interconnected resolution levels. Ethos uses full-image deep convolutional neural networks with tailored architectures that share many similarities with U-Net and DenseNet, which are widely used in image segmentation tasks. The network itself takes the full 3D iCBCT as an input and returns the same size of segmentation as an output. The neural network models used in the influencer segmentation process are static and do not continuously learn based on user input. This ensures the stability and performance of the algorithms over time.

Deep neural network model production

Neural network training was performed in a supervised learning setting using images and ground truth contours from several hundred patients. Data was acquired from multiple clinics across the Americas, Europe, Australia and Asia. Images for the training set were selected to represent a realistic spectrum of anatomical variety and typical image artifacts. Human anatomy experts created the ground truth contours as part of the algorithm development. A single set of contours was produced for each training image.

Training involves three separate datasets:

1. Training dataset. The training dataset is used to fit the model. This process involves utilizing a large set of consistently contoured data, which is used to perform the actual training of the neural networks. Contours in training data sets are randomly peer reviewed to ensure adherence to selected guidelines.
2. Validation dataset. The validation dataset is used to provide an unbiased evaluation of a model fit on the training dataset while tuning the model with hyperparameters. The validation data set is considered a subset of the training data set.
3. Test dataset. The test dataset is a smaller set of data used to provide an evaluation of a final model fit. Scans related to a patient that belongs to test set cannot be used for network training. Each image and contour in test set is reviewed by physicians for accuracy.

A neural network is trained using the classical backward-error-propagation algorithm. An error is computed at the output and distributed throughout the network layers. The gradient descent optimization algorithm uses back propagation to adjust the weight of neurons by calculating the gradient of the cost function. A cost function is a measure of how well a neural network performs with respect to the given training sample and the expected output. The cost function is typically expressed as a difference or distance between the predicted value and the actual value. It can be estimated by iteratively running the model to compare estimated predictions against the ground truth.

Hyperparameters in deep learning models

Hyperparameters are settings that can be tuned to control the behavior of a machine learning algorithm. Conceptually, they can be considered orthogonal to the learning model itself; although they live outside of the models, there is a direct relationship between them.

Examples of hyperparameters:

- Learning rate — the learning rate quantifies the learning progress of a model in a way that can be used to optimize its capacity.
- Number of hidden units — the number of hidden units is key to regulating the representation capacity of the model.
- Convolution kernel width — In CNNs, the kernel width influences the number of parameters in a model which, in turn, influences its capacity.

Hyperparameters may be tuned using two basic approaches: manual or automatic selection. Both approaches are technically viable but choosing between them typically represents a trade-off. The decision is related to the high

computation costs required for automatic selection algorithms. During training of the Ethos therapy deep learning models, a hyperparameter optimization is used to determine random weight initialization, as well as both the

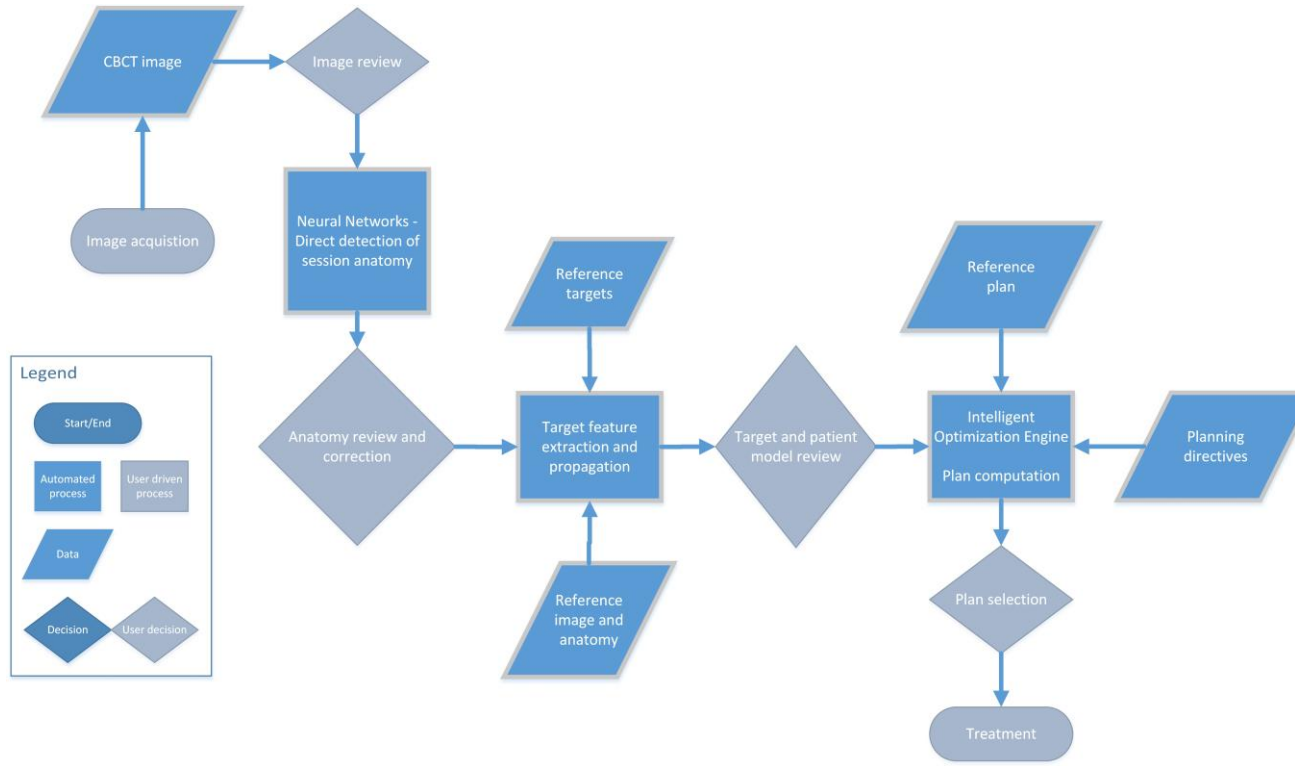

Fig. 1: Ethos therapy on-couch adaptive workflow

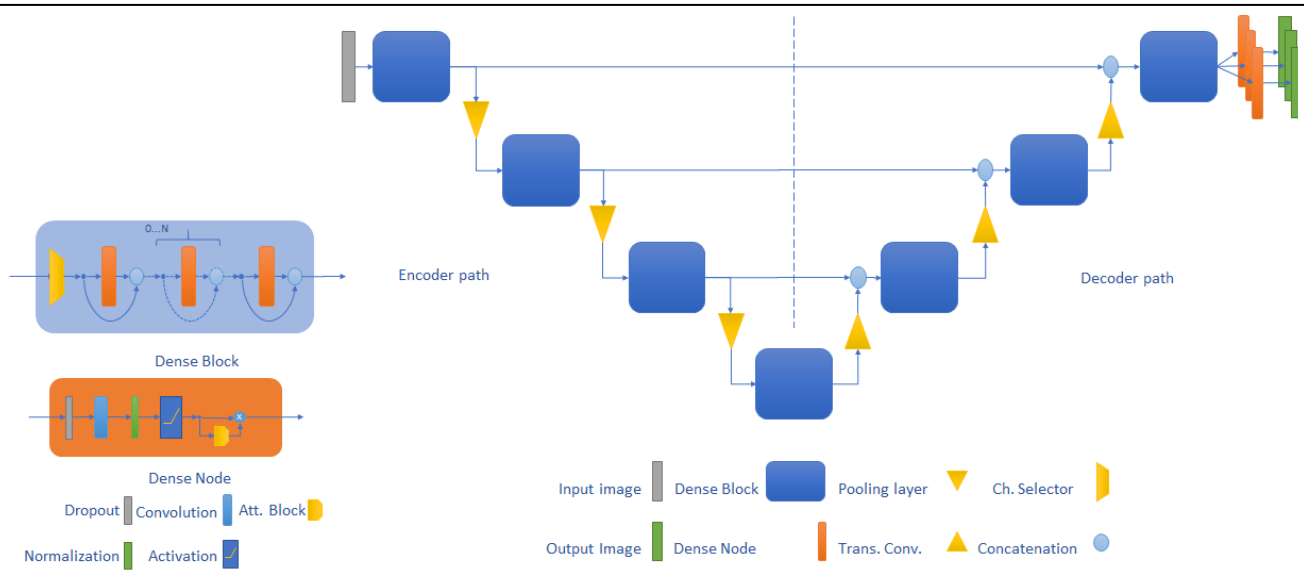

Fig. 2: An example of Varian's deep convolution neural network architecture.

loss function and the layer order.

Post-processing of convolutional neural network outputs

The segmentation output of the networks is passed through a post-processing module that ensures that the output matches the selected clinical guidelines. The processing operations include the removal of smaller segments or dislocated segments, smoothing the final contours as well as selecting terminal slices for segmentation.

Model verification and validation

Each trained neural network model undergoes verification tests that compare the obtained classification to the ground truth contours on multiple test sets. Several evaluation metrics are computed and evaluated during each verification test against passing criteria that are established for each of the evaluated structures. Models that pass verification tests qualify as candidates for validation tests. Because of low correlation between similarity metrics and measured correction and review times, the verification tests by themselves do not qualify models for deployment. Models that passed verification are then validated by clinicians in a test that better estimates the clinical review effort. Passing validation testing ensures that the model meets user needs, which qualifies the model for deployment. Performance of deployed models is monitored and compared to validation test results enabling algorithm improvement over time.

Propagation of targets

The generation and evaluation of a new treatment plans requires the fast generation of new target volume. The new target volume needs to be anatomically consistent with the initially defined target, that is, it must include the same areas of the body as the initial target. Usually, these areas are the primary tumor, the primarily affected organ, and regions where invasion of lymph nodes has been observed or is expected. The initial target volume is informed by many sources of clinical information which might include imaging, anatomical boundaries, or clinical disease spread knowledge, and therefore contains medical reasoning for which a human clinician is needed. Detecting it automatically on a new image consistently with the initial medical reasoning is thus not a straightforward process. However, finding a suitable geometric transformation that considers the large motion of organs and the partial rigidity due to anatomical circumstances at the same time, e.g, proximity to bone, is a feasible approach.

In Ethos therapy previously detected normal organ structures that strongly affect targets (so-called influencers) are used as a guidance together with partial rigidity constraints in a new deformable registration algorithm. This is a non-demons algorithm using the discretize-then-optimize approach. It is formulated as an optimization

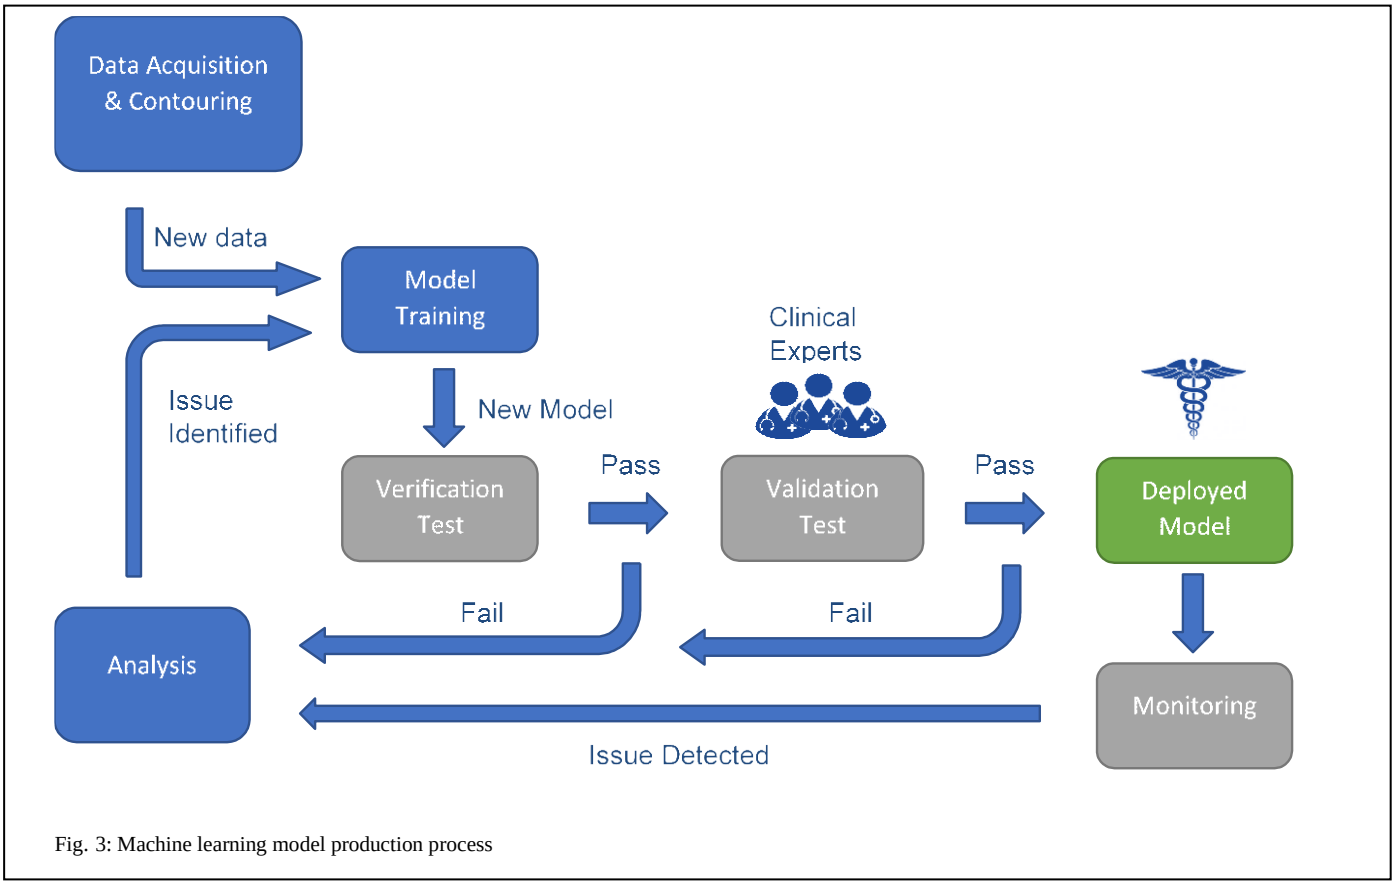

Fig. 3: Machine learning model production process

problem – NGF(CT(y),CBCT)+Curv(y)→min − using the normalized-gradient fields image similarity measure (NGF), which tries to align image edges, and the curvature regularizer (Curv), which utilizes the Laplacian to penalize large deformations y. Hence, we are solving for a deformation y so that CT(y) becomes similar to CBCT. The proposed algorithm extends the objective function with a set of penalty terms. For each delineated structure S_{CT} available on the planning CT and its corresponding structure S_{CBCT} on the CBCT, we can add a structure guidance term $\int (S_{CT} - S_{CBCT})^2 dx$ called sum-of-squared-differences. Those additional terms are driving the result of the deformable image registration towards a maximum overlap of the structures. This term enables an improved organ match for high magnitude deformations. The resulting deformable registration is then used to propagate the target from the planning CT to the new image.

X. On-couch adaptive treatment plan generation: intelligent optimization engine

The challenges involved in delivering on-couch adaptive therapy are addressed, in the Ethos system, through a re-planning workflow that has been reduced to well-defined and predictable clinical decision points in order to lower the cognitive load of the clinician.

Within Ethos therapy, the treatment planning is highly automated to allow the user to focus on the clinical aspects of the patient's therapy. In order to automate the plan generation, we introduce the Intelligent Optimization Engine (IOE), an algorithm that orchestrates the plan optimization. This algorithm aims to perform all the actions necessary to generate high-quality dose distributions that meet the clinical expectations for the plan and ensure that the plan is dosimetrically accurate. It sets up the optimization problem for the Photon Optimization algorithm and then controls and monitors the optimization process. The IOE is used in the Ethos Dose Preview workspace, which provides a fast, optimized dose distribution to check for potential clinical trade-offs, as well as in the automated plan generation, which produces IMRT and VMAT plans for a given set of inputs. In both cases, the IOE works as follows:

Pre-processing: Translation of goals to objective functions

The primary input of the IOE is an ordered list of clinical goals. The ordered list of clinical goals is created in the Ethos treatment management RT Intent module by the physician. IOE performs translation of the ordered list of clinical goals into objective functions for the Photon Optimizer and creates Quality-functions (Q-functions) to monitor and guide the progress of the optimization. Since the clinical goals from Ethos treatment management have an enforced syntax, the goal to objective function translation is straightforward. The Q-functions are described in more detail below.

Pre-processing: Overlap handling and objective setting

Prior to initiating the optimization and plan generation, the IOE performs a structure pre-processing step. In this step, the system examines the ordered goal list and the contoured organs and targets and assesses possible conflicts and overlaps between targets and organs, as well as between targets with different dose levels.

A common overlap situation that requires resolution prior to the plan generation occurs when a target overlaps with an organ and the user has specified goals for the target and the organ that conflict with one another. Due to the overlap these goals cannot be physically met. IOE uses the ordered clinical goal list to determine how to resolve these overlap situations. The IOE then creates a modified optimization structure set and adjusts the objective functions to account for the overlaps. In Dose Preview, the physician can investigate the effects of overlaps and fine-tune the clinical goal priority order prior to authorizing the RT Intent and starting the automated plan generation.

Figures 4 and 5 (p. 6) show some overlap examples and how they are resolved.

Optimization progress monitoring: Q-functions

The IOE establishes a set of piecewise continuous "quality" functions (Q-functions) for driving the plan quality optimization. The Q-functions are derived from a set of prototype functions per type of clinical goal. Examples of the different types are target lower dose (TLD) goals (goals which specify the minimum dose desired for a target), target upper dose (TUD) goals (goals which specify the tolerated maximum dose to the target), and organ upper dose (OUD) goals. The functional form of each prototype is based on the known features of a good dose distribution as described in the next chapter. The Q-functions are formed from the prototype functions by inserting the goal priority and relative goal value (dose or volume). This places the functions on the Priority-Quality-plane (P, Q) in such a way that the Q function goes through the goal point (Pi,Qi), where Pi is the priority for goal i and Qi is mapped to the relative goal value (volume or dose) for that goal.

The Q-functions for TLD goals are increasing for P<Pi. This advises the optimizer to improve the dose for the target if the plan quality metric is smaller than Pi. For P>Pi the TLD Q-functions are constant P=Qi. This signals the optimizer that there is no need to improve the achieved value for the goal once the goal is met. TUD and OUD Q-functions start as constant P=Qh for P<Pi-1, where Qh is a large Q-value. This ensures that while P is in this range, the goal plays no role in the optimization. After the constant part the functions have a steep decrease towards the goal value Qi. This advises the optimizer to work on these goals. Furthermore, for some P>Pi the TUD Q-functions decrease

towards the prescribed Q-value and the OUD Q-functions towards a Q-value that can be met without compromising higher priority goals. This guides the optimizer to try further improvements even if the goal is met.

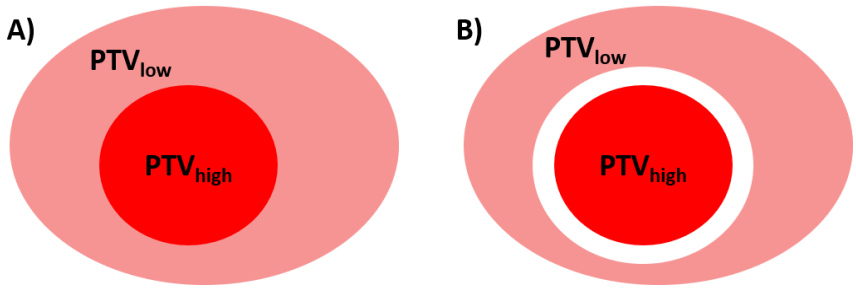

Fig 4. Conflict from multiple overlapping targets. **A)** High dose target (PTV_{high}) overlapping with lower dose target (PTV_{low}). **B)** User specifies a goal for maximum dose in PTV_{low} which is in direct conflict with the minimum dose goal for PTV_{high}. The system crops the PTV_{low} from PTV_{high} (with some margin for dosimetric fall-off) and then applies a maximum dose objective which follows the needed form for the PO algorithm (created from the input maximum dose goal) to the remainder of PTV_{low}.

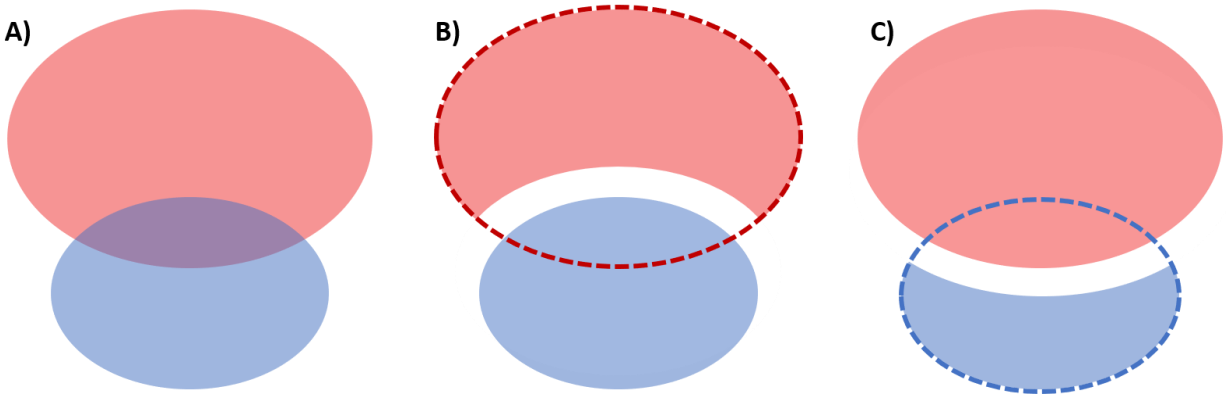

Fig 5. Conflict from organ (blue) overlapping with target (red). **A)** Target in shaded red with overlapping organ in shaded blue**. B)** Organ goal has higher priority than target goal: the target is cropped with the organ including a margin for dosimetric falloff. **C)** Target goal exists with higher priority than organ goal: the oar is cropped with the target, again with a margin for dosimetric falloff. The dashed lines show the unmodified structure outline.

Figure 6 (next page) provides an example of the Q functions for a set of 3 goals, one for each goal type: target upper, target lower, and organ upper. Once Q-functions are established, the system starts the optimization and then interrogates the achieved values for each goal at specified intervals (certain number of iterations) and then uses the associated Q-function to determine an achieved Pa value for the goal. The goal of the IOE is to maximize the collection of P values in a given optimization. Figure 7 illustrates how the Q-functions are utilized to monitor the progress of the optimization.

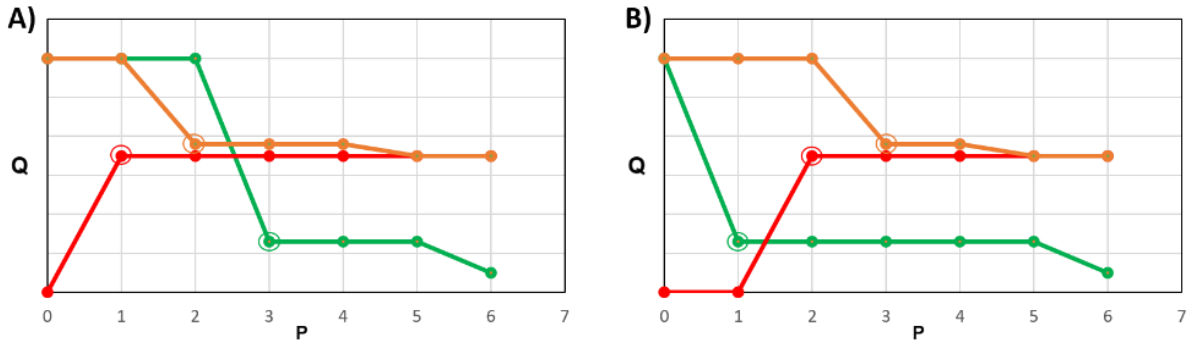

Fig 6. Intelligent Optimization Engine Q-functions (goal functions). In this example, the vertical axis in each panel is a quality measure for a goal. A goal meets its quality measure if the achieved value is the same or better than the goal value. The horizontal axis is a depiction of the relative priority of the set of goals. **A)** Red line is a TLD Q-function with priority 1, orange line is a TUD Q-function with priority 2, green line is an OUD Q-function with priority 3. In each, the circled point is the goal value in relative dose or volume. Note that organ goal functions have a decreasing component for P>Pi. This guides further improvement after the goal is met. **B)** Q-functions for a different set of goals: user has decided that the OUD goal has highest priority, and TUD has lowest priority.

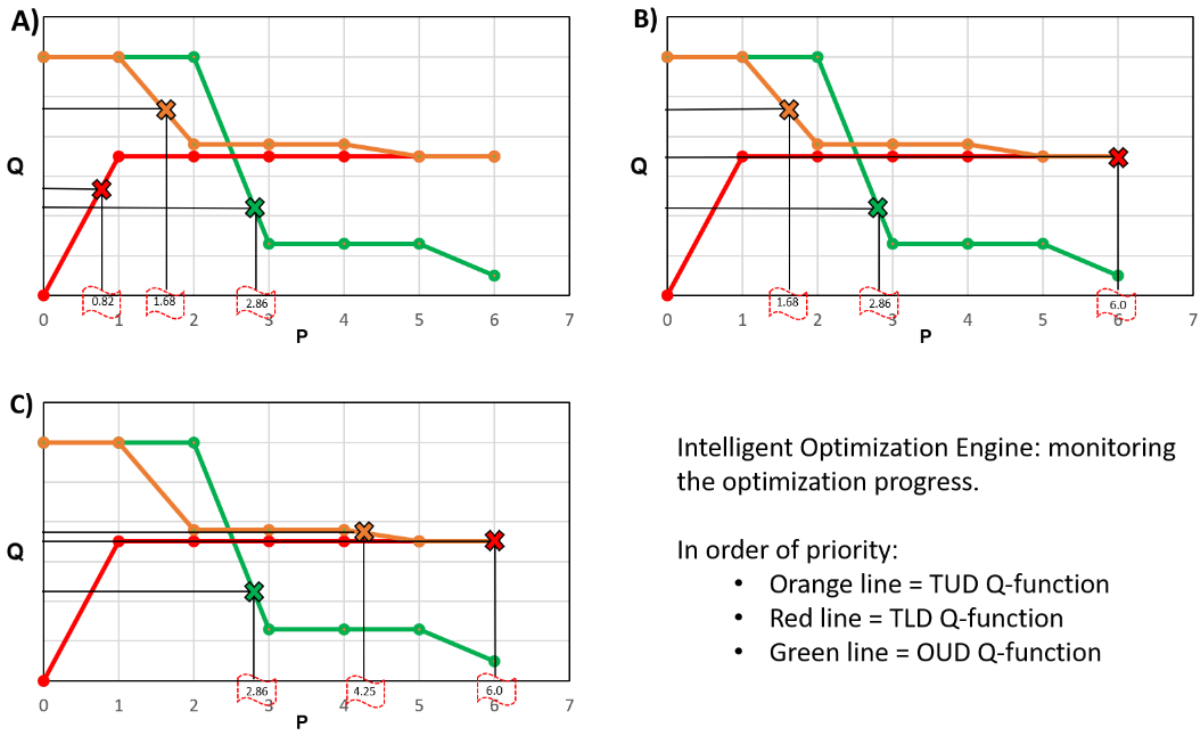

Fig 7. IOE, monitoring the optimization progress. Example with three goals. Optimization is started with fixed weightings for all objective functions. **A)** After *x* iterations, the dose distribution is interrogated, and P values obtained. IOE elevates cost for objective function with lowest P (TLD goal=*0.82*; red line). **B)** After another *x* iterations, dose distribution is interrogated again. TLD goal is met and P is at a maximum value (6.0 here). IOE elevates the cost for the goal with P=*1.68* (TUD). **C)** After another *x* iterations, dose distribution is interrogated a third time. Target upper goal is now met with P evaluated to *4.25*. IOE elevates cost for the goal that had P=*2.68* (OUD). Optimization continues until the collection of P is maximal and cannot further be improved.

The shape and location of the goal functions along the priority axis cause the optimization to progress similarly to how a human planner would work. Unmet goals with higher priority receive attention before unmet goals of lower priority. When goals are met, additional effort is expended to reduce the dose to organs where possible. The example in Figure 6 could be performed with the goal functions from panel B of Figure 7. In that case, the highest priority goal is the organ upper and would be achieved first, prior to focusing attention to the lower priority goals. Since this is goal-based optimization, the system does not stop when one goal is unmet; instead, the IOE detects the condition and re-baselines the goal function to higher (for organ or target upper goals) or lower (target lower) values.

Extra controls for clinical plan quality

Some aspects of clinical plan quality are not easily conveyed through clinical goals. To remedy this, the IOE adds some extra goals, structures, and optimization objectives for the PO algorithm.

Normal tissue dose is controlled using several methods. Firstly, hidden normal tissue optimization structures are created by the IOE. A ring-structure 1 cm away and 0.5 cm in thickness is created around every target. Another large normal tissue structure is created encompassing everything outward from the ring structures. All the normal tissue structures are given dose controlling clinical goals which the IOE will treat similarly as other clinical goals but with lower priority. The normal tissue dose is also controlled via Photon optimizer automatic normal tissue objective and a maximum dose objective that is assigned for the whole body.

The clinical quality for the dose coverage of PTV is also controlled. When the PTV is given a clinical goal that is not a minimum dose goal, the IOE adds a maximum of three helper objectives for the PTV. The objectives are placed along a parabola that has its maximum at V=100% and goes through the goal point (Di, Vi). Thus, the objectives have volume values between the goal volume Vi and full volume and the dose values are smaller than the goal dose Di. The addition of the objectives gives a tighter shape for the shoulder of the PTV's DVH curve. For example, a goal of D95%>50Gy will have extra objectives placed between 95% and 100% volume. The helper objectives are not added in cases where they would conflict with higher priority goals supplied by the user.

An extra control for clinical plan quality is also added for the plan complexity. This aspect is controlled and monitored during the optimization using an Photon optimizer smoothing objective. If the user has selected to use a RapidPlanTM model, the DVH estimates from the model are used as an additional guidance for the algorithm as described in a separate section below.

XI. RAPIDPLAN KNOWLEDGE-BASED PLANNING

Within the context of the aforementioned Intelligent Optimization Engine, the establishment of clinical goals compatible with the patient's unique geometry is not always intuitive. The proximity of a critical normal organ to the intended target may limit the potential for sparing and a compromise may be needed.

RapidPlan® knowledge-based planning, a machine learning tool that can potentially enhance the quality and efficiency of treatment planning based on historical patient data, addresses these challenges. The user builds models with RapidPlan by taking inputs from dosimetric and geometric parameters of the plans included in a training set. As an output, the models can generate predictions of the dose volume histograms for modeled structures and generate the optimization objectives needed to drive DVHs to those predictions. This permits the clinician to evaluate the predicted normal organ sparing prior to generation of a deliverable treatment plan, as well as incorporate the prediction to the plan quality evaluation.

Organ Partitioning

Organ volume partitioning is performed on each structure of every plan included in the training set and in application of a specific model to a new clinical case. The beam geometry is used to create the partitions, as the beam's eye view (BEV) from each field or control point is necessary to determine if a structure will receive any radiation dose at all. By combining the information obtained during partitioning, the software is able to predict dose volume histograms (DVH) for modeled structures. As shown in Figure 8, the organ partitions are:

- Out-of-field region – the region of the structure that receives only scattered radiation dose
- Leaf transmission region – the region where the structure is always covered by leaves from a multileaf collimator (MLC)
- Overlap of the organ with the target (or union of all targets)
- The in-field region – the region that that is distal or proximal to the target in the BEV and is not one the aforementioned regions. It represents the greatest contribution of dose to the modeled structure.

Partition modeling

Every case in the training set undergoes partitioning and RapidPlan extracts the average and standard deviations of the dose in the out-of-field partitions, leaf transmission partitions, and target overlap partitions. The in-field partition receives different treatment. This partition uses a supervised regression model of machine learning to infer characteristics that permit prediction of dose for this region. Combined with the result of the other three partitions, the entire DVH can be predicted.

Geometry Expected Dose

In order to extract information which connects the geometry of the patient to the observed radiation dose, we utilize the concept of Geometry Expected Dose (GED). The GED is a score for each voxel within the treatment volume based solely on basic photon beam characteristics and the relationship of the structure with the radiation fields.

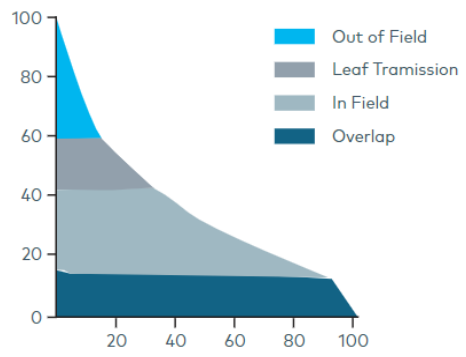

Fig 8. Organ volume partitions used in the machine learning DVH models. Each partition contributes to the sum total DVH as shown.

GED includes understanding the following properties:

- Field geometry
- Photon behavior
- Target geometry and dose levels
- Heuristics about which kind of beam arrangements lead to sparing of normal tissue

The GED can be calculated very quickly and only requires the field geometry, planning CT, and the structure geometry on the planning CT.

Supervised Regression

Once the GED for a given case is calculated, we can tabulate the GED volume histogram for the in-field partitions of modeled structures. These are considered to be highly correlated to the dose volume histograms for the same structures. This assumption is based on the observable fact that the geometry relation (proximity) to a target highly influences how well a structure can be spared in a given treatment plan. Over a population of similar treatment plans, or treatment plans from a similar anatomical site, the DVH of the in-field partition and the GED volume histogram will be highly correlated.

To extract the correlations, RapidPlan uses principle component analysis applied to both the collection of in-field DVH and GED volume histograms in a training data set. The coefficient obtained from the principle components can be arranged and analyzed through regression models to extract the correlation from a given GED volume histogram to an observed dose volume histogram. For any case for which a DVH prediction is desired, the dose volume histogram for the in-field partition is predicted from this regression obtained from the training data set.

RapidPlan compatibility with Ethos therapy

Any RapidPlan DVH estimation model can be imported and applied to a RT Intent within Ethos treatment management. If attached to an RT Intent, the DVH estimates for modeled structures are shown in the Ethos Dose Preview and in the Plan Review work areas. Additionally, the lower border of the DVH estimation band is used to derive a line objective which is applied during the plan generation process for both the initial planning and adaptive planning workflows in Ethos treatment planning.

Because there is not a known priority order for the line objective derived from the RapidPlan model, we cannot utilize the Intelligent Optimization Engine to effectively monitor or modify the strength of this line objective. The cost function derived from the line objective is added to the overall optimization, but at a level low enough not to overwhelm the objectives that the IOE determines, places, and monitors from the input goals and priority rankings. As such, its primary use in Ethos therapy is as a quality monitor. If the Ethos treatment planning Dose Preview or candidate plans from automated plan generation cannot achieve a result within the DVH predictions, the planner may need to add additional goals, change the order of goals, change the beam geometry, or determine that the case is not suitable for automated planning.

XII. CONCLUSION

Varian introduced and received CE mark for the Ethos therapy system in September of 2019; first patient treatments occurred later that month at Herlev Hospital in Denmark. The system received 510(k) clearance from the U.S. Food and Drug Administration in February 2020.

The Ethos therapy system incorporates technology that uses artificial intelligence and machine learning to create contours and generate adapted plans for physician review while a patient is on the treatment couch. The system offers radiation oncologists a set of simple tools that enable them to achieve their intention for each patient. The daily variation in a patient's anatomy, captured and visualized by iterative kV cone-beam CT (iCBCT) imaging, enables the on-couch adaptive workflow.

Ethos further allows a physician to choose which plan to deliver and to complete an adapted treatment within a typical 15-minute treatment time slot.

Clinical images at treatment delivery

Ethos therapy integrates multi-modality diagnostic images at the point of treatment on the treatment console. This means the daily re-planning sessions can utilize the

same multi-modality images that informed the initial planning stage. At each treatment, Ethos therapy shows:

- That day's anatomy with iCBCT images
- Registered CT, PET, and MR images
- The expected 3D radiation dose to the target and organs at risk for both the un-adapted and adapted plans

Decision-making guided by AI

The goal of Ethos therapy was to design a simple adaptive therapy workflow for both the initial planning and daily re-planning sessions.

During initial planning, Ethos therapy automatically produces several plan candidates with various beam geometries and techniques using prioritized target and organ at risk goals from the physician's intent. The clinician chooses the most suitable plan and authorizes it for delivery. This step provides confidence that the goals and patient geometry are compatible, and that plan automation can be performed each day. Each treatment day, once the daily anatomy is reviewed and accepted, Ethos therapy will prepare a new adapted plan using the beam geometry of the initial plan, the initial set of target and organ and risk goals, and give the clinician the choice of either the original or adapted plan for delivery.

The process is guided by the technology, as follows:

- A decision tree guides the entire adaptive therapy process
- Treatment management and treatment planning applications are tightly coupled and context-aware
- Clinician approvals move the process from one step to the next
- Every step of the workflow is optimized for speed and engineered for safety

Automated dose accumulation

Each day, the Ethos therapy system automatically reconstructs delivered dose in relation to today's anatomy. This capability:

- Demonstrates that the patient is receiving the intended dose
- Improves understanding of the treatment progress
- Helps identify when re-simulation may be required
- Simplifies off-line adaption

Familiar, efficient QA

QA for Ethos therapy follows a familiar workflow.

- The flexible workflow for pre-treatment QA accommodates phantom- or calculation-based QA methodologies.
- Initial planning and adaptive planning at the console use the same algorithms for consistency.
- Independent adaptive plan QA can be performed on-demand, without impeding treatment workflow.

This article derives from a Varian technical brief on Ethos™ therapy artificial intelligence.

Intended use summary: Varian Medical Systems' linear accelerators are intended to provide stereotactic radiosurgery and precision radiotherapy for lesions, tumors, and conditions anywhere in the body where radiation treatment is indicated.

Important safety information: Radiation treatments may cause side effects that can vary depending on the part of the body being treated. The most frequent ones are typically temporary and may include, but are not limited to, irritation to the respiratory, digestive, urinary or reproductive systems, fatigue, nausea, skin irritation, and hair loss. In some patients, they can be severe. Treatment sessions may vary.

Corresponding Author: Stephen Thompson
Varian Medical Systems
3100 Hansen Way, Palo Alto, California, 94089 USA
Email: stephen.thompson@varian.com