Radiotherapy and Oncology xxx (2017) xxx–xxx

Contents lists available at ScienceDirect

Radiotherapy and Oncology

journal homepage: www.thegreenjournal.com

Original article

Clinical decision support of radiotherapy treatment planning: A data-driven machine learning strategy for patient-specific dosimetric decision making

Gilmer Valdes [a,*], Charles B. Simone II [b], Josephine Chen [a], Alexander Lin [c], Sue S. Yom [a,d], Adam J. Pattison [e], Colin M. Carpenter [e], Timothy D. Solberg [a]

[a] Department of Radiation Oncology, University of California, San Francisco; [b] University of Maryland Medical Center, Baltimore; [c] Department of Radiation Oncology, University of Pennsylvania, Philadelphia; [d] Department of Otolaryngology-Head and Neck Surgery, San Francisco; and [e] Siris Medical, Redwood City, United States

ARTICLE INFO

Article history:
Received 31 May 2017
Received in revised form 10 October 2017
Accepted 10 October 2017
Available online xxxx

Keywords:
Decision support system
Knowledge-based planning
Dosimetric tradeoffs

ABSTRACT

Background and purpose: Clinical decision support systems are a growing class of tools with the potential to impact healthcare. This study investigates the construction of a decision support system through which clinicians can efficiently identify which previously approved historical treatment plans are achievable for a new patient to aid in selection of therapy.

Material and methods: Treatment data were collected for early-stage lung and postoperative oropharyngeal cancers treated using photon (lung and head and neck) and proton (head and neck) radiotherapy. Machine-learning classifiers were constructed using patient-specific feature-sets and a library of historical plans. Model accuracy was analyzed using learning curves, and historical treatment plan matching was investigated.

Results: Learning curves demonstrate that for these datasets, approximately 45, 60, and 30 patients are needed for a sufficiently accurate classification model for radiotherapy for early-stage lung, postoperative oropharyngeal photon, and postoperative oropharyngeal proton, respectively. The resulting classification model provides a database of previously approved treatment plans that are achievable for a new patient. An exemplary case, highlighting tradeoffs between the heart and chest wall dose while holding target dose constant in two historical plans is provided.

Conclusions: We report on the first artificial-intelligence based clinical decision support system that connects patients to past discrete treatment plans in radiation oncology and demonstrate for the first time how this tool can enable clinicians to use past decisions to help inform current assessments. Clinicians can be informed of dose tradeoffs between critical structures early in the treatment process, enabling more time spent on finding the optimal course of treatment for individual patients.

The ideal radiotherapy treatment plan should be personalized, delivering a potentially curative tumor dose while minimizing toxicity based on the individual patient's specific anatomy and underlying medical condition. Traditionally, treatment planning decisions are guided by high-quality scientific studies that map quantities of radiation dose, e.g. Dose to 20% of the Lung Volume (V20) or prescribed dose, to the likelihood of tumor control and normal tissue toxicity. While the challenge of dosimetrically-based planning is a solvable computational problem, the underlying clinical challenge lies in understanding the best treatment plan that can be achieved for a specific patient, related to differences in patient anatomy, tradeoffs in the weighting of planning constraints, and conscious and unconscious biases on the part of the prescribing physicians [1]. Moreover, the process of creating a treatment plan requires close communications between practitioners with different areas of expertise in clinical medicine (physician), radiation delivery (physicist), and treatment planning software (dosimetrist).

Clinical decision support systems leverage the history of past decisions by a clinical team and quickly provide reference data informed by past successes at a given clinic or shared between clinics. Combined with contemporary machine learning (also known as artificial intelligence) algorithms and large data stores, these expert systems have begun to impact clinical practice, with examples such as the triage of patients in the Emergency Department [2] or highlighting of calcifications in breast mammography

* Corresponding author at: Department of Radiation Oncology, University of California, San Francisco, San Francisco, CA, United States.
E-mail address: gilmer.valdes@ucsf.edu (G. Valdes).

https://doi.org/10.1016/j.radonc.2017.10.014
0167-8140/

[3]. A key element of these systems is the ability to augment clinicians' knowledge by processing previous decision records to identify those prior decisions and accompanying parameters that are relevant to the current patient. Together with new algorithm development, these systems promise to change the way decisions are made in medicine [4–6]. In radiation oncology, machine learning has been used in applications ranging from quality assurance to patient toxicity but clinical decision support systems (CDS) that empower physicians have not reached widespread use [7–13].

This paper demonstrates a clinical decision support system utilizing machine-learning for patient-specific treatment planning in radiation oncology with the purpose of assisting the radiotherapy planning team in making better treatment plan decisions by leveraging past data. The key differentiation of our system from other "knowledge based solutions" (KBS) to treatment planning is the focus on helping physicians navigate expectations about dosimetric tradeoffs before the treatment planning process. While other existing approaches focus on determining dose–volume histogram (DVH) expectations for individual organs at risk [14–16], our approach allows prospective, expectant navigation of the inherent tradeoffs between those expectations. The CDS described here is aimed at helping physicians and the clinical team to determine the best course of treatment before expending resources on a lengthy treatment planning process, and to better define expectations among the radiation oncology team during the course of plan development. While treatment planning is still required, the knowledge ascertained from a CDS has the potential to guide therapy and decrease the time needed to reach an acceptable plan.

Methods

Data collection

Data from 2009 to 2015 from the University of Pennsylvania Health System was used in this Institutional Review Board-approved retrospective study. The dataset was comprised of 104 consecutive early-stage lung cancer patients treated with stereotactic body radiation therapy (SBRT). Of these data, 81 received a prescription of 5000 cGy in 4 or 5 fractions to the planning target volume (PTV) (peripheral lesions), with the remaining receiving a prescription of 6000 cGy in 8 fractions (central lesions). An additional dataset was comprised of 40 patients with advanced-stage squamous cell carcinoma of the oropharynx who received postoperative proton radiotherapy. Of these data, 38 patients were prescribed between 6000 and 6600 cGy to the PTV (proton Radiobiological Effect Dose, RBE = 1.1). For each oropharynx patient, there also existed a volumetric modulated arc therapy (VMAT) clinical backup photon plan. Patients were identified through a database query (ARIA, Varian Medical Systems, CA). These data were exported, anonymized, and accumulated for processing by commercial software designed for the purposes of this study (QuickMatch, Siris Medical, CA).

Patient treatment plan classification

The goal of the CDS system is to match the current patient to previously treated patients with similar characteristics, such that previously achieved treatment plans and tradeoffs can be explored. This is represented schematically in Fig. 1. Current planning approaches either do not algorithmically use past data (Fig. 1A), or use past data to understand trends from DVH subpopulations (Fig. 1B), primarily as a quality assurance tool. In contrast, plan classification identifies discrete historical plans that can include dose tradeoffs between the target and various organs-at-risk. The requirements for plan classification are an accurate classification algorithm combined with a knowledge database of previous patient treatments. With a sufficiently large database, various achievable results will be proposed by the algorithm such that the clinical team has multiple reference points to use for optimally choosing the appropriate dose trade-offs for a given patient (Fig. 1C and E).

Consider a database of plans, $\boldsymbol{P}$, from which we seek matches to a specific plan p, where matching is defined according to a dissimilarity index of the dosimetric indices between plans. Addressing this problem as a classification problem similar to nearest neighbor classification, we would like to find the plans in $\boldsymbol{P}$ that are closer to p in terms of the dissimilarity index. A probability threshold, τ, is set for the dissimilarity index, and prior treatments that are within the threshold produce treatment plan matches. Formally, the subset $P_p \subseteq P$ of plans in $\boldsymbol{P}$ that matches plan p can be defined as:

$$P_p \subseteq P : \text{dissimilarity}(P_i, p | F_p = f_p, D_p = d_p, F_i = f_i, D_i = d_i) \leqslant \tau \tag{1}$$

where f are the features, d are the dosimetric indices, and the indices "i" denote different plans in the database. For the current patient, F_p and D_p are the features and dosimetric indices. A threshold, τ, is set for classification, and prior treatments that are within the threshold produce treatment plan matches. This probability threshold defines achievability and incorporates known sources of dosimetric variability in planning, including the repeatability of the treatment plan produced by the treatment planning system, v_{TPS}, and the variation in treatment planning preference between clinicians, v_C. The variability in the treatment planning system is found empirically by repeatedly running the treatment planning system for a given set of dosimetric objectives and priorities on exemplary treatment plans. The variability in clinician preference can be learned by calculating the variability in prediction for prediction models built on different subgroups of the dataset, stratified by, for example, treatment planner or physician. This threshold is defined as:

$$T = v_{TPS} + v_C \tag{2}$$

More specifically, the dissimilarity between the new patient plan p and a historical patient plan P_i is determined by calculating the difference between the j dosimetric indices of patient p and patient P_i. A historical patient is a match if,

$$\text{distance}\{d_{p_j}, d_{i_j}\} \leqslant T, \forall j \tag{3}$$

For an historical patient, d_i are the dosimetric indices from the historical plan. For the new patient,

$$d_p = \mathcal{F}(f_p) \tag{4}$$

where f_p are the generated features for the new patient. As noted above, a boosting framework is used to predict the dosimetric indices:

$$\mathcal{F}(f_p) = \sum_j T_j(f_p) \tag{5}$$

where $T(f)$ is a weak learner. Because of the modern machine learning approach that was used (boosting) and our extensive feature-set, the summation in Eq. (5) is over thousands of decision trees that take input from thousands of features. Therefore, a more detailed description of the function represented in Eq. (5) is not practical in this manuscript, and powerful computation is needed to obtain its value. Boosting is a well-known technique in modern machine learning and has been proven to be one of the most accurate, but powerful computing is required to generate the result [19].

Creating the features that account for data variability is a critical aspect of an accurate classifier; this process is often viewed as the most important aspect of a machine learning algorithm [19]. For an accurate feature set, we perform analysis on DICOM images,

Please cite this article in press as: Valdes G et al. Clinical decision support of radiotherapy treatment planning: A data-driven machine learning strategy for patient-specific dosimetric decision making. Radiother Oncol (2017), https://doi.org/10.1016/j.radonc.2017.10.014