PRECISION MEDICINE-COMMENTARY

Using Big Data Analytics to Advance Precision Radiation Oncology

Todd R. McNutt, PhD,* Stanley H. Benedict, PhD,† Daniel A. Low, PhD,‡
Kevin Moore, PhD,§ Ilya Shpitser, PhD,‖ Wei Jiang, MS,*
Pranav Lakshminarayanan, MSE,* Zhi Cheng, MD MPH,*
Peijin Han, MD MHS,* Xuan Hui, MD ScM,¶ Minoru Nakatsugawa, PhD,#
Junghoon Lee, PhD,* Joseph A. Moore, PhD,*
Scott P. Robertson, PhD,* Veeraj Shah,* Russ Taylor, PhD,‖
Harry Quon, MD MS,* John Wong, PhD,* and Theodore DeWeese, MD*

Departments of *Radiation Oncology and Molecular Radiation Sciences, and ‖Computer Science,
Johns Hopkins University, Baltimore, Maryland; †Department of Radiation Oncology, University of
California, Davis, Sacramento, California; ‡Department of Radiation Oncology, University of
California, Los Angeles, Los Angeles, California; §Radiation Medicine and Applied Sciences,
University of California, San Diego, La Jolla, California; ¶Department of Public Health Sciences,
University of Chicago, Chicago, Illinois; and #Canon Medical Systems Corporation

Summary

The conceptual model for decision support versus discovery with big clinical data analytics is different, and an overview of the implications of these differences is discussed in the context of precision medicine.

Big clinical data analytics as a primary component of precision medicine is discussed, identifying where these emerging tools fit in the spectrum of genomics and radiomics research. A learning health system (LHS) is conceptualized that uses clinically acquired data with machine learning to advance the initiatives of precision medicine. The LHS is comprehensive and can be used for clinical decision support, discovery, and hypothesis derivation. These developing uses can positively impact the ultimate management and therapeutic course for patients. The conceptual model for each use of clinical data, however, is different, and an overview of the implications is discussed. With advancements in technologies and culture to improve the efficiency, accuracy, and breadth of measurements of the patient condition, the concept of an LHS may be realized in precision radiation therapy. © 2018 Elsevier Inc. All rights reserved.

Reprint requests to: Todd R. McNutt, PhD, Department of Radiation Oncology, Johns Hopkins University, Harry and Jeanette Weinberg Bldg, 401 N Broadway, Ste 1440, Baltimore, MD 21231-2410. Tel: (410) 614-4594; E-mail: tmcnutt1@jhmi.edu

Introduction

The goal of precision medicine is to improve overall patient care and determine when and how to personalize patients' treatments. Currently, this is guided by a physician's understanding of the patient's condition by drawing from the physician's experience to align the specifics of care to the patient. Guidelines (1, 2) assist in the overall pathways for specific diseases, but for the most part, precision medicine is performed with finer granularity than the guidelines provide.

A learning health system (LHS) (3-5) is a concept where quantifiable diagnostic, treatment, and outcome data are captured from a continuous stream of patients and placed in a knowledge base. Knowledge is accessed by analytical tools that use statistical and machine learning algorithms to present trends and make predictions and causal inferences on outcomes. As more patient data are accumulated, the system continues to learn and improve on its models and ability to make specific predictions for individual patients.

When evaluating the possibilities of an LHS, it is important to recognize the difference between predictive modeling to assist in clinical decisions and knowledge discovery of the underlying mechanisms or causes of particular outcomes. In decision making, we decide on the most appropriate intervention for the patient, which may or may not be guided by complete knowledge of the underlying biological mechanisms. A new discovery, however, must uncover the biological understanding or derive hypotheses that may be further validated under more controlled studies. Clinical data complement pathology, genomics, and radiomics by providing details of the treatments and outcomes of patients for the advancement of precision medicine.

What Are Big Clinical Data?

The ability of big clinical data (6) to represent the real world with minimal bias, to accumulate assessments over time, to be linked with other databases, to be used and reused, and to enable a multidimensional understanding should all be considered to unlock the potential. Clinical data represent prior experience from patients and are captured through a multitude of methods, but limitations of our current protocols and pathways result in only a small fraction being used to make clinical decisions. For machine learning and statistical algorithms to take advantage of the entirety of the available data, medical records must adapt to support continuous feature extraction. Clinical data generally have a number of complications not found in typical cross-sectional study data sets. For example, clinical data exist in forms of free text to 3-dimensional volumes to structured data elements, all with longitudinal sampling. Clinical data also suffer from selective sampling, missingness, and measurement error.

Aside from lifestyle covariates, clinical data contain patient and disease status, treatment and symptom management, clinical and quality of life (QoL) outcomes, adverse effects, and survival. The key for enabling access is to extract meaningful information or features and store them in standardized ways (7).

Naturally, the level of precision in measuring outcomes dictates the quality of subsequent clinical conclusions. For instance, in current practice, a recurrence of a patient's cancer may be recorded but often without the specific location. This limits our understanding of whether the recurrence was coincident with the radiation treatment. Also, the measurement of a patient's clinical condition depends on available time and resources. For example, xerostomia can be scored by the clinician, evaluated through patient questionnaires, or measured with controlled stimulation methods, each with a corresponding increased time and cost.

Longitudinal assessment of patient status requires careful feature extraction. One can evaluate acute changes in toxicity such as taste disturbance or mucositis during treatment to understand a patient's ability to cope with treatment. Alternatively, evaluating longer-term toxicities provides a measure of permanent damage. Time to recovery of a particular function may also be measurable, as initial injury likely has different causal attributes than recovery of various irradiation-related toxicities.

Unlike standard cross-sectional studies, where treatments are binary and represent case and control groups, radiation therapy involves a 3-dimensional dose delivered over multiple days, yet protocol standards extract simplistic dose-volume features as efficient measures of treatment-plan quality. Dose-volume histograms (DVHs) leave out useful information and thus are insufficient on their own to support precision medicine (8). A DVH assumes each location within a region is equally sensitive to radiation and equally responsible for biological function. Advanced methods of extracting dose features and patterns would enable a better understanding of the impact on patient outcomes (9).

LHS and Predictive Modeling

A common goal of traditional statistical modeling is the discovery of the underlying mechanisms or cause of outcomes. Breiman (10) compared a "data model," where a statistical model is assumed to describe a relationship and validated with the data, with an "algorithmic model," where the mathematical model that relates the input variables to the outcomes is computationally determined through machine learning. Both approaches have benefits and flaws: The data models are usually hypothesis driven yet may not reflect the complexity of the true process, but they nonetheless enable improved understanding of the system. The algorithmic models, on the other hand, are hypothesis generating, presenting superior predictive accuracy, yet

make it challenging to uncover the dominant input variables and/or causal attributes.

Medical information is very complex and often aggregated into features that can mask important underlying details. Such dimension reductions are necessary but risk being insufficient. A good example is the selected points on a DVH, where we have essentially reduced 3-dimensional dose in a region to a single value of dose or volume. This data reduction may have a negative impact on the ability to build a model to predict organ function or disease control after treatment that may have spatial dependence. It is not easy to proactively determine whether this type of ad hoc feature will preserve or discard useful relationships between the features and outcomes. Developing and applying dimension reduction strategies that usefully preserve true relationships in the data may improve normal tissue complication models (11).

Considerations for predictive models must include the purpose of building them, whether they are to be used for decision support or for discovery of new knowledge. There is more than 1 tool, and selecting the right one to apply to the clinical question and purpose will be critical for making more precise patient care decisions.

Decision support

The goal of decision support is to provide the most appropriate intervention for the patient (12) and not necessarily to discover new knowledge. This begs the question of which outcome prediction models should be selected with what accuracy requirement.

The key to selecting the best performing model is understanding the decision and intervention to be made. For example, if the intervention is to use a feeding tube to prevent weight loss for head and neck cancer patients undergoing treatment with radiation therapy and chemotherapy, then it may not be necessary to know what combinations of toxicity caused the weight loss since the intervention is intended to treat the symptom instead of its underlying cause. Alternatively, if it is understood that, for a particular patient, taste disturbance would likely cause excessive weight loss, then the intervention may be to modify the radiation treatment to minimize the taste disturbance or to refer to a nutritionist to consult on nutritional support.

Figure 1 depicts a framework for decision support (3) where, at some time point in the care of a patient, a decision needs to be made. The inputs to the predictive model include the facts about the patient and potential interventions. Outcome predictions such as risk of a particular toxicity or probability of local disease control are presented to the clinician and patient with the specific attributes most influential to the prediction. These outputs could then be used to assist the decision making, whether it be selection of or change in the treatment course or an intervention to improve symptoms.

An evaluation of the dominant attributes of a specific prediction must consider an understanding of the individual patient and the existing knowledge of underlying causes. Predictive models often do not separate causation from association. Thus, interventions that depend on treating a causal attribute must consider the limitations of the predictive models.
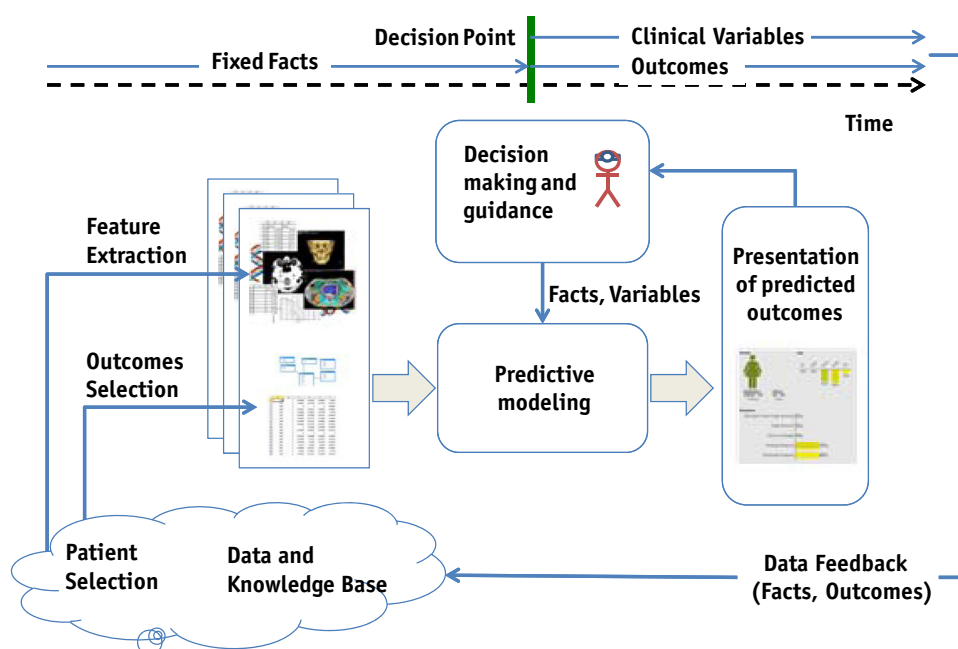

Fig. 1. Decision support framework to make predictions of outcomes of individual patients. The models, derived from the knowledge base, use the facts and clinical options and/or variables in making the predictions that are presented to the physician to assist in decision making.
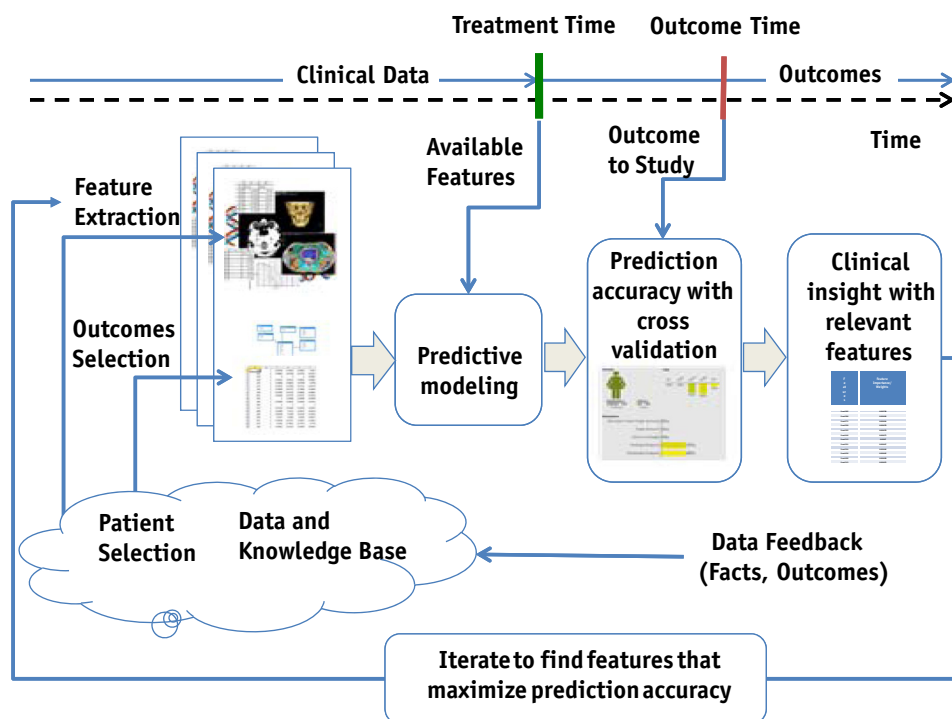
Fig. 2.   Hypothesis generation uses predictive modeling to maximize the prediction accuracy to uncover specific features of the patients and their treatments most correlated with an outcome. The features are derived from the raw clinical data.

Discovery and hypothesis derivation

An LHS also provides the opportunity to extend knowledge through discovery and hypothesis derivation. In essence, the goal is to both understand features most predictive of outcomes and uncover the underlying causes.

Figure 2 depicts a framework for discovery using the LHS. The process is to find features of the patients that most influence an outcome by generating predictive models and cross validating them with the available data to maximize prediction accuracy. In this approach, iterative exploration of an unlimited set of features seeks out those that maximize the predictive accuracy. After validation, a review of the relevant features can support hypothesis derivation and help uncover discoveries that can be further studied.

Aside from predictive modeling, cause-and-effect relationships between features and outcomes are important types of hypotheses and are often the most scientifically relevant. These types of hypotheses are most relevant for decision support, since making decisions based on purely associational criteria amounts "to an irrational policy of managing the news, and results, in practice, in replication failures and poor recommendations" (13). Identifying cause-effect relationships entails systematically adjusting for selection effects and confounding bias, using methods such as G-computation (14), propensity score matching (15), and inverse probability weighting (16). In addition, under strong assumptions, inferences about causal directionality underlying associational relationships between multiple variables are possible.

Although there is a large effort in machine learning and statistics to identify cause-and-effect relationships from observational data, all causal hypotheses generated by such methods must ultimately be validated by formal randomized controlled trials.

What is missing?

Both decision support and discovery are limited by the knowledge contained in the database. For example, one institution may have ancillary care pathways that differ from another institution's such that these differences impact patients' outcomes. If institution A uses a speech pathologist to provide routine swallow therapy and institution B does not, their outcomes for swallow function may be different. If the details on a patient's adherence to swallow therapy are not contained in the database for either institution, then the treating institution would be an aggregate variable that might correlate with a swallow function outcome.

This missing of data also manifests itself when models are validated between institutions. If a model is built from only institution A's data and validated with institution B, the unknown information may dominate, and the validation will fail. Alternatively, if a model is built with both institutions' data and institution selection is the most dominant variable, there may be little difference relative to having 2 models, 1 for each institution, since the prediction will mostly depend on the treating institution.

This has implications. When using the LHS for decision support, the goal is to have the most accurate prediction, and that may happen with models built using only patients treated at the institution where the patient is to be treated. For discovery, however, the goal is to uncover underlying mechanisms, and for this, interinstitutional validation becomes important, and completing missing information in the data is crucial to uncovering this new knowledge.

In addition to outright missing information, the knowledge base is limited within the norms of clinical care. With radiation treatments, for example, only the variability of the dose distributions present in the knowledge base is available (17). If a particular anatomic region of every patient received the same dose, then there is no possibility of learning the impact on the outcome of varying dose to that region. Since patients are treated with similar dose goals in planning, the data will inherently subdue the importance of the known dose goals, while it is potentially unethical to deviate from them. In essence, "Without deviation from the norm, progress is not possible," as stated by Frank Zappa (18). As the effects of irradiation on patients are explored, consideration of the existing knowledge and how much of it is inherently included—and thus subdued—is critical in any interpretations. Furthermore, as we trend toward standardized clinical guidelines, we risk further limitations on the knowledge contained in the data and on our ability to personalize care in the context of improved quality and safety.

Examples

Treatment-plan quality prediction

An early example of using big data tools in radiation therapy is the concept of geometry-driven or knowledge-based treatment planning (KBP) (19-24). KBP aligns with the LHS model in that it provides actionable predictions of dose goals for planning and continuously learns as more treatment-planning data are accumulated.

KBP analyzes a plurality of prior treatments to discover patient-specific anatomic features that precisely correlate to high-quality radiation dose delivery. With the model-based dose predictions, KBP can be used for treatment-plan quality control or outright plan automation. In its generalized form, KBP makes use of established machine learning techniques such as supervised inference engines to discover relevant geometric variables and their correlation to patient-specific dose prediction.

While KBP is already in routine clinical use at some institutions for the purposes of automated planning (25), one of the most important contributions from KBP has come in the combination of knowledge-based plan quality control with a cooperative group clinical trial to assess the frequency and clinical severity of suboptimal treatment planning in a diverse multi-institutional data set (26).

Incorporating toxicity outcomes and clinical intervention

The prediction of toxicities is also critical to a patient's ability to tolerate treatment and his or her long-term QoL. An example is weight loss prediction using a classification and regression tree for head and neck cancer patients (27). Two predictions at different time points were developed to predict weight loss at 3 months after treatment: (1) during planning using patient demographic and dosimetry data; and (2) at the end of treatment using additional on-treatment toxicities and patient-reported QoL data. During planning, the top 2 predictors of weight loss were tumor site and higher doses to the masticatory muscle, a potentially modifiable factor. By the end of the treatment, when irradiation-induced toxicities started to present, patient-reported oral intake, tumor site, and dose to the combined parotids were more predictive. Early identification of high risk of excessive weight loss may inform interventions such as feeding tube placement, referral to a speech pathologist for swallow function evaluation and exercise, or frequent monitoring early after treatment.

Another example is in the prediction of irradiation-induced xerostomia for head and neck cancer patients. A wide range of clinical, demographic, and dosimetric factors were evaluated by the algorithm and subsequently cross validated by the accruing data. In this example, a low-dose bath to the combined parotids and intermediate-level irradiation to the submandibular glands, alongside clinical factors of chemotherapy, human papillomavirus infection, feeding tube use, and baseline body mass index, were identified as crucial for patients prone to severe xerostomia. Downstream conditional predictive factors including age, alcohol use, age, and smoking were also attributable.

The LHS allows a comprehensive exploration of predictors for a variety of treatment-related toxicities beyond the single-organ DVH and simple normal tissue complication models and, furthermore, bridging of all other clinical and patient factors into an all-encompassing prediction model. Evidence from such models warrants the foundation for clinical decision support for the prevention and/or management of toxicities.

The exploration of dose distribution patterns and interorgan dependencies may be critical to precision medicine. Early studies have shown the spatial dependence of dose on xerostomia in the parotid glands (28-30) and dysphagia across the swallowing muscles (9, 31). Further exploration of these spatial and multiorgan dependencies will be enabled by the LHS and may improve our knowledge of the impact of irradiation on normal function (32).

Genomics, Pathology, and Radiomics

At a higher level, radiomics, genomics, and pathology are patient-specific data that are subjected to feature extraction

in clinical practice and for research (33-35). Radiomics is a clear example where a portion of a diagnostic image is identified and features of the voxel values—density, texture, and gradient—are calculated and presumed to reflect characteristics of the specific tissue being analyzed. These features are used to predict disease response to treatment or toxicity. The features themselves do not necessarily reflect underlying mechanisms or status of the tissue, but they might sample characteristics that reflect underlying differences between patients.

In contrast, pathology has had a long history of feature extraction where cell type, grade level, and differentiation are characterized from biopsy slides and the disease is classified with staging and grading models (36). This history has provided a means of communicating complex image and tissue characteristics, and these characteristics are used to classify patients for both research analysis and clinical decision pathways (37).

Genomics is another very complex data set used to seek out known features that are associated with particular outcomes. The dominant research looks to discover genetic predisposition to disease or response to treatment. Other work has suggested there may also be a predisposition to radiation toxicity based on genomic signatures (38). The LHS offers an opportunity to explore genomics in much greater detail and assist in uncovering genomic patterns that influence outcomes, which would otherwise be impossible to discover.

Uncovering how the features derived from images, pathology, and genomic signatures can inform clinical practice or discovery but ultimately relies on accurate measures of outcomes and treatment information. Thus, it is the combination of the clinical data with these measures that will advance their ability to provide precise treatment options for our patients.

Discussion

Just as machine learning is being used to drive autonomous vehicles, is it reasonable to expect similar successes in radiation oncology? At this point, self-driving cars focus on the rules of the road and respond to immediate detection of obstructions in their local environment. They, however, exhibit difficulty in defensive driving, where they must weigh the risks of the unknown and anticipate what might happen. Radiation oncology, though precise in the treatment, presents a similar situation with a few rules of the road and acute observations but may be dominated by unknowns and patterns of defensive practice. As such, our expectation for the foreseeable future should be one of improved risk or outcome prediction as a supplement to physician-based clinical decision making.

The key to success is to uncover and measure as many of the unknowns as possible. Is a future possible in which we accurately measure the critical aspects of a patient's outcomes and treatment? Computerization of health care is

advancing rapidly, and the societal culture evolving from having smartphones amplifies the likelihood that good measures of the continuous patient condition will only advance. As outcome measures improve, radiation oncology must do its part to accurately archive treatments in easily retrievable form, adhering to standard nomenclatures. It should be possible to query features of the patient's history, physical examination findings, radiographic studies, laboratory tests, and "delivered" dose for any patient from our clinical archive without significant processing. It should be part of our practice to be good stewards of the data and accurately record 3-dimensional delivery while capturing the clinical data, appreciating that these data ultimately will contribute to the LHS.

The presentation of a patient's condition is currently conveyed mostly in text and is typically presented in the absence of population-based information. A disease- and patient-specific presentation through modern human-computer interfaces coupled with population-based statistics can highlight how well a patient is doing in the context of his or her disease peers and can, in itself, aid in individualized decision making. Advances in such patient presentations offer the framework to present risk and outcome predictions in a form that is actionable.

The vision is a future where data are instinctively collected and each patient is provided a prediction of his or her disease outcomes and complications against the backdrop of his or her peer populations with treatments tailored to an individual's needs and sensitivities. Continuous learning of this LHS will open insights that involve patterns in data far more complex than our traditional evidence-based methods can uncover and will open the flood gates of knowledge.

References

1. American Society for Radiation Oncology. Clinical practice statements. Available at: https://www.astro.org/Clinical-Practice-Statements.aspx. Accessed November 21, 2017.
2. Jazieh AR, McClure JS, Carlson RW. Implementation framework for NCCN guidelines. *J Natl Compr Canc Netw* 2017;15:1180-1185.
3. McNutt TR, Moore KL, Quon H. Needs and challenges for big data in radiation oncology. *Int J Radiat Oncol Biol Phys* 2016;95:909-915.
4. Chen RC, Gabriel PE, Kavanagh BD, et al. How will big data impact clinical decision making and precision medicine in radiation therapy? *Int J Radiat Oncol Biol Phys* 2016;95:880-884.
5. Friedman CP, Allee NJ, Delaney BC, et al. The science of learning health systems: Foundations for a new journal. *Learn Health Sys* 2017;1:e10020.
6. Bellazzi R. Big data and biomedical informatics: A challenging opportunity. *Yearb Med Inform* 2014;9:8.
7. Tran T, Luo W, Phung D, et al. A framework for feature extraction from hospital medical data with applications in risk prediction. *BMC Bioinformatics* 2014;15:425.
8. Moiseenko V, Dyk JV, Battista J, et al. Limitations in using dose-volume histograms for radiotherapy dose optimization. In: The Use of Computers in Radiation Therapy. Berlin: Springer; 2000. p. 239-241. Available at: https://link.springer.com/chapter/10.1007/978-3-642-59758-9_90. Accessed November 10, 2017.

9. Monti S, Palma G, D'Avino V, et al. Voxel-based analysis unveils regional dose differences associated with radiation-induced morbidity in head and neck cancer patients. *Sci Rep* 2017;7:7220.

10. Breiman L. Statistical modeling: The two cultures. *Stat Sci* 2001;16:199-231.

11. Nabi R, Shpitser I. Semi-parametric causal sufficient dimension reduction of high dimensional treatments. Available at: https://arxiv.org/abs/1710.06727. Accessed November 21, 2017.

12. Lambin P, van Stiphout RG, Starmans MH, et al. Predicting outcomes in radiation oncology—Multifactorial decision support systems. *Nat Rev Clin Oncol* 2012;10:27-40.

13. Lewis D. Causal decision theory. *Australas J Philos* 1981;59:5-30.

14. Robins JM. A new approach to causal inference in mortality studies with sustained exposure periods—Application to control of the healthy worker survivor effect. *Math Model* 1986;7:1393-1512.

15. Rosenbaum PR, Rubin DB. The central role of the propensity score in observational studies for causal effects. *Biometrika* 1983;70:41-55.

16. Horvitz DG, Thompson DJ. A generalization of sampling without replacement from a finite universe. *J Am Stat Assoc* 1952;47:663-685.

17. Robertson SP, Quon H, Kiess AP, et al. A data-mining framework for large scale analysis of dose-outcome relationships in a database of irradiated head and neck cancer patients. *Med Phys* 2015;42:4329-4337.

18. Zappa F, Occhiogrosso P. The Real Frank Zappa Book. New York, NY: Poseidon Press; 1989.

19. Wu B, McNutt T, Zahurak M, et al. Fully automated simultaneous integrated boosted-intensity modulated radiation therapy treatment planning is feasible for head-and-neck cancer: A prospective clinical study. *Int J Radiat Oncol Biol Phys* 2012;84:647.

20. Appenzoller LM, Michalski JM, Thorstad WL, et al. Predicting dose-volume histograms for organs-at-risk in IMRT planning. *Med Phys* 2012;39:7446-7461.

21. Petit SF, Wu B, Kazhdan M, et al. Increased organ sparing using shape-based treatment plan optimization for intensity modulated radiation therapy of pancreatic adenocarcinoma. *Radiother Oncol* 2012;102:38-44.

22. Wang Y, Zolnay A, Incrocci L, et al. A quality control model that uses PTV-rectal distances to predict the lowest achievable rectum dose, improves IMRT planning for patients with prostate cancer. *Radiother Oncol* 2013;107:352-357.

23. Wu B, Ricchetti F, Sanguineti G, et al. Data-driven approach to generating achievable dose-volume histogram objectives in intensity-modulated radiotherapy planning. *Int J Radiat Oncol Biol Phys* 2011;79:1241-1247.

24. Wu B, Ricchetti F, Sanguineti G, et al. Patient geometry-driven information retrieval for IMRT treatment plan quality control. *Med Phys* 2009;36:5497-5505.

25. Li N, Carmona R, Sirak I, et al. Highly efficient training, refinement, and validation of a knowledge-based planning quality-control system for radiation therapy clinical trials. *Int J Radiat Oncol Biol Phys* 2017;97:164-172.

26. Moore KL, Schmidt R, Moiseenko V, et al. Quantifying unnecessary normal tissue complication risks due to suboptimal planning: A secondary study of RTOG 0126. *Int J Radiat Oncol Biol Phys* 2015;92:228-235.

27. Cheng Z, Nakatsugawa M, Hu C, et al. Evaluation of classification and regression tree (CART) model in weight loss prediction following head and neck cancer radiotherapy. *Adv Radiat Oncol* 2017. Available at: http://www.advancesradonc.org/article/S2452-1094(17)30229-4/fulltext.

28. Deasy JO, Moiseenko V, Marks L, et al. Radiotherapy dose-volume effects on salivary gland function. *Int J Radiat Oncol Biol Phys* 2010;76:S63.

29. Marungo F, Robertson S, Quon H, et al. Creating a data science platform for developing complication risk models for personalized treatment planning in radiation oncology. In: Proceedings of the Annual Hawaii International Conference on System Sciences. Washington, DC: IEEE Computer Society; 2015:3132-3140.

30. Quon H, Park S, Plishker W, et al. Preliminary clinical evidence of parotid subvolume radiosensitivity and the risk of radiation-induced xerostomia in head and neck cancer (HNC) patients. *Int J Radiat Oncol Biol Phys* 2016;96:E341.

31. Kumar R, Madanikia S, Starmer H, et al. Radiation dose to the floor of mouth muscles predicts swallowing complications following chemo-radiation in oropharyngeal squamous cell carcinoma. *Oral Oncol* 2014;50:65-70.

32. Lakshminarayanan P. "Radio-morphology: Parametric Shape-Based Features in Radiotherapy", Master's Thesis, Johns Hopkins University, Biomedical Engineering, December 2016 Readers: McNutt T, Taylor R, Siewerdsen J.

33. Aerts HJ, Velazquez ER, Leijenaar RT, et al. Decoding tumour phenotype by noninvasive imaging using a quantitative radiomics approach. *Nat Commun* 2014;5:4006.

34. El Naqa I, Grigsby P, Apte A, et al. Exploring feature-based approaches in PET images for predicting cancer treatment outcomes. *Pattern Recognit* 2009;42:1162-1171.

35. Gillies RJ, Kinahan PE, Hricak H. Radiomics: Images are more than pictures, they are data. *Radiology* 2016;278:563-577.

36. AJCC Cancer Staging Manual. New York, NY: Springer Science+Business Media; 2016.

37. Coley RY, Zeger SL, Mamawala M, et al. Prediction of the pathologic Gleason score to inform a personalized management program for prostate cancer. *Eur Urol* 2017;72:135-141.

38. West CM, Barnett GC. Genetics and genomics of radiotherapy toxicity: Towards prediction. *Genome Med* 2011;3:52.