

Machine learning for radiation outcome modeling and prediction

Yi Luo^{a)}

Department of Radiation Oncology, University of Michigan, Ann Arbor, MI 48103, USA

Shifeng Chen

Department of Radiation Oncology, University of Maryland School of Medicine, Baltimore, MD 21201, USA

Gilmer Valdes

Department of Radiation Oncology, University of California, San Francisco, CA 94158, USA

(Received 23 January 2019; revised 26 March 2019; accepted for publication 9 April 2019; published 15 May 2020)

Aims: This review paper intends to summarize the application of machine learning to radiotherapy outcome modeling based on structured and un-structured radiation oncology datasets.

Materials and methods: The most appropriate machine learning approaches for structured datasets in terms of accuracy and interpretability are identified. For un-structured datasets, deep learning algorithms are explored and a critical view of the use of these approaches in radiation oncology is also provided.

Conclusions: We discuss the challenges in radiotherapy outcome prediction, and suggest to improve radiation outcome modeling by developing appropriate machine learning approaches where both accuracy and interpretability are taken into account. © 2019 American Association of Physicists in Medicine [https://doi.org/10.1002/mp.13570]

Key words: accuracy, interpretability, machine learning, radiation outcome modeling, structured and unstructured datasets

1. INTRODUCTION

Machine learning (ML), which evolved from the study of pattern recognition and computational learning theory in artificial intelligence, intends to explore the study and construction of algorithms that can learn from data.¹ Recently, there is a tremendous increase in the use of ML in different areas of radiation oncology, such as treatment planning optimization,^{2,3} segmentation,⁴ radiation physics quality assurance,^{5–7} contouring or image-guided radiotherapy.^{8,9} In this paper, we focus on ML for radiation outcome modeling.^{10–12}

Radiation outcome modeling includes survival analysis, local tumor control probability (TCP), and normal tissue complication probability (NTCP) (e.g., radiation pneumonitis, cardiac toxicity, and esophagitis). In the past, these approaches relied on analytical models that assumed a defined relationship between the dose and the specific outcome (e.g., linear quadratic model for tumor control probability¹³ or the Lyman model for normal tissue complication probability¹⁴). Also, they were limited to only considering dosimetric factors. For analytical models, if the underlying assumptions and simplifications are not correct as it is usually the case, then prediction accuracy suffers. In contrast, with the use of ML algorithms, different relationships between dose and outcome can be automatically investigated. Additionally, nondosimetric factors such as comorbidities, age, performance status, imaging, and genetic information, etc., can be easily incorporated,^{10–12} and they have different representations, where age and imaging are described by continuous variables, dose-volume histogram (DVH) bins, performance status,

and single-nucleotide polymorphisms (SNPs) are denoted by ordinal variables, and gender and chemotherapy are indicated by binary variables. The radiation outcomes depend on all these characteristics included in them may result in more accurate models but this impose a challenge in terms of data size. Essentially, each patient is being characterized by a high-dimensional feature vector but the number of samples available to learn a valid relationship is comparatively smaller in radiation oncology compared to other fields. In order to enhance generalization or reduce overfitting, either redundant or irrelevant features may need to be removed without incurring much loss of the information, which is called “feature selection” in ML. Basically, wrappers, filters, and embedded methods are three main categories of feature selection algorithms.¹⁵ While wrapper approaches use a predictive model to score feature subsets and filter approaches use a proxy measure instead of the error rate to score a feature subset, embedded approaches are a catch-all group of techniques to perform feature selection as part of the model construction process. When filters and wrappers are used, a common mistake is to reuse the same observations used to select the variables to build and report the accuracy of the final model. Therefore, recently embedded methods like least absolute shrinkage and selection operator (LASSO) seem to be preferred because they perform feature selection internally and there is smaller chance of information leakage from the training to the testing datasets.¹⁶

For radiation outcome modeling, supervised ML is the main approach. In supervised learning, the ML algorithm is provided with input features (normally in the form of a matrix with each row corresponding to a patient and each

column corresponding to a feature) and outcome (in the form of a vector) as shown in Table I. With this data, a specific loss function is written and by minimizing it, the optimal mapping between the input features and the outcome is found. In this paper, we will review the challenges, pros and cons of using these algorithms to model outcome data in radiation oncology.

2. MACHINE LEARNING APPROACHES FOR RADIATION OUTCOME MODELING

Several ML algorithms have been investigated for radiation therapy outcome modeling. If the input features have meaning by themselves (called in the ML literature structured or tabular data), then ML algorithms like logistic regression with penalization, decision trees (DTs), Bayesian networks (BNs), random forests (RFs), and gradient boosting are preferred, although algorithms that do representational learning (design features or concepts) like neural networks can still be used in our field and outside for the analysis of the structured data. If the input variables are unstructured (images, text), then neural network-based deep learning (DL) algorithms are preferred.

2.A. Machine learning approaches for structured data

The observed radiation treatment outcome can be considered as the result of the interaction of several dosimetric, clinical, or biological variables.¹⁷ ML algorithms intend to develop data-driven models by fitting parameters using the collected clinical and dosimetric data.¹⁸ For structured dataset, ML approaches have already been employed for several treatment sites such as lung,¹⁹ prostate,²⁰ head & neck cancer,²¹ or meningioma.¹¹ In all cases, many different ML algorithms, such as linear regression, artificial neural network (ANN), support vector machine (SVM), BNs, DT, RFs, or gradient boosting machine (GBM), have been explored. Below, we will attempt to cover previous publications using these algorithms (nonexhaustive).

Linear regression models the relationship between the response and one or more independent variables as a linear equation. Systems biology is the computational and mathematical modeling of complex biological systems. In order to develop a systems biology understanding of radiosensitivity to enhance our ability to identify radiation-specific biomarkers for lung, colon, and breast cancers,²² linear regression was employed to correlate gene expression with survival. In

another study, the impact of overall treatment time on patient outcome in non-small-cell lung cancer was studied using linear regression and Cox proportional hazard models.²³

Artificial neural networks are multilayer nonlinear models that use gradient descent and back propagation to find the optimal coefficients. An ANN is based on a collection of connected nodes called artificial neurons. Each connection between nodes can transmit a signal from one neuron to another. There is a weight associated with each connection, and the weight increases or decreases the strength of the signal at a connection which is adjusted as learning proceeds. In a study to predict symptomatic lung injury based on pre-radiotherapy, biological, and physical data,²⁴ a nonlinear, feed-forward ANNs was trained to predict the radiation outcome. The ANN approach was also used to predict radiation pneumonitis following radiotherapy.²⁵ Three ANNs, each with three layers, were developed for this classification task. A growing/pruning-based ANN approach was developed to predict Grade 2 + radiation pneumonitis (RP2), starting from the smallest possible network until a satisfactory solution was found.²⁶ While the input of the first ANN was limited to the patient lung dose–volume data only, the second study showed that the addition of non-dosimetric features can significantly improve the generalization capability of the network.

A SVM is a class of ML algorithm that attempts to find a feature space where the data can be linearly separated using a kernel function to do the transformation. In a study to predict patient's complication risk for personalized treatment planning,²⁷ a nonlinear kernel-based SVM method was used to learn complex interactions between the observed toxicities and treatment, anatomical, and patient-related variables. In the same study, the value of feature transformation using principle component analysis (PCA) to visualize high-dimensional data and determine proper kernel type was also demonstrated. It was demonstrated that linear models probably work well and nonlinear methods are unnecessary when responses may be separated along a linear ridge in a PCA plot, and nonlinear features generated via kernel/SVM methods may improve the model prediction when there is no such linear ridge under PCA analysis.

Bayesian networks are a type of probabilistic graphical model to represent a set of variables and their conditional dependencies via a directed acyclic graph. In a new treatment planning model for more outcome-focused decision-making, a BN was employed to model the radiation therapy process of prostate cancer and treatment outcomes such as distant metastasis, rectal, and bladder complications using expert opinion, results of clinical trials, and published research. Then, the quality-adjusted life expectancy of a patient can be evaluated based on the BN.²⁸ In order to explore the biophysical relationship among radiation treatment, patients' characteristics (age, comorbidities, etc), and radiation outcomes, BN were employed to predict radiation pneumonitis,^{29,30} local control,³¹ or both.³² Figure 1 shows an example of a pretreatment joint BN to predict lung cancer tumor local control (LC) and radiation pneumonitis toxicity Grade 2 or above (RP2) simultaneously.³² The most important features for

TABLE I. Patients' input features and radiation treatment outcome.

Patients' ID	Local control	Tumor dose	Age	Chemo
Patient #1	1 (yes)	75.60	78	0 (no)
Patient #2	0 (no)	71.39	70	1 (yes)
Patient #3	1 (yes)	66.00	56	1 (yes)

radiation outcome prediction including tumor and lung generalized Equivalent Uniform Dose (gEUDs), one SNP (cxcr1-Rs2234671), two miRNAs (miR-20a-5p and miR-191-5p), one pretreatment PET radiomics feature metabolic tumor volume (MTV), and three pretreatment cytokines (IL-15, IL-4, and IL-10) were selected from a high-dimensional retrospective dataset as indicated by the nodes in the BN. The edges of the BN, identified with different colors, represent the biophysical relationships between the variables analyzed. BNs can be used to explore biophysical cause-effect relationship from retrospective dataset, but it is limited by inference complexity, and its computational cost grows exponentially as the number of nodes increases.

As mentioned above, one problem that limits the exploration of biological connections with BN is dataset size. Since most datasets in radiation oncology tend to be small, sharing data among different institutions might be necessary. Distributed learning is a technique that can allow the learning of predictive models on data originating from multiple hospitals without the data leaving the hospital. A BN model was adapted for distributed learning to predict dyspnea by combining data originating from multiple hospitals. It was shown that it is possible to use the distributed learning approach to train the BN model without the specific datasets leaving the individual hospital.³³ This results open a venue for data sharing among multiple hospitals without violating privacy laws.³³

Decision trees are constructed using recursive partitioning analysis to optimize successive dichotomisation of input variables. The resulted tree-like structure had been used to augment prediction of the classic Lyman NTCP,¹⁴ pneumonitis,^{10,34} chest wall pain,¹² salvage high-dose-rate brachytherapy (sHDRB)³⁵ and meningioma.¹¹ An advantage of DTs is that they are highly interpretable. Their biggest

disadvantage is that they are not the most accurate algorithm on expectation. MediBoost is a single tree technique based on boosting that attempts to improve accuracy while still building interpretable trees. Since it uses weighted versions of all cases to derive splits at each point in the tree, it can be advantageous for relatively small datasets. MediBoost has been used to predict both pneumonitis and biochemical failure after sHDRB surpassing RF performance in some cases but not all.^{11,35}

Random forest is considered one of the best out of the shelves ML algorithms. It improves trees' accuracy by building a large number of deep trees (ensemble) and averaging them which reduces variance and improves accuracy. In addition to predicting the above mentioned outcomes,^{10-12,35} the RF approach was also employed to predict rectal toxicity following prostate cancer radiotherapy,³⁶ and xerostomia after IMRT treatments.³⁷ Random survival forest (RSF) is a RF method for the analysis of right-censored survival data. Not only it incorporates all univariate and multivariate effects automatically, but also it can find influential covariates in highly correlated subsets, which is particularly useful in high-dimensional covariate selection problems.^{38,39} The RSF had been employed to explore single and multivariable models of overall survival (OS) and progression-free survival (PFS).⁴⁰

Gradient boosting machines (GBMs) is another popular ensemble ML technique usually use with DTs. It uses gradient boosting, a way to improve any ML model by iteratively training new models that specialize on the weak points of the previous ones. GBMs have been used in radiotherapy to predict radiation-induced pneumonitis¹⁰ and meningioma local failure.¹¹ A version of GBM, XGBoost, consistently wins most Kaggle competitions that involve structured data today.^{10,41} Kaggle is an online platform where private and public entities open data science projects for third parties to compete.

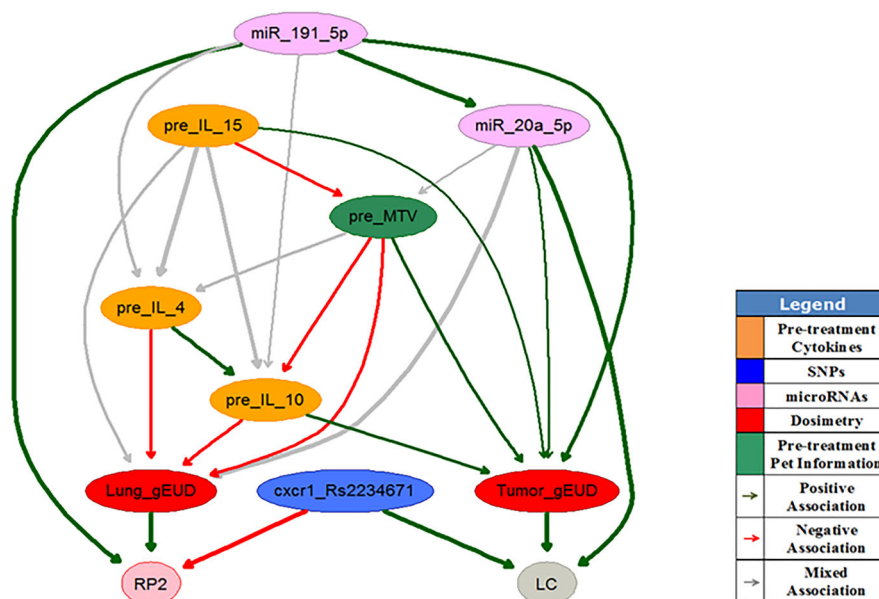


Fig. 1. A joint pretreatment Bayesian networks for local control and 2+ radiation pneumonitis (RP2) prediction.

With all these algorithms previously used, an important question is which algorithm or algorithms should be preferred? It is important to note that it has been proven mathematically that no algorithm is guaranteed to be best in all problems.⁴² However, general recommendations can be made to point out which algorithms are more likely to work better on average. If accuracy is the goal, and provided that there is enough data, RF and GBM are on average the best algorithms for the analysis of structured data.^{43,44} On the contrary, if interpretation and understanding are the goals, linear models, DTs, and BNs are inherently visual, and relationship between variables can be possibly interpreted.^{44,45} In the case of radiation oncology datasets, which tend to be smaller than those analyzed in other fields, it was shown that on average linear models are competitive to RF and GBM.⁴⁶ Therefore, it is our advised that at least these three algorithms should be explored.

2.B. Machine learning approaches for unstructured data (images and text)

Currently, many parts of the medical datasets collected in radiation oncology have unstructured format (e.g., medical notes, images). Understanding how much knowledge can be extracted from them is a challenging and exciting task.^{47–51} Deep learning (DL) is a specific subfield of ML that learns representations from data and it is perfectly suited to handle unstructured data like images or text. The “deep” in DL stands for the idea of the application of successive layers of representations (mathematical operations) as shown in Fig. 2.

Although we usually refer to them with one term, deep neural network (DNN) encompasses many different algorithms. Convolutional neural networks (CNNs), a special kind of DNN, are best known for image-related prediction. Their name stemmed from the successful application of the signal processing operation of convolution. Recurrent neural networks (RNNs) are another variant of DNN specifically useful for learning sequential data, such as voice or text data. The values of DNNs have also been explored in radiotherapy. A DNN was proposed to predict quality of life scores for prostate stereotactic body radiation therapy (SBRT) patients using DVH data.⁵² In a study predicting 2-yr local failure following SBRT,⁵³ a multi-DNN approach analyzing both patient CT simulation scans and clinical risk factors, was

shown to outperform logistic regression, RF, and SVM. To retrieve and classify multimodal medical images (MR, CT, PET) for 24 body organs with different levels of data extraction,⁵⁴ a CNN framework was developed, obtaining an excellent prediction accuracy performance. Additionally, CNNs were used to predict survival risks for rectal cancer patients based on imaging data showing improvement over other approaches.⁵⁵ Also, DNNs were used to explore state transition probabilities for building an autonomous clinical decision support system to adapt patient dose per fraction in a response-adapted treatment setting.⁵⁶ Finally, a DNN algorithm was also shown to outperform other approaches in recognizing clinical texts in medical documents⁵⁷ or learning radiomic features for survival prediction in lung and head & neck cancers.⁵⁸

In general, DL methods appear to be promising for outcome prediction in radiotherapy since it has relatively high prediction performance.⁵⁸ However, high-quality big data size is the key for successful application of the DL algorithms. For instance, Google brain explicitly recommended to not use DL with less than thousands of data points even when transfer learning, an approach that reuses parameters of a model trained on another task with more available data, is used.⁵⁹ These recommendations stem from the fact that DL models usually contain millions of parameters. With many more parameters to tune than other algorithms, overfitting in the cross-validation space is more likely to occur. Therefore, the field of radiotherapy should proceed with caution when these algorithms are explored in small datasets.⁵⁹

2.C. Interpretability and causal inference

The challenges of creating human-like intelligence in machines remain greatly underestimated. Current ML systems lack the essence of human intelligence: understanding the situations we experience or being able to grasp their meaning.⁶⁰ Most ML algorithms today find correlation and not causation.⁶¹ For instance, an ANN (a noninterpretable algorithm) that was developed to triage patients with pneumonia for hospital discharge was found to inadvertently label asthmatic patients as low risk.^{62–64} Colon cancer screening or abnormal breast finding were found to be highly correlated with the risk of having a stroke with no apparent clinical justification.⁶⁵ A CNN was found to be using the hospital system,




FIG. 2. A neural network structure for deep learning.

department or the label “portable” to improve their prediction of pneumonia while disregarding the true findings in the image.⁶⁶ Besides the issue of data quality, current state-of-the-art methods lack the interpretability that is essential for patient stratification and safe clinical decision-making, although ML techniques afford better-than-human performance in numerous domains.

Both clinicians and medical physicists should hesitate to apply prediction models based on black box algorithms.^{67,68} Unfortunately, there is a tradeoff between interpretability and accuracy in ML algorithms. DTs, BNs, and logistic regressions^{69–71} are among the preferred choices in medicine^{72–80} because they create models that are highly interpretable. However, they are consistently outperformed in terms of accuracy by “black box” algorithms used in many other fields, including ensemble methods, meta-algorithms, and DNNs.^{81–84} The tradeoff between accuracy and interpretability represents a long-standing problem in ML, and it is an active area of research from both practical and theoretical points of view.^{64,85–88} Additionally, techniques to interpret black boxes have been created but researchers and clinicians should be aware of their limitations. For instance, variable importance (or pixel importance for images) has limited capability to address how specific variables interact in a highly nonlinear function to build predictive model for a specific patient. Moreover, algorithms designed to explain the black box models do not provide a faithful representation of the original model. As such, the medical physics community should follow with attention recent theoretical developments on causality and interpretability of ML algorithms before we fully deploy the models mentioned above.

3. DISCUSSIONS AND CONCLUSION

Radiotherapy plays an increasingly dominant role in the comprehensive multidisciplinary management of cancer. About half of all cancer patients will receive radiotherapy either as a part of the initial treatment with curative intent or as palliative treatment.⁸⁹ Improved modeling accuracy for normal tissue complications and tumor control probability is highly important and still one of biggest challenges of physicists in radiation therapy today.^{35,90} In this paper, we reviewed a variety of ML approaches for radiation outcome modeling based on structured and unstructured data. Considering a relatively large number of features as the input, ML-based methods have the potential to improve radiation outcome modeling compared to analytical methods. Additionally, complex algorithms like RF, GBM, DL can outperform previously used models like DT, BN, and logistic regression in terms of accuracy. However, both clinicians and medical physicists should proceed with caution when employing these relative new algorithms due to their lack of interpretability. The accuracy on a test set is not a good representation of true performance because the complex algorithms can take advantage of spurious correlation (e.g., asthma status for pneumonia patients). Therefore, we would like to emphasize the need for the development of accurate and interpretable algorithms

as decision support tools for clinical radiation treatment decision making. The use of MediBoost in radiation oncology represents pioneering work in this direction.⁴⁴

We also need to highlight that limited clinical outcomes, inappropriate toxicity classification, genomics data, variation in dose, and fractionation directly affect the accuracy of current models. In fact, most models reported above have accuracy that oscillates around area under ROC curve of 0.70 which is suboptimal for their clinical use. Therefore, a conscientious effort needs to be done in the field to collect better and bigger datasets but many challenges wait ahead.^{91,92} Possible ways to remove these obstacles include, but are not limited to, standardization of processes and data structures, information sharing across institutions, distributed learning, transfer learning, or synthetic data generation using generative adversarial networks.⁹³ Imbalanced data is another intrinsic problem of applying ML approaches for radiation outcome prediction. For instance, if only 10% of patients develop certain toxicity, an algorithm that naively predicts that nobody develops the toxicity is 90% accurate. Balancing the imbalanced event rate during cross-validation and reporting balanced accuracy is necessary, and many techniques can be used to achieve this goal.^{94–96}

Additionally, imaging information is widely used in radiation oncology treatment planning, from dose calculation to implementation of the treatment. Quantitative features, such as radiomics and radiogenomics, extracted from these images can provide useful information about tumor characteristics habitats or evolution with the treatment. Therefore, these features have a potential to improve radiation outcome prediction for precision and personalized radiotherapy through advanced machine and deep learning techniques. As the number of features increases in the era of Big Data, feature selection and dataset size plays an even important role though.

Summarizing, the overwhelming evidence shows that the quality of reporting of prediction model studies is poor today. Although ML algorithms are a powerful tool, their complexity and higher chances of overfitting can result in poor replicability and generalization of studies. Then, model validation is critical for predictive models to be used in the clinical setting. While cross-validation and bootstrapping intend to avoid overfitting in developing the predictive models, external validation through reserved testing dataset or other hospital's datasets is necessary to evaluate their prediction performance. A set of recommendation, named transparent reporting of a multivariable prediction model for Individual Prognosis Or Diagnosis, has been developed.⁹⁷ We strongly encourage journals and researchers to follow these recommendations and clearly specify the level of validation of each study. Moreover, as it current stands, theoretical advances in current ML algorithms are also necessary for the wide acceptance of these models into clinical practice. Since most ML algorithms find correlation and not causation, careful examination and testing needs to be performed before deploying models into clinical practice. As such, further development on causal modeling and/or interpretability of algorithms is necessary.

CONFLICTS OF INTEREST

The authors have no relevant conflicts of interest to disclose.

^{a)}Author to whom correspondence should be addressed. Electronic mail: yiyiLuo@med.umich.edu.

REFERENCES

- Samuel AL. Some studies in machine learning using the game of checkers. *IBM J Res Dev*. 1959;3:210–229.
- 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;125:392–397.
- Shiraishi S, Moore KL. Knowledge-based prediction of three-dimensional dose distributions for external beam radiotherapy. *Med Phys*. 2016;43:378–387.
- Kearney V, Haaf S, Sudhyadhom A, Valdes G, Solberg TD. unsupervised convolutional neural network-based algorithm for deformable image registration. *Phys Med Biol*. 2018;63:185017.
- Interian Y, et al. Deep nets vs expert designed features in medical physics: an IMRT QA case study. *Med Phys*. 2018;45:2672–2680.
- Valdes G, et al. Use of TrueBeam developer mode for imaging QA. *J Appl Clin Med Phys*. 2015;16:322–333.
- Kearney V, et al. Correcting TG 119 confidence limits. *Med Phys*. 2018;45:1001–1008.
- Chen T, Jabbour SK, Qin S, Haffty BG, Yue N. Objected constrained registration and manifold learning: A new patient setup approach in image guided radiation therapy of thoracic cancer. *Med Phys*. 2013;40:041710.
- Peng LC, Huang P, Parekh V, et al. Application of machine learning and radiomics to distinguish true progression from treatment effect after stereotactic radiation therapy for brain metastases. *Int J Radiat Oncol Biol Phys*. 2018;102:S93–S94.
- Valdes G, Solberg TD, Heskell M, Ungar L, Simone CB. Using machine learning to predict radiation pneumonitis in patients with stage I non-small cell lung cancer treated with stereotactic body radiation therapy. *Phys Med Biol*. 2016;61:6105–6120.
- Gennatas ED, Wu A, Braunstein SE, et al. Preoperative and postoperative prediction of long-term meningioma outcomes. *PLoS ONE*. 2018;13:e0204161.
- Chao HH, Valdes G, Luna JM, et al. Exploratory analysis using machine learning to predict for chest wall pain in patients with stage I non-small-cell lung cancer treated with stereotactic body radiation therapy. *J Appl Clin Med Phys*. 2018;19:539–546.
- Held DK. Radiobiology for the radiologist, 6th ed., by Eric J. Hall and Amato J. Giaccia. *Radiat Res*. 2006;166:816–817.
- Lyman JT. Complication probability as assessed from dose volume histograms. *Radiat Res*. 1985;104:S13–S19.
- Guyon I, Elisseeff A. An introduction to variable and feature selection. *J Mach Learn Res*. 2003;3:1157–1182.
- Tang J, Alelyani S, Liu H. Feature selection for classification: a review. In: Aggarwal CC (ed.) *Data Classification: Algorithms and Applications*. Boca Raton, FL: CRC Press; 2014:37–64.
- El Naqa I, Bradley J, Blanco AI, et al. Multivariable modeling of radiotherapy outcomes, including dose-volume and clinical factors. *Int J Radiat Oncol Biol Phys*. 2006;64:1275–1286.
- Deasy JO, El Naqa I. Image-based modeling of normal tissue complication probability for radiation therapy. *Cancer Treat Res*. 2008;139:215–256.
- Das SK, Zhou S, Zhang J, Yin F-F, Dewhurst MW, Marks LB. Predicting lung radiotherapy-induced pneumonitis using a model combining parametric Lyman probit with nonparametric decision trees. *Int J Radiat Oncol Biol Phys*. 2007;68:1212–1221.
- Cosma G, Brown D, Archer M, Khan M, Graham Pockley A. A survey on computational intelligence approaches for predictive modeling in prostate cancer. *Expert Syst Appl*. 2017;70:1–19.
- Dean JA, Wong KH, Welsh LC, et al. Normal tissue complication probability (NTCP) modelling using spatial dose metrics and machine learning methods for severe acute oral mucositis resulting from head and neck radiotherapy. *Radiother Oncol*. 2016;120:21–27.
- Eschrich S, Zhang H, Zhao H, et al. Systems biology modeling of the radiation sensitivity network: a biomarker discovery platform. *Int J Radiat Oncol Biol Phys*. 2009;75:497–505.
- Chen M, Jiang G-L, Fu X-L, et al. The impact of overall treatment time on outcomes in radiation therapy for non-small cell lung cancer. *Lung Cancer*. 2000;28:11–19.
- Munley MT, Lo JY, Sibley GS, Bentel GC, Anscher MS, Marks LB. A neural network to predict symptomatic lung injury. *Phys Med Biol*. 1999;44:2241–2249.
- Su M, Miften M, Whiddon C, Sun X, Light K, Marks L. An artificial neural network for predicting the incidence of radiation pneumonitis. *Med Phys*. 2005;32:318–325.
- Chen SF, Zhou S, Zhang J, Yin F-F, Marks LB, Das SK. A neural network model to predict lung radiation-induced pneumonitis. *Med Phys*. 2007;34:3420–3427.
- El Naqa I, Bradley JD, Deasy J. *Nonlinear kernel-based approaches for predicting normal tissue toxicities*. In: Seventh International Conference on Machine Learning and Applications, Proceedings; 2008:539–544.
- Smith WP, Kim M, Holdsworth C, Liao J, Phillips MH. Personalized treatment planning with a model of radiation therapy outcomes for use in multiobjective optimization of IMRT plans for prostate cancer. *Radiat Oncol*. 2016;11:38.
- Oh JH, and El Naqa I. *Bayesian network learning for detecting reliable interactions of dose-volume related parameters in radiation pneumonitis*. In: Eighth International Conference on Machine Learning and Applications, Proceedings; 2009:484–488.
- Luo Y, El Naqa I, McShan DL, et al. Unraveling biophysical interactions of radiation pneumonitis in non-small-cell lung cancer via Bayesian network analysis. *Radiother Oncol*. 2017;123:85–92.
- Luo Y, et al. Development of a fully cross-validated Bayesian network approach for local control prediction in lung cancer. *IEEE Trans Radiat Plasma Med Sci*. 2019;3:232–241.
- Luo Y, McShan DL, Matuszak MM, et al. A multiobjective Bayesian networks approach for joint prediction of tumor local control and radiation pneumonitis in non-small-cell lung cancer (NSCLC) for response-adapted radiotherapy. *Med Phys*. 2018;45:3980–3995.
- Jochems A, Deist TM, van Soest J, et al. Distributed learning: developing a predictive model based on data from multiple hospitals without data leaving the hospital – a real life proof of concept. *Radiother Oncol*. 2016;121:459–467.
- Palma DA, Senan S, Tsujino K, et al. Predicting radiation pneumonitis after chemoradiation therapy for lung cancer: an international individual patient data meta-analysis. *Int J Radiat Oncol Biol Phys*. 2013;85:444–450.
- Valdes G, Chang AJ, Interian Y, et al. Salvage HDR brachytherapy: multiple hypothesis testing versus machine learning analysis. *Int J Radiat Oncol Biol Phys*. 2018;101:694–703.
- Ospina JD, Zhu J, Chira C, et al. Random forests to predict rectal toxicity following prostate cancer radiation therapy. *Int J Radiat Oncol Biol Phys*. 2014;89:1024–1031.
- Soares I, Dias J, Rocha H, et al. Predicting xerostomia after IMRT treatments: a data mining approach. *Health Technol*. 2018;8:159–168.
- Ishwaran H, Kogalur UB, Blackstone EH, et al. Random survival forests. *Ann Appl Stat*. 2008;2:841–860.
- Ishwaran H, Kogalur UB. Consistency of random survival forests. *Statist Prob Lett*. 2010;80:1056–1064.
- Jain R, Poisson LM, Gutman D, et al. Outcome prediction in patients with glioblastoma by using imaging, clinical, and genomic biomarkers: focus on the nonenhancing component of the tumor. *Radiology*. 2014;272:484–493.
- Chollet F, Allaire JJ. *Deep Learning with R*. Shelter Island, NY: Manning Publications; 2018.
- Wolpert DH. The lack of A priori distinctions between learning algorithms. *Neural Comput*. 1996;8:1341–1390.
- Caruana R, Niculescu-Mizil A. An empirical comparison of supervised learning algorithms. In: *ICML '06 Proceedings of the 23rd international conference on Machine learning*; 2006. Pittsburgh, Pennsylvania, USA: ACM New York, NY, USA

44. Valdes G, Luna JM, Eaton E, Simone CB, Ungar LH, Solberg TD. MediBoost: a patient stratification tool for interpretable decision making in the era of precision medicine. *Sci Rep*. 2016;6:37854.
45. El Naqa I, Li R, Murphy MJ. *Machine Learning in Radiation Oncology: Theory and Applications*. Switzerland: Springer; 2015.
46. Deist TM, Dankers FJWM, Valdes G, et al. Machine learning algorithms for outcome prediction in (chemo)radiotherapy: an empirical comparison of classifiers. *Med Phys*. 2018;45:3449–3459.
47. Ylijoki O, Porras J. Perspectives to definition of big data: a mapping study and discussion. *J Innov Manag*. 2016;4:69–91.
48. Birkhead GS, Klompas M, Shah NR. Uses of electronic health records for public health surveillance to advance public health. *Annu Rev Public Health*. 2015;36:345–359.
49. Thompson RF, Valdes G, Fuller CD, et al. Artificial intelligence in radiation oncology: a specialty-wide disruptive transformation? *Radiother Oncol*. 2018;129:421–426.
50. Feng M, Valdes G, Dixit N, Solberg TD. Machine learning in radiation oncology: opportunities, requirements, and needs. *Front Oncol*. 2018;8:110.
51. Morin O, Vallières M, Jochems A, et al. A deep look into the future of quantitative imaging in oncology: a statement of working principles and proposal for change. *Int J Radiat Oncol Biol Phys*. 2018;102:1074–1082.
52. Qi X, Neylon J, Santhanam A. Dosimetric predictors for quality of life after prostate stereotactic body radiation therapy via deep learning network. *Int J Radiat Oncol Biol Phys*. 2017;99:S167–S167.
53. Aneja S, Shaham U, Kumar RJ, et al. Deep neural network to predict local failure following stereotactic body radiation therapy: integrating imaging and clinical data to predict outcomes. *Int J Radiat Oncol Biol Phys*. 2017;99:S47.
54. Qayyum A, Anwar SM, Awais M, Majid M. Medical image retrieval using deep convolutional neural network. *Neurocomputing*. 2017;266:8–20.
55. Li H, Zhong H, Boimel Pj, Ben-Josef E, Xiao Y, Fan Y. Deep convolutional neural networks for imaging based survival analysis of rectal cancer patients. *Int J Radiat Oncol Biol Phys*. 2017;99:S183.
56. Tseng HH, Luo Y, Cui S, Chien J-T, Ten Haken RK, El Naqa I. Deep reinforcement learning for automated radiation adaptation in lung cancer. *Med Phys*. 2017;44:6690–6705.
57. Wu Y, Jiang M, Lei J, Xu H. Named entity recognition in chinese clinical text using deep neural network. *Stud Health Technol Inform*. 2015;216:624–648.
58. Jochems A, Deist TM, El Naqa I, et al. Developing and validating a survival prediction model for NSCLC patients through distributed learning across 3 countries. *Int J Radiat Oncol Biol Phys*. 2017;99:344–352.
59. Valdes G, Interian Y. Comment on 'deep convolutional neural network with transfer learning for rectum toxicity prediction in cervical cancer radiotherapy: a feasibility study'. *Phys Med Biol*. 2018;63:068001.
60. Mitchell M. Artificial Intelligence Hits the Barrier of Meaning; 2018. Available from <https://www.nytimes.com/2018/11/05/opinion/artificial-intelligence-machine-learning.html>
61. Meyer P, Noblet V, Mazzara C, Lallemand A. Survey on deep learning for radiotherapy. *Comput Biol Med*. 2018;98:126–146.
62. Cooper GF, Aliferis CF, Ambrosino R, et al. An evaluation of machine-learning methods for predicting pneumonia mortality. *Artif Intell Med*. 1997;9:107–138.
63. Cooper GF, Abraham V, Aliferis CF, et al. Predicting dire outcomes of patients with community acquired pneumonia. *J Biomed Inform*. 2005;38:347–366.
64. Caruana R, Lou Y, Gehrke J, Koch P, Sturm M, Elhadad N. Intelligible models for healthcare: Predicting pneumonia risk and hospital 30-day readmission. In: *The 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. Sydney: Australia; 2015.
65. Mullainathan S, Obermeyer Z. Does machine learning automate moral hazard and error? *Am Econ Rev*. 2017;107:476–480.
66. Zech JR, Badgeley MA, Liu M, Costa AB, Titano JJ, Oermann EK. Variable generalization performance of a deep learning model to detect pneumonia in chest radiographs: a cross-sectional study. *PLoS Medicine*. 2018;15:e1002683.
67. Kang J, Schwartz R, Flickinger J, Beriwal S. Machine learning approaches for predicting radiation therapy outcomes: a clinician's perspective. *Int J Radiat Oncol Biol Phys*. 2015;93:1127–1135.
68. Lynch CM, van Berkel VH, Frieboes HB. Application of unsupervised analysis techniques to lung cancer patient data. *PLoS ONE*. 2017;12:e0184370.
69. Quinlan JR. *C4.5: Programs for Machine Learning*. San Mateo: Morgan Kaufmann; 1993.
70. Breiman L, Friedman JH, Olshen RA, Stone CJ. *Classification and Regression Trees*. Boca Raton, FL: Chapman and Hall/CRC; 1984.
71. Pearl J. *Probabilistic Reasoning in Intelligent Systems*. San Mateo: Morgan Kaufmann; 1988.
72. Cain KP, McCarthy KD, Heilig CM, et al. An algorithm for tuberculosis screening and diagnosis in people with HIV. *N Engl J Med*. 2010;362:707–716.
73. Berlowitz DR, Ash AS, Hickey EC, et al. Inadequate management of blood pressure in a hypertensive population. *N Engl J Med*. 1998;339:1957–1963.
74. Haydel MJ, Preston CA, Mills TJ, Luber S, Blaudeau E, DeBlieux PMC. Indications for computed tomography in patients with minor head injury. *N Engl J Med*. 2000;343:100–105.
75. Chen HY, Yu S-L, Chen C-H, et al. A five-gene signature and clinical outcome in non-small-cell lung cancer. *N Engl J Med*. 2007;356:11–20.
76. Gilbert MR, Dignam JJ, Armstrong TS, et al. A randomized trial of bevacizumab for newly diagnosed glioblastoma. *N Engl J Med*. 2014;370:699–708.
77. Lionetti E, Castellana S, Francavilla R, et al. Introduction of gluten, HLA status, and the risk of celiac disease in children. *N Engl J Med*. 2014;371:1295–1303.
78. Spiegelhalter DJ. Probabilistic expert systems in medicine: practical issues in handling uncertainty. *Stat Sci*. 1987;2:3–44.
79. Shortliffe EH. The adolescence of Ai in medicine: will the field come of age in the 90s. *Artif Intell Med*. 1993;5:93–106.
80. Warner HR, Sorenson DK, Bouhaddou O. *Knowledge Engineering in Health Informatics*. Berlin, NY: Springer Verlag; 1997.
81. Breiman L. Random forests. *Mach Learn*. 2001;45:5–32.
82. Freund Y, Schapire RE. A decision-theoretic generalization of on-line learning and an application to boosting. *J Comput Syst Sci*. 1997;55:119–139.
83. Friedman J, Hastie T, Tibshirani R. Additive logistic regression: a statistical view of boosting – rejoinder. *Ann Stat*. 2000;28:400–407.
84. Friedman JH. Greedy function approximation: a gradient boosting machine. *Ann Stat*. 2001;29:1189–1232.
85. Cabitza F, Rasoini R, Gensini GF. Unintended consequences of machine learning in medicine. *JAMA*. 2017;318:517–518.
86. Castelvecchi D. Can we open the black box of AI? *Nat News*. 2016; 538:20–23.
87. Ribeiro MT, Singh S, Guestrin C. *Why should i trust you?: explaining the predictions of any classifier*. In: *The 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM: San Francisco. USA: CA; 2016.
88. Kim B, Shah J, Doshi-Velez F. *Mind the Gap: A Generative Approach to Interpretable Feature Selection and Extraction*. In: *Advances in Neural Information Processing Systems 28 (NIPS 2015)*; 2015:28.
89. Grau C, Hoyer M. High-precision Radiotherapy. *Business Briefing: European Oncology Review*; 2005.
90. Bortfeld T, Jeraj R. The physical basis and future of radiation therapy. *Br J Radiol*. 2011;84:485–498.
91. El Naqa I, Ruan D, Valdes G, et al. Machine learning and modeling: data, validation, communication challenges. *Med Phys*. 2018;45:E834–E840.
92. Lambin P, van Stiphout RGPM, Starman MHW, et al. Predicting outcomes in radiation oncology–multifactorial decision support systems. *Nat Rev Clin Oncol*. 2013;10:27–40.
93. Goodfellow IJ, Mirza M, Courville A, Bengio Y. *Generative adversarial nets*. In: *Advances in Neural Information Processing Systems 27 (NIPS 2014)*; 2014: 27.
94. Blagus R, Lusa L. SMOTE for high-dimensional class-imbalanced data. *BMC Bioinform*. 2013;14:106.
95. Parker BJ, Guenter S, Bedo J. Stratification bias in low signal microarray studies. *BMC Bioinform*. 2007;8:326.
96. Schiller TW, Chen Y, El Naqa I, Deasy JO. Modeling radiation-induced lung injury risk with an ensemble of support vector machines. *Neurocomputing*. 2010;73:1861–1867.
97. Collins GS, Reitsma JB, Altman DG, Moons KGM. Transparent reporting of a multivariable prediction model for individual prognosis or diagnosis (TRIPOD) response. *Ann Intern Med*. 2015;162:735–736.