Knowledge Models as Teaching Aid for Training Intensity Modulated Radiation Therapy Planning: A Lung Cancer Case Study

Matt Mistro [1,2], Yang Sheng [1]*, Yaorong Ge [3], Chris R. Kelsey [1], Jatinder R. Palta [4], Jing Cai [5], Qiuwen Wu [1], Fang-Fang Yin [1] and Q. Jackie Wu [1]

[1] Department of Radiation Oncology, Duke University Medical Center, Durham, NC, United States, [2] Medical Physics Graduate Program, Duke University, Durham, NC, United States, [3] Department of Software and Information Systems, University of North Carolina at Charlotte, Charlotte, NC, United States, [4] Department of Radiation Oncology, Virginia Commonwealth University, Richmond, VA, United States, [5] Department of Health Technology and Informatics, Hong Kong Polytechnic University, Hong Kong, China

Purpose: Artificial intelligence (AI) employs knowledge models that often behave as a black-box to the majority of users and are not designed to improve the skill level of users. In this study, we aim to demonstrate the feasibility that AI can serve as an effective teaching aid to train individuals to develop optimal intensity modulated radiation therapy (IMRT) plans.

Methods and Materials: The training program is composed of a host of training cases and a tutoring system that consists of a front-end visualization module powered by knowledge models and a scoring system. The current tutoring system includes a beam angle prediction model and a dose-volume histogram (DVH) prediction model. The scoring system consists of physician chosen criteria for clinical plan evaluation as well as specially designed criteria for learning guidance. The training program includes six lung/mediastinum IMRT patients: one benchmark case and five training cases. A plan for the benchmark case is completed by each trainee entirely independently pre- and post-training. Five training cases cover a wide spectrum of complexity from easy (2), intermediate (1) to hard (2). Five trainees completed the training program with the help of one trainer. Plans designed by the trainees were evaluated by both the scoring system and a radiation oncologist to quantify planning quality.

Results: For the benchmark case, trainees scored an average of 21.6% of the total max points pre-training and improved to an average of 51.8% post-training. In comparison, the benchmark case's clinical plans score an average of 54.1% of the total max points. Two of the five trainees' post-training plans on the benchmark case were rated as comparable to the clinically delivered plans by the physician and all five were noticeably improved by the physician's standards. The total training time for each trainee ranged between 9 and 12 h.

Conclusion: This first attempt at a knowledge model based training program brought unexperienced planners to a level close to experienced planners in fewer than 2 days. The proposed tutoring system can serve as an important component in an AI ecosystem that will enable clinical practitioners to effectively and confidently use KBP.

Keywords: knowledge model, lung cancer, machine learning, tutoring system, intensity modulated radiation therapy

INTRODUCTION

Knowledge models collect and extract important patterns and knowledge from high quality clinical plans and utilize them to predict clinically optimal solutions for new cases. For treatment planning, this comes in the form of selected beam angles, optimized collimator settings, predicted achievable dose-volume histogram (DVH) endpoints for inverse optimization, and combined multiple parameter predictions for a fully automated treatment planning process (Zhu et al., 2011; Breedveld et al., 2012; Zhang et al., 2012, 2018, 2019a,b; Good et al., 2013; Voet et al., 2013; Zarepisheh et al., 2014; Sheng et al., 2015, 2019; Yuan et al., 2015, 2018; Hazell et al., 2016). Knowledge models have been successfully used in the clinical workflow for fully automated planning for some simpler cancer sites like prostate (Voet et al., 2014), but for more complicated sites, there may yet be some hurdles to overcome. Due to the limitation of training samples and other factors, they are often simplified to improve generalizability by regulating the capability of handling a wide array of niche scenarios in which a human planner would be better fit to tackle. Despite this, there is a lot to be gained from investigating the implicit knowledge of these models. The simple, logical principles that most of these models are built upon can not only start a foundation for less experienced users to progress toward clinical reliability but also bridge the gap between human and model knowledge in what to look for in evaluation and identification of planning intricacies. The goal is to make a human-centered artificial intelligence (AI) system to exploit the strengths from both ends and efficiently train competent planners.

While extensive training and arduous hours of practice can certainly cultivate competent and professional planners, more effective training programs are urgently needed to help more planners become proficient in the clinic as technologies continue to become more advanced and more complex. Of course, there are aspects of planning that can only be obtained by years of nuanced planning, but plan quality is not always shown to be better in those who have more experience (Nelms et al., 2012). Some planners with planning experience may encounter a bottle-neck in improving their versatility in planning various scenarios, due to the lack of understanding of the underlying subtlety which can be readily provided and instructed by the knowledge-based models. In addition, training a planner to a highly proficient level in a traditional mentor-tutor fashion is expensive in time and resources, and sometime the limited training resources are dispatched to more entry level learners and/or regional centers. A person can quickly learn how to plan well if the teaching
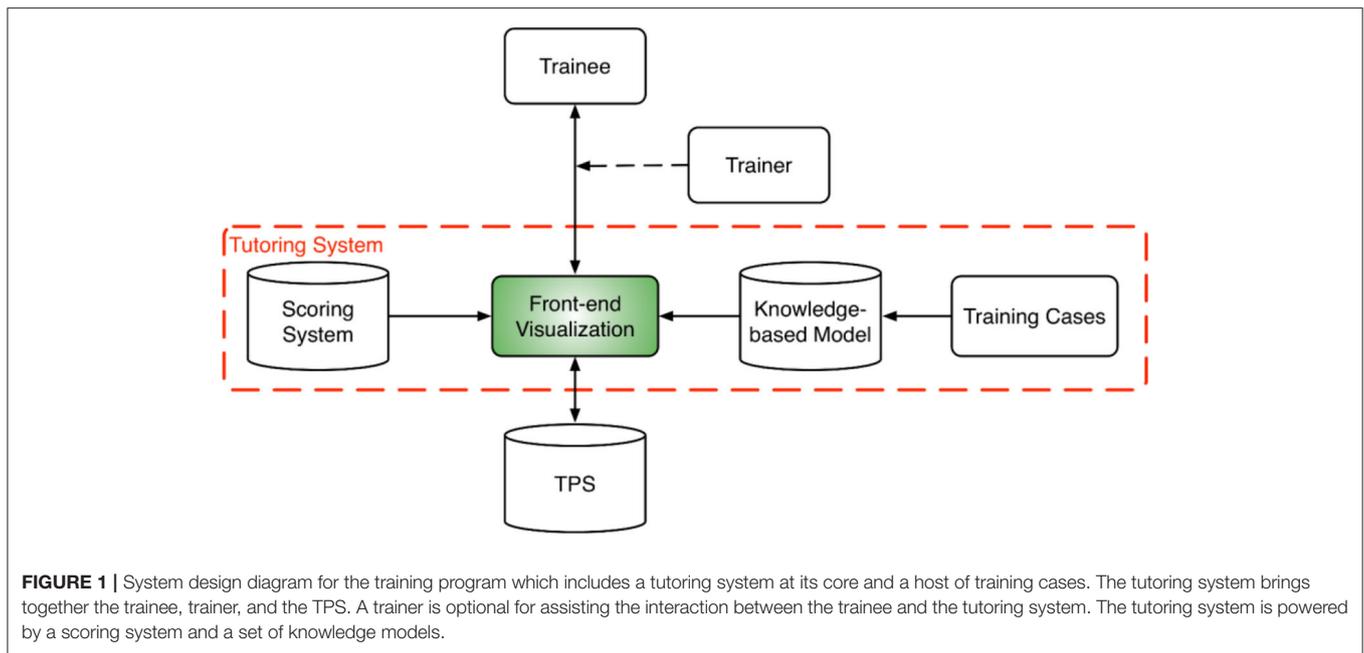
is well-thought out and provides the base for the person to build their own intuition. A training program that introduces the benefit of knowledge-based models can accomplish this and aid in tearing down the notion of these models being entirely a black box which has been restrictive to clinical usage of models. Such a program can be a catalyst to bring more models into routine clinical work by showing how they work and what the best practice is. This study examines the workflow and feasibility of a training program that takes advantage of two knowledge-based models (Yuan et al., 2012, 2018) with carefully developed scoring criteria to facilitate efficient and quality learning of lung IMRT treatment planning to help establish intuition to trainees with no previous clinical planning experience.

The proposed training program lays the foundation for an entirely self-sufficient training module that will be designed as a constraint-based intelligent tutoring system (ITS) (Mitrovic et al., 2007, 2013; Dermeval et al., 2018). The constraint-based approach supports the type of learning problem that does not have an explicit solution or path for a user to follow as is the case of IMRT planning. The constraints are defined in the form of the scoring system, and the end goal is for the user to learn the planning actions that optimize the scoring system to obtain the highest score possible. In this constraint based framework, the user has to forge their own path from the information that is directed to them, and two people can take entirely different strategies and arrive at good solutions. This is a proof-of-concept study to show that there is valuable information to be gained from the knowledge models and they can be effectively and efficiently used in training new planners and give them the ability to utilize these models to generate quality plans. Here, we define new planners as those who have completed adequate medical physics course work but have minimal clinical treatment planning practice. As such, they would have completed classroom instructions of radiation therapy physics and advanced treatment planning. They would have basic operational knowledge of the TPS system, but have no experience in planning real clinical cases.

METHODS AND MATERIALS

Training Program Design
Program Overview

The overall training program design is shown in **Figure 1**. At the core of the training program is the tutoring system which consists of a front-end visualization module powered by KBP models and a plan scoring system. The visualization module

FIGURE 1 | System design diagram for the training program which includes a tutoring system at its core and a host of training cases. The tutoring system brings together the trainee, trainer, and the TPS. A trainer is optional for assisting the interaction between the trainee and the tutoring system. The tutoring system is powered by a scoring system and a set of knowledge models.

(**Figure 2**) provides the vital interactive workspace for the trainee and trainer, while the KBP models and scoring system provides back-end knowledge support. The KBP models currently include a beam bouquet prediction model and a dose-volume histogram (DVH) prediction model, while the scoring system consists of physician chosen criteria for clinical plan evaluation as well as specially designed criteria for learning guidance. These additional specially designed criteria were designed to help trainee understand the full scope of treatment planning and eventually achieve the ability to create a high quality plan, especially focusing on the criteria that are often qualitatively evaluated by the physician such as the overall isodose line conformity. Further, the tutoring system works in concert with the clinical treatment planning system (TPS) as trainees learn to generate clinically plans in a realistic clinical planning environment. In this study, we use the Eclipse® TPS (Varian Medical Systems, Palo Alto, CA) which provides fluence map optimization and dose calculation.

The current training program utilizes six lung/mediastinum IMRT patient cases: one benchmark case (shown in **Figure 3**) and five training cases. Each case is composed of clinical images, structures, and a delivered plan which were de-identified before incorporated into the training program. The benchmark case is used to track skill development. The five training cases cover the complexity from easy (2), intermediate (1) to hard (2) in lung IMRT planning. The difficulty level is determined by an experienced planner who evaluated the prescription, tumor size, complexity of shape, and proximity to organs-at-risk (OARs). The benchmark case, considered "intermediate-to-hard," has a target volume of 762.8 cc and a prescription of 62 Gy; two "easy" training cases have an average target volume of 113.8 cc and prescription of 40 Gy (reduced dose due to prior treatment); the "intermediate" training case has a target volume of 453.0 cc and a
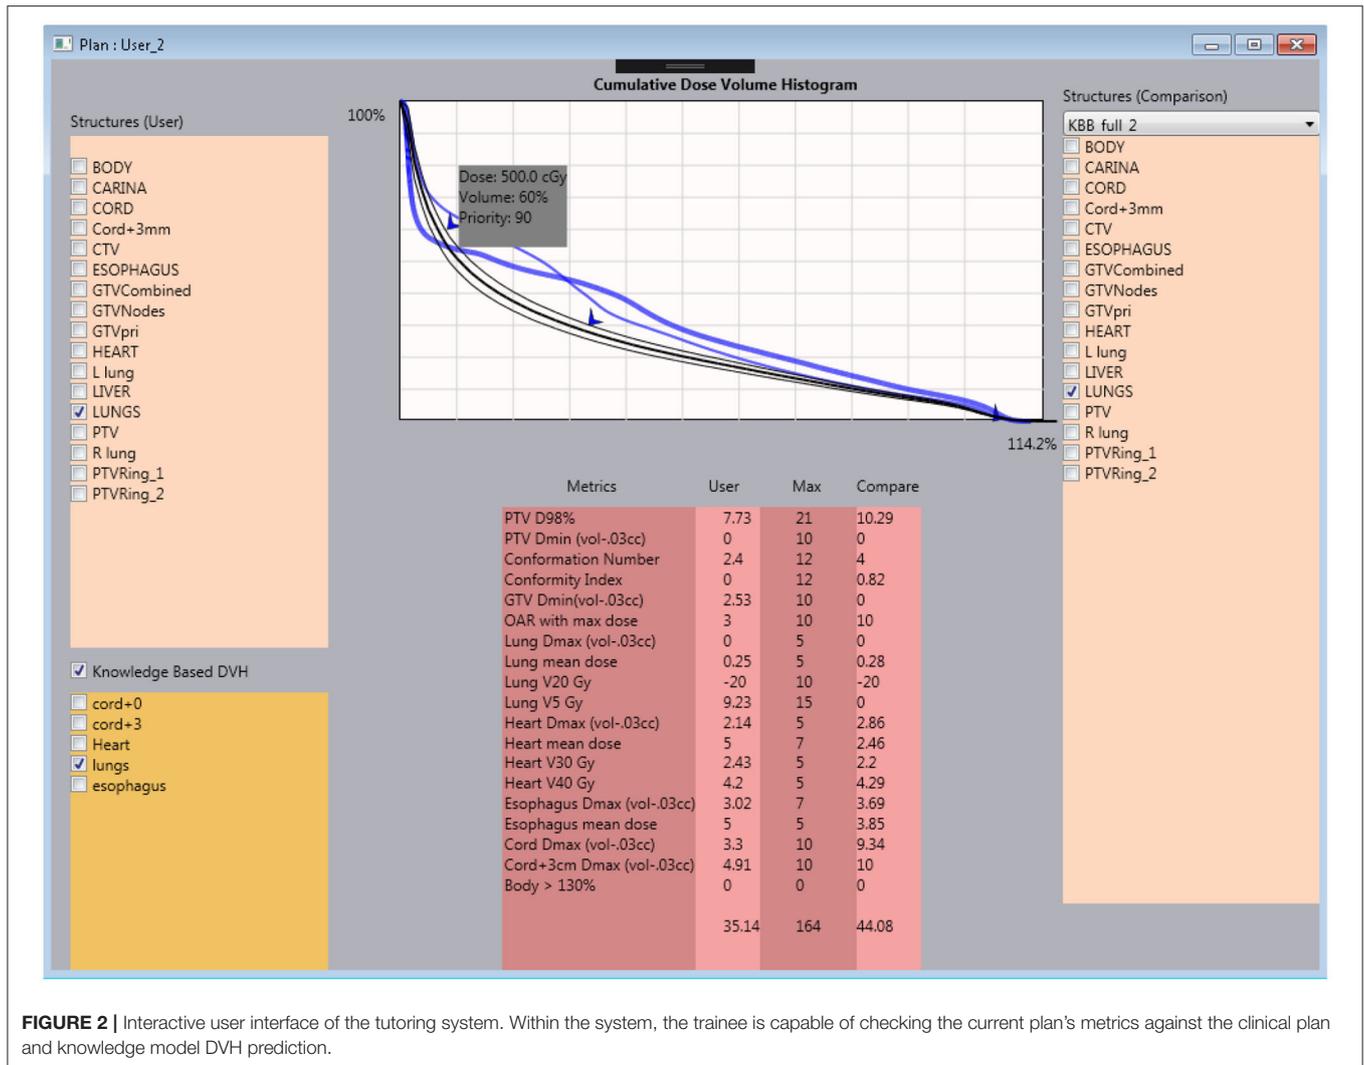
prescription of 60 Gy; two "hard" training cases have an average target volume of 845.7 cc and prescriptions of 60 Gy.

Before training begins, each trainee undergoes a benchmarking process to determine baseline score. In this process, the trainee is introduced to the treatment planning system with functionality they might not be familiar with as they have no prior experience. They are provided with the scoring metrics and asked to plan the benchmark case without any intervention from the trainer or the tutoring system (referred to as the baseline plan). The trainee is instructed that they have the choice of 6 or 10 MV beams and could have no more than 11 beams to align with current clinical practice.

Training Workflow

Figure 4 illustrates the typical training workflow (solid lines) for learning to plan one training case. The cases are selected sequentially from the easy ones to the difficult ones. For each case, a trainee goes through two phases of training: the beam selection phase and the fluence map optimization phase. In both phases, each training episode involves three main steps: (1) the trainee makes a decision (or takes an action); (2) the training program generates a plan corresponding to the decision and displays relevant dose metrics; (3) the training program then generates a comparison plan according to predictions from knowledge models and displays the same set of relevant dose metrics for comparison. The majority of interaction centers around the process with which the trainee learns to explain the differences between their plan and the comparison plan, as well as the resulting dosimetric implications of those differences.

During the beam selection phase, the trainee can choose the number of beams (seven to 11) and the direction of beams. The comparison plans are those generated with beams

Plan : User_2

Cumulative Dose Volume Histogram

100%

Dose: 500.0 cGy
Volume: 60%
Priority: 90

114.2%

Structures (User)

- BODY
- CARINA
- CORD
- Cord+3mm
- CTV
- ESOPHAGUS
- GTVCombined
- GTVNodes
- GTVpri
- HEART
- L lung
- LIVER
- ☑ LUNGS
- PTV
- R lung
- PTVRing_1
- PTVRing_2

☑ Knowledge Based DVH

- cord+0
- cord+3
- Heart
- ☑ lungs
- esophagus

Structures (Comparison)

KBB full 2

- BODY
- CARINA
- CORD
- Cord+3mm
- CTV
- ESOPHAGUS
- GTVCombined
- GTVNodes
- GTVpri
- HEART
- L lung
- LIVER
- ☑ LUNGS
- PTV
- R lung
- PTVRing_1
- PTVRing_2

Metrics	User	Max	Compare
PTV D98%	7.73	21	10.29
PTV Dmin (vol-.03cc)	0	10	0
Conformation Number	2.4	12	4
Conformity Index	0	12	0.82
GTV Dmin(vol-.03cc)	2.53	10	0
OAR with max dose	3	10	10
Lung Dmax (vol-.03cc)	0	5	0
Lung mean dose	0.25	5	0.28
Lung V20 Gy	-20	10	-20
Lung V5 Gy	9.23	15	0
Heart Dmax (vol-.03cc)	2.14	5	2.86
Heart mean dose	5	7	2.46
Heart V30 Gy	2.43	5	2.2
Heart V40 Gy	4.2	5	4.29
Esophagus Dmax (vol-.03cc)	3.02	7	3.69
Esophagus mean dose	5	5	3.85
Cord Dmax (vol-.03cc)	3.3	10	9.34
Cord+3cm Dmax (vol-.03cc)	4.91	10	10
Body > 130%	0	0	0
	35.14	164	44.08

FIGURE 2 | Interactive user interface of the tutoring system. Within the system, the trainee is capable of checking the current plan's metrics against the clinical plan and knowledge model DVH prediction.
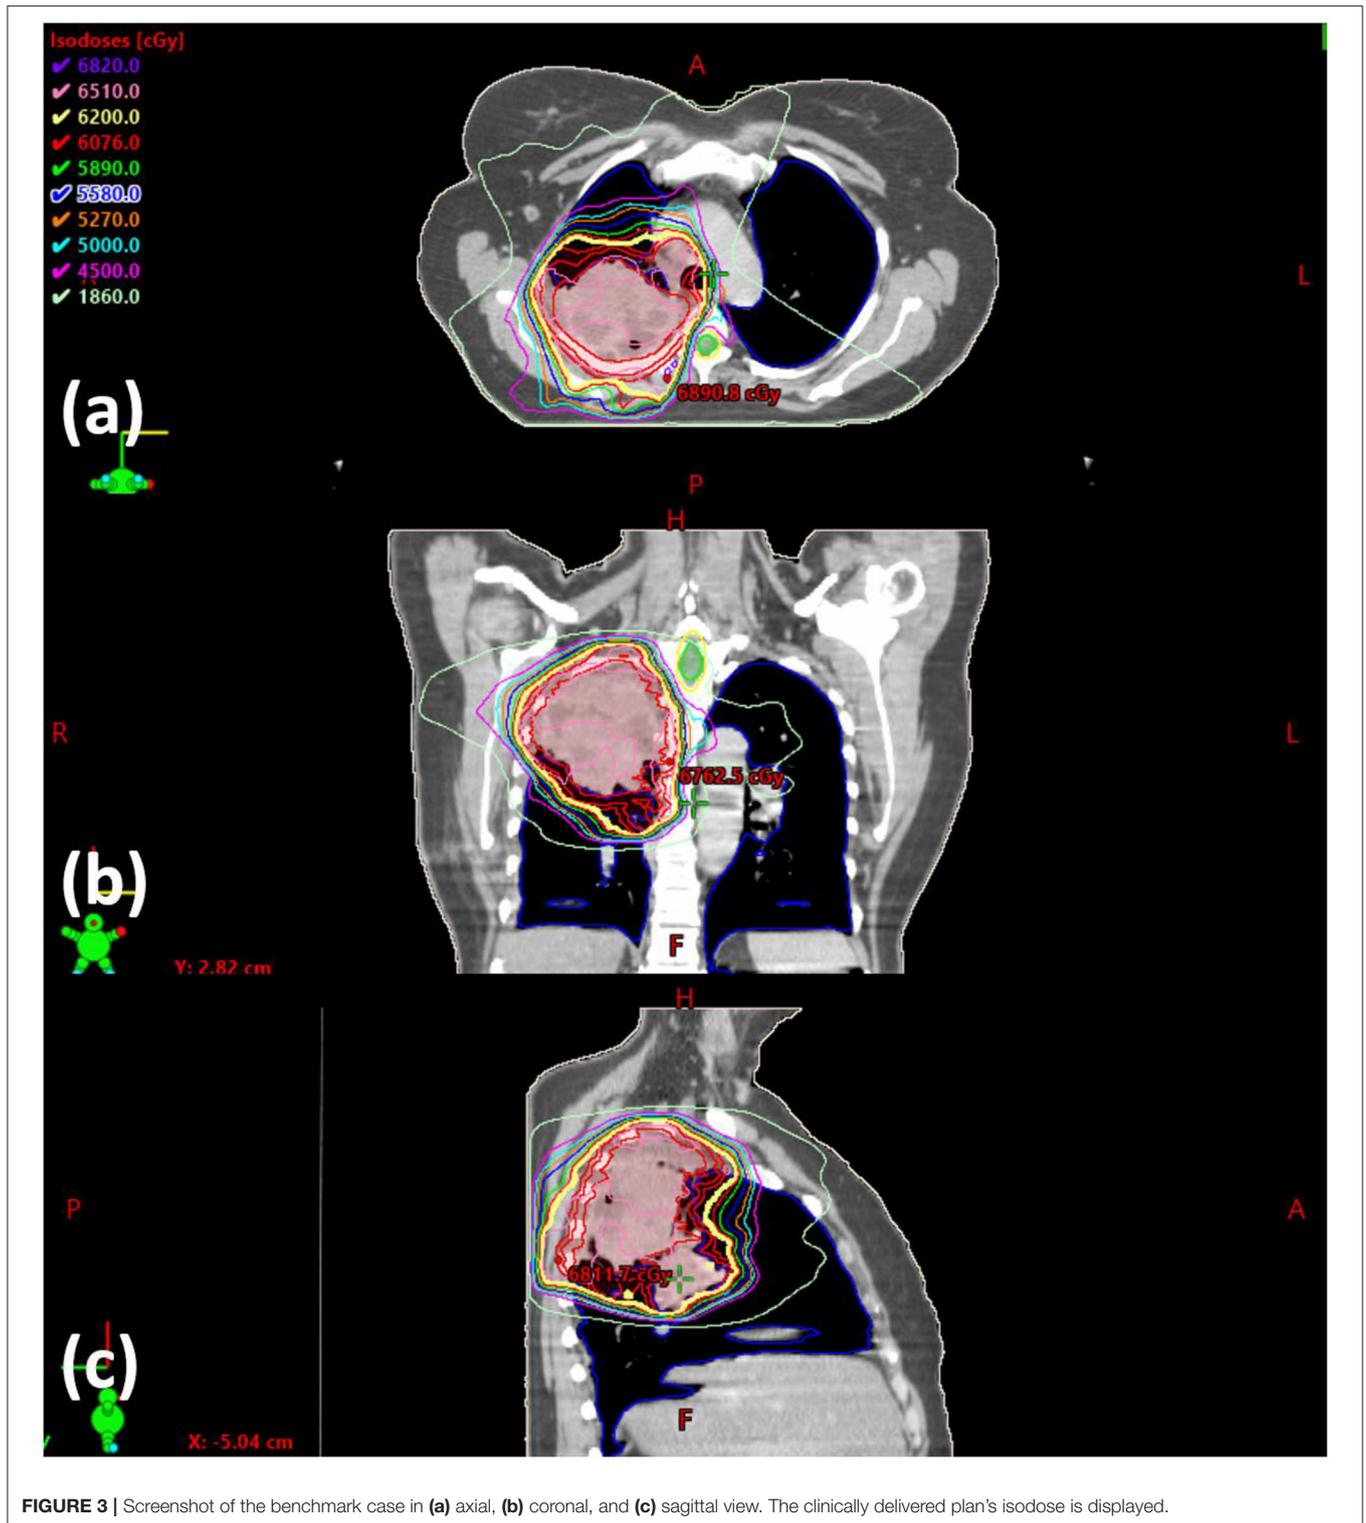
determined by the knowledge-based beam selection model. Both trainee plans and comparison plans are created by an automatic KBP algorithm. The trainee determines whether they prefer to move along the direction of model prediction or continue with their own direction. At the end of this phase, the final beam comparison provides an assessment of the expected dosimetric differences contributed by trainee's beam design.

When the optimization training phase begins, the trainee creates the initial optimization objectives and finishes the planning process. In parallel, a comparison plan is generated with the KBP beam setting using trainee's dose-volume constraints. Dosimetric comparisons between plans allow the trainee to appreciate whether the results aligned with their expectations during the aforementioned assessment, which builds a forward intuition on beam choice implications.

Following this, the KBP DVH model is imported and the trainee is able to compare their plan's DVHs and dose objectives to where the DVH model predicts they should be able to achieve. The trainee then makes changes based on what they see is

obtainable. After the changes are made and the plan is scored, a final comparison is done with the clinically delivered plan. The trainee works backwards by looking at the scoring and DVH of the clinical plan and ponders on how the clinically delivered plan might have been achieved. This is to further ingrain a backwards intuition for the metrics related to certain collective beam arrangements.

As shown in **Figure 4**, the training workflow also includes a few steps (dashed boxes) that are designed for people with little to no knowledge of treatment planning. These steps are optional when trainees are at more advanced stages during the training process. The first beam assessment is an initial guidance with the trainer about the best beam direction to select if they were to make a plan with only a single beam. This step encourages the trainee to think about how each individual beam will contribute to the final dose distribution. The second beam assessment helps the trainee make an optimal plan when only two beams are used. This helps planners understand how multiple beams interact with one another (i.e., the second best beam isn't necessarily the

FIGURE 3 | Screenshot of the benchmark case in **(a)** axial, **(b)** coronal, and **(c)** sagittal view. The clinically delivered plan's isodose is displayed.

best beam to work with the first). Lastly, the "basic constraint assessment" step is a simple check to ensure that the trainee has at least one objective for all the relevant structures and two for the target.

The current training program takes the trainee through the workflow described in **Figure 4** five times, one for each training case, in increasing order of difficulty. After completing all five cases, the trainee returns to the benchmark case and creates a
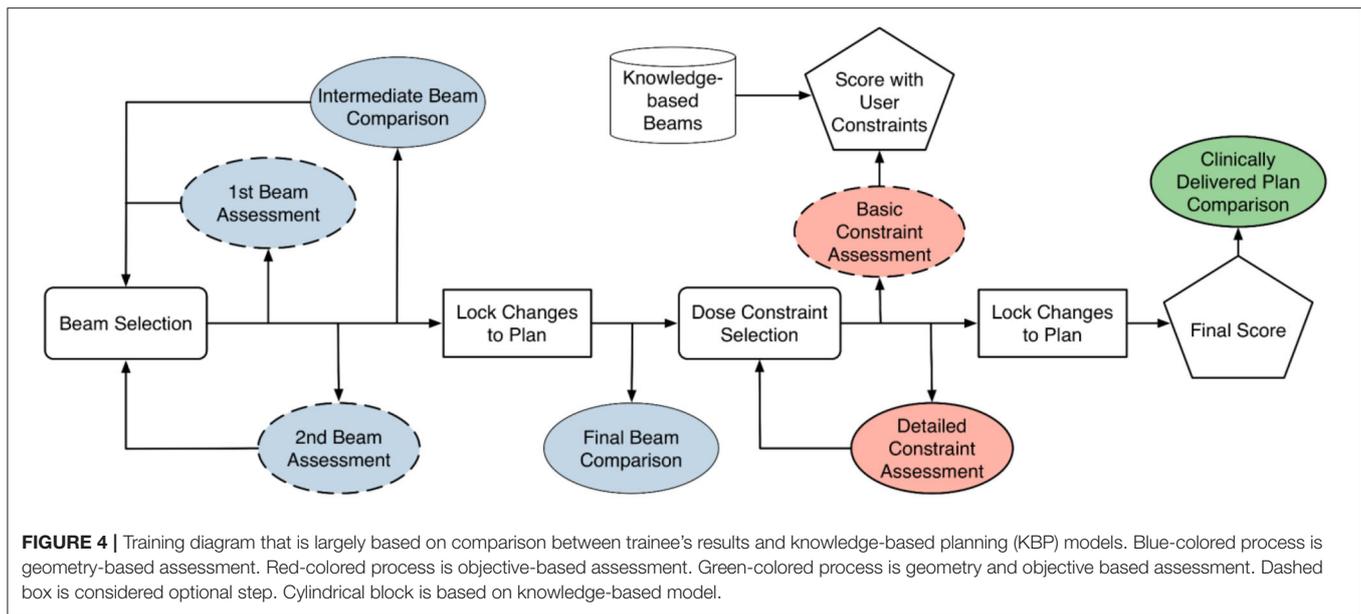
FIGURE 4 | Training diagram that is largely based on comparison between trainee's results and knowledge-based planning (KBP) models. Blue-colored process is geometry-based assessment. Red-colored process is objective-based assessment. Green-colored process is geometry and objective based assessment. Dashed box is considered optional step. Cylindrical block is based on knowledge-based model.

new plan entirely on their own without any intervention from the trainer, knowledge models, or the tutoring system. This post-training plan in comparison to the baseline plan on the same case provides an objective way to assess if there is any significant improvement in their planning ability.

Tutoring System Design

As introduced in the previous section, the current tutoring system includes three major components: a visualization module for user interaction, knowledge models for planning guidance, and a scoring system for plan assessment. The visualization module is integrated with the Eclipse® TPS and is currently implemented as a script using the Eclipse® API. In the following, we provide a brief description of the knowledge models and the scoring system.

Beam Angle Selection Model

The beam model (Yuan et al., 2015, 2018) predicts the best beam configuration for each new case, including the number of beams and the angle of the beams. It operates on a novel beam efficiency index that tries to maximize the dose delivered to a PTV and minimize the dose delivered to OARs based on a number of weighting factors. It also introduces a forced separation among good quality beams to cover sufficient co-planar space. The weighting factors and other parameters of the beam model are learned from a set of high quality prior clinical cases (Yuan et al., 2018). For the purposes of simplicity of introduction to new planners, all beams in the current training program are restricted to co-planar beams.

DVH Prediction Model

The DVH prediction model estimates the best achievable DVH of the OARs based on a number of anatomical features: distance-to-target histogram (DTH) principal components, OAR

volume, PTV volume, OAR-PTV overlapping volume, and out-of-field OAR volume (Yuan et al., 2012). The model is trained with a set of prior lung cases with a variety of tumor sizes and locations. For this study, the model predicts DVHs that are useful for the trainees during the learning and planning. Organs-at-risk included in each DVH are cord, cord+3 mm, lungs, heart, and esophagus.

Plan Scoring System

A plan scoring system was designed to help trainees understand the quality of different plans from the choices of beams and DVH parameters. Therefore, the scoring system incorporates both physician's clinical evaluation criteria and planning knowledge. The metrics with their respective max point values are shown in **Table 1**. As noted, since each case has its own unique anatomy and complexity, the most achievable points of a plan is always less than the total max points, while more difficult cases have lower best achievable points. The best achievable points of each plan are not normalized so the trainees are encouraged to rely on the actual planning knowledge to "do their best," rather than to get "100 percent score" or gaming the system. There were a total of 164 points, with which the raw score was normalized to represent the percentage score. A maximum scoring would have 100% percentage score. Normalization was performed after the training was done as a summary of the data. It is worth reiterating that the trainee was unaware of the maximally achievable score for each case so they couldn't game the system. Note that even clinically delivered plans may not be perfect in all categories, and therefore, may not achieve the highest possible scores.

An effective scoring system can be created in many ways. The current system starts with the logic of rewarding dosimetric endpoints that are clinically relevant as explained in RTOG reports (Chun et al., 2017), other clinical considerations (Kong et al., 2011; Baker et al., 2016) and planning competitions powered by ProKnow (ProKnow Systems, Sanford, FL;

TABLE 1 | Metrics chosen to be a part of the scoring system and their respective maximum point value.

Target	Max	Lung	Max	Heart	Max	Esophagus	Max	Spinal cord	Max
PTV D98%	21	Max dose	5	Max dose	5	Max dose	7	Cord max dose	10
PTV min dose	10	Mean dose	5	Mean dose	7	Mean dose	5	Cord+3 mm max dose	10
GTV min dose	10	V20 Gy	10	V30 Gy	5				
CN 95%	12	V5 Gy	15	V40 Gy	5				
CI 50%	12								
Location of max dose	10								

PTV, Planning Target Volume; GTV, Gross Tumor Volume; CN, Conformity Number; CI, Conformity Index.

www.proknowsystems.com). The conformity index (CI) (Knoos et al., 1998) aims to limit the isodose volume. The conformation number (CN) or Paddick conformity index (Paddick, 2000) follows similar logic to CI but focuses on the portion of that isodose volume within the target.

Training Program Assessment

To assess the effectiveness of the training program, five trainees who satisfy the criteria of new planners went through the entire training program. For all five trainees, the baseline and post-training plans of the benchmark case were scored and analyzed for evidence of learning. Furthermore, the post-training plans of all the training cases as well as all the clinically delivered plans were also scored and analyzed for trainee performance and potential knowledge gaps.

Moreover, to assess how the overall scores given by the scoring system closely reflect true plan quality in a real clinical scenario, a physician who specializes in the treatment of lung cancer evaluated each of the plans to provide an expert opinion on their clinical quality. For each trainee, the post-training plan of the benchmark case was first compared with the baseline plan of the same case and then against the clinically delivered plan by the physician. Each comparison was categorized on a simplified 5-point scale of (1) significantly worse, (2) moderately worse, (3) comparable, (4) moderately better, and (5) significantly better. The physician also evaluated the trainee's post-training plans on whether they could be approved for clinical delivery.

RESULTS

Scoring Results

Five trainees went through the training program and their scores are shown in **Figure 5**. All trainees went through multiple classroom courses on radiation physics, anatomy, radiation biology, and treatment planning/dosimetry. They also completed a basic practicum course to learn the essential operations of a treatment planning system. After training, the overall score of all trainees was unanimously improved from the baseline and was much closer to that of the clinically delivered plan (**Figure 5A**). Trainee 1 and 3 received a planning score point that was slightly above that of the clinically delivered plan, with an average of 54.4%, while the other three were marginally lower with an average of 50.1%. In comparison, the score of

the clinically delivered plan was 54.1%. Detailed scores are listed in **Table 2**. For the five cases used within the training program (**Figure 5B**), every trainee obtained a score in the final plan that was greater than that of the clinically delivered plan with the exception of case 3 for trainee 5 and there was an overall average of 12.6 raw planning score point improvement over the respective clinically delivered plans. Detailed breakdown of each trainee's performance on each training case is listed in **Table 3**.
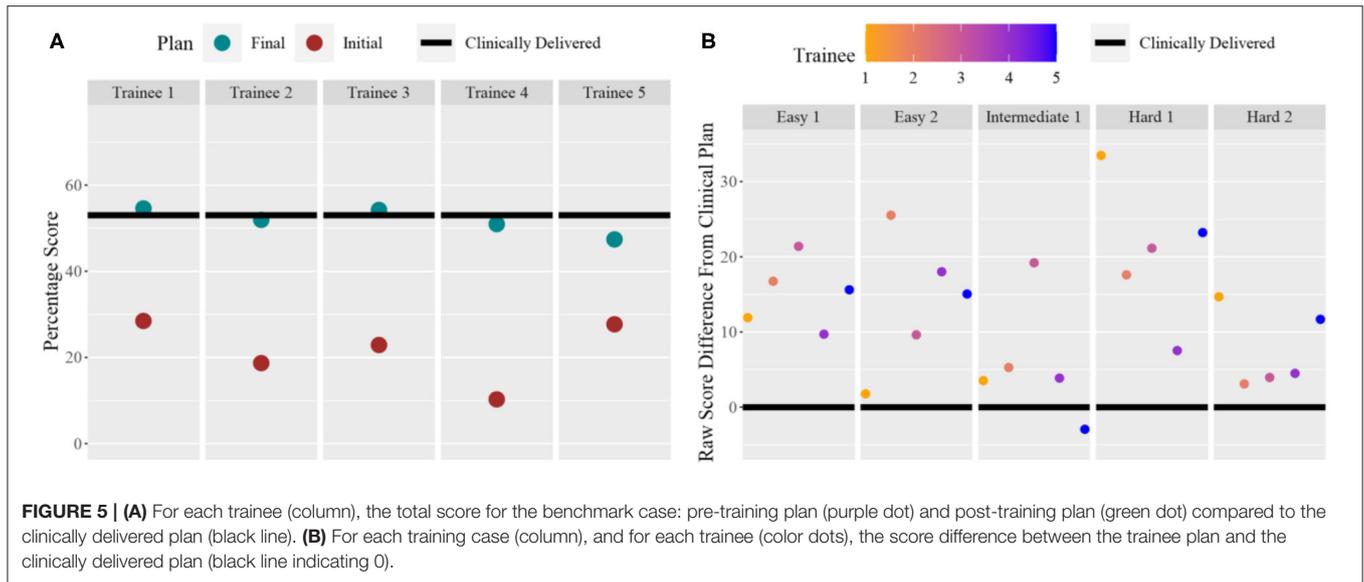
Physician Evaluation Results

Table 4 shows the physician evaluation of the trainee's post-training plans as compared to the benchmark plans and the clinically delivered plans for the benchmark case. Two plans designed by trainees #1 and #3 that scored slightly better than the clinical plan per the scoring system were deemed as comparable to the clinical plan by the physician. The other plans were rated as marginally worse. All trainee's post-training plans were rated moderately better than the initial benchmark plans. Only one of the trainee's plan was deemed appropriate for clinical use based on the physician's discretion.

DISCUSSION

This is the first attempt at developing an effective training program for IMRT planning that capitalizes on the implicit planning tactics that is built into knowledge models for lung IMRT. As trainees go through the training program, the prediction from knowledge models provides guidance at multiple steps and the carefully thought-out scoring objectives direct them toward appropriate choices or skills to create a clinically viable plan. The initial assessment indicates that the knowledge model based training program can substantially improve the planning knowledge of novice trainees in a short period of time (9–12 h in this study). Furthermore, for some trainees their knowledge may approach a clinical proficient level within this short period.

This training program demonstrates the feasibility that knowledge models can be effective teaching aids to help human planners understand the key steps toward generating a clinically viable plan. This is an important first attempt to use knowledge models in a human training process. We hypothesize that by giving trainees opportunities to compare and reflect on the predictions from knowledge models and their own understanding of the planning process, these human planners

FIGURE 5 | (A) For each trainee (column), the total score for the benchmark case: pre-training plan (purple dot) and post-training plan (green dot) compared to the clinically delivered plan (black line). **(B)** For each training case (column), and for each trainee (color dots), the score difference between the trainee plan and the clinically delivered plan (black line indicating 0).

TABLE 2 | Scores of the benchmark plan for each trainee before and ("Initial") after ("Final") training.

Trainee ID	Plan	Raw score	Percentage score
Trainee 1	Initial	46.69	28.47
Trainee 2	Initial	30.64	18.68
Trainee 3	Initial	37.53	22.88
Trainee 4	Initial	16.83	10.26
Trainee 5	Initial	45.42	27.70
Trainee 1	Final	89.54	54.60
Trainee 2	Final	85.18	51.94
Trainee 3	Final	88.96	54.24
Trainee 4	Final	83.56	50.95
Trainee 5	Final	77.74	47.40

will have a better and more concrete understanding of the knowledge models and thus have confidence in making their planning decisions rather than simply accepting the predicted results. While further research is needed to design more effective mechanisms for incorporating knowledge models in human learning, this study has shown that proper design of a plan scoring system provides one effective approach to helping trainees understand the effects of beams and constraints. Further development and testing of the scoring system are warranted since five cases are not likely to cover the possible case variations and review by only one physician may not be sufficient to cover variations in clinical considerations.

While the beam and DVH prediction models used in this study make for a good foundation, additional and more sophisticated knowledge models are needed to address the skills and knowledge that are currently provided by trainers throughout the training to produce clinically viable plans. Examples of important considerations during planning include

collimator optimization and strategies to fine-tune small regions that are less optimal.

In the current implementation, the plan scoring system serves multiple purposes. First, the total score should measure the overall quality of a plan. Second, the less than satisfactory scores should emphasize the most important metrics that require attention. Third, in an indirect way, we want the total score to measure a trainee's mastery of planning knowledge and the difference in scores on the same case to measure the trainee's level of improvement (i.e., learning). Scoring for the first purpose has been studied in quality assurance literature (Mayo et al., 2017). Unfortunately, this scoring will always have an *ad hoc* nature as physicians' preferences will vary, and one scoring system that is in perfect agreement with one physician may not hold true for another. Moreover, some metrics are prioritized conditionally depending on other metrics. One such scoring difficulty is in terms of the metrics that physicians utilize to make decisions based on seemingly minor differences. For example, in some cases, the esophagus may not be prioritized as highly as the lung or the heart, but if the other metrics are at an acceptable level then even small differences in the esophageal metrics may be considered more important than moderate differences in lung dose. This is because most people with locally-advanced lung cancer will experience some degree of esophagitis (Chapet et al., 2005) while a much smaller percentage will experience pneumonitis. This type of conditional prioritization poses significant challenges for scoring system design and require further investigation. Scoring systems for the latter two purposes have not been previously studied. One challenge that we faced is the exploitation of the scoring system by trainees. That is, poorly designed scoring systems tend to allow trainees to attain high scores without actually understanding planning knowledge and actually creating high quality plans. We have improved our scoring system iteratively by adjusting the priority (i.e., max point) assignments based on pilot testing

TABLE 3 | Scores of five plans during training for each of five trainees. Score difference is defined as the difference between trainee's plan's score vs. the clinical plan's score.

Trainee ID	Raw score	Score difference	Plan
1	129.02	11.9	Easy 1
2	133.84	16.74	Easy 1
3	138.49	21.39	Easy 1
4	126.8	9.7	Easy 1
5	132.71	15.61	Easy 1
1	120.51	1.77	Easy 2
2	144.27	25.53	Easy 2
3	128.36	9.62	Easy 2
4	136.75	18.01	Easy 2
5	133.79	15.05	Easy 2
1	54.75	3.53	Intermediate 1
2	56.49	5.27	Intermediate 1
3	70.42	19.2	Intermediate 1
4	55.08	3.86	Intermediate 1
5	48.27	-2.95	Intermediate 1
1	60.66	33.48	Hard 1
2	44.78	17.6	Hard 1
3	48.32	21.14	Hard 1
4	34.7	7.52	Hard 1
5	50.4	23.22	Hard 1
1	83.91	14.68	Hard 2
2	72.32	3.09	Hard 2
3	73.17	3.94	Hard 2
4	73.72	4.49	Hard 2
5	80.91	11.68	Hard 2

TABLE 4 | Physician evaluation of trainee post-training plan on 5-point scale (significantly worse to significantly better) and clinical feasibility rating.

Trainee #	Comparison to clinical	Comparison to benchmark	Clinically feasible
1	Comparable	Moderately better	No
2	Moderately worse	Moderately better	No
3	Comparable	Moderately better	Yes
4	Moderately worse	Moderately better	No
5	Moderately worse	Moderately better	No

results. It is also important not to adjust the scoring priority for every plan or trainee because there will always be new ways to exploit any scoring system. One possible solution to this is to have a progressive scoring system that adjusts priority when reaching certain thresholds. Another approach is to use entirely separate and different mechanisms for the latter two purposes. For example, instead of using a score to measure a trainee's knowledge, we may use a Bayesian model to assess the probability of the trainee's understanding of a case as is done in modern Intelligent Tutoring Systems (Santhi et al., 2013).

The current training program has many limitations. We can observe one example by comparing the left and right

of **Figure 5**. As seen in the right figure, after training using the knowledge models, all five trainees were able to generate plans that score higher than clinically delivered plans for all five training cases. However, as shown in the left figure, when the trainees returned to the benchmark case, only two trainees were able to achieve near or just at the level of the clinically delivered plans. It can be inferred that some of the trainees might not have fully absorbed the knowledge that was presented to them through the training program. It is also possible that the benchmark case requires special knowledge that is not well-presented to the trainees. In addition, we noticed that during physician plan evaluation, only one of two plans that outscored the benchmark case's clinical plan was deemed clinically acceptable. It is possible that additional plan quality related metric could be introduced in the scoring system to better quantify a plan's clinical applicability. Further research in all aspects of the program, including the knowledge models, the scoring system, the coverage of essential knowledge, and the selection of training cases, is necessary to improve the effectiveness of the training program. Finally, the current implementation is based on a specific commercial TPS platform and its existing application programming interface. While general principles of training workflow design are applicable to other commercial platforms, methods for adapting the proposed design to other planning technologies and platforms deserve further investigation.

Even though the current training program has shown encouraging results that demonstrate its feasibility, there are clearly much to be done to develop a truly effective training program for knowledge-based IMRT planning. The immediate next stage includes the need to enhance the scoring system, extend knowledge models, and expand to a larger study with more training cases and with a variety of sites beyond just the lung. Another important task is to conduct a larger study with more trainees and more physicians to fully evaluate the benefits of the training program centered around knowledge-based models. As discussed in the introduction, our ultimate goal is to develop the training program into a fully asynchronous intelligent tutoring system as we gain a better understanding of the essential components and algorithms that are required by such a training system. Having a human trainer in the current program will provide important feedback for future designs. With permission of trainees, all conversations can be recorded in order to find where and how best to provide certain learning materials and pertinent hints. An intelligent training system operating asynchronously may be invaluable for reducing costs of planner training, providing an educational resource to graduate programs, tearing down the black box mindset of knowledge models in clinical practices, and improving the quality of care in cancer centers across the world (Zubizarreta et al., 2015).

CONCLUSION

We have demonstrated that knowledge models can be effectively used as teaching aid in a training program to bring unexperienced

planners to a level close to experienced planners in a short period of time. The assessments indicate that the knowledge models helped trainees improve their knowledge and skills for producing higher quality plans. We believe this knowledge model based training program can serve as an important component of an AI ecosystem that will enable clinical practitioners to effectively and confidently use KBP in radiation treatment. Further efforts are needed to enhance, validate, and ultimately automate the training program.

DATA AVAILABILITY STATEMENT

The raw data supporting the conclusions of this article will be made available by the authors, without undue reservation.

AUTHOR CONTRIBUTIONS

MM and YS performed system design, trainee tutoring, and treatment plan comparison. CK reviewed plan as physician expert and provided feedback for plan quality. JC, JP, QW, and F-FY provide consultation, review experiment design, paper content, and statistical analysis. YG and QJW supervised the entire study and revised this paper. All authors contributed to the article and approved the submitted version.

FUNDING

REFERENCES

Baker, S., Dahele, M., Lagerwaard, F. J., and Senan, S. (2016). A critical review of recent developments in radiotherapy for non-small cell lung cancer. *Rad. Oncol.* 11:115. doi: 10.1186/s13014-016-0693-8

Breedveld, S., Storchi, P. R. M., Voet, P. W. J., and Heijmen, B. J. M. (2012). iCycle: Integrated, multicriterial beam angle, and profile optimization for generation of coplanar and noncoplanar IMRT plans. *Med. Phys.* 39, 951–963. doi: 10.1118/1.3676689

Chapet, O., Kong, F.-M., Lee, J. S., Hayman, J. A., and Ten Haken, R. K. (2005). Normal tissue complication probability modeling for acute esophagitis in patients treated with conformal radiation therapy for non-small cell lung cancer. *Radiother. Oncol.* 77, 176–181. doi: 10.1016/j.radonc.2005.10.001

Chun, S. G., Hu, C., Choy, H., Komaki, R. U., Timmerman, R. D., Schild, S. E., et al. (2017). Impact of intensity-modulated radiation therapy technique for locally advanced non–small-cell lung cancer: a secondary analysis of the NRG oncology RTOG 0617 randomized clinical trial. *Clin. Trial* 35, 56–62. doi: 10.1200/JCO.2016.69.1378

Dermeval, D., Paiva, R., Bittencourt, I. I., Vassileva, J., and Borges, D. (2018). Authoring tools for designing intelligent tutoring systems: a systematic review of the literature. *Int. J. Artificial Intelligence Edu.* 28, 336–384. doi: 10.1007/s40593-017-0157-9

Good, D., Lo, J., Lee, W. R., Wu, Q. J., Yin, F. F., and Das, S. K. (2013). A knowledge-based approach to improving and homogenizing intensity modulated radiation therapy planning quality among treatment centers: an example application to prostate cancer planning. *Int. J. Rad. Oncol. Biol. Phys.* 87, 176–181. doi: 10.1016/j.ijrobp.2013.03.015

Hazell, I., Bzdusek, K., Kumar, P., Hansen, C. R., Bertelsen, A., Eriksen, J. G., et al. (2016). Automatic planning of head and neck treatment plans. *J. Appl. Clin. Med. Phys.* 17, 272–282. doi: 10.1120/jacmp.v17i1.5901

Knoos, T., Kristensen, I., and Nilsson, P. (1998). Volumetric and dosimetric evaluation of radiation treatment plans: radiation conformity index. *Int. J. Rad. Oncol. Biol. Phys.* 42, 1169–1176. doi: 10.1016/S0360-3016(98)00239-9

Kong, F.-M., Ritter, T., Quint, D. J., Senan, S., Gaspar, L. E., Komaki, R. U., et al. (2011). Consideration of dose limits for organs at risk of thoracic radiotherapy: atlas for lung, proximal bronchial tree, esophagus, spinal cord, ribs, and brachial plexus. *Int. J. Rad. Oncol. Biol. Phys.* 81, 1442–1457. doi: 10.1016/j.ijrobp.2010.07.1977

Mayo, C. S., Yao, J., Eisbruch, A., Balter, J. M., Litzenberg, D. W., Matuszak, M. M., et al. (2017). Incorporating big data into treatment plan evaluation: development of statistical DVH metrics and visualization dashboards. *Adv. Radiat. Oncol.* 2, 503–514. doi: 10.1016/j.adro.2017.04.005

Mitrovic, A., Martin, B., and Suraweera, P. (2007). Intelligent tutors for all: the constraint-based approach. *IEEE Intelligent Syst.* 22, 38–45. doi: 10.1109/MIS.2007.74

Mitrovic, A., Ohlsson, S., and Barrow, D. K. (2013). The effect of positive feedback in a constraint-based intelligent tutoring system. *Comp. Educ.* 60, 264–272. doi: 10.1016/j.compedu.2012.07.002

Nelms, B. E., Robinson, G., Markham, J., Velasco, K., Boyd, S., Narayan, S., et al. (2012). Variation in external beam treatment plan quality: an inter-institutional study of planners and planning systems. *Pract. Rad. Oncol.* 2, 296–305. doi: 10.1016/j.prro.2011.11.012

Paddick, I. (2000). A simple scoring ratio to index the conformity of radiosurgical treatment plans. *J. Neurosurg.* 93:219. doi: 10.3171/jns.2000.93.supplement_3.0219

Santhi, R., Priya, B., Nandhini, J. M. J. A. (2013). Review of intelligent tutoring systems using bayesian approach. *[arXiv preprint].* arXiv:abs/1302.7081.

Sheng, Y., Li, T., Zhang, Y., Lee, W. R., Yin, F. F., Ge, Y., et al. (2015). Atlas-guided prostate intensity modulated radiation therapy (IMRT) planning. *Phys. Med. Biol.* 60, 7277–7291. doi: 10.1088/0031-9155/60/18/7277

Sheng, Y., Zhang, J., Wang, C., Yin, F.-F., Wu, Q. J., and Ge, Y. (2019). Incorporating case-based reasoning for radiation therapy knowledge modeling: a pelvic case study. *Technol. Cancer Res. Treat.* 18:1533033819874788. doi: 10.1177/1533033819874788

Voet, P. W., Dirkx, M. L., Breedveld, S., Fransen, D., Levendag, P. C., and Heijmen, B. J. (2013). Toward fully automated multicriterial plan generation: a prospective clinical study. *Int. J. Rad. Oncol. Biol. Phys.* 85, 866–872. doi: 10.1016/j.ijrobp.2012.04.015

Voet, P. W. J., Dirkx, M. L. P., Breedveld, S., Al-Mamgani, A., Incrocci, L., and Heijmen, B. J. M. (2014). Fully automated volumetric modulated arc therapy plan generation for prostate cancer patients. *Int. J. Rad. Oncol. Biol. Phys.* 88, 1175–1179. doi: 10.1016/j.ijrobp.2013.12.046

Yuan, L., Ge, Y., Lee, W., Yin, F. F., Kirkpatrick, J., and Wu, Q. (2012). Quantitative analysis of the factors which affect the inter-patient organ-at risk dose sparing variation in IMRT plans. *Med. Phys.* 39, 6868–6878. doi: 10.1118/1.4757927

Yuan, L., Wu, Q. J., Yin, F., Li, Y., Sheng, Y., Kelsey, C. R., et al. (2015). Standardized beam bouquets for lung IMRT planning. *Phys. Med. Biol.* 60, 1831–1843. doi: 10.1088/0031-9155/60/5/1831

Yuan, L., Zhu, W., Ge, Y., Jiang, Y., Sheng, Y., Yin, F., et al. (2018). Lung IMRT planning with automatic determination of beam angle configurations. *Phys. Med. Biol.* 63:135024. doi: 10.1088/1361-6560/aac8b4

Zarepisheh, M., Long, T., Li, N., Tian, Z., Romeijn, H. E., Jia, X., et al. (2014). A DVH-guided IMRT optimization algorithm for automatic treatment planning and adaptive radiotherapy replanning. *Med. Phys.* 41:061711. doi: 10.1118/1.4875700

Zhang, J., Ge, Y., Sheng, Y., Yin, F.-F., and Wu, Q. J. (2019a). Modeling of multiple planning target volumes for head and neck treatments in knowledge-based treatment planning. *Med. Phys.* 46, 3812–3822. doi: 10.1002/mp.13679

Zhang, J., Wu, Q. J., Ge, Y., Wang, C., Sheng, Y., Palta, J., et al. (2019b). Knowledge-based statistical inference method for plan quality quantification. *Technol. Cancer Res. Treat.* 18:1533033819857758. doi: 10.1177/1533033819857758

Zhang, J., Wu, Q. J., Xie, T., Sheng, Y., Yin, F.-F., and Ge, Y. (2018). An ensemble approach to knowledge-based intensity-modulated radiation therapy planning. *Front. Oncol.* 8:57. doi: 10.3389/fonc.2018.00057

Zhang, X., Li, X., Quan, E., Pan, X., and Li, Y. (2012). A methodology for automatic intensity-modulated radiation treatment planning for lung cancer. *Phys. Med. Biol.* 56:9. doi: 10.1088/0031-9155/56/13/009

Zhu, X., Ge, Y., Li, T., Thongphiew, D., Yin, F. F., and Wu, Q. (2011). A planning quality evaluation tool for prostate adaptive IMRT based on machine learning. *Med. Phys.* 38, 719–726. doi: 10.1118/1.3539749

Zubizarreta, E. H., Fidarova, E., Healy, B., and Rosenblatt, E. (2015). Need for radiotherapy in low and middle income countries – the silent crisis continues. *Clin. Oncol.* 27, 107–114. doi: 10.1016/j.clon.2014.10.006