

Automated Generation of Radiotherapy Treatment Plans Using Machine Learning
Methods
by

Wentao Wang

Graduate Program in Medical Physics
Duke University

Date: _____

Approved:

Q. Jackie Wu, Advisor

Qiuwen Wu

Yaorong Ge

Manisha Palta

Yang Sheng

Dissertation submitted in partial fulfillment of
the requirements for the degree of Doctor of Philosophy
in the Graduate Program in Medical Physics
in the Graduate School of Duke University

2021

ABSTRACT

Automated Generation of Radiotherapy Treatment Plans Using Machine Learning
Methods
by

Wentao Wang

Graduate Program in Medical Physics
Duke University

Date: _____

Approved:

Q. Jackie Wu, Advisor

Qiuwen Wu

Yaorong Ge

Manisha Palta

Yang Sheng

An abstract of a dissertation submitted in partial fulfillment
of the requirements for the degree of Doctor of Philosophy
in the Graduate Program in Medical Physics
in the Graduate School of Duke University

2021

Copyright by
Wentao Wang
2021

Abstract

With the development of medical linear accelerator technologies, the precision and complexity of external beam radiation therapy have increased tremendously over the years. The goal of radiation therapy has always been to push the limit to irradiate the target volume while preserving normal tissues. To achieve this goal, treatment planning for radiation therapy has become a labor-intensive and time-consuming task, which requires a high level of experience and knowledge from the planner. Therefore, automated treatment planning, or auto-planning, is of particular interest in radiation therapy research. The advantages of auto-planning are reduced planning time and increased plan quality consistency.

Since the treatment planning workflow has multiple steps, auto-planning includes the automation of different planning procedures, such as contouring, beam placement, and inverse optimization, which can be achieved in different approaches. The main approaches are knowledge-based planning, automated rule implementation and reasoning, and multicriteria optimization. We can generally consider such novel auto-planning applications as artificial intelligence (AI). This study primarily focuses on treatment plan generation using knowledge-based planning and machine learning techniques. The study includes two main projects: automated beam setting for whole

breast radiation therapy (WBRT) and fluence map prediction for intensity modulated radiation therapy (IMRT).

In WBRT planning, tangential beams are used to irradiate the entire breast volume and avoid the organs-at-risk (OARs) (i.e., the lungs and the heart) as much as possible. The placement of the beams is vital in determining the planning target volume (PTV) coverage and normal tissue sparing. Furthermore, planners need to take multiple clinical considerations into account, e.g., avoiding the contralateral breast and the heart, and use a variety of techniques to meet the demands.

Therefore, we developed an automated beam setting program which takes simple user settings and optimizes target coverage and OAR sparing. The program can be launched from the Eclipse Treatment Planning System (TPS) as a binary plug-in script, which generates a graphical user interface to accept user inputs. Several beam geometries are supported: tangential beams only (supine), tangential plus supraclavicular (SCV) beams (supine), and prone beams.

For all geometries, the program calculates the optimal gantry angles, collimator angles, isocenter location, jaw sizes, and MLC shapes. The borders of the SCV beams are also matched to the tangential beams by using couch kicks on the main tangential fields. For the supine geometries, a coefficient was learned from existing clinical plans to balance between the PTV and lung coverages. The program searches from an initial setting based on breast wires and finds the optimal setting. For the prone geometry, the

planner can set a margin to customize the coverage near the PTV-lung interface. The program has been implemented together with a WBRT fluence prediction program, which creates electronic compensation (ECOMP) plans from the given beam settings. This automated workflow can significantly reduce the workload of the forward planned ECOMP plans. The results showed that the AI plans achieved similar or better plan quality compared to the manual plans.

In IMRT planning, inverse optimization is the standard practice to create treatment plans. Dose-volume histogram (DVH) constraints and priorities are set by the planner to start the optimization and often continuously tuned throughout the planning process until the optimal dose distribution is achieved. The actual parameters to be optimized are fluence map intensities of the IMRT beams. Numerous efforts have been devoted in KBP to predict either the DVH or the dose of the optimal plan. The rationale is that, given the patient anatomy and the physician's prescription, the DVH or dose in the final plan can be predicted based on similar previous plans. The predicted DVH or dose can then be used as a reference to either evaluate the plan quality or generate new plans by converting them into inverse optimization objectives, which is a process also known as dose mimicking. However, most dose mimicking techniques are still in the development stage and not yet commercially available. We explored the feasibility to directly predict optimal fluence maps and generate IMRT plans without inverse optimization.

In order to achieve fluence map prediction, we first investigated the correlation between patient anatomy and fluence maps. A database of patient anatomy and fluence maps was built with pancreas SBRT cases. Treatment planning was done on 2D axial slices with in-house dose calculation and fluence optimization algorithms. For a new slice, an atlas matching method was developed to search for the most anatomically similar slice in the database and initialize the optimization with the existing fluence. The atlas-guided fluence optimization reduced the optimization cost and offered a small dosimetric improvement compared to uniform initialization.

With more training data, deep learning methods were experimented to predict fluence maps from patient anatomy. A deep learning framework consisting of two convolutional neural networks (CNN) was developed. As each plan has several beams, all beam doses must add up to the optimal plan's total dose, while each beam dose is deposited only by said beam's fluence map. Therefore, the BD-CNN predicts the individual beam doses (BD) for an IMRT plan, which tries to minimize the prediction error for both the beam doses and the total dose. Once the beam doses are available, each fluence map (FM) is generated separately by the FM-CNN. As the fluence maps exist in the beam's eye view (BEV), a projection of the 3D beam dose onto the 2D BEV is necessary. The resulting dose map is used as the input to the FM-CNN, which predicts the fluence map as the output. The predicted fluence maps are imported into the TPS for leaf sequencing and dose calculation, generating a deliverable plan.

These projects are retrospective studies using anonymized patient data for training and testing. The development of the deep learning framework was split into several stages: the initial test of the feasibility was conducted for pancreas stereotactic body radiation therapy (SBRT) with a single PTV, unified dose constraints, and a fixed 9-beam geometry; the networks were then modified to allow variable dose inputs and multiple PTVs for pancreas SBRT with simultaneous integrated boost (SIB); a transfer learning technique was applied to the training of the framework for adrenal SBRT plans with different beam settings and dose constraints, using the pancreas model as the base model. The framework has evolved to be more robust and support different sites and planning styles over time.

The AI plans with predicted fluence maps achieved similar plan quality as manual plans for most cases. For some cases with particularly challenging patient anatomies, the AI plans can struggle to reach the high standard of the expert plans. Fluence map prediction is a viable way to directly generate IMRT plans without inverse optimization. This application may be especially useful for adaptive treatment planning.

Contents

Abstract	iv
List of Tables	xiv
List of Figures	xv
List of Abbreviations	xviii
Acknowledgements	xxi
1. Introduction	1
1.1 Radiotherapy treatment planning	1
1.2 Existing automated treatment planning techniques	4
1.2.1 Knowledge-based planning	5
1.2.1.1 DVH prediction	6
1.2.1.2 Dose prediction	8
1.2.1.3 Plan generation via dose mimicking	11
1.2.2 Automated rule implementation and reasoning	12
1.2.3 Multicriteria optimization	13
1.3 Machine learning and radiotherapy	14
1.3.1 Traditional machine learning in radiotherapy	16
1.3.2 Deep learning in radiotherapy	18
1.4 Clinical significance of auto-planning	20
2. Study focus and the overall structure	22
3. Development of an automated beam setting program for whole breast radiotherapy	24

3.1 Introduction.....	24
3.1.1 Breast cancer and treatment.....	24
3.1.2 WBRT beam settings	25
3.1.3 WBRT fluence modulation.....	28
3.1.4 Automated WBRT planning	30
3.2 Materials and methods	31
3.2.1 Patient population	35
3.2.2 Initial beam calculation.....	36
3.2.3 Beam optimization	37
3.2.4 Plan evaluation	39
3.2.5 Prone beam setup	41
3.3 Results	42
3.3.1 Calculation time.....	42
3.3.2 Dosimetric evaluation.....	43
3.3.3 PTV and lung trade-off.....	43
3.3.4 Example cases	45
3.4 Discussion.....	49
3.5 Conclusion.....	53
4. Atlas-guided fluence initialization and optimization.....	55
4.1 Introduction.....	55
4.2 Materials and methods	57
4.2.1 Atlas database building	57

4.2.2 Atlas selection	58
4.3 Results	60
4.4 Discussion and conclusion	62
5. Deep learning-based fluence map prediction framework	63
5.1 Introduction.....	63
5.1.1 The limitations of dose prediction	63
5.1.2 Existing fluence map prediction methods	64
5.1.3 Beam dose and fluence map	65
5.2 Materials and methods	67
5.2.1 Patient selection and radiation therapy plan.....	67
5.2.2 Study workflow	69
5.2.3 Data preprocessing.....	70
5.2.4 Beam dose prediction.....	70
5.2.5 Fluence map prediction	73
5.2.6 BEV projection	75
5.2.7 Model assessment.....	76
5.3 Results	78
5.3.1 Model training	78
5.3.2 Model performance evaluation	80
5.4 Discussion.....	84
5.5 Conclusion.....	88
6. Fluence map prediction for pancreas SIB-SBRT	89

6.1 Introduction.....	89
6.2 Materials and methods	90
6.2.1 Materials	90
6.2.2 Deep learning framework	92
6.2.3 Individual beam dose prediction	94
6.2.4 Fluence map prediction	96
6.2.5 DL framework evaluation	98
6.2.6 Plan evaluation by physician.....	99
6.3 Results	101
6.3.1 Model training and prediction	101
6.3.2 Evaluation of predicted fluence maps.....	102
6.3.3 Evaluation of AI plan quality	103
6.3.4 Example cases	105
6.4 Discussion.....	110
6.5 Conclusions	113
7. Application of transfer learning in fluence map prediction for adrenal SBRT	114
7.1 Introduction.....	114
7.2 Materials and methods	116
7.2.1 Patient data.....	116
7.2.1.1 Pancreas SBRT patients	116
7.2.1.2 Adrenal SBRT patients	117
7.2.2 Beam geometries.....	118

7.2.2.1 Pancreas beam geometry	119
7.2.2.2 Adrenal beam geometry	119
7.2.3 Deep learning framework	120
7.2.4 Model training and testing.....	123
7.2.5 Evaluation of model performance	126
7.3 Results	129
7.3.1 Model training and prediction times.....	129
7.3.2 Plan quality scores in validation and testing.....	129
7.3.3 Example adrenal SBRT cases	131
7.4 Discussion.....	134
7.5 Conclusions	137
8. Conclusions.....	138
References	142
Biography.....	152

List of Tables

Table 1: Dose metrics comparison between AI and clinical plans.	43
Table 2: Model training and calculation details for the DL framework.	79
Table 3: Dose difference between all predicted plan groups and benchmark plans.	80
Table 4: Clinical protocol used for generating the benchmark plans.	91
Table 5: The patient statistics in the training set and the testing set.	91
Table 6: Plan statistics for adrenal SBRT cases.	117
Table 7: The plan quality scores in validation and testing.	129

List of Figures

Figure 1: The IMRT/VMAT manual treatment planning process.	3
Figure 2: Radiation therapy auto-planning, artificial intelligence, machine learning, and deep learning.	15
Figure 3: WBRT beam settings (supine position).	26
Figure 4: WBRT beam settings (prone position).	28
Figure 5: The WBRT study workflow.	33
Figure 6: The graphical user interface for the breast beam setting program.	34
Figure 7: PTV coverage (PTV_Eval $V_{95\%}$) and lung sparing (ipsilateral lung V_{20Gy}) comparison between AI plans (Auto) and clinical plans (Clinic).	45
Figure 8: Dose distribution comparison in the same axial slice of one example case between (A) the AI plan and (B) the clinical plan.	46
Figure 9: Dose distribution in the same axial slice of one example case in the AI plan.	47
Figure 10: Prone setup and dose example 1.	48
Figure 11: Prone setup and dose example 2.	49
Figure 12: The atlas-guided fluence optimization study workflow.	57
Figure 13: Illustration of two "dents" in the PTV contour.	59
Figure 14: The box plot for the cost function reduction for atlas-guided fluence initialization vs iteration number.	61
Figure 15: The averaged OAR DVH (left) and PTV DVH (right) comparisons.	61
Figure 16: An example of total dose (top) and three of the beam doses (bottom) in a 9-beam IMRT plan.	66
Figure 17: The plan generation workflow of the proposed deep learning framework.	69
Figure 18: The network architecture of the BD-CNN in the pancreas DL framework.	71

Figure 19: The network architecture of the FM-CNN in the pancreas DL framework.....	73
Figure 20: The number of training cases vs. testing loss for both CNNs.....	80
Figure 21: Examples of AI fluence map and dose comparison with benchmark in one test case.....	82
Figure 22: The example of PTV (solid) and OAR (dashed) DVH comparison.....	83
Figure 23: Test set distributions of PTV (left) and OAR (right) dose metrics comparing benchmark plans and AI plans.....	84
Figure 24: The proposed DL framework for fluence map prediction compared with manual planning.....	93
Figure 25: The network architecture of the BD-CNN in the DL framework for pancreas SIB-SBRT.....	94
Figure 26: The network architecture of the FM-CNN in the DL framework for pancreas SIB-SBRT.....	97
Figure 27: Fluence map comparisons.....	103
Figure 28: The probability density plots for dose metric deviations of AI plans from benchmark plans.....	104
Figure 29: One test case received grade A by the physician.....	106
Figure 30: One test case received grade B by the physician.....	107
Figure 31: One test case received grade B after renormalization (original grade C).....	108
Figure 32: One test case received grade D by the physician due to inadequate PTV33 coverage.....	109
Figure 33: Beam configurations for pancreas and adrenal plans.....	119
Figure 34: The neural network architecture of the BD-CNN in the adrenal DL framework.....	122
Figure 35: The neural network architecture of the FM-CNN in the adrenal DL framework.....	123

Figure 36: The fluence prediction workflow for the DL framework and model training and testing process.....	124
Figure 37: The pancreas and adrenal cases used for model training, validation, and testing.	126
Figure 38: Plan quality score design.....	128
Figure 39: Dose comparison for adrenal SBRT case 1.....	132
Figure 40: Dose comparison for adrenal SBRT case 2.....	134

List of Abbreviations

AI	Artificial intelligence
ANN	Artificial neural network
BCS	Breast-conserving surgery
BCT	Breast-conserving therapy
BD	Beam dose
BEV	Beam's eye view
CNN	Convolutional neural network
CT	Computed tomography
DL	Deep learning
DTH	Distance-to-target histogram
DVH	Dose-volume histogram
ECOMP	Electronic compensation (a delivery technique)
ESAPI	Eclipse Scripting Application Programming Interface
FM	Fluence map
GAN	Generative adversarial network
GI	Gastrointestinal
GUI	Graphical user interface
HU	Hounsfield Unit
IMN	Internal mammary node

IMRT	Intensity-modulated radiation therapy
KBP	Knowledge-based planning
MAE	Mean absolute error
MCO	Multicriteria optimization
ML	Machine learning
MLC	Multi-leaf collimator
MU	Monitor unit
OAR	Organ-at-risk
PCA	Principal component analysis
PTV	Planning target volume
RF	Random forest
RL	Reinforcement learning
ROI	Region of interest
RT	Radiotherapy/radiation therapy
SBRT	Stereotactic body radiation therapy
SCV	Supraclavicular
SIB	Simultaneous integrated boost
SSD	Source to skin distance
TL	Transfer learning
TPS	Treatment planning system

VMAT	Volumetric modulated arc therapy
WBRT	Whole breast radiation therapy
3D CRT	Three-dimensional conformal radiation therapy

Acknowledgements

Firstly, I would like to thank my PhD dissertation committee. I would like to express my deepest appreciation to my advisor, Dr. Jackie Wu, for your guidance and unparalleled support throughout my graduate study. Your knowledge and vision in radiation therapy steered me through these exciting research projects and shaped me into an independent researcher. I am extremely grateful to Dr. Qiuwen Wu, who served as my committee chair, for your insightful comments on my research and teachings on the fundamentals of radiation therapy. These projects would not have been possible without Dr. Yang Sheng, who guided me through every step of the way. Thank you for setting up such a high standard as a role model for me to work towards. I would like to thank Dr. Yaorong Ge, for your expertise in machine learning, shaping the methodology in my projects, and helping with publishing these works. I very much appreciate Dr. Manisha Palta for providing the knowledge and insightful suggestions from a radiation oncologist's perspective, which greatly enhanced the clinical significance of these works.

I owe my achievements to the Department of Radiation Oncology and Medical Physics Graduate Programs at Duke University and Duke Kunshan University. I would like to express my sincere gratitude to Dr. Fang-Fang Yin for overseeing the research projects and providing valuable comments and support. I would like to thank Dr. Sua Yoo and Dr. Rachel Blitzblau for their knowledge and guidance on the breast project. I am also grateful to Dr. David Huang, Dr. Jim Bowsher, and Claire Luo for their support

during my master's study at DKU. I also had the great pleasure of working with Dr. Mark Oldham, Dr. Justus Adamson, Dr. Oana Craciunescu, and Dr. Sheridan Meltsner on the radiation therapy and brachytherapy courses. Many thanks to Dr. Anuj Kapadia, Dr. Olga Baranova, and Katherine Hand, for helping me navigate through my graduate study at Duke. I will forever cherish the friendships and fun memories from all these years.

I would like to say a big thank you to my colleagues/friends in our lab: Dr. Chunhao Wang, for the helpful discussions on research direction and career development; Dr. Jiahan Zhang, for helping me with scripting and patient data collection and sharing your knowledge in KBP; Xuejun Sun, Phillip Antoine, and Christopher Buarotti, for the IT support and troubleshooting for the Eclipse system; Yushi Chang, James Jensen, Xinyi Li, Hunter Stephens, Ruilin Li, and Jiang Zhang, for the helpful discussions and collaborations on research projects.

Finally, I would like to thank my mom and dad, who have always been unconditionally supportive of me in anything I choose to do. To my girlfriend Xinyue, thank you for being my motivation and cheering me up across the oceans during these difficult times.

The research projects in this study were partially supported by an NIH grant and a master research grant from Varian Medical Systems.

1. Introduction

1.1 *Radiotherapy treatment planning*

Radiotherapy (RT) treatment delivery techniques have evolved rapidly over the years. From the conventional 2D radiotherapy, where targets are delineated on orthogonal x-ray images, to image-based 3D conformal radiation therapy (3D CRT), multi-leaf collimator (MLC)-based intensity modulated radiation therapy (IMRT) and volumetric modulated arc therapy (VMAT), the delivery of radiation dose to the target has become more and more accurate and conformal to the target. The motivation of such technological advancements lies in the clinical objectives of maximizing tumor control probability (TCP) and minimizing normal tissue complication probability (NTCP) at the same time. An ideal radiotherapy plan delivers the full prescribed dose to the planning target volume (PTV) uniformly and no dose to the surrounding organs-at-risk (OARs). However, due to the nature of radiation beams, either photon, electron, or proton, a compromise needs to be made between PTV coverage and OAR sparing. Treatment planning is the process of designing such plans to achieve the ideal compromise for the specific patient.

As the complexity of treatment plans grows, the task of treatment planning has become more labor and resource intensive. Although there are many treatment modalities in radiotherapy, the main focus of this dissertation is IMRT, which is widely used in today's clinics. In IMRT, a nonuniform fluence is delivered to the patient from

any given position of the treatment beam to optimize the composite dose distribution (Khan & Gibbons, 2014, pp. 430-435), which is a non-convex optimization problem. In VMAT, the radiation beam is on as the gantry continuously rotates around the isocenter, i.e., arc therapy, with the MLC leaf positions and dose rates being optimized simultaneously for all gantry angles in the arcs (Khan & Gibbons, 2014, pp. 438-439). With a large number of gantry angles and the physical constraints of the MLC movement, the VMAT optimization process is more complicated and time-consuming than IMRT, but it provides higher delivery efficiency (Khan & Gibbons, 2014, p. 439).

Figure 1 shows the IMRT/VMAT manual planning process. The current treatment planning workflow includes the delineation of targets and organs in anatomical images, designing the radiation beam geometry, setting dose-volume constraints and priorities based on physician's prescriptions, running the inverse optimization program, and iteratively finetuning the dose-volume constraints until a satisfactory dose distribution is achieved. The human planner determines the settings and inputs for the inverse optimization engine based on the patient prescription, institutional guidelines, and previous planning experiences.




Figure 1: The IMRT/VMAT manual treatment planning process.

Within the inverse optimization engine, three main elements are involved: fluence map, dose distribution, and dose-volume histogram (DVH). As indicated by its name, the intensity of the beam fluence map is modulated in IMRT plans. The modulation is achieved by a series of MLC control points, where each control point contains the beam aperture shape and the amount of monitor units delivered. A beam's fluence map is the superposition of the 2D fluence intensities from all MLC control points. With the fluence map of one beam, the dose from this beam can be calculated in the treatment planning system (TPS) using the beam data. This dose calculation process typically uses fast but less accurate algorithms. IMRT plans use multiple (e.g., 7, 9, or more) beams with different gantry angles to achieve the desired dose distribution in the target and avoid high skin dose. Therefore, the total plan dose is the summation of all calculated beam doses. As the optimization objective is typically given as a series of

dose-volume constraints, the dose-volume histogram is calculated from the total plan dose in order to obtain the objective function value.

In each iteration, the fluence intensities in all beams are updated to reduce the objective function value until a local minimum is reached. Since the objective function is non-convex, there is no guarantee that the global minimum can be obtained without an exhaustive search, which is computationally intractable given the large degree of freedom with all beams' fluence maps. In modern treatment planning systems, the planner can view the approximate dose distribution and interactively modify the dose objectives during optimization. When the optimization converges, the fluence maps are considered optimal and used to calculate a more accurate dose distribution for the planner to review. If the plan dose is acceptable to the planner and the physician, the plan is approved and ready to be delivered. Otherwise, manual finetuning of the dose constraints and/or beam settings is needed to improve the dose distribution. In this process, multiple runs of inverse optimization are usually needed to achieve the ideal dose distribution, which can be time-consuming and labor-intensive. Furthermore, it requires extensive treatment planning knowledge and experience from the planner. This creates the inconsistency in plan quality among different planners.

1.2 Existing automated treatment planning techniques

To reduce the treatment planning time and increase the plan quality consistency, numerous efforts have been made to automate different parts of the treatment planning

workflow. Automated treatment planning, or auto-planning for short, has become a prominent topic in radiation oncology research in recent years. Auto-planning can be implemented in several ways, including but not limited to: knowledge-based planning (KBP), automated rule implementation and reasoning, and multicriteria optimization (MCO) (C. Wang, Zhu, Hong, & Zheng, 2019). The main focus of this dissertation is knowledge-based planning using machine learning, which will be described in detail. Automated rule implementation and reasoning is featured in the breast project. Multicriteria optimization will be briefly summarized to provide context and comparison with the KBP approach.

1.2.1 Knowledge-based planning

The central aim of knowledge-based planning is to quantify, summarize, and utilize the intangible “knowledge” of treatment planning from existing plans to assist in the treatment planning process for new patients. This knowledge encompasses many aspects of treatment planning, such as the selection of modality/energy, the design of beam geometry, and the ideal dose distribution. It is mainly used for more complex treatment modalities, such as IMRT and VMAT, since they often require a high level of treatment planning skill. Individual planners accumulate their treatment planning knowledge through years of training and reviewing existing plans. Knowledge-based planning fits in the clinical workflow by providing the planners with an easy access to a condensed version of treatment planning knowledge. It can reduce the need for some

time-consuming tasks, such as manual finetuning of optimization objectives, and improve the plan quality consistency among different planners. As KBP models do not typically include sensitive patient information, except for database plan matching techniques, it is more feasible for different institutions to share their planning knowledge and practice through these models. A new or smaller hospital can implement and improve on well-established KBP models without a large number of existing plans as training data.

1.2.1.1 DVH prediction

Knowledge-based planning can be implemented in several ways. The early KBP models focused on the prediction of DVH objectives. As DVH is the one-dimensional representation of the three-dimensional dose distribution, it is a summary of the dose in the form of a curve and is commonly used to evaluate plan quality in the clinics.

Cumulative integral dose volume histograms are more commonly used to represent the dose to individual structures, which plots the relative/absolute volume receiving a certain dose or higher (y axis) as a function of dose (x axis). The less used differential DVH plots the volume receiving a dose within a dose interval. In the following, DVH refers to cumulative DVH unless otherwise specified.

In IMRT optimization, dose constraints are often given in the form of points or lines in the DVH plot. Therefore, if the optimal DVHs can be predicted prior to obtaining the actual plan, they can facilitate in the inverse optimization process (Yuan et

al., 2012). The predicted DVHs can also be compared to the actual DVHs as a plan quality evaluation tool to identify suboptimal plans (Zhu et al., 2011).

The OAR DVHs can be used to evaluate normal tissue toxicity quantitatively. For example, in whole breast planning, the ipsilateral lung V_{20Gy} is a commonly used lung dose indicator to be minimized. OAR dose can be related to a variety of factors, but the most direct one is the spatial relationship between the OAR and the PTV. Similar to the DVH concept, the distance-to-target histogram (DTH) describes the distance between OAR voxels to the PTV with a curve. Zhu et al. (2011) used the OAR volumes and DTH to characterize the OAR anatomical information. Principal component analysis (PCA) was used to reduce the dimensions of DTH and DVH. Support vector regression (SVR) was trained to correlate the anatomical feature vectors with the DVH feature vectors. Yuan et al. (2012) further identified some patient anatomical factors contributing to OAR dose sparing, which included the median distance between OAR and PTV, the portion of OAR volume within an OAR specific distance range, the fraction of OAR volume which overlaps with PTV, and the portion of OAR volume outside the primary treatment field.

Once the DVHs are predicted, they could be used in the inverse optimization process to create KBP plans. In 2014, the model developed by Yuan et al. (2012) has been modified and adopted into the Eclipse Treatment Planning System (Varian Medical Systems, Palo Alto, CA) as a commercially available KBP module known as

RapidPlan™. RapidPlan can build a DVH prediction model given a set of existing clinical plans for one treatment site. The predicted OAR DVHs are then used to generate dose-volume constraints for inverse optimization, where the user can configure some settings include the number, type, and priorities of objectives. The system has been tested in a number of treatment sites, such as pelvis (Hussein et al., 2016), head and neck (Fogliata et al., 2017; Tol, Delaney, Dahele, Slotman, & Verbakel, 2015), rectum (Wu, Jiang, Yue, Li, & Zhang, 2016), lung (Snyder et al., 2016), and spine (Foy et al., 2017).

1.2.1.2 Dose prediction

Since the DVH is the 1D representation of the 3D dose, some spatial information of the dose is lost in the process. The same DVH curve could be derived from two different dose distributions, which may have clinically relevant differences. As physicians use both the DVH and dose to evaluate plan quality, dose prediction could provide some advantages over DVH prediction alone. Furthermore, the predicted dose could be used to create deliverable plans through dose mimicking algorithms.

Since dose is essentially a 3D matrix, the prediction of dose can be achieved voxel by voxel, slice by slice, or in 3D volumes. Following the success of DVH prediction, voxel-wise dose prediction methods use the anatomical information of a voxel and its surrounding area to estimate the dose for that single voxel, which is repeated for all voxels to generate the volumetric dose. The input and output are limited to a small area and have low dimensions. Therefore, traditional machine learning

techniques have been used to predict voxel-wise dose. The anatomical features are hand-crafted to suit the specific disease site and treatment modality. The success of such models often depends on the data preprocessing step. For example, Liu et al. (2015) developed a voxel dose prediction model for spine stereotactic body radiation therapy (SBRT) through stepwise multiple regression. The PTV and OAR contours were characterized by an active shape model, while the dose was described by an active optical flow model. Artificial neural networks (ANN) can also be used to establish the correlation between patient anatomical features and dose. Shiraishi and Moore (2016) developed an ANN to predict achievable 3D dose distributions based on patient-specific geometric and planning parameters.

Instead of feature engineering, convolutional neural networks (CNN) are well suited for directly processing high-dimensional image data, such as CT images and contour masks. As dose has a direct relationship with patient anatomy, CNNs have been used to predict dose distribution from patient anatomy in recent years. Kearney, Chan, Haaf, Descovich, and Solberg (2018) developed a novel CNN architecture called DoseNet, which was a three-dimensional fully convolutional network with CT and organ contours as input and predicted dose as output. DoseNet was trained and tested on prostate SBRT plans using non-coplanar beam delivery geometries. DoseNet, like most other dose prediction networks, draws inspiration from U-Net, which was initially developed for medical image segmentation (Ronneberger, Fischer, & Brox, 2015). For

example, Nguyen, Long, et al. (2019) used a modified seven-level hierarchy U-Net to predict the dose distribution for 7-field prostate IMRT plans. The prediction error for the max and mean dose of major structures were within 5% of the prescription dose, and the average dice coefficient between predicted and true isodose volumes was 0.91. This network was trained on single slices of the patient, thus lacked the patient anatomical information in the superior-inferior direction.

In order to account for the superior-inferior anatomy variation, 3D networks were used to predict dose in 3D volumes. Unlike 2D networks which use individual slices as training data, 3D networks use volumetric image patches sampled from the patient volume as their base operating space. Theoretically, 3D networks can produce more realistic dose falloff at the superior and inferior borders for coplanar beams and are more suited for non-coplanar beam settings. The U-Net architecture, with an additional dimension, remains a popular choice for 3D dose prediction model. However, compared to 2D networks, 3D networks typically have much a higher demand on computer memory and computational power. Data augmentation techniques are usually needed to increase the training data size given the same number of patients. Ana María Barragán - Montero et al. (2019) designed a 3D hierarchically densely connected U-Net with the input image size of $96 \times 96 \times 64$. In addition to 9 input channels for anatomical information, the model had one channel for beam setup information in order to account

for beam geometry variations. The beam setup information was provided by the approximate total dose from all beams with unmodulated fluence maps.

1.2.1.3 Plan generation via dose mimicking

One of the major drawbacks of dose prediction is that the predicted dose is not associated with a deliverable plan. Therefore, it is unknown whether the predicted dose can be achieved or even approximated by a plan. Dose prediction is essentially a mathematical operation which rarely involves any physics principles such as the inverse square falloff and tissue scatter/attenuation that determine the dose profile of an actual photon beam. The resemblance of the predicted dose with the actual dose comes from the training data and the model's ability to learn what a dose distribution should look like. Assuming the predicted dose by an existing model has a very small deviation of the optimal clinical dose, the next step then becomes how to create a deliverable plan that produces such a dose distribution. This process can be generally referred to as "dose mimicking", which can be achieved in several ways.

Given the predicted dose distribution, an objective function can be constructed for the inverse optimization process. Babier, Boutilier, Sharpe, McNiven, and Chan (2018) proposed an inverse optimization model to estimate objective function weights from clinical DVHs. It was applied by Mahmood, Babier, McNiven, Diamant, and Chan (2018) as the plan generation step after dose prediction via a generative adversarial network (GAN). This process converted the predicted dose back to DVH, thus losing the

spatial dose information. Another approach is to use all predicted dose as the ground truth and set the dose error as the objective function. In the study by Fan et al. (2019), a residual network was designed to predict IMRT dose for head and neck (H&N) patients. After dose prediction, the authors used the voxel-by-voxel dose squared error as objective function to optimize the fluence maps through the matRad software developed by Wieser et al. (2017). However, without assigning voxel importance, the objective function may overlook important features, such as PTV coverage, while trying to match the dose values for places with little clinical significance. Moreover, current commercial treatment planning systems do not yet support voxel-level dose objective functions, which hinders the clinical implementation of such techniques.

1.2.2 Automated rule implementation and reasoning

Some treatment planning decisions are made by the planner based on certain rules with straightforward criteria. For example, in whole breast radiation therapy, the tangential beams are typically half open with the central axis aligned to the breast wires. These procedures are straight forward to an experienced planner, but they may require specific human knowledge and experience to perform a series of actions in the TPS. For example, in the *mdaccAutoPlan* system developed by X. Zhang, Li, Quan, Pan, and Li (2011), the initial selection of beam angles for lung IMRT plans were achieved by a set of rules involving the tumor location in the lung, which was derived from expert experience. To automatically perform certain actions, most commercial treatment

planning systems offer programmable interfaces for user-designed scripting, such as Varian's Eclipse Scripting Application Programming Interface (ESAPI). Through scripting, many parts of the treatment planning process can be automated, which can potentially reduce the treatment planning time and increase the clinical efficiency.

Automated rule implementation and reasoning with scripting can be combined with complex machine learning models or other programs to provide an automated treatment planning workflow with only a few clicks. For example, Kim et al. (2018) developed an one-click solution for automated breast field-in-field planning by combining ESAPI and user-executed programming. Huang et al. (2018) designed a script to find the optimal VMAT jaw settings and evoke a RapidPlan model to generate patient-specific optimization objectives for inverse planning. J. Zhang et al. (2020) implemented a reinforcement learning (RL) planning bot via ESAPI to generate clinically acceptable plans for pancreas SBRT. Each case took an average of 7.3 minutes for the planning bot to generate a deliverable plan from a set of contours.

1.2.3 Multicriteria optimization

In DVH-based inverse optimization, there are usually several evaluation criteria for treatment plans. The objective function, also known as cost function, is defined by the weighted sum of these criteria. The weighting coefficients represent the relative importance of the criteria, which can be adjusted by the planner to achieve the clinically optimal plan for the specific patient. However, finding the optimal weight coefficients is

not a trivial task and requires trial-and-error and planner's experience. Multicriteria optimization (MCO) was proposed to generate many alternative plans, called anchor plans, based on different weighting coefficients. For example, PTV coverage and OAR sparing are usually conflicting objectives, meaning that no single plan exists that simultaneously optimizes both objectives. There exists a number of Pareto optimal plans, which cannot improve on one objective without degrading the other. These plans form a hypersurface in the possible plan space, known as the Pareto surface. The Pareto surface can be approximately interpolated from anchor plans, offering clinicians the freedom to navigate the Pareto surface and choose the optimal plan based on their clinical judgement.

1.3 Machine learning and radiotherapy

Artificial intelligence (AI) generally refers to intelligence demonstrated by machines, as opposed to human intelligence. Computer programs that mimic human behavior such as decision-making and learning can be considered as AI. For radiotherapy, when a new program is developed that can automate some treatment planning tasks previously only done or not even achievable by humans, it can be considered as an AI application. Machine learning (ML) is a part of artificial intelligence, which has seen rapid development in recent years. Machine learning algorithms are able to capture patterns and knowledge from existing data, i.e., the learning process, and make decisions or predictions without the traditional programmed logic. The keys to

developing robust machine learning algorithms are high-quality data and fast computers. In general, more complex models require more training data and computational power. With more knowledge learned, these models are able to make more precise and complex predictions. Within machine learning, a subset of algorithms is known as deep learning (DL), which typically involves building an artificial neural network (ANN), as “deep” refers to the multiple layers used in the network. For clarity, machine learning algorithms that are not deep learning are referred to as *traditional* machine learning. DL generally requires much more training data than traditional ML and has superior capability in fields such as computer vision (CV) and natural language processing (NLP). The relationship between RT auto-planning, AI, ML, and DL is illustrated in Figure 2.




Figure 2: Radiation therapy auto-planning, artificial intelligence, machine learning, and deep learning.

Depending on the learning methodology, machine learning algorithms can be broadly divided into three categories: supervised learning, unsupervised learning, and

reinforcement learning. In supervised learning, each training data sample consists of an input and an output, also known as labeled data, and the model learns the correlation between the input and the output. For a test data sample (previously unseen by the researcher and the model), the model predicts the output from the input. In unsupervised learning, the training data sample is not labeled, and the model is tasked with finding patterns from the dataset, e.g., clustering data samples into different groups. Reinforcement learning trains intelligent agents to take actions in an environment, such as the treatment planning system, to maximize a reward function. Unlike supervised learning, reinforcement learning does not require labeled data for training; instead, it learns by trial and error. A notable example of reinforcement learning is *AlphaGo Zero* developed by Google DeepMind, which learned the game of Go from self-play without any human knowledge (Silver et al., 2017).

1.3.1 Traditional machine learning in radiotherapy

Some of the early efforts in automated radiotherapy treatment planning used traditional ML algorithms. In the study by Yuan et al. (2012), principal component analysis (PCA) was first applied to DVHs and DTHs. PCA is widely used in traditional ML studies for dimensionality reduction. For data with high dimensions, direct input into traditional ML models may lead to overfitting, meaning that irrelevant information from the training data, i.e., noise, is learned by the model so that it does not generalize well for unseen data. PCA transforms the data to a new coordinate system, with the

greatest variance on the first coordinate, which is called the first principal component, and the second greatest variance on the second coordinate, and so on. Yuan et al. (2012) selected the first three components as anatomical features to correlate with dose volume features via stepwise multiple regression. Stepwise regression chooses the predictive variables by an automatic procedure, where the variables are added if they improve the model fit and eliminated if they do not. Multiple regression indicates that the higher order terms of the variables are included in the regression model.

The selection of traditional ML algorithms depends on the specific task at hand, e.g., regression and classification. To estimate fluence intensities for tangential beams in whole breast radiation therapy, Sheng et al. (2019) utilized a random forest model to correlate shape-based input features with pixel-wise fluence intensities. Other commonly used algorithms include support vector machine (Zhao et al., 2012), least absolute shrinkage and selection operator (LASSO) (T.-F. Lee et al., 2014), decision tree (Surucu et al., 2016), and k-medoids clustering (Yuan et al., 2015). Generally speaking, model selection and finetuning is achieved by validation, where the model is tested on data held out from the training process, providing a somewhat unbiased model performance evaluation. Validation could also facilitate in the selection of useful features, which remains a key issue for traditional ML application in radiotherapy. In comparison, feature selection/engineering is less of an issue for deep learning algorithms, which automatically learn useful representations of the data.

1.3.2 Deep learning in radiotherapy

DL applications in different fields of computer science have gained traction over the last decade. Similar tasks in radiotherapy have quickly adopted the powerful tool of DL. For image prediction/generation, two main approaches are widely used: convolutional neural network and generative adversarial network. U-Net, a fully convolutional network, has been developed for organ segmentation (Ronneberger et al., 2015) and adopted for dose prediction from patient anatomy. The input and output of a U-Net are images with a high resolution. The model architecture consists of several resolution levels connected by downsampling layers and upsampling layers. At the same resolution level, layers are connected by regular convolutions with an activation function such as the rectified linear unit (ReLU). On the left side of the U-Net (contracting path), the image size decreases, and the channel number increases. Local features are captured by the convolutional layers in the high dimensional images. After each downsampling operation, which can be max pooling or strided convolution, the feature map resolution is halved, and the convolutional layers with the same kernel size can capture image information on a larger scale, i.e., global image features are captured. On the right side of the U-Net (expansive path), the image size increases, and the channel number decreases. The upsampling operation is transposed convolution, also known as up-convolution, which doubles the feature map resolution. Between the two sides, feature maps are copied and concatenated from the left side to the right side, also

known as skip connections, which preserves the local features in the reconstruction path. For dose prediction, the input layers are typically anatomical contour masks, and the output is the total dose of an ideal plan. Theoretically, the Hounsfield unit (HU) number in the CT images should affect the deposition of dose in the tissue; however, most studies did not incorporate CT images as an input. Since dose deposition involves complex physical interactions, convolutional layers could not efficiently capture these features.

Generative adversarial network, or GAN, on the other hand, is a DL framework that employs two competing neural networks in order to improve the quality of the generated output. Originally, GAN was developed as a generative model for unsupervised learning, e.g., creating images similar to a pool of training images. GAN has proved useful for supervised learning and reinforcement learning. For example, pix2pix, a conditional GAN, is able to transform image labels to real images (Isola, Zhu, Zhou, & Efros, 2017). GAN consists of a generator, which generates a “fake” image from noise or some inputs, and a discriminator, which tries to identify the “fake” images from the “real” images. In terms of dose prediction, the generator of GAN can use anatomical contour masks as input and generate the dose distribution very much like a U-Net. In fact, many GAN applications use the U-Net architecture for their generator. The fundamental difference lies in the loss function: while a U-Net typically uses pixel-wise prediction error as its loss function, GAN tries to increase the error rate of the

discriminator by producing realistic outputs indistinguishable from the true data. Therefore, GANs can potentially create more realistic looking images than U-Nets, which often suffer from blurry images due to the absolute/squared error loss.

Due to the complex network architecture, DL algorithms are commonly looked at as “black boxes”. Only the input and output are comprehensible to human observers, while the functions of the middle layers are largely conjecture. The black box nature of DL algorithms can lead to mistrust of the technology and hinder the future development due to the lack of understanding. The design of network architecture and the finetuning of the hyperparameters depend heavily on empirical evidence rather than well-established theories, as some would suggest DL is more like art than science. Therefore, efforts have been made to create interpretable DL models in the context of auto-planning (J. Zhang et al., 2020). In comparison, traditional ML algorithms can be analyzed and understood to some extent.

1.4 Clinical significance of auto-planning

From hand calculation to computer-based planning, the level of accuracy and automation in radiation therapy has been constantly improving. In the 21st century, auto-planning has provided convenience to clinicians and shifted the paradigm of the treatment planning process. AI-assisted auto-planning has the potential to simplify the clinical workflow by shortening/eliminating certain tasks and the back-and-forth communication between radiation oncology team members. Some tasks, such as 3D dose

prediction, are not feasible by human planners but can be done by an AI agent, which can lead to new areas of research and ways to create treatment plans. As the research matures, commercial products will provide easy access to newly developed technologies. Clinical workflow will adapt to the new technologies and improve efficiency, ultimately leading to better and more accessible patient care.

Auto-planning can facilitate the teaching process for inexperienced planners. For example, Mistro et al. (2020) designed a tutoring system using knowledge models (beam angle prediction and DVH prediction) and a scoring system to train individuals to develop optimal lung IMRT plans. In fewer than 2 days, participants in the study significantly improved their planning skills to a level close to experienced planners. As cross-institution patient data sharing is hindered by institutional restrictions, auto-planning models have the potential to bring the treatment planning knowledge, guidelines, and expertise of large medical institutions to smaller clinics.

2. Study focus and the overall structure

The study focuses on developing automated treatment planning tools to generate radiotherapy plans. At Duke University Medical Center, whole breast radiation therapy (WBRT) is delivered with tangential beams and electronic compensation for modulated fluence. The beam settings for WBRT heavily influence the dose distribution and require user experience to balance between PTV coverage and OAR sparing. With the beam settings, fluence maps are “painted” by the planner in a forward planning fashion, which is very time-consuming. Therefore, an automated workflow to generate the beam settings and fluence maps can greatly improve the treatment planning efficiency.

In the IMRT inverse planning workflow, fluence maps are optimized to minimize DVH-based objective functions. Manual planning often involves iterative finetuning of the dose constraints and priorities to produce clinically optimal plans. While more studies focused on the prediction of dose distribution, research on the direct prediction of fluence maps is still in the initial exploratory phase. Accurate prediction of fluence maps can be directly used in the plan generation without inverse optimization.

The study consists of two main projects. The first project describes the development of an automated beam setting program for whole breast radiation therapy, which is described in Chapter 3. The second project describes the development of a deep learning framework for fluence map prediction. The project is split into several smaller projects, which include building an atlas database for fluence initialization (Chapter 4),

the initial development of the deep learning framework and proof-of-concept testing for pancreas SBRT (Chapter 5), adding multiple PTV prescriptions for pancreas SIB-SBRT (Chapter 6), allowing variable beam geometries for adrenal SBRT with transfer learning (Chapter 7). The entire study is summarized in the last chapter.

3. Development of an automated beam setting program for whole breast radiotherapy

3.1 Introduction

3.1.1 Breast cancer and treatment

Breast cancer is estimated to account for 30% of new cancer cases in women in the United States in 2019, the highest number for any single site (Siegel, Miller, & Jemal, 2019). For women in the U.S., it accounts for an estimated 41,760 deaths (14.6% of the total deaths) in 2019, second only to lung and bronchus cancer (Siegel et al., 2019). The early detection of breast cancer plays an important role in reducing its morbidity and mortality (Oeffinger et al., 2015).

Up to 20% to 30% of all screen-detected breast cancers are ductal carcinoma in situ (DCIS), which is an early, non-invasive form of breast cancer once considered a rare disease (Ernster et al., 2002; Nakhlis & Morrow, 2003). If treated properly, patients with early-stage breast cancer usually have low recurrence rates and high quality of life (Pignol et al., 2008). The standard treatment for early-stage breast cancer is breast-conserving therapy (BCT), which consists of breast conserving surgery (BCS) followed by adjuvant radiation therapy (Pignol et al., 2008). BCS is also known as lumpectomy or partial mastectomy and consists of surgical removal of the primary tumor with negative margins while leaving as much normal tissue as possible. BCT has proven to result in equivalent long-term survival rates as mastectomy, and it is often the treatment of choice for women with relatively small breast cancers (Veronesi et al., 2002). Compared with

mastectomy, in which a woman's entire breast is surgically removed, BCS can obtain a better cosmetic result and improve the patient's satisfaction (Al-Ghazal, Fallowfield, & Blamey, 2000). After the lumpectomy, adjuvant external beam radiation therapy to the whole breast is considered standard of care.

Radiation therapy was firstly used in breast cancer treatment to reduce the locoregional recurrence rates following radical mastectomy (Bentel, 1996). A number of randomized clinical trials and meta-analyses (Bijker et al., 2006; Fisher et al., 2002; Group, 2006; Veronesi et al., 2002) demonstrate that adjuvant whole breast radiation therapy significantly reduces ipsilateral breast tumor recurrence (IBTR) rate for both invasive and non-invasive breast cancers and improves survival in patients with invasive breast cancer. Partial breast radiation therapy (PBRT) may also sometimes be used after BCS.

3.1.2 WBRT beam settings

In the clinic, the initial beams for WBRT are manually set by the dosimetrist and subsequently adjusted and approved by the treating physician. The setup of a beam includes its isocenter, gantry angle, collimator angle, couch angle, field size, and MLC/block shape. Two parallel opposed tangential beams are used to irradiate the whole breast (Figure 3). The gantry angles of the two beams are separated by approximately 180 degrees (small adjustments may be made for some cases), and they are angled to prevent the beams from diverging into the lung. The anterior field borders

are open 2~3 cm distal to the skin, which is also known as the skin flash region, to account for breast movement due to respiration and daily setup variation.




Figure 3: WBRT beam settings (supine position). The two beams are shown in an axial slice of the simulation CT images. The tissues within the field borders (the thin yellow lines) are irradiated. The deep field border determines the irradiated volume (the part of the patient covered by the beams). The orange double arrow marks the skin flash margin.

In the beam's eye view (BEV) of the tangential beams, in order to reduce the irradiated normal tissue volume, collimator rotation and/or MLC blocks can be used to align the deep field border along the chest wall curvature. To achieve opposed tangential beams, the two beams have the opposite collimator rotation. Depending on patient anatomy, more or less adjustment of collimator rotations or MLC may be needed and are not always required.

For patients with nodal involvement, the regional lymph nodes, including the axillary, supraclavicular (SCV), and internal mammary nodes (IMNs), may also be

irradiated. The IMNs and lower axillary nodes are treated with the whole breast or chest wall within the opposed tangential fields or with a separate electron beam, while the SCV and upper axillary lymph nodes are treated with a third field adjacent to the tangential beams (Bentel, 1996). This anterior oblique field is called the supraclavicular field or SCV field. The inferior field border of the SCV field is carefully matched to the superior border of the tangential fields to avoid hot or cold spots. This field-matching problem can be solved by setting the same isocenter for all three fields with half-beam blocking, or by using collimator rotation and couch rotation for the tangential fields to match to a half-beam blocked SCV field.

For some patients with larger breasts, prone setup (Figure 4) has some advantages over supine setup, including eliminating skin folds, better sparing of the lungs and the heart (Formenti, DeWyngaert, Jozsef, & Goldberg, 2012; Formenti et al., 2007), and less respiratory motion (Morrow, Stepaniak, White, Wilson, & Li, 2007). A randomized trial comparing prone and supine setup for hypofractionated large breast irradiation found that prone treatment resulted in overall improved dose coverage and homogeneity and decreased skin toxicity (Mulliez et al., 2013).

However, prone setup is less reproducible than supine setup (Varga et al., 2009). The setup process is more difficult for both the patients and the therapists and requires daily imaging. With the use of prone breast board, patients may be less comfortable and more susceptible to head/neck/shoulder and rib pain. As some parts of the radiation

beams go through the breast board and/or couch, some dosimetric uncertainty is introduced in the prone setup. The combination of the bulky breast board and large patients requires special attention in setting up the beams to avoid collision with the gantry head.

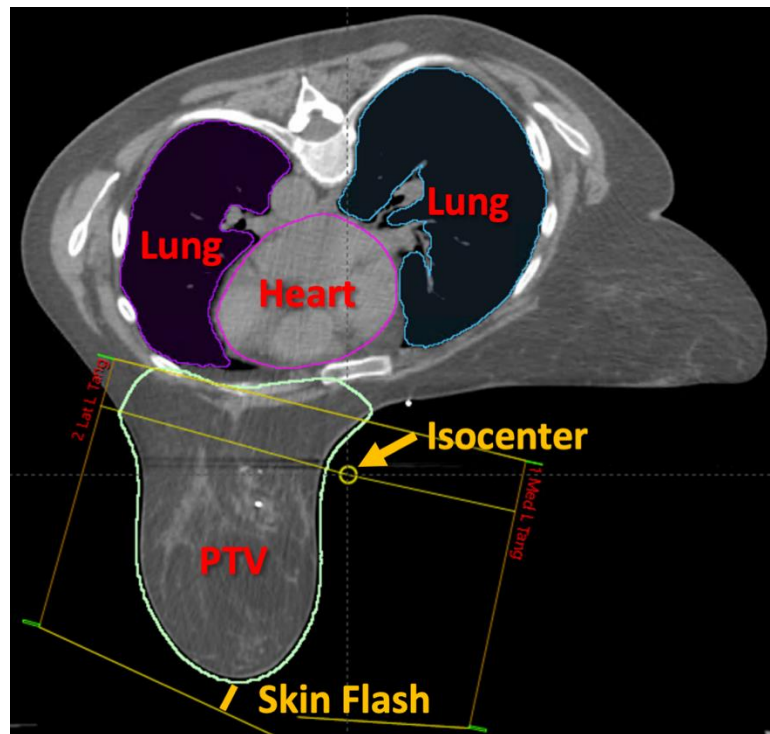


Figure 4: WBRT beam settings (prone position). Two tangential beams are also used for this setup. Due to gravity, the separation of breast and the chest wall is larger, especially for large breasted patients. The elongated breast tissues can also result in better dose homogeneity.

3.1.3 WBRT fluence modulation

Historically, WBRT utilized wedges to compensate for the missing tissue in the anterior breast seen from the BEV and improve dose homogeneity in the breast. Without wedges, the uniform photon beam fluence would result in large high dose regions in the

shallow part of the breast due to the anatomy of the breast. In the BEV, the depth of the breast tissue decreases distally, so a wedge compensates for the missing tissue.

However, a wedge can only modulate the fluence in one direction. To compensate for anterior-posterior direction and superior-inferior direction, multiple wedges should be used. Plus, the wedge angles and the field size using a wedge are limited, and breasts are also often shaped irregularly with skin folds, excessive tissue, and surgical areas. Therefore, it is difficult to achieve optimal dose homogeneity using wedges alone. More complex dose modulation is needed to account for irregular tissue variations in both anterior-posterior and superior-inferior directions.

A technique called Electronic Compensation (ECOMP) has been widely used for dose modulation in WBRT in recent years, allowing customization for individual patient anatomy. Similar to intensity-modulated radiation therapy (IMRT), which uses non-uniform beam intensities (fluence) determined by computer-based optimization techniques (Ezzell et al., 2003), ECOMP also modulates the intensity of each beam by using dynamic multileaf collimators (Friend, 2014). ECOMP uses the forward planning technique: the fluence maps are manually modified by the planner to compensate for the missing tissue in a process called dose painting, which is time-consuming and requires a high level of planner experience. Fluence editing and dose calculation will be repeated many times until the dose distribution is clinically acceptable. Thus, the treatment planning process usually takes 1 to 4 hours.

3.1.4 Automated WBRT planning

In previous studies by Purdie, Dinniwell, Letourneau, Hill, and Sharpe (2011) and Purdie, Dinniwell, Fyles, and Sharpe (2014), an automated treatment planning methodology was developed for inversely planned tangential breast intensity-modulated radiation therapy. For beam geometry, the authors utilized radiopaque breast wires for identifying the breast tissue and optimizing the beam. A number of clinically focused options were provided for the user to set. The algorithm optimizes both gantry and collimator angle based on lung and normal tissue volumes, where various gantry angle/isocenter combinations are scored until it is clinically acceptable (Purdie et al., 2011). The automated step-and-shoot IMRT plans were optimized via direct machine parameter optimization in the Pinnacle treatment planning system. The optimization parameters were automatically specified according to the target volume. The automated treatment planning process took around 8 minutes on standalone workstations and 5 minutes on commercial server hardware (Purdie et al., 2014).

Zhao et al. (2012) developed a method using support vector machine (SVM) to optimize the tangential beam placement for left-sided WBRT. The posterior treatment field plane, which separates the target volumes and the OARs, was found by applying an SVM on surface points of targets and OARs, which was assigned with different values. With the isocenter fixed, the beam placement was controlled by three parameters: gantry angle, collimator angle, and jaw size. The beam parameters were

derived from the separating plane generated by the SVM. While the SVM algorithm was used in this study, it was not used to generalize treatment planning knowledge from prior cases, but rather separating targets and OAR on a case-by-case basis. IMRT optimization was not included in this study.

As part of our effort to automate WBRT planning, a beam setting program has been developed as an ESAPI binary plug-in application. The program aims to achieve the optimal balance between PTV coverage and OAR sparing by learning the beam settings in existing clinical plans. Combined with the in-house fluence optimization program developed by Sheng et al. (2019), an automated treatment planning workflow for WBRT has been established. The goal of the workflow is to generate high-quality WBRT plans using the ECOMP delivery technique in a fraction of the time needed for manual planning. The majority of the beam setting program development has been published in *Technology in Cancer Research & Treatment* (W. Wang et al., 2019).

3.2 Materials and methods

The WBRT study workflow is illustrated in Figure 5. The beam setting program consists of two main components: initial beam generation and beam optimization. Fluence optimization was achieved with an in-house program, and dose evaluation was conducted between AI plans and clinical plans. The workflow includes two beam geometries: tangents only and tangents plus SCV. The prone setup was later added in the program as an update. The program was written as an ESAPI binary plug-in

program in Eclipse version 15.6, which could write the AI plan directly in the TPS. The three types of beam geometries share one main graphical user interface (GUI), with unique options for each geometry (Figure 6).

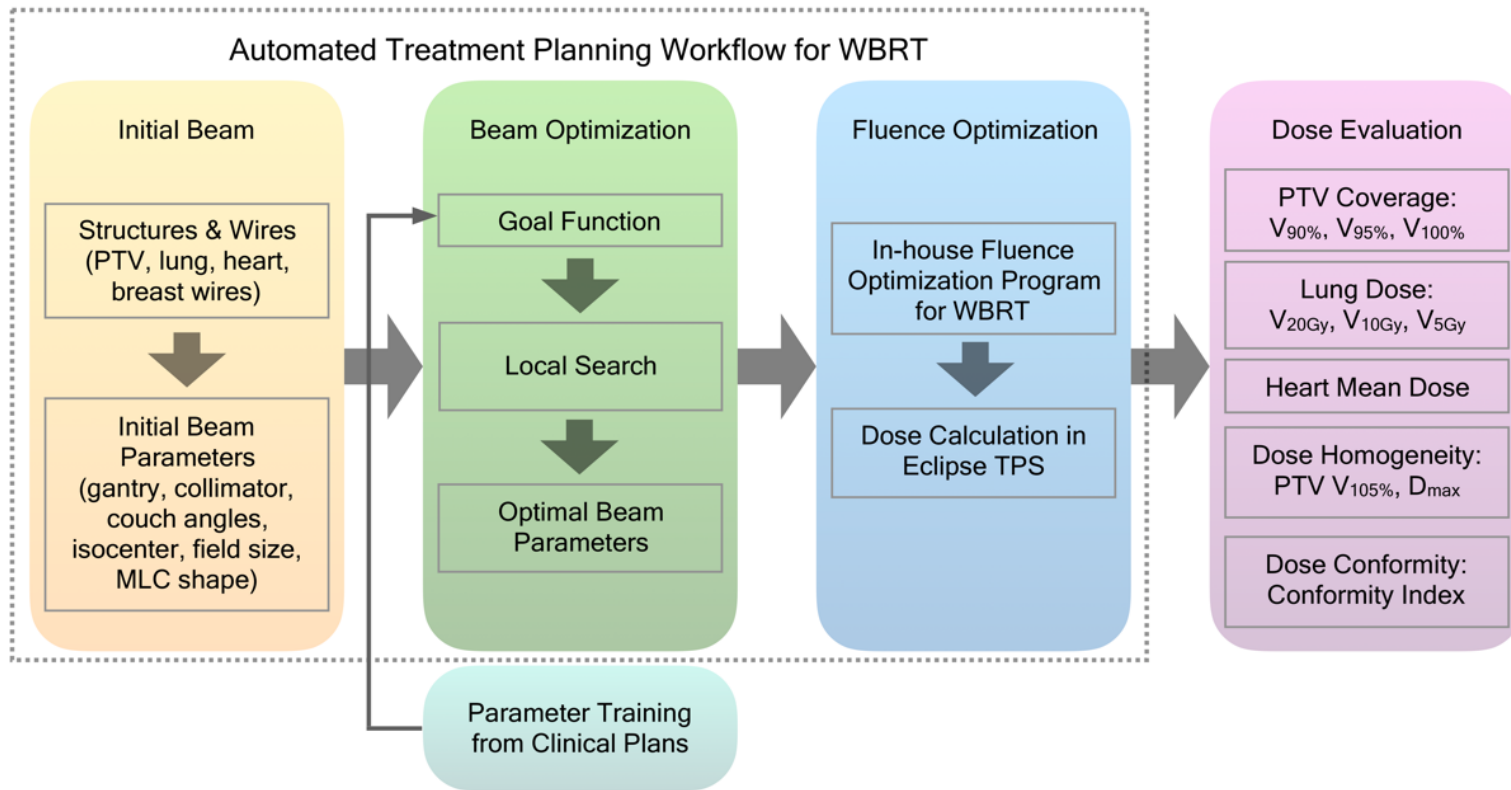


Figure 5: The WBRT study workflow. The automated WBRT treatment planning workflow is enclosed in the dashed lines.

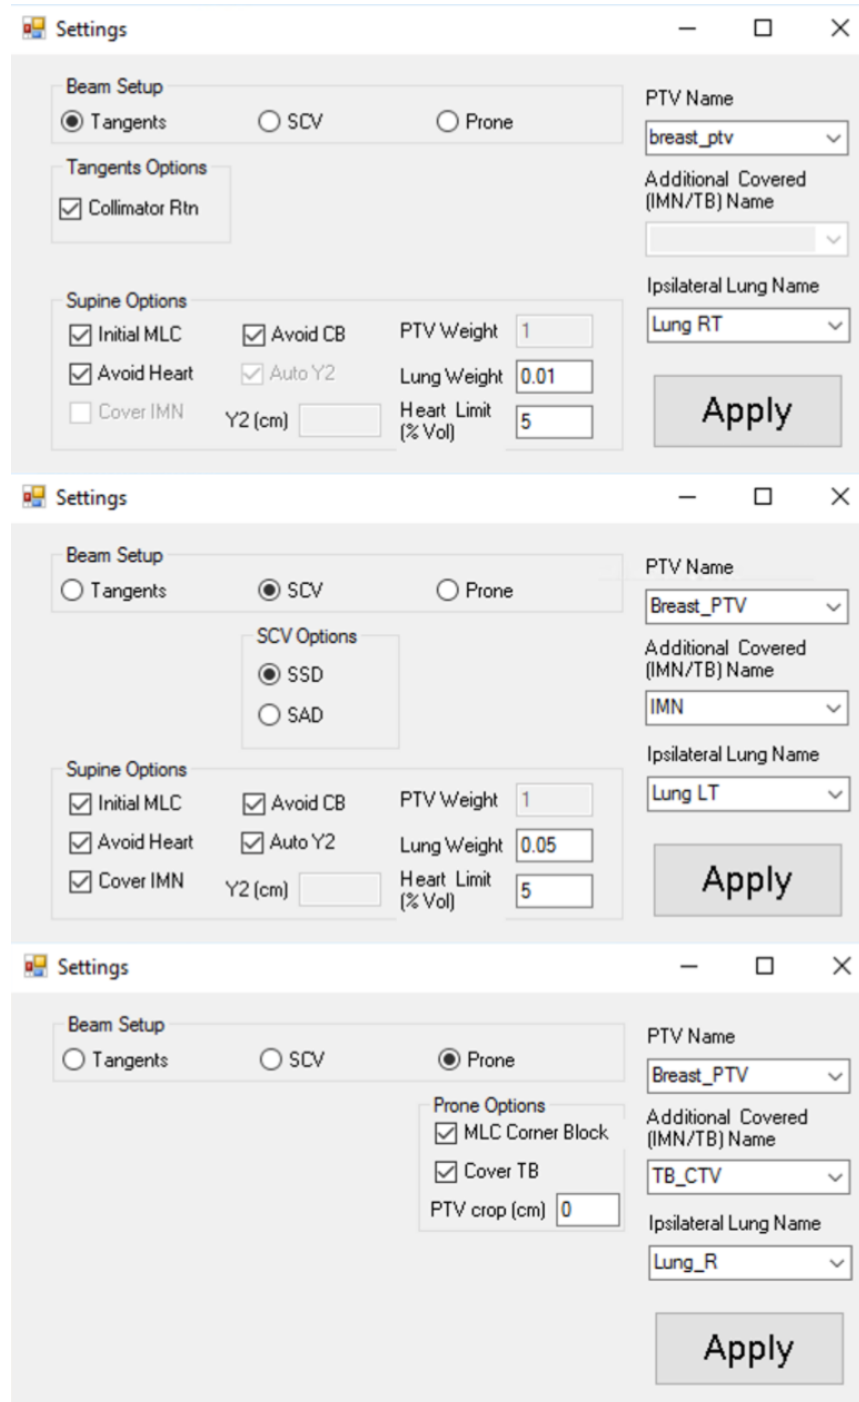


Figure 6: The graphical user interface for the breast beam setting program. Three types of beam geometries are shown (top: tangents only, middle: tangents plus SCV, bottom: prone), and each geometry has its own unique options for the planner to choose.

3.2.1 Patient population

Forty breast cancer patients treated with tangential WBRT in the supine setup at DUMC were selected for this retrospective study after institutional review board (IRB) approval. All patients had radiopaque demarcating wires placed to indicate the superior, inferior, midline, and lateral borders of the breast. Among them, 20 patients were treated with only tangential beams to cover breast PTV (“tangent only geometry”); 9 patients were treated for the right breast, and 11 for the left breast. The other 20 patients, including 12 postmastectomy cases indicated by the use of chest wall PTV, were treated with tangential beams plus supraclavicular (SCV) beams, which covered SCV PTV, axillary PTV, and internal mammary node (IMN) PTV (“tangent plus SCV geometry”) in addition to breast/chest wall PTV; 8 patients were treated for the right breast and 12 for the left breast. The average breast separation (the patient thickness along the central axis of the treatment tangential beams) of all 40 patients is 22.6 cm and the standard deviation is 3.2 cm. From the 20 patients of the same beam geometry, 10 were randomly selected as the training set, and the other 10 were reserved to test the program.

All patients include structure contours of breast/chest wall clinical target volume (CTV), breast/chest wall PTV, ipsilateral lung, heart, and midline wire. In addition, a structure for main PTV dose evaluation called “PTV_Eval” was also drawn. Breast/chest wall CTV were manually drawn by the attending physicians. The PTV was generated

with 5 mm margin from the CTV, and the PTV_Eval was generated from the PTV by removing 5 mm skin (to remove build-up region) and excluding chest wall.

3.2.2 Initial beam calculation

The program requires a patient-specific set of initial beam parameters to serve as a starting point for beam optimization. With the current clinical procedures, the patient-specific anatomical information needed for the initial beam calculation was acquired through wires placed during simulation and contoured PTV. The midline and lateral wires were used to determine the initial gantry angles and isocenter location. The CT0 is defined at the middle of the breast tissue range in the superior–inferior direction. During simulation, 3 radiopaque fiducial markers are placed on the patient for laser alignment, which are used to locate the origin point (CT0) of the CT coordinates. On the axial slice of the CT0, the treatment isocenter was set at the midpoint of the midline and the lateral wires. The gantry angles were determined by the beam central axis which goes through the 2 wires. The midline wire was available as a contoured structure. The lateral wire was automatically detected from the computed tomography (CT) image using thresholding methods on local image patches.

The field size includes 4 independent jaws (X_1 , X_2 , Y_1 , and Y_2). The X jaws were opened from the beam central axis and extended 2 cm beyond the apex of the breast as the skin flash region. In the tangent only geometry, the Y jaws were opened to cover the main PTV volume, with superior and inferior field margins; in the tangent plus SCV

geometry, the superior Y jaw (Y2) was instead opened to the clavicle head (the field matching location with SCV beam) regardless of the PTV, or it could also be specified by the user. In the tangent only geometry, the collimator angles were calculated based on the slope of the midline wire seen in the BEV and the couch angles were set as 0. In the tangent plus SCV geometry, the collimator and couch angles were solved analytically given the gantry angle and Y2. The initial MLC determines the shape of the beam aperture. The initial MLC leaf positions were calculated based on the PTV and OAR projections in the BEV. Several MLC options could be set by the user, including covering the IMN and blocking the heart.

3.2.3 Beam optimization

With the initial beam parameters set in the initial beam calculation, the program starts to optimize the beams to fit the PTV and OARs. The beam optimization is based on 2 parameters: the gantry angle and the isocenter location for tangential beams. Since the medial and lateral beams are parallel opposed, only one gantry angle is needed for optimization; here, the gantry angle of the medial beam was used. In theory, the isocenter location has 3 degrees of freedom: X, Y, and Z coordinates. The isocenter's Z coordinate was fixed at zero, and its location in the X-Y plane was restricted along the perpendicular line of the initial central axis, essentially reducing the dimension from three to one. The optimization goal was to find the gantry angle-isocenter pair that minimized the goal function that penalizes PTV under coverage and lung volume being

irradiated. The program automatically identifies the structures from the structure list using a template of their conventional names while also allowing the user to manually select the PTVs. The jaw size was fixed during optimization, as well as the collimator and couch angles for the tangent only geometry. Multi-leaf collimators (if used for shaping the aperture), the collimator and couch angles for the tangent plus SCV geometry were dynamically updated to match the tangent field for each gantry angle-isocenter pair.

The goal function Ω was based on the out-of-field PTV volume ($V_{ptv_{out}}$) and in-field ipsilateral lung volume ($V_{lung_{in}}$). It is written as

$$\Omega = V_{ptv_{out}}^2 + w_{lung} \cdot V_{lung_{in}}^2 + \infty \cdot H(V_{heart_{in}} - V_{th}) \quad (3.1)$$

$H(x)$ is the Heaviside function which takes on values of 1 and 0 for positive and negative x respectively, and 0.5 for $x = 0$. The $V_{heart_{in}}$ is the percentage volume of the heart in the field, and V_{th} is the user-defined threshold, which is 5% by default. The infinity sign (∞) is introduced to penalize non-negative input, which restricts the value of $V_{heart_{in}}$ to be smaller than the threshold.

The lung weighting factor w_{lung} was trained from clinical plans for the tangent only geometry and tangent plus SCV geometry separately. For each patient in the training set, the percentage volumes were calculated for the clinical beam setting and AI beam settings with different gantry angles and isocenters. The optimal w_{lung} was identified, which minimized the training set's sum of goal function differences between

the clinical beam setting and the optimum in the AI beam settings. The training process essentially averaged the clinical choices of the PTV-OAR trade-offs in the training set. The trained lung weighting factors are set as default values for the program, which could be easily adjusted by the user to achieve a desired trade-off between PTV and lung. In this study, the default value for each beam geometry was used for beam optimization.

The search space was a 2-dimensional grid centered at the initial beam setting. The grid size for the gantry angles and the isocenter location was 1 and 1 mm, respectively. Due to the design of the goal function and the spatial relationship between the PTV and the lung, the search space was convex with only one global minimum, which was established when an exhaustive search method was firstly experimented on the training sets. To accelerate the optimization, a local search scheme was employed in the program, where the current beam setting was moved to the optimal setting among its 4 nearest neighbors until a local optimum was reached. After the beam is optimized, the isocenter can be slightly shifted along the central axis direction to adjust the SSDs of the 2 tangential beams. Possible collision between the patient and the gantry head can be detected, and a warning would be sent to the user.

3.2.4 Plan evaluation

For plan evaluation, the beam setting choices, including the use of initial MLC, IMN coverage, the use of collimator rotation for the tangent only geometry, and the field

matching location for the tangent plus SCV geometry in the AI plans were consistent with the clinical plans.

After beam optimization, the in-house fluence optimization program (Sheng et al., 2019) built with ESAPI was executed to produce optimal fluence maps. Either single-energy (6 MV) beams or mixed energy (6 MV and 15 MV) beams were used for each patient, which was clinically determined by the dosimetrist and agreed by the attending physician. The beam energy choices of the AI plans followed the clinical choice. For plan dose comparison, the prescription doses (Rx) for all patients were set at 200 cGy per fraction for 25 fractions with the total dose of 5000 cGy. For tangent plus SCV geometry, only the dose contribution from the tangential beams was compared between the clinical and AI plans, excluding the SCV beams.

Several dose metrics were recorded from the TPS, including PTV_Eval $V_{105\%}$, $V_{100\%}$, $V_{95\%}$, $V_{90\%}$, ipsilateral lung V_{20Gy} , V_{10Gy} , V_{5Gy} , heart mean dose, maximum dose of body, and conformity index (CI). The CI (Paddick, 2000) is defined as the ratio of the absolute volume (of the body) covered by the 95% isodose over the absolute volume of the PTV_Eval, which is expressed as

$$CI = \frac{V_{95\%}^{body}}{V_{PTV_Eval}} \quad (3.2)$$

The mean values and standard deviations of the dose metrics were calculated for the clinical and AI plans of the training and test sets. Wilcoxon signed-rank tests were conducted with a significance level of .05.

3.2.5 Prone beam setup

As a small portion of the WBRT patients are treated in the prone position, a prone beam setup option was included in the program. As the prone setup is significantly different than the supine setup, a different beam optimization logic is used. For the prone setup, the isocenter location is either fixed at CT0 or shifted a few (3/5/7) centimeters anteriorly due to the jaw opening limit. Therefore, the only major parameter to optimize is the gantry angles. The program automatically sets the isocenter based on the X jaw size required for the optimized gantry angles.

First, the ipsilateral lung and the heart is combined into one “OAR” structure. Then, the OAR is gradually expanded until it overlaps with the breast PTV (overlap volume > 10 cc). For different gantry angles (medial beam within 30° range outwards from the horizontal position), the overlap volume is projected along the beam’s eye view into a 2D shadow. Due to the anatomy of the breast and the OAR, the optimal beam angle is selected by the minimal shadow area. As the BEV projection calculation is relatively fast, an exhaustive search is conducted for the entire gantry angle range with a step size of 1°.

As can be seen from Figure 6, the prone setup has the option to specify a PTV crop distance. This distance is selected by the planner to avoid irradiating a portion of the PTV closest to the OARs, which is one of the considerations for the beam setup in Figure 4. This option is left for the planner to choose the aggressiveness of the treatment

dose coverage. When the tumor bed (TB) is close to the chest wall, the option to “cover TB” will override the PTV crop, ensuring adequate dose coverage to the TB. For both the supine and prone setup, possible collision between the patient or the breast board and the gantry head would be alerted to the user.

The prone beam setting was developed with 9 existing cases for testing and debugging. It is considered an automated rule implementation and reasoning approach of automated planning. The testing for this functionality is ongoing in the clinic for new patients. As the program is tested with more patients, continuous updates will increase the robustness of the program.

3.3 Results

3.3.1 Calculation time

For each supine patient, the calculation time for beam geometry optimization varied between 10 seconds and 120 seconds, depending on the patient anatomy. For each prone patient, the calculation was near real-time, which took less than 5 seconds to generate the setup plan. The program was able to produce a set of beam parameters for every patient. The total planning time of the automated treatment planning process, including fluence optimization, was less than 5 minutes for each patient, which reflected a significant decrease from manual treatment planning.

3.3.2 Dosimetric evaluation

For the supine test patient set, the PTV_Eval $V_{95\%}$, ipsilateral lung V_{20Gy} , heart mean dose \bar{D}_H , maximum dose of body, and CI were tabulated for comparison between clinical and AI plans in Table 1. For all testing cases, the mean value for all dose metrics except lung dose was comparable between two plan groups, with AI plan achieving lower lung mean dose than the clinical plans. It should be noted from Table 1 that the mean lung V_{20Gy} for both beam geometries were lower for the AI plans, and that the mean PTV_Eval $V_{95\%}$ for tangent plus SCV geometry were higher for the AI plans, which represents improvement in both target coverage and OAR sparing.

Table 1: Dose metrics comparison between AI and clinical plans. The numbers in each cell are reported as the mean value (standard deviation).

		PTV_Eval $V_{95\%}$, %	Lung V_{20Gy} , %	Heart \bar{D}_H , % Rx	Maximum Dose, % Rx	CI
Tangent Test Set	Clinical	94.4 (3.5)	12.9 (5.4)	2.2 (1.4)	108.7 (1.4)	1.40 (0.16)
	Auto	94.5 (2.7)	9.8 (5.1)	2.0 (1.0)	109.1 (0.6)	1.32 (0.10)
	p-value	0.846	0.049*	0.416	0.516	0.062
SCV Test Set	Clinical	82.6 (7.8)	23.0 (8.6)	2.4 (1.5)	109.2 (2.9)	1.43 (0.52)
	Auto	87.5 (6.2)	20.7 (8.0)	2.7 (1.6)	108.5 (1.4)	1.45 (0.54)
	p-value	0.037*	0.232	0.049*	0.316	0.977
All Test Set	Clinical	88.5 (8.4)	17.9 (8.7)	2.3 (1.4)	108.9 (2.2)	1.41 (0.38)
	Auto	91.0 (5.9)	15.2 (8.6)	2.4 (1.4)	108.8 (1.1)	1.38 (0.38)
	p-value	0.079	0.025*	0.466	0.727	0.239

*indicates statistical significance.

3.3.3 PTV and lung trade-off

To evaluate the combined performance of the PTV coverage and lung sparing, which are typically trade-off criteria, the ipsilateral lung V_{20Gy} difference is plotted case

by case against the PTV_Eval $V_{95\%}$ difference for all testing cases in Figure 7. The goal driven beam optimization can improve lung sparing without compromising PTV coverage. With different gantry angles and isocenter location from the clinical plan, the irradiated normal tissue volume was reduced. For the 6 patients in the upper-left quadrant, trade-off favored more lung sparing than PTV coverage for the AI plan. For the 6 patients in the lower-right quadrant, lung dose was higher while better PTV coverage was provided. No AI plan showed inferior PTV coverage and worse lung sparing. The two beam geometries are separated by color, with red being tangent only geometry and blue being tangent plus SCV geometry. All but three tangents plus SCV geometry cases achieved similar PTV coverage with PTV $V_{95\%}$ differences within $\pm 5\%$. The three outlier cases had considerably higher ($> 9\%$) PTV $V_{95\%}$ for the AI plans, and two of them had increased lung dose.

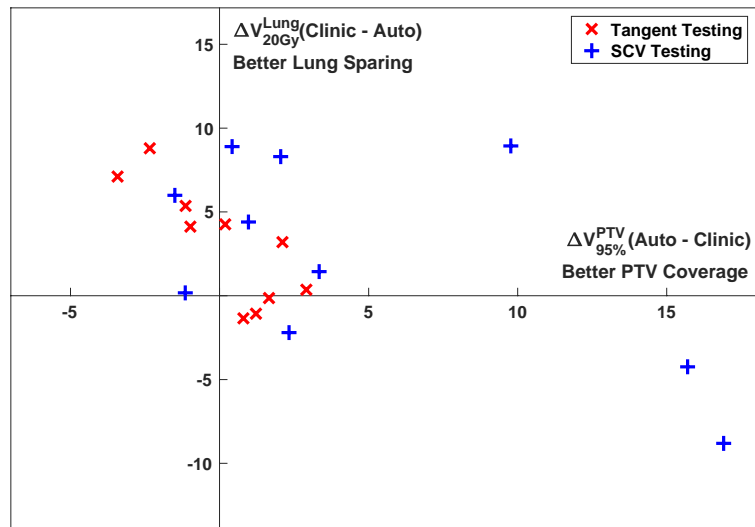


Figure 7: PTV coverage (PTV_Eval $V_{95\%}$) and lung sparing (ipsilateral lung V_{20Gy}) comparison between AI plans (Auto) and clinical plans (Clinic). The horizontal axis is the difference (Auto – Clinic) of PTV_Eval $V_{95\%}$, and positive direction on the axis means better PTV coverage; the vertical axis is the difference (Clinic – AI) of ipsilateral lung V_{20Gy} , and positive direction on the axis means better lung sparing. Different patient groups are separated and denoted with different markers as shown in the legend.

3.3.4 Example cases

As can be seen, the AI plan achieved both superior PTV coverage and better lung sparing for 8 of 20 patients (upper-right quadrant). An example case in this quadrant is shown in Figure 8.




Figure 8: Dose distribution comparison in the same axial slice of one example case between (A) the AI plan and (B) the clinical plan. The thick green line is the 95% isodose line, and the thin purple line is the PTV_Eval contour on this slice. The white arrows point to the same position at the edge of the lung (note the distance to the 95% isodose line). This case shows how the AI plan improves the sparing of the lung without sacrificing the PTV coverage.

One example case for demonstrating PTV coverage and lung sparing trade-off is shown in Figure 9. This case falls in the lower-right quadrant in Figure 7. It illustrates how beam geometry can affect the dose distribution and the PTV-lung balance. The posterior field border in the AI plan (Figure 9A) was placed deeper into the ipsilateral lung than the clinical plan (Figure 9B), in exchange of better PTV coverage. By increasing the lung weighting factor from the default 0.05 to 0.2, the program produced another AI plan (Figure 9C), which was in between the previous two plans (Figure 9A and Figure 9B) in terms of PTV-lung balance.




Figure 9: Dose distribution in the same axial slice of one example case in the AI plan with default lung weighting factor of 0.05 (A), the clinical plan (B), and the AI plan with increased lung weighting factor of 0.2 (C). The thick green line is the 95% isodose line, and the thin orange line is the PTV_Eval contour on this slice. The arrows point to the same position at the edge of the PTV_Eval contour. This case shows how the AI plan improves the PTV coverage at the expense of increasing the lung dose and how the lung weighting factor affects the dose coverage.

Two examples of prone setup comparison between AI plans and clinical plans are shown in Figure 10 and Figure 11. In Figure 10, the breast PTV is mostly covered in the clinical plan and the AI plan. However, the clinical beams are more tilted towards the patient lateral side, resulting in a larger lung volume being irradiated than the AI plan. The PTV coverages in the two plans are similar.




Figure 10: Prone setup and dose example 1. The breast PTV in both plans is mostly covered. The AI plan (left) achieved similar PTV coverage as the clinical plan (right) without irradiating as much lung volume.




Figure 11: Prone setup and dose example 2. The portion of the breast PTV adjacent to the ipsilateral lung and the heart is not covered. Note that the PTV crop margin is manually set in the AI plan (left) based on the planner's judgement.

In Figure 11, a significant portion of the PTV closest to the lung and the heart is not covered by the clinical beams. Therefore, for this type of cases, the planner has the option to use the PTV crop option in the program to setup the AI beams. As each calculation is near real-time, the user can easily create multiple setup plans with different PTV crop margins. The final beam setting should be based on the clinical judgement of the planner or the physician.

3.4 Discussion

This study developed an automated beam geometry optimization program that enables the user to quickly generate a customizable and clinically feasible beam

geometry for WBRT treatment planning. The customizable user settings add flexibility to the beam geometry and can easily adapt to most clinical scenarios. The adjustable lung weighting allows the user to fine-tune the balance between PTV coverage and lung sparing. The study is of substantial clinical significance given that it dramatically reduces the time required for WBRT treatment planning and improves the clinical efficiency. This work also reduces inter-human variation, as its underlying parameters reflect clinical importance. Together with the automated fluence optimization program, an automated treatment planning workflow for WBRT can be integrated into the Eclipse treatment planning system to replace the routine breast treatment planning manual process.

From the data in Table 1, the tangent only geometry cases generally have better PTV coverage and lung sparing compared to the tangent plus SCV geometry cases. The reasons for this observation are two folds. Firstly, if portions of the PTV extend superior to the clavicle head in the tangent plus SCV geometry, they are covered by the SCV beam(s) instead of the tangential beams. Since the dose evaluation only included the tangential beams, this part of the PTV receives no primary dose from the plans in this study. Secondly, most clinical plans of the tangent plus SCV geometry cover the PTV including the IMN, which extends more medially than the PTV in tangent only geometry. This part of the target significantly overlaps with the OARs (lung and heart) seen from the BEV, resulting in higher lung and heart dose.

In many clinical plans of the tangent only geometry, the entire PTV was covered by the beams, because the OAR dose constraints were more easily satisfied as the PTV-OAR overlaps in the BEV were smaller. The irradiated lung volume was not an issue for these easier cases. In most clinical plans of the tangent plus SCV geometry, however, PTV coverage was compromised in order to restrict OAR doses. For these more difficult cases, the planner was required to achieve a balance between PTV coverage and OAR sparing. In the clinical plans, the average percentage out-of-field PTV volume and in-field ipsilateral lung volume are 0.93% and 13.12% for the tangent only geometry, and 4.53% and 29.85% for the tangent plus SCV geometry. The freedom to adjust the lung weighting factor allows the planner to choose different priorities as shown in Figure 9. Therefore, the target differences justify the use of different lung weighting factors in the goal function for the two beam geometries.

The program is designed to be customizable and flexible. It is feasible to handpick a group of similar training cases to train the lung weighting factor which tailors to a specific clinical scenario or an individual planner. Since the patients in this study were randomly selected from the same patient cohort, the clinical plans include a small range of clinical preferences (PTV-OAR balance). As the trained lung weighting factor reflects the averaged clinical preference, the size of the training dataset does not significantly affect the training result in this study. To validate the default lung weighting factor and investigate the training data size impact on the model

performance, for each beam geometry, the training process was repeated 10 times on 19 cases (10 original training cases and 9 testing cases, with one testing case left out each time). The average lung weighting factors (standard deviation) trained from 19 cases were 0.012 (0.002) for tangent only geometry and 0.059 (0.004) for tangent plus SCV geometry, while the default values trained from 10 cases are 0.02 and 0.05. A planner can also explore the automatic plan space by experimenting with different parameters and settings within minutes and choose the optimal one based on clinical judgement. However, the MLC shape algorithm is currently not affected by the lung weighting factor, which limits the variability of the beam aperture shape given different weighting factors. In general, the optimization result is not sensitive to small changes in the lung weighting factor but depends heavily on the beam geometry settings (e.g., use of MLC, heart block). Even if the planner is not completely satisfied with the automatic plan, it is easy to manually fine-tune the beam parameters and/or MLC shapes within the TPS before fluence map optimization.

One important feature of the program is its dependence on the structure contours. The program calculates the 3-dimensional spatial relationship between the beams and relevant structures. For the same patient anatomy, the automatically generated beam geometry could be different if the target and OAR contours were drawn differently. As the contours were manually delineated, even if the same guideline is followed, there could be variation in contours among different operators, especially for

the PTV. Therefore, automatic contouring will further improve planning efficiency and the consistency of plan quality. One limitation of current form of beam optimization is that it is looking for breast wires to start with. Breast wire is placed during CT simulation to mark the boundary of the breast tissue. However, variations of the wire placement would not significantly affect the beam optimization result as they are only used to determine the optimization starting point and the collimator angles in the tangent only geometry. In order to deploy this work, additional steps of wire placement and PTV contouring need to be added if not previously implemented. It would be of interest to develop beam optimization without wire, which offers greater efficiency, and research along this line is warranted. The proposed implementation of beam optimization does not waive the necessity of other procedures in the treatment chain, such as CT simulation, chart checking, treatment simulation or IMRT QA.

3.5 Conclusion

An automated goal-driven beam setting optimization program for whole breast radiation therapy is developed in this study. The program is able to produce optimized beams for all patients. Together with the fluence optimization program, automated treatment planning for WBRT is made possible. Plans with automatically generated beam settings achieved comparable plan quality as manually generated clinical plans in terms of PTV coverage, OAR sparing, dose homogeneity, and dose conformity. This program offers a valuable tool for WBRT treatment planning, as it provides clinically

relevant solutions based on previous clinical practice as well as patient specific anatomy under a substantially faster time frame.

4. Atlas-guided fluence initialization and optimization

4.1 Introduction

In IMRT planning, the optimal fluence maps of all beams determine the final dose distribution of the patient. Inverse optimization is the process that iteratively modifies the fluence intensities to achieve the desired optimal dose. For pancreas SBRT, the challenging task of respecting luminal OAR constraints while maximizing target coverage is dependent on the planner's experience to finetune the optimization parameters. Therefore, treatment planning time can be prolonged to find the optimal solution to meet clinical needs.

In order to aid the inverse optimization process, many studies have attempted to infer the dose distribution from patient anatomy. While the total dose distribution is intuitively related to the patient anatomical contour, i.e., generally higher dose to the PTV and lower dose to the OAR, the relationship between fluence maps and patient anatomy is not as straightforward. The reasons include: (1) multiple beams simultaneously contribute dose to the same region in the patient, i.e., the cross-firing nature of IMRT plans; (2) each beam's path is unique in terms of depth and tissue attenuation, affecting the individual dose to the target volume.

In automated treatment planning, previously treated plans can be assembled into a plan database (pool). Based on patient anatomy similarity, a reference (atlas) plan can be found for the new patient to aid the treatment planning process. T. Li et al. (2011)

created a patient-specific plan database for adaptive image guided radiation therapy (AIGRT) for prostate patients. The database consisted of the original plan and re-optimized plans from daily CBCT images. An automatic plan selection algorithm attempted to match the treatment day CBCT with previous images in the database and re-position the patient for the matched plan delivery. When no existing plans satisfied the coverage criteria, a new plan was re-optimized and added to the database. Sheng et al. (2015) developed an atlas-based IMRT planning technique for prostate cancer, which used a k-medoids clustering analysis to find anatomy pattern variations in the 70-case prostate plan dataset. Deformable registration was performed on the best matched atlas plan dose to the new patient anatomy. Fully automated inverse planning for the new patient was guided by the deformed atlas dose distribution.

This study tests the hypothesis that there is an inherent relationship between patient anatomy and fluence maps for a set of fixed beam angles. The optimal fluence maps can thus be deduced from previous plans with similar anatomy. The study investigated the feasibility of using atlas matching to initialize and accelerate fluence optimization for pancreas SBRT. It provided key anatomical features related to fluence optimization and set the foundation for deep learning-based fluence map prediction. The study was presented at the 2019 AAPM annual meeting (W Wang et al., 2019).

4.2 Materials and methods

The study workflow is summarized and illustrated in Figure 12. Thirty-three pancreas SBRT patients were retrospectively included in the study, which were split into 30 patients for training and 3 patients for testing.

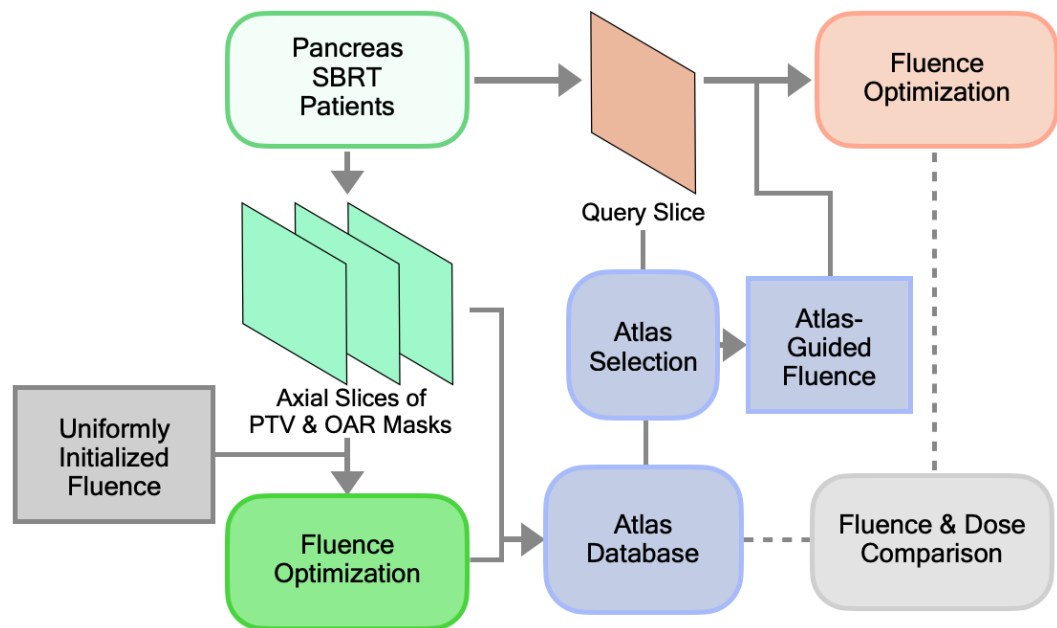


Figure 12: The atlas-guided fluence optimization study workflow. The atlas database was built with pancreas SBRT plans inversely optimized from uniformly initialized fluence maps. The atlas selection for a new patient was performed slice by slice, and the best matched fluence was used to initialize the fluence optimization.

4.2.1 Atlas database building

For the 30 training cases, the PTV and OAR contours of all 2D axial slices containing the PTV were extracted. The OAR considered in this study is the duodenum (C-Loop). As a feasibility study, the problem was reduced to 2D treatment planning. An IMRT plan was created for each axial slice using 9 equally spaced beams. Fluence

optimization was performed on these 2D slices using the in-house dose calculation algorithm and fluence optimizer.

The PTV prescription was set at 33 Gy in 5 fractions, and the OAR limit was 25 Gy maximum dose. The resulting fluence maps for each slice were nine one-dimensional arrays. To increase the training data size, data augmentation techniques were used on the training cases. Random shift in location and scaling of structure size were applied to the PTV and OAR, increasing the number of slices to 5944. Each atlas in the database included the PTV, OAR contours and the optimized fluence maps for 9 beams.

4.2.2 Atlas selection

A scoring system was developed for atlas selection, which matched the query slice with an anatomically similar atlas slice. First, each slice was assigned to one of three groups according to its PTV-OAR minimum distance:

Group A: The slice only contains the PTV without any OAR.

Group B: The PTV overlaps with the OAR, or their minimum distance is smaller than 10 mm.

Group C: The minimum distance between the PTV and the OAR is larger than 10 mm.

For a query slice, we only searched for a match within the same group. The matching process is determined by a scoring system for a pair of patient contours. The scoring system considers the PTV and its two largest “dents”, which are regions outside

the PTV contour but enclosed in its convex hull (Figure 13). The dents are numbered according to their sizes.




Figure 13: Illustration of two "dents" in the PTV contour. Dent 1 and 2 are the largest and second largest dents of a concave PTV. The angle θ is calculated with the centroids of PTV and the two dents, and it is between 0 and 180 °.

For each shape (including PTV and its dents), the ratio (PTV: r_0 , dent 1: r_1 , dent 2: r_2) of the minor and major axis lengths of the ellipse with the same second central moment are computed. The area ratios (dent 1: a_1 , dent 2: a_2) of the dents and the PTV, and the angle θ between the two dents with respect to the PTV centroid (Figure 13) are also calculated. The score is the weighted sum of the "distances" between the query and atlas PTV shapes, which should be minimized. It can be expressed as

$$\begin{aligned} \text{Score} = & w_0 \times \frac{\|r_0^{(m)} - r_0^{(q)}\|_2}{(r_0^{(q)})^2} + w_1 \times \left(\|a_1^{(m)} - a_1^{(q)}\|_1 + \|r_1^{(m)} - r_1^{(q)}\|_1 \right) \\ & + w_2 \times \frac{90^\circ - \|\theta^{(m)} - \theta^{(q)}\|_1}{90^\circ} \times \left(\|a_2^{(m)} - a_2^{(q)}\|_1 + \|r_2^{(m)} - r_2^{(q)}\|_1 \right) \end{aligned} \quad (4.1)$$

where the superscript (q) stands for query slice, and (m) stands for matched atlas slice.

The score is a weighted sum of three terms, and w_i is the weighting factor for each shape. When either $\theta^{(m)}$ or $\theta^{(q)}$ does not exist, the angle θ term is taken as 1.

The atlas slice with the minimum score was considered the matched slice. The fluence map from the matched slice was then scaled to the query PTV's size and used for fluence initialization in the inverse optimization of the query slice.

During model training, leave-one-out cross-validation was performed to tune the scoring system. After the weights were finalized, all 30 training cases were included in the atlas database, which was tested on the 3 test cases. For comparison, uniform fluence was used to initialize the inverse optimization.

4.3 Results

The optimization objectives remained the same for two optimization schemes. The average initial cost function values for the inverse optimization were 2097 for uniform initialization and 715 for atlas-guided initialization. Atlas-guided fluence optimization reduced the mean initial cost function value by 69.7% ($p < 0.01$). Figure 14 is the box plot showing the relationship between cost function reduction and iteration number for the test set. The optimized DVH comparisons (average for all test slices) between the two optimization schemes are shown in Figure 15. For dosimetric endpoints, the average OAR mean dose and PTV V100% were 14.07Gy and 37.1% for atlas-guided plans, and 14.21Gy ($p < 0.01$) and 35.7% ($p < 0.01$) for uniformly initialized plans.




Figure 14: The box plot for the cost function reduction for atlas-guided fluence initialization vs iteration number. The red central mark on each box indicates the median, and the bottom and top edges of the box indicate the 25th and 75th percentile, respectively. The largest reduction occurs at iteration 0, which corresponds to the initial cost function reduction with atlas-guided fluence initialization. As the iteration number increases, the differences between the two optimization schemes diminish rapidly.

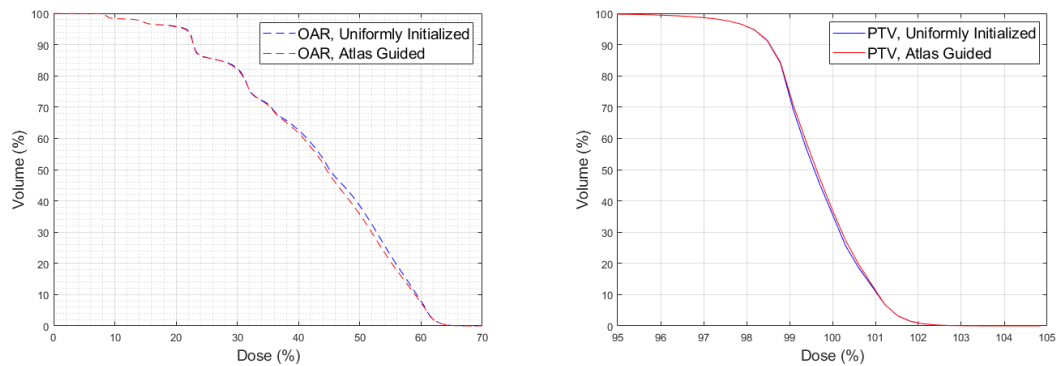


Figure 15: The averaged OAR DVH (left) and PTV DVH (right) comparisons between uniformly initialized optimization (blue lines) and atlas-guided optimization (red lines). Atlas-guided optimization achieved slightly lower OAR doses and higher PTV V100%.

4.4 Discussion and conclusion

The preliminary data showed that atlas matching could help quickly approach optimal fluence in IMRT inverse optimization. With the simple constraints used in this study, improved dosimetric outcomes were observed for atlas-guided fluence optimization. The results showed feasibility to use prior fluence from plans with similar anatomical patterns to guide pancreas SBRT fluence optimization.

However, the limitation of scoring systems and database size was the major constraints in finding a high-quality fluence map for the given anatomy. This approach would also face challenges when translating to 3D planning, as the similarity of treatment volumes becomes much more complex. The superior-inferior shape variation would limit the implementation of this technique due to beam divergence and scattered dose. In addition, the fluence at the edges depends heavily on the structure information from adjacent slices if no contour is present in that slice. It can be anticipated that structure shape matching in the 3D space would require a significantly larger atlas database. As the current study used a fixed set of coplanar IMRT beams, the use of different beam angles and non-coplanar beams necessitates a different approach. Nonetheless, we were motivated by the results and continued to explore deep learning techniques for generating optimal fluence maps for pancreas SBRT.

5. Deep learning-based fluence map prediction framework

5.1 Introduction

5.1.1 The limitations of dose prediction

Deep learning has been implemented in automated treatment planning mainly through dose prediction. The correlation between structure contours (PTV and OAR) and the optimal dose distribution can be effectively learned by deep neural networks. The actual dose distribution is a compromise between the idea dose (full prescription dose to the PTV and no dose to the OAR) and the physical limitations of the linac and its radiation beam (entry and exit dose, scatter/leakage dose, etc.). As the idea dose information can be obtained as the input to the deep neural network, the task of deep learning is the imitation of the achievable dose, considering all the physical limitations and the resulting clinical tradeoffs.

While the predicted dose is valuable in evaluating existing clinical plans and assisting knowledge-based planning, a dose mimicking step, which commonly use inverse optimization, is required to create a deliverable plan. For IMRT plans, the dose is delivered by multiple beams with fixed gantry angles and modulated fluence maps. Thus, the direct prediction of fluence maps for IMRT beams can bypass the inverse optimization and generate deliverable plans.

5.1.2 Existing fluence map prediction methods

A few studies have focused on the direct prediction of fluence maps. In the previously described automated planning workflow for WBRT, the fluence maps for the optimized beams were generated by a machine learning method developed by Sheng et al. (2019). The model utilized the digitally reconstructed radiographs (DRR) of the beams as input. PCA was used to reduce the dimension of the gray level histogram of the DRRs and extract useful features. Then, initial fluence maps were predicted pixel-wise by a random forest model with shape-based features as inputs. The RF generated fluence maps were not optimal for the specific patient anatomy. Therefore, a fluence finetuning step using centrality correction was added to provide patient-specific coverage and hot spot reduction.

H. Lee et al. (2019) proposed a deep learning-based fluence prediction method, which used the dose distribution per beam and organ contours as input to predict fluence maps for 7-beam prostate IMRT plans. All input data (3D organ contours and dose) are viewed from the BEV of each beam and divided into 64 slices (each one 10 mm thick) using the cone beam geometry. The network architecture was a modified U-Net, using the structural similarity (SSIM) index and the mean absolute error (MAE) with equal weightings as cost functions. Each fluence map was predicted independently without influence from other beam angles. The authors acknowledged that a separate network using only the dose distribution as input achieved similar overall performance

as the proposed network (H. Lee et al., 2019). This suggests that beam dose provides the majority of the necessary information to predict fluence maps, which is consistent with the physical dose deposition process. While the study achieved decent prediction accuracy for fluence maps (mean value of the fluence map MAE: 9.95×10^{-4}), the input dose distributions were extracted from existing clinical plans, which could not be directly used for new patients.

X. Li et al. (2020) developed a deep learning based fluence prediction algorithm (AIP-SFFP) for 9-beam prostate IMRT plan generation. In this method, the patient structure information was first projected onto the 2D beam's eye view. The input projections included infrastructure projections of the PTVs and interface projections of the PTVs and the OARs. These projections were derived from the photon attenuation coefficient integration in or between critical planning structures and stacked for all beam angles. The nine fluence maps were predicted by a Dense-Res Hybrid Network from these input projections and post-processed by a Gaussian-based 2D de-convolution operation for edge information enhancement. The authors reported higher 3D maximum dose and similar PTV coverage compared to the KBP plans (ground truth).

5.1.3 Beam dose and fluence map

While current dose prediction networks can predict the total dose of the optimal plan with relatively high accuracy, no study has proposed a solution for the prediction of individual beam's dose from patient anatomy. Beam dose prediction faces some

unique challenges compared to total dose prediction. An example axial slice of the total dose vs beam dose in an IMRT plan is shown in Figure 16. As can be seen in the figure, depending on the beam angle, the beam dose follows a straight path in a certain direction with beam divergence effect. All dose distributions are bounded by the patient body. Each radiation beam is attenuated when interacting with the tissue and has a sharp fall off at the edge of the beam due to the linac secondary collimator (jaw). Since the source to skin distance (SSD) and tissue composition are variable for different beams, the same fluence intensity will result in different dose values at the isocenter. Inside the primary beam, the dose distribution in the transverse direction is determined by the fluence intensity and scattered dose.




Figure 16: An example of total dose (top) and three of the beam doses (bottom) in a 9-beam IMRT plan.

Even though beam dose does not share the shape similarity with structure contours, we believe that, with standardized dose constraints and beam settings, the beam doses of an IMRT plan could be predicted from patient anatomy via deep learning. The beam doses can then be used to reconstruct the fluence maps, which enables direct generation of IMRT plans. The underlying assumption is that the inversely optimized fluence maps of a plan are fixed given the patient anatomy, beam setting, and dose constraints. In this feasibility study, we developed a novel deep learning framework for direct fluence map prediction (a.k.a. direct plan generation) and demonstrate its performance on pancreas SBRT cases with a single PTV. The study has been published in *Frontiers in Artificial Intelligence* (W. Wang et al., 2020).

5.2 Materials and methods

5.2.1 Patient selection and radiation therapy plan

One hundred pancreatic cancer patients previously treated with SBRT at Duke University Medical Center between 2014 and 2019 were included in this retrospective study. This study was approved by the institutional review board. In clinical plans, the dose prescription to the PTV was 25 Gy, often with simultaneous integrated boost to the internal gross tumor volume (iGTV) with 33 or 40 Gy. The GI OAR (stomach, C-loop/duodenum, and bowels) dose constraints varied in maximum dose and maximum volume with different physician preferences.

We aim to develop a model that is capable of generating clinical quality pancreas SBRT plans using multiple IMRT beams. In this feasibility study, each case was re-planned by experienced clinical physicists specialized in GI SBRT using unified planning objectives and a standardized IMRT protocol with single prescription level. The prescription for both the PTV and iGTV were 33 Gy in 5 fractions. All plans were designed with 9 equally spaced coplanar 10-MV photon beams (gantry angles: 20°, 60°, ..., 340°). Stomach, C-loop/duodenum, and bowels were combined and referred to as the OAR. The maximum dose for the OAR was limited to 25 Gy (0.1 cc). This protocol creates the scenario of inverted relationship of target and OAR dose prescription, a clinical scenario that often has to be handled manually by an experienced planner for each case.

In the following, we will refer to the resulting standardized plans as the benchmark plans, which were used to train the model. The same beam orientations, including gantry angles, beam shape definition via its open field dose, referred to as beam templates, were also included as input for the DL models. The 100 patient cases were divided randomly into 85/15 ratio for training/testing. All treatment plans were generated in a research only Eclipse® Treatment Planning System (Varian Medical Systems, Palo Alto, CA) version 13.7 with volume dose calculated by the Analytical Anisotropic Algorithm version 13.7.14. A Varian Millennium 120 multi-leaf collimator

(MLC) was used to deliver the modulated fluence maps. The leaf sequencing algorithm used was Smart LMC version 13.7.14.

5.2.2 Study workflow




Figure 17: The plan generation workflow of the proposed deep learning framework.

The plan generation workflow of the proposed deep learning framework is illustrated in Figure 17. The proposed framework adopts a pipeline structure, where two CNNs make consecutive predictions to generate the complete plan with fluence maps. The input into the pipeline includes the beam template as well as contours of the PTV and the OAR. The beam template refers to the beam doses in the beam template plan, which uses open beams (unmodulated fluence) in the 9 specified gantry angles. The beam dose CNN (BD-CNN) predicts the 9 beam doses in 9 output channels. Next, each 3D beam dose is projected along the beam's eye view (BEV), generating the 2D BEV dose map. Finally, the fluence map CNN (FM-CNN) predicts the fluence map for each beam from the corresponding BEV dose map. The predicted fluence maps are imported in the treatment planning system to generate the AI plan. The AI plans were compared with the benchmark plans to evaluate the model performance.

The two CNNs were implemented in Keras with Tensorflow backend and trained separately. The entire model was trained on a workstation with an Intel Xeon E5 v4 processor, 64 GB of RAM, and a NVIDIA Quadro M4000 graphics card.

5.2.3 Data preprocessing

All plan data were exported from the Eclipse TPS as DICOM files, except for the benchmark fluence maps, which were text files. As the original plans have different spatial resolutions, resampling was performed on dose and contour images with 1 mm axial resolution and 2 mm slice thickness. Linear interpolation was used to increase the resolution of dose distributions to facilitate more accurate dose prediction. Relative values were used in field doses, with the prescription dose of 33 Gy normalized to 100%. Axial slices were cropped to a 192×192 pixels image centered at the isocenter. Fluence maps and other BEV projections had a resolution of 2.5×2.5 mm² at isocenter plane, which was the default resolution in Eclipse. All the training data were shuffled before holding out a validation set.

5.2.4 Beam dose prediction

The objective of BD-CNN is to predict field doses from CT images and structure contours. The network architecture of BD-CNN is illustrated in Figure 18, and it operates on a slice-by-slice basis.

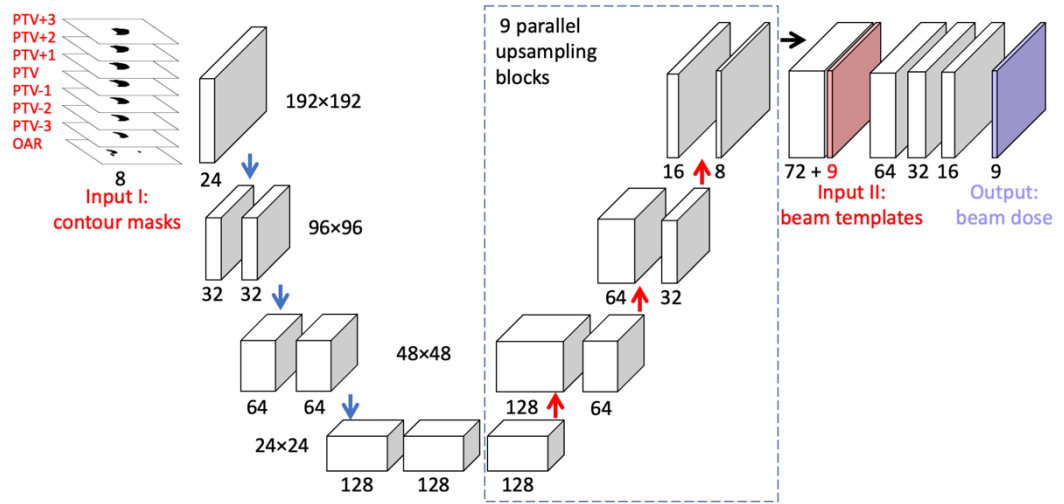


Figure 18: The network architecture of the BD-CNN in the pancreas DL framework. The number below each layer is the channel number. Each adjacent layer is connected by a convolution from left to right. Blue arrows indicate the downsampling operation, while red arrows indicate the upsampling operation. The upsampling block (enclosed in the dashed line) is duplicated 9 times in parallel, and their outputs are concatenated (black arrow).

To predict beam dose in one query slice, the main input includes seven PTV slices (the query slice and six adjacent slices in the superior and inferior directions) and the OAR query slice, which are all 192×192 binary masks. The adjacent PTV slices were included to account for PTV shape change in the Z direction. In the downsampling block, the contour masks were downsampled three times using strided 2D convolution to produce 128 channels of 24×24 feature images. An upsampling block produced 8-channel feature images, using strided 2D transposed convolution three times to restore the 192×192 resolution. There were 9 parallel upsampling blocks, which generated the 72-channel output. CT images were incorporated in the form of beam templates (Input II in Figure 2) calculated by the TPS and concatenated to the 72-channel upsampling

output. A final convolution block was applied to produce nine beam doses for the nine equally spaced beam angles. The prediction region was limited to a region-of-interest (ROI), which was the PTV expanded by 1 cm.

The Swish activation function (Ramachandran, Zoph, & Le, 2017) was used in the network to introduce nonlinearity. Other commonly used activation functions include rectified linear unit (ReLU), leaky ReLU, and Softmax. Swish is the product of an identity function and a sigmoid function, which can be expressed as

$$\mathbf{Swish}(x) = \frac{x}{e^{-x} + 1} \quad (5.1)$$

In predicting all beam doses, the total dose was acquired automatically by summation. The loss function of BD-CNN (L_{BD}) was the sum of two parts: beam dose (BD) error and total dose (TD) error in the ROI, which is formulated as

$$L_{BD} = \frac{1}{N(\mathbf{ROI})} \left[\sum_{\text{beam}} \sum_{\mathbf{ROI}} (\mathbf{BD}_{\text{bench}} - \mathbf{BD}_{\text{pred}})^2 + \mu \cdot \sum_{\mathbf{ROI}} (\mathbf{TD}_{\text{bench}} - \mathbf{TD}_{\text{pred}})^2 \right] \quad (5.2)$$

$N(\mathbf{ROI})$ is the number of ROI pixels. $\mathbf{BD}_{\text{bench}}$ and $\mathbf{TD}_{\text{bench}}$ are the benchmark plan beam and total dose. $\mathbf{BD}_{\text{pred}}$ and $\mathbf{TD}_{\text{pred}}$ are the predicted beam and total dose. The beam dose and total dose error terms were summed with regularization term of μ as tuned by validation. All slices with ROI were used to predict beam dose by the BD-CNN. For each patient, all the predicted 2D dose slices were stacked together to form the predicted 3D dose distributions of a given beam.

In total, there were 3238 slices from all 85 training cases. Benchmark plan's beam dose is used as the ground truth for model training. Ten percent of the training slices were held out for validation. The BD-CNN was trained using Adam optimizer (Kingma & Ba, 2014) with a learning rate of 0.001 and early stopping with patience of 8 epochs (training terminates when validation loss does not improve for 8 epochs).

5.2.5 Fluence map prediction

The second DL model is the FM-CNN, which predicts one fluence map from each 3D beam dose. The network architecture of FM-CNN is illustrated in Figure 19. It adopts a customized U-Net shape, which includes three resolution hierarchies (96, 48, and 24 pixels).

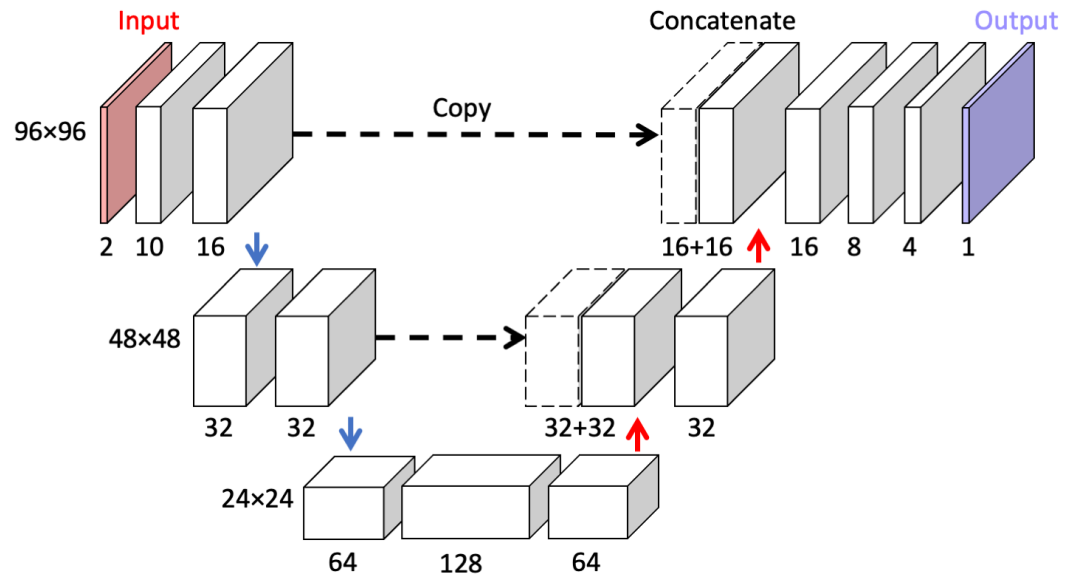


Figure 19: The network architecture of the FM-CNN in the pancreas DL framework. The number below each layer is the channel number. Each adjacent layer is connected by a convolution from left to right. Blue arrows indicate the downsampling operation, while red arrows indicate the upsampling operation. Layers

are copied and concatenated (dashed arrows) from the downsampling side (left) to the upsampling side (right).

The inputs of FM-CNN are the BEV dose map and the BEV PTV map, and the output is the fluence map. For one beam, the BEV dose map is the projection of the predicted field dose along the BEV, and the BEV PTV map is the binary projection of the PTV contour along the BEV. The upsampling and downsampling were achieved with strided 2D convolution and strided 2D transposed convolution, respectively. The BEV dose maps and fluence maps of the benchmark plans serve as ground truth for model training. The loss function of FM-CNN (L_{FM}) is a modified mean absolute error (MAE), which is formulated as

$$L_{FM} = (1 + \lambda) \frac{\sum |y_{bench} - y_{pred}|}{N(y_{bench} > 0)} \quad (5.3)$$

y_{bench} and y_{pred} are the benchmark and predicted values of the fluence map, and $N(y_{bench} > 0)$ is the count of benchmark fluence map pixels with non-zero values. The factor λ is the regularization term to prevent FM-CNN from over- or underestimating the fluence maps overall. It is expressed as

$$\lambda = \frac{|N(y_{bench} - y_{pred} > y_{th}) - N(y_{bench} - y_{pred} < -y_{th})|}{N(y_{bench} > 0)} \quad (5.4)$$

Since fluence intensity is directly linked to field dose, the fluence prediction error should have a mean value close to zero in order to avoid overdosing and underdosing. The threshold value y_{th} was set to 0.001 to ignore small deviations. Therefore, this

regularization factor is added to control the mean value of prediction error for all pixels and keep the numbers of positive and negative errors at the same level.

The total training data size was 765 for 85 patients, of which 10% were held out for validation. The model was trained using Adam optimizer with a learning rate of 0.001 and early stopping with patience of 15 epochs. In the final validation step, these predicted fluence maps were subsequently imported into the TPS for leaf sequencing and dose calculation. The resulting plans are referred as AI plans; and are compared to the benchmark plans for overall performance.

5.2.6 BEV projection

BEV projection is an interim analytical step between two neural networks, which converts 3D beam doses to BEV dose maps, or 3D contours to BEV contour maps. Each beam has a BEV plane, which is defined at the isocenter and perpendicular to the beam's central axis. A point in a BEV plane represents a beamlet, and its beam dose profile could be obtained by tracing from the source to the point through the 3D field dose volume. The maximum dose value in the profile is largely determined by the fluence intensity and the SSD of the beamlet, with contributions from nearby beams' scattered dose. To correct dose falloff and tissue inhomogeneity, the beam dose profile was divided by the corresponding open beam dose profile; therefore, the normalized dose value should be approximately the same within one profile. The SSD effect was also cancelled out due to the same beam geometry in both situations. The beam dose profile

values were then averaged to obtain the BEV dose map value in this pixel. Three-dimensional PTV contours were projected onto the BEV plane to produce PTV maps in a similar manner as dose projection. The resulting PTV map is a binary mask defined in the BEV plane, which is also used as an input in the subsequent fluence map prediction step. The PTV map provided the shape/boundary of the PTV in the BEV, which could help the fluence prediction.

5.2.7 Model assessment

For model evaluation, the benchmark plan is considered as the ground truth. Each of the two models are evaluated separately and then collectively for dosimetric quality and deliverability. The BD-CNN field dose is compared with the corresponding field dose of the benchmark plan to evaluate BD-CNN's performance. To evaluate FM-CNN's performance, a special plan, the FM-CNN plan, is generated by FM-CNN using the field dose from the benchmark plan, thus eliminating error contamination from the first CNN model. The AI plan is the final plan created with the fluence map predicted by the complete model (both CNNs), and thus evaluates the overall performance the framework.

The 15 cases not included in model training were used as an independent test set, which consists of 638 slices and 135 fluence maps. For each test case, an BD-CNN field dose, an FM-CNN plan, and an AI plan were created. The voxel-wise percentage dose difference ΔD is calculated as

$$\Delta D(V) = \frac{1}{N(V)} \sum_{i \in V} \left| \frac{D_{bench}^{(i)} - D_{pred}^{(i)}}{D_{prescription}} \right| \times 100 \quad (5.5)$$

V is the calculation volume, and $N(V)$ is the number of voxels in this volume.

Several dosimetric endpoints were also used for assessment. These include PTV max dose (0.1 cc), mean dose, and $D_{95\%}$ for the PTV and mean dose and max dose (0.1 cc) for the OARs. To provide a direct assessment of fluence map prediction, MAEs were calculated between FM-CNN fluence maps and benchmark fluence maps.

In the Eclipse TPS, optimal fluence maps, generated by inverse optimization or deep learning models, are converted to actual fluence maps by leaf sequencing algorithms to enable delivery on the machine. Unrealistic optimal fluence map features, such as extremely heterogeneous regions or high transmission value at a single pixel, could potentially result in a large discrepancy between optimized dose and delivered dose. Therefore, the deliverability of fluence maps was measured by the gamma index between optimal (before leaf sequencing) and actual fluence maps (after leaf sequencing) for both benchmark plans and AI plans. We employed the gamma analysis in a similar fashion and intent as IMRT quality assurance. Here, a high gamma passing rate indicates that the optimal fluence map is physically realistic and could be achieved by the leaf sequencing algorithm. Gamma analysis was performed using an in-house program with a 3%/3 mm criterion. Total monitor units (MUs) from benchmark plans and AI plans were compared.

After the DL framework was completely trained and tested, we reduced the training cases for both CNNs and calculated the loss values on the test set. In addition, a series of ablation studies were conducted, where certain CNN components were removed to test the model performance. For the BD-CNN model, we removed the input of one, two, or three pairs of adjacent PTV slices or beam templates. For the FM-CNN model, we removed the input of PTV map. The reduced models were evaluated on the same test set and compared with the original models.

5.3 Results

5.3.1 Model training

The model training details are summarized in Table 2. BD-CNN has 3.35 million trainable parameters and took 3 hours to train. FM-CNN has a much less complex architecture with 0.20 million trainable parameters and took 4 minutes to train. The projection of field dose and PTV along the BEV is relatively time-consuming compared to CNN predictions. On average, prediction of 9 fluence maps for each patient took 7.1 seconds, including 1.10 seconds for BD-CNN prediction, 5.97 seconds for BEV projection, and 0.03 seconds for FM-CNN prediction. In the entire workflow, the computation time of the model is typically less than that of TPS dose calculation. An AI plan was generated within 1 to 2 minutes including calculating the AI plan dose in TPS, as compared to the traditional manual planning between 1 and 3 hours.

Table 2: Model training and calculation details for the DL framework. The training details of two CNNs include number of trainable parameters, training epochs, and training time. BEV projection is a deterministic process which requires no training. The calculation times listed are average prediction time of CNNs and average calculation time of BEV projection.

	Trainable parameters	Training data size	Epochs	Training time	Calculation time per image	Calculation time per patient
BD-CNN	3,351,185	3238	48	3 hr	0.026 s	1.100 s
BEV projection	n/a	n/a	n/a	n/a	0.663 s	5.966 s
FM-CNN	203,621	765	134	4 min	0.003 s	0.030 s

The model performance of BD-CNN and FM-CNN were plotted against the number of training cases used, as shown in Figure 20. It could be seen that the testing loss of BD-CNN plateaued after 55 cases, while FM-CNN required only 35 cases to achieve reasonably good performance. The ablation study showed that, for BD-CNN, removing the beam template input would increase the testing loss by 20%; using only 4, 2, and 0 adjacent PTV slices would increase the testing loss by 7%, 25%, and 66%, respectively. For FM-CNN, removing the PTV map input would only slightly increase the testing loss by 1%.

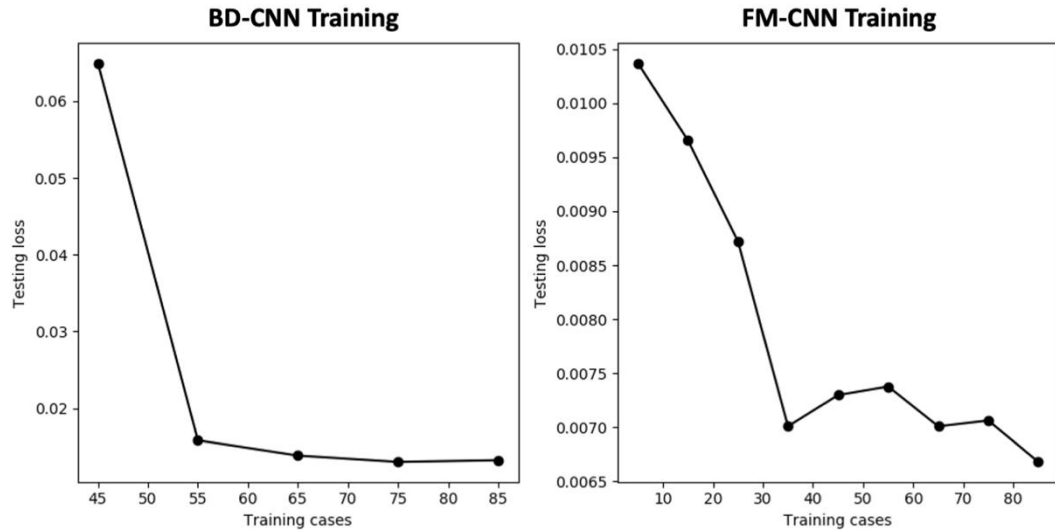


Figure 20: The number of training cases vs. testing loss for both CNNs. The testing loss stabilized when using 55 or more training cases for BD-CNN (left) and 35 or more cases for FM-CNN (right).

5.3.2 Model performance evaluation

The dosimetric evaluation results are summarized in Table 3. Here, the ground truth is the dose from the benchmark plans. AI plans have the largest dose differences among the three evaluation plans as they represent the overall performance of the workflow. In the deliverable plans, i.e., FM-CNN plans and AI plans, PTV and OAR (stomach, C-loop/duodenum, and bowels combined) maximum dose errors are larger than mean dose errors.

Table 3: Dose difference between all predicted plan groups and benchmark plans. AI plans exhibited larger dose differences than BD-CNN doses and FM-CNN plans.

Plan Type	Dose Type	Region	Voxel Dose Difference [%]	D_{mean} Difference [%]	D_{max} Difference [%]
		ROI	1.79 ± 2.21	0.41 ± 0.28	0.48 ± 0.31

BD-CNN Dose	Total Dose (CNN)	PTV	0.91 ± 0.79	0.57 ± 0.25	0.48 ± 0.31
		ROI – PTV	2.65 ± 2.75	0.51 ± 0.34	0.49 ± 0.54
	Field Dose (CNN)	ROI	1.25 ± 1.11	0.51 ± 0.42	1.82 ± 1.44
FM-CNN Plan	Total Dose (TPS)	PTV	1.22 ± 0.96	0.88 ± 0.65	1.46 ± 1.19
		OAR	0.86 ± 0.75	0.30 ± 0.17	0.86 ± 0.52
AI Plan	Total Dose (TPS)	PTV	2.41 ± 1.87	1.24 ± 0.74	4.10 ± 2.35
		OAR	2.70 ± 2.45	0.94 ± 0.65	4.77 ± 2.84

Figure 21 compares the total dose distribution between the AI plan and the benchmark plan of an example case. Figure 22 compares the DVH for the same case. As shown in the figure, the predicted fluence map achieves similar fluence modulation as the fluence map of the benchmark plan. Further, the TPS-calculated dose distribution of the AI plan exhibits small differences from the corresponding benchmark plan, indicating highly similar plan quality.

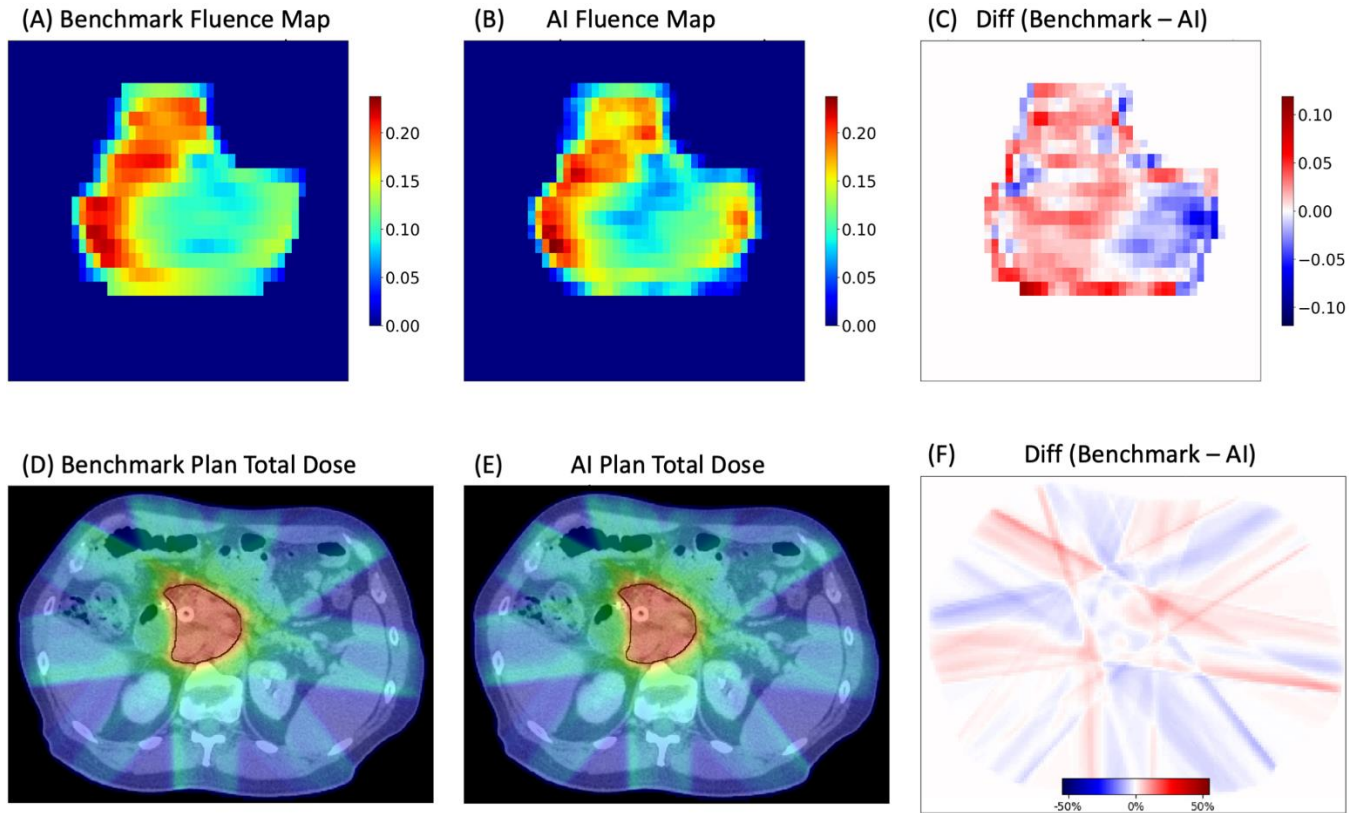


Figure 21: Examples of AI fluence map and dose comparison with benchmark in one test case. The AI fluence map recreated the fluence contrast in the benchmark. The AI plan achieved similar total dose as the benchmark. The first row shows the benchmark (A) and AI (B) fluence maps of one beam, and the difference (C). The second row shows one axial slice of the dose distribution of benchmark plan (D) and AI plan (E), and the dose difference (F). The PTV contour is marked with black lines in (D) and (E).

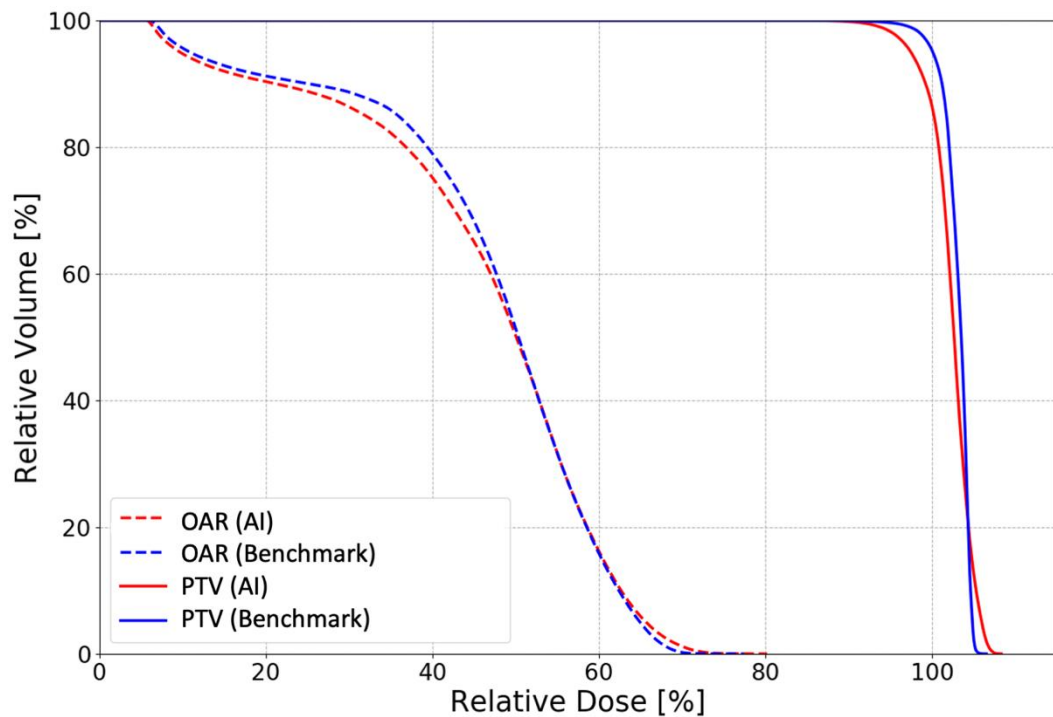


Figure 22: The example of PTV (solid) and OAR (dashed) DVH comparison in one test case between benchmark plan (blue) and AI plan (red). The benchmark plan has slightly better PTV homogeneity than the AI plan.

The test set distributions of PTV and OAR dose metrics of benchmark plans and AI plans are plotted in Figure 23. As shown in the plots, the AI plans had higher PTV maximum dose and OAR maximum dose than the benchmark plans. The PTV and OAR mean doses in the two plan groups were similar, while the PTV $D_{95\%}$ was lower in the AI plans. The overall quality of the AI plans was slightly inferior to the benchmark plans, which was expected considering the prediction errors in the two CNNs.

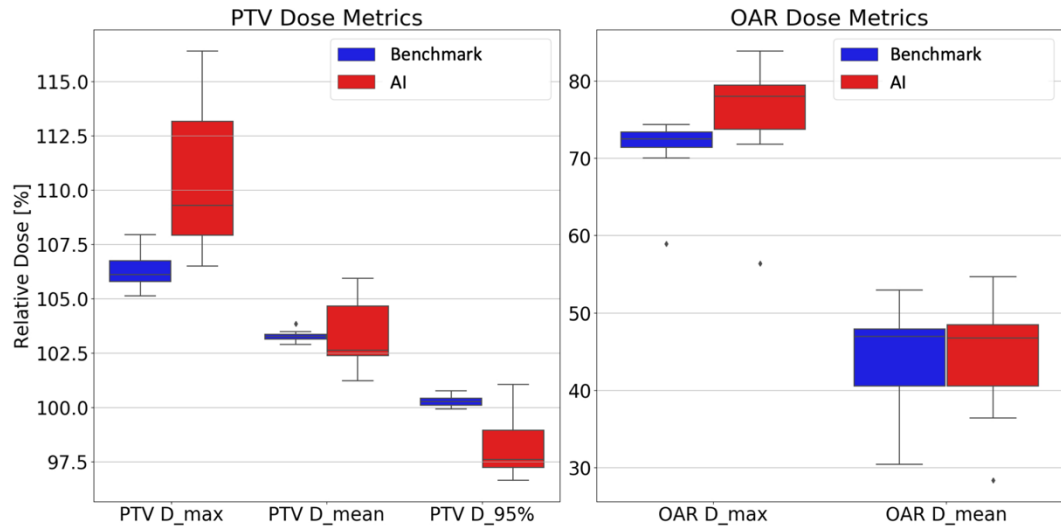


Figure 23: Test set distributions of PTV (left) and OAR (right) dose metrics comparing benchmark plans and AI plans. AI plans have higher PTV and OAR maximum dose and lower D_95% than the other plan groups. Dose values are reported as percentage of the prescription dose.

In terms of fluence map deliverability, the average \pm standard deviation of gamma passing rate was $98.14\% \pm 0.74\%$ for benchmark plans and $97.69\% \pm 0.96\%$ for AI plans, respectively, which demonstrates highly similar deliverability. The average \pm standard deviation of total monitor units (MU) per patient is 2265 ± 373 in benchmark plans and 2122 ± 281 in AI plans.

5.4 Discussion

We developed a novel deep learning framework to generate clinical-quality pancreas SBRT plans in seconds. It offers the advantage of bypassing the lengthy optimization during which the planner needs to adjust optimization objectives and aims to achieve similar performance as the human expert exercising inverse optimization. This study demonstrates the novel approach of AI-driven treatment planning via

predicting fluence maps, thus providing a more complete approach to generating deliverable high-quality plans, which has not been sufficiently addressed in previous studies (Ana M Barragán - Montero et al., 2019; X. Chen, Men, Li, Yi, & Dai, 2019; Kearney et al., 2018; Liu et al., 2015; Nguyen, Jia, et al., 2019; Nguyen, Long, et al., 2019; Skarpman Munter & Sjolund, 2015).

The aim of the proposed DL solution is to learn the knowledge from existing plans and generate deliverable plans for new patients, which falls under the broad KBP vision. Our approach directly predicts fluence maps rather than predicting achievable DVH/dose in other KBP approaches. More specifically, we used CNNs to establish the correlation between patient anatomy patterns and individual beam's dose/fluence map, which has not been investigated in previous KBP studies. This approach is built upon beamlet-based fluence optimization, where a subsequent leaf sequencing process converts the fluence maps to MLC motion parameters. By replacing the FM-CNN, the proposed approach could also be employed along with direct aperture optimization (DAO) (Shepard, Earl, Li, Naqvi, & Yu, 2002) in step-and-shoot IMRT, which would offer the advantage of fewer segments and MUs. This is a potential area of study that warrants future effort. It would also be of interest to compare the proposed approach with other KBP-based plan generation methods in future studies.

We redesigned the entire radiation therapy treatment planning workflow by incorporating DL models for dose prediction and fluence map generation, thereby

completing an AI-driven plan generation in seconds. Our results demonstrate that such AI-driven plans have a similar quality compared to manually generated inversely optimized plans, while more importantly a ready-to-deliver plan is generated with no further human intervention needed. In dose prediction, the total dose in PTV predicted by BD-CNN achieved similar level of accuracy as existing deep learning-based dose prediction models. The input of adjacent PTV slices provided the superior-inferior contour change information efficiently, which significantly reduced the testing loss while maintaining lower memory consumption than 3D networks. The downsampling block in the BD-CNN extracts feature information at different resolution levels from the input contours. As the overall image size decreases, more global features are extracted since the convolution kernel size remains the same. The upsampling blocks produce the beam doses from the extracted features and restore the image resolution. The parallel upsampling blocks were designed to produce dose distributions from different gantry angles, and the final convolutional layers combined the information from all angles to generate the predicted beam doses. The beam templates, i.e., open beam doses, provided the information of photon interaction with the patient body, which depended on the beam data and the patient anatomy. Since the information is specific for each beam angle, they were introduced into the network before the final convolutional layers. Therefore, it is feasible to train models for different beam energies. With the appropriate training data, we can potentially train one model that works for multiple beam energies,

which increases the robustness of the framework. The second step of our framework, i.e., fluence map prediction from existing field dose, directly converts individual field dose to its corresponding fluence map, which eliminates interplay among beams and requires no optimization or intermediate dose calculation. The U-Net architecture used by the FM-CNN was inspired by its successful application in organ segmentation and dose prediction, where the overall shapes of the input and output images are similar. In this case, the input (dose map) is a blurred version of the output (fluence map) because of the photon interactions (scattering and leakage) in the patient body. The fluence maps also have a distinct pattern due to the MLC leaf size (5 mm for the central leaves) and the fluence pixel size (2.5 mm). Therefore, the FM-CNN restored the fluence pattern and reversed the dose deposition process. With existing ground truth field dose, we achieved similar fluence map MAE (mean value: 2.06×10^{-3}) as H. Lee et al. (2019) (median value: 9.95×10^{-4}). With similar fluence map prediction accuracy, our proposed framework is capable of directly predicting fluence map from contour and CT alone, which H. Lee et al. (2019) has yet to achieve.

We used standardized 9-beam IMRT plans as benchmark in this study, which increased consistency in plan quality and reduced the need for a large amount of training data. The training data generated by human experts is optimal in terms of the endpoints of target coverage and luminal structure maximum dose. One limitation of the model is that the training and testing cases must have the same beam arrangement,

dose prescription level, and physician preferences. The models in the following sections will address these limitations.

5.5 Conclusion

We developed a deep learning framework utilizing two CNNs to directly generate clinical-quality IMRT plan from CT images and contours for pancreas SBRT. This framework changes the traditional approach of inverse treatment planning by replacing the inverse optimization engine with the intelligent neural networks. The proposed method has great potential to improve clinical efficiency and plan quality consistency for challenging treatment sites.

6. Fluence map prediction for pancreas SIB-SBRT

6.1 Introduction

In pancreas SBRT planning, a higher prescription dose can be concurrently delivered to the gross tumor volume (GTV) within the PTV, i.e., simultaneous integrated boost (SIB) (Brown et al., 2006; Koay et al., 2020; Yang et al., 2015). The reporting of promising local control and moderate toxicity rates has encouraged further dose escalation approaches to improve treatment outcomes (Brunner, Nestle, Grosu, & Partridge, 2015; Holmlund et al., 2019).

However, the complex anatomy of multiple PTVs and OARs, including the stomach and the duodenum/bowel, poses a significant challenge. When utilizing SIB techniques, the boost target (GTV) dose usually exceeds the dose limit of GI luminal OARs. Although achieving this high dose gradient is a major objective in planning, a trade-off is usually required to deliver as much high dose to the target as possible while respecting normal organ dose constraints. This challenging task generally requires extensive and superior treatment planning skills to achieve the desired goals and adds significant time and cost to the planning process. Therefore, plan quality can be highly dependent on the planner's experience and time resources available (Nelms et al., 2012).

We have previously proposed a deep learning framework to predict fluence maps for pancreas SBRT with a single PTV (W. Wang et al., 2020). In this study, we propose using the modified deep learning framework to predict fluence maps for plans

with multiple PTV prescriptions. We hypothesize that, given standardized beam orientation and prescription doses, deep neural networks can be trained to predict the optimal fluence maps directly from patient anatomy and thus lead to high quality plans without inverse optimization. We have applied this DL framework in automated treatment planning approaches for pancreas SBRT to evaluate its feasibility in challenging SIB scenarios.

This work has been published in *Advances in Radiation Oncology* (W. Wang et al., 2021).

6.2 Materials and methods

6.2.1 Materials

One hundred pancreatic cancer patients previously treated with SBRT at Duke University Medical Center were randomly selected for this retrospective study following institutional review board approval. Eighty cases were randomly selected for training, with the remaining 20 cases for testing. Since SBRT is a rapidly evolving treatment modality for pancreatic cancer, the prescriptions to the boost target volume, the dose limits to the GI structures, as well as the treatment beam setting, have varied over time. Therefore, in this study, a set of benchmark plans were designed by clinical physicists specialized in GI treatment using the unified prescription template shown in Table 4, partially based on a prior phase II study utilizing 5 fractions (Herman et al., 2015). The

patient statistics, including the PTV25 volume, PTV33 volume, and the benchmark total MU, are listed in Table 5.

Table 4: Clinical protocol used for generating the benchmark plans.

Planning structure	Structure name	Prescription
Elective target volume	PTV25	25 Gy to >95% volume
Boost target volume	PTV33	33 Gy ideally to >95% volume Yield to GI max dose limit
Duodenum	OAR	Maximum dose (0.1 cc) < 29 Gy
Stomach		
Bowels		
Bilateral kidney	Lt Kidney & Rt Kidney	$V_{15Gy} < 15\%$
Liver	Liver	$V_{15Gy} < 10\%$

Table 5: The patient statistics in the training set and the testing set.

Statistic	Patient Group	Average	Range	Standard Deviation
PTV25 Volume (cc)	Training	323.3	32.6 – 1134.1	226.2
	Testing	234.4	37.1 – 550.0	158.6
PTV33 Volume (cc)	Training	51.6	1.1 – 178.5	41.3
	Testing	44.4	11.35 – 180.6	39.5
Benchmark Total MU	Training	1668.2	1219.2 – 2494.3	274.3
	Testing	1480.3	1260.7 – 2081.6	182.0

All benchmark plans were IMRT plans generated in the Eclipse® treatment planning system (Version 15.6, Varian Medical Systems, Palo Alto, CA) using nine equally spaced beams and created for Varian TrueBeam® Linear Accelerator with Millennium 120 multi-leaf collimator. All plans were reviewed and deemed clinically acceptable before they were used for training and validation in this study.

Resampling was performed on dose and contour images (axial slices) with 1 mm axial resolution and 1 mm slice thickness using linear interpolation. The BEV images retained the native resolution of 2.5 mm in Eclipse. The networks were trained on a server with an Intel Xeon W-2195 processor and 256 GB of RAM, using one Nvidia Quadro RTX 8000 graphics card. Model testing was performed on a workstation with an Intel Xeon E5 v4 processor, 64 GB of RAM, and a Nvidia Quadro M4000 graphics card.

6.2.2 Deep learning framework

The proposed DL framework aims to eliminate the inverse optimization process, as illustrated in Figure 24. In manual planning, the optimization engine produces optimal fluence maps with many iterations to minimize the objective function. In contrast, the DL framework directly predicts fluence maps with two CNNs, fully replacing the inverse optimization process. The prescription information is added as input compared to the previous version. The predicted fluence maps are then sent to the TPS to finalize the plan. Although the two CNNs were trained and validated independently, they were integrated as a whole framework to replace the manual iterative inverse planning process.




Figure 24: The proposed DL framework for fluence map prediction compared with manual planning. The 9-beam IMRT benchmark plans are generated manually with the traditional inverse planning workflow. In the DL framework, the BD-CNN predicts beam dose from the anatomy and prescription. The predicted beam dose is used as the input for FM-CNN to predict the fluence map. Both benchmark plans (manual) and AI plans (DL) are finalized in the TPS using nine fluence maps.

The beam dose CNN (BD-CNN) predicts the 3D dose contribution of each beam. This component of the DL framework takes the patient anatomy as input and predicts as output the dose contribution of each beam, which we refer to as *Predicted Beam Dose*. All nine beam doses are predicted simultaneously. In the next step, these 3D predicted beam doses are projected along the beam’s eye view and converted to 2D beam dose maps. Additionally, two 3D PTV volumes (primary and boost) are similarly converted to 2D PTV maps as the other input. In the final step, the 2D beam dose map and the two 2D PTV maps associated with each beam are then taken together as the input to the fluence map CNN (FM-CNN) to predict this beam’s fluence map. A single FM-CNN is used to separately predict fluence maps for different beam angles.

Just like the fluence maps generated from standard inverse planning process, the predicted fluence maps were imported into the Eclipse TPS to create a deliverable plan via leaf sequencing (Smart LMC version 15.6.03) and dose calculation (Analytical Anisotropic Algorithm version 15.6.03). The import and calculation steps in the TPS were executed by automated scripts.

6.2.3 Individual beam dose prediction

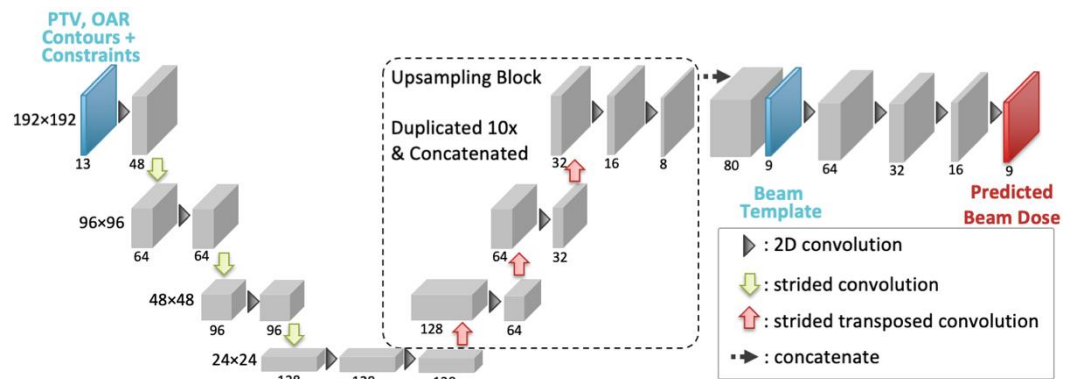


Figure 25: The network architecture of the BD-CNN in the DL framework for pancreas SIB-SBRT.

The BD-CNN is a convolutional neural network with an encoder-decoder structure, as shown in Figure 25. The PTV and OAR contours are converted to masks and partitioned to their corresponding dose constraints. The combined contour mask is fed to the CNN network in 13 consecutive axial slices. As the slice thickness is 1 mm, the field of view in the Z direction is 13 mm. Compared with the previous network architecture, the OAR and PTV contours are included in the same image, increasing the OAR input information from only the center slice to all slices. The number of

upsampling blocks is increased by one, providing greater model complexity. The slice thickness is halved, increasing the resolution in the superior-inferior direction. With twice the number of slices as input, the physical field of view stays the same, while providing greater structure contour details especially on the edges.

Downsampling and upsampling are achieved with strided convolutional layers and strided transposed convolutional layers, respectively. The resolution of the input layer is 192×192 . This network extracts information from the high-resolution contour masks into low-resolution features (encoding) and reconstructs the high-resolution beam doses (decoding). The downsampling block produces 128 channels with a resolution of 24×24 . The upsampling block is duplicated in parallel for 10 times, each with a different set of weights and producing 8 channels with the 192×192 resolution. The outputs of the 10 upsampling blocks are stacked with the 9-channel beam templates. Finally, four convolutional layers reduce the channel number to 9, corresponding to 9 beam doses.

Nine beam dose distributions for the central slice are generated, one in each output channel. This CNN design allows the independent prediction of each axial slice while incorporating contour variation in the superior-inferior direction. This process was repeated for all PTV slices iteratively; then, the predicted beam dose slices from the same angle were stacked to create a 3D beam dose for fluence map prediction.

The loss function of the BD-CNN is defined as a weighted sum of beam dose (BD) error and total dose (TD) error, with the hyperparameter α tuned during validation. The prediction errors are calculated in a region-of-interest (ROI), which is the PTV25 expanded by 1 cm. The loss function is expressed as

$$L_{BD} = \frac{1}{N(ROI)} \sum_{ROI} \left[(BD_{true} - BD_{pred})^2 + \alpha (TD_{true} - TD_{pred})^2 \right] \quad (6.1)$$

Of all the training data, 10% were held out as a validation set in order to fine-tune the model architecture and hyperparameters, including the loss function weight α . The model was trained with early stopping based on validation loss and a maximum of 150 epochs.

6.2.4 Fluence map prediction

The FM-CNN adopts a U-net (Ronneberger et al., 2015) shape with three resolution levels and predicts the fluence map for each beam individually, as shown in Figure 26. The input includes a beam dose map and the projected contour maps of PTV25 and PTV33. The beam dose map is each beam's predicted dose contribution (output of BD-CNN) projected along the BEV. The output is the fluence map of the corresponding beam.

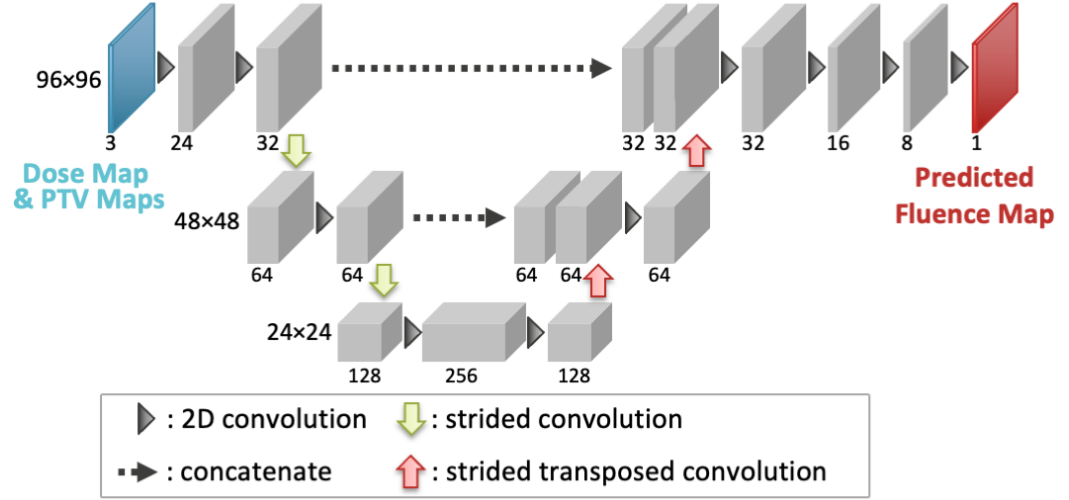


Figure 26: The network architecture of the FM-CNN in the DL framework for pancreas SIB-SBRT.

Both input and output have a dimension of 96×96 pixels and a pixel size of 2.5 mm, which is the standard fluence map resolution in our clinic. Strided convolutional layers and strided transposed convolutional layers are used for downsampling and upsampling similar to the BD-CNN. Concatenation of layers from the left to the right side of the U-net adds skip connection to the FM-CNN.

The loss function of the FM-CNN (L_{FM}) is a modified mean absolute error (MAE), which is formulated as

$$L_{FM} = (1 + \lambda) \frac{\sum |y_{true} - y_{pred}|}{N(y_{true} > 0)}, \quad (6.2)$$

where y_{true} and y_{pred} are the ground truth (benchmark) and predicted values of the fluence map, and $N(y_{true} > 0)$ is the count of ground truth pixels with non-zero values.

λ is a coefficient which prevents FM-CNN from over- or underestimating the fluence maps. It is expressed as

$$\lambda = \frac{|N(y_{true} - y_{pred} > y_{th}) - N(y_{true} - y_{pred} < -y_{th})|}{N(y_{true} > 0)}, \quad (6.3)$$

where the fluence error threshold y_{th} was selected during the validation stage.

Similar to BD-CNN training, 10% of the training data were used for validation to fine tune the model architecture and hyperparameters. The model was trained with early stopping based on validation loss and a maximum of 150 epochs.

6.2.5 DL framework evaluation

The framework is evaluated using the 20 test cases in three aspects: (1) planning time, (2) fluence map prediction, and (3) dosimetric quality of the final plans generated from the predicted fluence maps.

Each step of the DL framework was timed and averaged among all test cases.

The prediction error of each fluence map was calculated as

$$Err(F_{true}, F_{pred}) = \frac{\sum |F_{true} - F_{pred}|}{\sum F_{true}}. \quad (6.4)$$

In addition to the fluence prediction error, we also compared the similarity between the predicted fluence map and that of the benchmark plan, normalized cross correlation was calculated for each beam. A cutoff threshold of 20% of the maximum fluence value was used to define the area of interest, i.e., values below the threshold were set to 0.

Although the direct output of the DL framework is fluence maps, the quality of the final plans ultimately determines the usefulness of the model in the clinic. Thus, the predicted fluence maps were imported into the Eclipse TPS to generate the AI plan for each test case. The AI plans were compared with the benchmark plans based on clinically relevant dosimetric endpoints, including PTV33 D_{\max} , PTV33 $D_{95\%}$, PTV33 $V_{33\text{Gy}}$, PTV33 D_{mean} , PTV25 $V_{25\text{Gy}}$, PTV25-33 D_{mean} , OAR D_{\max} , OAR $D_{1\text{cc}}$, and OAR $D_{2\text{cc}}$. The dice coefficients between prescription (25 or 33 Gy) isodose and PTV contours were calculated to represent the dose conformity of the benchmark and AI plans.

6.2.6 Plan evaluation by physician

Finally, the AI plans were evaluated and compared with the corresponding benchmark plans by an attending physician specializing in GI cancers. All plans were directly generated from DL prediction and no renormalization was performed. The AI plans were assigned one of four grades (A, B, C, D) by the physician (Cox, Teckie, Kapur, Chou, & Potters, 2020). Plans with grade A or B were considered clinically acceptable, while plans with grade C or D were renormalized and re-evaluated. The grades served as a quantitative measure of clinical acceptability of the AI plans. The pooled grade point average (GPA) of all test cases were used to measure the overall performance of the proposed framework.

The physician rated the AI plans with letter grades as follows: A (4.0) for plans with overall quality highly comparable to benchmark; B (3.0) for plans with overall

quality slightly worse but acceptable for treatment. Frequently, the inferior aspects of the AI plans are very small in numerical values and has little to no impact on the DVH results, such as at sharp corners; C (2.0) for plans with overall quality of grade A and B but within 1% of meeting the PTV/OAR criteria, e.g., PTV $V_{100\%}$ within 94% to 94.9%; D (1.0) for plans that require large modifications for treatment, or with overall quality of grade A and B but over 1% of meeting the PTV/OAR criteria.

In general, when plans had slight undercoverage to the PTVs (94%-94.9%), a simple renormalization (less than 1% scaling adjustment) would bring the plan to meet prescription without violating OAR prescription limits. However, if the renormalization scale was somewhat larger, the risk of OAR exceeding prescription limits becomes higher, hence those plans would be graded D even if the overall dose distributions are highly conformal.

Once the physician completed grading, a renormalization was performed on eligible grade C and D plans, which were re-evaluated for potential upgrading. Five grade C plans and one grade D plan were improved to grade B after renormalization with increased PTV coverage and acceptable OAR sparing. One grade C plan did not change grade because renormalization was not feasible (i.e., would exceed the OAR limit). The rationale is that, even though the renormalization is an extra step outside the DL framework, it is a very simple process to be automated and integrated into the

workflow. Furthermore, renormalization is always one of the options to adjust the overall dose distribution and plan quality during the clinical planning process.

6.3 Results

6.3.1 Model training and prediction

For the BD-CNN, 7718 slices from 80 cases were used for training. The BD-CNN has the capacity of 4.4 million trainable parameters, and the training process took approximately 6 hours to finish 100 epochs with early stopping to avoid overtraining. The loss function weight α was 0.1. Compared with the BD-CNN, the FM-CNN has fewer trainable parameters (0.8 million). There were 720 fluence maps (9 each from 80 cases) for the FM-CNN training. The training process took 5 minutes to finish 75 epochs with early stopping as well. The fluence error threshold y_{th} was 0.005.

After model training, the DL framework takes little time to generate plans for new cases. Figure 3 illustrates the average time taken by the major steps of the proposed. For the 20 test cases, the PTV25 volume is an excellent indicator of the input size (average: 234.4 cc, range: 37.1 cc – 550.0 cc, standard deviation: 162.7 cc). The total planning time is 107.2 s on average (range: 78.2 s – 142.6 s, standard deviation: 18.3 s). Since all 3D planning elements were upsampled to $1 \times 1 \times 1$ mm³ voxel size, data preprocessing took the longest time (average: 51.7 seconds, range: 29.6 s – 94.5 s, standard deviation: 16.1 s) in the workflow. On average, individual beam's dose projection to the BEV plane took 10.8 seconds per patient (range: 3.5 s – 19.5 s, standard

deviation: 5.6 s), and the combined prediction time of the two CNNs was only 4.73 seconds (range: 3.2 s – 7.2 s, standard deviation: 1.3 s). The plan finalization step in the TPS took 40 seconds per patient, which includes leaf sequencing and final dose calculation. This step also takes place in the manual planning workflow (after inverse optimization) and consumes the same amount of time. In total, a deliverable plan could be generated within 2 minutes, which is significantly faster than manual planning.

6.3.2 Evaluation of predicted fluence maps

The predicted fluence maps had an error (Eq. 4) of $4.0\% \pm 1.0\%$ (mean \pm standard deviation) relative to the mean values of benchmark fluence maps. The normalized cross correlation between the benchmark fluence map and predicted fluence map was 0.949 ± 0.022 (mean \pm standard deviation), with the range (0.878, 0.992). Figure 27 shows the predicted fluence maps side-by-side with the benchmark fluence maps from three randomly selected test cases. The results highlight the high similarity of the fluence map pattern between the predicted and benchmark fluence maps.

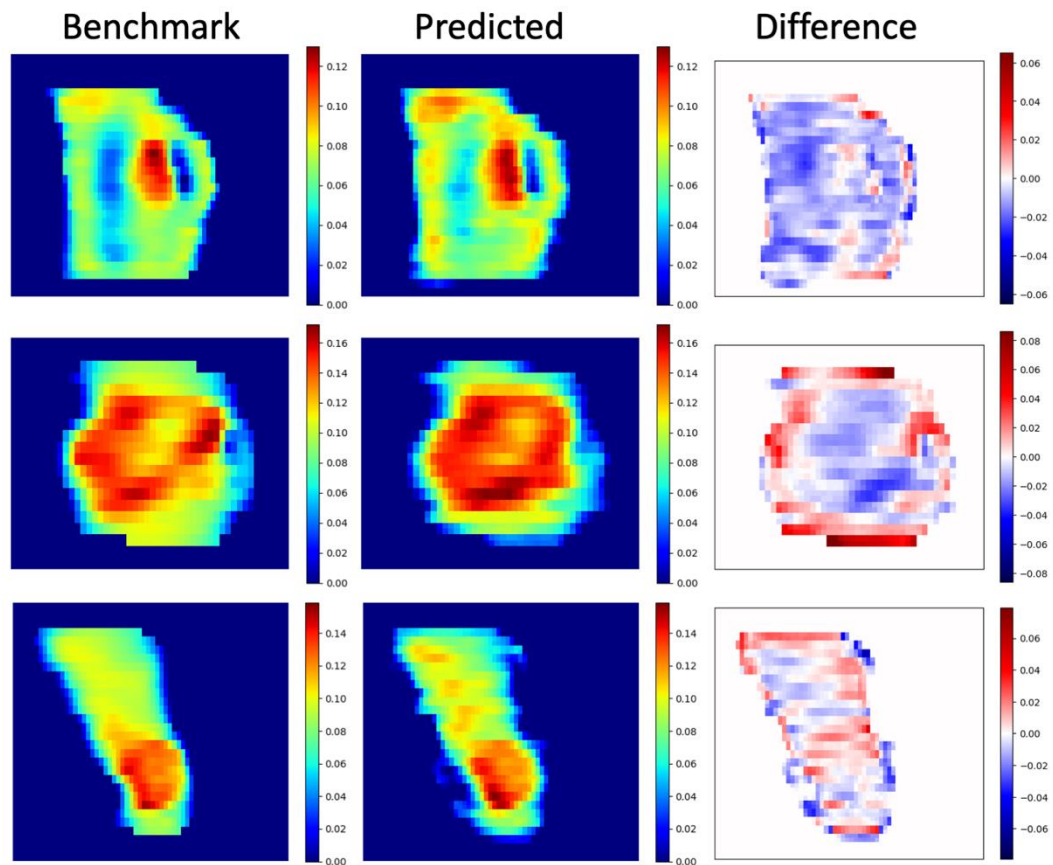


Figure 27: Fluence map comparisons (left column: benchmark, center column: predicted, right column: difference (benchmark - predicted)). The fluence map pairs are randomly selected from 3 of the 20 test cases. Each pair uses the same color map. The predicted fluence maps exhibit similar patterns as the benchmark fluence maps, especially in high fluence regions.

6.3.3 Evaluation of AI plan quality

To illustrate the AI plans' dosimetric deviations from the benchmark plans and the general trend of AI plan quality, we show in Figure 5 probability density plots of the deviations measured in the 20 test plans for each dosimetric endpoint. A plot with a narrower peak, such as PTV25-33 D_{mean} and PTV25 $V_{25\text{Gy}}$, corresponds to smaller variation between the AI and benchmark plans. The AI plans' PTV33 and OAR

maximum doses are 1.03 Gy and 0.95 Gy higher than those in the benchmark plans on average, while the other dose metrics have relatively smaller deviations. The mean deviation values are denoted on the plots as vertical dashed lines. Bilateral kidneys and liver dose were all well below protocol constraint for both plan groups.




Figure 28: The probability density plots for dose metric deviations of AI plans from benchmark plans. As all dose metrics are relative values of dose or volume, the deviation values (X axis) are in percentage differences with the benchmark plans. The probability density plots display the deviation distributions for all 20 test cases. All plots have the same scale in Y axis, which denotes the relative likelihood of the deviation. For each dose metric, the mean deviation value is denoted by the dashed line.

The total monitor units (mean \pm standard deviation) of the AI plans (1480.3 ± 182.0) and benchmark plans (1529.4 ± 235.6) are comparable, suggesting similar fluence

map complexity. The dice coefficients between the 25 Gy isodose line and PTV25 contour are 0.924 ± 0.061 (benchmark) and 0.918 ± 0.058 (AI) (paired t-test $p=0.08$). The dice coefficients between the 33 Gy isodose line and PTV33 contour are 0.902 ± 0.051 (benchmark) and 0.876 ± 0.050 (AI) (paired t-test $p<0.01$).

In the initial evaluation of the AI plans, the physician assigned 8 plans the grade of A; 5 grade B; 6 grade C; and 1 grade D. After renormalization, 19 out of 20 AI plans were deemed acceptable by the physician for clinical treatment, with the overall GPA improved from 3.0 to 3.35.

6.3.4 Example cases

The case shown in Figure 29 received grade A. The PTV25 volume is 83.2 cc. The dice coefficient between the 25 Gy isodose (cyan) and the PTV25 (magenta) is 0.93 for the AI plan and 0.97 for the benchmark plan. The PTV33 volume is 26.2 cc. The dice coefficient between the 33 Gy isodose (yellow) and the PTV33 (red) is 0.80 for the AI plan and 0.84 for the benchmark plan. All dosimetric evaluation endpoints are met, and the dose distribution is highly conformal. Given the near identical plan quality, the physician graded the AI plan A.




Figure 29: One test case received grade A by the physician. The isodose lines are drawn at 33 Gy (yellow), 29 Gy (blue), and 25 Gy (cyan). The structures are displayed as segments for PTV33 (red), PTV25 (magenta), and GI OAR (green). The two plans demonstrate no clinically significant dose differences.

The case shown in Figure 30 received grade B. The PTV25 volume is 391.1 cc. The dice coefficient between the 25 Gy isodose (cyan) and the PTV25 (magenta) is 0.87 for the AI plan and 0.86 for the benchmark plan. The PTV33 volume is 16.6 cc. The dice coefficient between the 33 Gy isodose (yellow) and the PTV33 (red) is 0.90 for the AI plan and 0.92 for the benchmark plan. All the DVH constraints are still satisfied, even at sharp OAR corners (pointed by the blue arrow). However, there is visibly more undercoverage of sharp PTV corners (pointed by the red arrow) in the AI plan. Manual

planning can cover sharp edges and corner better than DL, which also leads to improved dice coefficient numbers (especially for small PTV33 volumes).




Figure 30: One test case received grade B by the physician. The isodose lines are drawn at 33 Gy (yellow), 29 Gy (blue), and 25 Gy (cyan). The structures are displayed as segments for PTV33 (red), PTV25 (magenta), and GI OAR (green). All the DVH constraints are still satisfied, even at sharp OAR corners (pointed by the blue arrow). However, there is visibly more undercoverage of sharp PTV corners (pointed by the red arrow) in the AI plan. Thus, the physician graded the AI plan B.

The case shown in Figure 31 received grade B, which was upgraded from C after renormalization. Since the normalization factor (99.7%) is very close to 1, only the renormalized AI plan is shown. The PTV25 volume is 257.4 cc. The dice coefficient between the 25 Gy isodose (cyan) and the PTV25 (magenta) is 0.95 for the AI plan and

0.95 for the benchmark plan. The PTV33 volume is 92.3 cc. The dice coefficient between the 33 Gy isodose (yellow) and the PTV33 (red) is 0.92 for the AI plan and 0.94 for the benchmark plan. The original AI plan's PTV33 $V_{33\text{Gy}}$ was 94.4%, which was improved to 95% after renormalization. The renormalized plan satisfied all dose constraints but has slightly worse PTV coverage than the benchmark plan. Thus, the physician graded B for the renormalized plan.




Figure 31: One test case received grade B after renormalization (original grade C). The isodose lines are drawn at 33 Gy (yellow), 29 Gy (blue), and 25 Gy (cyan). The structures are displayed as segments for PTV33 (red), PTV25 (magenta), and GI OAR (green). The normalization factor was 99.7%. The renormalized plan (left) satisfied all dose constraints but has slightly worse PTV coverage than the benchmark plan (pointed by the red arrows). Thus, the physician graded B for the renormalized plan.

The case shown in Figure 32 received grade D for the original AI plan (left column) and grade B after renormalization (middle column). The normalization factor is 99.1%. The PTV25 volume is 460.4 cc. The dice coefficient between the 25 Gy isodose (cyan) and the PTV25 (magenta) is 0.96 for the AI plan and 0.96 for the benchmark plan. The PTV33 volume is 70.2 cc. The dice coefficient between the 33 Gy isodose (yellow) and the PTV33 (red) is 0.89 for the AI plan and 0.91 for the benchmark plan. The original AI plan's PTV33 V_{33Gy} was 91.6%, which was improved to 95% after renormalization. The renormalized plan satisfied all dose constraints but has slightly worse PTV coverage than the benchmark plan. Thus, the physician graded B for the renormalized plan.




Figure 32: One test case received grade D by the physician due to inadequate PTV33 coverage (left column). The renormalized plan (middle column, normalization factor 99.1%) has a PTV33 V_{33Gy} of 95%. The physician determined that the renormalized plan was slightly inferior to the benchmark plan (right column) in terms of PTV coverage (pointed by the red arrows), thus grading it B. The OAR constraints were satisfied in all plans.

6.4 Discussion

In this study, we have investigated a deep learning framework designed for pancreas IMRT planning with simultaneous boost targets. This framework transforms 3D anatomical images to 2D fluence maps, a process traditionally solved by the time-consuming inverse optimization. We split the problem into two independent tasks, each solved by a deep neural network. The BD-CNN operates in the 3D anatomy/dose space, while the FM-CNN operates in the 2D beam's eye view/fluence space. Combining the two networks, our framework can produce a deliverable plan in 2 minutes, allowing fast plan generation without inverse optimization. This plan generation process is also free from human intervention, which avoids inconsistency in plan quality due to planner experience. With more optimized code and higher computation power, the time for data preprocessing and dose projection could be reduced even more, offering the potential of real time automated planning for clinical practice.

The dice coefficients suggest that both planning methods can generate prescription isodose lines highly conformal to PTV targets, with benchmark plans having slightly better PTV33 conformity than AI plans. While the overall PTV coverage and OAR sparing have achieved clinical quality, dose in certain local areas may be overlooked by the DL framework when attempting to reduce the total error during model training. This could have a significant impact on certain dose metrics such as the maximum point dose.

In this study, we are focusing on high dose conformity (i.e., dose gradient) around the PTV, especially in the PTV-OAR overlapping regions, with a L2-norm penalty function designed in the BD-CNN. Hence, large dose gradient variations are heavily penalized in model training while maximum point dose variation in the PTV is not. Such balance helps us achieve overall dosimetry quality in the AI plans. In addition, in the context of SBRT planning, higher maximum dose within the PTV (compared to standard IMRT/VMAT planning) is often acceptable with high precision imaging guidance implemented for those treatments (IGRT).

The attending physician assessed the AI plans both qualitatively and quantitatively. Qualitatively, all AI plans were considered clinically acceptable based on prevailing clinical guidelines. Quantitatively, the AI plans were evaluated in comparison to benchmark plans that adhere to more stringent criteria adopted for this study. Here, the GI constraint ($V_{29\text{Gy}} < 0.1 \text{ cc}$) used to evaluate AI plans is considerably stricter than the constraint of GI OAR $V_{29\text{Gy}} < 1 \text{ cc}$ used during early clinical implementations of pancreas SBRT or $V_{33\text{Gy}} < 1 \text{ cc}$ currently adopted by many clinical practices. The 0.1 cc dose limit is equivalent to maximum point dose, which simplifies model training, and hence adopted for this DL network design. The renormalization step is used in all manual planning. Although not necessary for AI plans, renormalization was performed on all low-grade cases using an automated script, which improved the GPA from 3.0 to 3.35. The renormalization of AI plans only scaled them by 0.15% to 1%, which was

sufficient to bring the PTV dose to meet the $V_{100\%} > 95\%$ limit without affecting other dose qualities.

Grade B is defined as “slightly worse than the benchmark plan but acceptable for treatment”, which means all dosimetric parameters met the physician’s evaluation criteria. Compared with the benchmark plans designed by expert physicists and representing highest plan quality, the AI plans fall short in certain small areas which are difficult to cover. These are usually in spots where PTV or OAR has sharp contour changes or edges. Human experts can pay extra attention to these regions during manual inverse planning, and they usually add extra manipulation specifically targeting the coverage for these corners. The AI plans for such cases received grade B because the DL framework currently cannot achieve the high standard of human experts. However, the DVHs of these plans show little differences and are deemed clinically acceptable, hence the average GPA of 3.35 represents very stringent evaluation of these plans.

The structure of the BD-CNN model restricts the beam setting to a certain number of fixed beam angles (e.g., 9 in this study). The AI plans should have the same beam settings as the training plans. While this limitation has a relatively smaller impact on pancreas SBRT planning with its naturally central disease locations, the BD-CNN’s model architecture will require modification for sites such as the lung, where variable beam angles are essential.

This study is an important step towards automating fluence map prediction and plan generation for pancreas SBRT treatment planning. Translating this technique to clinical deployment may require additional methods in place in the future to further enhance the versatility of the tool in tailoring the dose for a specific patient. Such areas may include customizable hotspot volume, adjustable dose gradient level to enhance or relax OAR dose, tradeoff toggle between PTV coverage and OAR sparing, etc. Inter-fraction contour variation should also be taken into account as the pancreas SBRT plans focus on highly conformal dose to the PTV and sharp dose gradient falloff near the OARs. These additional steps would further enhance the user experience when the tool is deployed clinically.

6.5 Conclusions

A novel deep learning framework was developed to directly predict fluence maps and has demonstrated feasibility for pancreas SIB-SBRT. The framework utilizes two CNNs to perform beam dose prediction and fluence map prediction, which bypasses the time-consuming inverse optimization process. This approach enables rapid IMRT plan generation which provides a valuable tool for a high throughput clinic.

7. Application of transfer learning in fluence map prediction for adrenal SBRT

7.1 Introduction

There has been increasing evidence and interest in the treatment of oligometastatic disease with ablative radiotherapy (Gomez et al., 2019; Palma et al., 2019). The adrenal glands are a common site of metastases from lung cancer, renal cell carcinoma, and melanoma. Treatment with stereotactic body radiotherapy (SBRT) is associated with high rates of local control with overall low toxicity (W. C. Chen et al., 2020). The proximity of gastrointestinal organs at risk, such as the duodenum, stomach and other small bowel, often limits the amount of dose that can be safely delivered to the planning target volume. Therefore, treatment planning for adrenal SBRT is complex and time-consuming. Plan quality often depends on the experience of the planner.

Compared to diagnosis and segmentation, radiotherapy plan datasets are often limited by treatment machine or modality, physician preference, and planning style. The lack of large number of high-quality plan datasets is often the limiting factor for deep learning applications in radiotherapy, as most DL models require hundreds of high-quality plans for the same site and using the same technique. SBRT is also a relatively new and rapidly evolving technology, which further limits the availability of plan data. Furthermore, even with the same DL model design, training a specific model for each disease site could be time and labor intensive. Transfer learning could potentially be a viable way to overcome the barrier of insufficient training data.

Transfer learning refers to a machine learning concept which takes advantage of a pre-trained model or knowledge from a large dataset and applies to a different yet related scenario, usually with a much smaller dataset in the latter. Knowledge learned in the source domain (large dataset) is used to train a new model in the target domain (small dataset), which is usually achieved by transferring parts of the network weights. In medical image processing, networks trained on everyday-life image datasets, such as ImageNet (Deng et al., 2009), have been used in transfer learning to detect lesions (Bar et al., 2015; Van Ginneken, Setio, Jacobs, & Ciompi, 2015) in x-ray/CT images. In automated radiotherapy treatment planning, a large plan set could potentially be used to facilitate model training for a different site or planning style with that has a relatively small dataset. For example, Kandalan et al. (2020) used transfer learning to train several U-Net models for prostate dose prediction. In this study, a source model was trained with 108 prostate VMAT plans. Four target models, including three different planning styles and one external institution, were trained with 14-29 cases per style. The target models achieved higher mean dice similarity coefficient for isodose volumes and more accurate DVH parameter predictions in their respective styles than the source model.

In this work, we modified the previously proposed deep learning framework to generate adrenal SBRT plans using transfer learning with a base model trained on pancreas SBRT cases. The motivations are two-fold: First the number of available adrenal SBRT plans for DL modeling are often limited, with about 20-30 per year within

our institution. Data limitation is often the first reason to use transfer learning. Second, both sites share similar emphasis on the sparing of gastrointestinal organs-at-risk (GI-OARs), including the stomach, duodenum, and small bowel otherwise. However, the two sites also have differences in other planning styles. Most significant ones are that adrenal SBRT cases often have different beam settings and dose constraints from pancreatic cases. Therefore, such transfer learning process will need to adjust the overall planning knowledge and acquire new beam setting features during the process. The manuscript of this project is in preparation for publication at the time of writing.

7.2 Materials and methods

7.2.1 Patient data

7.2.1.1 Pancreas SBRT patients

One hundred patient cases treated with pancreas SBRT were used to train the base model. Fifty patients had manually generated plans with a single PTV with a 33 Gy prescription dose and a GI-OAR maximum dose limit of 25 Gy. Another fifty patients had manually generated simultaneous integrated boost (SIB) plans with two PTVs with 25 Gy (primary) and 33 Gy (boost) prescription doses, and a GI-OAR maximum dose limit of 29 Gy. These standardized benchmark plans were all generated by the same expert physicist for building DL models from previous studies (W. Wang et al., 2021; W. Wang et al., 2020).

7.2.1.2 Adrenal SBRT patients

Forty-five patients treated with adrenal SBRT were retrospectively included in this study. Each patient has one or two PTV prescription levels with total prescribed dose ranging from 24 Gy to 50 Gy (31 cases: 50 Gy in 5 or 10 fractions; 6 cases: 45 Gy in 3 fractions; 4 cases: 40 Gy in 5 or 10 fractions; 2 cases: 35 Gy in 5 fractions; 2 cases: 24 Gy in 3 or 4 fractions) and varying GI OAR constraints. These cases were randomly split into 20 cases for model training, 10 cases for model validation, and 15 cases for final model testing. The detailed statistics of the 45 adrenal SBRT cases are listed in Table 6.

Table 6: Plan statistics for adrenal SBRT cases. The statistics are listed for all cases as well as the training, validation, and testing set separately. Note: since not all cases have GI-OAR constraints, the number of cases with the specific organ constraint is listed as n.

Variable	All Cases (n = 45)	Training Set (n = 20)	Validation Set (n = 10)	Testing Set (n = 15)
Date Range	2016/1 – 2020/11	2016/1 – 2020/6	2017/5 – 2020/6	2017/4 – 2020/11
Sidedness [Left/Right]	23/22	10/10	3/7	10/5
Clinical Plan Modality [VMAT/IMRT]	36/9	17/3	8/2	11/4
PTV Volume (cc) [mean ± SD]	116.9 ± 149.6	159.3 ± 199.6	104.5 ± 89.8	68.6 ± 61.8
PTV Prescription [50 Gy/under 50 Gy]	31/14	16/4	8/2	7/8
Bowel/duodenum Limit (Gy) [median ± SD (n)]	35 ± 7.4 (n = 40)	40 ± 6.2 (n = 19)	38 ± 7.8 (n = 9)	33.5 ± 7.0 (n = 12)
Stomach Limit (Gy) [median ± SD (n)]	35 ± 7.5 (n = 24)	40 ± 8.2 (n = 14)	40 ± 7.5 (n = 5)	32 ± 2.2 (n = 5)

Clinical plans used a mixed modality of VMAT and IMRT approaches. Therefore, for the 20 adrenal SBRT training cases, the benchmark 9-beam IMRT plans were generated by an experienced physicist to mimic the clinical plan dose. The benchmark planning is the same process we have used in the pancreas SBRT DL learning modeling. It helps the training dataset to reduce noise and variations in the optimization/inverse planning process. The combined pancreas and adrenal benchmark plans used for training are referred to as benchmark plans. The plans created by the DL framework are referred to as AI plans. For model validation and testing, these AI plans are directly compared to clinical plans in the validation and testing case pool.

7.2.2 Beam geometries

The pancreas SBRT plans and adrenal SBRT plans not only differ in prescription and anatomy, the two sites also differ in general beam geometries. Therefore, the transfer learning task in this study not only needs to consider dose prescription differences, but also needs to translate the learned beam-to-beam interaction to the new model. The beam settings for the pancreas and adrenal plans are shown in Figure 33.




Figure 33: Beam configurations for pancreas and adrenal plans. The large ellipse represents the body outline in an axial slice in the caudal-cranial direction. Nine equally spaced coplanar beams are used in pancreas SBRT plans. While in adrenal SBRT plans, beams are concentrated on one side of the body. The beam settings for the left adrenal (scenario 1) and right adrenal (scenario 2) are symmetrical.

7.2.2.1 Pancreas beam geometry

All pancreas benchmark plans used 9 equally spaced IMRT beams with beam angles of 20°, 60°, 100°, 140°, 180°, 220°, 260°, 300°, and 340°. Pancreas SBRT shares similar GI-OAR dose constraints as adrenal SBRT. However, the anatomical location of the pancreas is generally more central than the adrenal glands. The pancreas model uses the beam setting of nine equally spaced beams over 360 degrees. For adrenal cases, it is often more desirable to use a beam setting with gantry angles concentrated on one side of the body. Thus, the pancreas model would produce less optimal results if the same setting is used for adrenal cases.

7.2.2.2 Adrenal beam geometry

For each of the 20 training cases, a benchmark plan was created by a medical physicist using the clinical dose constraints and a set of nine IMRT beams to mimic the

clinical plan dose. For left-sided cases, the beam angles were 25°, 60°, 90°, 120°, 150°, 180°, 210°, 325°, and 345°; for right-sided cases, the beam angles were 15°, 35°, 150°, 180°, 210°, 240°, 270°, 300°, and 335°. In model training, the two scenarios were identical after mirroring the CT images and dose from left to right with respect to the isocenter. With the small number of training cases, such unified beam setting helps the DL learning model capture beam interaction features.

For the validation and test cases, no benchmark plans were made. The AI plans were directly compared with the clinical plans. The AI plans used the fixed 9 IMRT beam setting mentioned above, while the clinical plans all have similar beam orientation preference, but with some degree of variations among planners. Some clinical plans used VMAT technique, with the arc span similar to the IMRT beam coverage range. In two test cases, all gantry angles were rotated (by 60° and 90° respectively) in the AI plans to mimic the clinical beam settings and gain extra avoidance of the OARs. The predictions were achieved by rotating the input by the gantry angle offset value in the opposite direction. The gantry angle offset could offer planner some flexibility of beam arrangements, as long as the relative separation among the 9 beams is fixed.

7.2.3 Deep learning framework

The entire framework includes data preparation, BD-CNN, dose projection, and FM-CNN. The BD-CNN predicts the beam dose for all IMRT beams using the structure

contours and dose constraints as input. The FM-CNN predicts the fluence map for each IMRT beam using the projected beam dose on the beam's eye view as input.

In the current implementation based on Eclipse Scripting Application Programming Interface (Varian Medical Systems), all steps in this workflow are completely automated except for the structure selection and dose constraint input at the beginning of the data preparation step. We have developed a graphical user interface (GUI) that can be launched within the TPS to select the PTVs and OARs and specify the dose constraint for each structure. Since the OAR sparing is prioritized over PTV coverage, the PTVs are automatically cropped to avoid the OARs by a margin as planning structures. The margin size is determined by the ratio of PTV prescription and OAR dose limit, which is designed to take physical dose gradient into account and reflect common clinical practice principles. A beam template plan is generated using nine unmodulated beams with the given gantry angles, which cover the PTVs with a 1-cm margin. The selected structures and their dose constraints, along with the beam template dose, are exported and upsampled to $1\times 1\times 1$ cm³ voxel size.

The model architecture of BD-CNN has a customized encoder-decoder design (Figure 34), which consists of a downsampling block, upsampling blocks, and a convolutional output block. For the BD-CNN, each axial slice (after upsampling) generates one data sample. The input images (dimension: $192\times 192\times 26$) includes the PTV contours and OAR contours in the current slice as well as 6 slices in both the superior

and inferior directions. The contours are assigned the values of the PTV prescription dose or the OAR dose limit, all normalized by the highest prescribed PTV dose. The output images (dimension: $192 \times 192 \times 9$) are 9 beam dose distributions for the current slice. The prediction is repeated for each axial slice, and all predicted dose slices for one beam are stacked to form a 3D volume.

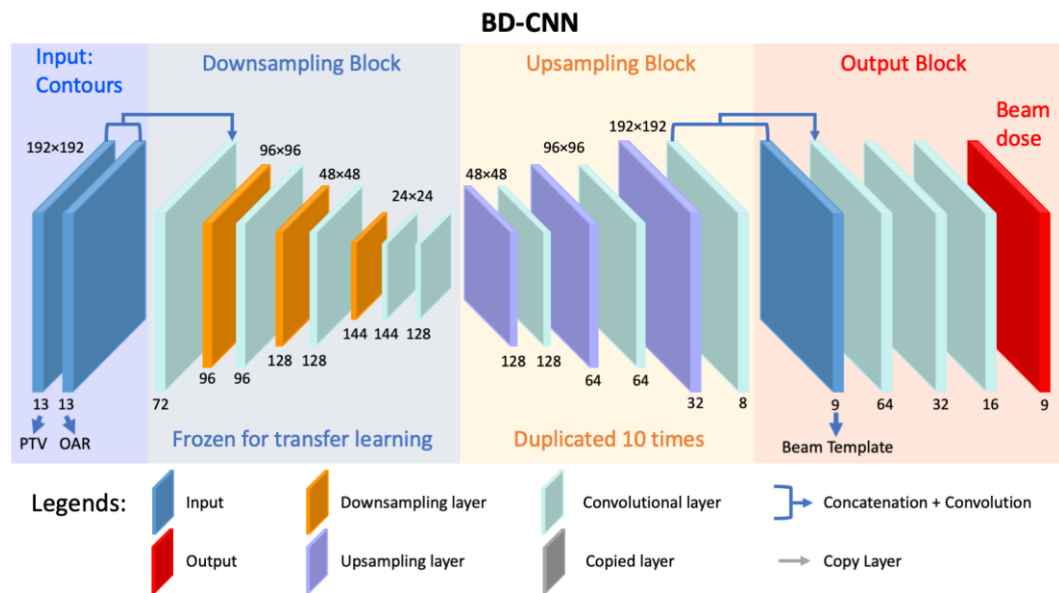


Figure 34: The neural network architecture of the BD-CNN in the adrenal DL framework. The BD-CNN predicts 9 beam doses from the contours of PTV and OAR. It is the only network retrained in transfer learning. The number of channels is labeled under each layer, and the image size is indicated above.

Each beam dose volume is projected along the beam's eye view into a 2D dose map, after taking out the effect of dose attenuation with the beam template dose. Each 2D dose map is then used as the input (dimension: $96 \times 96 \times 1$) for the FM-CNN to predict the fluence map (dimension: $96 \times 96 \times 1$). The FM-CNN uses a 3-resolution level U-Net

architecture (Figure 35). Finally, the predicted fluence maps are imported back into the TPS for leaf sequencing and dose calculation, which generates the AI plan.

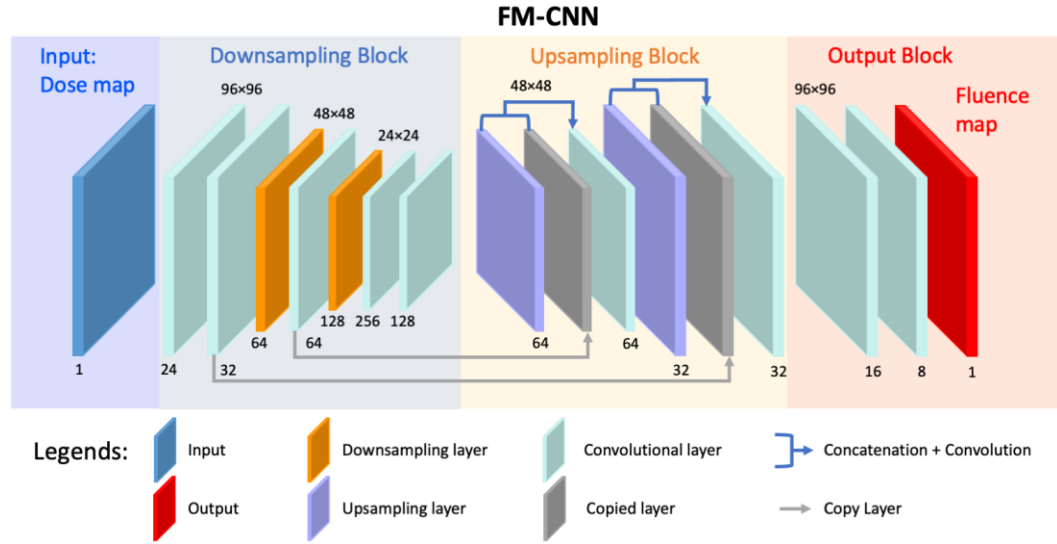


Figure 35: The neural network architecture of the FM-CNN in the adrenal DL framework. The FM-CNN predicts the fluence map from the dose map of one beam. The number of channels is labeled under each layer, and the image size is indicated above.

7.2.4 Model training and testing

Transfer learning reduces the data demand for training complex deep neural networks. To investigate this effect, we created several comparison models trained with varying number of training samples. As each complete DL framework consists of a BD-CNN and a FM-CNN, we added a descriptor in parentheses after the model name such as BD-CNN(X).

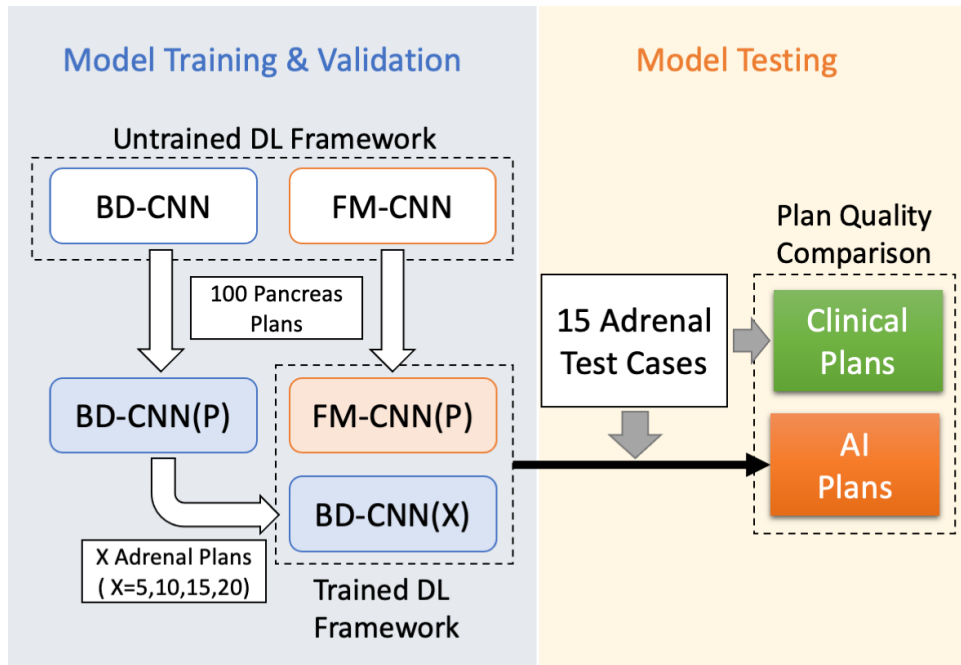


Figure 36: The fluence prediction workflow for the DL framework and model training and testing process.

The overall fluence prediction workflow and model training steps are illustrated in Figure 36. The two CNNs in the DL framework were first trained from random initialization with the 100 pancreas cases. These pancreas base models are referred to as BD-CNN(P) and FM-CNN(P), respectively. Since the FM-CNN predicts the fluence map for each beam individually, it is not site-specific and does not involve the interplay between different beams. Thus, it can be directly used for adrenal cases without further training. Although not designed for adrenal cases, BD-CNN(P) and FM-CNN(P) were tested on the validation set to provide a baseline performance.

We speculate that the downsampling block in BD-CNN is responsible for feature extraction from the patient contours, while the upsampling blocks and the convolutional

block shapes the dose distributions for all beam angles. Hence, our design of transfer learning aims to leverage the pre-trained feature extraction layers for the new model so that the training data size can be reduced. To realize this design, the BD-CNN(P) was retrained with the adrenal training set after fixing the downsampling block of the base model, i.e., the layer weights did not update during training. The model was trained with the Adam optimizer and a learning rate of 0.001. 10% of the training data were held out to calculate a validation loss after each epoch. Model training was set for a maximum of 30 epochs, and early stopping was employed by monitoring the validation loss. After this transfer learning phase of training was finished, the previously fixed downsampling layers were included in another round of training for model finetuning. The learning rate for this phase was set at 0.0001, and the number of epochs was 15.

To explore the number of cases required to train an adrenal model with transfer learning, we used 5, 10, 15, and 20 cases to train four adrenal models. The four resulting models are referred to as BD-CNN(5), BD-CNN(10), BD-CNN(15), and BD-CNN(20). These models were compared in the validation stage to determine the number of cases necessary to train a robust model. Figure 37 shows how the cases were split into training, validation, and testing.




Figure 37: The pancreas and adrenal cases used for model training, validation, and testing.

Finally, to illustrate the necessity of transfer learning with limited number of cases, BD-CNN and FM-CNN were trained with the 20 adrenal training cases from randomly initialized weights, i.e., without transfer learning. We refer to the models as BD-CNN(R) and FM-CNN(R), where R stands for random initialization. The model is evaluated together with the four transfer learning models to show the model performance without transfer learning.

7.2.5 Evaluation of model performance

Each CNN was directly evaluated by calculating the loss function on the validation data. For the BD-CNN, the loss function is defined as

$$L_{BD} = \frac{1}{N(ROI)} \sum_{ROI} \left[(BD_{true} - BD_{pred})^2 + \alpha (TD_{true} - TD_{pred})^2 \right]. \quad (7.1)$$

This loss function penalizes the beam dose (BD) and total dose (TD) mean squared error in the region-of-interest, which is defined as the PTVs expanded by 1 cm. The coefficient α in the summation was determined during model validation and kept the same for all models.

For the FM-CNN, the loss function is defined as

$$L_{FM} = \frac{\sum |y_{true} - y_{pred}|}{N(y_{true} > 0)}. \quad (7.2)$$

This is the mean absolute error for the predicted fluence map. The summed absolute error is divided by the number of non-zero fluence values in the ground truth fluence map to eliminate the effect of different field sizes.

Traditionally, deep learning model validation is conducted by calculating the loss function values of competing models on a validation dataset. However, in this study, the final product of the DL framework is the AI plan. Therefore, the framework is evaluated by the overall plan quality rather than the individual loss function values of the CNNs. To evaluate the AI plans, we designed a plan quality score based on dosimetric endpoints, which was adapted from Mistro et al. (2020). This score allowed us to evaluate plans quickly and objectively. The scoring system is illustrated in Figure 38.




Figure 38: Plan quality score design. The PTV/GTV metrics (red text) are combined to provide a PTV score. The OAR metrics (blue text) are combined to provide an OAR score. The total score is the sum of the PTV score and the OAR score.

The plan quality score takes several dose metrics into account. These metrics include PTV $V_{100\%}$, PTV $D_{98\%}$, PTV D_{min} , GTV D_{min} , maximum dose location, OAR D_{max} , OAR D_{2cc} , and OAR V_c (volume at dose constraint). Furthermore, the plan quality score can be split into a PTV score, which is the sum of first five components, and an OAR score, which is the sum of the last three components. When multiple OAR constraints are present in one case, the OAR scores are summed for all constraints. Due to differences in anatomy and dose constraints, the raw plan scores are patient-dependent and cannot be compared among different patients. Therefore, for each patient, the clinical plan score is considered the gold standard and the base score, and the AI plan

score is reported as a percentage of the clinical plan score (raw AI plan score/ raw clinical plan score \times 100). When the AI plan score approaches 100, its quality is closest to clinical plans.

7.3 Results

7.3.1 Model training and prediction times

The training time of the BD-CNN(P) was approximately 9 hours. The training time of the adrenal models were 23 minutes for BD-CNN(5), 31 minutes for BD-CNN(10), 50 minutes for BD-CNN(15), and 74 minutes for BD-CNN(20). The FM-CNN(P) was trained for 5 minutes on pancreas patients, which was then directly used in the planning workflow for adrenal cases. BD-CNN(R) and FM-CNN(R) were trained from scratch using only 20 adrenal plans. Their training times were 77 minutes and 1 minute, respectively.

Once trained, the two networks were used sequentially to predict fluence maps for the validation and test adrenal SBRT cases. The prediction times were not significantly different for different models. To generate a complete SBRT plan for one patient, the average (test set) prediction times of the BD-CNN(10) and FM-CNN(P) were 4.6 s and 0.4 s, respectively.

7.3.2 Plan quality scores in validation and testing

The plan quality scores in model validation and testing are listed in Table 7.

Table 7: The plan quality scores in validation and testing. Six different DL frameworks were used to generate plans for the validation set (10 patients). The

framework with the highest total score (10-case model) was used for the final test (15 patients). The range of each score is shown in square brackets as [minimum, maximum].

DL Framework	BD-CNN	FM-CNN	Validation Score [Min, Max]	Test Score [Min, Max]
Pancreas model	BD-CNN(P)	FM-CNN(P)	62.1 [30.0, 80.2]	-
Randomly initialized model	BD-CNN(R)	FM-CNN(R)	78.2 [58.5, 101.9]	-
5-case TL model	BD-CNN(5)	FM-CNN(P)	81.5 [61.0, 99.5]	75.3 [21.3, 101.7]
10-case TL model	BD-CNN(10)	FM-CNN(P)	89.9 [61.4, 106.5]	79.4 [33.7, 105.0]
15-case TL model	BD-CNN(15)	FM-CNN(P)	88.2 [59.9, 108.0]	82.0 [23.8, 112.1]
20-case TL model	BD-CNN(20)	FM-CNN(P)	86.7 [58.1, 106.6]	82.3 [18.5, 127.6]

In model validation, the percentage scores were averaged of all validation plans for the four TL models (in the order of increasing number of adrenal training cases shown in the parentheses), and were 81.5 (5), 89.9 (10), 88.2 (15), and 86.7 (20). The validation score peaked with the 10-case TL model and had a slight reduction with more training cases. As expected, the pancreas model performed poorly when directly used on adrenal cases, scoring the lowest of all models. More importantly, without transfer learning, the randomly initialized adrenal model, although trained with all 20 adrenal cases, achieved a lower score than even the 5-case TL model.

The TL models were used on the test set to produce the final results. There is a significant score increase of 4.1 from the 5-case model to the 10-case model, while the 15-

case model and 20-case model only achieved 2.6 and 2.9 improvement from the 10-case model. The results suggest that this transfer learning technique requires at least 10 training cases to achieve good performance for the BD-CNN. Additional cases beyond 10 achieved only marginal improvements.

7.3.3 Example adrenal SBRT cases

Dose comparisons are shown in Figure 39 and Figure 40 for two adrenal SBRT test cases. Both clinical plans in the two cases used VMAT as the beam setting, with two half arcs (Figure 39) and three full arcs (Figure 40), while the AI plans (10-case TL model) used the left adrenal 9-beam IMRT setup shown in Figure 33 (scenario 1).

In Figure 39, the test case had two PTVs with prescribed doses of 40 Gy (80%) and 50 Gy (100%). The AI plan achieved adequate dose coverage for both PTVs and similar GI-OAR sparing as the clinical plan. The clinical plan's dose was more conformal to the PTV, especially shown in the coronal view (highlighted by the red arrow); however, there were more hot spots (>110%) inside the PTV than the AI plan.




Figure 39: Dose comparison for adrenal SBRT case 1. The clinical plan dose (left column) is shown side by side with the AI plan dose (right column) in the same three views (row 1: axial view; row 2: coronal view; row 3: sagittal view). The two PTVs are prescribed to 50 Gy (100%) and 40 Gy (100%), with the GI-OAR limit of 50 Gy (100%).

In Figure 40, the test case had one large PTV with close proximity to the GI-OAR. In the coronal and sagittal views, it can be seen that the dose gradient is sharper at the superior PTV edge in the clinical plan (highlighted by the red arrows), which leads to

better PTV coverage than the AI plan. The dose coverage difference in the PTV edge was mainly due to the cropping rules in the automated planning workflow. Compared with manual planning using VMAT, the avoidance margin from OAR to PTV was larger in the AI plan. The errors in the deep learning networks limited the achievable dose gradient, and the 9-beam IMRT plans lack the modulation complexity of multiple arc VMAT plans. Therefore, the PTV cropping was more conservative in the AI plans than in the clinical plans.

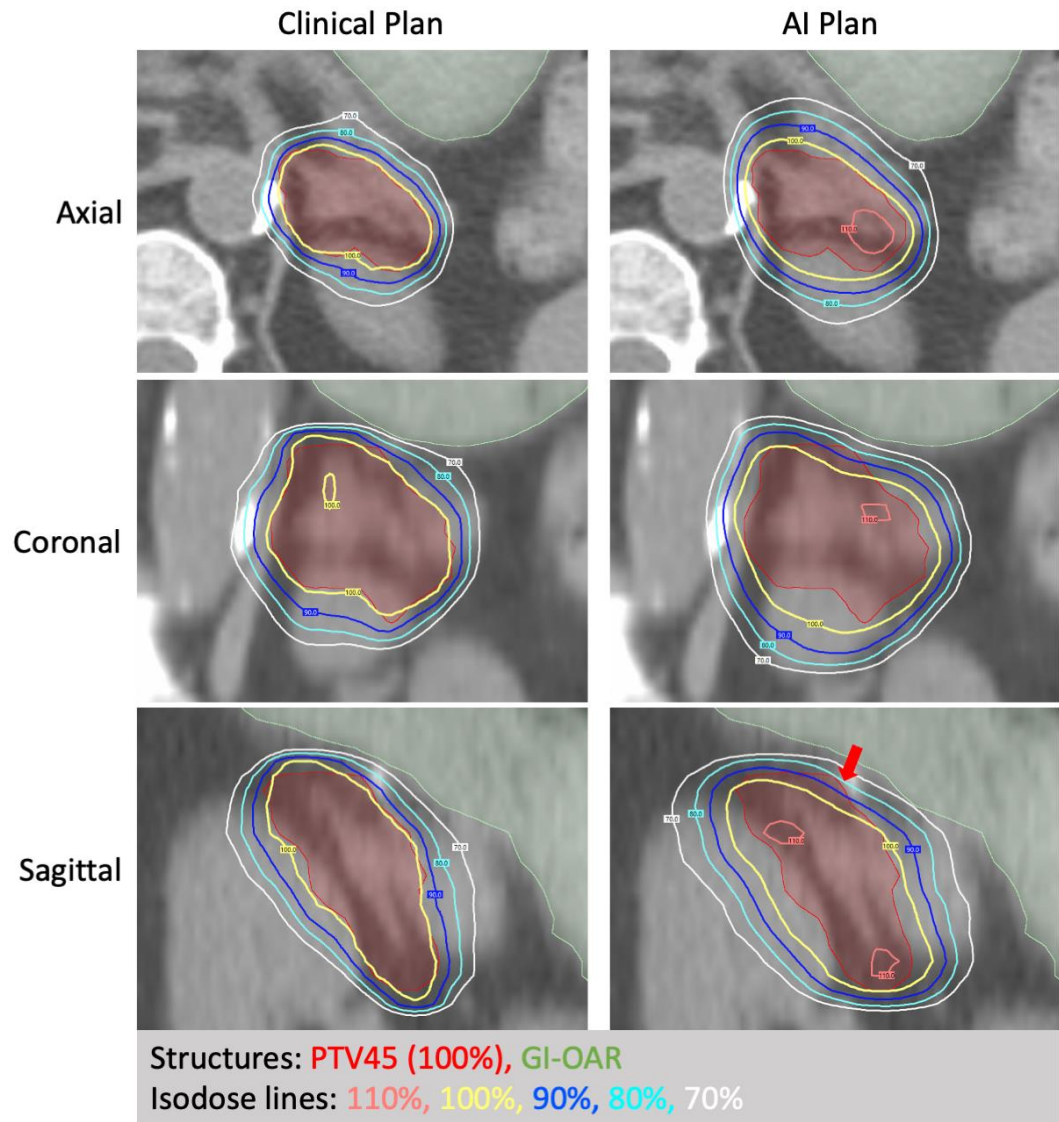


Figure 40: Dose comparison for adrenal SBRT case 2. The clinical plan dose (left column) is shown side by side with the DL plan dose (right column) in the same three views (row 1: axial view; row 2: coronal view; row 3: sagittal view). The PTV is prescribed to 45 Gy (100%) with the GI-OAR limit of 30 Gy (66.7%).

7.4 Discussion

The transfer learning technique effectively reduced the number of cases required to train the deep learning model. In our previous study (W. Wang et al., 2020), we found

that the BD-CNN required at least 45 cases to start to converge and around 55 cases to reach reasonable prediction accuracy. For the pancreas model we developed in previous study, we used 85 training cases. In this study, we show that with a prior model, a high-quality BD-CNN model can be trained with only 10 to 15 cases. The training times for the TL models were also significantly shorter than that of the base model due to the smaller training data size.

In the transfer learning process, we used the common approaches of freezing layers and model fine-tuning (Morid, Borjali, & Del Fiol, 2020). In model validation, we experimented with freezing different parts of the network and using different learning rates, and the current method produced the best results. The downsampling layers of the base pancreas model contain the knowledge of feature extraction from structure contours, while the upsampling layers decode these features and produce beam doses, which are specific to their gantry angles. Since the adrenal cases used a different set of gantry angles from the pancreas cases, only the downsampling layers were frozen during the first model training step. Once the first step was completed, we enabled all layers in the fine-tuning step with 10% of the previous learning rate. During fine-tuning, the downsampling layers were adjusted along with the upsampling layers to further optimize the model for adrenal cases, which featured different prescription levels and target locations.

One of our goals is to reduce the data demand for training such deep learning models. In this study, we trained several models with different training cases. We found that BD-CNN(10) had the highest validation plan score. This suggests that at a minimum, around 10 cases are required as training data to train a robust model. The decrease in validation scores after 10 cases could be attributed to statistical fluctuation, as the test scores increased with more training cases. Transfer learning effectively reduced the data demand to train such deep learning models. It may be useful when there are not enough cases to train a new model from scratch. In general, however, one should use as many high-quality data as possible to train such DL models. Furthermore, for two test cases, we rotated the entire beam bouquet by 60 and 90 degrees to better align with the OAR avoidance. The AI plans for these two cases had similar plan quality as the rest of the DL plans in the test pool and slightly improved test scores compared with the standard beam settings for the same cases. Therefore, the DL model may be able to offer the versatility by allowing different beam angles.

With our previous experience on fluence map prediction for pancreas SBRT, we observed that the DL plans could achieve similar overall dose distribution as the benchmark plans (W. Wang et al., 2020). The results in this study did not reach the same degree of fidelity. As shown in the validation and test cases, the plan quality scores of the DL plans are slightly lower than that of the clinical plans, especially in terms of the OAR score. Two main sources of uncertainty may have contributed to lower prediction

quality. First, the transfer learning model could be further refined to adapt to the unique features of adrenal cases. Second, unlike pancreas SBRT cases which carried a unified prescription dose and OAR constraints, the adrenal cases in this study had a variety of prescriptions as well as OAR constraint preferences. As many clinical plans were VMAT plans with several arcs, the 9-beam IMRT plans were at an inherent disadvantage in terms of plan complexity. Future research work will explore the separation of the causes and attempt to improve the accuracy.

7.5 Conclusions

Transfer learning was used to train a deep learning framework for fluence prediction in adrenal SBRT planning. The AI plans have comparable plan quality as clinical plans in terms of dosimetric endpoints. With only a limited number of training cases, the proposed DL framework can be used to generate plans for different treatment sites and beam orientations. Considering that the dose prescription and beam geometry are different between the two sites, transfer learning has great potential in reducing the data and time demand for training dose prediction networks for rare disease sites and/or by smaller institutions.

8. Conclusions

This study explored different aspects of automated treatment planning using machine learning or deep learning. Many machine learning models have been developed to generalize and utilize prior treatment planning knowledge; however, there is a gap between the developed models and the actual clinical implementation. The direct generation of treatment plans remains difficult for most KBP models or requires additional human intervention. The first project provided a practical solution for WBRT beam placement, which was part of a fully automated planning workflow. The subsequent projects focused on the prediction of IMRT fluence maps, which laid the foundation for an inverse-optimization-free IMRT planning workflow.

The automated beam setting program for WBRT provided a practical solution to beam placement for WBRT. By calculating the geometric field coverage of the PTV and OARs, optimized beams can be generated around one minute. Working in conjunction with the fluence optimization program, this auto-planning system was able to learn from previous clinical plans and create a high-quality breast ECOMP plan adapted to the patient's anatomy. The planning time of the auto-planning workflow (several minutes) was significantly reduced compared with manual planning (several hours), and the plan quality is more consistent. The integration with the commercial TPS further enhanced the clinical feasibility of this approach. Future development of the program

should focus on improving the versatility and robustness to deal with all planning scenarios and patient anatomy.

While this program is designed specifically for WBRT, the methodology is applicable to other disease sites sensitive to the beam placement. As scripting becomes more capable and pervasive in commercial TPS, such treatment planning tools will provide convenience and efficiency to clinics around the world.

The proposed deep learning framework for fluence map prediction is the first complete model to predict IMRT beam doses and fluence maps. The combination of radiation physics and treatment planning knowledge is the key to the success of the model. Data preprocessing is an important step in model development since only useful information should be extracted from existing treatment plans. The multi-slice input allows the model to have the extended field of view in the Z direction, which accounts for beam divergence and scattered dose. Instead of using raw CT images as input, open beam dose is automatically generated in the TPS, providing beam geometry information and patient-specific dose deposition information. The BEV projection serves as a bridge between two image spaces: 3D contour masks and 2D fluence maps. The FM-CNN reverses the dose deposition process by converting beam dose to fluence map. With the fast treatment planning process, the DL framework can be potentially used for adaptive treatment planning, which designs a customized plan for each fraction with daily imaging, thus mitigating the effect of inter-fraction anatomical changes.

The pipeline structure of the DL framework allows the independent development and testing for each component. However, when used in conjunction, the prediction errors are compounded together in the final plan. In the current implementation, the majority of the prediction error comes from beam dose prediction, which can be attributed to the blurring artifacts of the CNN. The resulting plan may lack the sharp dose gradient and higher maximum dose, which is also common in other fluence prediction studies. This effect can be mitigated in several ways, e.g., more training data and more sophisticated loss functions or training techniques like GAN.

There are a few factors that limit the direct clinical implementation of the proposed DL model. The data transfer involves importing and exporting files from the treatment planning system, and the DL prediction is done separately. A built-in program with a graphical user interface would increase the efficiency and lower the technical barrier in using the DL model. In addition to maximum OAR dose, dose-volume constraints, e.g., $V_{15\text{Gy}} < 15\%$, are needed to increase the applicability for different treatment sites. In our evaluation, the model performance can be suboptimal for uncommon patient anatomy, i.e., outlier cases. High model robustness is therefore key in the clinical deployment, which should be improved with continued model development and more training data.

Deep learning models are often considered as “black boxes” compared to more traditional machine learning models. The exact meaning and function of the hidden

layers in a neural network are often unclear since the layer outputs are highly abstract and nonsensical to a human observer. The black box nature of such models can be a limiting factor in their clinical implementations: The users may be less confident in the DL plans compared to the manual plans; The finetuning of the model architecture is a time-consuming trial-and-error process in order to improve the model performance. Therefore, comprehensive quality assurance of the DL plans by physicians and medical physicists or dosimetrists should be implemented before clinical use. Ideally, the DL plans should achieve similar if not better plan quality as the manual plans.

The projects in this study all have limited operating space, i.e., they are designed for a specific site and treatment modality. With more advanced machine learning techniques and available data, it is of great interest to develop more generally applicable AI-assisted treatment planning methods. Such technology will bring about a significant shift and improvement to the standard practice in radiation therapy.

References

- Al-Ghazal, S. K., Fallowfield, L., & Blamey, R. W. (2000). Comparison of psychological aspects and patient satisfaction following breast conserving surgery, simple mastectomy and breast reconstruction. *Eur J Cancer*, 36(15), 1938-1943. Retrieved from <https://www.ncbi.nlm.nih.gov/pubmed/11000574>
- Babier, A., Boutilier, J. J., Sharpe, M. B., McNiven, A. L., & Chan, T. C. (2018). Inverse optimization of objective function weights for treatment planning using clinical dose-volume histograms. *Physics in Medicine & Biology*, 63(10), 105004.
- Bar, Y., Diamant, I., Wolf, L., Lieberman, S., Konen, E., & Greenspan, H. (2015). *Chest pathology detection using deep learning with non-medical training*. Paper presented at the 2015 IEEE 12th international symposium on biomedical imaging (ISBI).
- Barragán - Montero, A. M., Nguyen, D., Lu, W., Lin, M., Norouzi - Kandalan, R., Geets, X., . . . Jiang, S. (2019). Three - dimensional dose prediction for lung IMRT patients with deep neural networks: robust learning from heterogeneous beam configurations. *Medical Physics*.
- Barragán - Montero, A. M., Nguyen, D., Lu, W., Lin, M.-H., Norouzi - Kandalan, R., Geets, X., . . . Jiang, S. (2019). Three - dimensional dose prediction for lung IMRT patients with deep neural networks: robust learning from heterogeneous beam configurations. *Medical Physics*, 46(8), 3679-3691.
- Bentel, G. C. (1996). *Radiation therapy planning* (2nd ed.). New York: McGraw-Hill, Health Professions Division.
- Bijker, N., Meijnen, P., Peterse, J. L., Bogaerts, J., Van Hoorebeeck, I., Julien, J.-P., . . . Fentiman, I. S. (2006). Breast-conserving treatment with or without radiotherapy in ductal carcinoma-in-situ: ten-year results of European Organisation for Research and Treatment of Cancer randomized phase III trial 10853—a study by the EORTC Breast Cancer Cooperative Group and EORTC Radiotherapy Group. *Journal of Clinical Oncology*, 24(21), 3381-3387.
- Brown, M. W., Ning, H., Arora, B., Albert, P. S., Poggi, M., Camphausen, K., & Citrin, D. (2006). A dosimetric analysis of dose escalation using two intensity-modulated radiation therapy techniques in locally advanced pancreatic carcinoma. *International Journal of Radiation Oncology* Biology* Physics*, 65(1), 274-283.

- Brunner, T. B., Nestle, U., Grosu, A.-L., & Partridge, M. (2015). SBRT in pancreatic cancer: What is the therapeutic window? *Radiotherapy and Oncology*, 114(1), 109-116.
- Chen, W. C., Baal, J. D., Baal, U., Pai, J., Gottschalk, A., Boreta, L., . . . Raleigh, D. R. (2020). Stereotactic body radiation therapy of adrenal metastases: A pooled meta-analysis and systematic review of 39 studies with 1006 patients. *International Journal of Radiation Oncology* Biology* Physics*, 107(1), 48-61.
- Chen, X., Men, K., Li, Y., Yi, J., & Dai, J. (2019). A feasibility study on an automated method to generate patient-specific dose distributions for radiotherapy using deep learning. *Med Phys*, 46(1), 56-64. doi:10.1002/mp.13262
- Cox, B. W., Teckie, S., Kapur, A., Chou, H., & Potters, L. (2020). Prospective peer review in radiation therapy treatment planning: Long-term results from a longitudinal study. *Practical radiation oncology*, 10(4), e199-e206.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., & Fei-Fei, L. (2009). *Imagenet: A large-scale hierarchical image database*. Paper presented at the 2009 IEEE conference on computer vision and pattern recognition, Miami, FL.
- Ernster, V. L., Ballard-Barbash, R., Barlow, W. E., Zheng, Y., Weaver, D. L., Cutter, G., . . . Kerlikowske, K. (2002). Detection of ductal carcinoma in situ in women undergoing screening mammography. *Journal of the National Cancer Institute*, 94(20), 1546-1554.
- Ezzell, G. A., Galvin, J. M., Low, D., Palta, J. R., Rosen, I., Sharpe, M. B., . . . committee, A. R. T. (2003). Guidance document on delivery, treatment planning, and clinical implementation of IMRT: report of the IMRT Subcommittee of the AAPM Radiation Therapy Committee. *Med Phys*, 30(8), 2089-2115. doi:10.1118/1.1591194
- Fan, J., Wang, J., Chen, Z., Hu, C., Zhang, Z., & Hu, W. (2019). Automatic treatment planning based on three - dimensional dose distribution predicted from deep learning technique. *Medical Physics*, 46(1), 370-381.
- Fisher, B., Anderson, S., Bryant, J., Margolese, R. G., Deutsch, M., Fisher, E. R., . . . Wolmark, N. (2002). Twenty-year follow-up of a randomized trial comparing total mastectomy, lumpectomy, and lumpectomy plus irradiation for the treatment of invasive breast cancer. *New England Journal of Medicine*, 347(16), 1233-1241.

- Fogliata, A., Reggiori, G., Stravato, A., Lobefalo, F., Franzese, C., Franceschini, D., . . . Cozzi, L. (2017). RapidPlan head and neck model: the objectives and possible clinical benefit. *Radiation Oncology*, 12(1), 1-12.
- Formenti, S. C., DeWyngaert, J. K., Jozsef, G., & Goldberg, J. D. (2012). Prone vs supine positioning for breast cancer radiotherapy. *JAMA*, 308(9), 861-863. Retrieved from https://jamanetwork.com/journals/jama/articlepdf/1356348/jlt120130_861a_863.pdf
- Formenti, S. C., Gidea-Addeo, D., Goldberg, J. D., Roses, D. F., Guth, A., Rosenstein, B. S., & DeWyngaert, K. J. (2007). Phase I-II trial of prone accelerated intensity modulated radiation therapy to the breast to optimally spare normal tissue. *Journal of Clinical Oncology*, 25(16), 2236-2242.
- Foy, J. J., Marsh, R., Ten Haken, R. K., Younge, K. C., Schipper, M., Sun, Y., . . . Matuszak, M. M. (2017). An analysis of knowledge-based planning for stereotactic body radiation therapy of the spine. *Practical radiation oncology*, 7(5), e355-e360.
- Friend, M. (2014). An overview of Electronic tissue Compensation (ECOMP) for breast Radiotherapy. *Journal of Medical Imaging and Radiology Oncology*, 58, 231.
- Gomez, D. R., Tang, C., Zhang, J., Blumenschein Jr, G. R., Hernandez, M., Lee, J. J., . . . Camidge, D. R. (2019). Local consolidative therapy vs. maintenance therapy or observation for patients with oligometastatic non-small-cell lung cancer: long-term results of a multi-institutional, phase II, randomized study. *Journal of Clinical Oncology*, 37(18), 1558.
- Group, E. B. C. T. C. (2006). Effects of radiotherapy and of differences in the extent of surgery for early breast cancer on local recurrence and 15-year survival: an overview of the randomised trials. *The Lancet*, 366(9503), 2087-2106.
- Herman, J. M., Chang, D. T., Goodman, K. A., Dholakia, A. S., Raman, S. P., Hacker - Prietz, A., . . . Pai, J. S. (2015). Phase 2 multi - institutional trial evaluating gemcitabine and stereotactic body radiotherapy for patients with locally advanced unresectable pancreatic adenocarcinoma. *Cancer*, 121(7), 1128-1137.
- Holmlund, J., Brookes, M., Colbert, L. E., Moningi, S., Bhutani, M. S., Katz, M. H., . . . Taniguchi, C. M. (2019). Adaptive Dose Escalation Trial of Stereotactic Body Radiation Therapy (SBRT) in combination with GC4419 in pancreatic cancer. In: American Society of Clinical Oncology.

- Huang, Y., Yue, H., Wang, M., Li, S., Zhang, J., Liu, Z., & Zhang, Y. (2018). Fully automated searching for the optimal VMAT jaw settings based on Eclipse Scripting Application Programming Interface (ESAPI) and RapidPlan knowledge - based planning. *Journal of applied clinical medical physics*, 19(3), 177-182.
- Hussein, M., South, C. P., Barry, M. A., Adams, E. J., Jordan, T. J., Stewart, A. J., & Nisbet, A. (2016). Clinical validation and benchmarking of knowledge-based IMRT and VMAT treatment planning in pelvic anatomy. *Radiotherapy and Oncology*, 120(3), 473-479.
- Isola, P., Zhu, J.-Y., Zhou, T., & Efros, A. A. (2017). *Image-to-image translation with conditional adversarial networks*. Paper presented at the Proceedings of the IEEE conference on computer vision and pattern recognition.
- Kandalan, R. N., Nguyen, D., Rezaeian, N. H., Barragán-Montero, A. M., Breedveld, S., Namuduri, K., . . . Lin, M.-H. (2020). Dose Prediction with Deep Learning for Prostate Cancer Radiation Therapy: Model adaptation to Different Treatment Planning Practices. *Radiotherapy and Oncology*, 153, 228-235. doi:10.1016/j.radonc.2020.10.027
- Kearney, V., Chan, J. W., Haaf, S., Descovich, M., & Solberg, T. D. (2018). DoseNet: a volumetric dose prediction algorithm using 3D fully-convolutional neural networks. *Phys Med Biol*, 63(23), 235022. doi:10.1088/1361-6560/aaef74
- Khan, F. M., & Gibbons, J. P. (2014). *Khan's the physics of radiation therapy* (5 ed.): Lippincott Williams & Wilkins.
- Kim, H., Kwak, J., Jung, J., Jeong, C., Yoon, K., Lee, S.-w., . . . Cho, B. (2018). Automated field-in-field (FIF) plan framework combining scripting application programming Interface and user-executed program for breast forward IMRT. *Technology in cancer research & treatment*, 17, 1533033818810391.
- Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Koay, E. J., Hanania, A. N., Hall, W. A., Taniguchi, C. M., Rebuena, N., Myrehaug, S., . . . Herman, J. M. (2020). Dose-escalated radiation therapy for pancreatic cancer: a simultaneous integrated boost approach. *Practical radiation oncology*.

- Lee, H., Kim, H., Kwak, J., Kim, Y. S., Lee, S. W., Cho, S., & Cho, B. (2019). Fluence-map generation for prostate intensity-modulated radiotherapy planning using a deep-neural-network. *Sci Rep*, 9(1), 15671. doi:10.1038/s41598-019-52262-x
- Lee, T.-F., Chao, P.-J., Ting, H.-M., Chang, L., Huang, Y.-J., Wu, J.-M., . . . Lan, J.-H. (2014). Using multivariate regression model with least absolute shrinkage and selection operator (LASSO) to predict the incidence of xerostomia after intensity-modulated radiotherapy for head and neck cancer. *PLoS One*, 9(2), e89700.
- Li, T., Thongphiew, D., Zhu, X., Lee, W. R., Vujaskovic, Z., Yin, F.-F., & Wu, Q. J. (2011). Adaptive prostate IGRT combining online re-optimization and re-positioning: a feasibility study. *Physics in Medicine & Biology*, 56(5), 1243.
- Li, X., Zhang, J., Sheng, Y., Chang, Y., Yin, F.-F., Ge, Y., . . . Wang, C. (2020). Automatic IMRT planning via static field fluence prediction (AIP-SFFP): a deep learning algorithm for real-time prostate treatment planning. *Physics in Medicine & Biology*, 65(17), 175014.
- Liu, J., Wu, Q. J., Kirkpatrick, J. P., Yin, F. F., Yuan, L., & Ge, Y. (2015). From active shape model to active optical flow model: a shape-based approach to predicting voxel-level dose distributions in spine SBRT. *Phys Med Biol*, 60(5), N83-92. doi:10.1088/0031-9155/60/5/N83
- Mahmood, R., Babier, A., McNiven, A., Diamant, A., & Chan, T. C. (2018). *Automated treatment planning in radiation therapy using generative adversarial networks*. Paper presented at the Machine Learning for Healthcare Conference.
- Mistro, M., Sheng, Y., Ge, Y., Kelsey, C. R., Palta, J. R., Cai, J., . . . Wu, Q. J. (2020). Knowledge Models as Teaching Aid for Training Intensity Modulated Radiation Therapy Planning: A Lung Cancer Case Study. *Frontiers in Artificial Intelligence*, 3, 66.
- Morid, M. A., Borjali, A., & Del Fiol, G. (2020). A scoping review of transfer learning research on medical image analysis using ImageNet. *Computers in Biology and Medicine*, 104115.
- Morrow, N. V., Stepaniak, C., White, J., Wilson, J. F., & Li, X. A. (2007). Intra-and interfractional variations for prone breast irradiation: an indication for image-guided radiotherapy. *International Journal of Radiation Oncology* Biology* Physics*, 69(3), 910-917.

- Mulliez, T., Veldeman, L., Van Greveling, A., Speleers, B., Sadeghi, S., Berwouts, D., . . . Van den Broecke, R. (2013). Hypofractionated whole breast irradiation for patients with large breasts: a randomized trial comparing prone and supine positions. *Radiotherapy and Oncology*, *108*(2), 203-208.
- Nakhlis, F., & Morrow, M. (2003). Ductal carcinoma in situ. *Surgical Clinics*, *83*(4), 821-839.
- Nelms, B. E., Robinson, G., Markham, J., Velasco, K., Boyd, S., Narayan, S., . . . Sobczak, M. L. (2012). Variation in external beam treatment plan quality: An inter-institutional study of planners and planning systems. *Pract Radiat Oncol*, *2*(4), 296-305. doi:10.1016/j.prro.2011.11.012
- Nguyen, D., Jia, X., Sher, D., Lin, M. H., Iqbal, Z., Liu, H., & Jiang, S. (2019). 3D radiotherapy dose prediction on head and neck cancer patients with a hierarchically densely connected U-net deep learning architecture. *Phys Med Biol*, *64*(6), 065020. doi:10.1088/1361-6560/ab039b
- Nguyen, D., Long, T., Jia, X., Lu, W., Gu, X., Iqbal, Z., & Jiang, S. (2019). A feasibility study for predicting optimal radiation therapy dose distributions of prostate cancer patients from patient anatomy using deep learning. *Sci Rep*, *9*(1), 1076. doi:10.1038/s41598-018-37741-x
- Oeffinger, K. C., Fontham, E. T., Etzioni, R., Herzig, A., Michaelson, J. S., Shih, Y. C., . . . American Cancer, S. (2015). Breast Cancer Screening for Women at Average Risk: 2015 Guideline Update From the American Cancer Society. *JAMA*, *314*(15), 1599-1614. doi:10.1001/jama.2015.12783
- Paddick, I. (2000). A simple scoring ratio to index the conformity of radiosurgical treatment plans - Technical note. *Journal of Neurosurgery*, *93*, 219-222. doi:DOI 10.3171/jns.2000.93.supplement_3.0219
- Palma, D. A., Olson, R., Harrow, S., Gaede, S., Louie, A. V., Haasbeek, C., . . . Yaremko, B. P. (2019). Stereotactic ablative radiotherapy versus standard of care palliative treatment in patients with oligometastatic cancers (SABR-COMET): a randomised, phase 2, open-label trial. *The Lancet*, *393*(10185), 2051-2058.
- Pignol, J. P., Olivetto, I., Rakovitch, E., Gardner, S., Sixel, K., Beckham, W., . . . Paszat, L. (2008). A multicenter randomized trial of breast intensity-modulated radiation therapy to reduce acute radiation dermatitis. *J Clin Oncol*, *26*(13), 2085-2092. doi:10.1200/JCO.2007.15.2488

- Purdie, T. G., Dinniwell, R. E., Fyles, A., & Sharpe, M. B. (2014). Automation and intensity modulated radiation therapy for individualized high-quality tangent breast treatment plans. *Int J Radiat Oncol Biol Phys*, 90(3), 688-695. doi:10.1016/j.ijrobp.2014.06.056
- Purdie, T. G., Dinniwell, R. E., Letourneau, D., Hill, C., & Sharpe, M. B. (2011). Automated planning of tangential breast intensity-modulated radiotherapy using heuristic optimization. *Int J Radiat Oncol Biol Phys*, 81(2), 575-583. doi:10.1016/j.ijrobp.2010.11.016
- Ramachandran, P., Zoph, B., & Le, Q. V. (2017). Swish: a self-gated activation function. *arXiv preprint arXiv:1710.05941*, 7.
- Ronneberger, O., Fischer, P., & Brox, T. (2015). *U-Net: Convolutional Networks for Biomedical Image Segmentation*. Paper presented at the International Conference on Medical Image Computing and Computer-Assisted Intervention, Cham.
- Sheng, Y., Li, T., Yoo, S., Yin, F. F., Blitzblau, R., Horton, J. K., . . . Wu, Q. J. (2019). Automatic Planning of Whole Breast Radiation Therapy Using Machine Learning Models. *Front Oncol*, 9, 750. doi:10.3389/fonc.2019.00750
- Sheng, Y., Li, T., Zhang, Y., Lee, W. R., Yin, F. F., Ge, Y., & Wu, Q. J. (2015). Atlas-guided prostate intensity modulated radiation therapy (IMRT) planning. *Physics in Medicine & Biology*, 60(18), 7277.
- Shepard, D. M., Earl, M. A., Li, X. A., Naqvi, S., & Yu, C. (2002). Direct aperture optimization: A turnkey solution for step-and-shoot IMRT. 29(6), 1007-1018. doi:10.1118/1.1477415
- Shiraishi, S., & Moore, K. L. (2016). Knowledge-based prediction of three-dimensional dose distributions for external beam radiotherapy. *Med Phys*, 43(1), 378. doi:10.1118/1.4938583
- Siegel, R. L., Miller, K. D., & Jemal, A. (2019). Cancer statistics, 2019. *CA Cancer J Clin*, 69(1), 7-34. doi:10.3322/caac.21551
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., . . . Bolton, A. (2017). Mastering the game of go without human knowledge. *nature*, 550(7676), 354-359.

- Skarpman Munter, J., & Sjolund, J. (2015). Dose-volume histogram prediction using density estimation. *Phys Med Biol*, 60(17), 6923-6936. doi:10.1088/0031-9155/60/17/6923
- Snyder, K. C., Kim, J., Reding, A., Fraser, C., Gordon, J., Ajlouni, M., . . . Chetty, I. J. (2016). Development and evaluation of a clinical model for lung cancer patients using stereotactic body radiotherapy (SBRT) within a knowledge-based algorithm for treatment planning. *J Appl Clin Med Phys*, 17(6), 6429. Retrieved from <https://www.ncbi.nlm.nih.gov/pubmed/27929499>
- Surucu, M., Shah, K. K., Mescioglou, I., Roeske, J. C., Small Jr, W., Choi, M., & Emami, B. (2016). Decision trees predicting tumor shrinkage for head and neck cancer: Implications for adaptive radiotherapy. *Technology in cancer research & treatment*, 15(1), 139-145.
- Tol, J. P., Delaney, A. R., Dahele, M., Slotman, B. J., & Verbakel, W. F. (2015). Evaluation of a knowledge-based planning solution for head and neck cancer. *Int J Radiat Oncol Biol Phys*, 91(3), 612-620. doi:10.1016/j.ijrobp.2014.11.014
- Van Ginneken, B., Setio, A. A., Jacobs, C., & Ciompi, F. (2015). *Off-the-shelf convolutional neural network features for pulmonary nodule detection in computed tomography scans*. Paper presented at the 2015 IEEE 12th International symposium on biomedical imaging (ISBI), New York, NY.
- Varga, Z., Hideghéty, K., Mező, T., Nikolényi, A., Thurzó, L., & Kahán, Z. (2009). Individual positioning: a comparative study of adjuvant breast radiotherapy in the prone versus supine position. *International Journal of Radiation Oncology* Biology* Physics*, 75(1), 94-100.
- Veronesi, U., Cascinelli, N., Mariani, L., Greco, M., Saccozzi, R., Luini, A., . . . Marubini, E. (2002). Twenty-year follow-up of a randomized study comparing breast-conserving surgery with radical mastectomy for early breast cancer. *New England Journal of Medicine*, 347(16), 1227-1232.
- Wang, C., Zhu, X., Hong, J. C., & Zheng, D. (2019). Artificial Intelligence in Radiotherapy Treatment Planning: Present and Future. *Technol Cancer Res Treat*, 18, 1533033819873922. doi:10.1177/1533033819873922
- Wang, W., Sheng, Y., Palta, M., Czito, B., Willett, C., Hito, M., . . . Wu, Q. J. (2021). Deep Learning-based Fluence Map Prediction for Pancreas Stereotactic Body Radiation Therapy with Simultaneous Integrated Boost. *Advances in Radiation Oncology*, 100672.

- Wang, W., Sheng, Y., Wang, C., Zhang, J., Ge, Y., Yin, F., & Wu, Q. (2019). *Atlas-Guided Fluence Initialization and Optimization for Pancreas SBRT: A Feasibility Study*. Paper presented at the Medical Physics.
- Wang, W., Sheng, Y., Wang, C., Zhang, J., Li, X., Palta, M., . . . Wu, Q. J. (2020). Fluence Map Prediction Using Deep Learning Models – Direct Plan Generation for Pancreas Stereotactic Body Radiation Therapy. *Frontiers in Artificial Intelligence*, 3(68). doi:10.3389/frai.2020.00068
- Wang, W., Sheng, Y., Yoo, S., Blitzblau, R. C., Yin, F. F., & Wu, Q. J. (2019). Goal-Driven Beam Setting Optimization for Whole-Breast Radiation Therapy. *Technol Cancer Res Treat*, 18, 1533033819858661. doi:10.1177/1533033819858661
- Wieser, H. P., Cisternas, E., Wahl, N., Ulrich, S., Stadler, A., Mescher, H., . . . Burigo, L. (2017). Development of the open - source dose calculation and optimization toolkit matRad. *Medical Physics*, 44(6), 2556-2568.
- Wu, H., Jiang, F., Yue, H., Li, S., & Zhang, Y. (2016). A dosimetric evaluation of knowledge - based VMAT planning with simultaneous integrated boosting for rectal cancer patients. *Journal of applied clinical medical physics*, 17(6), 78-85.
- Yang, W., Reznik, R., Fraass, B. A., Nissen, N., Hendifar, A., Wachsman, A., . . . Tuli, R. (2015). Dosimetric evaluation of simultaneous integrated boost during stereotactic body radiation therapy for pancreatic cancer. *Medical Dosimetry*, 40(1), 47-52.
- Yuan, L., Ge, Y., Lee, W. R., Yin, F. F., Kirkpatrick, J. P., & Wu, Q. J. (2012). Quantitative analysis of the factors which affect the interpatient organ - at - risk dose sparing variation in IMRT plans. *Medical Physics*, 39(11), 6868-6878.
- Yuan, L., Wu, Q. J., Yin, F., Li, Y., Sheng, Y., Kelsey, C. R., & Ge, Y. (2015). Standardized beam bouquets for lung IMRT planning. *Phys Med Biol*, 60(5), 1831-1843. doi:10.1088/0031-9155/60/5/1831
- Zhang, J., Wang, C., Sheng, Y., Palta, M., Czito, B., Willett, C., . . . Wu, Q. (2020). An interpretable planning bot for pancreas stereotactic body radiation therapy. *International Journal of Radiation Oncology* Biology* Physics*.
- Zhang, X., Li, X., Quan, E. M., Pan, X., & Li, Y. (2011). A methodology for automatic intensity-modulated radiation treatment planning for lung cancer. *Physics in Medicine & Biology*, 56(13), 3873.

- Zhao, X., Kong, D., Jozsef, G., Chang, J., Wong, E. K., Formenti, S. C., & Wang, Y. (2012). Automated beam placement for breast radiotherapy using a support vector machine based algorithm. *Medical Physics*, 39(5), 2536-2543.
- Zhu, X., Ge, Y., Li, T., Thongphiew, D., Yin, F. F., & Wu, Q. J. (2011). A planning quality evaluation tool for prostate adaptive IMRT based on machine learning. *Med Phys*, 38(2), 719-726. doi:10.1118/1.3539749

Biography

Wentao Wang earned his bachelor's degree in June 2016 with an astronomy major in Kuang Yaming Honors School of Nanjing University and received several scholarships. In August 2014, he studied physics for one semester as an exchange student at Duke University with a full scholarship. In August 2016, he joined the medical physics graduate program at Duke Kunshan University with Chancellor's scholarship and conducted research in Dr. Jackie Wu's lab at Duke University. He received his Master of Science degree in medical physics in 2018 and joined the medical physics graduate program at Duke University as a PhD student later in the same year. He has published peer-reviewed articles "Goal-Driven Beam Setting Optimization for Whole-Breast Radiation Therapy" in *Technology in Cancer Research & Treatment*, "Fluence Map Prediction Using Deep Learning Models–Direct Plan Generation for Pancreas Stereotactic Body Radiation Therapy" in *Frontiers in Artificial Intelligence*, "Deep Learning-based Fluence Map Prediction for Pancreas Stereotactic Body Radiation Therapy with Simultaneous Integrated Boost" in *Advances in Radiation Oncology* as the first author, as well as a co-authored article in *Medical Physics*. He has nine conference presentations, including 6 oral presentations and 4 ePosters in AAPM annual and clinical meetings, ASTRO annual meetings, RSS annual meetings. He is a student member of the AAPM, ASTRO, and RSS.