

Quantum Inspired Machine Learning Algorithms for Adaptive Radiotherapy

by

Julia M. Pakela

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Applied Physics)
in the University of Michigan
2021

Doctoral Committee:

Professor Issam El Naqa, Co-chair
Associate Professor Martha M. Matuszak, Co-chair
Professor J. Brian Fowlkes
Professor S. Sandeep Pradhan
Professor Randall K. Ten Haken

Julia M. Pakela

jpakela@umich.edu

ORCID ID: 0000-0001-5506-187X

© Julia M. Pakela 2021

Acknowledgements

I have accomplished this endeavor through the help of many people. My advisor Professor Issam El Naqa has been an incredible mentor to me, and I thank him for his patience, guidance, and for believing in me (sometimes more than I believed in myself). I have been fortunate to have been part of a wonderful research group whose questions and discussions have helped to improve the quality of my work. I also would like to thank my committee members, in particular Professors Randall Ten Haken and Martha Matuszak, for their mentorship and encouragement. The staff at Argus has been an incredible help to me, and in particular I owe thanks to Janell Dow for her patient instruction on the use of treatment planning software, Dr. Jim Irrer for his help with all things DICOM related, and Gabriel Dotson for assisting with many logistical issues. I also owe thanks to my Applied Physics family, in particular Professor Cagliyan Kurdak and Cynthia McNabb, for helping me to navigate any questions and hurdles that came up over the course of my degree. Finally, I would like to thank my parents-in-law Mark and Karen, my parents Stacey and Gregg, and my husband Christopher for their unconditional love and support.

Table of Contents

Acknowledgements	ii
List of Tables	vii
List of Figures.....	ix
List of Acronyms and Abbreviations	xiii
Abstract.....	xvii
Chapter 1 Introduction.....	1
1.1 Adaptive Radiotherapy	1
1.2 Machine Learning and Deep Learning	3
1.3 Quantum-Inspired Algorithms and Quantum Computing	5
1.4 Motivation and Contributions of this Study	6
1.4.1 Treatment Planning Optimization	7
1.4.2 Adaptation of Radiotherapy Plans to Geometrical Changes	7
1.4.3 Radiotherapy Outcomes and Feature Identification	8
1.5 Dissertation Organization	8
1.6 Accomplishments.....	9
1.6.1 Awards.....	9
1.6.2 Peer-reviewed Publications	9
1.6.3 Book Chapters	11
1.6.4 Presentations.....	11
1.7 References.....	13
Chapter 2 Background	18
2.1 Radiation Therapy.....	18
2.1.1 Types of External Beam Radiation Therapy	19
2.1.2 Treatment Planning	20
2.1.3 Adaptive Radiotherapy for Head and Neck Cancer	21

2.2	Handcrafted Features	23
2.2.1	Radiomic Features	23
2.3	Optimization Algorithms	24
2.3.1	Gradient Descent	24
2.3.2	Simulated Annealing	25
2.4	Classical Machine Learning Algorithms	25
2.4.1	Logistic Regression	25
2.4.2	Cox Proportional Hazards Model.....	26
2.5	Deep Learning.....	27
2.5.1	Fully Connected Neural Networks	28
2.5.2	Training a Neural Network.....	29
2.5.3	Hyperparameters Associated with Training	32
2.5.4	What Makes a Neural Network Deep?.....	34
2.6	Model Evaluation and Benchmarking	35
2.6.1	Performance Evaluation: Area Under the Receiver Operating Characteristic Curve and Harrel's C-Index.....	35
2.6.2	Cross Validation	36
2.7	Quantum Algorithms and Quantum Computing.....	37
2.7.1	Overview of Quantum Computing	37
2.7.2	Postulates of Quantum Mechanics	40
2.7.3	Quantum Cognition and Decision Theory.....	43
2.7.4	Quantum Annealers	43
2.8	References.....	46
Chapter 3 A Quantum-Inspired Approach to Radiotherapy Optimization.....		51
3.1	Introduction.....	51
3.2	Methods	53
3.2.1	Quantum Tunnel Annealing	53
3.2.2	Calculation of Barrier Width.....	55
3.2.3	Simulated Annealing (SA)	58
3.2.4	IMRT Case Selection	59
3.2.5	Objective Function	61
3.2.6	Extension to Influence-based Direct Aperture Optimization	64
3.2.7	Criteria for Convergence	64
3.2.8	Computing Environment	65

3.3	Results.....	68
3.3.1	Case 1	68
3.3.2	Case 2: Redefined Smoothing Filter	69
3.3.3	Case 2: Barrier Width Schedule Effect	71
3.3.4	Case 2: Annealing Schedule Effect	72
3.3.5	Case 2: Optimization Stability	73
3.3.6	Aperture-Weight Optimization via Influence-Based DAO	77
3.4	Discussion.....	79
3.5	Conclusions.....	83
3.6	References.....	84
Chapter 4 Dynamic Stochastic Deep Learning Approaches for Predicting Geometric Changes in Head and Neck Cancer		87
4.1	Introduction.....	87
4.1.1	Adaptive Radiotherapy.....	87
4.1.2	Predictive Models for ART	89
4.1.3	Quantum Predictive Models and Quantum Processes.....	90
4.1.4	Classical Markov Processes	92
4.1.5	Deep Learning and Recursive Neural Networks.....	93
4.1.6	Study Objectives.....	94
4.2	Methods	95
4.2.1	Patient Population and Data Acquisition.....	95
4.2.2	Data Preprocessing	96
4.2.3	Mapping to Discrete Orthonormal State Vectors	98
4.2.4	Generation of Additional Synthetic Data via Generative Adversarial Network (GAN).....	100
4.2.5	Quantum-based Prediction Algorithm.....	102
4.2.6	Markov-based Prediction Algorithm (MRNN)	105
4.2.7	Model Architecture and Training	106
4.3	Results.....	106
4.3.1	Overall Results	107
4.3.2	4-state Predictions	110
4.3.3	6-state Predictions	112
4.3.4	8-state Predictions	114
4.3.5	10-state Predictions	116

4.3.6	External Testing Results.....	118
4.4	Discussion.....	119
4.5	Conclusion	123
4.6	References.....	123
Chapter 5	Prediction of Adaptation Decisions and Outcomes in Larynx Cancer Patients	
	126	
5.1	Introduction.....	127
5.1.1	Bioselection as a Method for Adaption in Larynx Cancer Patients	127
5.1.2	Data-driven Outcomes Prediction in Head and Neck Cancers.....	127
5.1.3	Study Objectives.....	128
5.2	Methods	129
5.2.1	Patient Cohort.....	129
5.2.2	Feature Acquisition and Processing	130
5.2.3	Statistical Regression Models Design	134
5.2.4	Nested Cross Validation for Feature Selection in Classical Models.....	136
5.2.5	Deep Learning Model Design	137
5.3	Results.....	138
5.3.1	Classical Model Results	138
5.3.2	Deep Learning Model Results	139
5.4	Discussion.....	141
5.5	Conclusion	142
5.6	References.....	142
Chapter 6	Discussion and Future Directives	146
6.1	Discussion.....	146
6.2	Current challenges and limitations	147
6.2.1	Data Acquisition, Standardization, and Quantization	147
6.2.2	Clinical Implementation	149
6.3	Future work.....	150
6.3.1	Further Validation and Incorporation of Biological Objectives for QTA.....	150
6.3.2	Identifying Problem-spaces Tailored to Quantum-based Frameworks	150
6.3.3	Incorporation of PET-CT and Identification of Robust Features.....	151
6.4	References.....	152

List of Tables

Table 3.1: DVH constraints applied to objective function for Case 1.	61
Table 3.2: DVH constraints applied to objective function for Case 2.	61
Table 3.3: Convergence times (in seconds) for QTA with different barrier width schedules as well as SA.	72
Table 3.4: Mean convergence times (in seconds) for QTA and SA with perturbations to the initial beamlet-weight values.	76
Table 3.5: Parameter changes and convergence times (in seconds) for QTA and SA with perturbations to the original dose constraints.	77
Table 4.1: Validation AUCs averaged over 5-folds of cross validation recorded for MRNN and QRNN models after 100 epochs. (mean \pm 95% confidence interval [CI]), as well as z-scores and corresponding <i>p</i> -values from the DeLong test.	108
Table 4.2: Number of folds omitted from each fraction during AUC score calculation due to lack of samples for each possible class.	110
Table 4.3: Testing AUC scores calculated as the highest achieved over the course of 100 epochs of training. Results include models trained with real and synthetic data as well as models trained without synthetic data (labeled as “no GAN”).	118
Table 5.1: Demographics of the training and testing dataset.	129
Table 5.2: Radiomic imaging features.	132
Table 5.3: Feature selection for model building.	139

Table 5.4: Classical Model performance.	139
Table 5.5: Deep learning model performance.....	140

List of Figures

Figure 1.1: Incidence of deep-learning-themed papers in biomedical literature as well as in medical physics and radiation oncology. Data was obtained through the advanced search function on PubMed. Search criteria used were (all fields: deep learning), (all fields: deep learning) AND (all fields: medical physics OR all fields: radiation oncology) and (all fields: deep learning) AND (all fields: radiology) respectively.	5
Figure 2.1: Visualization of radiotherapy workflow highlighting some of the primary questions and uncertainties surrounding the implementation of ART techniques.....	22
Figure 2.2: Conventional “shallow” machine learning (top) versus deep learning algorithms, where image data representation and classification are handled within the same framework.	27
Figure 2.3: A three-layer neural network (input layer, hidden layer, and output layer). The output of each node is determined by performing a non-linear transformation (known as the activation function) on the sum of the weighted inputs plus an additional bias term.	29
Figure 2.4: Common nonlinear activation functions used in neural networks.	31
Figure 3.1: Figurative illustration of a particle (represented by its wave-function, Ψ) tunneling through a potential energy barrier (in region B) in a 1-dimensional energy landscape.....	55
Figure 3.2: (a) Barrier width rate extrapolated from metal organic chemical vapor deposition (MOCVD) studies by Leys and Veenvliet. ²⁵ (b) Barrier width function calculated via numerical integration of (a). (c) Additional width functions explored in this study.	57

Figure 3.3: CT scans show contours for structures optimized for Case 1 (a) and Case 2 (b), respectively. Case 1 features a PTV that is roughly spherical in shape and far from major organs. (with the exception of the liver). Case 2 features a PTV with convex geometrical features and close proximity to both the liver and the stomach. 60

Figure 3.4: Process of calculating the convergence iteration number from a representative QTA optimization. The saved energy history is represented by (a), and is used to calculate the energy gradient ((b)). From this gradient a MAM of width 100 was calculated ((c)). The black vertical line is plotted at the maximum iteration number j for $\max_j(|MAM(j)| > c^{tol})$ 65

Figure 3.5: Quantum Tunnel Annealing (QTA) algorithm for intensity modulated radiation therapy (IMRT) optimization. 67

Figure 3.6: Optimization results for QTA and SA applied to Case 1. (a) and (b) display DVH curves for 10 separate optimizations using QTA and SA, respectively. (c) and (d) display the PE trajectories for the 10 QTA and SA optimizations. (e) and (f) display representative dose distributions calculated in Eclipse using fluence values from the 10th QTA and SA optimization. 69

Figure 3.7: (a) displays the fluence map results for a single beam in Case 2 resulting from the QTA optimization without refined smoothing. (b) displays the fluence map results from QTA optimization with refined smoothing. 71

Figure 3.8: Displays the annealing schedule functions tested for QTA and SA. Note that T4 was the annealing schedule used for all remaining studies. (b) displays box and whisker plots of the convergence results for QTA and SA, respectively, for each annealing schedule. 73

Figure 3.9: DVH bands, PE trajectories, and representative dose distributions for stochastic optimizations ($N = 10$, 500000 iterations) with the initial beamlet-weight vector set to random

values uniformly distributed between 0 and 22 on a challenging SBRT liver case for QTA ((a), (c), and (e)) and SA ((b), (d), and (f)), respectively. 75

Figure 3.10: DVH bands, PE trajectories, and representative dose distributions for stochastic optimizations ($N = 10$, 250000 iterations) of aperture weights on a challenging SBRT liver case for QTA ((a), (c), and (e)) and SA ((b), (d), and (f)), respectively. 78

Figure 4.1: Schematic of data collected from each patient across treatment fractions. Primary CTV volumes were acquired from daily CBCTs and daily table shifts (described by the coordinate space shown in the treatment system) were collected for each available fraction and used to define an orthonormal state vector. Linear accelerator image adapted from Varian official website, 98

Figure 4.2: Distribution of lateral shifts recorded across all fractions from the original dataset and after encoding to 4 and 10 states, as well as a subset of the lateral shifts recorded from the raw, unquantized dataset (orange) and after encoding (blue). 100

Figure 4.3: Assessment of synthetic data generated by TimeGAN for the 10-state vector encoding trained using original data from fold 5: (a) shows the distribution of the original 10-state encoded data, (b) shows the distribution of the synthetic, (c) displays discriminative and predictive scores for the HN synthetic data (H&N) and the scores from the lung cancer pathways data (Lung*) reported in the original TimeGAN paper, (d) and (e) display t-SNE and PCA plots respectively for the original 10-state data and the synthetic 10-state data. 102

Figure 4.4: Schematic of QRNN Algorithm. A set number of initial states are fed into a recurrent neural network, which produces a vector of parameters necessary to build a Hamiltonian matrix. The Hamiltonian is then used to construct a transition matrix, which in turn is used in conjunction with the oldest initial state to predict the probability of each remaining fraction existing in a given

state. The fraction state predictions are then used in conjunction with the ground truth to calculate the model loss. 105

Figure 4.5: MRNN and QRNN validation AUC scores vs fraction number for each system state representation. Shaded region indicates the 95% confidence interval..... 109

Figure 4.6: QRNN and MRNN model results for 4-state system predicting fractions 15-35: AUC scores (mean and 95% CI) (a) & (b), model loss (c) & (d), and confusion matrix metrics: sensitivity, specificity, positive predictive value (PPV), and negative predictive value (NPV) averaged across 5-folds after 100 epochs of training (e) & (f). 111

Figure 4.7: QRNN and MRNN model results for 6-state system predicting fractions 15-35: AUC scores (mean and 95% CI) (a) & (b), model loss (c) & (d), and confusion matrix metrics: sensitivity, specificity, positive predictive value (PPV), and negative predictive value (NPV) averaged across 5-folds after 100 epochs of training (e) & (f). 113

Figure 4.8: QRNN and MRNN model results for 8-state system predicting fractions 15-35: AUC scores (mean and 95% CI) (a) & (b), model loss (c) & (d), and confusion matrix metrics: sensitivity, specificity, positive predictive value (PPV), and negative predictive value (NPV) averaged across 5-folds after 100 epochs of training (e) & (f). 115

Figure 4.9: QRNN and MRNN model results for 10-state system predicting fractions 15-35: AUC scores (mean and 95% CI) (a) & (b), model loss (c) & (d), and confusion matrix metrics: sensitivity, specificity, positive predictive value (PPV), and negative predictive value (NPV) fractions averaged across 5-folds after 100 epochs of training (e) & (f). 117

List of Acronyms and Abbreviations

3DCRT: 3-Dimensional Conformal Radiation Therapy

ANN: Artificial Neural Network

API: Application Program Interface

ART: Adaptive Radiation Therapy/Adaptive Radiotherapy

AUC: Area Under the Receiver Operating Characteristic Curve

CAD: Computer-aided Diagnosis

CBCT: Cone-beam Computed Tomography

Cox: Cox Proportional Hazards

CI: Confidence Interval

CT: Computed Tomography

CTV: Clinical Target Volume

CV: Cross Validation

DAO: Direct Aperture Optimization

DIR: Deformable Image Registration

DL: Deep Learning

DVH: Dose Volume Histogram

EBRT: External Beam Radiation Therapy

FPR: False Positive Rate

GaAs: Gallium Arsenide

GAN: Generative Adversarial Network
GLCM: Gray-Level Co-occurrence Matrix
GLRLM: Gray-Level Run-Length Matrix
GLSZM: Gray-Level Size Zone Matrix
GPU: Graphics Processing Unit
GRU: Gated Recurrent Unit
GTV: Gross Tumor Volume
HNC: Head and Neck Cancer
HPV: Human Papilloma Virus
IC: Induction Chemotherapy
ICR: Response to Induction Chemotherapy
IG-ART: Image-guided Adaptive Radiotherapy
IMRT: Intensity Modulated Radiation Therapy
IRB: Institutional Review Board
LASSO: Least Absolute Shrinkage and Selection Operator
LFS: Laryngectomy Free Survival
Linac: Linear Accelerator
LMR: Lymphocyte-Monocyte Ratio
LOG: Laplace of Gaussian
LOOCV: Leave One Out Cross Validation
LR: Logistic Regression
MAM: Moving Average Mean
MLC: Multi-leaf Collimator

MLP: Multi-layer Perceptron

MOCVD: Metal Organic Vapor Deposition

MU: Monitor Unit

NCV: Nested Cross Validation

NGTDM: Neighborhood Gray-Tone Difference Matrix

NLR: Neutrophil-Lymphocyte Ratio

NMR: Nuclear Magnetic Resonance

NPV: Negative Predictive Value

NTCP: Normal Tissue Complication Probability

OAR: Organ At Risk

OS: Overall Survival

PCA: Principal Component Analysis

PET: Positron Emission Tomography

PIMV: Plan Intensity Map Variation

PPV: Positive Predictive Value

PTV: Planning Target Volume

QA: Quantum Annealing

QTA: Quantum Tunnel Annealing

Qubit: Quantum Bit

ReLU: Rectified Linear Unit

RNN: Recursive Neural Network

ROC: Receiver Operating Characteristic

ROI: Region of Interest

RT: Radiotherapy

SBRT: Stereotactic Body Radiation Therapy

SG: Savitzky-Golay filter

TCP: Tumor Control Probability

Time-GAN: Time-series Generative Adversarial Network

TPR: True Positive Rate

t-SNE: t-Distributed Stochastic Neighbor Embedding

VMAT: Volumetric Arc Therapy

WKB: Wentzel-Kramers-Brillouin approximation

Abstract

The quest for innovations that can maximize the curative potential of radiation therapy while minimizing harmful dose to critical organs has been an active research topic for over 100 years. In recent decades, adaptive radiotherapy (ART) has been recognized as an important step towards realizing this goal. ART refers to the modification of treatment plans in response to patient changes, such as weight-loss, tumor shrinkage, or daily anatomical variations over the course of treatment. The rationale for adaptive radiotherapy is that as patients (who typically undergo weeks of fractionated treatment) experience geometric or anatomical changes near the tumor site, the parameters of their treatment plan (defined by a computed tomography [CT]) simulation taken prior to the start of treatment) may no longer provide adequate target coverage or normal tissue sparing days or weeks later. This dissertation explores the novel application of quantum physics principles and deep machine learning techniques to address three challenges towards the clinical implementation of ART: (1) efficient calculation of optimal treatment parameters, (2) adaptation to geometrical changes over the treatment period while mitigating associated uncertainties, and (3) understanding the relationship between individual patient characteristics and clinical outcomes. Applications of quantum and machine learning modeling in other fields support the potential of this novel, interdisciplinary approach.

For efficient optimization, we developed and tested a quantum-inspired, stochastic algorithm for intensity-modulated radiotherapy (IMRT): quantum tunnel annealing (QTA). By modeling the likelihood probability of accepting a higher energy solution after a particle tunneling through a potential energy barrier, QTA features an additional degree of freedom (the barrier width, w) not

shared by traditional stochastic optimization methods such as simulated annealing (SA). QTA was validated on two liver cancer patients for beamlet weight optimization for IMRT and direct aperture optimization for volumetric arc therapy. QTA was found to achieve convergence up to 46.6% (26.8%) faster than SA for beamlet-weight optimization and direct aperture optimization, respectively. The results of this study suggest that the additional degree of freedom provided by QTA can improve convergence rates and achieve a more efficient and, potentially, effective treatment planning process.

For geometrical adaptation, we investigated the feasibility of predicting patient changes across a fractionated treatment schedule using two approaches. The first was based on a joint framework (referred to as QRNN) employing quantum mechanics in combination with deep recurrent neural networks (RNNs). The second approach was developed based on a classical framework (MRNN), which modelled patient anatomical changes as a Markov process. We evaluated and compared these two approaches' performance characteristics using a dataset of 125 head and neck cancer patients who received radiotherapy as part of their treatment and had daily cone-beam CT (CBCT) imaging. The MRNN framework was found to have slightly better performance than the QRNN framework, with MRNN(QRNN) validation area under the receiver operating characteristic curve (AUC) scores (\pm 95% confidence interval) of 0.742 ± 0.021 (0.675 ± 0.036), 0.709 ± 0.026 (0.656 ± 0.021), 0.724 ± 0.036 (0.652 ± 0.044), and 0.698 ± 0.016 (0.605 ± 0.035) for system state vector sizes of 4, 6, 8, and 10, respectively. Of these, only the results from the two higher order states were found to have statistically significant differences ($p < 0.05$). A similar trend was also observed when the fully trained models were applied to an external testing dataset of 20 patients, yielding MRNN(QRNN) AUC scores of 0.707 (0.623), 0.687 (0.608), 0.723 (0.669), and 0.697 (0.609) for states vectors sizes of 4, 6, 8, and 10, respectively. These results

suggest that these stochastic models provide added value in predicting patient changes during the course of adaptive radiotherapy.

Towards understanding the relationship between patient characteristics and clinical outcomes, we performed a series of studies which investigated the use of quantitative patient features for predicting clinical outcomes in laryngeal cancer patients who underwent treatment in a bioselection paradigm based on surgeon-assessed response to induction chemotherapy. Among the features investigated from CT scans taken before and after induction chemotherapy, two (gross tumor volume change between pre- and post-induction chemotherapy, and nodal stage) had prognostic value for predicting patient outcomes using standard regression models. Artificial neural networks did not improve predictive performance in this case.

Taken together, the significance of these studies lies in their contribution to the body of knowledge of medical physics and in their demonstration of the use of novel techniques which incorporate quantum mechanics and machine learning as a joint framework for treatment planning optimization and prediction of anatomical patient changes over time.

Chapter 1 Introduction

“Nothing is too wonderful to be true, if it be consistent with the laws of nature”

– *Michael Faraday*

1.1 Adaptive Radiotherapy

Adaptive radiotherapy (ART) refers to a radiation therapy treatment regimen in which a patient’s treatment plan is modified mid-treatment in response to measurements on relevant patient features. Adaptive radiotherapy is an attractive alternative to standard radiotherapy, in which the treatment plan remains static over the entire course of treatment, because patients do not remain static with respect to the medical images used to design their treatment plans. Over the course of treatment, patients may experience anatomical variations such as tumor shrinkage or weight loss, posture and daily setup variations, as well as anatomical changes due to bodily functions such as respiration or digestion.¹ These changes in the “state” of the patient represent an uncertainty in tumor coverage which traditionally has been handled by adding a margin around the clinical target volume (CTV) to create the planning target volume (PTV). ART promises to provide a personalized approach handling anatomical variations which provides greater tumor coverage and sparing of normal tissues.

Since its inception in the late 1990s, much progress has been made towards the implementation of adaptive radiotherapy.² Advancements in imaging such as the development of

cone-beam computed tomography (CBCT) have allowed for standard use of daily on-board imaging, in which the patient is imaged directly from the treatment couch either immediately before or during treatment. The use of daily imaging allows for the opportunity to perform image-guided adaptive radiotherapy, in which daily imaging is used to inform necessary treatment modifications. In addition, many studies have sought to address questions such as: 1) what are the expected margins of anatomical variation during treatment,³⁻⁸ 2) how much do such variations result in dose deviations from the original treatment plan,⁹ and 3) what are the potential clinical impacts of such deviations?¹⁰ Finally, studies have also presented predictive frameworks to identify patients who would benefit from ART based on expected anatomical or dosimetric changes.¹¹

Nevertheless, many challenges still remain for broad implementation of ART in the clinical environment. Because treatment plan design and optimization are time-consuming and resource-intensive processes, efficient and robust optimization algorithms are necessary for plan adaptations to be feasible under the constraints of a busy clinical environment. In addition, there is still room for the development of improved predictive frameworks which can pre-emptively identify which patients are likely to experience significant anatomical changes during the course of their treatment, and such frameworks need to be robust to uncertainties related to CBCT image quality and restricted field of view. Finally, there is a general lack of consensus among the radiation oncology community regarding what kinds of patient changes or features that should be utilized for guiding adaptive radiotherapy.¹² Further work is therefore necessary to better understand how individual patient characteristics and anatomical changes can relate to meaningful clinical outcomes such as survival or radiation-induced toxicity. The studies presented in this dissertation

seek to provide additional tools and insights to address these challenges, especially in the most common geometrical deformations that may impact the quality of the delivered radiotherapy.

1.2 Machine Learning and Deep Learning

Machine learning refers to data-driven algorithms which use observations in a training dataset to learn a functional mapping from inputs to desired outputs (labels). Deep learning refers to a subclass of machine learning algorithms that are capable of learning higher order representations of data through the application of multiple layers on non-linear operations performed on the input data.

The foundational groundwork for deep learning was laid decades prior to its rise in popularity: first in the 1950s when Rosenblatt introduced the concept of the perceptron (a precursor to the hidden nodes in today's neural networks) and later in the 1980s with Hinton and Sejnowski's invention of the Boltzmann Machine (a multilayer network similar in design to modern neural networks).^{13,14} However, for several decades deep learning methods were not viable for solving real-world problems. This was partly due to limitations on the computer processing power necessary to train the multi-layer networks envisioned by deep learning researchers and also due to the need for further algorithmic innovations (such as dropout and stochastic optimization schemes) to improve training efficiency. By the early 2010s, such discoveries, in combination with improvements in computer parallel processing power and graphics processing units (GPUs), helped to provide the necessary conditions for deep learning methods to achieve widespread success and recognition. This phenomenon arguably began with the startling 2012 victory in the ImageNet competition (an annual image classification challenge which sets the bar for the state-of-the-art in computer vision) with AlexNet: a deep convolutional neural network architecture.¹⁵ Deep learning has since been applied with great success to many diverse challenges, including

language translation, speech recognition, training of self-driving cars, and even stock market predictions.¹⁶

Since its recent rise in popularity, deep machine learning has also made a significant impact on the medical field, specifically in the areas of diagnostic (radiology) and therapeutic (radiation oncology) radiological sciences.¹⁷ Applications in radiology and computer-aided diagnosis (CAD) have been at the forefront of application of machine and deep learning in medicine since the 1980s.¹⁸⁻²³ These applications included using artificial neural networks (ANNs) generally²⁴ as well as convolutional neural networks (CNNs) for breast cancer detection and diagnosis.²⁵⁻²⁷ This pioneering work has led to several Food-and-Drug-Administration-approved systems including the QuantX Advanced system to aid in breast cancer diagnosis (CADx), developed originally by Giger and colleagues.²⁸⁻³² Today, deep learning techniques are touching every aspect of radiology—from enhancing image quality by improving current image reconstruction and filtering of modalities such as MRI,³³ CT,³⁴ and ultrasound³⁵; to image segmentation³⁶ and registration³⁷; and to precision medicine and the derivation of reproducible imaging biomarkers.³⁸

As another innovative and data-heavy field, radiation oncology is uniquely positioned to experience an explosion of deep learning applications. The number of radiation oncology and medical physics papers published which feature deep learning has increased steadily over the past five years, with a wide variety of applications including treatment planning,³⁹⁻⁴² adaptive radiotherapy,^{35,43-46} quality assurance,⁴⁷⁻⁵⁰ and outcomes modeling.⁵¹⁻⁵⁷ **Figure 1.1** provides a visualization of the rise in deep learning both in biomedical literature and in publications specific to medical physics and radiation oncology.

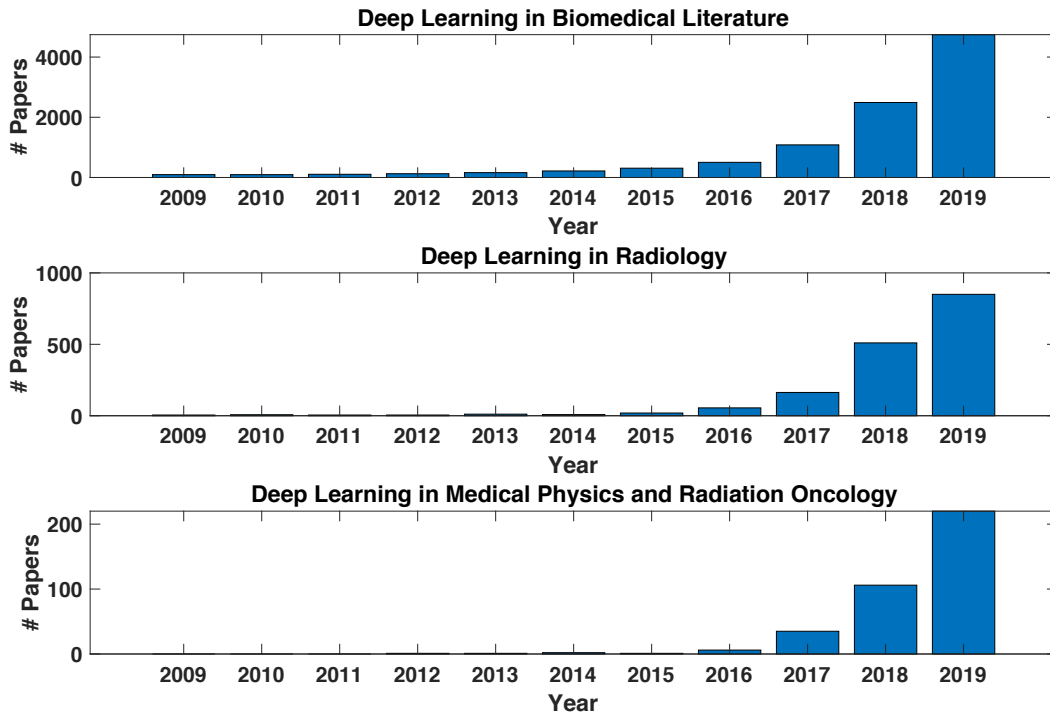


Figure 1.1: Incidence of deep-learning-themed papers in biomedical literature as well as in medical physics and radiation oncology. Data was obtained through the advanced search function on PubMed. Search criteria used were (all fields: deep learning), (all fields: deep learning) AND (all fields: medical physics OR all fields: radiation oncology) and (all fields: deep learning) AND (all fields: radiology) respectively.

1.3 Quantum-Inspired Algorithms and Quantum Computing

Quantum mechanics was developed as a predictive framework to describe phenomena which violate the laws of classical physics. More specifically, quantum mechanics can be used to describe the behavior of systems in nature which exhibit parameters on the order of their de Broglie wavelength—such systems are referred to as “quantum systems.”

In recent decades there have been efforts across multiple fields to harness quantum mechanics for application to problems outside of fundamental physics. Perhaps the most prominent example can be found in quantum computing, which seeks to develop algorithms designed to perform on specially built hardware in which quantum bits (qubits) perform computational

operations while in a quantum state. Quantum computers can leverage distinctly quantum processes such as superposition and entanglement to achieve superior performance over a classical computer for certain types of problems.⁵⁸

Quantum mechanics has also been used as a predictive framework to describe abstract processes which are not inherently “quantum” in nature. Quantum annealing (QA) was initially proposed by Kadowaki and Nishimori as an optimization method which models a quantum system under adiabatic conditions (i.e., no transfer of energy or mass between a process and its surroundings) to find the global minimum of an objective function; it has since been explored as an optimization method for radiation therapy treatment planning.⁵⁹ There has also been success in utilizing quantum mechanics as a framework for modeling human cognition and decision making.⁶⁰ It has been shown that quantum models are advantageous for modeling scenarios in which human decision-making violates the laws of classical probability, which have been found to occur under conditions of uncertainty or when sequential decisions are subject to order effects.⁶¹ Based on these previous works, a motivation in the studies of this dissertation were, when appropriate, to investigate the use of quantum mechanics in conjunction with deep learning as a unified framework for addressing the challenges of ART.

1.4 Motivation and Contributions of this Study

This dissertation seeks to address some of the pressing challenges associated with wider implementation of adaptive radiotherapy—specifically the need for efficient and robust methods for treatment plan optimization, the need for robust frameworks for predicting patient anatomical variations, and finally the need to understand what kinds of patient characteristics can play a role in clinical outcomes. The following sections provide a brief description of the studies presented as well as their limitations and contributions.

1.4.1 Treatment Planning Optimization

The first step in a radiotherapy process after consultation is the development of an appropriate treatment plan using virtual simulations to ensure its accurate delivery. This can be a laborious and time-consuming process, especially for ART applications where the plan may be updated/adapted several times during the course of the treatment. Hence, the focus of this study was to develop and investigate the use of a quantum-inspired algorithm, quantum tunnel annealing (QTA), for more efficient radiotherapy treatment plan optimization. We found that QTA converged up to 26.8% and 46.6% faster than the traditional simulated annealing when applied to direct aperture optimization and beamlet weight optimization, respectively. The results suggest that QTA could be an attractive optimization method in radiotherapy under scenarios that are currently hindered by lack of robust and efficient optimization. However, further studies will be necessary to better understand the performance of QTA under a wider variety of conditions, including different patient populations, additional structure sets, and objectives.

1.4.2 Adaptation of Radiotherapy Plans to Geometrical Changes

Patients undergoing radiotherapy may experience physiological and anatomical changes over the course of their treatment. Thus, the focus of this study was to develop and investigate two deep machine learning frameworks (one quantum-based and the other Markov-based) for predicting these anatomical changes during fractionated radiotherapy so that an optimal outcome can be achieved. This was evaluated in a population of head and neck cancer patients, which represent a population that typically experiences anatomical variations requiring plan adaption. The MRNN framework was found to have slightly better performance than the QRNN framework, with MRNN(QRNN) validation AUC scores of 0.742 ± 0.021 (0.675 ± 0.036), 0.709 ± 0.026 (0.656 ± 0.021), 0.724 ± 0.036 (0.652 ± 0.044), and 0.698 ± 0.016 (0.605 ± 0.035) for system state vector

sizes of 4, 6, 8, and 10, respectively. Of these, only the results from the two higher order states were found to have statistically significant differences ($p < 0.05$). A similar trend was also observed when the fully trained models were applied to an external testing dataset of 20 patients, yielding MRNN(QRNN) AUC scores of 0.707 (0.623), 0.687 (0.608), 0.723 (0.669), and 0.697 (0.609) for states vectors sizes of 4, 6, 8, and 10, respectively. These results suggest that stochastic models have potential value in predicting patient changes during the course of adaptive radiotherapy.

1.4.3 Radiotherapy Outcomes and Feature Identification

It is recognized that different patients respond differently to radiotherapy. Thus, in the quest to develop personalized treatment plans, it is necessary to predict beforehand how an individual person would respond to radiotherapy. Hence, in this study we developed machine learning and deep learning models to identify patient clinical and imaging features which could be tied to relevant clinical outcomes in a larynx cancer population. Two features (gross tumor volume change between pre- and post-induction chemotherapy, and nodal stage) had prognostic value for predicting patient outcomes using standard regression models. The use of deep learning models was not found to improve model performance, but future studies incorporating more patient data could help to overcome model variance in this case.

1.5 Dissertation Organization

The remainder of this dissertation is organized as follows: Chapter 2 introduces relevant background on the topics and methodologies featured in this work, including fundamentals of radiation therapy, treatment planning, and adaptive radiotherapy; features relevant to models in radiation therapy; optimization methods; machine and deep learning models; and best practices for model evaluation. Chapter 3 presents a study on a quantum-inspired optimization algorithm for

radiotherapy treatment planning. Chapter 4 presents a study on a stochastic predictive framework for patient anatomical variations, and Chapter 5 presents a study on identifying relevant patient features predictive of outcomes in larynx cancer patients. Finally, Chapter 6 provides a general discussion of the contributions and limitations of the work presented in this dissertation.

1.6 Accomplishments

1.6.1 Awards

- Medical Physics Distinguished Reviewer (2021)
- Alternate for Rackham Predoctoral Fellowship (2020)
- Medical Physics Editor’s Choice (2020): First author of article “Quantum-inspired algorithm for radiotherapy planning optimization”
- Early Career Medical Physicist Scholarship Award (2019): Awarded at the 2019 Winter Institute of Medical Physics
- University of Michigan Regents’ Fellowship (2016-2018)

1.6.2 Peer-reviewed Publications

- Pakela, J. M., Matuszak, M. M., Ten Haken, R. K., McShane, D. L., El Naqa, I. (expected early 2021) Dynamic Stochastic Deep Learning Approaches for Predicting Geometric Changes in Head and Neck Cancer. (submitted).

- Gharzai, L. A., Pakela, J. M., Jaworski, E., El-Naqa, I., Schonewolf, C. A., Hawkins, P. G., Wilkie, J. R., Spector, M. E., Bradford, C. R., Chinn, S. B., Hogikyan, N., Malloy, K., Kupfer, R., Shuman, A., Stucken, C. L., Prince, M., Shah, J., Swiecicki, P. L., Casper, K., Eisbruch, A., Wolf, G., Worden, F., Srinivasan, A., Mierzwa, M. L. (expected 2020/early 2021) Imaging response versus operative laryngoscopy assessment of induction chemotherapy response in an induction bioselection approach to larynx cancer. (submitted).
- Lee, S. Y., Pakela, J. M., Na, K., Shi, J., McKenna, B. J., Simeone, D. M., Yoon, E., Scheiman, J. M., Mycek, M-A. (2020), Needle-compatible miniaturized optoelectronic sensor for pancreatic cancer detection, *Sci. Adv.*, 6(47): eabc1746.
- Cui, S., Tseng, H.-H., Pakela, J., Ten Haken, R. K. and El Naqa, I. (2020), Introduction to machine and deep learning for medical physicists. *Med. Phys.*, 47: e127-e147.
- Pakela J. M., Tseng H. H., Matuszak M. M., Ten Haken R. K., McShan D. L., El Naqa I. (2020), Quantum-inspired algorithm for radiotherapy planning optimization. *Med. Phys.*, 47(1): 5-18.
- Lee, S. Y.*, Pakela, J. M.*, Helton, M. C., Vishwanath, K., Chung, Y. G., Kolodziejcki, N. J., Stapels, C. J., McAdams, D. R., Fernandez, D. E., Christian, J. F., O'Reilly, J., Farkas, D., Ward, B. B., Feinberg, S. E., Mycek, MA. (December 2017). Compact dual-mode diffuse optical system for blood perfusion monitoring in a porcine model of microvascular tissue flaps. *J. Biomed. Opt.*, 22(12): 121609(1-14).

* indicates co-first authorship

1.6.3 Book Chapters

- Pakela J., El Naqa, I. (anticipated early 2021). Overview of Deep Machine Learning Methods. In I. El Naqa & M. J. Murphy (Eds.), *Machine Learning in Radiation Oncology: Theory and Applications* (2nd ed.). Springer International Publishing.
- Niraula D., Jamaluddin J., Pakela J., (anticipated early 2021). Quantum Computing for Machine Learning. In I. El Naqa & M. J. Murphy (Eds.), *Machine Learning in Radiation Oncology: Theory and Applications* (2nd ed.). Springer International Publishing.

1.6.4 Presentations

- Pakela, J. M., Gharzai, L. A., Jaworski, E., Schonewolf, C. A., Hawkins, P. G., Bradford, C. R., Chinn, S. B., Hogikyan, N., Malloy, K., Kupfer, R., Spector, M. E., Shuman, A., Stucken, C. L., Prince, M. E., Shah, J. L., Srinivansan, A., Swiecicki, P. L., Casper, K., Worden, F., Wolf, G. T., Mierzwa, M. L., El Naqa, I. (October 2020). *Roles of clinical and image-based features for outcomes in larynx cancer patients*.
Poster presentation at the American Society for Radiation Oncology (ASTRO) Annual Meeting, (virtual).
- Pakela, J. M., Ten Haken, R. K., Mcshan, D. L., Matuszak, M. M., El Naqa, I. (July 2020). *A quantum-inspired approach to predicting geometric changes in head and neck cancer*.
Blue ribbon poster presentation at the American Association for Physicists in Medicine (AAPM) Annual Meeting, (virtual).

- Jaworski, E., Gharzai, L. A., Pakela, J. M., El Naqa, I., Schonewolf, C. A., Hawkins, P. G., Bradford, C. R., Chinn, S. B., Hogikyan, N., Malloy, K., Kupfer, R., Spector, M. E., Shuman, A. Stucken, C. L., Prince, M. E., Srinivasan, A., Swiecicki, P. L., Casper, K., Worden, F., Mierzwa, M. L. (February 2020). *Imaging response versus operative laryngoscopy assessment of induction chemotherapy response in an induction bioselection approach to larynx cancer.*
Oral presentation at the Multidisciplinary Head and Neck Symposium, Scottsdale, Arizona.
- Pakela, J. M., Tseng, H. -H., Matuszak, M. M., Ten Haken, R. K., McShan, D. L., El Naqa, I. (June 2019). *A quantum inspired algorithm for radiotherapy planning optimization.*
Oral presentation at the International Conference on the Use of Computers in Radiation Therapy (ICCR), Montreal, Canada.
- Pakela, J. M., Tseng, H. -H., Matuszak, M. M., Ten Haken, R. K., McShan, D. L., El Naqa, I. (February 2019). *Quantum-inspired algorithms for adaptive radiotherapy.*
Poster presentation at the Winter Institute of Medical Physics (WIMP), Breckenridge, CO.
- Pakela, J. M., El Naqa, I., Matuszak, M. M., Ten Haken, R. K., McShan, D. L. (July 2018). *Quantum Annealing for IMRT Optimization.*
Oral presentation at the American Association for Physicists in Medicine (AAPM) Annual Meeting, Nashville, TN.

- Pakela, J. M., Lee, S. Y., Hedrick, T. L., Vishwanath, K., Helton, M. C., Chung, Y. G., Kolodziejcki, N. J., Staples, C. J., McAdams, D. R., Fernandez, D. E., Christian, J. F., O'Reilly, J., Farkas, D., Ward, B. B., Feinberg, S. E., Mycek, M.-A. (February 2017). *In vivo preclinical verification of a multimodal diffuse reflectance and correlation spectroscopy system for sensing tissue perfusion*. Oral presentation at SPIE Photonics West, San Francisco, CA.
- Pakela, J. M., Hedrick, T. L., Lee, S. Y., Vishwanath, K., Zanfardino, S., Chung, Y. G., Helton, M. C., Kolodziejcki, N. J., McAdams, D. R., Fernandez, D. E., Christian, J. F., Feinberg, S. E., Mycek, M.-A. (February 2017). *Design verification of a compact system for detecting tissue perfusion using bimodal diffuse optical technologies*. Poster presentation at SPIE Photonics West, San Francisco, CA.

Underline indicates presenter

1.7 References

1. Sonke J-J, Aznar M, Rasch C. Adaptive Radiotherapy for Anatomical Changes. *Seminars in Radiation Oncology*. 2019;29(3):245-257.
2. Brock KK. Adaptive Radiotherapy: Moving Into the Future. *Seminars in Radiation Oncology*. 2019;29(3):181-184.
3. van Kranen S, van Beek S, Rasch C, van Herk M, Sonke J-J. Setup uncertainties of anatomical sub-regions in head-and-neck cancer patients after offline CBCT guidance. *International Journal of Radiation Oncology* Biology* Physics*. 2009;73(5):1566-1573.
4. Barker Jr JL, Garden AS, Ang KK, et al. Quantification of volumetric and geometric changes occurring during fractionated radiotherapy for head-and-neck cancer using an integrated CT/linear accelerator system. *International Journal of Radiation Oncology* Biology* Physics*. 2004;59(4):960-970.
5. Sonke J-J, Belderbos J. Adaptive radiotherapy for lung cancer. 2010.
6. Wang JZ, Li JB, Wang W, et al. Detection of interfraction displacement and volume variance during radiotherapy of primary thoracic esophageal cancer based on repeated four-dimensional CT scans. *Radiation Oncology*. 2013;8(1):1-7.
7. Langerak T, Mens JW, Quint S, et al. Cervix motion in 50 cervical cancer patients assessed by daily cone beam computed tomographic imaging of a new type of marker. *International Journal of Radiation Oncology* Biology* Physics*. 2015;93(3):532-539.

8. Dees-Ribbers HM, Betgen A, Pos FJ, Witteveen T, Remeijer P, van Herk M. Inter-and intra-fractional bladder motion during radiotherapy for bladder cancer: A comparison of full and empty bladders. *Radiotherapy and Oncology*. 2014;113(2):254-259.
9. Lee C, Langen KM, Lu W, et al. Assessment of parotid gland dose changes during head and neck cancer radiotherapy using daily megavoltage computed tomography and deformable image registration. *International Journal of Radiation Oncology* Biology* Physics*. 2008;71(5):1563-1571.
10. Zhao L, Wan Q, Zhou Y, Deng X, Xie C, Wu S. The role of replanning in fractionated intensity modulated radiotherapy for nasopharyngeal carcinoma. *Radiotherapy and Oncology*. 2011;98(1):23-27.
11. McDonald BA, Vedam S, Yang J, et al. Initial Feasibility and Clinical Implementation of Daily MR-guided Adaptive Head and Neck Cancer Radiotherapy on a 1.5 T MR-Linac System: Prospective R-IDEAL 2a/2b Systematic Clinical Evaluation of Technical Innovation. *International Journal of Radiation Oncology* Biology* Physics*. 2020.
12. Brouwer CL, Steenbakkers RJHM, Langendijk JA, Sijtsema NM. Identifying patients who may benefit from adaptive radiotherapy: Does the literature on anatomic and dosimetric changes in head and neck organs at risk during radiotherapy provide information to help? *Radiotherapy and Oncology*. 2015;115(3):285-294.
13. Sejnowski TJ. The unreasonable effectiveness of deep learning in artificial intelligence. *Proceedings of the National Academy of Sciences*. 2020:201907373.
14. Wang H, Raj B. On the origin of deep learning. *arXiv preprint arXiv:170207800*. 2017.
15. Krizhevsky A, Sutskever I, Hinton GE. Imagenet classification with deep convolutional neural networks. 2012.
16. Sejnowski TJ. *The Deep Learning Revolution*. MIT Press; 2018.
17. El Naqa I, Haider MA, Giger ML, Ten Haken RK. Artificial Intelligence: reshaping the practice of radiological sciences in the 21st century. *Br J Radiol*. 2020;93(1106):20190855.
18. Doi K. Computer-aided diagnosis in medical imaging: historical review, current status and future potential. *Comput Med Imaging Graph*. 2007;31(4-5):198-211.
19. Doi K, ML G, Nishikawa R, MacMahon H, Schmidt R. Artificial intelligence and neural networks in radiology: Application to computer-aided diagnostic schemes. In: Hendee W, Trueblood J, eds. *Digital Imaging*. Vol 2. AAPM Medical Physics Monograph; 1993:301-322.
20. Giger M, Huo Z, Kupinski M, Vyborny C. Computer-aided diagnosis in mammography. In: Sonka M, Fitzpatrick M, eds. *Handbook of Medical Imaging*. Vol 2. SPIE; 2000:915-1004.
21. Giger ML. Future of breast imaging. Computer-aided diagnosis. In: Haus A, Yaffe M, eds. *AAPM/RSNA Categorical Course on the Technical Aspects of Breast Imaging*. 1992:257-270.
22. Giger ML. Computer-aided diagnosis in radiology. *Academic radiology*. 2002;9(1):1-3.
23. Swett H, Giger M, Doi K. Computer vision and decision support. In: Hendee W, Wells P, eds. *Perception of Visual Information*. Springer-Verlag; 1993:272-315.
24. Wu Y, Doi K, Giger ML, Nishikawa RM. Computerized detection of clustered microcalcifications in digital mammograms: applications of artificial neural networks. *Medical physics*. 1992;19(3):555-560.

25. Chan HP, Lo SC, Sahiner B, Lam KL, Helvie MA. Computer-aided detection of mammographic microcalcifications: pattern recognition with an artificial neural network. *Medical physics*. 1995;22(10):1555-1567.
26. Sahiner B, Chan HP, Petrick N, et al. Classification of mass and normal breast tissue: a convolution neural network classifier with spatial domain and texture images. *IEEE transactions on medical imaging*. 1996;15(5):598-610.
27. Zhang W, Doi K, Giger ML, Wu Y, Nishikawa RM, Schmidt RA. Computerized detection of clustered microcalcifications in digital mammograms using a shift-invariant artificial neural network. *Medical physics*. 1994;21(4):517-524.
28. Chen W, Giger ML, Newstead GM, et al. Computerized assessment of breast lesion malignancy using DCE-MRI robustness study on two independent clinical datasets from two manufacturers. *Academic radiology*. 2010;17(7):822-829.
29. Chen W, Giger ML, Bick U. A fuzzy c-means (FCM)-based approach for computerized segmentation of breast lesions in dynamic contrast-enhanced MR images. *Academic radiology*. 2006;13(1):63-72.
30. Bhooshan N, Giger ML, Jansen SA, Li H, Lan L, Newstead GM. Cancerous breast lesions on dynamic contrast-enhanced MR images: computerized characterization for image-based prognostic markers. *Radiology*. 2010;254(3):680-690.
31. Yuan Y, Giger ML, Li H, Bhooshan N, Sennett CA. Multimodality computer-aided breast cancer diagnosis with FFDM and DCE-MRI. *Academic radiology*. 2010;17(9):1158-1167.
32. IonQ Inc. A true quantum leap. <https://ionq.co/>. Published 2018. Accessed August 2018.
33. Lundervold AS, Lundervold A. An overview of deep learning in medical imaging focusing on MRI. *Zeitschrift für Medizinische Physik*. 2019;29(2):102-127.
34. Akagi M, Nakamura Y, Higaki T, et al. Deep learning reconstruction improves image quality of abdominal ultra-high-resolution CT. *European radiology*. 2019;29(11):6163-6171.
35. Huang P, Yu G, Lu H, et al. Attention-aware fully convolutional neural network with convolutional long short-term memory network for ultrasound-based motion tracking. *Medical physics*. 2019;46(5):2275-2285.
36. Graffy PM, Sandfort V, Summers RM, Pickhardt PJ. Automated Liver Fat Quantification at Nonenhanced Abdominal CT for Population-based Steatosis Assessment. *Radiology*. 2019:190512.
37. Haskins G, Kruger U, Yan P. Deep learning in medical image registration: a survey. *Machine Vision and Applications*. 2020;31(1):8.
38. Giger ML. Machine Learning in Medical Imaging. *Journal of the American College of Radiology : JACR*. 2018;15(3 Pt B):512-520.
39. Fan J, Wang J, Chen Z, Hu C, Zhang Z, Hu W. Automatic treatment planning based on three-dimensional dose distribution predicted from deep learning technique. *Medical physics*. 2019;46(1):370-381.
40. Liu F, Yadav P, Baschnagel AM, McMillan AB. MR-based treatment planning in radiation therapy using a deep learning approach. *Journal of Applied Clinical Medical Physics*. 2019;20(3):105-114.
41. Shen C, Gonzalez Y, Klages P, et al. Intelligent inverse treatment planning via deep reinforcement learning, a proof-of-principle study in high dose-rate brachytherapy for cervical cancer. *Physics in Medicine & Biology*. 2019;64(11):115013.

42. Liu Y, Lei Y, Wang T, et al. MRI-based treatment planning for liver stereotactic body radiotherapy: validation of a deep learning-based synthetic CT generation method. *The British Journal of Radiology*. 2019;92(1100):20190067.
43. Elmahdy MS, Jagt T, Zinkstok RT, et al. Robust contour propagation using deep learning and image registration for online adaptive proton therapy of prostate cancer. *Medical physics*. 2019;46(8):3329-3343.
44. Tseng H-H, Luo Y, Cui S, Chien J-T, Ten Haken RK, Naqa IE. Deep reinforcement learning for automated radiation adaptation in lung cancer. *Medical physics*. 2017;44(12):6690-6705.
45. Chun J, Zhang H, Gach HM, et al. MRI super-resolution reconstruction for MRI-guided adaptive radiotherapy using cascaded deep learning: In the presence of limited training data and unknown translation model. *Medical physics*. 2019;46(9):4148-4164.
46. Kurz C, Maspero M, Savenije MHF, et al. CBCT correction using a cycle-consistent generative adversarial network and unpaired training to enable photon and proton dose calculation. *Physics in Medicine & Biology*. 2019;64(22):225004.
47. Nyflot MJ, Thammasorn P, Wootton LS, Ford EC, Chaovalitwongse WA. Deep learning for patient-specific quality assurance: Identifying errors in radiotherapy delivery by radiomic analysis of gamma images with convolutional neural networks. *Medical physics*. 2019;46(2):456-464.
48. Tomori S, Kadoya N, Takayama Y, et al. A deep learning-based prediction model for gamma evaluation in patient-specific quality assurance. *Medical physics*. 2018;45(9):4055-4065.
49. Galib SM, Lee HK, Guy CL, Riblett MJ, Hugo GD. A fast and scalable method for quality assurance of deformable image registration on lung CT scans using convolutional neural networks. *Medical physics*. 2020;47(1):99-109.
50. Kimura Y, Kadoya N, Tomori S, Oku Y, Jingu K. Error detection using a convolutional neural network with dose difference maps in patient-specific quality assurance for volumetric modulated arc therapy. *Physica Medica*. 2020;73:57-64.
51. Hosny A, Parmar C, Coroller TP, et al. Deep learning for lung cancer prognostication: A retrospective multi-cohort radiomics study. *PLoS Med*. 2018;15(11):e1002711-e1002711.
52. Wei L, Osman S, Hatt M, El Naqa I. Machine learning for radiomics-based multimodality and multiparametric modeling. *The quarterly journal of nuclear medicine and molecular imaging : official publication of the Italian Association of Nuclear Medicine (AIMN) [and] the International Association of Radiopharmacology (IAR), [and] Section of the So*. 2019;63(4):323-338.
53. Bibault J-E, Giraud P, Housset M, et al. Deep Learning and Radiomics predict complete response after neo-adjuvant chemoradiation for locally advanced rectal cancer.
54. Goodfellow I, Pouget-Abadie J, Mirza M, et al. Generative adversarial nets. 2014.
55. Shen W-C, Chen S-W, Wu K-C, et al. Prediction of local relapse and distant metastasis in patients with definitive chemoradiotherapy-treated cervical cancer by deep learning from [18F]-fluorodeoxyglucose positron emission tomography/computed tomography. *European radiology*. 2019;29(12):6741-6749.
56. Cui S, Luo Y, Hsin Tseng H, Ten Haken RK, El Naqa I. Artificial Neural Network with Composite Architectures for Prediction of Local Control in Radiotherapy. *IEEE Trans Radiat Plasma Med Sci*. 2019;3(2):242-249.

57. Cui S, Luo Y, Tseng H-H, Ten Haken RK, El Naqa I. Combining handcrafted features with latent variables in machine learning for prediction of radiation-induced lung damage. *Medical physics*. 2019;46(5):2497-2511.
58. Nielsen MA, Chuang IL. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press; 2010.
59. Kadowaki T, Nishimori H. Quantum annealing in the transverse Ising model. *Physical Review E*. 1998;58(5):5355.
60. Busemeyer JR, Bruza PD. *Quantum Models of Cognition and Decision*. Cambridge University Press; 2012.
61. Bruza PD, Wang Z, Busemeyer JR. Quantum cognition: a new theoretical approach to psychology. *Trends in cognitive sciences*. 2015;19(7):383-393.

Chapter 2 Background

This chapter contains excerpts from two book chapters and one review paper:

1. Pakela J., El Naqa, I. (anticipated early 2021). Overview of Deep Machine Learning Methods. In I. El Naqa & M. J. Murphy (Eds.), *Machine Learning in Radiation Oncology: Theory and Applications* (2nd ed.). Springer International Publishing.
2. Niraula D., Jamaluddin J., Pakela J., (anticipated early 2021). Quantum Computing for Machine Learning. In I. El Naqa & M. J. Murphy (Eds.), *Machine Learning in Radiation Oncology: Theory and Applications* (2nd ed.). Springer International Publishing.
3. Cui, S., Tseng, H.-H., Pakela, J., Ten Haken, R. K. and El Naqa, I. (2020), Introduction to machine and deep learning for medical physicists. *Med. Phys.*, 47: e127-e147.

2.1 Radiation Therapy

Radiation therapy is a form of cancer treatment in which ionizing radiation is used to target and eradicate cancerous cells. Because ionizing radiation can also damage normal tissues, a major challenge of radiation therapy is to maximize the therapeutic window—i.e., maximize probability of local tumor control while minimizing the probability of normal tissue complications.¹ Radiation therapy can be delivered in the form of implantable ionizing sources (brachytherapy) or through external beams of ionizing particles such as photons, electrons, or protons (external beam radiotherapy). An ideal treatment would involve delivering as much dose as possible to the tumor region with negligible dose to the surrounding normal tissue. Because it is not physically possible to perfectly localize radiation to only the tumor, in practice there is a tradeoff between delivering

as much dose as is necessary to eradicate the cancer while avoiding harmful dose levels to critical structures that could lead to side effects.

Advancements in computing over the past several decades have helped to revolutionize radiotherapy (RT) treatment. In the late 1990s and early 2000s, the ability to efficiently perform optimizations on objective functions with hundreds or thousands of degrees of freedom allowed for the adaption of multileaf collimators (MLCs) into the treatment delivery process, leading to the rise of intensity modulated radiation therapy (IMRT) and volumetric arc therapy (VMAT). Both IMRT and VMAT achieve dose distributions which conform tightly around irregularly shaped/concave target structures, providing improved sparing of normal tissues and allowing the tumor to be safely treated with higher doses.² These dosimetric improvements have resulted in IMRT methods becoming a regularly used standard of care.³ Advancements in computing have also allowed for the rise of image-guided radiotherapy—in which imaging is performed directly on the treatment couch right before or during therapy, allowing for improved alignment and motion management during treatment.⁴

2.1.1 Types of External Beam Radiation Therapy

As mentioned in the previous section, external beam radiotherapy (EBRT) involves targeting the tumor with beams of high energy photons, electrons, or protons. Due to their energy deposition properties, photons and protons are used for deep seated tumors, while electrons are reserved for superficial disease. EBRT is most commonly delivered in the form of photons due the prevalence of deep-seated tumors as well as the comparatively higher costs associated with proton therapy.⁵ IMRT and VMAT represent two major modalities for photon delivery. Unlike 3-dimensional conformal radiation therapy (3DCRT), which utilizes multiple static beams, IMRT and VMAT beams are dynamically shaped during treatment using a device called a multi-leaf collimator

(MLC) which consists of movable pairs of radiation-blocking tungsten leaves. VMAT can be considered a type of IMRT in which the gantry also moves during treatment delivery across a pre-determined arc.⁶ In addition to being capable of achieving highly conformal dose distributions associated with IMRT techniques, VMAT allows for each fraction of treatment to be delivered over a shorter time span.⁷

2.1.2 Treatment Planning

Treatment planning describes the process in the radiotherapy workflow during which a treatment plan is designed to meet the dosimetric constraints prescribed by the radiation oncologist. Typically, the dosimetric constraints consist of lower dose limits for diseased tissue and upper dose limits for normal tissues and critical structures. There are three categories of volumes which receive lower dose limits in radiotherapy. The gross tumor volume (GTV) is defined as the region in which the tumor can be physically observed via imaging. In order to account for uncertainties regarding the presence of microscopic disease spreading from the edges of the viewable tumor, a margin is drawn around the GTV, and the volume contained within this margin is referred to as the clinical target volume (CTV). Finally, to account for random motion uncertainties associated with daily setup and anatomical variations, an additional margin is drawn around the CTV; the volume within is referred to as the planning target volume (PTV).⁸

For EBRT, treatment planning requires the selection or optimization of several different types of parameters including the number of beams to use, the orientation of those beams (in the case of VMAT, the arc path through which the beam traverses), as well as fluence vectors or aperture shapes (which determine how much radiation the MLC allows to flow through each beam aperture and where). While certain parameters (such as number of beams, beam orientations, or arc paths) can either be optimized or selected by an experienced treatment planner, others (such as

fluence vectors/aperture shapes) represent thousands of degrees of freedom and thus require the use of computational optimization algorithms. The time and resource constraints associated with a busy clinical environment as well as the growing interest in incorporating complex, biologically driven objectives result in a need for robust, efficient optimization algorithms.

2.1.3 Adaptive Radiotherapy for Head and Neck Cancer

The concept of adaptive radiotherapy (ART) was first proposed in the literature by Yan et al. in 1997 as a means of delivering personalized treatment and safely performing dose escalation through the modification of radiotherapy plans mid-treatment in response to measured changes in the patient.⁹ ART methods are categorized by the timeframe over which the adaptation is made. In *online* ART, plan adaptations are made immediately before a treatment while the patient is immobilized on the treatment couch and are performed over a timescale of minutes. In *offline* ART, plan adaptations are made between treatment fractions and can occur over a timescale ranging from hours to days. Because CBCT imaging has become a ubiquitous standard of care to ensure proper daily alignment with the treatment plan, it is common practice to utilize CBCT images to perform *geometric adaptations*, in which adaptations are made based on measured changes to the patient's geometry.¹⁰ There has also been interest in utilizing functional imaging such as Positron Emission Tomography (PET) or diffusion-weighted magnetic resonance imaging (MRI) to perform *biological adaptations*, in which biomarkers guide ART.¹¹

Figure 2.1 provides a visualization of the radiotherapy workflow. The teal boxes represent the standard (i.e., non-adaptive) RT workflow while the white boxes represent an adaptive workflow, highlighting the primary questions and uncertainties which still need to be addressed for clinical implementation, including: “What types of information should be used for determining

the need for adaptation?”, “What should be the necessary criteria to trigger a plan adaption?”, and “What kind of adaption should occur, when, and how often?”.

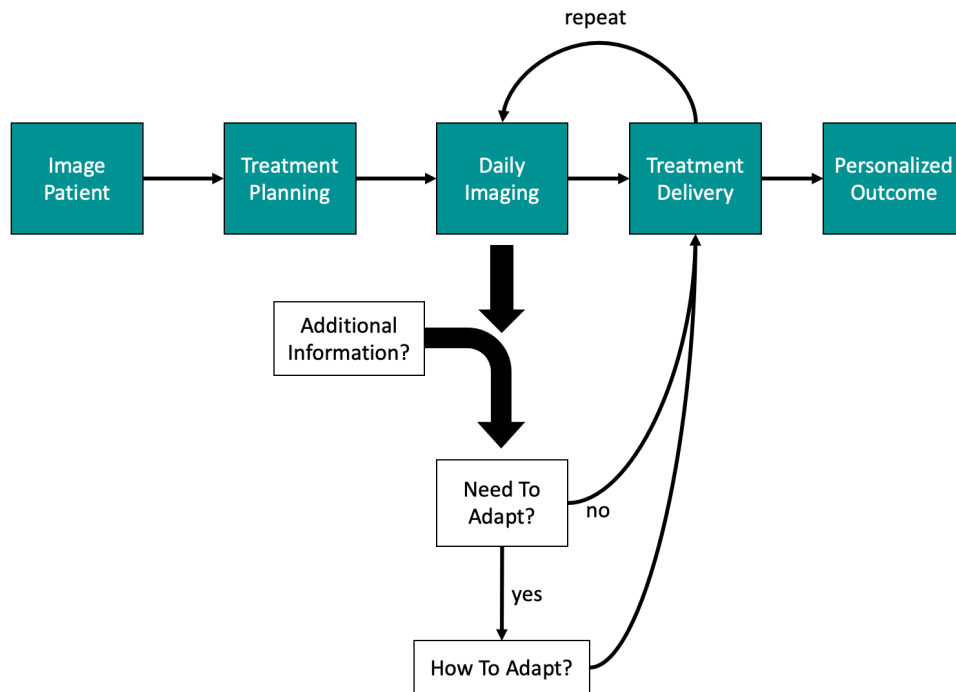


Figure 2.1: Visualization of radiotherapy workflow highlighting some of the primary questions and uncertainties surrounding the implementation of ART techniques.

Given that the clinical implementation of ART represents a complex, multifaceted problem, it has been the subject of many different areas in research. One broad area of research is the identification of patient populations who are most likely to benefit from ART procedures.^{12,13} Another critical area of ART research is to develop technology which allows ART to be implemented within time and resource constraints of a clinical setting. This technology can take the form of algorithms for improving daily pre-treatment image quality,¹⁴⁻¹⁶ improved image registration techniques for accurate mapping of structures,¹⁷⁻¹⁹ algorithms for faster treatment plan optimization and dose estimation,²⁰⁻²² and outcome prediction as well as optimal treatment decision models to aid in clinical decision making.²³⁻²⁹

2.2 Handcrafted Features

In machine learning, a *feature* refers to a quantitative, measurable property which serves as an input to a model.³⁰ Handcrafted features are variables which are selected manually using expert knowledge. Commonly used handcrafted features in radiotherapy modeling include clinical and anatomical patient features, including demographics such as age and gender, genetic information (genomics), disease-specific features such as tumor site and disease stage, treatment-specific features such as dose prescription, and quantitative radiological imaging features (radiomics) which are discussed in further detail below.

2.2.1 Radiomic Features

Radiomic features are quantitative variables extracted from medical-image data (typically radiology images such as computed tomography [CT], positron emission tomography [PET], or magnetic resonance imaging [MRI]).³¹ Radiomic features include morphological features such as volume or surface area as well as texture features which are calculated using the intensity levels within a region of interest (ROI) (such as the 3D CT contour of the GTV) of an isotropic image. Texture features can be further classified into global (first order) features and local (higher order) features. Global texture features can be calculated directly from the normalized intensity distribution histogram p and the average gray-level. Given an isotropic volume contained within an ROI, and $P(i)$ representing the number of voxels with gray level, i , the i th entry of the normalized intensity distribution histogram can be defined as:

$$p(i) = \frac{P(i)}{\sum_{i=1}^{N_g} P(i)} \quad 2.1$$

Where N_g is the number of gray level bins.³²

Higher order texture features include Gray-Level Co-occurrence Matrix (GLCM) features, Gray-Level Run-Length Matrix (GLRLM) features, Gray-Level Size Zone Matrix (GLSZM) features, and Neighborhood Gray-Tone Difference Matrix (NGTDM) features. Higher order texture features differ from global texture features in that each of the matrices used to calculate these features include a distance parameter in the computation, allowing them to capture spatial patterns and interrelationships between voxels.³¹ Image processing and feature extraction for radiomic features in Chapter 5 were performed using the Radiomics toolbox in MATLAB.³³

2.3 Optimization Algorithms

Optimization problems can be formulated as the goal of finding the minimum value of an objective function $E(x)$ with respect to a set of optimization parameters, \mathbf{w} :

$$\min_{\mathbf{w}} E(\mathbf{w}) \quad 2.2$$

In radiotherapy treatment planning, the optimization parameters, \mathbf{w} , can represent a number of different variables, such as fluence weights over the beam's cross sections, MLC aperture shapes, or beam orientations. One challenge in radiotherapy treatment planning optimization is that the objective function often contains terms such as dose volume penalties which render it non-convex, making optimization more difficult because the solution space may contain many local minima and is not guaranteed to have a single global minimum.

2.3.1 Gradient Descent

Gradient descent is an iterative optimization method which searches for the local minimum of an objective function by updating the solution via the gradient of the objective function with respect to the optimization parameters:

$$\mathbf{w}^{(\tau+1)} = \mathbf{w}^{(\tau)} - \eta \nabla E(\mathbf{w}^{(\tau)}) \quad 2.3$$

Where η serves as the learning rate coefficient and is usually chosen as a small number. One limitation of gradient descent is that the solution can become trapped in local minima for non-convex optimization functions.³⁴

2.3.2 Simulated Annealing

Simulated annealing (SA) is an iterative, stochastic optimization algorithm which approximates a global minimum of a function.³⁵ After initializing the optimization parameters as a first guess, a new potential solution, w' , is randomly selected from the neighborhood of the current solution space. If the potential solution yields a lower objective function score than the current solution ($E(w') < E(w)$), the new solution is accepted and set as the current solution. If the potential solution is found to be worse than the current solution ($E(w') > E(w)$), the potential solution is accepted as the current solution with probability $P = \exp\left(\frac{-\Delta E}{T}\right)$, where T is an annealing parameter that represents the “temperature” of the system and is gradually decreased so that it approaches 0 by the end of the algorithm. By allowing the algorithm to occasionally accept worse solutions, SA is able to avoid becoming trapped in local minima. It has been shown theoretically that if T is decreased slowly enough, SA is guaranteed to find the global minimum of E .

2.4 Classical Machine Learning Algorithms

Classical machine learning refers to any machine learning algorithm which use handcrafted features (also called feature vectors) to identify patterns in data. This section provides a brief overview of classical machine learning methods utilized in this work.

2.4.1 Logistic Regression

Logistic regression (LR) is a binary classification model which can be used for outcome prediction, disease classification, and as a way to estimate the relative importance between various patient

features and clinical outcomes. LR assumes the natural log of the odds ratio (also called the logit) is a linear function of the feature vector \mathbf{x} :

$$\ln \left(\frac{\pi}{1 - \pi} \right) = \boldsymbol{\beta} \cdot \mathbf{x} \quad 2.4$$

Where π represents the probability of outcome $Y = 1$, $1 - \pi$ represents the probability of the alternative outcome $Y = 0$, and $\boldsymbol{\beta}$ represents regression coefficients learned during model training (typically using the maximum likelihood estimation).³⁶ The posterior probability of a sample with feature vector \mathbf{x} can thus be written as:

$$P(Y = 1|\mathbf{x}) = \sigma(\boldsymbol{\beta} \cdot \mathbf{x}) \quad 2.5$$

Where $\sigma(a)$ is the logistic sigmoid function:

$$\sigma(a) = \frac{1}{1 + \exp(-a)} \quad 2.6$$

2.4.2 Cox Proportional Hazards Model

For certain kinds of medical data, it is also relevant to consider the “time to event” associated with binary outcomes. For a set of patients each characterized by a feature vector (or covariates), \mathbf{x} , the Cox proportional hazards (Cox) model assumes the risk of an event occurring for a given patient at time, t , to take the form of the hazard function:

$$h(t) = h_0(t) * \exp(\boldsymbol{\beta} \cdot \mathbf{x}) \quad 2.7$$

Where $h_0(t)$ is the baseline hazard, and the set of covariate coefficients, $\boldsymbol{\beta}$, are fit during model training and represent the relative importance of each feature in \mathbf{x} .³⁷ The values $\exp(\beta_i)$ are referred to as hazard ratios and describe how each covariate contributes to the risk of event, with $\exp(\beta_i)$ greater than (less than) 1 indicating that the risk to the patient increases (decreases) as x_i increases.

2.5 Deep Learning

Deep machine learning or “deep learning” refers to a class of machine learning methods which take raw data as inputs and, through training, learn multiple layers of relevant latent features to map the raw inputs to the desired output space; the desired mapping is defined by either a reward or a loss function of the outputs for detection or classification tasks.³⁸ This is in contrast to shallow/classical machine learning, where the features are manually crafted and do not contain multiple layers of abstraction. Conceptually then, deep learning can be applied to any machine learning technology as depicted in **Figure 2.2**, but as of this time it has been practically shown to be most effective with deep neural networks.^{39,40}

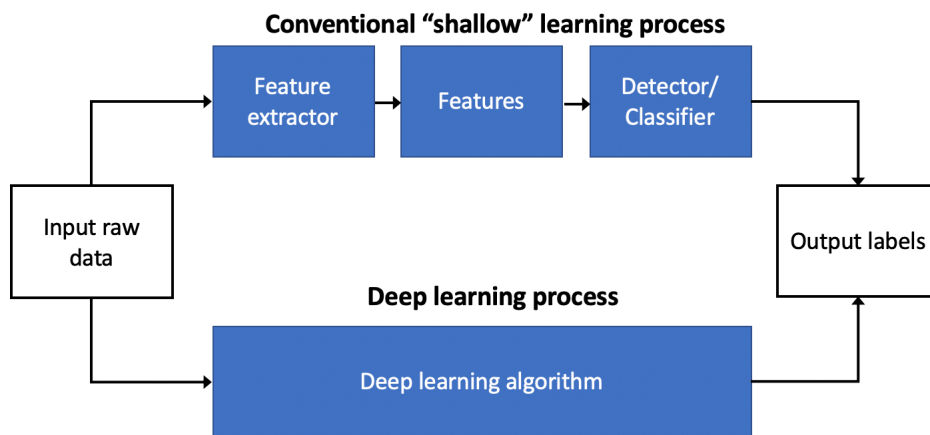


Figure 2.2: Conventional “shallow” machine learning (top) versus deep learning algorithms, where image data representation and classification are handled within the same framework.

Deep learning algorithms typically take the form of artificial neural networks (ANNs) with multiple hidden layers; however, it is important to recognize that deep learning is defined by the ability to automatically learn relevant features (data representations) from raw inputs rather than any particular structure scheme.⁴⁰

2.5.1 Fully Connected Neural Networks

A standard “fully connected” or “vanilla” neural network is one of the primary building blocks of many neural network models. A fully connected neural network consists of layers of neurons (also called “units” or “nodes”). In a given layer, each node has weighted connections to every node in the previous layer and every node in the next layer. Nodes do not share connections within the same layer. This forward-directed flow of information is why vanilla neural networks are also referred to as “feedforward neural networks.” In literature, one may also see the term “multilayer perceptron” used to refer to fully connected feedforward neural networks.

The structure of a three-layer vanilla neural network is visualized in **Figure 2.3**. The first layer in the neural network is the input layer, $\mathbf{x}^{(1)}$. The input layer takes raw data inputs and propagates them to the next layer. The final layer, $\mathbf{x}^{(3)}$, is the output layer. The results from the output layer represent the network’s output and are used to define a loss function. The width (i.e., number of nodes) for the input and output layers are typically determined by inherent characteristics of the data and the task the network performs. For example, a network designed to use grayscale images with 28×28 pixels will have 784 nodes in its input layer, where each node represents the intensity of a pixel in the image. If the task for the network is to classify handwritten digits, then the output layer will have a width of 10, with each node outputting the raw prediction score (to be later normalized into a probability) of a given digit. Hidden layers ($\mathbf{x}^{(2)}$) consist of all layers between the input and output. The width and total number of hidden layers are *hyperparameters*, meaning their values are chosen by the user and are not updated during training. The act of selecting/tuning hyperparameters for a given model is an active area of research and considered an art in and of itself.^{41,42}




Figure 2.3: A three-layer neural network (input layer, hidden layer, and output layer). The output of each node is determined by performing a non-linear transformation (known as the activation function) on the sum of the weighted inputs plus an additional bias term.

2.5.2 Training a Neural Network

The process of training a neural network involves training a set of weights and biases which act like “knobs” to control the flow of information across the network. The number of weights and biases needed to train depends on the size of the network: the example neural network shown in **Figure 2.3** has 128 weights and 12 biases, but a typical network can have hundreds of millions of weights.⁴⁰ As mentioned in Section 2.5.1, each node in a neural network receives information from each of the nodes in the previous layer (this can be visualized through the color-coded arrows in **Figure 2.3**). The output for each node is a function of the weighted sum of each incoming signal plus a bias term. We can therefore write the output for a given node, $x_j^{(l+1)}$, within layer, $l + 1$, as:

$$x_j^{(l+1)} = \sigma \left(\sum_i W_{ij} \cdot x_i^{(l)} + b_j \right) \quad 2.8$$

Where W_{ij} represents the weight associated with the connection from node $x_i^{(l)}$ to $x_j^{(l+1)}$, b_j is the bias associated with node $x_j^{(l+1)}$, and σ is the activation function. In a biological neural network, in order for a neuron to fire, it needs to receive enough electrical signal from other neurons to overcome a threshold known as the activation potential. In artificial neural networks, the *activation function* performs a similar role to the activation potential: it determines whether the node has received enough “signal” to fire. Common choices for activation functions include the hyperbolic tangent, sigmoid, softmax, and Rectified Linear Unit (ReLU) functions. Of these three, ReLU is arguably the most popular as it has been found to make networks more easily trainable.^{40,43} **Figure 2.4** displays equations and graphical representations for each of these functions. Importantly, one trait that all neural network activation functions share is that they are non-linear functions: they perform a nonlinear operation or transformation on input features to produce an output. This trait is significant because it allows neural networks to model complex, non-linear relationships between the input data and the desired output.




Figure 2.4: Common nonlinear activation functions used in neural networks.

At the start of training, the weights and biases of the network are initialized. A common method for initialization which has been found to work well is the *Xavier* initialization, in which the biases are initially set to 0 and the weights for a given node are randomly sampled from a normal distribution and bounded by:

$$\left[-\frac{6}{\sqrt{n_j + n_{j+1}}}, \frac{6}{\sqrt{n_j + n_{j+1}}} \right] \quad 2.9$$

Where n_j and n_{j+1} are the number input and output connections to the node, respectively.⁴⁴

Another popular initialization technique is *He* initialization,⁴⁵ which instead bounds uniformly sampled weights from:

$$\left[-\sqrt{\frac{2}{n_j}}, \sqrt{\frac{2}{n_j}} \right] \quad 2.10$$

After initialization, data is fed into the network, which produces an output. This output is then fed into a cost function (also called a loss function) to calculate a loss. The weights and the values of the network are then updated via *gradient descent*. The gradient of the loss function is calculated with respect to the weights and biases of every node in the network using the chain rule. Each of these parameters is then updated by adding the gradient term multiplied by a fractional learning rate typically on the order of hundredths or thousandths. This process is then repeated in an iterative fashion until a stopping criterion is met—either the loss reaches a minimum criterion, or a max number of iterations is performed. For a large network, calculating the gradient of the loss function with respect to millions of parameters is a computationally expensive endeavor. A major breakthrough in the deep learning community was the introduction of the backpropagation algorithm, which provided an efficient means of calculating the loss gradient with respect to network parameters.^{38,46}

2.5.3 Hyperparameters Associated with Training

Four significant hyperparameters associated with the training process are *batch size*, number of *epochs*, *learning rate*, and *dropout rate*. The *batch size* is the number of data samples fed to the network before the weights and biases are updated. Batch size can range from 1 (the network is updated after a single sample) to the size of the training dataset (the network is only updated after it has seen every possible data sample). Training schemes with a batch size of 1 are referred to as stochastic gradient descent, while training schemes with a batch size equal to the number of training samples are referred to as just gradient descent or batch gradient descent. If the batch size is a number between 1 and the training sample size, the algorithm is said to undergo minibatch gradient descent. Of these three training schemes, batch gradient descent is the least noisy because it uses every data sample when calculating the gradient—meaning it isn't going to be impacted by

variations within the dataset. However, for very large datasets (on the order of millions of training examples) it is computationally expensive to pass all of the data through the network before each update, which increases the overall training time. Stochastic gradient descent leads to a much noisier learning process because the network is updated after only a single data sample, which may not be representative of the dataset as a whole. A benefit of stochastic gradient descent is that in updating after only a single sample, the loss function may approach a value close to the global minimum at a faster rate. However, this also means the network must be updated more frequently, which is also computationally expensive. Minibatch gradient descent serves as a compromise between these two extremes by using a batch size large enough to be somewhat representative of the entire dataset, which minimizes training noise, but small enough that the network is able to update more frequently, leading to faster overall training times. During training, it is standard to normalize the network inputs on a batch by batch basis—a procedure called *batch normalization*—in order to improve the training speed and model stability.⁴⁷

An *epoch* refers to an instance in which the entire training dataset has been passed through the network. The *number of epochs* is thus a hyperparameter which describes how many times the entire dataset is passed through the network during training. Typically, a network requires the entire training dataset to be passed through multiple times before the loss function converges to a minimum.

The *learning rate* is a fractional coefficient applied to the loss function gradient and can be thought of as the step size used when updating the weights and biases to minimize the loss function. A popular choice for neural network optimization is to use an adaptive learning rate defined using the Adam stochastic optimization algorithm.⁴⁸

Dropout is a standard practice used during neural network training to prevent overfitting of the model. It consists of randomly selecting a set number of nodes in the network and setting their output to 0, effectively dropping them (and their connections) from the network. The fraction of neurons in a given layer which are dropped during training is called the *dropout rate*.⁴⁹

2.5.4 What Makes a Neural Network Deep?

Deep machine learning was defined in 2.5 as any machine learning method which learns multiple layers of latent features from raw data. Deep neural networks are by far the most successful deep learning architecture to date—so much so that the phrase *deep learning* is sometimes used interchangeably with *neural networks*. This subsection aims to provide some intuition as to why and under what conditions neural networks perform so well.

One important property of neural networks (already mentioned in 2.5.2) is the use of nonlinear activation functions, which allow the model to learn nonlinear relationships between data and outputs. With respect to modeling capabilities, it has been proven that neural networks are universal approximators: a feedforward neural network with as few as one hidden layer and arbitrary bounded, non-constant activation functions can be used to approximate any real, continuous function on a closed and bounded subset of \mathbb{R}^n to any accuracy.⁵⁰ This result is often referred to as the *universal approximation theorem* and tells us that a multilayer feedforward neural network can represent nearly any function given a sufficient number of hidden nodes. The result has undergone multiple iterations to clarify what conditions the activation function must meet for the theorem to be true—notably showing that the universal approximation theorem extends to the highly successful ReLU activation function.^{51,52}

A key limitation of the universal approximation theorem is that it does not tell us how large (in terms of layers or neurons) the network needs to be in order to achieve universal approximation,

nor does it guarantee that a sufficiently large network can be successfully and efficiently trained to be generalizable—i.e., a network that is large enough to represent any function may be trained to perform with high accuracy on the training dataset but still perform poorly on the validation data.⁵² While a shallow neural network with only one hidden layer is capable of being a universal approximator, such a network requires so many nodes in its hidden layer that efficient and effective training is infeasible: it is too wide.^{52,53} It has been found that neural networks with more hidden layers reduce the necessary width for each layer and ultimately the total number of nodes needed for high accuracy, providing the best tradeoff between training efficiency and model performance. Generally, a neural network is considered *deep* if it has multiple hidden layers.

2.6 Model Evaluation and Benchmarking

This section describes both performance metrics for evaluating data-driven models as well as frameworks for applying these metrics. A main goal of evaluating data-driven models is the assessment of their *generalizability*, i.e., after the model has learned a functional mapping from given inputs to outputs using training data, how well does that model then perform in the real world on unseen data? In order to assess generalizability, it is standard practice to reserve a portion of the available data from training to use for evaluation of the model.

2.6.1 Performance Evaluation: Area Under the Receiver Operating Characteristic Curve and Harrel's C-Index

Two standard measures of model performance are the area under the receiver operating characteristic curve (AUC) and Harrel's c-statistic or c-index. For binary classification models, the model output consists of a score representing the probability of being in class 1 (subtracting the score from 1 produces the probability of being in class 0). A threshold value, t , is selected such that any output whose score is greater than or equal to the t is assigned to class 1; otherwise, it is

assigned to class 0. The resulting classification predictions can be used in conjunction with the ground truth labels to calculate the model sensitivity or true positive rate (TPR), which is defined as the ratio of the number of correctly classified samples from class 1 to the total number of samples from class 1, and the false positive rate (FPR), which is defined as the ratio of samples incorrectly classified as belonging to class 0 to the total number of samples in class 0. The receiver operating characteristic (ROC) curve is defined as the curve obtained from plotting the TPR against the FPR for all possible values of the threshold, t , ranging from 0 to 1. The AUC is defined as the area under this curve and ranges in value from 1 (indicating a perfect classifier) to 0 (indicating the classifier is wrong 100% of the time). An AUC score of 0.5 indicates the model performs no better than random guessing. The AUC is a valuable performance metric because it is independent of the choice of classification threshold and can also be generalized to multi-class models.⁵⁴

The c-index provides a global assessment of performance for survival models. It is defined as the probability that for two randomly sampled patients, the patient with the higher survival score will outlive the patient with the lower survival score⁵⁴; this can be considered a generalization of the AUC.⁵⁵ Similar to the AUC, a c-index value of 1 indicates that the model correctly assigns risk to each possible pair of patients 100% of the time, while a c-index value of 0.5 indicates the model performs no better than random guessing.

2.6.2 Cross Validation

Cross validation is a specific procedure for model evaluation which helps to minimize *selection bias*—i.e., bias introduced by the choice of how the data is split between the training and validation. The premise of cross validation is to randomly split the available dataset into k equally sized subsets referred to as “folds.” The model is then trained using the first $k - 1$ folds and

evaluated on the remaining “hold-out” fold. This process is then repeated $k - 1$ times, until every combination of training and hold-out folds has been exhausted. The model’s final validation score is then calculated as the average of the k validation scores using a chosen performance measure. When k is equal to the total number of data samples, each hold-out set contains only one sample, and the procedure is referred to as leave one out cross validation (LOOCV). For smaller values of k , the procedure is called k -fold cross validation. Cross validation is particularly useful when the available dataset is small, as smaller datasets are more vulnerable to selection bias.

Another variant of cross validation utilized in this work is nested cross validation. In k -nested cross validation, the model undergoes two loops of cross validation. The outer-loop splits the original dataset into k folds, where training partitions are made up of $k - 1$ folds, with the remaining fold reserved for validation. During each iteration of the outer loop, prior to training each training partition is further split into n folds, which then undergo n -fold cross validation. This n -fold cross validation is referred to as the inner loop. Nested cross validation is a valuable tool for hyperparameter tuning and feature selection in a cross validation paradigm because it allows for model selection to occur independently for each outer validation, reducing selection bias.⁵⁶

2.7 Quantum Algorithms and Quantum Computing

2.7.1 Overview of Quantum Computing

Quantum computing refers to the use of quantum mechanical systems to perform computational tasks. The original inspiration for the concept of quantum computing is credited to physicist and Nobel prize laureate Richard Feynman, who postulated in the 1970s that a quantum computer could solve problems that classical computers cannot. This idea followed a proposal by Paul Benioff of a quantum mechanical model of the Turing machine.⁵⁷

In classical computing, the basic unit of information is called a *bit*, which can take the value of either 0 or 1. In quantum computing, the quantum counterpart of a bit is called a *qubit*; unlike classical bits, qubits can exist simultaneously in both states at once (also called a *mixed state*), which can be represented by a linear superposition:

$$|\Psi\rangle = c_1|0\rangle + c_2|1\rangle \quad 2.11$$

Where the probability amplitudes c_1 and c_2 are complex numbers. Only upon an event (i.e., a quantum measurement) will the qubit $|\Psi\rangle$ collapse to a binary quantum state (either $|0\rangle$ or $|1\rangle$ with probability $|c_1|^2$ or $|c_2|^2$, respectively). By extension, a quantum system containing n qubits (called an n -qubit register) can represent 2^n states simultaneously. For instance, a 2-qubit system defined as a uniform superposition state, $|\Psi\rangle = \frac{1}{2}(|00\rangle + |01\rangle) + \frac{1}{2}(|10\rangle + |11\rangle)$ represents 4 states simultaneously. This property of being able to exist simultaneously in multiple states at once is known as *quantum superposition*, which further allows for *quantum parallelism*—providing quantum computing with a much faster computational speed than its classical counterpart.

Quantum parallelism is the ability to simultaneously operate on all the superimposed states in parallel. Since the number of states represented by a qubit register grows exponentially with the number of registers, an operation performed on a quantum computer would have taken an exponential number of operations on a classical computer with the same numbers of registers. One tradeoff of quantum parallelism is that the probability of measuring one particular state out of 2^n states also decreases exponentially.

Another quantum mechanical phenomenon relevant to quantum computing is *quantum entanglement*. Quantum entanglement refers to a phenomenon in which 2 or more qubits in a multi-qubit system are correlated with each other. For example, in a maximally entangled 2-qubit state (also called a Bell state) given by: $|\Psi\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$, $|\Psi\rangle$ can be found upon measurement

to be in either state $|00\rangle$ or state $|11\rangle$, each with a probability of 50%. What this means is that if the first qubit is measured to be in the $|1\rangle$ state, the second qubit will also collapse to the $|1\rangle$ state; the outcome of the second qubit is intrinsically tied to the first, and vice-versa. The entanglement phenomenon provides quantum computing with a resource absent in classical computing that makes classically impossible processes, like super-dense coding and quantum teleportation, possible. An actively sought out real-life application of quantum communication is to create an eavesdrop-proof channel for key distribution in cryptography.⁵⁸ For practical purposes, this translates to an efficient representation of highly correlated information, which is still a challenge in classical computation.

A quantum algorithm refers to a set of instructions for performing a calculation via operations on qubits. While it is possible for quantum algorithms to be simulated on a classical computer, often any theoretically hoped-for computational advantage can only be realized using a *quantum computer*—a hardware device which allows calculations to be performed on physical (rather than simulated) quantum states. Along with quantum information theory, quantum hardware is an area of active research with growing interest from both the private sector and the government.⁵⁹⁻⁶¹ A quantum computer is built from quantum circuits and quantum gates, just like classical computers are made up of electrical circuits and logic gates. The physical realization of quantum circuits has been achieved through different means: trapped ions,⁶² nuclear magnetic resonance (NMR), linear optical systems, and superconducting solid-state system.⁶³ Currently, superconducting solid-state systems are the main focus in commercial sectors and companies— with Google, IBM, Microsoft, and Intel racing to build a universal quantum computer—whereas nuclear magnetic resonance (NMR), optical, and trap ion computing systems are most actively researched in academic and governmental institutes.

2.7.2 Postulates of Quantum Mechanics

Quantum mechanics is a mathematical framework which was invented as a means of describing and understanding physical systems in nature which were observed to disobey the laws of classical physics. Quantum mechanics can be summarized as a set of postulates (or logical assumptions), which are presented briefly in this section.

1. The state of a quantum mechanical system is completely represented by a normalized ket, $|\psi\rangle$.

Mathematically, kets are vectors that reside in an inner-product vector space over complex number field called Hilbert Space. While Hilbert Space can represent infinite dimensional space, the dimensionality of the Hilbert space is determined by the physics of a problem. Vector spaces satisfy the closure property, i.e., the linear combination of any two vectors of a vector space must lie in the same vector space. This implies that the superposition of two states is also a state of the system: if $|\psi_0\rangle$ and $|\psi_1\rangle$ are two possible states of a system, then so is $|\psi\rangle = c_0|\psi_0\rangle + c_1|\psi_1\rangle$, where c_0 and c_1 are complex numbers. In Dirac notation, a 'bra' ($\langle |$) is the dual of a vector 'ket' ($| \rangle$) and together in the bra-ket succession define an inner product. Mathematically, 'bra' denotes a linear functional $f: V \rightarrow \mathbb{C}$ that maps vector 'ket' to a number in the complex plane. The Dirac notation simplifies the inner product notation especially for states of continuous variables, which are represented in terms of wave functions, i.e.,

$$\langle \Psi | \phi \rangle = \int dx \Psi^*(x) \phi(x) \quad 2.12$$

2. Every physical observable attribute of a quantum mechanical system is described by Hermitian operator, \hat{O} , that acts on kets describing that system.

Observables are physically measurable attributes such as position, linear and angular momentum, energy, and so on. An operator is Hermitian if it is equal to its conjugate transpose, i.e., $O = O^\dagger$.

This property guarantees real eigenvalues necessary for the observables to be physical. An operator \hat{O} acting on a ket $|\psi\rangle$ is denoted by left multiplication i.e., $|\psi'\rangle = \hat{O}|\psi\rangle$. In general, the operation changes the state of the quantum system.

3. The only possible result of a measurement of an observable O is one of the eigenvalues of the corresponding operator \hat{O} .

This postulate describes the source of the word "quantum." If the observable is of continuous spectrum, like position or momentum, then the measurement will give us classical results. However, if the observable has a discrete spectrum, like the angular momentum of an orbiting electron, the measurement will yield discrete set of values in multiples of Planck's constant.

4. Upon a measurement of the observable O on a quantum mechanical system in the state $|\psi\rangle$, the probability of obtaining the eigenvalue O is given by the square of the inner product of $|\psi\rangle$ with the eigenstate $|o_n\rangle$, i.e., $|\langle o_n|\psi\rangle|^2$.

Besides having real eigenvalues, the eigenstates of a Hermitian operator are orthogonal, i.e., $\langle o_i|o_j\rangle = \delta_{ij}$ where the Kronecker delta $\delta_{ij} = 1$ for $i = j$ and $\delta_{ij} = 0$ otherwise. The orthogonal eigenstate spans the space of states and forms a basis. This means that the state of a quantum mechanical system can be expanded as a linear combination of the eigenstates of a Hermitian operator with complex coefficients.

$$|\psi\rangle = \sum_n c_n |o_n\rangle \tag{2.13}$$

The complex coefficients, $\langle o_n|\psi\rangle = c_n$, represent the *probability amplitude* associated with the eigenstate $|o_n\rangle$. This means that $|\psi\rangle = \sum_n \langle o_n|\psi\rangle |o_n\rangle$ and hence $\langle\psi|\psi\rangle = \sum_n |\langle o_n|\psi\rangle|^2$. Since $\langle\psi|\psi\rangle$ is the total probability and sums to 1, $|\langle o_n|\psi\rangle|^2$ is the individual probability associated with measuring eigenstate o_n . The probabilistic interpretation requires the normalization of kets, i.e.,

$\langle\psi|\psi\rangle = 1$, which can be achieved by dividing $|\psi'\rangle$ by its norm, $\sqrt{\langle\psi'|\psi'\rangle}$, where prime notation represents the state before normalization.

5. Immediately after the measurement of the observable O that yielded the value o_n , the state of the system collapses to the normalized eigenstate $|o_n\rangle$.

This postulate is "counterintuitive" with respect to quantum mechanics. Measurements of observable O after preparing several identical quantum mechanical systems in state $|\psi\rangle$ can yield different results. This is to be expected, as quantum mechanical states are probabilistic in nature and can exist in a superimposed state. However, carrying out a second measurement immediately on a system that yielded o_n value in the first measurement will always yield o_n .

6. The time evolution of a quantum mechanical system preserves the normalization of the associated ket. The time evolution of the state is a unitary transformation described by
$$|\psi(t)\rangle = \hat{U}(t, t_0)|\psi(t_0)\rangle.$$

The preservation of normalization during time evolution of a quantum mechanical system implies conservation of probability: the probability of finding a system in an eigenstate summed over all possible eigenstates must be 1. This is ensured by setting the time evolution operator as a unitary transformation. Mathematically, unitary transformations preserve inner products: the inner product of kets are equal to the inner products of the transformed kets, i.e., $\langle\psi|\psi\rangle = \langle\psi|U^\dagger U|\psi\rangle$ which is true when $U^\dagger U = 1$, implying $U^\dagger = U^{-1}$. By symmetry of the inner product operations, $U^\dagger U = I = U U^\dagger$. Practically, unitary transformations are reversible operations such as rotation and reflection, i.e., an inverse transformation of equal magnitude will nullify the forward transformation and send a state back to its original state.

2.7.3 Quantum Cognition and Decision Theory

Although quantum mechanics was originally invented as a predictive framework to describe the behavior of physical systems in nature, in recent decades it has also been adapted by the field of psychology as a framework for modeling human behavior and decision making.⁶⁴ The motivation for applying a quantum framework to model human cognition/decision making comes from the observation of scenarios in which human behavior violates classically based Kolmogorov probability theory. One important example is the violation of the “sure-thing” principle, which states that if a person prefers action A over action B under the state of world X, and also prefers action A over action B under the complimentary state of world X’, then the person should prefer action A over action B even if they do not know the state of the world.⁶⁵ Researchers have experimentally shown that humans often violate the “sure-thing” principal when asked to make decisions under conditions of uncertainty,^{66,67} leading to a violation of the law of total probability, which states:

$$P(A) = P(A|X) * P(X) + P(A|X') * P(X') \quad 2.14$$

The benefit of adopting a quantum framework for probabilistic modeling of human behavior is that because quantum theory is not bound by the law of total probability, it can more accurately predict human behavior under certain conditions—namely scenarios in which certain aspects of a system are uncertain to the decision-maker or subject to order-effects.⁶⁸ These studies help to provide insight into how quantum mechanics may be utilized for modeling other abstract systems under conditions of uncertainty.

2.7.4 Quantum Annealers

Quantum annealers refer to a class of quantum hardware which solve optimization problems through the process of quantum annealing (QA). First proposed by Kadowaki and Nishimori,⁶⁹

QA exploits the result of the quantum adiabatic theorem which states that when a quantum system undergoes a gradual change in its total energy from H_i to H_f , if it starts out in the n^{th} eigenstate of H_i then it will end up in the corresponding n^{th} eigenstate of H_f . This means that a quantum system can be constructed to start in the ground state of a known (or solvable) objective function and then gradually shifted into the ground state of the objective function of interest using the annealing coefficient, Γ :

$$H(t) = H_f + \Gamma(t)H_i \quad 2.15$$

At the start of the optimization process, the annealing variable Γ is very large such that $H(0) \cong \Gamma(0)H_i$. By the end of the optimization process, $\Gamma(t)$ approaches 0 and the system is now in the lowest energy eigenstate of the Hamiltonian defined by the objective function of interest, $H(t_{end}) \cong H_f$. Thus, similar to simulated annealing, QA is guaranteed to find the global optimal solution if allowed to run long enough, and theoretical and experimental results suggest that QA boasts performance benefits over its classical counterpart, simulated annealing.^{70,71}

Initial studies on QA were performed as simulations on classical computers⁷²; however, classical simulation of this quantum process is computationally expensive and therefore impractical for higher order objective functions—such as those which would be necessary for many real-world problems. Quantum annealers are devices which use quantum hardware to physically realize the QA process. The major developer of quantum annealers to date is DWave INC. DWave quantum annealers use qubits which are made up of superconducting loops of current—the direction of the resulting magnetic field (down vs. up) defines the classical binary states of the qubit (0 vs. 1), but as quantum objects they can also exist in both states simultaneously. As of this time, DWave hardware supports a maximum of 2,000 qubits, which can represent 2^{2000} unique states.

The earliest known application of quantum computing to medical physics was a study published in 2015 which investigated QA as a method for IMRT treatment plan optimization.⁷³ In this study, IMRT beamlet weights for two prostate cancer cases were optimized using quantum annealer hardware and compared against two optimization methods—Tabu search and simulated annealing implemented on a standard classical computer.

Each algorithm was run for the same number of objective function evaluations, and performance was assessed by computational speed (defined as wall clock time), by the final objective function value, and by the overall quality of the treatment plans generated. For both cases, QA had a wall clock time that was more than twice as fast as simulated annealing and more than three times as fast as Tabu search. Simulated annealing was found to produce the highest quality plans for both cases, while the QA technique came in second and third, respectively.

One unique aspect of performing calculations on quantum-hardware is that problems must be formulated such that inputs and measured outputs are binary. In the case of the above study, each beamlet weight value was represented as a 5-bit vector, $w_b \in \{[00000], [10000], [01000], \dots, [11111]\}$, allowing for $2^5 = 32$ levels of discretization. In total, each beamlet required 7 qubits for optimization: 5 to represent the numeric value and 2 for functional smoothing. It is important to note that when performing optimization with quantum hardware, the size of the solution search space is limited by the total number of qubits that the hardware can support. The QA device used in the above study supported 512 qubits. The authors therefore defined the beamlet dimensions such that ~ 70 beamlet weights were optimized for each plan. For clinical applications, however, a typical IMRT plan can have on the order of thousands to tens of thousands of beamlets—due to the use of more treatment beams (the study above used only 5) as well as smaller beamlet sizes. In addition, beamlet weights are usually represented as

(nearly) continuous variables, meaning each beamlet weight has a much larger range of potential values it can take. The most advanced quantum annealers as of the year 2020 support up to 2,000 qubits; thus, technology is not yet at the point where it can support the full complexity of these types of optimization problems in radiation oncology. However, if quantum annealers are able to double the number of qubits they support every 2-4 years (as has historically been the case), then this may cease to be an issue within the next decade.

2.8 References

1. Chang DS, Lasley FD, Das IJ, Mendonca MS, Dynlacht JR. Therapeutic Ratio. In: Chang DS, Lasley FD, Das IJ, Mendonca MS, Dynlacht JR, eds. *Basic Radiotherapy Physics and Biology*. Cham: Springer International Publishing; 2014:277-282.
2. Bortfeld T. IMRT: a review and preview. *Phys Med Biol*. 2006;51(13):R363-379.
3. Shumway DA, Griffith KA, Pierce LJ, et al. Wide Variation in the Diffusion of a New Technology: Practice-Based Trends in Intensity-Modulated Radiation Therapy (IMRT) Use in the State of Michigan, With Implications for IMRT Use Nationally. *Journal of Oncology Practice*. 2015;11(3):e373-e379.
4. Verellen D, Ridder MD, Storme G. A (short) history of image-guided radiotherapy. *Radiotherapy and Oncology*. 2008;86(1):4-13.
5. DeLaney TF. Proton Therapy in the Clinic. *Frontiers of Radiation Therapy and Oncology*. 2011;43:465-485.
6. Otto K. Volumetric modulated arc therapy: IMRT in a single gantry arc. *Medical physics*. 2008;35(1):310-317.
7. Teoh M, Clark CH, Wood K, Whitaker S, Nisbet A. Volumetric modulated arc therapy: a review of current literature and clinical use in practice. *The British journal of radiology*. 2011;84(1007):967-996.
8. Burnet NG, Thomas SJ, Burton KE, Jefferies SJ. Defining the tumour and target volumes for radiotherapy. *Cancer Imaging*. 2004;4(2):153-161.
9. Yan D, Vicini F, Wong J, Martinez A. Adaptive radiation therapy. *Physics in Medicine and Biology*. 1997;42(1):123-132.
10. Lim-Reinders S, Keller BM, Al-Ward S, Sahgal A, Kim A. Online Adaptive Radiation Therapy. *Int J Radiat Oncol Biol Phys*. 2017;99(4):994-1003.
11. Hunt A, Hansen VN, Oelfke U, Nill S, Hafeez S. Adaptive Radiotherapy Enabled by MRI Guidance. *Clinical Oncology*. 2018;30(11):711-719.
12. Brouwer CL, Steenbakkers RJ, Langendijk JA, Sijtsema NM. Identifying patients who may benefit from adaptive radiotherapy: Does the literature on anatomic and dosimetric changes in head and neck organs at risk during radiotherapy provide information to help? *Radiother Oncol*. 2015;115(3):285-294.
13. Bosch MV, Oellers M, Reymen B, Elmpt WV. Automatic selection of lung cancer patients for adaptive radiotherapy using cone-beam CT imaging. *Physics and Imaging in Radiation Oncology*. 2017;1:21-27.

14. Oyama A, Kumagai S, Arai N, et al. Image quality improvement in cone-beam CT using the super-resolution technique. *J Radiat Res.* 2018;59(4):501-510.
15. Yang C-C, Chen F-L, Lo Y-C. Improving Image Quality of On-Board Cone-Beam CT in Radiation Therapy Using Image Information Provided by Planning Multi-Detector CT: A Phantom Study. *PLoS One.* 2016;11(6):e0157072-e0157072.
16. Tien H-J, Yang H-C, Shueng P-W, Chen J-C. Cone-beam CT image quality improvement using Cycle-Deblur consistent adversarial networks (Cycle-Deblur GAN) for chest CT imaging in breast cancer patients. *Scientific Reports.* 2021;11(1):1133.
17. Zhang T, Chi Y, Meldolesi E, Yan D. Automatic delineation of on-line head-and-neck computed tomography images: toward on-line adaptive radiotherapy. *Int J Radiat Oncol Biol Phys.* 2007;68(2):522-530.
18. Castadot P, Lee JA, Parraga A, Geets X, Macq B, Grégoire V. Comparison of 12 deformable registration strategies in adaptive radiation therapy for the treatment of head and neck tumors. *Radiother Oncol.* 2008;89(1):1-12.
19. Wang H, Dong L, O'Daniel J, et al. Validation of an accelerated 'demons' algorithm for deformable image registration in radiation therapy. *Phys Med Biol.* 2005;50(12):2887-2905.
20. Poon M, Holborn C, Cheng KF, Fung WWK, Chiu G. Evaluation of deformed image-based dose calculations for adaptive radiotherapy of nasopharyngeal carcinoma. *Med Dosim.* 2017;42(4):273-281.
21. Xing Y, Zhang Y, Nguyen D, Lin M-H, Lu W, Jiang S. Boosting radiotherapy dose calculation accuracy with deep learning. *Journal of Applied Clinical Medical Physics.* 2020;21(8):149-159.
22. Li N, Zarepisheh M, Uribe-Sanchez A, et al. Automatic treatment plan re-optimization for adaptive radiotherapy guided with the initial plan DVHs. *Phys Med Biol.* 2013;58(24):8725-8738.
23. Tseng H-H, Luo Y, Ten Haken RK, El Naqa I. The Role of Machine Learning in Knowledge-Based Response-Adapted Radiotherapy. *Frontiers in Oncology.* 2018;8:266.
24. Tseng H-H, Luo Y, Cui S, Chien J-T, Ten Haken RK, Naqa IE. Deep reinforcement learning for automated radiation adaptation in lung cancer. *Medical physics.* 2017;44(12):6690-6705.
25. Wilkie JR, Mierzwa ML, Casper KA, et al. Predicting late radiation-induced xerostomia with parotid gland PET biomarkers and dose metrics. *Radiother Oncol.* 2020;148:30-37.
26. Rosen BS, Hawkins PG, Polan DF, et al. Early Changes in Serial CBCT-Measured Parotid Gland Biomarkers Predict Chronic Xerostomia After Head and Neck Radiation Therapy. *International journal of radiation oncology, biology, physics.* 2018;102(4):1319-1329.
27. Cui S, Luo Y, Tseng H-H, Ten Haken RK, El Naqa I. Combining handcrafted features with latent variables in machine learning for prediction of radiation-induced lung damage. *Medical physics.* 2019;46(5):2497-2511.
28. Luo Y, El Naqa I, McShan DL, et al. Unraveling biophysical interactions of radiation pneumonitis in non-small-cell lung cancer via Bayesian network analysis. *Radiotherapy and Oncology.* 2017;123(1):85-92.
29. Luo Y, McShan D, Ray D, et al. Development of a Fully Cross-Validated Bayesian Network Approach for Local Control Prediction in Lung Cancer. *IEEE Trans Radiat Plasma Med Sci.* 2019;3(2):232-241.

30. Bishop CM. *Pattern recognition and machine learning*. New York : Springer, [2006] ©2006; 2006.
31. Wei L, El Naqa I. Feature extraction and qualification. *Radiomics and Radiogenomics: Technical Basis and Clinical Applications*. 2019:121.
32. Vallières M, Freeman C, Skamene S, El Naqa I. Vallières2015 PMB suppData. In:2015.
33. Vallières M, Freeman CR, Skamene SR, El Naqa I. A radiomics model from joint FDG-PET and MRI texture features for the prediction of lung metastases in soft-tissue sarcomas of the extremities. *Physics in Medicine & Biology*. 2015;60(14):5471.
34. Joan LB. Techniques for avoiding local minima in gradient-descent-based ID algorithms. Paper presented at: Proc.SPIE1997.
35. Kirkpatrick S, Gelatt CD, Vecchi MP. Optimization by Simulated Annealing. *Science*. 1983;220(4598):671.
36. Peng C-YJ, Lee KL, Ingersoll GM. An Introduction to Logistic Regression Analysis and Reporting. *The Journal of Educational Research*. 2002;96(1):3-14.
37. Cox DR. Regression models and life-tables. *Journal of the Royal Statistical Society: Series B (Methodological)*. 1972;34(2):187-202.
38. Rumelhart DE, Hinton GE, Williams RJ. Learning representations by back-propagating errors. *Nature*. 1986;323(6088):533-536.
39. Hinton GE, Salakhutdinov RR. Reducing the Dimensionality of Data with Neural Networks. *Science*. 2006;313(5786):504-507.
40. LeCun Y, Bengio Y, Hinton G. Deep learning. *nature*. 2015;521(7553):436-444.
41. Hutter F, Lücke J, Schmidt-Thieme L. Beyond Manual Tuning of Hyperparameters. *KI - Künstliche Intelligenz*. 2015;29(4):329-337.
42. Luo G. A review of automatic selection methods for machine learning algorithms and hyper-parameter values. *Network Modeling Analysis in Health Informatics and Bioinformatics*. 2016;5(1):18.
43. Xavier G, Antoine B, Yoshua B. Deep Sparse Rectifier Neural Networks. 2011/06/14, 2011.
44. Glorot X, Bengio Y. Understanding the difficulty of training deep feedforward neural networks. 2010.
45. Hochreiter S. The vanishing gradient problem during learning recurrent neural nets and problem solutions. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*. 1998;6(02):107-116.
46. Nielsen MA. *Neural networks and deep learning*. Vol 2018: Determination press San Francisco, CA, USA.; 2015.
47. Ioffe S, Szegedy C. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:150203167*. 2015.
48. Kingma DP, Ba J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:14126980*. 2014.
49. Srivastava N, Hinton G, Krizhevsky A, Sutskever I, Salakhutdinov R. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*. 2014;15(1):1929-1958.
50. Hornik K. Approximation capabilities of multilayer feedforward networks. *Neural Networks*. 1991;4(2):251-257.

51. Leshno M, Lin VY, Pinkus A, Schocken S. Multilayer feedforward networks with a nonpolynomial activation function can approximate any function. *Neural networks*. 1993;6(6):861-867.
52. Goodfellow I, Bengio Y, Courville A. *Deep Learning*. MIT Press; 2016.
53. Bengio Y. *Learning deep architectures for AI*. Now Publishers Inc; 2009.
54. Hand DJ, Till RJ. A Simple Generalisation of the Area Under the ROC Curve for Multiple Class Classification Problems. *Machine Learning*. 2001;45(2):171-186.
55. Hanley JA, McNeil BJ. The meaning and use of the area under a receiver operating characteristic (ROC) curve. *Radiology*. 1982;143(1):29-36.
56. Cawley GC, Talbot NLC. On over-fitting in model selection and subsequent selection bias in performance evaluation. *The Journal of Machine Learning Research*. 2010;11:2079-2107.
57. Benioff P. The computer as a physical system: A microscopic quantum mechanical Hamiltonian model of computers as represented by Turing machines. *Journal of Statistical Physics*. 1980;22(5):563-591.
58. Yin J, Li Y-H, Liao S-K, et al. Entanglement-based secure quantum cryptography over 1,120 kilometres. *Nature*. 2020;582(7813):501-505.
59. Arute F, Arya K, Babbush R, et al. Quantum supremacy using a programmable superconducting processor. *Nature*. 2019;574(7779):505-510.
60. U.S. Subcommittee On Quantum Information Science COS, National Science & Technology Council. National Strategic Overview for Quantum Information Science. In:2018.
61. Policy USOOsaT. Artificial Intelligence and Quantum Information Science R&D Summary: Fiscal Years 2020-2021. In:2020.
62. Steane A. The ion trap quantum information processor. *Applied Physics B*. 1997;64(6):623-643.
63. Huang H-L, Wu D, Fan D, Zhu X. Superconducting quantum computing: a review. *Science China Information Sciences*. 2020;63(8):180501.
64. Bruza PD, Wang Z, Busemeyer JR. Quantum cognition: a new theoretical approach to psychology. *Trends Cogn Sci*. 2015;19(7):383-393.
65. Savage LJ. *The foundations of statistics*. Courier Corporation; 1972.
66. Busemeyer JR, Wang Z, Lambert-Mogiliansky A. Empirical comparison of Markov and quantum models of decision making. *Journal of Mathematical Psychology*. 2009;53(5):423-433.
67. Shafir E, Tversky A. Thinking through uncertainty: Nonconsequential reasoning and choice. *Cognitive Psychology*. 1992;24(4):449-474.
68. Busemeyer JR, Wang Z, Shiffrin RM. Bayesian model comparison favors quantum over standard decision theory account of dynamic inconsistency. *Decision*. 2015;2(1):1-12.
69. Kadowaki T, Nishimori H. Quantum annealing in the transverse Ising model. *Physical Review E*. 1998;58(5):5355.
70. Morita S, Nishimori H. Mathematical foundation of quantum annealing. *Journal of Mathematical Physics*. 2008;49(12):125210.
71. Mukherjee S, Chakrabarti BK. Multivariable optimization: Quantum annealing and computation. *The European Physical Journal Special Topics*. 2015;224(1):17-24.

72. Farhi E, Goldstone J, Gutmann S, Lapan J, Lundgren A, Preda D. A quantum adiabatic evolution algorithm applied to random instances of an NP-complete problem. *Science*. 2001;292(5516):472-475.
73. Nazareth DP, Spaans JD. First application of quantum annealing to IMRT beamlet intensity optimization. *Phys Med Biol*. 2015;60(10):4137-4148.

Chapter 3 A Quantum-Inspired Approach to Radiotherapy Optimization

This chapter discusses the design and evaluation of a new quantum inspired optimization algorithm for radiotherapy treatment planning optimization and is based on the paper: Pakela J.M., Tseng H.H., Matuszak M.M., Ten Haken R.K., McShan D.L., El Naqa I. (2020), Quantum-inspired algorithm for radiotherapy planning optimization. *Med. Phys.*, 47(1): 5-18.

3.1 Introduction

Radiation therapy has been established as one of the primary modalities for cancer treatment, used either exclusively or in combination with other techniques such as chemotherapy or surgery.^{1,2} A critical challenge for radiation therapy (and all cancer therapies) is to deliver an adequate dose to the tumor to ensure curative or palliative results while minimizing the dose delivered to normal tissues. Intensity modulated radiation therapy (IMRT) is a type of external beam radiation therapy in which each beam is subdivided into a grid of beamlets whose intensities are determined by dynamic shielding via a multi-leaf collimator (MLC). Because IMRT and other radiation therapy techniques which rely on dynamic intensity modulation (such as Volumetric Arc Therapy (VMAT)) are capable of creating concave-shaped dose distributions, they are particularly effective for challenging cases in which the tumor volume is irregular and near critical organs at risk (OARs).^{3,4} The intensity modulations determined from this dynamic shielding optimization are characterized by aperture or beamlet weights. The challenge of calculating optimal weights for a treatment plan often represents a non-convex,⁵ large-scale optimization problem that must be

solved within a clinically reasonable timeframe. The ability to quickly perform robust optimizations is particularly significant in online adaptive radiotherapy, in which a patient's plan may be reoptimized several times during the treatment course to account for changes such as tumor shrinkage or organ deformation.⁶

Quantum computing research is believed to hold promise for achieving computational speedup for certain types of problems.⁷ In quantum computing, classical bits (whose two states are often represented by 0 and 1) are replaced by quantum bits (qubits) which may exist in any linear superposition of 0 and 1.⁸ This allows quantum computers to explore multiple solutions simultaneously, and quantum algorithms can take advantage of this to achieve a significant computational speedup.^{8,9} However, the direct use of quantum computers is still limited by challenges related to creating a proper hardware environment where qubits are maintained in quantum coherence⁷ and the number of qubits deployed is still limited (11-2,000)¹⁰⁻¹⁵ to effectively handle large scale optimization problems like planning optimization. On the other hand, quantum-inspired algorithms also hold promise for achieving computational speedup of complex optimization problems. Such algorithms are not necessarily quantum processes per se (though some can be formulated to run on a quantum computer); rather, they are quantum simulations designed to run on a classical computer.

The idea of incorporating quantum-inspired techniques into stochastic algorithms was first proposed by de Falco et al. in 1989.¹⁶ A few years later, Kadowaki and Nishimori demonstrated the use of quantum annealing (QA) on an Ising model of atomic spins by applying a transverse field, which was annealed to 0° and numerically solving the time-dependent Schrödinger equation for small systems; they found that the probability of reaching the ground state was consistently

higher for QA than simulated annealing (SA). Many studies have since ensued that have demonstrated QA's potential for a variety of problems.¹⁶⁻²⁰

While QA holds theoretical promise for certain problem classes with limited dimensionality,²¹ its implementation on a classical computer is impractical for IMRT optimization¹⁹ and deployment on a quantum computer is currently hindered by the limited number of qubits built into existing hardware systems.^{19,21} To avoid these computational limitations, we have implemented another quantum-inspired optimization scheme that models the exploration of higher energy solutions based on the probability of a particle tunneling through a one-dimensional potential energy barrier. We refer to our algorithm as Quantum Tunnel Annealing (QTA) to distinguish it from the QA algorithms described by de Falco and others.¹⁶⁻²² In this paper, we present a proof-of-concept study that (1) demonstrates the behavior of QTA when applied to beamlet intensity and direct aperture optimization for IMRT treatment planning, and (2) compares QTA performance with that of SA as a representative benchmark of traditional optimization methods.

3.2 Methods

3.2.1 Quantum Tunnel Annealing

QTA works by modeling an optimization problem as a biased random walk over a fixed number of iterations. During each iteration, a new potential solution (e.g., beamlet-weight vector) is selected from within the neighborhood of the current solution. The energy associated with the new potential solution, given by the objective function, is then calculated and compared against that of the current solution. Potential solutions with lower energies are immediately accepted and set as the current solution. A significant challenge associated with non-convex optimization problems is that the algorithm can become stuck in a local minimum before it has a chance to reach the globally

optimal solution. To avoid this pitfall and ensure adequate exploration of the solution space, QTA simulates quantum fluctuations, allowing the algorithm to accept a worse solution with some probability P . In this process, consider a quantum-particle with energy E , traversing through a one-dimensional potential energy landscape, $V(x)$. The particle's wavefunction, $\Psi(x)$, obeys the time-independent Schrödinger equation:

$$H \Psi(x) = E \Psi(x), \quad 3.1$$

Where the Hamiltonian operator, H , is a function of the particle's potential V and kinetic energy T :

$$H = T + V. \quad 3.2$$

Figure 3.1 illustrates such a particle encountering a potential energy barrier (denoted V^B). The particle's wavefunction prior, during, and after encountering the barrier can be expressed as:

$$\Psi(x) = \begin{cases} Ae^{ikx} + A'e^{-ikx}, & \text{in region A} \\ Be^{\kappa x} + B'e^{-\kappa x}, & \text{in region B} \\ Ce^{ik'x}, & \text{in region C} \end{cases}, \quad V(x) = \begin{cases} V^A, & \text{in region A} \\ V^B, & \text{in region B} \\ V^C, & \text{in region C} \end{cases} \quad 3.3$$

With wave-numbers:

$$k = \sqrt{\frac{2m}{\hbar^2}(\Gamma - V^A)}, \quad \kappa = \sqrt{\frac{2m}{\hbar^2}(V^B - \Gamma)}, \quad k' = \sqrt{\frac{2m}{\hbar^2}(\Gamma - V^C)}. \quad 23 \quad 3.4$$

A positive exponent represents the particle traveling to the right, and a negative exponent represents the particle traveling to the left. Thus, $A(A')$ and $B(B')$ represent the amplitudes for the incident (reflected) waves in regions A and B, respectively, and C is the amplitude of the wave transmitted through the barrier. The probability of tunneling through the barrier is given by the transmission coefficient $T = \frac{k'}{k} \left| \frac{C}{A} \right|^2$. This value has suggested by Mukherjee and Chakrabarti to be

on the order of $e^{-\frac{w\sqrt{V^B-V^A}}{\Gamma}}$ using a Wentzel-Kramers-Brillouin (WKB) approximation.^{18,23}




Figure 3.1: Figurative illustration of a particle (represented by its wave-function, Ψ) tunneling through a potential energy barrier (in region B) in a 1-dimensional energy landscape.

Hence, the probability of QTA accepting a worse solution can be redefined to be proportional to $\exp\left(\frac{-w*\sqrt{V_{new}-V_{old}}}{\Gamma}\right)$, where Γ is the kinetic energy of the system (an annealing variable synonymous with the temperature, T , in SA), V_{new} is the potential energy of the system at new solution defined by the objective function, and w is the width of the barrier being tunneled through. This barrier width is a dynamic parameter, which serves as an additional degree of freedom that is not present in the SA formalism.

3.2.2 Calculation of Barrier Width

As stated in Section 3.2.1, the barrier width represents an additional degree of freedom, which QTA can use to obtain an optimal solution in a shorter timeframe. The expected trend in the barrier width's evolution over the course of the optimization can be derived from the following argument: At the start of the optimization, energy barriers that the system encounters have finite widths; as the system approaches its global minimum, the widths of any barriers encountered would grow increasingly large.

In the interest of modeling the barrier width after a physical system in nature, one of the common barrier width schedules tested was modeled after the growth rate of Gallium Arsenide (GaAs) during the process of metal organic chemical vapor deposition (MOCVD). A typical MOCVD setup consists of a reaction chamber and a substrate material on a heated platform. As the substrate is heated by the platform, chemical reactions take place in the gas of the reaction chamber, leading to the growth of thin films upon the surface of the substrate. In a horizontal-type reaction chamber, the reactants are passed through the chamber horizontally. One of the most common semiconductors grown using MOCVD is GaAs.²⁴ The growth of semiconductors using MOCVD is a complex process influenced by many parameters. It was shown experimentally that GaAs's growth rate is proportional to the square root of the gas velocity.²⁵ Given that kinetic energy is also proportional to the square root of velocity (i.e. $\Gamma = \frac{1}{2}mv^2$), we can express the growth rate as:

$$\frac{dw}{dt}(t) \propto \sqrt[4]{\Gamma(t)} \quad 3.5$$

Where t represents the annealing time defined as the iteration number, and Γ is the kinetic energy of the annealing system, defined in this study as:

$$\Gamma(t) = 10 \times \left(1 - \frac{\log(t)}{\log(N)}\right) \quad 3.6$$

With N defined as the total number of iterations performed during the optimization.

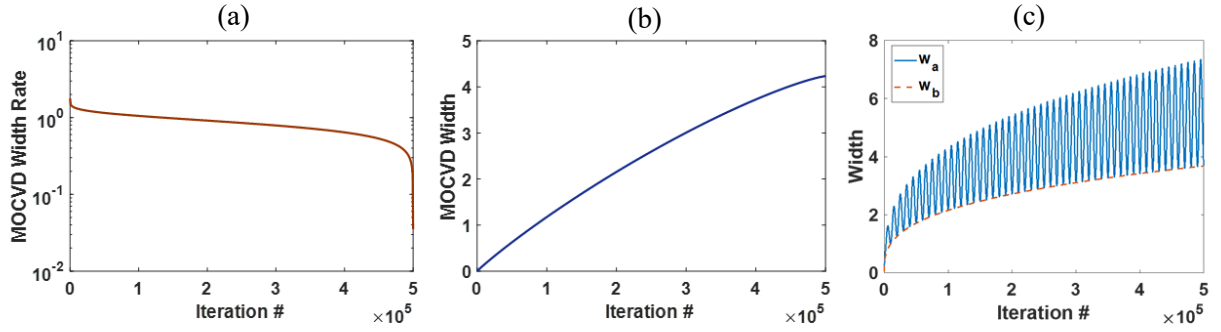


Figure 3.2: (a) Barrier width rate extrapolated from metal organic chemical vapor deposition (MOCVD) studies by Leys and Veenvliet.²⁵ (b) Barrier width function calculated via numerical integration of (a). (c) Additional width functions explored in this study.

The values corresponding to $w(t)$ were obtained using MATLAB's numerical integration function, "integral()" and applying a proportionality factor (k); through trial and error, this was found to work well with $k=1 \times 10^{-5}$. Both $\frac{dw}{dt}$ and $w(t)$ with $t \leq N = 5 \times 10^5$ iterations are displayed in **Figure 3.2**(a) and (b), respectively.

Because QTA occasionally accepts worse solutions, it stands to reason that the barrier width does not grow continuously but rather experiences local width fluctuations combined with a globally increasing trend. Therefore, in addition to the MOCVD-inspired barrier width schedule, another schedule was also tested, defined as:

$$w^a(t) = w^b(t) \left(\sin^2 \frac{50\pi t}{N} + 1 \right) \quad 3.7$$

Where:

$$w^b(t) = 10 \times \sqrt[3]{(w' * t)} \quad 3.8$$

With $w' > 0$ used as a tunable parameter to control how quickly the width increases over the course of the optimization. The form of w^a was chosen to introduce more local variations in the

barrier width schedule in addition to the global trend of increasing width at a decreasing rate. This was done by coupling a fractional power function (given by w^b) with a sinusoidal function. A squared sine function was chosen to ensure that the width was always at least as large as the global trend. For an annealing schedule where $N = 5 \times 10^5$, the period of 10,000 corresponded to 10 full cycles during the search time.

3.2.3 Simulated Annealing (SA)

For comparison purposes, we used SA, a stochastic search algorithm, which was first introduced for IMRT optimization by Webb in 1989.^{1,26} Like QTA, SA models the optimization problem as a system which undergoes a biased random walk. Over the course of the random walk, the system will always accept new solutions, which improve on the old solution. In order to avoid getting trapped in local minima, the system accepts worse solutions with a probability proportional to $\exp\left(\frac{-(V_{new} - V_{old})}{T}\right)$, where T is the temperature of the system that is annealed (decreased) over the course of the algorithm search. Mathematically, SA was proven to converge to a global optimal solution with minor assumptions on the cooling schedule and appropriate conditions on irreducibility, aperiodicity and reversibility of the induced Markov chain.^{27,28} Because SA has a long history of use in our clinic and the literature, it served as our benchmark algorithm for evaluating the success of QTA.²⁹ The annealing schedule for T was identical to the schedule used for the QTA annealing variable, Γ , and is defined in Equation 3.6. Note that while the formalism of QTA shares many similarities with SA, the probability of accepting a worse solution in QTA differs from SA in two key respects: (1) reduced dependence on the potential energy difference between the current and new solution and (2) the presence of an additional dynamic parameter in the barrier width. These differences provide QTA with more freedom to explore the solution space.

3.2.4 IMRT Case Selection

To analyze the performance of our quantum-inspired algorithm, we compared QTA and SA on two stereotactic body radiation therapy (SBRT) liver cases chosen from the University of Michigan Radiation Oncology Department's clinical database.

Case 1, a 12-field three-dimensional IMRT liver plan, was selected as an “easy” test case to confirm that both QTA and SA were performing properly. This case was not expected to pose a significant challenge for either optimization algorithm because it featured a minimal amount of overlap between the planning target volume (PTV) and the liver, and no overlap with other structures. For simplicity, the structures selected for optimization from Case 1 were the PTV and liver exclusive of the gross tumor volume (Liver – GTV) as shown **Figure 3.3(a)**. Influence matrices for these structures were calculated using built-in functions defined in the Eclipse scripting application program interface (API). The voxel size used was 2 mm and the beamlet size was 5 mm x 5 mm, for a total of 158,720 voxels, 768 beamlets, and 1,602,504 nonzero elements in the dose influence matrices.

Case 2 served as a “challenge” case to determine if the additional degree of freedom associated with QTA facilitated better results—such as plan quality, robustness, or speed—for more clinically relevant and difficult optimization problems. Designed as a 5-field IMRT plan, Case 2 was selected because it had significant overlap between the PTV, stomach, and liver structures as shown in **Figure 3.3(b)**. Because this was a proof of concept study, only a subset of structures from the original treatment plan were included in our optimization. The structures were selected based on the priority assigned to them in the original clinical treatment plan. In addition, the dose volume histogram (DVH) constraints were also inspired by those used clinically. The influence matrices for these structures (3 mm voxel size, 2.5 mm x 5 mm beamlet size) were again

calculated using built-in functions available in the Eclipse scripting API. Note that smaller beamlet sizes were purposefully chosen for Case 2 to increase the number of optimization parameters. Case 2 contained 79,977 voxels, 4166 beamlets, and 1,558,612 nonzero elements in the dose influence matrices. Because Case 2 contained more than four times more beamlet weights, it also represented a more challenging optimization problem than Case 1. The DVH constraints used in the optimization of Case 1 and Case 2 can be viewed in **Table 3.1** and **Table 3.2** respectively.




Figure 3.3: CT scans show contours for structures optimized for Case 1 (a) and Case 2 (b), respectively. Case 1 features a PTV that is roughly spherical in shape and far from major organs. (with the exception of the liver). Case 2 features a PTV with convex geometrical features and close proximity to both the liver and the stomach.

Table 3.1: DVH constraints applied to objective function for Case 1.

Case 1 DVH Constraints					
Structure	Constraint Type	Limit	Volume (%)	Dose [Gy]	Penalty
PTV	DVH Point	Lower	100	29.7	50
	DVH Point	Lower	95	30	50
	DVH Point	Upper	0	60	100
	Max Dose Range	N/A	N/A	[30 42]	100
Liver - GTV	DVH Point	Upper	0	42	100

Table 3.2: DVH constraints applied to objective function for Case 2.

Case 2 DVH Constraints					
Structure	Constraint Type	Limit	Volume (%)	Dose [Gy]	Penalty
PTV	DVH Point	Lower	99	33	100
	DVH Point	Lower	95	30	200
	DVH Point	Lower	100	28	200
	DVH Point	Upper	0	48	160
GTV Deformed MR	DVH Point	Lower	100	43	100
	DVH Point	Upper	0	48	160
Liver - GTV	Mean	N/A	N/A	4	50
Stomach	DVH Point	Upper	0.001	28	150
Stomach PRV	DVH Point	Upper	0.003	25	300

3.2.5 Objective Function

The objective function used for both SA and QTA IMRT optimization is defined by:

$$\begin{aligned} & \min_{\mathbf{b}} E(\mathbf{b}) \\ & \text{subject to } \mathbf{b} \geq 0 \end{aligned} \tag{3.9}$$

Where:

$$\begin{aligned} E(\mathbf{b}) = & \sum_{n=1}^N \frac{\lambda_n \| \mathbf{D}_n - \mathbf{d}_n(\mathbf{b}) \|^2}{J_n} \\ & + \beta \sum_{m=1}^M \sum_{ij} |(L \times B_m)_{ij}|^2 + \sum_{n=1}^N \alpha_n P_n(\mathbf{b}, \text{DVH constraints}) \end{aligned} \tag{3.10}$$

And:

$$\mathbf{d}_n(\mathbf{b}) = \mathbf{I}_n * \mathbf{b}. \tag{3.11}$$

The first term in the objective function represents the mean squared error between the prescribed dose, \mathbf{D}_n , and the delivered dose, $\mathbf{d}_n(\mathbf{b})$, for each structure n of N structures. $\mathbf{d}_n(\mathbf{b})$ represents the dose delivered to each voxel in structure n , and is defined in Equation 3.11 as the product of the structure's influence matrix, \mathbf{I}_n , and the beamlet-weight vector, \mathbf{b} . J_n is the number of voxels in structure n . The influence matrices for each structure were calculated using the Eclipse Scripting API's built-in "CalculateInfluenceMatrixToMemory()" function. The point cloud which was input into this influence matrix function was calculated using an in-house script that generates a normally distributed random set of point locations whose average distance is the cube root of the desired voxel size.

For an influence matrix \mathbf{I} , matrix element I_{ij} is defined as the dose contribution to voxel i from beamlet j . Any given beamlet is expected to contribute primarily to the voxels it overlaps with and their nearest neighbors. However, due to scattering effects Eclipse-generated influence matrices contain no non-0 values; they contain a subset of elements whose values are orders of magnitude smaller than the largest values in the matrix—corresponding to a beamlet's contribution to a distant voxel. To facilitate faster optimization, a tolerance value was defined below which

influence values were deemed negligibly small and reset to 0. This allowed for the influence matrices to be saved as sparse matrices, reducing calculation times. An acceptable tolerance value was determined by trial and error to be 0.015—this was done by loading fluence vectors that were optimized using filtered influence matrices into the Eclipse scripting API, performing MLC leaf sequencing and dose volume calculation, and comparing the resulting DVH histograms with those produced in-house.

The second term in objective function represents a smoothing penalty which was implemented to ensure the treatment plans could be delivered efficiently. In order to determine the optimal filter, L , a series of QTA optimizations were performed on Case 2 using a number of different filter types—including median, Savitzky-Golay (SG), plan intensity map variation (PIMV), and Laplacian and Laplace of Gaussian (LOG) filters with kernels of sizes 3, 5, 7, 9, and 15, respectively.³⁰ For the smoothing filters, a penalty value was defined as the squared difference between the original and smoothed fluence map. For the PIMV-type filter, the square of the PIMV value for each beam was used as the penalty. For the edge-finding filters of kernel size n , the filter kernel was convolved with the beamlet matrix B_m (reshaped from the beamlet weight vector) for each beam. The squared sum of the indices of the resulting matrix yielded a value correlated to the degree of irregularity for each beamlet matrix. With the exception of the Laplace filter and the PIMV filter, all filters tested were implemented using MATLAB built-in functions. Each filter's performance was evaluated by visually inspecting fluence maps and comparing the total number of monitor units (MUs) necessary after MLC leaf sequencing.

The third term in Equation 3.10 represents additional penalties based on DVH constraints associated with each structure. The dose constraints and penalties, P_n , used in each case can be viewed in **Table 3.1** and **Table 3.2**. For Case 1, simple Boolean conditions were used to assign

penalty values (for example, if 99% of the PTV receives < 33 Gy, add 100 to the DVH penalty). The weighting factors α_n used in Case 1 were set to 1 for all structures. For the more challenging Case 2, we found it necessary to adjust the calculation of the DVH penalty. Specifically, for Case 2, penalties for missed DVH constraints were assigned as the penalty value (listed in the last columns of **Table 3.1** and **Table 3.2**) multiplied by the absolute difference between the DVH constraint and the actual metric achieved. For example, if 99% of the PTV volume received ≥ 29 Gy, the penalty for that constraint would be $(30-29) \times 100$. Because the constraint type is designated as “lower,” no penalty is assigned if 99% of the PTV volume receives > 33 Gy. Finally, for Case 2, α_{PTV} and $\alpha_{Liver-GTV}$ were set to 9 and 10, respectively.

3.2.6 Extension to Influence-based Direct Aperture Optimization

In addition to fluence map optimization, the objective function described in Section 3.2.5 can be generalized to directly optimize apertures (defined by MLC leaf positions) and their weights using a method known as influence-based direct aperture optimization (DAO).^{31,32} This is accomplished by defining the fluence weights as a function of the MLC leaf segment positions and aperture weights, which for small beamlets can be written as:

$$\mathbf{b}(\mathbf{l}, \mathbf{w}) = \sum_i \mathbf{T}(\mathbf{l})_i \times \mathbf{w}_i, \quad 3.12$$

Where \mathbf{l} defines the MLC leaf positions, \mathbf{w}_i is the weight assigned to aperture i , and $\mathbf{T}(\mathbf{l})_i$ is a transmission matrix whose values represent the fraction of each beamlet unobstructed by the MLC leaf segments for aperture i .³¹

3.2.7 Criteria for Convergence

In order to compare QTA and SA’s performance in a faithful manner, it is necessary to develop a quantitative method for defining convergence. For both optimization methods, the energy at each iteration t was saved in a vector, $\mathbf{E}(t)$. The gradient of $\mathbf{E}(t)$ was calculated numerically in

MATLAB. From this gradient, a moving average mean (MAM) with width 100 was then calculated. A tolerance value, c^{tol} , was selected by trial and error, and the largest index position, j (for which $|\text{MAM}(j)| > c^{\text{tol}}$) was identified. The convergence point for the algorithm was then defined as iteration $j+1$. **Figure 3.4** displays the process of finding the j (and thus $j+1$) from $\mathbf{E}(t)$. An appropriate value for c^{tol} was found to be 0.1.




Figure 3.4: Process of calculating the convergence iteration number from a representative QTA optimization. The saved energy history is represented by (a), and is used to calculate the energy gradient ((b)). From this gradient a MAM of width 100 was calculated ((c)). The black vertical line is plotted at the maximum iteration number j for $\max_j(|\text{MAM}(j)| > c^{\text{tol}})$.

3.2.8 Computing Environment

All beamlet-weight optimizations described in this paper were performed using MATLAB scripts with graphics processing unit (GPU) acceleration on the University of Michigan’s High-Performance Computing Linux-based cluster, Flux (central processing unit (CPU): Intel Haswell, GPU: Nvidia K40). Each job was submitted with 2 CPU cores (4 GB/core) and 1 GPU.

MLC leaf sequencing and dose volume calculations used for final plan visualization were performed using clinical software (Varian Medical Systems, Inc. Eclipse Treatment Planning System: Varian Leaf Motion Calculator Version 13.6.23, Anisotropic Analytical Algorithm Version 15.5.11).

The complete QTA algorithm for IMRT optimization is summarized in **Figure 3.5**. The maximum possible number of iterations performed in each run was defined as $N = 5 * 10^5$. Because the parameter N was used as a variable in both the annealing schedule (T or Γ) and the barrier width schedule (w), its value was not altered over the course of the reported studies. Therefore, in order to vary the actual number of iterations performed, an additional break parameter was defined which forced the algorithm to end early at iteration $n = n^{\text{break}}$. This break parameter was implemented both to shorten the duration of optimizations when it was clear an optimal solution had been reached prior to N as well as to confirm that the convergence iteration numbers—whose calculation was described in Section 3.2.7—represented clinically acceptable plans.

Algorithm 1: QTA

```
Initialize:  $b_{solution} \leftarrow b_0$ ;  
 $V(b) := \sum_{n=1}^N \lambda_n \|D_n - d_n(b)\|^2 + \beta \sum_{m=1}^M \sum_{ij} |(L * B_m)_{ij}|^2 +$   
 $\sum_{n=1}^N \alpha_n P_n(b, DVHconstraints)$  ;  
 $d(b) := I * b$ ;  
 $w \leftarrow w_0$ ;  
 $\mathcal{T} \leftarrow \mathcal{T}_0$ ;  
for  $I_{ij}$  in  $I$  do  
  if  $I_{ij} < tol$  then  
     $I_{ij} = 0$   
  end  
end  
for  $n \leftarrow 1$  to  $N_{max}$  do  
   $\mathcal{T}_n \leftarrow \mathcal{T}_0(1 - \log(n)/\log(N))$ ;  
   $w_n \leftarrow w_{n+1}$ ;  
   $b_n \leftarrow \max(0, b_n + r_{neighbor})$ ;  
  if  $n = n_{break}$  then  
     $b_{solution} \leftarrow SG_{filter}(b_{solution})$ ;  
    for  $b_i$  in  $b_{solution}$  do  
      if  $b_i < 0$  then  
         $b_i = 0$   
      end  
    end  
    break;  
  end  
  if  $V(b_n) - V(b_{n-1}) < 0$  then  
     $b_{solution} \leftarrow b_n$ ;  
  else  
    if  $\exp\left(\frac{-w\sqrt{E(b_n)-E(b_{n-1})}}{\mathcal{T}}\right) > rand$  then  
       $b_{solution} \leftarrow b_n$ ;  
    end  
  end  
  if  $n = N$  then  
     $b_{solution} \leftarrow SG_{filter}(b_{solution})$ ;  
    for  $b_i$  in  $b_{solution}$  do  
      if  $b_i < 0$  then  
         $b_i = 0$   
      end  
    end  
    break;  
  end  
end  
end
```

Figure 3.5: Quantum Tunnel Annealing (QTA) algorithm for intensity modulated radiation therapy (IMRT) optimization.

3.3 Results

3.3.1 Case 1

Preliminary studies on a geometrically simple case, designated “Case 1,” confirmed that the QTA and SA algorithms were performing properly. **Figure 3.6(a)** through (d) displays the DVH and potential energy (PE) trajectory results acquired by running the QTA and SA algorithms 20 times each for $N = 5 \times 10^5$ iterations and no premature breaks (i.e., $n^{\text{break}} > N$). **Figure 3.6(e)** through (f) displays representative dose distributions for QA and SA, respectively, which were calculated in Eclipse using optimized beamlet-weights from the tenth run. For Case 1, the incorporation of a Laplace edge-finding filter with a kernel size of 3 into the objective function was found to yield sufficiently deliverable plans. Beamlet-weights generated from both QTA and SA were found to consistently yield plans that satisfied the DVH constraints.

The DVH curves for QTA (**Figure 3.6(a)**) and SA (**Figure 3.6(b)**) indicate that for case 1, QTA exhibited greater stability over SA with respect to the quality of the final plan. SA converged to a solution with worse PTV coverage 60% of the time. **Figure 3.6(c)** and (d) display the PE trajectories for the QTA and SA runs, respectively. The PE trajectories for QTA indicate that QTA explored higher energy solutions prior to sudden extreme drops around the $(n = 1 \times 10^5)^{\text{th}}$ iteration, whereas SA featured a more linear decrease. The resulting dose distributions were found to be similar between both algorithms and featured reasonable tumor coverage while minimizing the dose to the surrounding normal tissues.

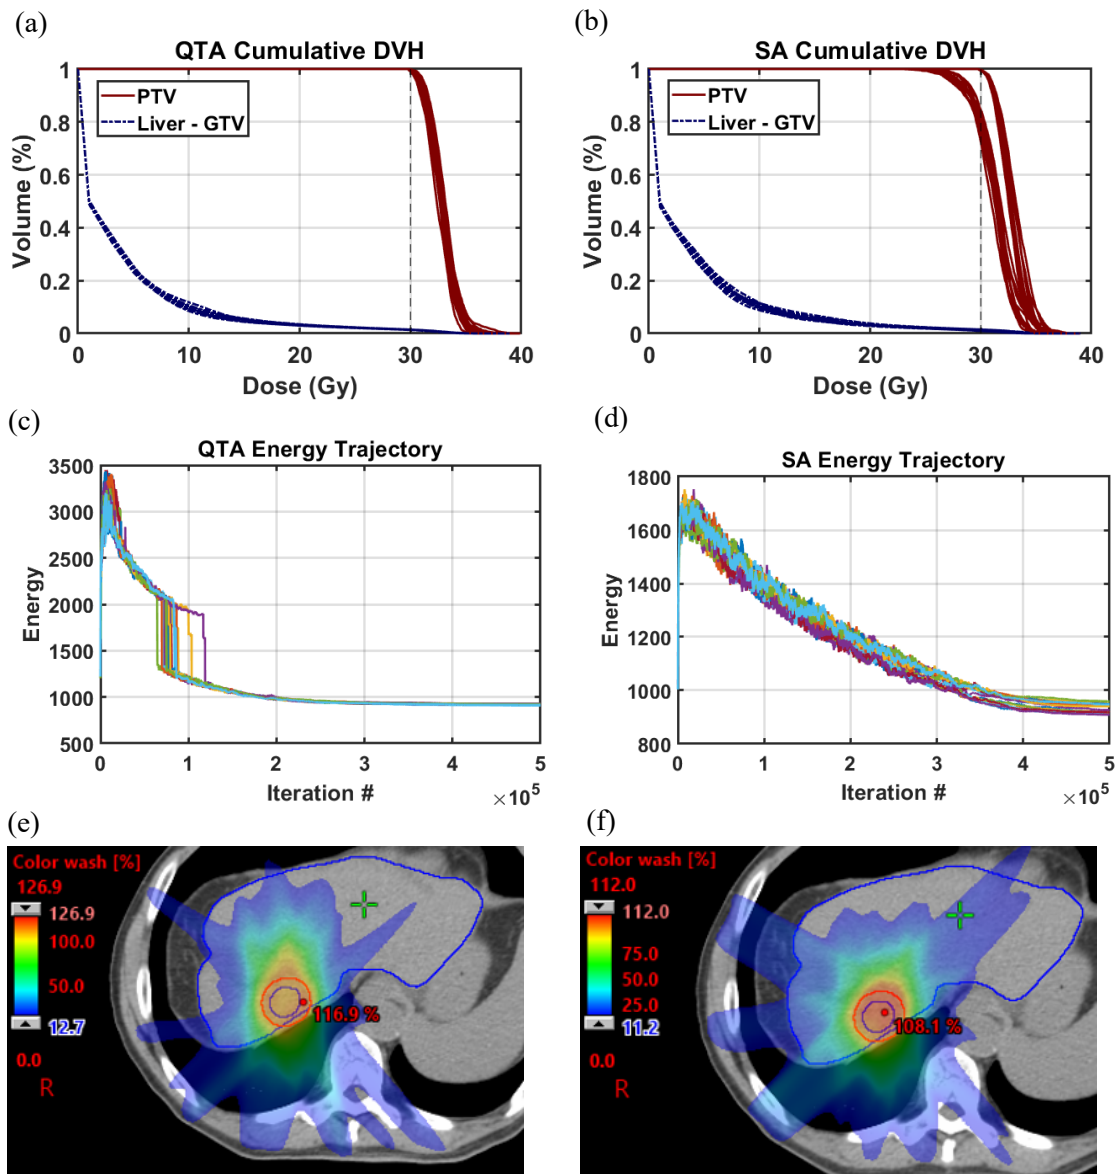


Figure 3.6: Optimization results for QTA and SA applied to Case 1. (a) and (b) display DVH curves for 10 separate optimizations using QTA and SA, respectively. (c) and (d) display the PE trajectories for the 10 QTA and SA optimizations. (e) and (f) display representative dose distributions calculated in Eclipse using fluence values from the 10th QTA and SA optimization.

3.3.2 Case 2: Redefined Smoothing Filter

In the pursuit of designing an objective function that can produce clinically acceptable and deliverable plans, a comprehensive study (described in detail in Section 3.2.5) was performed to determine the optimal measure of smoothness for use as a penalty in the objective function.

Smoothness was assessed qualitatively using the fluence maps and quantitatively using the total MU required (summed over each beam). **Figure 3.7(a)** displays the optimized fluence map for one of the Case 2 beams using a LOG filter within the objective function. The speckled appearance of **Figure 3.7(a)** suggests that smoothing within the objective function alone is not sufficient, and the MU necessary for this plan was more than 20% larger than predicted for an Eclipse-optimized plan which met the same DVH constraints. Adjustments to the size of the kernel and the type of filter used within the objective function did not yield discernable improvement to fluence regularity or total MU.

We also explored directly applying a smoothing filter to the beamlet weights outside of the objective function. We found that the optimal smoothing process consisted of the 7x7 LOG filter within the objective function, combined with a two-dimensional SG filter applied to the beamlet-weights during the final iteration of the algorithm. The optimized fluence map using this refined smoothing filter is displayed in **Figure 3.7(b)** and appears markedly smoother than the LOG-filter alone. This refined smoothing filter resulted in a total of 2877 MU, which was 34% lower than the LOG filter alone and more than 20% lower than the Eclipse-optimized plan. The plan quality, as gauged by DVH constraints, experienced only a slight reduction.

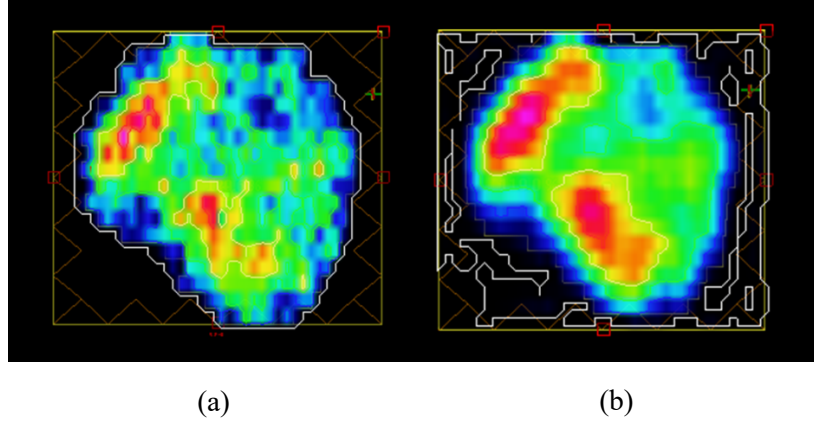


Figure 3.7: (a) displays the fluence map results for a single beam in Case 2 resulting from the QTA optimization without refined smoothing. (b) displays the fluence map results from QTA optimization with refined smoothing.

3.3.3 Case 2: Barrier Width Schedule Effect

As discussed in 3.2.2, four different barrier width schedules were investigated for QTA. One was inspired by the growth rate of GaAs in MOCVD, while the remaining three were designed to allow for local fluctuations in the barrier width within a globally increasing trend. **Table 3.3** lists convergence rates calculated for QTA optimizations using the four barrier-width schedules as well as optimizations for SA. Three of the four barrier widths tested yielded convergence faster than SA. The optimal barrier width schedule was found to be the w^a function with $w' = 1 \times 10^{-5}$, and it reached convergence in less than half the time required for SA.

Table 3.3: Convergence times (in seconds) for QTA with different barrier width schedules as well as SA.

Case 2 Convergence Results QTA Barrier Width Testing	
Algorithm type	Convergence (s)
SA	1062.5
QTA, $w^a, w' = 1 \times 10^{-5}$	528.6
QTA, $w^a, w' = 1 \times 10^{-7}$	637.2
QTA, $w^a, w' = 1 \times 10^{-9}$	1762
QTA, MOCVD	874.2

3.3.4 Case 2: Annealing Schedule Effect

Each algorithm's sensitivity to the choice of annealing schedule was assessed by comparing their performance across five different functions (shown in **Figure 3.8(a)**): T1, a linear function; T2, a sigmoidal function; T3, an exponential function; T4, a logarithmic function; and T5, a power law function with fractional exponent. **Figure 3.8(b)** displays box and whisker plots of the convergence rates for QTA and SA respectively for each annealing schedule. For schedules T1, T2, and T3, SA failed to reach convergence prior to the breakpoint at $n^{break} = 2.5 \times 10^5$, resulting in the tight spread of data for SA at these schedules. QTA exhibited lower average convergence rates for all five annealing schedules. Note that for this paper, T4 (defined by Equation 3.6) served as the default annealing schedule because it was found to yield reasonable optimization results during early on in the project.

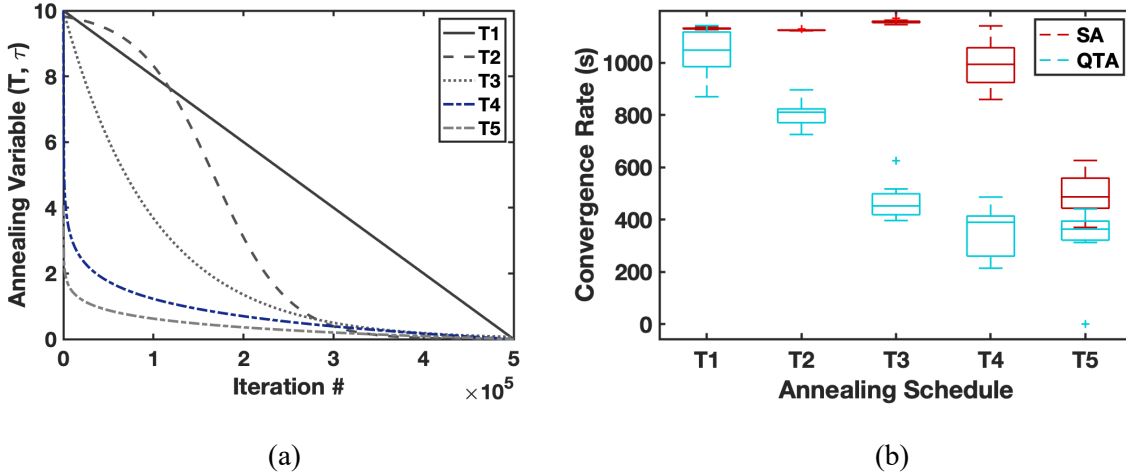


Figure 3.8: Displays the annealing schedule functions tested for QTA and SA. Note that T4 was the annealing schedule used for all remaining studies. (b) displays box and whisker plots of the convergence results for QTA and SA, respectively, for each annealing schedule.

3.3.5 Case 2: Optimization Stability

QTA's stability was verified by performing a series of optimization tests using different starting beamlet-weight vectors (10 runs per initial beamlet-weight vector tested, $N = 5 \times 10^5$ iterations, which ran until $n^{break} = 2.5 \times 10^5$). As a comparison, SA optimizations were also performed under the same conditions. The optimizations began with initial beamlet-weight vector values set to 0, 11, and 20, respectively. These values represent the minimum, average, and maximum fluence values expected for the optimized beamlet-weight vector. In addition, tests were also conducted using an initial beamlet-weight vector whose values were randomly distributed over a range from 0-20. In order to assess whether QTA is primarily advantageous later in the annealing schedule after the algorithm has become stuck in local minima, additional tests were performed on a hybrid SA-QTA algorithm, which ran SA for the first 5×10^4 iterations after which the algorithm switched to QTA. The initial beamlet-weights used for the hybrid tests were also randomly distributed over a range from 0-22.

Figure 3.9 displays the results for 10 QTA and SA optimizations using the randomly distributed initial beamlet-weight vector. The DVH curves for QTA (**Figure 3.9(a)**) and SA (**Figure 3.9(b)**) suggest that both reached final solutions with nearly identical dose coverage. This finding was found to hold for all iterations regardless of the initial beamlet-weights used. The energy trajectories for QTA (**Figure 3.9(c)**) and SA (**Figure 3.9(d)**) are plotted on a Log scale to highlight differences in the shape of the curves. Like Case 1, the QTA PE trajectories for Case 2 feature a region of rapid descent, located just after the 10^3 iteration. All QTA and SA runs (for $b_0 = 0, 1, 20$, and rand) required $\geq 1.1 \times 10^5$ iterations to reach convergence. **Figure 3.9(e)** and (f) display the Eclipse-calculated dose distributions from the tenth optimization for QTA and SA, respectively. The final dose distributions were found to be nearly identical and exhibited reasonable dose coverage (in comparison to Eclipse-optimized plans).

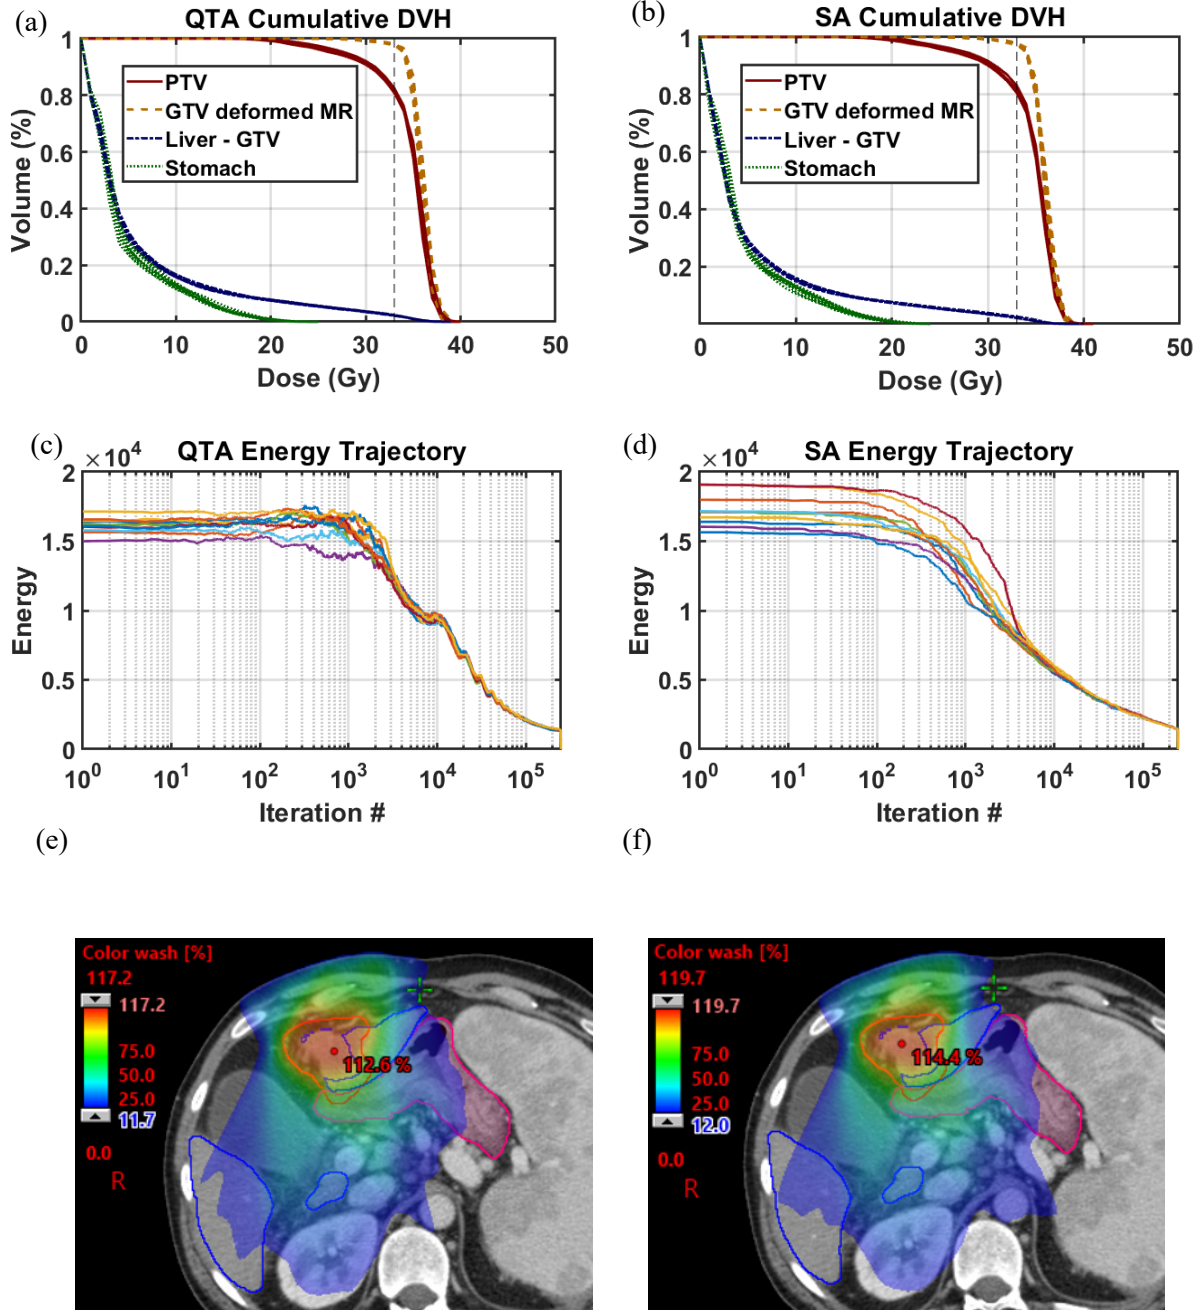


Figure 3.9: DVH bands, PE trajectories, and representative dose distributions for stochastic optimizations ($N=10$, 500000 iterations) with the initial beamlet-weight vector set to random values uniformly distributed between 0 and 22 on a challenging SBRT liver case for QTA ((a), (c), and (e)) and SA ((b), (d), and (f)), respectively.

Table 3.4 displays the mean convergence rates (in seconds) for QTA, SA, and the hybrid SA-QTA algorithm. QTA consistently exhibited faster convergence rates and had smaller standard

deviations than SA in all but one case ($b_0 = 11$). The convergence rates of the SA-QTA hybrid algorithm were similar to the performance of SA.

Table 3.4: Mean convergence times (in seconds) for QTA and SA with perturbations to the initial beamlet-weight values.

Case 2 Convergence Results for Initial Beamlet-Weight Testing			
b_0	QTA (s)	SA (s)	SA-QTA hybrid (s)
0	637.9 ± 63.2	982.4 ± 96.3	N/A
11	644.8 ± 84.4	987.4 ± 82.1	N/A
22	693.2 ± 75.9	1103 ± 84.0	N/A
Randomly distributed	611.0 ± 72.4	996.4 ± 103.0	953.1 ± 65.5

The stability of QTA and SA was also assessed by making perturbations in the original dose constraints. For each of these tests, a perturbation was made to a single constraint while all others were held constant. Each optimization was run for $N = 5 \times 10^5$ iterations and stopped at $n^{break} = 2.5 \times 10^5$.

Table 3.5 summarizes the perturbations tested and the corresponding convergence rates (in seconds). For all perturbation types, QTA exhibited faster convergence. However, the percent difference in the perturbed convergence rates from the original convergence rate ranged from 5.95%-43.7% for QTA and 4.1%-5.1% for SA, indicating that QTA may exhibit higher sensitivity than SA.

Table 3.5: Parameter changes and convergence times (in seconds) for QTA and SA with perturbations to the original dose constraints.

Case 2 Convergence Results for Parameter Testing					
Parameter	Organ	from	to	QTA (s)	SA (s)
DVH	Stomach	Max dose = 28Gy	Max dose = 18Gy	788.8	1038.9
DVH	Liver	Mean dose = 4Gy	Mean dose = 2Gy	951.6	1108.2
DVH and Dose	PTV	Target dose = 33Gy	Target Dose = 43Gy	701.2	1028.6
Original	Original	N/A	N/A	661.8	1083.3

3.3.6 Aperture-Weight Optimization via Influence-Based DAO

Influence-based DAO was performed on QTA and SA for 10 runs per initial beamlet-weight vector tested ($N = 5 \times 10^5$ iterations running until $n^{break} = 2.5 \times 10^5$) using the fluence approximation formalism described in 3.2.6. For these optimizations, leaf segment information was extracted from a pre-existing Eclipse-optimized IMRT plan (with a total of 431 apertures) on Case 2, and aperture weights were optimized with the starting weight of each aperture set to 0. **Figure 3.10** displays the resulting cumulative DVHs ((a) and (b)), potential energy trajectories ((c) and (d)), and representative dose distributions ((e) and (f)) for QTA and SA, respectively.

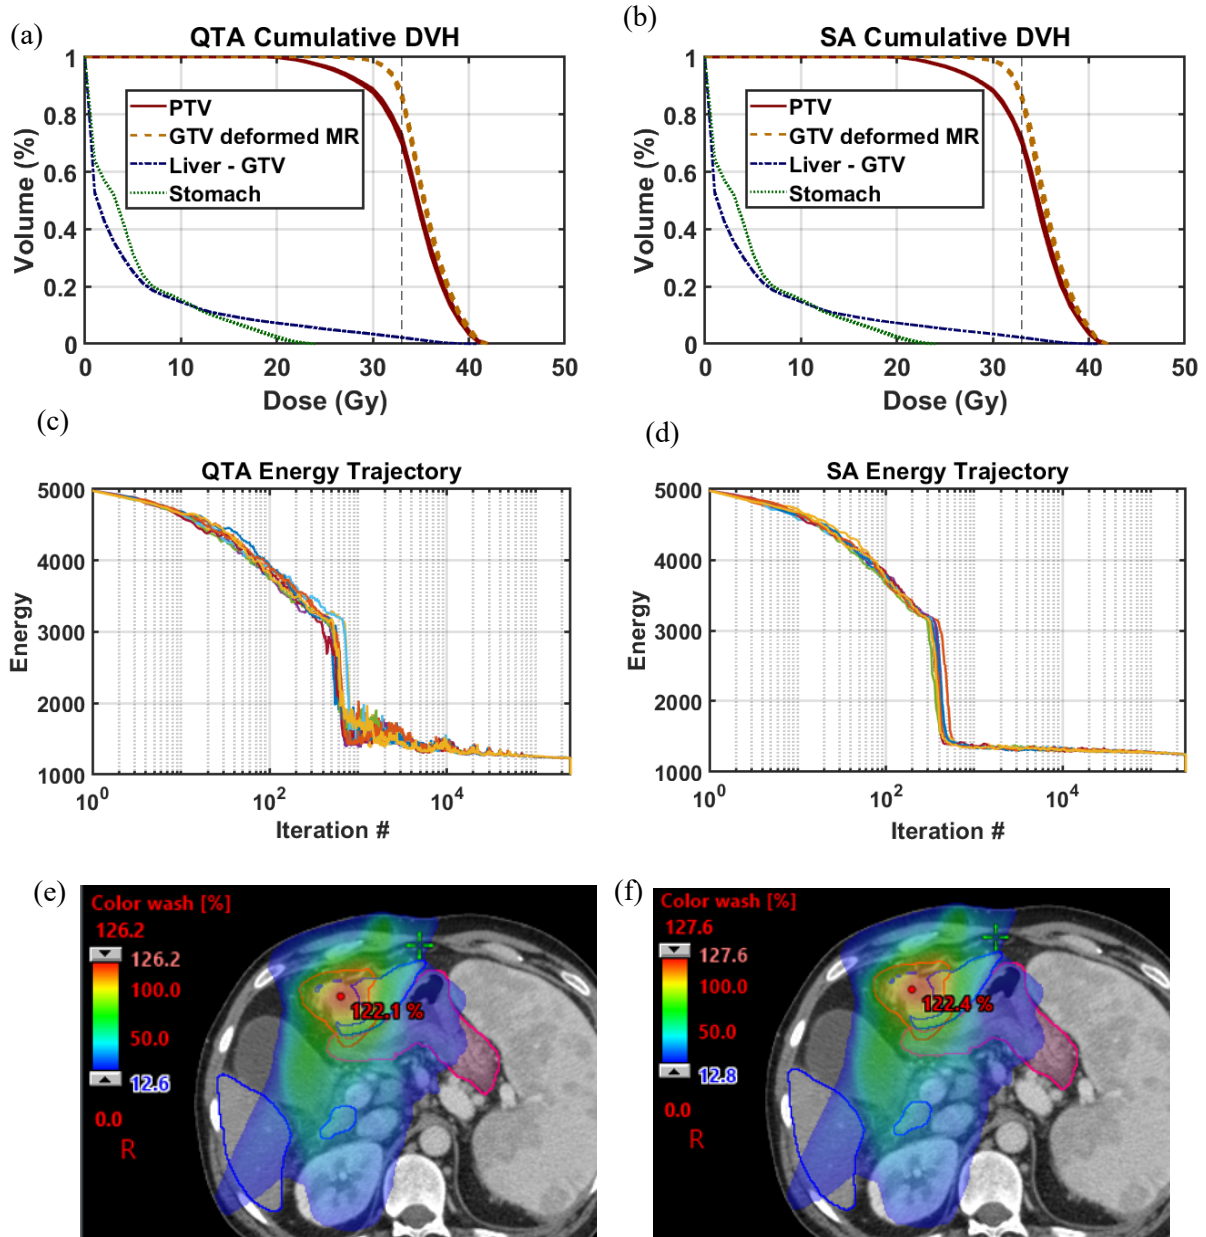


Figure 3.10: DVH bands, PE trajectories, and representative dose distributions for stochastic optimizations ($N = 10$, 250000 iterations) of aperture weights on a challenging SBRT liver case for QTA ((a), (c), and (e)) and SA ((b), (d), and (f)), respectively.

The DVHs displayed in **Figure 3.10(a)** and (b) indicate that QTA and SA achieved comparable tumor coverage and OAR sparing. While the energy trajectories in **Figure 3.10(c)** and (d) indicate that QTA exhibited more stochastic exploration of the solution space early on in the optimization, QTA converged within $5,234 \pm 622.4$ (s) on average while SA had an average

convergence rate of $7,151 \pm 504.5$ (s). **Figure 3.10**(e) and (f) show that both algorithms also produced similar dose distributions.

3.4 Discussion

The optimization results from Case 1 confirmed that both algorithms were capable of delivering clinically acceptable results. QTA was found to be more stable than SA with regard to the quality of the final solution to which it converged, as SA converged to a worse solution 60% of the time.

Because it was more geometrically complex, Case 2 was used to characterize QTA's performance. One of the ways QTA distinguishes itself from SA is that the probability of accepting a worse solution during the course of the optimization is a function of the estimated width of the potential energy barrier, providing an additional degree of freedom with which to explore the solution space. We tested several expressions which were heuristically selected to represent the barrier-width function. Adjusting the form of the barrier-width function did not influence the quality of the final plan if the algorithm was allowed to run for its fully allotted time. However, the form of the barrier-width function did influence how quickly the algorithm reached convergence. The convergence results listed in **Table 3.3** suggest that the barrier-width function can be used as a tunable parameter to achieve faster convergence. While further tests are warranted to determine an optimal expression for the barrier width, the majority of the functions tested yielded faster convergence rates than SA.

The convergence rates of both algorithms were found to be dependent on the annealing schedule chosen. For three of the five functions tested, SA failed to converge 30-80% of the time, while QTA reached convergence for all five evaluated functions. In addition, QTA had faster mean convergence rates for all five annealing functions tested. These results suggest that QTA is more robust against the choice of annealing schedule. Another way to conceptualize this advantage is to

interpret QTA as having a modified annealing schedule in which the barrier width function serves as an additional time-dependent, tunable parameter, coupled with a dampened dependence on the energy difference between the current and new solution.

Testing the sensitivity of QTA with respect to changes in the initial beamlet-weights, b_0 , is useful for determining whether the algorithm can reliably deliver clinically acceptable plans under conditions where a “good” first guess is unknown. In initial beamlet-weight tests (described in Section 3.2.4) we found that QTA consistently achieved faster convergence times over SA across all variations of b_0 .

Unlike Case 1, **Figure 3.9(a)** and (b) suggest that both QTA and SA consistently achieved final solutions of nearly identical plan quality for Case 2. These findings held even after varying the initial starting guess. These results may seem surprising given that Case 2 represented the more challenging case. The explanation lies in the difference between the objective functions used for Cases 1 and 2, which are described in detail in Section 3.3.5. Case 1 penalties based on the DVH constraints were assigned using Boolean conditions. Implementing the DVH constraint portion of the objective function was found to be insufficient for Case 2 because it could not provide sufficient PTV coverage without delivering an excessive dose to the organs at risk. Therefore, when we began working on Case 2, it was necessary to adjust the objective function so that penalties based on the DVH constraints were weighted more heavily as plan results strayed farther from the objectives. The difference in results between Case 1 and Case 2 suggest that the additional constraints applied to Case 2’s objective function narrowed the solution space available to the algorithms. In light of this point, the combined results from both cases suggest that QTA is more robust than SA to changes in the formulation of the objective function.

To assess QTA's sensitivity to changes in treatment plan goals, a series of optimizations were run for QTA and SA in which perturbations were made to the PTV dose prescription and to OAR dose constraints. It was found that while QTA continued to achieve faster convergence rates, those rates exhibited greater variation from the original, unperturbed convergence rate.

For Case 2, it was found that implementing smoothing only within the objective function was insufficient for producing plans with clinically deliverable fluence maps. This is likely due to the algorithms' stochastic nature and the fact that Case 2 contained more than four times the number of beamlets as Case 1. Ultimately, a refined smoothing technique was developed which combined a LOG filter—used to define an irregularity penalty in the objective function, with a two-dimensional SG filter that was applied to each beamlet map during the final iteration. The resulting fluence maps for these plans had total MU values which were more than 20% less than those for an Eclipse-optimized plan. It is perhaps unorthodox to include a smoothing filter outside of the objective function, as this can compromise plan quality.³⁰ However, we found that implementing the SG filter during the optimization's final iteration had only a small impact on plan quality, and all plans generated using this technique were comparable in quality to plans generated using Eclipse-based IMRT optimization.

In order to further investigate the potential of QTA over SA, it is necessary to test additional optimization formalisms with known ill-behavior. One such representative approach is to estimate the aperture weights directly using the influence-based DAO approach described in Section 3.2.6. DAO was evaluated on the more complex Case 2. The results from these optimizations (presented in Section 3.3.6) indicate that on average QTA converged up to 26.8% faster than SA. DAO is a more complex optimization problem than fluence optimization. While the results of this study

example may suggest that the performance gap between QTA and SA seemingly becomes narrower, QTA still exhibits notable benefits over SA overall.

The limitations of this study are summarized as follows: Because only two patient cases were considered, our knowledge of the algorithm's sensitivity to different cases is still limited. We chose to only optimize the most challenging and critical structures in each case; for this reason, the convergence times reported are not representative of a full treatment plan. In addition, the expression used in QTA to define the probability of a particle tunneling through a potential energy barrier contains weaknesses in its assumptions about the size of the annealing variable, Γ . Due to these assumptions, while the formulation for QTA can be described as quantum-inspired, it does not represent a true simulation of a quantum process. Nevertheless, QTA was found to exhibit several qualities that suggest it might be an attractive candidate for applications which necessitate rapid optimization of complex or challenging treatment plans. QTA consistently performed faster than SA across multiple types of perturbations and yielded treatment plans of equal quality. Furthermore, a hybrid SA-QTA algorithm was found to perform only slightly better than SA alone, reinforcing the merit of the full QTA algorithm. The presence of an additional degree of freedom represented by the barrier width schedule leaves open the possibility that this parameter might be further fine-tuned to achieve even faster results.

The results of this study suggest that the extra degree of freedom associated with QTA's barrier-width schedule allowed for the algorithm to be better "tuned" to converge at faster rates than SA. Natural future directions for this work include performing QTA optimizations on full IMRT treatment plans and VMAT plans, which represent a larger optimization problem. Based on QTA's computational speedup and ability to escape local minima, it may be a useful tool for computationally demanding adaptive radiotherapy applications. Finally, QTA would be a valuable

tool for implementing more complex (typically non-convex) objectives based on biological optimization objectives combining imaging and molecular biomarkers with dose-response functions derived via multiple outcome and utility modeling methods,^{33,34} which as of now are hindered in clinical implementation by a lack of efficient and robust optimization techniques.

In addition to further studies applying QTA to more challenging treatment problems, we would also like to explore whether implementing QTA on a quantum computer could lead to greater computational speedup. In their 2015 study, Nazareth and Spaans reported on the first use of a quantum annealing computer for IMRT beamlet weight optimization; they found that while SA consistently produced higher quality plans, optimizations performed on a quantum annealing device (using a modified version of Tabu Search as the optimization algorithm) were >2.5 times faster than SA.¹⁹ At the time of their study, the researchers were limited to a 512 qubit device, which restricted the complexity of the treatment problems they could tackle. For reference, if the beamlets in Case 1 were discretized using the same method used by Nazareth and Spaans, 5,376 qubits would be required. In early 2019, the development of a 5,000 qubit commercial quantum annealing computer was announced,³⁵ which would better allow QA to be scaled to higher variable optimization problems but practical clinical application remains a subject for future studies. We believe QTA would be an exciting candidate for quantum computing because it has already shown promise over SA when run on a classical computer.

3.5 Conclusions

In this study we have explored the behavior of a novel algorithm inspired by quantum tunneling, QTA, for use in IMRT beamlet-weight optimization on two SBRT liver cases. We compared QTA's performance with classical SA, an algorithm which has historically been used for this application. On the easier case, QTA exhibited greater stability than SA. On the challenging case,

when allowed to run for the fully allotted number of iterations, both algorithms performed well and exhibited stability with respect to plan quality. With regards to the differences observed between Case 1 and Case 2, it is worth noting that the primary benefit of QTA in a more constrained solution space is the speedup at which it converges, while in a larger (i.e., less constrained) solution space, QTA appears to achieve both faster convergence and plans of more robust quality. These results suggest that QTA may have promise in application to more complex objective functions, such as those defined by biological predictive models. Extension to DAO is demonstrated to be feasible with similar performance, suggesting potential application of QTA for VMAT type applications as well.

3.6 References

1. Bortfeld T. IMRT: a review and preview. *Phys Med Biol.* 2006;51(13):R363-379.
2. Citrin DE. Recent Developments in Radiotherapy. *N Engl J Med.* 2017;377(22):2200-2201.
3. Cho B. Intensity-modulated radiation therapy: a review with a physics perspective. *Radiat Oncol J.* 2018;36(1):1-10.
4. Bortfeld T, Schmidt-Ullrich R, Neve W, Wazer DE, SpringerLink. Image-Guided IMRT. In: Berlin, Heidelberg: Springer Berlin Heidelberg,; 2006.
5. Zhang Y, Merritt M. Dose-volume-based IMRT fluence optimization: A fast least-squares approach with differentiability. *Linear Algebra and its Applications.* 2008;428(5):1365-1387.
6. Lim-Reinders S, Keller BM, Al-Ward S, Sahgal A, Kim A. Online Adaptive Radiation Therapy. *Int J Radiat Oncol Biol Phys.* 2017;99(4):994-1003.
7. Preskill J. Quantum computing and the entanglement frontier. *arXiv preprint arXiv:12035813.* 2012.
8. Nielsen MA, Chuang IL. Quantum Computation and Quantum Information: 10th Anniversary Edition. In: Cambridge: Cambridge University Press; 2010: <http://csis.pace.edu/ctappert/cs837-18spring/QC-textbook.pdf>. Accessed 1/29/2019.
9. Grover LK. From Schrödinger's equation to the quantum search algorithm. *American Journal of Physics.* 2001;69(7):769-777.
10. D-Wave Systems Inc. D-Wave Announces D-Wave 2000Q Quantum Computer and First System Order. <https://www.dwavesys.com/press-releases/d-wave%C2%A0announces%C2%A0d-wave-2000q-quantum-computer-and-first-system-order>. Published 2017, January 24. Accessed January 2019.
11. IBM. Quantum devices and simulators. <https://www.research.ibm.com/ibmq/technology/devices/>. Accessed January 2019.
12. IonQ Inc. A true quantum leap. <https://ionq.co/>. Published 2018. Accessed August 2018.

13. Intel. Reinventing data processing with quantum computing. <https://www.intel.com/content/www/us/en/research/quantum-computing.html>. Accessed January 2019.
14. rigetti. QPU Specifications. <https://rigetti.com/qpu>. Published 2019. Accessed January 2019.
15. Google AI. A Preview of Bristlecone, Google's New Quantum Processor. <https://ai.googleblog.com/2018/03/a-preview-of-bristlecone-googles-new.html>. Published 2018, March 5. Accessed January 2019.
16. Apolloni B, Carvalho C, de Falco D. Quantum stochastic optimization. *Stochastic Processes and their Applications*. 1989;33(2):233-244.
17. Farhi E, Goldstone J, Gutmann S, Lapan J, Lundgren A, Preda D. A quantum adiabatic evolution algorithm applied to random instances of an NP-complete problem. *Science*. 2001;292(5516):472-475.
18. Mukherjee S, Chakrabarti BK. Multivariable optimization: Quantum annealing and computation. *The European Physical Journal Special Topics*. 2015;224(1):17-24.
19. Nazareth DP, Spaans JD. First application of quantum annealing to IMRT beamlet intensity optimization. *Phys Med Biol*. 2015;60(10):4137-4148.
20. Santoro GE, Martoňák R, Tosatti E, Car R. Theory of quantum annealing of an Ising spin glass. *Science*. 2002;295(5564):2427-2430.
21. Morita S, Nishimori H. Mathematical foundation of quantum annealing. *Journal of Mathematical Physics*. 2008;49(12):125210.
22. Kadowaki T, Nishimori H. Quantum annealing in the transverse Ising model. *Physical Review E*. 1998;58(5):5355.
23. Zettili N. *Quantum Mechanics : Concepts and Applications*. New York, UNITED KINGDOM: John Wiley & Sons, Incorporated; 2009.
24. Kuech TF. Metal-organic vapor phase epitaxy of compound semiconductors. *Materials Science Reports*. 1987;2(1):1-49.
25. Leys MR, Veenfliet H. A study of the growth mechanism of epitaxial GaAs as grown by the technique of metal organic vapour phase epitaxy. *Journal of Crystal Growth*. 55(1):145-153.
26. Webb S. Optimization of Conformal Radiotherapy Dose Distributions by Simulated Annealing. *Physics in Medicine and Biology*. 1989;34(10):1349-1370.
27. Bertsimas D, Tsitsiklis J. Simulated Annealing. *Statistical Science*. 1993;8(1):10-15.
28. Hajek B. Cooling Schedules for Optimal Annealing. *Mathematics of Operations Research*. 1988;13(2):311-329.
29. Kessler ML, McShan DL, Epelman MA, et al. Costlets: A Generalized Approach to Cost Functions for Automated Optimization of IMRT Treatment Plans. *Optimization and Engineering*. 2005;6(4):421-448.
30. Matuszak MM, Larsen EW, Fraass BA. Reduction of IMRT beam complexity through the use of beam modulation penalties in the objective function. *Medical physics*. 2007;34(2):507-520.
31. MacFarlane M, Hoover DA, Wong E, Goldman P, Battista JJ, Chen JZ. A fast inverse direct aperture optimization algorithm for intensity-modulated radiation therapy. *Medical physics*. 2019;46(3):1127-1139.
32. Härdenmark B, Liander A, Reh binder H, Löf J, Robinson D. P3IMRT: Direct machine parameter optimization. *Pinnacle White Paper*. 2004(4535):983.

33. Luo Y, McShan DL, Matuszak MM, et al. A multiobjective Bayesian networks approach for joint prediction of tumor local control and radiation pneumonitis in nonsmall-cell lung cancer (NSCLC) for response-adapted radiotherapy. *Medical physics*. 2018;45(8):3980-3995.
34. Tseng H-H, Luo Y, Cui S, Chien J-T, Ten Haken RK, Naqa IE. Deep reinforcement learning for automated radiation adaptation in lung cancer. *Medical physics*. 2017;44(12):6690-6705.
35. D-Wave Systems Inc. D-Wave Previews Next-Generation Quantum Computing Platform. <https://www.dwavesys.com/press-releases/d-wave-previews-next-generation-quantum-computing-platform>. Published 2019, February. Accessed January 2020.

Chapter 4 Dynamic Stochastic Deep Learning Approaches for Predicting Geometric Changes in Head and Neck Cancer

This chapter discusses the design and validation of two novel frameworks employing quantum mechanics and Markov models, in combination with deep recurrent neural networks, for predicting patient geometric changes across fractionated radiotherapy. It is based on the paper: Pakela, J. M., Matuszak, M. M., Ten Haken, R. K., McShane, D. L., El Naqa, I. (expected early 2021) Dynamic Stochastic Deep Learning Approaches for Predicting Geometric Changes in Head and Neck Cancer. (under processing).

4.1 Introduction

4.1.1 Adaptive Radiotherapy

Radiation therapy (RT) has been widely established as the standard of care for numerous cancer types, with as many as 50% of cancer patients who receive RT as part of their treatment each year.¹ Innovations in treatment delivery, such as the use of multi-leaf collimators during intensity modulated radiation therapy (IMRT) and volumetric arc therapy (VMAT) to dynamically move and shape the beam during treatment, have made it possible to achieve dose distributions which conform tightly around the tumor with steep dose gradients, allowing for the delivery of higher, more curative doses to treat the disease while minimizing the dose delivered to the surrounding healthy tissues.^{2,3}

Despite these advancements, a challenge remains in that over the course of fractionated treatment, factors such as weight loss, tumor shrinkage, or daily anatomical variations can cause geometrical changes in the patient's anatomy such that the original treatment plan is no longer optimal—either due to inadequate target coverage or due to increased exposure to organs at risk (OARs). Clinics often track anatomical changes through daily cone-beam computed tomography (CBCT) images taken prior to each treatment; if a patient's plan is deemed to stray too far from the original treatment objectives, alterations can be made, which at their most extreme involve a full replan. The process of using imaging data to make adaptations to a patient's plan during treatment is called *image-guided adaptive radiotherapy* (IG-ART).

While IG-ART has been shown to improve the accuracy and quality of patient treatment plans,^{4,5} several challenges continue to inhibit wider implementation:

1. Clinical time and resource constraints necessitate that IG-ART methods be efficient.⁶
2. In addition, because CBCT images suffer from higher noise and a limited field of view,⁷ daily anatomical changes can be difficult to quantify with accuracy, necessitating IG-ART plans to be robust to measurement uncertainties.
3. Finally, while it has been shown that IG-ART can improve treatment accuracy, there exists no standardized set of criteria for how much a patient's plan needs to deviate from its intended objectives in order to trigger the use of adaptive radiotherapy (ART) methods.⁸

A consequence of these challenges and their time-consuming processes is that some patients do not undergo necessary plan adaptations who might otherwise benefit from them to optimize their treatment outcomes. One approach which addresses the first of these obstacles is to develop a framework which, based on relevant patient characteristics, predicts the likelihood of patient

anatomical changes occurring in future fractions. Such a framework would allow for a proactive approach to ART in contrast to passive image deformation approaches, effectively buying time for busy clinics by allowing them to anticipate and prepare for plan adaptations before they actually occur.

4.1.2 Predictive Models for ART

The development of predictive models to aid in decision making for ART is an active research topic. Predictive models typically utilize imaging data (image-guided models) and/or clinical, dosimetric, biological, and radiomic data (knowledge-based models).⁹ Given the prevalence of daily imaging for patient setup, models based on image guidance are arguably the easiest to implement into a clinical workflow.

There is currently no set consensus regarding what types of predictions are most valuable or necessary for an ART paradigm. Recent studies have reported on models for predicting tumor control probability (TCP) / normal tissue complication probability (NTCP),^{10,11} dosimetric deviations from the original treatment plan,¹² as well as incidence of radiation-induced toxicity¹³ and tumor failure.¹⁴ This variety of approaches is due in part because the objectives for implementing ART may differ across cancer type and/or treatment clinic and can vary from improving dosimetric accuracy, reducing tumor failure, or preventing specific radiation-induced toxicities such as radiation pneumonitis or xerostomia. Adding to this challenge is that there is still an incomplete understanding between the relationship of anatomical/dosimetric deviations during treatment and clinically meaningful outcomes.¹⁵

Recent studies have also reported on the use of geometric anatomical changes (measured via computed tomography [CT] imaging) as a condition for triggering plan adaptation and linked anatomical changes to significant, undesirable dosimetric outcomes.^{16,17} However, to date there

has been limited work towards trying to predict geometric changes during radiotherapy, and the only known existing study is limited both in sample size (N=23) and the ability to predict beyond the next day's fraction.¹⁸

4.1.3 Quantum Predictive Models and Quantum Processes

Recent work has found that quantum mechanics can be applied successfully as a predictive framework for systems outside of fundamental physics—namely human cognition and decision making.¹⁹ In the next two sections we provide a brief review of the quantum predictive framework as well as the Markov predictive framework.

At the turn of the 20th century, scientists began to observe phenomena in nature which violated the laws of classical physics—notably the behavior of light and thermodynamic systems.^{20,21} Quantum mechanics is the mathematical framework which was developed as a means to describe and understand this strange behavior.

The first three postulates of quantum mechanics can be formulated as follows²²:

1. A quantum system can be fully described by its unitary state vector, which can be written as a weighted sum of orthonormal vectors:

$$|\Psi\rangle = \sum_i a_i |v_i\rangle \quad 4.1$$

Where the set of vectors $|v_{i \in N}\rangle$ form an orthonormal basis in a finite N-dimensional Hilbert space. The notation " $| \rangle$ " is called a *ket* and represents a column vector in a complex vector space. Conversely, " $\langle |$ " is called a *bra* and represents a row vector, such that $|v\rangle$ is the conjugate transpose of $\langle v|$ (i.e., $\langle v|^\dagger = |v\rangle$). The coefficients, a_i in Equation 4.1 are called probability amplitudes because the probability of a quantum system, $|\Psi\rangle$, being observed in state $|v_i\rangle$ is given by $|a_i|^2$. In order to obey the law of total probability:

$$\sum_i |a_i|^2 = 1 \quad 4.2$$

2. The time evolution of a stationary quantum state is described by the time-independent Schrödinger equation:

$$\frac{\partial |\Psi(t)\rangle}{\partial t} = -i\hbar |\Psi(t)\rangle \quad 4.3$$

Whose solution takes the form:

$$|\Psi(t+n)\rangle = U(t+n)|\Psi(t)\rangle \quad 4.4$$

$$U(t+n) = e^{-i(t+n)H} \quad 4.5$$

Where $|\Psi(t)\rangle$ is the state at time, t , and $U(t+n)$ is a unitary transition matrix defined by a Hermitian operator, H , called the Hamiltonian. In physics, the Hamiltonian represents the total energy of the system. (Note that in Equation 4.5, Plank's constant (denoted by \hbar) has been absorbed into H .)

3. A measurement on a quantum system is conducted by applying a measurement operator, M_m , to the system state vector. The index, m , spans each of the possible outcomes of the measurement such that:

$$\sum_m M_m^\dagger M_m = I \quad 4.6$$

Where I is the identity matrix. The probability of obtaining a measurement outcome m on a quantum system described by $|\Psi\rangle$ is given by:

$$P(m) = |M_m|\Psi\rangle|^2 = \langle\Psi|M_m^\dagger M_m|\Psi\rangle \quad 4.7$$

And the state of the quantum system after applying measurement M_m is given by:

$$\frac{M_m |\Psi\rangle}{\sqrt{\langle \Psi | M_m^\dagger M_m | \Psi \rangle}} \quad 4.8$$

4.1.4 Classical Markov Processes

A *Markov process* refers to a classical stochastic process which has no memory of its past and is influenced only by its current state.²³ Given a finite number, N , of potential states, the state probability distribution can be written as an $N \times 1$ column vector, $M(t)$. Each entry of $M(t)$ represents the probability of finding the system in a given state at time t , such that:

$$\sum_i M_i(t) = 1 \quad 4.9$$

Similar to a quantum process, a Markov process evolves over time according to a transition matrix, $P(t)$, such that:

$$M(t) = P(t)M(0) \quad 4.10$$

The time evolution of this state transition matrix is described by the Kolmogorov forward and backward equations:

$$\frac{d}{dt}P(t) = P(t)Q \quad 4.11$$

$$\frac{d}{dt}P(t) = QP(t) \quad 4.12$$

An expression for $P(t)$ can therefore be written as:

$$P(t) = e^{tQ} \quad 4.13$$

Where Q is referred to as the generator matrix, which satisfies the conditions:

$$Q_{ij} \geq 0 \quad \forall i \neq j \quad 4.14$$

$$\sum_j Q_{ij} = 0 \quad 4.15$$

The results from Equations 4.10 and 4.13 can be used to describe the time evolution of the Markov system:

$$M(t + n) = e^{Q(t+n)}M(t) \quad 4.16$$

4.1.5 Deep Learning and Recursive Neural Networks

Deep learning refers to a broad class of machine learning methods which are capable of learning latent/higher order characteristics of a dataset in order to achieve a user-defined goal (typically the minimization of an objective function). While there are no structural requirements for deep learning algorithms, the most widely successful implementations of deep learning have been realized through artificial neural networks (ANNs). Inspired by and named after biological neural networks, ANNs consist of layers of nodes, with each node receiving input through weighted connections to other nodes and producing an output determined by a non-linear activation function. During training, a neural network learns to optimize the weights of the connections between nodes through a process called backpropagation, in which parameters are iteratively updated by calculating the gradient of the loss function with respect to each weight in the network.²⁴ ANN-based deep learning algorithms have enjoyed widespread success across a diverse range of applications in medicine and hold the potential to revolutionize the field of radiation oncology.²⁵

Recursive neural networks (RNNs) are neural networks which can handle sequence data.²⁶ RNNs differ from traditional ANNs in that all of the nodes in an RNN share the same weights and biases. In addition, because they take an entire sequence as input, the width of an RNN (number

of nodes per layer) is usually set as the number of events (or timesteps) per sequence. Each node in an RNN takes two categories of information as input: information about the current state of the sequence (either from the input data or from the output of the node in the previous layer) and information about the previous states, represented by a hidden state, h_{t-1} . The flow of information in an RNN can thus be thought of as traveling both *across* the sequence and *upwards* through the layers.

4.1.6 Study Objectives

In this study, we present a novel method of predicting patient changes over the course of treatment which models the patient as a stationary quantum state whose timepoints are discretized by fraction number. The motivation for modeling the dynamics of patient anatomy as a quantum system is informed by the nature of quantum systems. Specifically, it is impossible for human observers to predict the behavior of a quantum system with absolute certainty. Similarly, in a radiation oncology paradigm, the aforementioned limitations in CBCT image quality in combination with daily setup variations, physiological changes such as tumor response, and anatomical motion, mean that there are inherent uncertainties in knowing a patient's "state" during the course of treatment. The non-deterministic nature of these two systems lends to the notion that modeling patient changes after the evolution of a quantum system may lead to more accurate predictions than those which could be achieved using a classical predictive framework.

To test this hypothesis, we designed a predictive framework which models patient changes as a stationary quantum system defined by the time-independent Schrodinger equation evolving according to a patient-specific Hamiltonian. The parameters of the Hamiltonian are derived from an RNN, which takes data from the first several fractions of a patient as input in order to predict later fractions. In addition, we also developed a Markov-based predictive framework which models

patient changes according to the Kolmogorov backward equation and a patient-specific generator matrix. We trained and evaluated both frameworks on a head and neck cancer (HNC) radiotherapy dataset, which was selected because HNCs represent a patient population which commonly experiences geometric changes leading to plan adaptations.⁷

4.2 Methods

4.2.1 Patient Population and Data Acquisition

Training data for this study were acquired retrospectively from a dataset of 128 HNC patients who received VMAT at one institution between the years of 2013 and 2016. Of the original 128, three patients were excluded due to abnormalities which occurred during their treatment sequence. The remaining 125 patients consisted of 102 (23) males (females) with an average age of 58.3 (range = [10, 84]). (Note that this gender imbalance is typical for HNC populations.²⁷) Patients received a target prescription dose between 60-70 Gy delivered over 30-35 fractions, with as many as 17 patients undergoing replanning. Oropharyngeal cancers which are positive for human papilloma virus (HPV) have been shown to have higher radiosensitivity.^{28,29} Of the patients whose tumor site was located in the oropharynx (n=68), 52 were recorded as positive for human papilloma virus (HPV), 5 were negative, and 1 was unknown. Patient immobilization of the upper body was achieved using 5 point masks and daily setup errors were corrected using CBCT image guidance.

An additional 20 HNC patients treated with 30-35 fractions of VMAT radiotherapy between 2018-2020 were selected to serve as an external testing dataset. The patients consisted of 14 (7) males (females) with an average age of 60 (range = [22, 86]). Deformable image registration was performed on daily CBCT images for two fractions per treatment week, and volume/table shifts were calculated from these fractions in the same manner as the training dataset.

4.2.2 Data Preprocessing

The data used to define the patient state at each treatment fraction consisted of fractional primary clinical target volumes (CTVs) (with respect to the volume at the first fraction) as well as daily table shifts (lateral, longitudinal, vertical) used for patient alignment (**Figure 4.1**). Patient rotation angles (defined as the yaw angle of the treatment couch) were initially recorded but were ultimately omitted from the model because the majority of patients did not undergo rotational shifts during treatment. Pitch and roll angles were also omitted because all of the patients in the training dataset were treated on 4D couches (which perform translations and rotations around the yaw axis).

Deformable image registration (DIR) was performed by an experienced dosimetrist between daily CBCT images, and the planning CT as described previously.¹³ The majority of patients utilized CBCTs from every treatment fraction, while the rest included at least two fractions per treatment week. CTV volumes were then obtained from the daily CBCT images and normalized with respect to the volume of the first fraction in their respective treatment course. The motivation for reporting fractional (or relative) volumes as opposed to raw volumes was in part due to the wide range of volume sizes across patients [2.5cc - 456cc] and the need to relate meaningful tumor volume changes between patients (i.e., depending on tumor size, a 2cc reduction in tumor volume could represent an 80% reduction in tumor volume in one patient and only a 0.5% reduction in another). Daily table shifts and rotations were calculated as the difference between table coordinates recorded for the daily CBCT image and the DICOM treatment record (RTRECORD).

Fractions within a treatment sequence were also omitted from the model if they did not meet the inclusion criteria described below. The recorded CTV volumes and table shifts occasionally contained anomalies which did not represent meaningful changes across the treatment

fractions. For example, some CTV volumes were measured which appeared as large spikes in the overall fraction sequence. Upon inspection of the original CBCT structure sets in Aria (Varian), it was found that these spikes in volume could be attributed to an error in the deformable image registration. To resolve this issue, fractions were discarded whose volume differed from both the fraction immediately prior and immediately preceding by > 10 cc. Abnormal table shifts were defined as a shift which exceeded 15 mm in any direction. These instances were attributed to large setup deviations from the original planning CT. Fractions with shifts exceeding this cut-off were excluded when they occurred sparsely throughout a treatment regimen. In two patients, CBCT shifts were recorded which exceeded 40 mm for every fraction. Because these represented a consistent setup variation from the planning CT, the shifts in these fractions were not discarded and were instead normalized by subtracting off the shift in the first fraction to more faithfully represent position modifications that were not due to setup error. This ultimately resulted in the exclusion of 99 fractions across 35 patients for volume spikes and 8 fractions across 7 patients for table shift abnormalities (out of 3657 total fractions). Data preprocessing was performed in the same manner for the external testing dataset, with a total of 7 fractions from this dataset discarded (from a total of 678 fractions) for failing to meet the inclusion criteria defined above.

After cleaning, the dataset underwent data augmentation, in which missing fractions were filled in by their nearest neighbor. Even fractions were discarded such that the final dataset used for training and validation had a sequence length of 18 and spanned fractions 1-35.

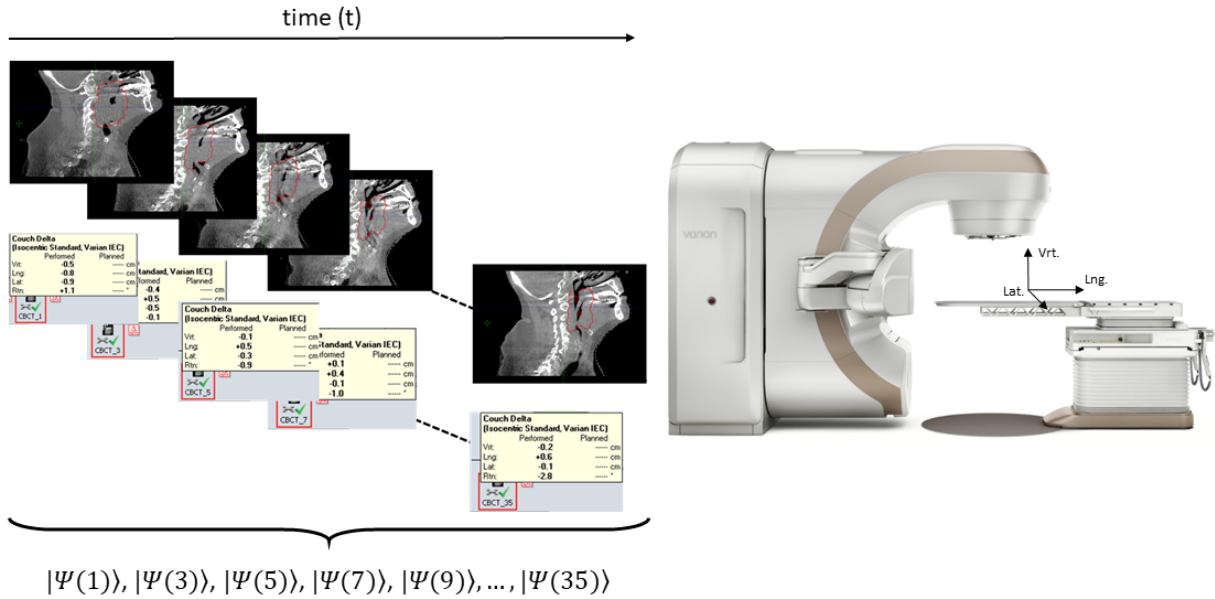


Figure 4.1: Schematic of data collected from each patient across treatment fractions. Primary CTV volumes were acquired from daily CBCTs and daily table shifts (described by the coordinate space shown in the treatment system) were collected for each available fraction and used to define an orthonormal state vector. Linear accelerator image adapted from Varian official website, https://varian.widen.net/view/pdf/hujb29bryb/TrueBeam2.7_ProductBrief_RAD10437A_June2020.pdf?u=wefire

4.2.3 Mapping to Discrete Orthonormal State Vectors

In order to model patient anatomical changes as a quantum system, the patient characteristics (volume + table shift values) at each point in time (fraction #) were mapped to an orthonormal state vector, Ψ_i , where $\Psi_{i \in [1, N]}$ form a basis in an N -dimensional Hilbert space, and N is the number of potential states the patient can be in at a given time. It is desirable to select a value for N so as to balance the need to ensure that the Hilbert space contains enough dimensions to meaningfully capture the complexity of the patient population but is not so large as to result in large numbers of sparsely used discrete states. The dataset was encoded using several different values of N to assess the predictive models' performance under this tradeoff.

The primary CTV volumes and daily table shifts were converted to orthonormal state vectors by first mapping them to discrete state values using vector quantization.³⁰ A generalization of Lloyd-Max scalar quantization, vector quantization is a data reduction technique which finds an optimal codebook of vectors such that for a given dataset, the total difference between each multivariate sample and its nearest codebook neighbor is minimal. Once the codebook has been calculated, each sample vector mapped to its nearest codebook neighbor, and the dimensionality of the data is effectively reduced to the number of vectors (or states) contained within the codebook. Vector quantization was performed via the Vector Quantization Design Tool (vqdtol) available through MATLAB's DSP system toolbox (Mathworks, Inc.), after which each state was encoded as a one-hot vector. **Figure 4.2** displays an example of the distribution of lateral shifts recorded across all fractions from the original dataset (a) and after encoding to 4 states (b) and 10 states (c) as well as a subset of the lateral shifts recorded from the original dataset (orange) and after encoding (blue) (e & f). It is clear that the 10-state encoding, while computationally more expensive, provides a closer approximation to the original data. The external testing dataset was quantized using the same codebook as the training data. Note that the 10-state encoding only includes 6 discrete states for the lateral shifts (**Figure 4.2(c)**). This is because the vector quantizer looks for the optimal representation of the distribution of vectors within a multidimensional dataset, and a single parameter may thus appear multiple times within a codebook of unique vectors.

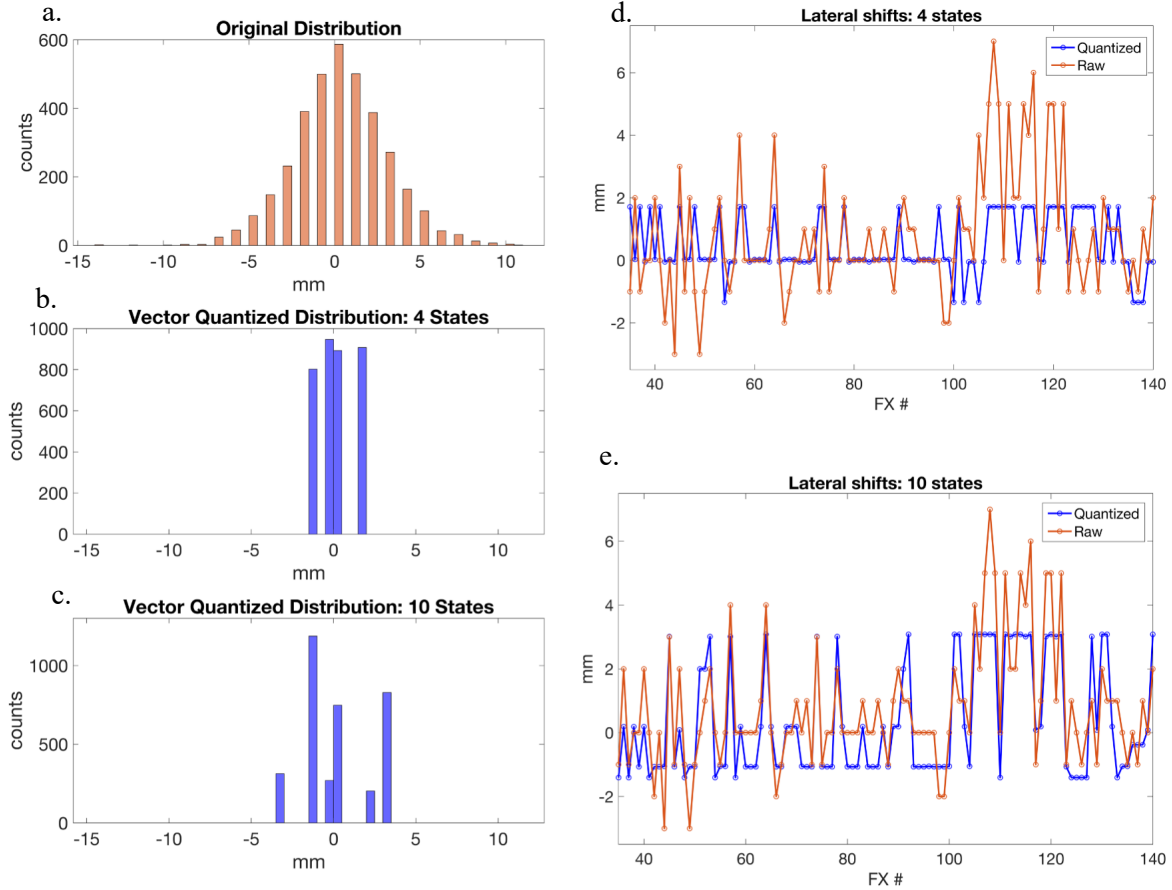


Figure 4.2: Distribution of lateral shifts recorded across all fractions from the original dataset and after encoding to 4 and 10 states, as well as a subset of the lateral shifts recorded from the raw, unquantized dataset (orange) and after encoding (blue).

4.2.4 Generation of Additional Synthetic Data via Generative Adversarial Network (GAN)

One challenge for this study was the limited number of patients available to train the proposed models. To address this challenge, a Time-series generative adversarial network (TimeGAN) was employed to generate additional synthetic data.³¹ Generative adversarial networks describe a class of deep learning algorithms which learn the latent distributions of a dataset by pitting two competing networks against each other. The generator network seeks to generate realistic “fake” data samples to fool a discriminator network, which seeks to distinguish real samples from fakes. Training is successful when the generator has reached the point where it can produce fake samples

which a well-trained discriminator identifies correctly only 50% of the time. TimeGAN learns the distributions of features across time by utilizing RNNs for both the generator and the discriminator as well as introducing a supervised loss term in addition to the traditional unsupervised adversarial loss. Finally, TimeGAN also utilizes an additional embedding network which learns a reversible mapping between latent space and features. Our implementation of TimeGAN followed the same architecture and hyperparameters as the original publication.³¹ Synthetic data was generated in the form of discrete state values produced from vector quantization (i.e., after discrete state mapping but prior to one-hot encoding). In order to prevent model leakage, TimeGAN was trained on five separate training folds used to train outer loop of the predictive models during nested cross validation. This resulted in an additional 200 samples per fold, for a total of 325 patients for model evaluation.

In order to assess the quality of the synthetic data, histograms counting all instances of each discrete state across all fractions and patients were calculated, and t-distributed stochastic neighbor embedding (t-SNE) as well as principal component analysis (PCA) analyses were performed (flattened in the time dimension) to visually compare the original and synthetic data. In addition, two quantitative performance scores, a discriminative score and a predictive score, were obtained by training two separate neural networks. The discriminative score was calculated by training a discriminator (in the form of an RNN) to classify real vs. synthetic samples, and then reporting the classification error from a holdout testing dataset. The predictive score was calculated by training an RNN network using the synthetic dataset to predict the state in the next time step, and then reporting the mean absolute error of the model when tested on the original data. **Figure 4.3** displays the results of the data visualizations (a, b, d, & e) as well as the testing scores for the discriminator and predictive networks (c), which are compared against the results of the model trained on another

sequence-based, medical dataset consisting of a private lung-cancer pathways dataset as reported in Yoon et al, 2019.³¹

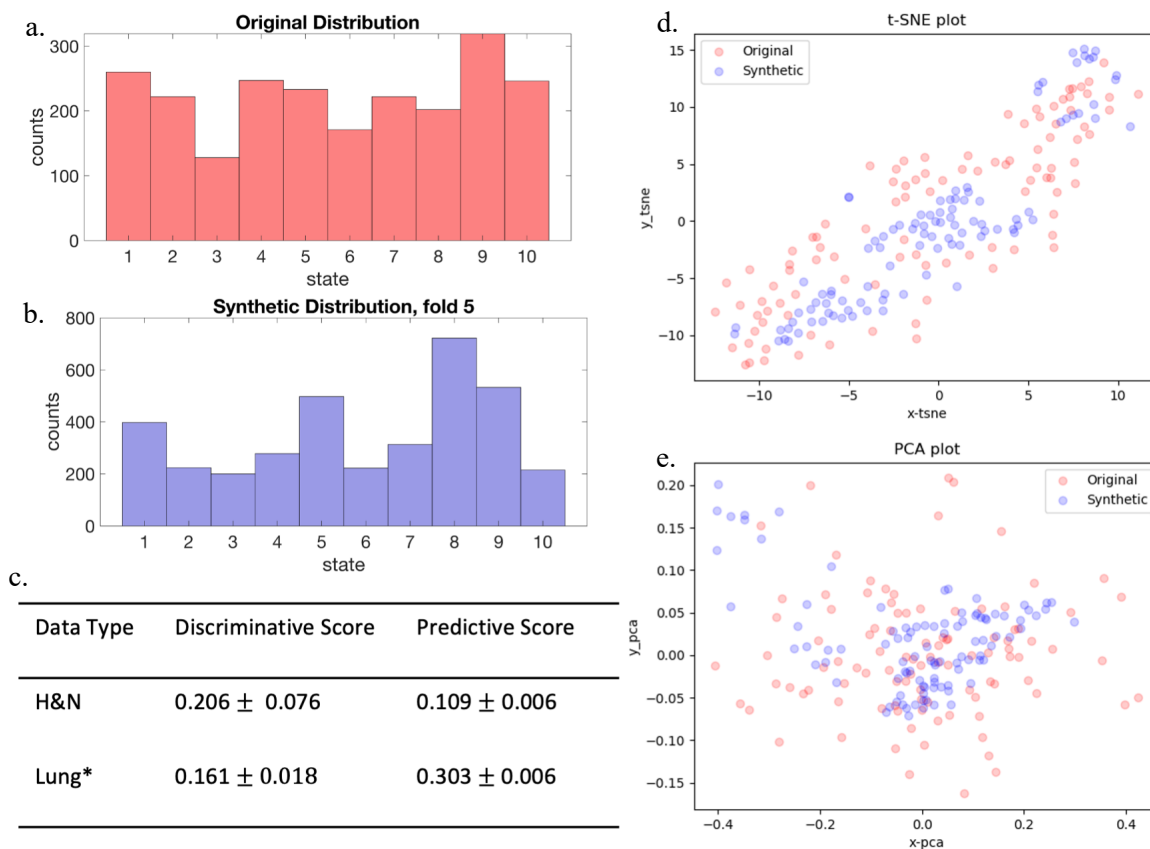


Figure 4.3: Assessment of synthetic data generated by TimeGAN for the 10-state vector encoding trained using original data from fold 5: (a) shows the distribution of the original 10-state encoded data, (b) shows the distribution of the synthetic, (c) displays discriminative and predictive scores for the HN synthetic data (H&N) and the scores from the lung cancer pathways data (Lung*) reported in the original TimeGAN paper, (d) and (e) display t-SNE and PCA plots respectively for the original 10-state data and the synthetic 10-state data.

4.2.5 Quantum-based Prediction Algorithm

After mapping patient attributes to state vectors, the remaining key piece of information necessary to predict future patient states is to acquire an expression for the Hamiltonian, which determines the time evolution of the quantum system according to Equations 4.4 and 4.5. The Hamiltonian

was derived from the output of a gated recurrent unit (GRU) neural network,³² which we refer to as the *QRNN*. In order to preserve the unity length of predicted states as well as ensure that the squared probability amplitudes sum to 1, the predicted Hamiltonian must be Hermitian (i.e., equal to its conjugate transpose). However, because PyTorch is currently limited in its use of complex numbers within ANNs, we assumed a real-valued Hamiltonian and therefore needed only to build a symmetric matrix from the GRU output. State vectors from the several fractions are used as input to the GRU, and the network outputs parameters necessary to construct the Hamiltonian in the form of vector of length $(N + 1) \times N/2$, where N is the total possible number of states. This parameter vector was then used to construct a symmetric (and by default Hermitian) matrix according to Section 1.2.7 in Golub and Van Loan, 2013.³³ In order to calculate the unitary transition matrix defined by Equation 4.5, the expression for the transition matrix was reformulated using Euler's formula and spectral decomposition was used to perform the necessary operations on the Hamiltonian matrix:

$$e^{-iHt} = \cos(Ht) - i\sin(Ht) \quad 4.17$$

$$f(H) = v^T f(\lambda) v \quad 4.18$$

Where v and λ are the respective eigenvectors and eigenvalues of H . The transition matrix was split into its real and complex components such that:

$$\begin{cases} Re(U) = v^T \cos(\lambda t) v \\ Im(U) = v^T \sin(\lambda t) v \end{cases} \quad 4.19$$

Where t is expressed as an integer and determines which future fraction to be predicted. (Note that we can parameterize t , as $t = 2x$ such that $|\Psi(t_0 + t)\rangle = |\Psi(t_0 + 2x)\rangle$ to yield odd numbered fractions.) Future fraction predictions, $|\hat{\Psi}(t_0 + t)\rangle$ were acquired by calculating the probability

amplitude of each state after applying the unitary transition matrix to the known system state of fraction t_0 :

$$|\widehat{\Psi}(t_0 + t)\rangle = U(t)|\Psi(t_0)\rangle \quad 4.20$$

$$= v^T \cos(\lambda t) v |\Psi(t_0)\rangle - i v^T \sin(\lambda t) v |\Psi(t_0)\rangle \quad 4.21$$

$$P(|\widehat{\Psi}(t_0 + t)\rangle = |v_{i \in \{1, n\}}\rangle) = a_i^2 \quad 4.22$$

Where a_i^2 is the i^{th} element in the vector given by:

$$a^2 = (|\widehat{\Psi}(t_0 + t)\rangle)^2 \quad 4.23$$

$$= (\text{Re}(|\Psi(t_0 + t)\rangle))^2 + (\text{Im}(|\Psi(t_0 + t)\rangle))^2 \quad 4.24$$

Where “.2” refers to element-wise multiplication. The predicted future state was then defined as the basis vector with the largest probability. The model loss function was then defined as the negative log likelihood of the algorithm predicting the correct future state:

$$QLoss = \log(|\Psi(t)|^2) - \log(|U(t)|\Psi(t_0)|^2) \quad 4.25$$

A schematic of the flow of information through the neural network is shown in **Figure 4.4**.

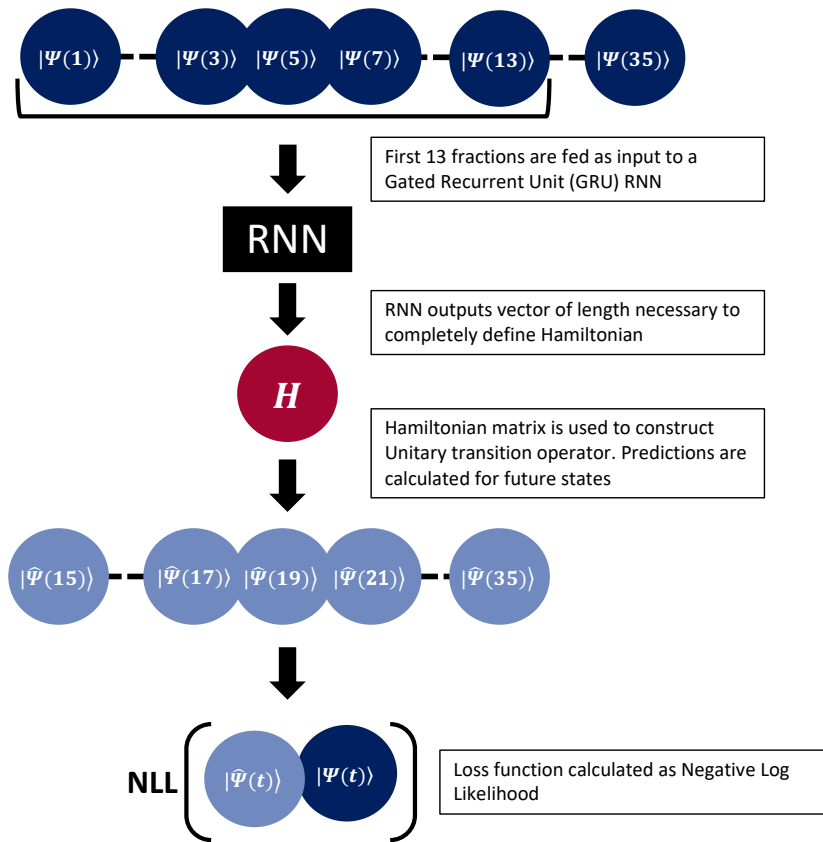


Figure 4.4: Schematic of QRNN Algorithm. A set number of initial states are fed into a recurrent neural network, which produces a vector of parameters necessary to build a Hamiltonian matrix. The Hamiltonian is then used to construct a transition matrix, which in turn is used in conjunction with the oldest initial state to predict the probability of each remaining fraction existing in a given state. The fraction state predictions are then used in conjunction with the ground truth to calculate the model loss.

4.2.6 Markov-based Prediction Algorithm (MRNN)

Similar to the quantum model, the Markov model also utilized a GRU neural network (which we refer to as the *MRNN*) and took the first 13 odd fractions as its input. The output of the MRNN was a vector of length $l = N \times (N - 1)$, where again N is the total number of possible states. A generator matrix, Q , was then built from this parameter vector such that Equations 4.14 and 4.15 were satisfied. The probability transition matrix defined by Equation 4.13 was then calculated.³⁴

The probability distribution of future states was thus calculated as:

$$M(t_0 + t) = e^{Qt}M(t_0) \quad 4.26$$

And the predicted future state was assigned as the state with the maximum probability. The loss function for the MRNN was then defined as the negative log likelihood between the Markov-predicted state and the known future state:

$$MLoss = \log(M(t_0)) - \log(e^{Qt}M(t_0)) \quad 4.27$$

4.2.7 Model Architecture and Training

The models were built and trained with Python using the PyTorch library on a Windows 10 desktop computer (AMD Ryzen 5 3600 6-Core Processor, 3.59 GHz, NVIDIA GeForce GTX 1660 SUPER). The neural networks for both models utilized Adam optimization with a learning rate of 0.001 and a batch size of 30. The QRNN and MRNN models were trained and assessed using a nested cross validation (NCV) scheme to mitigate overfitting effects and assess generalizability.³⁵ NCV involves performing an *inner loop* of cross validation (CV) on each training fold for parameter tuning in order to avoid introducing additional bias into the model assessment. The inner loop consisted of 5-fold CV and was trained for 50 epochs in order to train the number of layers in the GRU as well as the dropout values. Optimal parameters were based on the minimum loss model loss achieved over every combination of parameter values. The outer loop consisted of 5-fold CV and was trained for 100 epochs, which was determined experimentally. Model loss, area under the receiver operating characteristic curve (AUC), and confusion matrices were recorded to assess model performance.

4.3 Results

For each system state representation $N \in \{4, 6, 8, 10\}$, models were trained using the first 13 odd fractions as input to predict the remaining odd fractions (15-35). To assess model performance, AUC scores were calculated for each model using the scikit-learn library's `roc_auc_score()`

function with parameters set such that the metrics for each label were weighted by the number of true instances for each label, and the AUC reported was the average AUC of all possible pairwise combinations of classes.³⁶ AUC scores were computed both across all fractions (15-35) and on a per fraction basis in order to assess both general model performance as well as model performance at specific timepoints.

4.3.1 Overall Results

Table 4.1 displays the validation AUC scores calculated using all future fractions (15-35) averaged over the 5 folds of cross validation (outer fold of nested CV) as well as z-scores and corresponding p-values calculated from the DeLong test for each system state representation.³⁷ From the results in **Table 4.1**, the MRNN model was observed to have a generally higher performance on the validation set than the QRNN model; however, these differences were only statistically significant ($p < 0.05$) for the 8-state and 10-state representations. In addition, both models seemed to experience reduced performance with increasing state-vector size (with the exception of the 8-state MRNN results).

Table 4.1: Validation AUCs averaged over 5-folds of cross validation recorded for MRNN and QRNN models after 100 epochs. (mean \pm 95% confidence interval [CI]), as well as z-scores and corresponding p -values from the DeLong test.

Validation AUC after 100 Epochs				
State-vector size	MRNN	QRNN	Z-score	p -value
4	0.742 \pm 0.021	0.675 \pm 0.036	1.35	0.177
6	0.709 \pm 0.026	0.656 \pm 0.021	1.83	0.068
8	0.724 \pm 0.036	0.652 \pm 0.044	2.53	0.011
10	0.698 \pm 0.016	0.605 \pm 0.035	2.67	0.007

Figure 4.5 displays the AUC scores on the internal validation dataset of the outer loop of NCV validation after 100 epochs of training. The results in **Figure 4.5** suggest that for both the MRNN and QRNN models, the AUC score appears to decrease for increasing system state sizes. In addition, there also appears to be a slight decrease in AUC score for fractions at later timepoints.

One challenge of calculating AUC scores on a per-fraction basis is that this leads to fewer samples used to determine each score. In order to calculate the AUC, it is necessary for sample set to contain instances (true labels) of every possible class. Because of this requirement and the fact that it is not possible to stratify patients across folds such that each fraction has an equal distribution of state values, it is possible for some folds to have undefined AUC values at specific fractions. For such instances, the mean AUC values and CIs were reported based on the remaining folds. This resulted in models sometimes artificially appearing to have smaller measurement uncertainties—for example, the MRNN and QRNN models trained using system state vectors of size 10 had 4 out of 5 folds which contained undefined AUC scores at fraction 25. **Table 4.2** summarizes the number of folds left out for each fraction during the per-fraction AUC calculations.

The higher order system state representations (8 and 10) were found to suffer from a higher number of undefined folds.

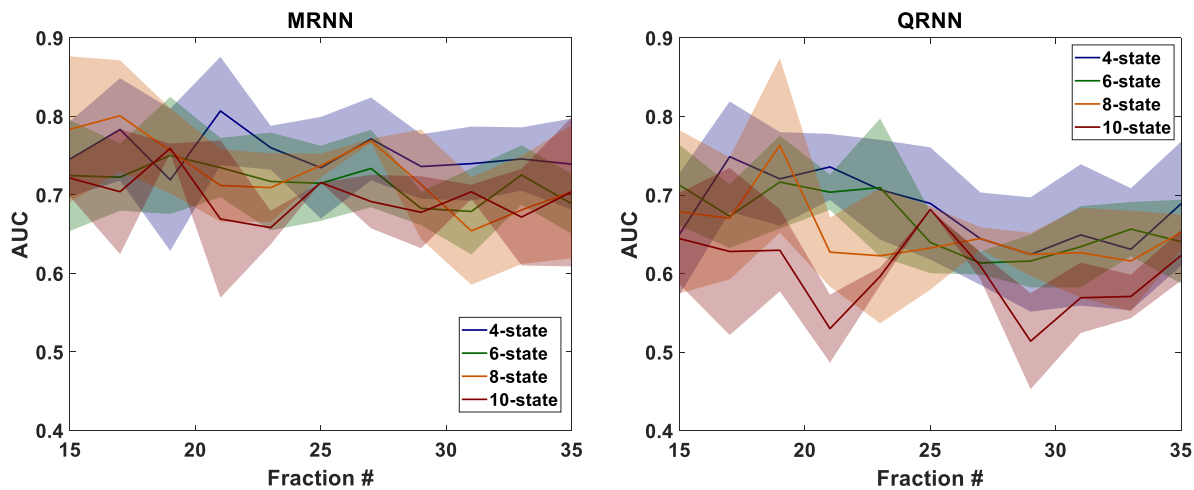


Figure 4.5: MRNN and QRNN validation AUC scores vs fraction number for each system state representation. Shaded region indicates the 95% confidence interval.

Table 4.2: Number of folds omitted from each fraction during AUC score calculation due to lack of samples for each possible class.

Number of folds omitted from AUC calculation				
Fraction #	4-state	6-state	8-state	10-state
15	0	1	2	2
17	0	1	1	1
19	0	1	2	3
21	0	1	3	3
23	0	1	2	3
25	0	0	1	4
27	0	0	2	1
29	0	0	1	2
31	0	0	1	1
33	0	1	0	3
35	0	0	3	2

4.3.2 4-state Predictions

Figure 4.6 displays the results for the 4-state quantum and Markov models, including the mean training and testing AUCs recorded over 100 epochs during the outer loop of 5-fold cross validation (a & b), the corresponding model loss (c & d), as well as confusion matrix metrics: sensitivity, specificity, positive predictive value (PPV), and negative predictive value (NPV) across each of the 4-states (e & f).

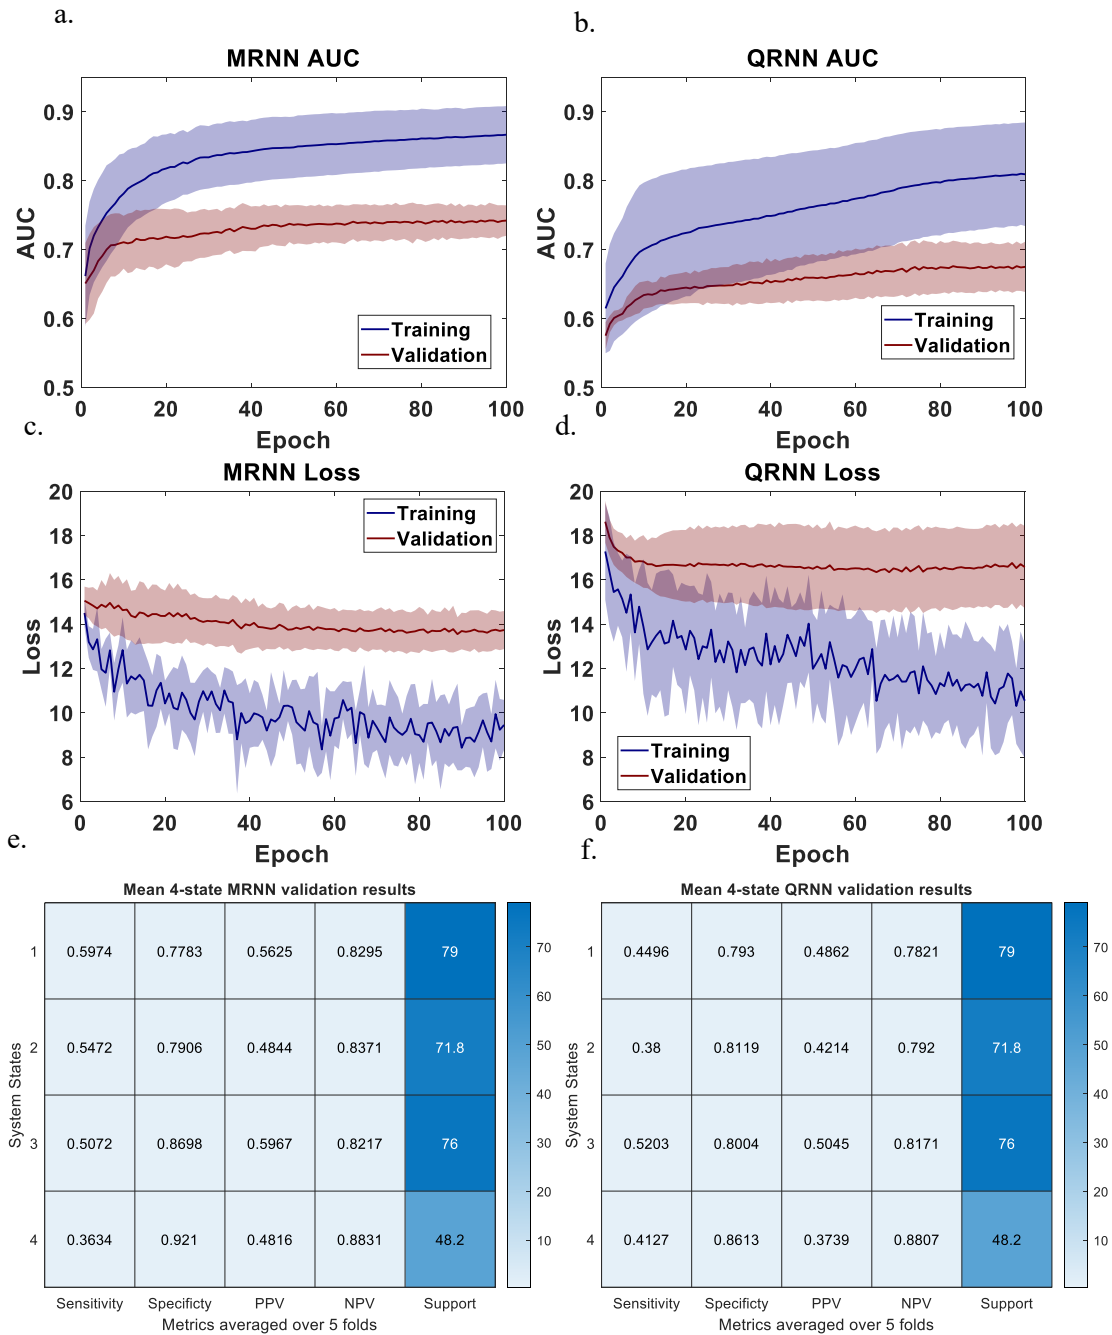


Figure 4.6: QRNN and MRNN model results for 4-state system predicting fractions 15-35: AUC scores (mean and 95% CI) (a) & (b), model loss (c) & (d), and confusion matrix metrics: sensitivity, specificity, positive predictive value (PPV), and negative predictive value (NPV) averaged across 5-folds after 100 epochs of training (e) & (f).

4.3.3 6-state Predictions

Figure 4.7 displays the results for the 6-state quantum and Markov models, including the mean training and testing AUCs recorded over 100 epochs during the outer loop of 5-fold cross validation (a & b), the corresponding model loss (c & d), as well as confusion matrix metrics: sensitivity, specificity, PPV, and NPV across each of the 6-states (e & f).

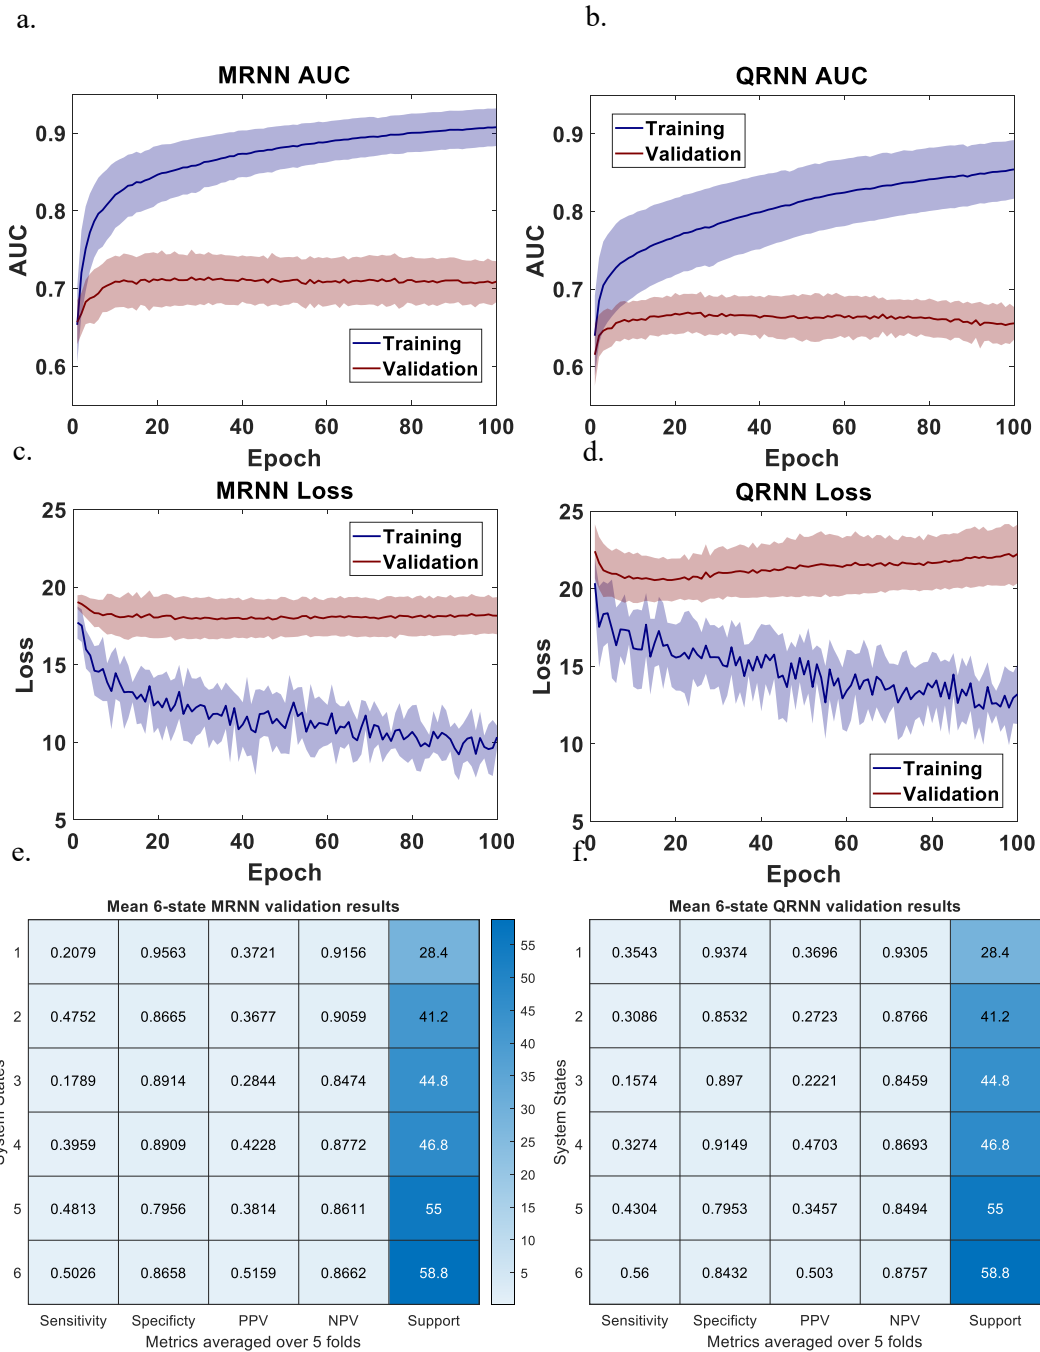


Figure 4.7: QRNN and MRNN model results for 6-state system predicting fractions 15-35: AUC scores (mean and 95% CI) (a) & (b), model loss (c) & (d), and confusion matrix metrics: sensitivity, specificity, positive predictive value (PPV), and negative predictive value (NPV) averaged across 5-folds after 100 epochs of training (e) & (f).

4.3.4 8-state Predictions

Figure 4.8 displays the results for the 8-state quantum and Markov models, including the mean training and testing AUCs recorded over 100 epochs during the outer loop of 5-fold cross validation (a & b), the corresponding model loss (c & d), as well as confusion matrix metrics: sensitivity, specificity, PPV, and NPV across each of the 8-states (e & f).

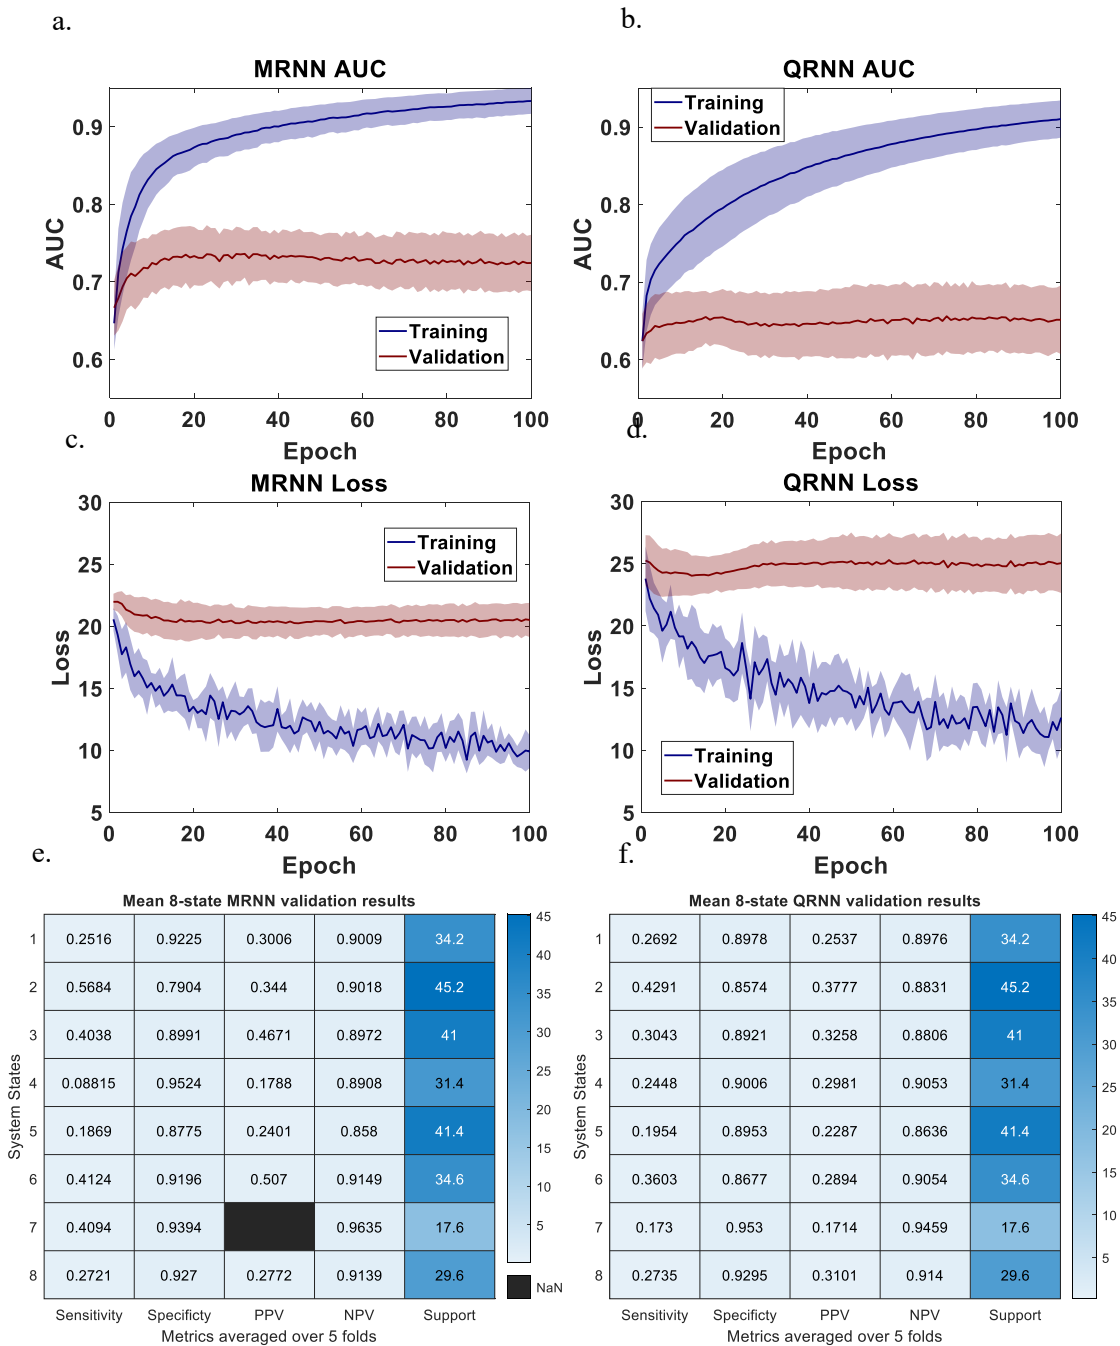


Figure 4.8: QRNN and MRNN model results for 8-state system predicting fractions 15-35: AUC scores (mean and 95% CI) (a) & (b), model loss (c) & (d), and confusion matrix metrics: sensitivity, specificity, positive predictive value (PPV), and negative predictive value (NPV) averaged across 5-folds after 100 epochs of training (e) & (f).

4.3.5 10-state Predictions

Figure 4.9 displays the results for the 10-state quantum and Markov models, including the mean training and testing AUCs recorded over 100 epochs during the outer loop of 5-fold cross validation (a & b), the corresponding model loss (c & d), as well as confusion matrix metrics: sensitivity, specificity, PPV, and NPV across each of the 8-states (e & f).

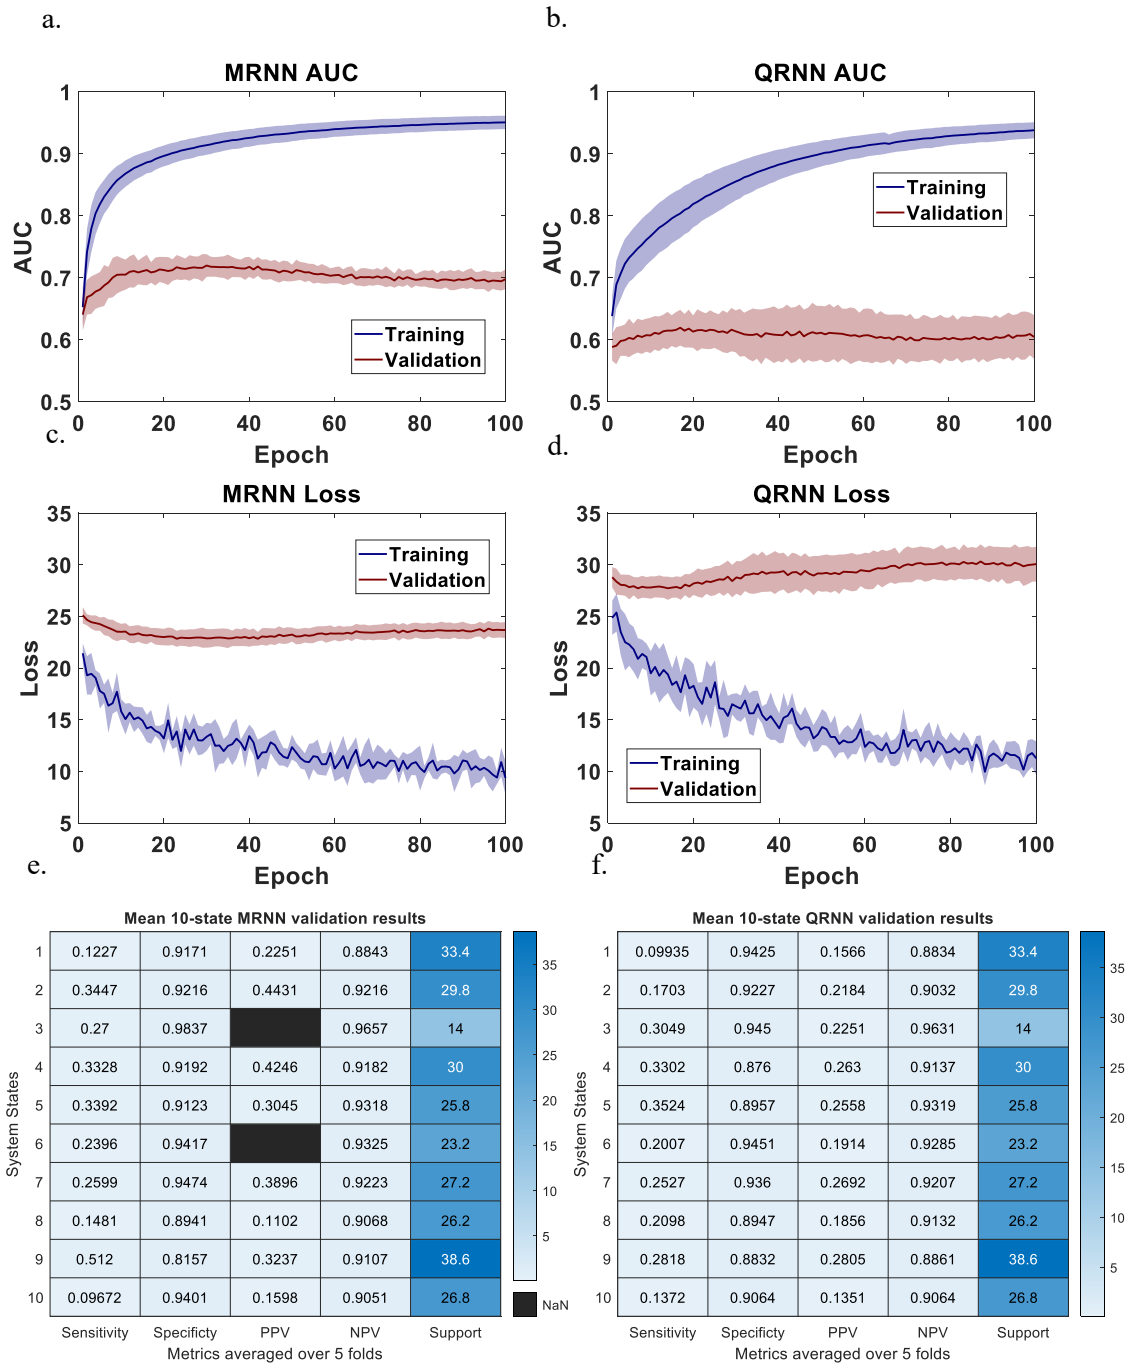


Figure 4.9: QRNN and MRNN model results for 10-state system predicting fractions 15-35: AUC scores (mean and 95% CI) (a) & (b), model loss (c) & (d), and confusion matrix metrics: sensitivity, specificity, positive predictive value (PPV), and negative predictive value (NPV) fractions averaged across 5-folds after 100 epochs of training (e) & (f).

4.3.6 External Testing Results

Final models for each system-state representation were trained using all of the real training data ($N = 125$) as well as 250 synthetic samples (50 samples randomly selected from each of the 5 outer cross validation folds) so that the ratio of real to synthetic training data was the same as used during the cross validation process (1:2). For comparison, an additional set of final models were also trained using only real data. The hyperparameters for each model were selected as those chosen most frequently during nested cross validation. All models had an RNN depth of 2 layers and dropout coefficients, which ranged from 0.1-0.5. Model testing was performed using the external dataset of 20 patients described in Section 4.2.1. **Table 4.3** displays the final testing AUC scores for each model; these are consistent with the training cross-validation results, where the MRNN approach was found to achieve higher AUC scores than the QRNN, though model performance was relatively less sensitive to the order of the system state representation.

Table 4.3: Testing AUC scores calculated as the highest achieved over the course of 100 epochs of training. Results include models trained with real and synthetic data as well as models trained without synthetic data (labeled as “no GAN”).

Testing AUC				
State-vector size	MRNN	QRNN	MRNN (no GAN)	QRNN (no GAN)
4	0.729	0.650	0.707	0.623
6	0.677	0.651	0.687	0.608
8	0.716	0.683	0.723	0.669
10	0.671	0.607	0.697	0.609

4.4 Discussion

In this work we presented two approaches for radiotherapy geometric adaptation based on quantum computing and Markov methods within a deep learning framework. Overall, the Markov predictive model was found to exhibit slightly higher performance (though not always statistically significant) than the quantum predictive model—both across all fractions (**Table 4.1**) and on a per-fraction basis (**Figure 4.5**). One possible explanation for this outcome is the difference in the number of parameters necessary to construct the quantum-based Hamiltonian and the Markov generator matrix. The Markov generator matrix is required to have positive off-diagonal values and rows which sum to 0, requiring $N \times (N - 1)$ parameters for N states. The Hamilton matrix is required to be Hermitian (in this case symmetric due to the additional assumption that it contained only real values), requiring $N/2 \times (N + 1)$ parameters for N states. This resulted in the Markov model having two additional learned parameters for the 4-state representation and 35 additional parameters for the 10-state. In future studies, (once complex numbers are supported by the autograd function in Pytorch) we plan to reformulate the QRNN model such that the Hamiltonian can more realistically contain any value in the complex plane. This would represent an increase in the degrees of freedom for the QRNN model and could help to better improve its performance compared to the MRNN model.

For external testing, final MRNN and QRNN models were trained in one of two ways – either with both real and synthetic data, or with only real data. For the MRNN testing results, it was found that the models trained using only real data outperformed those with synthetic data for three of the four system state vectors, while the opposite trend was observed in the QRNN model results. However, if we estimate that the uncertainty associated with the testing results is at least

as large as the 95% confidence intervals reported for the cross validation then these trends do not appear to be significant.

A key challenge for the MRNN and QRNN predictive models is to determine an adequate size of the system state vector. Higher dimensional vectors are preferable from a clinical standpoint because they can more faithfully represent the original data as demonstrated by **Figure 4.2** (d) and (e). In addition, higher dimensional system state vectors can more easily allow for the incorporation of additional types of patient data—such as biomarkers, radiomics, or dosimetric information—allowing for a more complete and individualized representation of patients. However, the results of this study suggest that there is a tradeoff between model performance and state vector complexity. This observation is not surprising: because the problem is formulated as a multi-class classification problem, an increase in system-state vector size increases the number of classes which the model must learn to distinguish without increasing the number of samples—effectively forcing each class to be more sparsely represented. This challenge makes it particularly difficult to draw confident conclusions about both models’ predictive abilities across different fractions (**Figure 4.5**). Based on the 4-state and 6-state results it appears there may be a slight reduction in model performance as the fraction number to be predicted grows further from the last initial state. However, the sparsity of representation across the possible states meant that the higher order representations (8-state and 10-state) had many fractions in which folds were omitted due to an inability to calculate the AUC using the samples. Finally, it is worth noting that of all of the system state vector lengths investigated in this study, only the 4-state models continued to learn over the entire course of 100 epochs (**Figure 4.6**). The other three higher order models all reached their maximum performance within the first 20-50 epochs and began overfitting from that point afterwards.

The current workflow for implementing plan adaptations on HNC patients involves identification of anatomical changes through visual inspection of daily CBCT images. When dramatic changes are identified by the clinical team, a new CT simulation can be performed to better quantify those changes and to assist in adaptive replanning. This current ART workflow is hindered by the time constraints inherent in a clinical environment. Models which can predict patient changes during treatment could help to alleviate this challenge by allowing clinical teams to anticipate and prepare for adaptations before they are necessary. Several previous studies have reported on the use of predictive models for HN cancer. Rosen et al. reported a model which predicted radiation-induced xerostomia using dose volume histogram metrics along with CBCT image information from the last two weeks of treatment ¹³. In another study, McCulloch et al. reported a model which predicted the need for replan based on estimated dose deviations calculated using daily CBCT images from the 5th, 10th, and 15th fractions ¹². To our knowledge, the models presented in this study are the first which predict anatomical changes beyond the next treatment fraction using early treatment data (the first 13 fractions). This approach allows for predictions to be made far enough in advance for clinical intervention and could also be used to improve future prediction studies. For example, geometric predictions could be coupled with daily dose calculations to provide a more accurate estimation of dosimetric deviations. These models further distinguish themselves from previous studies in that they can be formulated to encode any aspect of the patient state (geometrical, physiological, dosimetric, or higher order latent features).

The most accurate data driven models for outcome prediction are deep learning models because they can recognize complex nonlinear relationships in data. However, machine learning is sensitive to noise in the data, and in our case because there are inherent uncertainties in knowing the state of the patient during any given fraction, we can only offer noisy data to the algorithm for

training. The motivation for utilizing Markov/quantum models is that both of rely on probabilistic transition matrices to predict changes over time and are formulated deal with systems that have a stochastic component/inherent uncertainties.

The limitations associated with this study are summarized as follows: the patient variations predicted in this study were limited to geometric anatomical changes over the course of treatment. While there is limited consensus on which exact patient changes are relevant and/or necessary to track for the purpose of ART, dosimetric information (in particular dose deviations from the original treatment plan) represents another category of data commonly used for informing ART decisions. In addition, because many HNC patients experience xerostomia induced by radiation toxicity, it would be valuable to also include data related to the parotid glands in future studies. Due to a number of factors, including the CTV structure boundaries ending outside the field of view of the daily CBCTs as well as occasional failure of the DIR algorithm, it was not possible to include every treatment fraction for each patient. While generally at least 2 fractions per week were included, the RNN models might be more successful at learning relevant time-dependent patterns in the patient data if all fractions could be included. Finally, this study did not consider the impact of “measurements” on the patient system. In a quantum framework, any measurement on the system changes the state of the system itself (often referred to as “collapsing the systems wave function”). While the values predicted in this study were treated as somewhat artificial measurements (primarily because the quantitative value of the day-to-day volume of the primary CTV and treatment couch position, while potentially valuable, are not currently explicitly used in the clinical workflow for ART decision-making), it would be interesting in future studies to consider how these models can predict the impact of clinical decisions (such as the choice to

modify a plan) over the course of treatment. Such a study may reveal additional benefits that are more unique to the quantum-based QRNN model.

4.5 Conclusion

In this study we report on the design and assessment of two novel deep machine learning frameworks for predicting anatomical changes in radiotherapy patients. A comparison between the performance of the quantum-inspired and Markov-inspired models indicated that the Markov models tend to exhibit higher performance. In addition, it was found that there was a tradeoff between model performance and the size of the system state vector. A major challenge of this study was that as the system state dimension increased, so did the sparsity of class representation, leading to a decrease in model performance. Further studies utilizing both greater numbers of patients with diverse treatment outcomes, information regarding clinical decisions over the course of treatment, as well as additional datatypes beyond basic imaging features such as tumor volume are necessary to fully evaluate the predictive power of these models.

4.6 References

1. Baskar R, Lee KA, Yeo R, Yeoh K-W. Cancer and radiation therapy: current advances and future directions. *Int J Med Sci.* 2012;9(3):193-199.
2. Bortfeld T. IMRT: a review and preview. *Physics in Medicine and Biology.* 2006;51(13):R363-R379.
3. Teoh M, Clark CH, Wood K, Whitaker S, Nisbet A. Volumetric modulated arc therapy: a review of current literature and clinical use in practice. *Br J Radiol.* 2011;84(1007):967-996.
4. Yan D, Vicini F, Wong J, Martinez A. Adaptive radiation therapy. *Physics in Medicine and Biology.* 1997;42(1):123-132.
5. Schwartz DL, Garden AS, Thomas J, et al. Adaptive Radiotherapy for Head-and-Neck Cancer: Initial Clinical Outcomes From a Prospective Trial. *International Journal of Radiation Oncology* Biology* Physics.* 2012;83(3):986-993.
6. Wu QJ, Li T, Wu Q, Yin F-F. Adaptive Radiation Therapy: Technical Components and Clinical Applications. *The Cancer Journal.* 2011;17(3).
7. Lim-Reinders S, Keller BM, Al-Ward S, Sahgal A, Kim A. Online adaptive radiation therapy. *International Journal of Radiation Oncology* Biology* Physics.* 2017;99(4):994-1003.

8. Sonke J-J, Aznar M, Rasch C. Adaptive Radiotherapy for Anatomical Changes. *Seminars in Radiation Oncology*. 2019;29(3):245-257.
9. Tseng H-H, Luo Y, Ten Haken RK, El Naqa I. The Role of Machine Learning in Knowledge-Based Response-Adapted Radiotherapy. *Front Oncol*. 2018;8:266-266.
10. Zhen X, Chen J, Zhong Z, et al. Deep convolutional neural network with transfer learning for rectum toxicity prediction in cervical cancer radiotherapy: a feasibility study. *Physics in Medicine & Biology*. 2017;62(21):8246.
11. Luo Y, El Naqa I, McShan DL, et al. Unraveling biophysical interactions of radiation pneumonitis in non-small-cell lung cancer via Bayesian network analysis. *Radiotherapy and oncology : journal of the European Society for Therapeutic Radiology and Oncology*. 2017;123(1):85-92.
12. McCulloch MM, Lee C, Rosen BS, et al. Predictive models to determine clinically relevant deviations in delivered dose for head and neck cancer. *Practical radiation oncology*. 2019;9(4):e422-e431.
13. Rosen BS, Hawkins PG, Polan DF, et al. Early Changes in Serial CBCT-Measured Parotid Gland Biomarkers Predict Chronic Xerostomia After Head and Neck Radiation Therapy. *Int J Radiat Oncol Biol Phys*. 2018;102(4):1319-1329.
14. Vallières M, Kay-Rivest E, Perrin LJ, et al. Radiomics strategies for risk assessment of tumour failure in head-and-neck cancer. *Scientific Reports*. 2017;7(1):10117.
15. Heukelom J, Fuller CD. Head and Neck Cancer Adaptive Radiation Therapy (ART): Conceptual Considerations for the Informed Clinician. *Seminars in Radiation Oncology*. 2019;29(3):258-273.
16. Møller DS, Holt MI, Alber M, et al. Adaptive radiotherapy for advanced lung cancer ensures target coverage and decreases lung dose. *Radiotherapy and Oncology*. 2016;121(1):32-38.
17. Bhide SA, Davies M, Burke K, et al. Weekly volume and dosimetric changes during chemoradiotherapy with intensity-modulated radiation therapy for head and neck cancer: a prospective observational study. *International Journal of Radiation Oncology* Biology* Physics*. 2010;76(5):1360-1368.
18. Kamrani A, Azimi M, Nasr EA. Predictive modeling of tumors using RP. 2015.
19. Busemeyer JR, Wang Z, Pothos E. Quantum models of cognition and decision. *Oxford handbook of computational and mathematical psychology*. 2015:369-389.
20. Rigden JS. Einstein's revolutionary paper. *Physics World*. 2005;18(4):18-19.
21. Kragh H. Max Planck: the reluctant revolutionary. *Physics World*. 2000;13(12):31-36.
22. Nielsen MA, Chuang IL. Quantum computation and quantum information.
23. Norris JR. *Markov Chains*. Cambridge University Press; 1998.
24. LeCun Y, Bengio Y, Hinton G. Deep learning. *nature*. 2015;521(7553):436-444.
25. Cui S, Tseng H-H, Pakela J, Ten Haken RK, El Naqa I. Introduction to machine and deep learning for medical physicists. *Medical Physics*. 2020;47(5):e127-e147.
26. Lipton ZC, Berkowitz J, Elkan C. A critical review of recurrent neural networks for sequence learning. *arXiv preprint arXiv:1506.00019*. 2015.
27. Mourad M, Jetmore T, Jategaonkar AA, Moubayed S, Moshier E, Urken ML. Epidemiological Trends of Head and Neck Cancer in the United States: A SEER Population Study. *Journal of Oral and Maxillofacial Surgery*. 2017;75(12):2562-2572.

28. Fakhry C, Westra WH, Li S, et al. Improved survival of patients with human papillomavirus-positive head and neck squamous cell carcinoma in a prospective clinical trial. *J Natl Cancer Inst.* 2008;100(4):261-269.
29. Lindel K, Beer KT, Laissue J, Greiner RH, Aebbersold DM. Human papillomavirus positive squamous cell carcinoma of the oropharynx: a radiosensitive subgroup of head and neck carcinoma. *Cancer.* 2001;92(4):805-813.
30. Gersho A, Gray RM. *Vector quantization and signal compression.* Vol 159: Springer Science & Business Media; 2012.
31. Yoon J, Jarrett D, van der Schaar M. Time-series generative adversarial networks. 2019.
32. Cho K, Van Merriënboer B, Gulcehre C, et al. Learning phrase representations using RNN encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078.* 2014.
33. Golub GH, Van Loan CF. *Matrix Computations.* Johns Hopkins University Press; 2013.
34. Skafté N. Cuda_exmp. *GitHub Repository.* 2018.
35. Cawley GC, Talbot NLC. On over-fitting in model selection and subsequent selection bias in performance evaluation. *The Journal of Machine Learning Research.* 2010;11:2079-2107.
36. Hand DJ, Till RJ. A Simple Generalisation of the Area Under the ROC Curve for Multiple Class Classification Problems. *Machine Learning.* 2001;45(2):171-186.
37. DeLong ER, DeLong DM, Clarke-Pearson DL. Comparing the Areas under Two or More Correlated Receiver Operating Characteristic Curves: A Nonparametric Approach. *Biometrics.* 1988;44(3):837-845.

Chapter 5 Prediction of Adaptation Decisions and Outcomes in Larynx Cancer Patients

This chapter features two related studies conducted through collaboration with physicians at Michigan Medicine. The candidate's role in these studies involved processing image data and calculation of radiomic/imaging features, as well as training, validation, and testing of the classical and deep learning models. The physicians' role included contouring of computed tomography (CT) images, acquisition of relevant clinical features/outcomes from patient charts, and providing clinical expertise in interpretation of results. Results from the first study are currently undergoing peer review in:

Gharzai, L. A., Pakela, J. M., Jaworski, E., El-Naqa, I., Schonewolf, C. A., Hawkins, P. G., Wilkie, J. R., Spector, M. E., Bradford, C. R., Chinn, S. B., Hogikyan, N., Malloy, K., Kupfer, R., Shuman, A., Stucken, C. L., Prince, M., Shah, J., Swiecicki, P. L., Casper, K., Eisbruch, A., Wolf, G., Worden, F., Srinivasan, A., Mierzwa, M. L. (expected 2021), Imaging response versus operative laryngoscopy assessment of induction chemotherapy response in an induction bioselection approach to larynx cancer. (submitted to Oral Oncology).

Another manuscript is under processing with a focus on deep learning application that the candidate will be leading. The presented text which appears in this chapter was rewritten by the candidate.

5.1 Introduction

5.1.1 Bioselection as a Method for Adaption in Larynx Cancer Patients

In larynx cancer, the use of bioselection treatment strategies based on tumor response to induction chemotherapy (IC) has been tied to improved patient outcomes—including overall survival (OS) and larynx preservation.^{1,2} The current bioselection paradigm involves patients undergoing direct laryngoscopy after their first round of chemotherapy (i.e., IC) in order to assess their tumor's response to the treatment. During the laryngoscopy, a surgeon visually inspects the primary tumor and provides a clinical assessment of its response. Patients who experience $\geq 50\%$ tumor response (according to surgical assessment) advance to chemoradiation; otherwise, patients undergo a laryngectomy—a surgical removal of the larynx which critically alters or impairs a patient's ability to speak, taste, smell, and swallow.^{3,4} The use of a bioselection approach is therefore valuable not only for its improvement in patient survival rates but also in its ability to preserve patient quality of life. Despite data in support of this approach, barriers exist related to the direct laryngoscopy procedure (including financial cost, patient risks from intubation/anesthetization, and surgeon availability) which limit its wider implementation. In addition, because this assessment is made via physical observation by a surgeon, it is largely subjective and may suffer from limitations in accuracy and intra-observer variability.

5.1.2 Data-driven Outcomes Prediction in Head and Neck Cancers

Data-driven outcomes prediction is an active field in radiation oncology that plays a role in helping clinicians better understand the relationship between patient features and clinical outcomes, allowing for implementation of personalized medicine—in which a patient's unique characteristics and biomarkers play a role in identifying their optimal treatment. In head and neck cancer, there exists a substantial number of studies evaluating the use of many different types of clinical features

as prognostic factors for patient outcomes.⁵⁻¹⁰ For larynx cancer in particular, the incorporation of data extracted from radiological medical images (radiomics) has been recognized to be of potential value for better understanding and management of a cancer characterized by complex regional and functional anatomy.¹¹ There has also been increasing interest in the use of advanced machine learning methods such as deep learning (DL) models for outcomes prediction, particularly for the incorporation of higher dimensional feature sets, including radiomics and genomics, and for the ability to engineer features directly from medical images.^{12,13}

5.1.3 Study Objectives

In the case of laryngeal cancer patients, a better understanding of the relationship between quantitative patient features and clinical outcomes could help to make the bioselection process less subjective. Additionally, a fully verified model could potentially serve in place of, or in supplement to, a surgeon's expert assessment: particularly in non-academic centers with greater resource constraints. To this end, the primary goal of this study was to retrospectively investigate the use of quantitative, non-invasive features in the form of imaging radiomics, lab, and clinical data as predictors for two types of outcomes: overall survival (OS) and laryngectomy free survival (LFS) for patients undergoing bioselection treatment. We also investigated the use of these quantitative features to predict the surgeon-assessed response to IC (ICR). The secondary goal of this study was to investigate whether a particular class of data-driven model (traditional linear models or DL) was better suited for this task.

5.2 Methods

5.2.1 Patient Cohort

The inclusion criteria for this study consisted of locally advanced laryngeal cancer patients who were treated with a bioselection paradigm and had contrast-enhanced CT images taken immediately prior to and after induction chemotherapy as well as valuable laboratory data. The training dataset for this study consisted of 93 patients who were treated at a single institution and belonged to one of two prospective clinical trials (UMCC 9520 and NCT 01633541). An additional 22 laryngeal cancer patients who also underwent bioselection off-trial were selected as an external testing dataset. The demographics and treatment information of both datasets are summarized in **Table 5.1**.

Table 5.1: Demographics of the training and testing dataset.

Patient Demographics		
	Training Dataset (n=93)	Testing Dataset (n=22)
	N (%)	N (%)
<i>Clinical Information</i>		
Age (years, median, range)	57 (19-82)	58.5 (29-77)
Pre-treatment tumor size (cc, mean, range)	23.85 (1.1-74.6)	13.26 (2-37.9)
Site		
	Supraglottis	70 (75.3%)
	Glottis	14 (63.6%)
	Hypopharynx	5 (22.7%)
		3 (13.6%)
T-stage		
	T1	0 (0%)
	T2	1 (4.5%)
	T3	3 (3.2%)
	T4	39 (41.9%)
		10 (45.5%)
		9 (40.9%)
N-stage		

Patient Demographics		
	Training Dataset (n=93)	Testing Dataset (n=22)
N0	33 (35.5%)	8 (36.4%)
N1	11 (11.8%)	1 (4.5%)
N2a	2 (2.2%)	0 (0%)
N2b	22 (23.7%)	4 (18.2%)
N2c	25 (26.9%)	6 (27.3%)
N3	0 (0%)	3 (13.6%)
Median Neutrophil-Lymphocyte Ratio (median, range)	2.88 (0.88-56)	2.84 (0.44-10.85)
Median Lymphocyte-Monocyte Ratio (median, range)	1.55 (0.28-5.67)	2.58 (0.70-12.14)
<i>Treatment Information</i>		
Number of days between CT imaging and surgeon assessment of response (mean, standard deviation)	1.8 (2.2)	1.1 (0.97)
Tumor Response (by surgeon assessment)		
<50%	33 (35.5%)	4 (18.2%)
≥50%	60 (64.5%)	18 (81.8%)
Average Tumor Reduction by imaging, % (standard deviation)	35.45% (34.78%)	55.98% (26.95%)
Treatment received		
Total laryngectomy	23 (24.7%)	4 (18.2%)
Chemoradiation	70 (75.3%)	18 (81.8%)

5.2.2 Feature Acquisition and Processing

Clinical variables were acquired from patient electronic medical records and included tumor site label, T-stage, and N-stage.¹⁴ Laboratory-based features included neutrophil-lymphocyte ratio (NLR) and lymphocyte-monocyte ratio (LMR) calculated from pre-treatment complete blood counts. NLR and LMR were incorporated into the models in one of two ways: 1) as log-transformed, continuous variables or 2) as binary variables determined by cut-points reported

previously in literature to have prognostic value for IC outcome in the 9520-patient population.¹⁵ Specifically, NLR variables were set to 0 for NLR values ≤ 2.8 and set to 1 otherwise, while LMR variables were set to 0 for LMR values ≥ 2.8 and set to 1 otherwise.

To acquire image-based features, primary gross tumor volume (GTV) structures were contoured in pre- and post-IC CT images by experienced radiation oncologists using clinical software (Varian Medical Systems, Inc. Eclipse Treatment Planning System). CT-based morphological variables investigated were pre-IC primary GTV volume (GTVp) as well as the percent reduction in GTV volume between pre and post IC (Δ GTV).

Radiomic texture features calculated from pre- and post-IC GTV images were also included in this study. Using the RADIOMICS toolbox for MATLAB,¹⁶ imaging data underwent preprocessing, including isotropic resampling to the planar resolution of the CT image and the creation of image masks consisting of 3D arrays of 0s and 1s which define the region of interest (ROI) specified by the primary GTV contour data.

First-order texture features were calculated from 64 bin histograms of the gray level image data from each ROI. Higher order texture features were calculated using 32 gray-level Lloyd-Max quantization. Radiomic features were input into the models as *delta radiomics* (the difference between the radiomic values calculated before and after IC). The 43 radiomics features used for this study are summarized in **Table 5.2**, and a complete description of their definitions and properties can be found in the reference manual of the Image Biomarker Standardization Initiative.¹⁷

Table 5.2: Radiomic imaging features.

Radiomic Features	
Global Textures	<ul style="list-style-type: none"> Variance Skewness Kurtosis
Gray-Level Co-occurrence Matrix (GLCM)	<ul style="list-style-type: none"> Energy Contrast Correlation Homogeneity Variance Sum Average Entropy Dissimilarity Autocorrelation
Gray-Level Run-Length Matrix (GLRLM)	<ul style="list-style-type: none"> Short Run Emphasis (SRE) Long Run Emphasis (LRE) Gray-Level Nonuniformity (GLN) Run-Length Nonuniformity (RLN) Run Percentage (RP) Low Gray-Level Run Emphasis (LGRE) High Gray-Level Run Emphasis (HGRE) Short Run Low Gray-Level Emphasis (SRLGE) Short Run High Gray-Level Emphasis (SRHGE) Long Run Low Gray-Level Emphasis (LRLGE) Long Run High Gray-Level Emphasis (LRHGE) Gray-Level Variance (GLV) Run Length Variance (RLV)
Gray-Level Size Zone Matrix (GLSZM)	<ul style="list-style-type: none"> Small Zone Emphasis (SZE) Large Zone Emphasis (LZE) Gray-Level Nonuniformity (GLN) Zone-Size Nonuniformity (ZSN) Zone Percentage (ZP) Low Gray-Level Zone Emphasis (LGZE) High Gray-Level Zone Emphasis (HGZE) Small Zone Low Gray-Level Emphasis (SZLGE) Small Zone High Gray-Level Emphasis (SZHGE) Large Zone Low Gray-Level Emphasis (LZLGE) Large Zone High Gray-Level Emphasis (LZHGE) Gray-Level Variance (GLV)

Neighbourhood Gray-Tone Difference Matrix (NGTDM)	Zone Size Variance (ZSV)
	Coarseness
	Contrast
	Busyness
	Complexity
	Strength

During the later development of the DL models, a decision was made to re-calculate the radiomic features using a clinical target volume (CTV) in place of the GTV in order to capture additional tumor information from microscopic disease not visible to the human eye. The CTVs were defined by adding a 1 cm margin to each primary GTV structure. This was accomplished by applying image dilation with a spherical kernel to each ROI. During this process, it was discovered that a small subsection of CT images (from 10 patients) contained discrepancies in the metadata between different slices of CT scan occurring within the ROI. Specifically, the DICOM attribute, image position (patient)—which details the x, y, and z coordinates of the upper left corner of each slice in mm—was found in some patient images to vary along an axis that was not the image orientation axis. This is believed to have been caused in the situation where patient imaging was paused and restarted—resulting in two images being effectively spliced together. While this shift appeared to be handled without issue internally by the treatment planning software based on visual inspection of the images, it was not communicated during export of the image data to DICOM files, resulting in patient structures not lining up correctly across the CT slices. Patients who had CT scans with this discrepancy were discarded from the deep learning study because the local intensity information of their GTV structures could not be faithfully represented. This resulted in the DL models being trained on 83 samples with external testing occurring on 21 samples.

5.2.3 Statistical Regression Models Design

A logistic regression (LR) model¹⁸ was used to predict the surgeon-reported ICR, where ICR was defined as $\geq 50\%$ shrinkage of the tumor volume. LR models the log odds for binary outcomes as a linear function of the feature set, \mathbf{x} :

$$\log\left(\frac{p}{p-1}\right) = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \dots \quad 5.1$$

Where the values, $\boldsymbol{\beta} \in \{\beta_0, \beta_1, \beta_2, \dots\}$ are model parameters which represent the relative importance of each feature in \mathbf{x} . From this expression, the probability of event $p = P(Y = 1|\mathbf{x})$ occurring can be written as:

$$p = \frac{e^{\mathbf{x}^T \cdot \boldsymbol{\beta}}}{1 + e^{\mathbf{x}^T \cdot \boldsymbol{\beta}}} \quad 5.2$$

The values e^{β_i} are referred to as odds ratios and can be used to understand the relative importance of each feature x_i to the model. In order to improve model performance, feature selection was performed using elastic-net regularization¹⁹ and validated using a 3-layer nested cross validation (CV) scheme described in Section 5.2.4. The objective function for the regularized LR model is represented by:

$$\operatorname{argmin}_{\boldsymbol{\beta}} \left(-\frac{1}{N} \left[\sum_{i=1}^N y_i \cdot (\mathbf{x}_i^T \cdot \boldsymbol{\beta}) - \log(1 + e^{\mathbf{x}_i^T \cdot \boldsymbol{\beta}}) \right] + \lambda \left[\frac{1}{2} (1 - \alpha) \|\boldsymbol{\beta}\|_2^2 + \alpha \|\boldsymbol{\beta}\|_1 \right] \right) \quad 5.3$$

Where for patient i , y_i and \mathbf{x}_i represent the binary ICR outcome (defined as the surgeon-reported tumor response) and the feature vector, respectively; N is the total number of patients. The term within the first bracket is the negative log likelihood of ICR, and the second term is the elastic-net regularization penalty. Elastic-net regularization is a regression method defined as a weighted combination of ridge regression and least absolute shrinkage and selection operator (LASSO) regression. The parameter λ represents the regularization penalty, while α is a mixing parameter which takes a value between 0 and 1 with $\alpha = 0$ representing pure ridge regression and $\alpha = 1$

representing pure LASSO regression. For each of the three outcome models, α was set to 1 as a default. In the event that pure LASSO regression was too parsimonious, α was relaxed to allow the model to select more features.

Cox proportional hazards (Cox) models were used to predict LFS and OS, where survival times were defined as the time between the pre-induction CT scan and the event of interest. Both models were right-censored to time of last follow up. Cox models are survival models that estimate the risk of an event occurring at time t as:

$$h(t) = h_0(t) * \exp(\boldsymbol{\beta} \cdot \mathbf{x}) \quad 5.4$$

Where $h_0(t)$ is a baseline hazard function and $\boldsymbol{\beta}$ is a vector of feature coefficients fit during model training and represents the relative importance of each feature in \mathbf{x} .²⁰ The objective function for the regularized Cox model is represented by:

$$\underset{\boldsymbol{\beta}}{\operatorname{argmin}} \left(-\log \left[\prod_{i=1}^m \frac{e^{\mathbf{x}_{j(i)}^T \boldsymbol{\beta}}}{\sum_{j \in R_i} e^{\mathbf{x}_j^T \boldsymbol{\beta}}} \right] + \lambda \left[\frac{1}{2} (1 - \alpha) \|\boldsymbol{\beta}\|_2^2 + \alpha \|\boldsymbol{\beta}\|_1 \right] \right) \quad 5.5$$

Where m represents the total number of unique survival times, \mathbf{x}_j represents a vector of co-variates for patient j , and $\boldsymbol{\beta}$ is a vector of feature coefficients; the index $j(i)$ represents the index of the patient who experienced failure (or was censored) at time, t_i ; and R_i represents subset of patient indices who are still at risk of failure at time t_i . The first term in the expression is the negative log of the Cox partial likelihood—defined as the product of the probabilities of each event at time, i , out of m event times given R_i patients at risk of failure at t_i . The second term is the elastic-net regularization. As with ICR, feature selection for both OS and LS was implemented via elastic-net and the model validated using a 3-layer nested cross validation scheme described in the next section. Note that all cross validation used in this study (for both classical and DL models) was stratified by outcome.

5.2.4 Nested Cross Validation for Feature Selection in Classical Models

The motivation for using nested cross validation is to avoid introducing bias into the models which would lead to overly optimistic performance evaluations.²¹ In the innermost loop, 10-fold cross validation was performed using the `cvglmnet()` function from the GLMNET package in MATLAB to tune the regularization penalty, λ . The optimal λ was chosen based on the minimal partial likelihood loss for the Cox models and the minimal deviance for LR.

In the middle loop, 10 iterations of 5-fold cross validation were performed to select features. Performing multiple iterations of CV was done in order to minimize variance caused by sensitivity to how the data is split. For each training fold, the optimal λ was acquired from the innermost loop. The Cox (or LR) model was then fit to the training data. The optimal set of features was selected based on which set of features was most frequently selected out of 50 different folds and returned to the outer loop.

In the outermost loop, 10 iterations of 5-fold cross validation were performed to evaluate the performance of the model-building method. For each training fold, an optimal feature set was returned from the middle loop. A non-regularized Cox or LR model was then trained on these features and c-index performance was measured on the test fold. The mean and standard deviations of the c-index values across the total of 50 folds was recorded to assess the quality of the model building method.

A final model (i.e., feature set) was obtained by choosing the feature set which appeared most frequently in the 50 folds of the outer CV loop. The Cox (or LR) model was then trained on the full training dataset using these features and evaluated on the test dataset to assess model generalizability.

5.2.5 Deep Learning Model Design

In order to assess potential higher order, nonlinear relationships between covariates and patient outcomes, two deep learning models were also developed.

DeepSurv²² is a Cox proportional hazards deep neural network which was utilized to predict OS and LFS outcomes. Recall the expression for the hazard function of the Cox survival model given by Equation 5.4 DeepSurv takes advantage of the fact that the expression in the exponent, $\boldsymbol{\beta} \cdot \mathbf{x}$, can be generalized to any unknown function, $h(\boldsymbol{\beta}, \mathbf{x})$.²⁰ DeepSurv takes the patient features, \mathbf{x} , as inputs and propagates them through a set of fully connected hidden layers, each with dropout and non-linear activation function.²² The network's output consists of a single node, $\hat{h}(\boldsymbol{\beta}, \mathbf{x})$, which represents an estimate of the generalized exponential term in the Cox hazard function. The networks loss function is similar to that used for the classical implementation of the Cox model but uses ridge regularization instead of elastic net:

$$l = -\log \left[\prod_{i=1}^m \frac{e^{\hat{h}(\boldsymbol{\beta}, \mathbf{x}_{j(i)})}}{\sum_{j \in R_i} e^{\hat{h}(\boldsymbol{\beta}, \mathbf{x}_j)}} \right] + \lambda \|\boldsymbol{\beta}\|_2^2 \quad 5.6$$

Two DeepSurv models were trained to predict OS and LFS outcomes using a 2-layer nested cross validation framework, in which the model hyperparameters (regularization penalty, dropout, learning rate, number of hidden layers, and number of nodes per hidden layer) were tuned during the inner loop of 5-fold CV and model training and validation occurred in the outer loop of 5-fold CV.

A feedforward neural network was designed to predict the probability of ICR outcomes. The network took patient features as inputs and, after propagation through a set number of hidden layers (each undergoing rectified linear unit [ReLU] activation and dropout), produced an output consisting of a single node which was passed through a sigmoid activation function to generate a final output between 0 and 1. The model was trained using a binary cross entropy loss function:

$$l = (y \cdot \log(x) + (1 - y) \cdot \log(1 - x)) \quad 5.7$$

Similar to DeepSurv, a 2-layer nested cross validation was performed for hyperparameter tuning (dropout, learning rate, number of layers, and number of nodes per layer).

5.3 Results

5.3.1 Classical Model Results

Both classical models (Cox and LR) were unable to train when the delta radiomic variables were included in the feature set. **Tables 5.3** and **5.4** summarize the selected features and model performance when the 43 delta radiomics features were excluded from the feature set.

For the ICR model (implemented through LR), two features were selected through the nested-cross validation procedure, Δ GTV and N-stage, with odds ratios (p -values) of 5.78 (<0.001) and 1.64 (0.0027) respectively. The ICR model performance was assessed using the area under the receiver operating characteristic curve (AUC). The mean AUC was 0.828 (95% CI: 0.803-0.853) for cross validation and 0.847 on the external testing dataset.

For the LFS model (implemented through Cox proportional hazards), a single feature was selected through nested cross validation, Δ GTV, with a hazard ratio of 0.58 and a p -value of <0.0001. The LFS model performance was assessed using Harrel's c-index, which was measured to be 0.724 (95% CI: 0.699-0.749) during validation and 0.721 on the external testing dataset. An additional feature, NLR, was retrospectively identified as contributing significantly to the LFS model ($p = 0.006$) with a hazard ratio of 2.8. However, when this feature was included with Δ GTV, the resulting c-index was 0.71 and therefore did not improve the model performance.

For the OS model (implemented through Cox proportional hazards), two features were selected through nested cross validation, GTVp and N-stage, with hazard ratios (p -values) of 1.3 (0.174) and 1.48 (0.039) respectively. The OS model was evaluated using Harrel's c-indices, with

a mean c-index of 0.601 (95% CI: 0.567-0.635) for cross validation and 0.552 on the external testing dataset.

Table 5.3: Feature selection for model building.

Model Feature Selection			
Model	Feature set	Odds Ratio	P-value
ICR	Δ GTV	5.78	<0.001
	N-stage	1.64	0.0027
Hazard Ratio			
LFS	Δ GTV	0.58	<0.0001
OS	GTVp	1.30	0.174
	N-stage	1.48	0.039

Table 5.4: Classical Model performance.

Model Performance AUC/Harrel's C-indices (95% CI range)		
	Cross Validation	Testing
ICR (AUC)	0.828 (0.803-0.853)	0.847
LFS (c-index)	0.724 (0.699-0.749)	0.721
OS (c-index)	0.601 (0.567-0.635)	0.552

5.3.2 Deep Learning Model Results

The DL models were trained on three sets of features—clinical only, radiomics only, and clinical + radiomics—in order to assess whether the radiomics features could be used in complement to or in replacement of the clinical features. It was expected that the greater complexity of the deep learning models would allow them to train with a greater number of features than the LR and Cox

models. This was found to be true to the extent that for all three sets of features, both DL models were able learn during model training. However, for all three outcomes (LFS, OS, and ICR), the DL models exhibited the best performance when only the clinical features were included. **Table 5.5** summarizes the model performance achieved using DL methods for the clinical features. The radiomics features, and combined clinical and radiomics features are also included in **Table 5.5** for completeness. Like the classical models, performance was assessed using Harrel’s c-statistic for the survival models and the AUC for the classification model.

Table 5.5: Deep learning model performance.

Model Performance AUC/Harrel’s C-indices (95% CI range)		
	Cross Validation	Testing
Clinical and Radiomics		
ICR (AUC)	0.522 (0.298-0.745)	0.603
LFS (c-index)	0.467 (0.367-0.567)	0.625
OS (c-index)	0.619 (0.420-0.8192)	0.533
Radiomics Only		
ICR (AUC)	0.525 (0.424-0.625)	0.485
LFS (c-index)	0.421 (0.336-0.506)	0.58
OS (c-index)	0.523 (0.375-0.6706)	0.567
Clinical Only		
ICR (AUC)	0.789 (0.683-0.896)	0.838
LFS (c-index)	0.674 (0.596-0.752)	0.693
OS (c-index)	0.585 (0.460-0.710)	0.578

5.4 Discussion

In this study, we evaluated the use of quantitative, non-invasive features for predicting outcomes for laryngeal cancer patients in a bioselection treatment paradigm. Using classical machine learning methods, we identified two features (Δ GTV, and N-stage) which contributed significantly to the models' predicted outcome. Given that 30% of patients had discordant results between surgeon assessment and actual tumor volume response, it is possible that the number of nodes present inadvertently influences the surgeon's assessment of the primary tumor volume. Our results suggest that these features could serve as a quantitative metric in assisting with the decision process for patient treatment selection.

We did not identify any radiomic texture features in this study which could contribute positively to model performance. In the case of the classical models, the models were unable to learn using the nested cross validation framework when the additional 43 texture features were added. For the 3-level nested cross validation scheme, feature selection occurred in the second layer over training on 80% of the full training dataset (~75 patients). This challenge was likely due to the fact that when the radiomics features were included, the total number of features (50) approached the number of training samples during feature selection via LASSO. For the DL models, the addition of the radiomics features was found to have a negative impact on model performance. These findings suggest that the radiomic features as processed did not play a complimentary role to the clinical features. This could be explained by the fact that radiomic features have been found to have high correlation with tumor volume.²³

Deep learning methods are particularly attractive for use as predictive models because they do not require time-consuming feature engineering and are also capable of modeling complex, non-linear relationships between features and outcomes. However, in this study the deep learning

models trained using clinical features did not exhibit markedly different performance than the classical models and in some cases performed slightly worse. This may suggest that there simply wasn't sufficient training data to overcome model variance.

One limitation of this study was the small sample size. The original 93 patients used in the classical model training were reduced to only 67 for the very inner loop of cross validation. In addition, the patient dataset for this study ranged in treatment year from 2003-2019, resulting in some patient scans being of lower image quality; this may have impacted the quality and robustness of the radiomic features. Finally, there is evidence emerging that advanced imaging modalities such as PET, CT perfusion, and MRI hold prognostic value for HCN.²⁴⁻³⁶ For this study we were interested in including data from PET imaging but ultimately were unable to due to the low number of PET scans available in the patient datasets. Future studies should aim to increase the size of the patient dataset and seek to perform a thorough pre-analysis of radiomic features to ensure repeatability and robustness prior to incorporation into the model.

5.5 Conclusion

In this study, we investigated the role of CT-based, clinical, and lab features as predictors for OS, LFS, and surgeon-assessed response to IC. Initial results indicate that several features investigated may have value as predictors for patient outcomes. Due to limited data size, further studies with larger datasets are necessary to validate these results and further explore the potential of these variables.

5.6 References

1. Wolf GT, Bellile E, Eisbruch A, et al. Survival rates using individualized bioselection treatment methods in patients with advanced laryngeal cancer. *JAMA Otolaryngology–Head & Neck Surgery*. 2017;143(4):355-366.

2. Urba S, Wolf G, Eisbruch A, et al. Single-Cycle Induction Chemotherapy Selects Patients With Advanced Laryngeal Cancer for Combined Chemoradiation: A New Treatment Paradigm. *Journal of Clinical Oncology*. 2006;24(4):593-598.
3. McAuliffe MJ, Ward EC, Bassett L, Perkins K. Functional Speech Outcomes After Laryngectomy and Pharyngolaryngectomy. *Archives of Otolaryngology–Head & Neck Surgery*. 2000;126(6):705-709.
4. Mumovic G, Hocevar-Boltezar I. Olfaction and gustation abilities after a total laryngectomy. *Radiol Oncol*. 2014;48(3):301-306.
5. Bradford CR, Wolf GT, Carey TE, et al. Predictive markers for response to chemotherapy, organ preservation, and survival in patients with advanced laryngeal carcinoma. *Otolaryngology–Head and Neck Surgery*. 1999;121(5):534-538.
6. Hoban CW, Beesley LJ, Bellile EL, et al. Individualized outcome prognostication for patients with laryngeal cancer. *Cancer*. 2018;124(4):706-716.
7. Zhang H, Graham CM, Elci O, et al. Locally Advanced Squamous Cell Carcinoma of the Head and Neck: CT Texture and Histogram Analysis Allow Independent Prediction of Overall Survival in Patients Treated with Induction Chemotherapy. *Radiology*. 2013;269(3):801-809.
8. El Naqa I, Grigsby P, Apte A, et al. Exploring feature-based approaches in PET images for predicting cancer treatment outcomes. *Pattern Recognit*. 2009;42(6):1162-1171.
9. Tseng Y-J, Wang H-Y, Lin T-W, Lu J-J, Hsieh C-H, Liao C-T. Development of a Machine Learning Model for Survival Risk Stratification of Patients With Advanced Oral Cancer. *JAMA Network Open*. 2020;3(8):e2011768-e2011768.
10. Staton J, Robbins KT, Newman L, Samant S, Sebelik M, Vieira F. Factors Predictive of Poor Functional Outcome after Chemoradiation for Advanced Laryngeal Cancer. *Otolaryngology–Head and Neck Surgery*. 2002;127(1):43-47.
11. Chiesa-Estomba CM, Echaniz O, Larruscain E, Gonzalez-Garcia JA, Sistiaga-Suarez JA, Graña M. Radiomics and Texture Analysis in Laryngeal Cancer. Looking for New Frontiers in Precision Medicine through Imaging Analysis. *Cancers (Basel)*. 2019;11(10).
12. Diamant A, Chatterjee A, Vallières M, Shenouda G, Seuntjens J. Deep learning in head & neck cancer outcome prediction. *Scientific Reports*. 2019;9(1):2764.
13. Zhao Z, Li Y, Wu Y, Chen R. Deep learning-based model for predicting progression in patients with head and neck squamous cell carcinoma. *Cancer Biomark*. 2020;27(1):19-28.
14. Cancer Staging. National Cancer Institute. <https://www.cancer.gov/about-cancer/diagnosis-staging/staging>. Published 2015. Accessed 3/11/2021, 2021.
15. Risch Z, Bellile E, Townsend R, et al. Association of low neutrophil-to-lymphocyte ratios (NLR) and high lymphocyte-to-monocyte ratios (LMR) with overall survival and response to induction chemotherapy (IC) when used to select patients (pts) with locally advanced squamous cell of the larynx (LSCC) for combined chemoradiation (CRT). *Journal of Clinical Oncology*. 2019;37(15_suppl):e17539-e17539.
16. Vallières M, Freeman CR, Skamene SR, El Naqa I. A radiomics model from joint FDG-PET and MRI texture features for the prediction of lung metastases in soft-tissue sarcomas of the extremities. *Physics in Medicine and Biology*. 2015;60(14):5471-5496.
17. Zwanenburg A, Vallières M, Abdalah MA, et al. The Image Biomarker Standardization Initiative: Standardized Quantitative Radiomics for High-Throughput Image-based Phenotyping. *Radiology*. 2020;295(2):328-338.

18. Bishop CM. *Pattern recognition and machine learning*. New York : Springer, [2006] ©2006; 2006.
19. Zou H, Hastie T. Regularization and variable selection via the elastic net. *Journal of the royal statistical society: series B (statistical methodology)*. 2005;67(2):301-320.
20. Cox DR. Regression models and life-tables. *Journal of the Royal Statistical Society: Series B (Methodological)*. 1972;34(2):187-202.
21. Cawley GC, Talbot NLC. On over-fitting in model selection and subsequent selection bias in performance evaluation. *The Journal of Machine Learning Research*. 2010;11:2079-2107.
22. Katzman JL, Shaham U, Cloninger A, Bates J, Jiang T, Kluger Y. DeepSurv: personalized treatment recommender system using a Cox proportional hazards deep neural network. *BMC Medical Research Methodology*. 2018;18(1):24.
23. Vallières M, Kay-Rivest E, Perrin LJ, et al. Radiomics strategies for risk assessment of tumour failure in head-and-neck cancer. *Scientific Reports*. 2017;7(1):10117.
24. Hatakenaka M, Nakamura K, Yabuuchi H, et al. Pretreatment apparent diffusion coefficient of the primary lesion correlates with local failure in head-and-neck cancer treated with chemoradiotherapy or radiotherapy. *Int J Radiat Oncol Biol Phys*. 2011;81(2):339-345.
25. Lambrecht M, Van Calster B, Vandecaveye V, et al. Integrating pretreatment diffusion weighted MRI into a multivariable prognostic model for head and neck squamous cell carcinoma. *Radiother Oncol*. 2014;110(3):429-434.
26. Cao Y, Aryal M, Li P, et al. Predictive Values of MRI and PET Derived Quantitative Parameters for Patterns of Failure in Both p16+ and p16- High Risk Head and Neck Cancer. *Front Oncol*. 2019;9:1118.
27. Driessen JP, van Bommel AJ, van Kempen PM, et al. Correlation of human papillomavirus status with apparent diffusion coefficient of diffusion-weighted MRI in head and neck squamous cell carcinomas. *Head Neck*. 2016;38 Suppl 1:E613-618.
28. de Perrot T, Lenoir V, Domingo Ayllon M, Dulguerov N, Pusztaszeri M, Becker M. Apparent Diffusion Coefficient Histograms of Human Papillomavirus-Positive and Human Papillomavirus-Negative Head and Neck Squamous Cell Carcinoma: Assessment of Tumor Heterogeneity and Comparison with Histopathology. *AJNR Am J Neuroradiol*. 2017;38(11):2153-2160.
29. Cao Y, Popovtzer A, Li D, et al. Early prediction of outcome in advanced head-and-neck cancer based on tumor blood volume alterations during therapy: a prospective study. *Int J Radiat Oncol Biol Phys*. 2008;72(5):1287-1290.
30. Hermans R, Meijerink M, Van den Bogaert W, Rijnders A, Weltens C, Lambin P. Tumor perfusion rate determined noninvasively by dynamic computed tomography predicts outcome in head-and-neck cancer after radiotherapy. *Int J Radiat Oncol Biol Phys*. 2003;57(5):1351-1356.
31. Schwartz DL, Harris J, Yao M, et al. Metabolic tumor volume as a prognostic imaging-based biomarker for head-and-neck cancer: pilot results from Radiation Therapy Oncology Group protocol 0522. *Int J Radiat Oncol Biol Phys*. 2015;91(4):721-729.
32. Kikuchi M, Koyasu S, Shinohara S, et al. Prognostic value of pretreatment 18F-fluorodeoxyglucose positron emission tomography/CT volume-based parameters in patients with oropharyngeal squamous cell carcinoma with known p16 and p53 status. *Head Neck*. 2015;37(10):1524-1531.

33. Floberg JM, DeWees TA, Chin RI, et al. Pretreatment metabolic tumor volume as a prognostic factor in HPV-associated oropharyngeal cancer in the context of AJCC 8th edition staging. *Head Neck*. 2018;40(10):2280-2287.
34. Moan JM, Amdal CD, Malinen E, Svestad JG, Bogsrud TV, Dale E. The prognostic role of 18F-fluorodeoxyglucose PET in head and neck cancer depends on HPV status. *Radiother Oncol*. 2019;140:54-61.
35. Chepeha DB, Sacco AG, Oxford LE, et al. Advanced squamous cell carcinoma of the oropharynx: efficacy of positron emission tomography and computed tomography for determining primary tumor response during induction chemotherapy. *Head Neck*. 2009;31(4):452-460.
36. de Bree R, Wolf GT, de Keizer B, et al. Response assessment after induction chemotherapy for head and neck squamous cell carcinoma: From physical examination to modern imaging techniques and beyond. *Head Neck*. 2017;39(11):2329-2349.

Chapter 6 Discussion and Future Directives

6.1 Discussion

In this study, we designed and utilized stochastic and quantum-based algorithms in conjunction with classical and deep machine learning algorithms to address challenges in radiation oncology related to clinical implementation of image-guided adaptive radiotherapy (ART). To address the need for efficient, robust optimization techniques for treatment planning, we developed a quantum-inspired stochastic optimization algorithm, quantum tunnel annealing (QTA), which models a particle tunneling through a one-dimensional potential energy barrier. QTA was found to exhibit faster convergence rates than the previously used simulated annealing (SA) across different dose constraints and for both beamlet weight and direct aperture optimization. As such, QTA may be a promising candidate for optimization scenarios with stronger time constraints such as during adaptive replanning.

ART is a cost- and labor-intensive process which may not provide added benefit for all patients.¹ Because it is not possible (or clinically necessary) to perform ART for every patient, there is a need to develop tools and workflows which can identify which patients stand to benefit from ART. We designed and tested two stochastic deep learning frameworks (one quantum-based and the other Markov based) for predicting anatomical changes in head and neck cancer patients over the course of 30-35 fractions of radiotherapy (RT) treatment. Both algorithms showed promising performance in their ability to predict future patient states defined from vector-quantized distributions of volume changes and couch shifts.

Finally, an additional challenge related to ART implementation is that we do not have a full understanding of how different individual patient characteristics relate to clinical outcomes. We performed a study on laryngeal cancer patients to investigate the use of quantitative, noninvasive patient features for predicting patient outcomes as well as predicting clinical decision-making in a bioselection treatment paradigm. Our study identified two variables (nodal stage [N-stage] and change in primary gross tumor volume [GTV] before and after induction chemotherapy) to significantly contribute models predicting laryngectomy free survival and surgeon assessment of tumor response. Interestingly, no radiomics features were identified as contributing significantly to any of the models developed.

6.2 Current challenges and limitations

6.2.1 Data Acquisition, Standardization, and Quantization

In the era of big data, radiation oncology presents a unique data science challenge in that it contains an incredible volume of information on the individual patient level—including radiomic data, genomic data, disease classification, dosimetric data, etc.—but patient sample sizes themselves are relatively small (on the order of tens to hundreds) whereas other machine learning endeavors typically utilize training sample sizes orders of magnitude larger. Particularly for the case of deep learning, the property of increased model complexity results in models which are vulnerable to fitting noise when the dataset is too small because the deep learning models memorize the training dataset rather than learning the true relationship between the data and the output.² The studies presented in this work involving machine learning were limited to patients treated at a single institution, and thus training datasets were ultimately limited to a range of 83-125 samples. When possible, additional synthetic data was generated to assist with model training; however, efforts

being undertaken to make datasets publicly available to medical researchers from multiple institutions represent an important task for the future of machine learning in medical physics.³

Another challenge associated with data collection in these studies was related to a lack of standardization of labeling across patients. These instances tended to be an inherent aspect of the clinical workflow. For example, in the study presented in Chapter 4, it was necessary to extract the volume information from the primary clinical target volume (CTV) structures of each patient. Head and neck cancers typically involve a main tumor structure as well as one or more additional nodes, resulting in a primary—as well as secondary or tertiary—CTV structures. In our patient data, the contour names assigned to different CTV structures were inconsistent across patients because the names chosen by the dosimetrist during the treatment planning were descriptive of the unique geometries and/or treatment objectives associated with each patient. Primary CTV labels included variations of “CTV High,” “CTV 70,” “CTV Left,” “CTV Right.” Due to ambiguity (it is not possible to know whether “CTV Left” or “CTV Right” is the primary or secondary structure), it was necessary to manually review each patient’s treatment plan to create a list of keys identifying the name of the primary CTV for each patient. In addition, during data processing many patients were identified who had multiple cone-beam computed tomography images (CBCTs) acquired on the same fraction day. It was important therefore to ensure that only information from the CBCT taken immediately prior to the recorded treatment time-stamp was used. However, it was found that some patients experienced treatment fractions in which treatment began in one couch position, was paused, and then resumed after adjustments. In these instances, such fractions were discarded from the training model. These inconsistencies required careful, time-consuming combing of the dataset for irregularities. While the original intent of creating medical data is to perform functions

related to the clinical workflow, a few simple guidelines to support the standardization of how data is recorded could generally save hundreds of hours spent in medical research.

Finally, an additional challenge associated with data processing in this study was related to the transformation of multivariate, continuous patient data into orthonormal state vectors for the quantum and Markov-based recursive neural networks (RNNs). A tradeoff was found between the number of potential states that could be encoded in the state vector and the overall model performance. During the study, the data input into the model consisted of four variables: volume and the x, y, and z values denoting change in table position. Future studies which seek to encode more complex or complete descriptions of the patient state (through the inclusion of additional features) may require more sophisticated data reduction techniques to adequately capture the patient status at each fraction.

6.2.2 Clinical Implementation

A key challenge to future clinical implementation for the models presented in this study (and for any model used in medical practice) is verification of the accuracy and trustworthiness of the model. The need for accuracy can perhaps best be explained by amending the phrase attributed to statistician George P. Box: “All models are wrong, but some models (if accurate enough) are useful.”⁴ The need for trustworthiness stems from ethical issues surrounding the use of an inanimate algorithm in medical processes in which a human life is at stake. In order for clinicians to follow their oath to do no harm, they need to not only be confident in the ability of the tools at their disposal but also understand the inner workings of those tools in enough detail to identify when a malfunction has occurred. In general, machine learning algorithms suffer from a tradeoff between accuracy and interpretability. Due to their complexity, deep learning algorithms, despite being the most accurate class of machine learning algorithm, can be characterized as black boxes:

we see what is input into the model and what comes out, but exactly how the algorithm reaches its decision is difficult to quantify.

6.3 Future work

6.3.1 Further Validation and Incorporation of Biological Objectives for QTA

For the quantum-inspired optimization algorithm featured in Chapter 3, future studies should seek to incorporate additional disease locations and treatment modalities (such as brachytherapy) for further validation of the model's capabilities. In addition, it would be interesting to evaluate the model's performance capabilities using more complex, biological optimization objectives combining imaging and molecular biomarkers with dose-response functions derived via multiple outcome and utility modeling methods.^{5,6} Finally, if QTA is found to be robust to a variety of different radiotherapy applications, it would also be worth exploring its utility as an optimizer for training machine learning algorithms.

6.3.2 Identifying Problem-spaces Tailored to Quantum-based Frameworks

One challenge in developing quantum-based algorithms is to identify problem spaces in which a particular quantum algorithm will achieve higher performance than a classical counterpart.⁷ In Chapter 4, we found that when applied to the task of predicting patient state changes across treatment changes, the quantum-based model had slightly lower performance than the Markov-based model. The quantum model design in this study assumed that the patient states evolved under a stationary, real Hamiltonian and that no measurements were made on the state during the course of treatment. (Note that the assumption that no measurements were made on the patient's "system state" during treatment was made because changes in structure volume were not explicitly measured during treatment from CBCT data.) There are a few interesting modifications that could

be made to this framework which could result in advanced performance for the quantum model. The first would be to take into account potential measurements on the patient system, which would result in a change or “collapse” of the quantum wave state. Such measurements could be defined as instances in which the patient undergoes additional computed tomography (CT) scanning partway through their treatment (a phenomenon which was noted to have occurred for many patients in our dataset, likely to assess their anatomical response to treatment). In addition, at the time of implementation, the deep learning library utilized (PyTorch) was not yet fully able to handle complex values within the backpropagation function. This meant that we were required to define our Hamiltonian as a real-valued matrix, reducing the potential solution space. Given that complex functions have been an area of active development for PyTorch as of 2020, it would be interesting in the near future to assess the performance of the quantum model when the Hamiltonian is allowed (more realistically) to contain any value in the complex plane.

6.3.3 Incorporation of PET-CT and Identification of Robust Features

In Chapter 5, we did not include positron emission tomography (PET) imaging data into our model because of a lack of available PET images in the patient dataset. However, PET radiomics features have been shown to have some prognostic value for head and neck cancer outcomes.⁸⁻¹⁰ For future studies it would be helpful to curate additional patients into our current dataset with a focus on those who have PET imaging available so that these variables can be assessed under the bioselection paradigm. Finally, given that a portion of the CT radiomics in this study came from older, lower quality scans, future studies might benefit from performing additional analysis to identify the subset of radiomic features which are robust to variations in image quality or other imaging parameters such as reconstructions settings.¹¹

6.4 References

1. Brouwer CL, Steenbakkens RJ, van der Schaaf A, et al. Selection of head and neck cancer patients for adaptive radiotherapy to decrease xerostomia. (1879-0887 (Electronic)).
2. Cui S, Tseng H-H, Pakela J, Ten Haken RK, El Naqa I. Introduction to machine and deep learning for medical physicists. *Medical physics*. 2020;47(5):e127-e147.
3. Mayo CS, Phillips M, McNutt TR, et al. Treatment data and technical process challenges for practical big data efforts in radiation oncology. *Medical physics*. 2018;45(10):e793-e810.
4. Wasserstein R. George Box: a model statistician. *Significance*. 2010;7(3):134-135.
5. Luo Y, McShan DL, Matuszak MM, et al. A multiobjective Bayesian networks approach for joint prediction of tumor local control and radiation pneumonitis in nonsmall-cell lung cancer (NSCLC) for response-adapted radiotherapy. *Medical physics*. 2018;45(8):3980-3995.
6. Tseng H-H, Luo Y, Cui S, Chien J-T, Ten Haken RK, Naqa IE. Deep reinforcement learning for automated radiation adaptation in lung cancer. *Medical physics*. 2017;44(12):6690-6705.
7. Nielsen MA, Chuang IL. Quantum Computation and Quantum Information: 10th Anniversary Edition. In: Cambridge: Cambridge University Press; 2010: <http://csis.pace.edu/ctappert/cs837-18spring/QC-textbook.pdf>. Accessed 1/29/2019.
8. Guezennec C, Robin P, Orhac F, et al. Prognostic value of textural indices extracted from pretherapeutic 18-F FDG-PET/CT in head and neck squamous cell carcinoma. *Head Neck*. 2019;41(2):495-502.
9. Ulrich EJ, Menda Y, Boles Ponto LL, et al. FLT PET Radiomics for Response Prediction to Chemoradiation Therapy in Head and Neck Squamous Cell Cancer. *Tomography*. 2019;5(1):161-169.
10. Bogowicz M, Riesterer O, Stark LS, et al. Comparison of PET and CT radiomics for prediction of local tumor control in head and neck squamous cell carcinoma. *Acta Oncol*. 2017;56(11):1531-1536.
11. Wei L, Cui C, Xu J, Kaza R, El Naqa I, Dewaraja YK. Tumor response prediction in 90Y radioembolization with PET-based radiomics features and absorbed dose metrics. *EJNMMI Physics*. 2020;7(1):74.