

**A METHOD FOR ONLINE DOSE PREDICTION FOR IMAGE GUIDED
PROTON THERAPY**

A Thesis
Presented to
The Academic Faculty

By

William Andrew LePain

In Partial Fulfillment
of the Requirements for the Degree
Masters of Science in the
School of Mechanical Engineering
Department of Nuclear & Radiological Engineering and Medical Physics

Georgia Institute of Technology

August 2021

© William Andrew LePain 2021

A METHOD FOR ONLINE DOSE PREDICTION FOR IMAGE GUIDED PROTON THERAPY

Thesis committee:

Dr. Stella Flampouri
Emory Proton Therapy Center
Emory University

Dr. C-K Chris Wang
Nuclear & Radiological Engineering and
Medical Physics
Georgia Institute of Technology

Dr. Tianye Niu
Nuclear & Radiological Engineering and
Medical Physics
Georgia Institute of Technology

Date approved: July 27, 2021

All have their worth and each contributes to the worth of the others.

-J. R. R. Tolkien, The Silmarillion

For those affected by Glenn Coe Campbell's struggle with cancer. Especially his wife:

Margaret O'Neal Campbell

ACKNOWLEDGMENTS

To my Committee: Thank you for being patient on this wild ride through COVID and the delays that it brought. Dr. Stella Flampouri, I couldn't have asked for a better advisor/mentor. Without your input and willingness to help I would never have made it this far. You were always my advocate when trying to coordinate the bureaucracies of GT and Emory during this time of lockdown. Dr. Tianye Niu, always available when I needed to ask questions despite the 12 hour time difference. Dr. Chris Wang, likely the most gifted teacher I'll ever have both in and out of the classroom. Few have such a mastery of the art of making the complicated seem simple, and even fewer the art of teaching with a test.

My lovely wife, Madison LePain: You have been supportive and understanding throughout this entire process, and I hope that I can do the same for you when the time comes for your thesis. You are more than I could ever have asked for.

My parents, Anthony & Kathleen LePain: Your unwavering support through all my insane endeavors has been invaluable.

My friends and fellow students, Clara Glassman, Duncan Bohannon, & John DeMoor: Going through this process together has made the frustration and delays bearable.

TABLE OF CONTENTS

Acknowledgments	v
List of Tables	viii
List of Figures	ix
List of Acronyms	x
Summary	xii
Chapter 1: Introduction	1
1.1 Background	1
1.1.1 Computed Tomography	1
1.1.2 Relevant Physics of Proton Therapy	4
1.2 Motivation	8
1.2.1 Dose Verification	8
1.2.2 CBCT proton dose calculations	8
Chapter 2: Methods	11
2.1 Input Data Set	11
2.2 Dataflow	12
2.3 Dose Stretching	14

2.4	Proof-of-Principle	15
2.5	CBCT Corrections	16
2.5.1	Histogram Matching	16
2.5.2	Smoothing	17
2.5.3	Application of the RSP Curve	17
2.6	Dose Smoothing	17
2.7	Evaluation	17
Chapter 3: Results		19
3.1	Overview	19
3.2	Smoothing Strength	21
3.3	Sensitivity to Variation	22
3.4	Gamma Evaluations per Field	23
Chapter 4: Discussion		24
4.1	Visualization	24
4.2	CBCT Corrections	25
4.2.1	Application of the RSP Curve	25
4.2.2	Histogram Matching Improvements	26
4.2.3	Gaussian Smoothing	27
4.3	Dose Smoothing	28
Chapter 5: Conclusion		29
References		30

LIST OF TABLES

2.1	Input Images and Dose Distrobutions	11
3.1	Dose Accuracy Evaluations	20
3.2	Comparitave Significance	20
3.3	Smoothing Strength	21
3.4	Global Gamma Passrates per Field	23
4.1	Smoothing Effects	27

LIST OF FIGURES

1.1	Bragg Peak	4
1.2	Spread Out Bragg Peak	5
1.3	RSP conversion curve	6
1.4	Response to Range Changes	7
1.5	Effects of Histogram Matching	9
2.1	Data Pipeline	12
2.2	Algorithm Overview	13
2.3	Dose Stretching	15
3.1	Smoothing Strength	21
3.2	Gamma Passrates vs Time Since computed tomography (CT) Simulation . .	22
4.1	Dose Visualization	24
4.2	CTV DVH before and after RSP curve application	25
4.3	Before and After Dose Smoothing	28

LIST OF ACRONYMS

BEV	beam's eye view
CBCT	cone beam CT
cHU	corrected hounsfield units
CT	computed tomography
CTV	clinical target volume
DICOM	Digital Imaging and Communications in Medicine
DVH	dose volume histogram
EEGE	exponential edge gradient effect
EPTC	Emory Proton Therapy Center
GPU	graphics processing unit
H&N	head and neck
HU	hounsfield unit
IGRT	image guided radiotherapy
LET	Linear Energy Transfer
OAR	organ at risk
PET	positron emission tomography
QACT	quality assurance CT
RefCBCT	reference CBCT
ROI	region of interest
RSP	relative stopping power
SecCBCT	secondary CBCT
SecCT	secondary CT

SIB simultaneous integrated boost

SOBP spread out bragg peak

TPCT treatment planning CT

TPS treatment planning system

SUMMARY

This project was an attempt to design a deterministic algorithm for dose prediction using cone beam CTs (CBCTs) that could output a prediction on the time-scale of patient setup during proton image guided radiotherapy (IGRT). We hypothesized that a comparative analysis of a reference CBCT (RefCBCT) matching the anatomy of a treatment planning CT (TPCT) and a secondary CBCT (SecCBCT) taken as part of an IGRT setup sequence could be used to consistently and accurately predict the daily dose distribution using the original treatment plan as a guide. In the interest of simplicity and speed of prediction, only a small number of simple corrections were applied to the CBCTs to determine their effectiveness, and in an attempt to minimize the number of corrections necessary for accuracy. The corrections used included histogram matching both to the TPCT and the most recent quality assurance CT (QACT), the standard hounsfield unit (HU) to relative stopping power (RSP) curve used in the treatment planning system (TPS), as well as a simple gaussian smoothing to lessen the local effects of CBCT artifacts.

The comparison itself involved interpolating each line dose along the beam path from the cumulative corrected hounsfield units (cHU) RefCBCT to the cumulative cHU SecCBCT. After all corrections were applied on a sample set of 21 oropharynx patients with 4-5 QACTs each making up 103 prediction sets, an average global gamma passrate of 96.75 % was achieved with an average prediction time of 21.89 s. Predictions were compared to dose distributions calculated by the TPS on deformed CTs matching the anatomy of the SecCBCT used in the prediction process. Gamma values were determined using parameters of 3 %/3 mm with a dose cutoff of 2 Gy, or one fraction.

CHAPTER 1

INTRODUCTION

1.1 Background

1.1.1 Computed Tomography

As an x -ray beam pass through some material, photons are removed from it in proportion to the amount of photons currently in the beam.

$$I \propto \frac{dI}{dx} \quad (1.1)$$

This basic relation gives rise to the attenuation equation.

$$\frac{I}{I_0} = e^{-\mu x} \quad (1.2)$$

Where μ is the attenuation coefficient, x is the depth in the material from the surface, and I/I_0 is the intensity of the beam relative to the surface. The attenuation coefficient in turn can be related to the electron density (ρ_e) within whatever medium the beam is passing through[1][2]. The attenuation coefficient in this form is an amalgamation of many different interaction coefficients describing many different types of interactions. The two important categories of interactions for the purposes of this work are catastrophic interactions where a photon is removed entirely from the beam and scattering interactions where a photon changes energy or direction.

For medical imaging the attenuation coefficient is often expressed in terms of the hounsfield unit (HU), defined below[1].

$$HU = 1000\left(\frac{\mu}{\mu_{water}} - 1\right) \quad (1.3)$$

The primary reason for this definition is for data storage purposes over the range of attenuation coefficients normally found in medical imaging. Each voxel within a computed tomography (CT) image is associated with an average HU value. These values are obtained through a few different algorithms, that take the known intensity of an x -ray beam and compare it to measured intensities after the beam has passed through the patient.

The two types of CTs important for this work are fan-beam or standard CTs, and cone beam CTs (CBCTs). A fan beam CT, as its name implies, uses a fan shaped geometry for the beam and makes multiple passes around the patient where each pass roughly corresponds to an image slice. A CBCT on the other hand uses a cone shaped beam so that the third dimension is covered by the beam instead of taken sequentially. The largest benefit of the standard fan-beam CT method is that it avoids scatter along that third dimension, and therefore can provide a much less noisy and far more accurate representation of the HU values within the imaged object. A CBCT on the other hand is much faster, and requires less mechanical support as the detector and x -ray generator only need to rotate once. These distinctions make the standard CT (nearly) perfect for applications where exact HU values are necessary, for instance when attempting to measure electron density, and the CBCT perfect for applications that just need a general geometry check.

Cone beam CTs have a number of artifacts that make their image quality less than ideal.

Scattered radiation in all its forms is the bane of radiation physicists everywhere. The same holds true for medical imaging. A scattered photon in the center of the beam of a CBCT has a high chance of making it to the detector, and adding to the measured beam intensity at some other location. This artificially decreases the calculated attenuation through all voxels within whatever projection passes through the detector pixel in question. As the area of the detector increases, the larger the probability a scattered photon will be detected, and the image as a whole shows greater noise and overall less attenuation. Since there is no simple way to tell what photons have been scattered and what photons have been unaffected in the medium, the mathematics of image reconstruction assume that all interactions

are catastrophic.

The energy spectrum of a standard x-ray tube is generated using bremsstrahlung interactions. This gives a somewhat broad energy distribution, and due to the fact that lower energy photons are attenuated more strongly the energy distribution of the beam changes as it passes through an object. The imaging system expects a more consistent energy distribution, and therefore has a tendency to under-estimate the amount of attenuation in areas behind highly attenuating objects. A thick object, such as a torso, can also show this effect as a darkening in the center of the image. This effect is known as “beam hardening.”

Many implants used in dentistry or other reconstructive/replacement applications are made of highly attenuating material. In addition to the high amount of scatter these objects introduce, they also heavily decrease beam intensities. Beam intensity can suffer behind these types of objects so much that the detector is unable to measure the final intensity. This can be thought of as an extreme case of beam hardening, and is known as extinction. In areas where the boundary between one of these heavily absorbing objects and low attenuating material exists there tends to be bright and dark streaks that fan out around these boundaries. Partial volume averaging is to blame here as the attenuated and non-attenuated sections of the beam are averaged across one pixel after another. This effect is known as the exponential edge gradient effect (EEGE).

Rings may appear in CBCTs around the center of the image if a single detector pixel is dead or otherwise mis-calibrated. They are most notable in areas of homogenous composition, but can appear anywhere.

Aliasing is an artifact that looks like rays diverging from the center of the image, and is a consequence of under-sampling of the spatial domain within the periphery of the image. Due to the fact that a CBCT uses a diverging source, the spatial frequency detectable by the imaging system changes along the beam axis. This often leads to a violation of the Nyquist sampling theorem when structures with a high spatial frequency pass close by the detector. This is most likely to happen near the image boundaries where beam divergence is more

pronounced than in the center.

For a more detailed discussion of CBCT artifacts and their causes see Schulze *et al.*[3].

1.1.2 Relevant Physics of Proton Therapy

One of the most important equations in the field of proton therapy is the Bethe-Bloch stopping power formula:

$$\left(\frac{dT}{dx}\right)_{coll} = 307.1 \frac{\rho Z z^2}{A \beta^2} \left[\ln \left(\frac{1.022 \text{ MeV} \beta^2}{I(1 - \beta^2)} \right) - 2\beta^2 \right] \frac{\text{keV}}{\text{cm}} \quad (1.4)$$

Where Z/A is the electrons/amu of the target material, z is the charge of the particle, and β is that particle's relativistic velocity (v/c). The variable I is the mean excitation potential of the medium respectively. Mean excitation potential is the geometric-mean value of all ionization and excitation potentials of the medium the charged particle finds itself in.

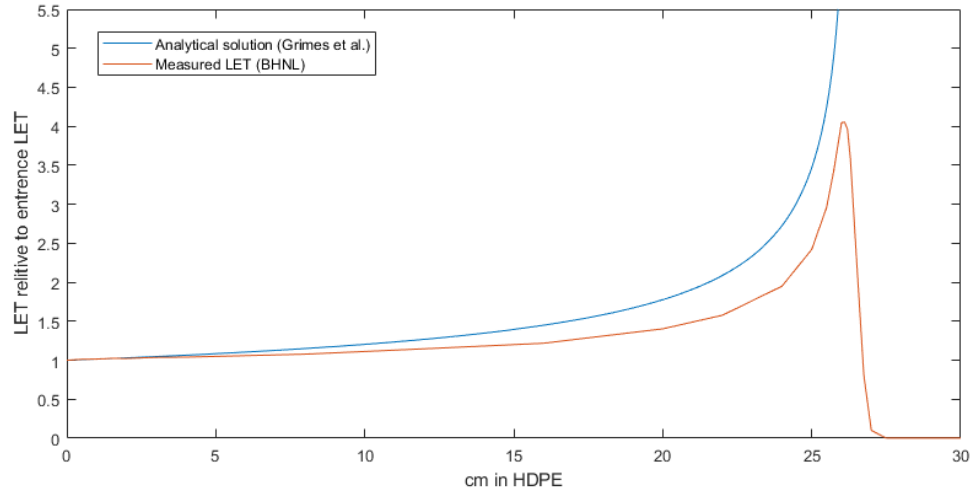


Figure 1.1: Bragg Peak

Calculated Linear Energy Transfer (LET) from a stopping power approximation by Grimes *et al.* plotted against a visualization of LET data for a 205 MeV proton in high density polyethylene taken from the Brookhaven National Laboratory website. Differences at the end of the particle track between the two plots are primarily due to the fact that the calculated stopping power begins to exceed the kinetic energy of the particle.

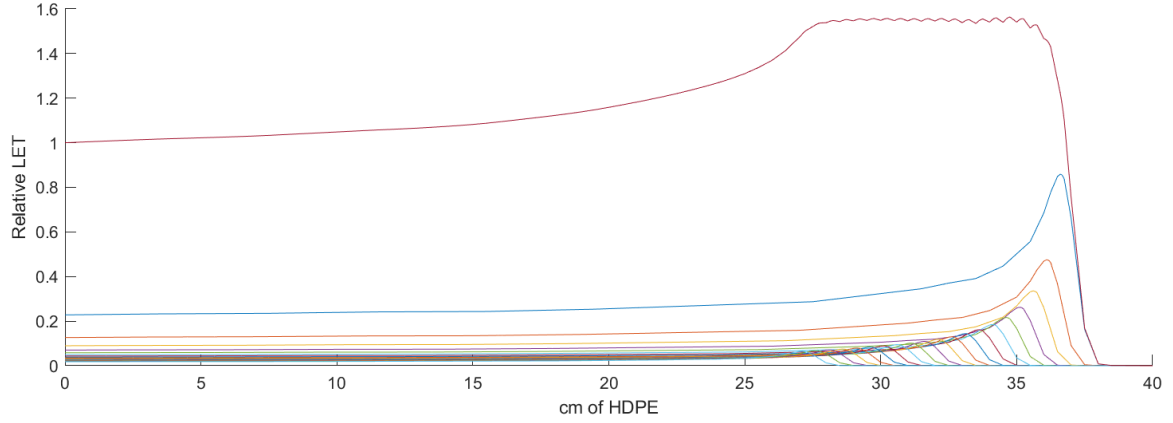


Figure 1.2: Spread Out Bragg Peak

spread out bragg peak (SOBP) constructed from multiple simulated bragg peaks using LET data from Figure 1.1. The SOBP is the curve in red, and is constructed by summing contributions from all other curves in the figure.

This equation describes the aggregate effect of all collisional interactions on a charged particle traveling through any medium. These interactions may be atomic excitations, ionization events, or even the production of δ -rays, a form or secondarily generated electrons with enough energy to cause ionization events on their own. Regardless of the form it takes, the energy removed from the traveling charged particle eventually makes its way into the surrounding medium where it is dissipated. This energy deposition and its controlled distribution within the body is the primary goal of any radiation therapy modality.

Importantly this formulation shows us three important dependencies for stopping power. There is a linear dependency on electron density ($\rho Z/A$) of the medium. The atomic composition of the stopping medium determines how tightly these electrons are bound, and therefore determine the mean excitation potential (I). And the kinetic energy of the charged particle ($T \propto \beta^2$) determines how strongly the particle interacts with its surroundings. The first two remain constant throughout the particle's path within a single material as they are entirely dependent only on the properties of the stopping medium. The last dependency shows us that as the particle loses energy the stopping power acting on it must also change, and owing to the fact that the β^2 is in the denominator we can also see that as the energy of the particle diminishes the strength of its interactions increase.

This positive feedback gives rise to a catastrophic deposition of energy at the end of the particle's path, a phenomenon known as the bragg peak. Seen in Figure 1.1, the bragg peak is the fundamental building block of proton therapy. Energy modulation can be used to stretch the bragg peak out creating a plateau known as a SOBP shown in Figure 1.2. As discussed above two of the primary dependancies of Equation 1.4 are electron density,




Figure 1.3: RSP conversion curve
Curve used by RayStation to convert HU to relative stopping power (RSP).

and electron binding energy. It is rather simple to define the electron density of a medium using attenuation coefficients (subsection 1.1.1) so long as the x -ray beam has well defined energy. The mean excitation potential, however, is not directly measurable in every voxel of a CT. As fate would have it, most tissues within the body with similar densities have similar atomic compositions. This means that, at least approximately, there is a one-to-one relationship between HU and mean excitation potential. Given this, we can construct a one-to-one mapping from HU to stopping power, or as it is stored in treatment planning systems (TPSs): relative stopping power to water.




Figure 1.4: Response to Range Changes

The effects of small range changes on the order of 0.5 cm on the dose profile. One can see that the photon dose barely changes at depth, where the proton dose nearly entirely misses it's mark.

While the fundamental mechanisms of dose absorption for both photons and protons are the same, namely the fact that only charged particles deliver dose directly, the steps to get there for each method are quite different. Dose delivery for an x-ray beam is carried out by the secondarily generated charged particles, however the shape and general characteristics of the beam are governed mostly by the physics of non-charged particle interactions. Proton beams on the other hand are governed by the physics of charged particle interactions, and being a charged particle themselves are capable of directly depositing dose within a volume. This difference is important when considered from the perspective of uncertainties along the beam path. As seen in Figure 1.4, a small change in density won't have much effect on a *x*-ray beam at depth due to its relatively flat local dose distribution. Protons on the other hand are highly sensitive to changes owing to the steep dose falloff at the distal end of the beam. The irony is that the very thing that makes proton radiotherapy theoretically superior to other forms of radiotherapy is also its greatest weakness.

1.2 Motivation

1.2.1 Dose Verification

On-line and off-line methods for dose verification has been a warm topic for some time [5]. Techniques have used anything from measurements of Cherenkov radiation[6], positron emission tomography (PET) from activated β^+ emitters[7], thermally generated sound waves[8] to calculations on setup CBCTs[9]. This interest derives from the effect of range uncertainties in Proton therapy. Due to the steep falloff on the distal end of the SOBP, any minor range uncertainties can become major dose uncertainties. These uncertainties inevitably lead to planning techniques that are unable to take full advantage of the benefits of proton therapy.

Any technique that measures dose proxies (Cherenkov/PET/Ultra Sound) necessitates specially designed equipment, and therefore adds a significant amount of expense when compared to CBCT based calculations for any clinic attempting to implement them. Additionally, these methods do not allow for prediction, only verification. The advantage of CBCT dose calculations is that they provide predictive power, and therefore may be used to show an expected dose distribution based on the geometry of a particular setup before beam-on. A major obstacle here is time/computational power due to the fact that most methods using CBCTs run dose calculations from scratch.

The method proposed in this work is designed to piggy-back on previously generated dose calculations and interpolate dose from one CBCT to another thereby bypassing the computational time required for a full dose calculation.

1.2.2 CBCT proton dose calculations

A significant amount of work has gone into using image guided radiotherapy (IGRT) setup CBCTs to calculate dose[9]. In a study by Giacometti *et al.*, the authors reviewed sixty-nine studies on the topic and found that no single dose calculation method seemed to be

superior. Of the methods discussed, a subset of the “HU override” technique that used histogram matching to match a CBCT to a corresponding CT was of particular interest to us. Histogram matching is a fairly quick process, and doesn’t require any additional input in the form of clinical target volumes (CTVs), organ at risk (OAR), or any other delineated structure[10, 11]. This makes histogram matching a prime candidate for use in quick automatic dose calculations.

The purpose of a histogram matching algorithm is to take the cumulative histogram of all grayscale values within an image or region of interest and increase/decrease their values until the cumulative histogram for the image matches that of a reference image. The assumption is that the relative proportions of each type of tissue, and therefore relative proportion of HU values remains unchanged. This is not strictly true, especially for patients

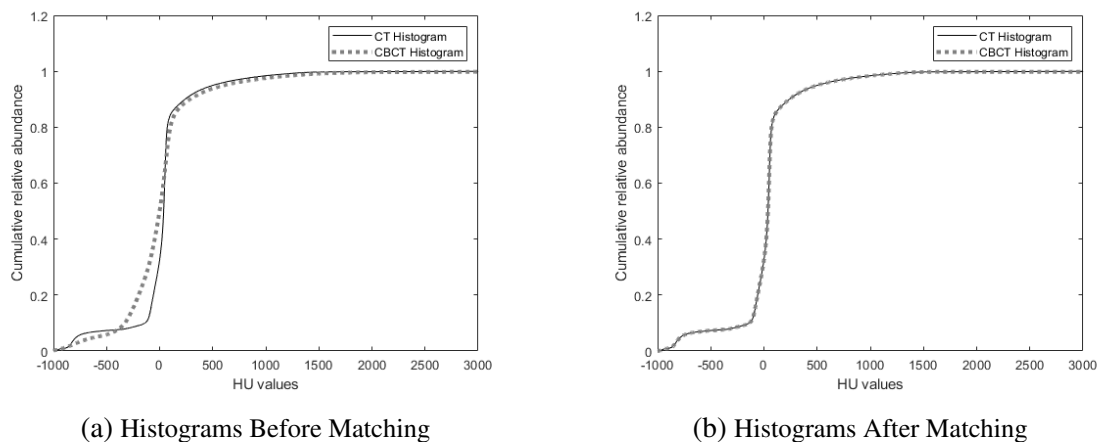


Figure 1.5: Cumulative histogram for a CT and a CBCT both before and after the matching process.

that have lost or gained significant weight. The goal however, is not to perfectly recreate a CT, but to get reasonably close. A visual representation of Histogram Matching can be seen in Figure 1.5.

One of the major problems with histogram matching a CBCT to a CT is the former’s propensity for beam hardening artifacts. Including a large volume affected by this artifact would drastically effect the relative frequencies of each HU value. For this reason, histogram matching alone can only be applied to relatively thin structures, and even then care

must be taken to ensure the regions being compared are, anatomically, as similar as possible. Head and neck (H&N) CBCTs fit this criterion provided the shoulders are avoided during histogram matching.

CHAPTER 2

METHODS

2.1 Input Data Set

All patients included in this study were treated using a 4 or 5 field simultaneous integrated boost (SIB) technique for oropharyngeal cancer at Emory Proton Therapy Center (EPTC). 14 out of 21 of the patients received a replan due to anatomical changes. Every patient received the first day’s CBCT between 10 to 23 days after CT simulation.

Table 2.1: Input Images and Dose Distributions

Name	Derived from	Deformed to	Dose	Purpose
SecCT	QACT	SecCBCT	Calculated in TPS	Evaluation
RefCBCT	first day CBCT	TPCT	Planned dose	Prediction reference
SecCBCT	CBCT of the day	–	Predicted	Dose prediction

Each patient had 4 to 5 quality assurance CTs (QACTs) throughout treatment allowing for the creation of deformed CTs that matched the anatomy of some of the CBCTs taken during setup. These deformed CTs were called secondary CTs (SecCTs), and were used to generate expected or “evaluation” dose distributions for the prediction model using the TPS.

Each prediction would need a baseline image for reference when interpolating dose onto a new CBCT. The first treatment day’s IGRT setup CBCT was deformed to match the anatomy of the treatment planning CT (TPCT) so that the dose predictor would be comparing CBCT to CBCT. This synthetic CBCT was known as the reference CBCT (RefCBCT).

The CBCTs used for prediction were referred to as secondary CBCTs. Each secondary CBCT (SecCBCT) was the setup CBCT with the minimum number of days between it and some QACT. As these were supposed to be a direct input for the algorithm they were not

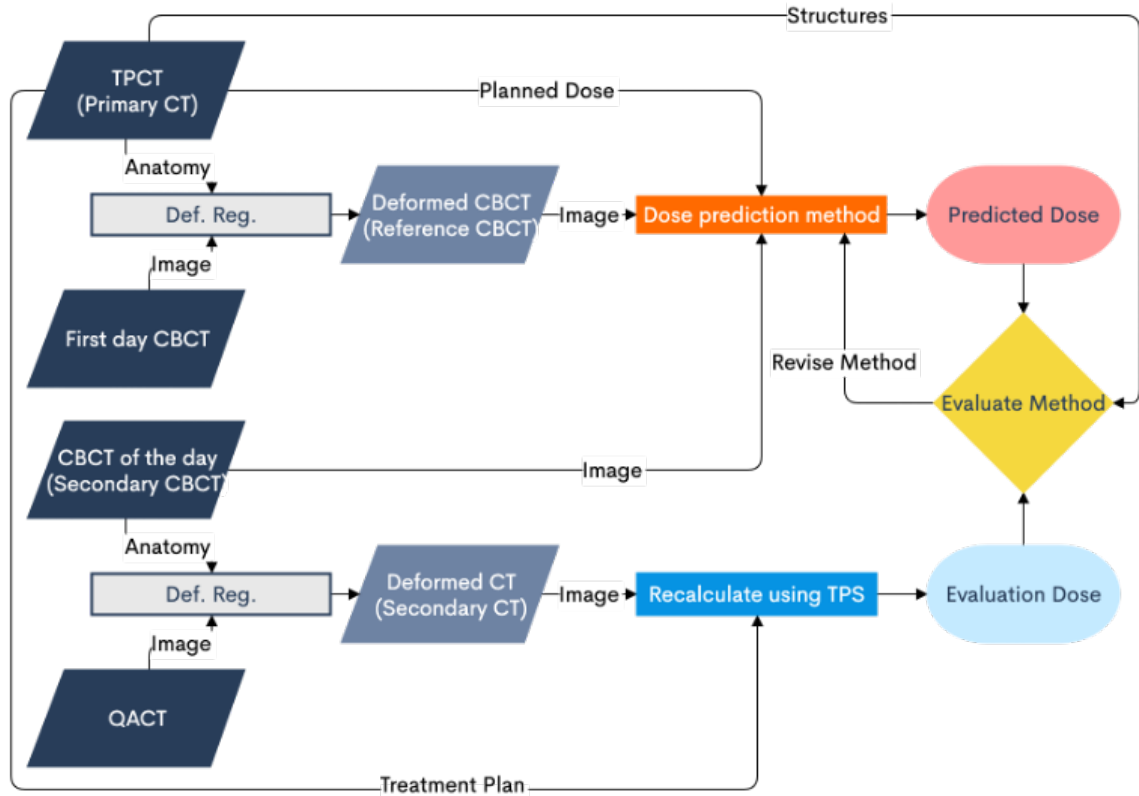


Figure 2.1: Data Pipeline

Visualization of the prediction and evaluation process. This figure shows how each piece of data was generated, but does not represent any temporal relation between them.

modified in any way before being exported from the TPS

These CTs, their purpose, and their respective dose distributions are listed in Table 2.1 for ease of reference, and their relationship to other images and dose distributions are displayed in Figure 2.1. All CTs and dose files along with the structure set generated for the TPCT were exported as Digital Imaging and Communications in Medicine (DICOM) files and imported to Matlab.

2.2 Dataflow

Each set of TPCT/RefCBCT and the corresponding dose files were loaded into a single data structure. The RefCBCT then had any of the desired correction methods applied to it; see section 2.5 & section 4.2. The values in each voxel were called corrected hounsfield

units (cHU) to distinguish them from the HU values measured in CBCT acquisition, and to avoid using the term RSP as absolute RSP is never calculated for any voxel. Within this data structure, the beam dose and a copy of the RefCBCT for each field were rotated to beam's eye view (BEV) and the reverse transformation was saved for use after prediction. The rotation was necessary to align the voxel columns with the beam direction, and avoid

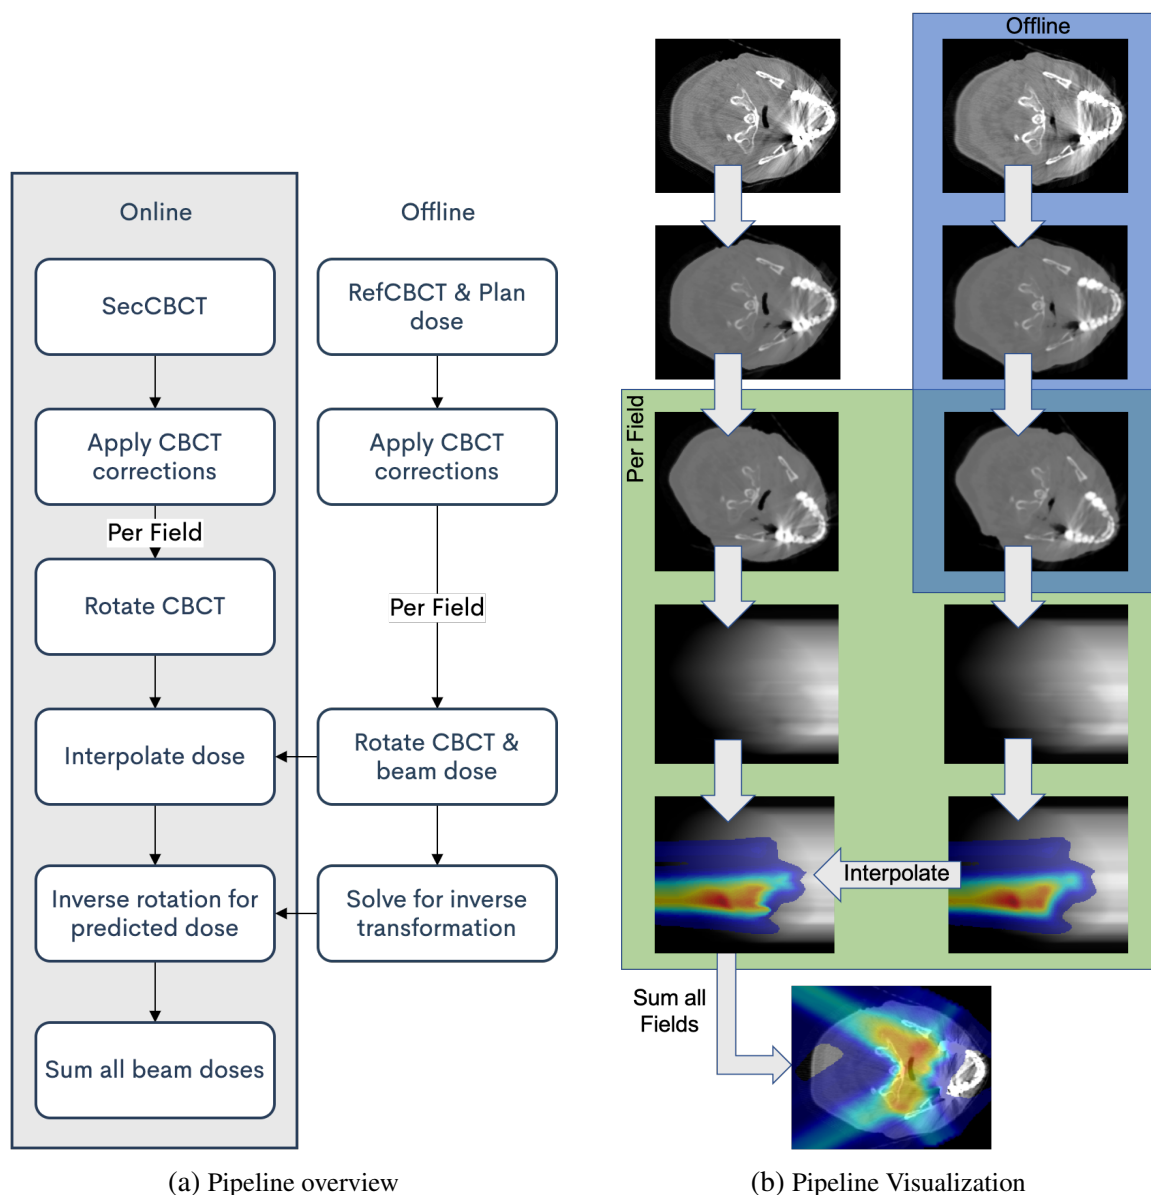


Figure 2.2: Visualization of the steps within the prediction algorithm. Any calculations that can be done offline are completed and loaded into memory before the SecCBCT is taken.

dose interpolation along oblique angles through the images. These steps were completed

offline to minimize the amount of calculations required during the online section of the algorithm.

The second input for our dose predictor was the DICOM files for the SecCBCT and it's IGRT setup sequence registration file. As this image is created with the patient on the treatment couch, the rest of the algorithm must be run against the clock.

The SecCBCT DICOMs were loaded into the algorithm, its registration file was used to place it on the same coordinate system as the TPCT, and the chosen CBCT corrections were applied; see section 2.5 & section 4.2 for correction details. Once loaded the function would then rotate the SecCBCT to BEV just as had been done with the RefCBCT and call the dose predictor described in section 2.3 for each beam. The outputs of the dose predictor were rotated back to the original coordinate system using the saved reverse transformations from the reference data set, and added together to form a predicted dose distribution. The prediction time per field as well as the total time for all corrections and the time for the entire prediction process were recorded, and are displayed in Table 3.1.

Both global and local gamma evaluations as well as dose volume histograms (DVHs) for the CTV were used as a measure of method evaluation. The pass-rates generated by the gamma evaluations were used to flag problem runs, debug/modify our method, and also to evaluate how similar our predicted plans were to the evaluation plans generated by the TPS.

2.3 Dose Stretching

After the CBCTs were processed, the first step in the prediction was to find the distal end of the beam in each voxel column. The last voxel containing dose was used as a proxy for beam end. A cumulative cHU map of the rotated RefCBCT was created for each voxel column from patient surface to the beam end, and the same was done for the SecCBCT stopping once the targeted cHU for the beam end was reached or the algorithm exited the patient. An equation representation is given below, and a more visual representation is

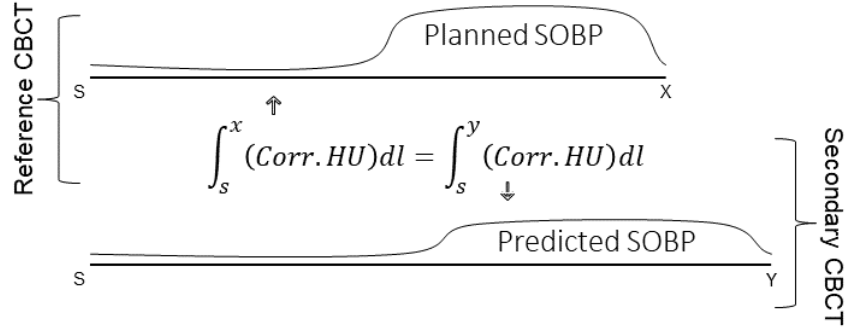


Figure 2.3: Dose Stretching

The planned SOBP from patient surface (s) to beam end (x) is stretched or contracted using interpolation on the cumulative cHU for both the RefCBCT & SecCBCT to fit within the new beam size (s to y).

given in Figure 2.3.

$$\sum_{i=s}^y cHU_i(refCBCT) = \sum_{i=s}^x cHU_i(secCBCT) \quad (2.1)$$

In both Equation 2.1 & Figure 2.3 s is the voxel corresponding to the patient surface within the column in question, and x & y are the beam ends of the RefCBCT & SecCBCT respectively.

To create the predicted SOBP in Figure 2.3 the dose from the cumulative RefCBCT was interpolated over the cumulative cHU of the SecCBCT within the defined beam. Any columns who's line integral dose was less then 0.5 Gy or fell outside the region of interest were ignored, and simply transferred, unchanged, to the prediction dose map.

2.4 Proof-of-Principle

Proof-of-principle was carried out using direct TPCT to QACT comparison for a subset of 6 patients for a total of 21 predictions. The TPCT was used as both the TPCT & RefCBCT while the QACT was used as the SecCT & SecCBCT. The only correction applied was the HU to RSP curve discussed in subsection 2.5.3. Results for the proof-of-principle run are shown in Table 3.1

2.5 CBCT Corrections

Tested permutations of the corrections below and their respective gamma passrates and prediction times are displayed in Table 3.1.

2.5.1 Histogram Matching

When employed, the histogram matching algorithm used some defined beam region of interest (ROI) on both the CT and CBCT in question to construct a cumulative histogram of HU values and modify the values of the CBCT such that its cumulative histogram matched that of the CT. Both CBCTs were matched to the TPCT. There was an attempt to match the SecCBCT to the most recent QACT, however the data-set used in this work was not conducive to this alteration. The results from that run were thrown out, and it will be investigated as an option in future work.

A few different ROIs were chosen for the generation of the cumulative histograms employed during the matching process. Using the entirety of the body was the simplest approach, and this ROI mask was constructed using edge detection. This ROI was not tested directly, however it was used as the base structure on which to build more complicated ROIs. One of these more complicated masks was the intersection of the body mask and the union of each beam mask; in Table 3.1 this is the base histogram correction. Due to the decrease in accuracy of the histogram matching process discussed further in subsection 4.2.2 it was determined that more complicated ROIs were needed, and the masks labeled as “sans shoulders (-sh)” and “sans Extreme HU’ (-ex)” were created. The first was the previous beam path mask with the first 55 slices removed in an attempt to remove the shoulders from the histogram ROI. The second was a more dynamic mask involving the removal of voxels from the Beam Path ROI if its HU value in the CBCT was outside the range -700 to 2000 . The purpose of this HU exclusion was to partially remove the effects of artifacts from the histogram matching process. Results for this section are displayed in

Table 3.1.

2.5.2 Smoothing

A gaussian smoothing technique was also employed on some iterations to smooth out the contributions of CBCT artefacts. A 3D smoothing function (smooth3 in Matlab) with a kernel with a radius of 7 mm and a standard deviation of 0.75 mm was used initially for both the RefCBCT and the SecCBCT. Various kernel radii and standard deviations were tested; results shown in Table 3.3.

2.5.3 Application of the RSP Curve

Conversion from HU to RSP was done through lookup table with a resolution of integer HU values. The RSP conversion curve from RayStation was used, as it is the TPS used at EPTC.

2.6 Dose Smoothing

After the most accurate method for CBCT correction was determined, it was realized that the dose prediction method used thus far was operating on a $1 \times 1 \times 1mm$ voxel size, but the TPS calculates dose over a voxel size of $3 \times 3 \times 3mm$. This realization lead to the application of gaussian smoothing on the per-field dose distributions before they were summed together with a kernal radius of 3 voxels and a standard deviation of 4 voxels.

2.7 Evaluation

Each prediction run was compared to a corresponding evaluation dose map generated using deformable registration of a QACT to the anatomy of the SecCBCT in question. Local and global 3D gamma evaluations were performed at 3%/3 mm and a dose cutoff of one fraction (2 Gy). A DVH was also constructed for the CTV. The CTV was lifted directly, undeformed, from the original treatment plan as if this method is to be used in a clinic there

will be no time to deform the target. A short investigation was done to see if the DVH could be used as a proxy for a gamma analysis, but the primary function of the DVH evaluation was to quickly visualize when there was a problem with the prediction method.

CHAPTER 3

RESULTS

3.1 Overview

Correction labels are as follows:

- S = Gaussian Smoothing
- H = Histogram Matching
- R = RSP curve
- -ex = Ignore “extreme” HU values during histogram matching
- -sh = Ignore shoulder region during histogram matching
- -sm = Dose is smoothed per beam before gamma analysis

When more than one is applied they are applied in the order listed above.

For the 21 patients in this study, the best dose interpolation was the one that used all three correction methods while removing extreme HU values outside the bounds -700 to 2000 giving a global gamma passrate of 96.75% with a standard deviation of 1.29% after dose smoothing. Table 3.1 shows how this, and other correction permutations fared against the proof-of-principal run discussed in section 2.4. Table 3.2 shows both the paired T-test and Wilcoxon signed-rank test significance values for each permutation of corrections vs every other. This is intended to show if a correction had a significant effect on the accuracy of the prediction method.

Table 3.1: Dose Accuracy Evaluations

Gamma Passrates (%)			Prediction Time (s)		
Corrections	Local	Global	Per field	Correctioin	Total
Proof of Princ.	94.69 \pm 1.35	97.79 \pm 0.79	4.49 \pm 0.42	0.92 \pm 0.18	22.57 \pm 2.09
None	52.59 \pm 3.89	59.06 \pm 3.57	4.47 \pm 0.73	0.4 \pm 0.13	22.29 \pm 3.92
R	92.16 \pm 1.82	95.85 \pm 1.39	4.34 \pm 0.74	0.84 \pm 0.24	21.69 \pm 3.96
RH	89.46 \pm 4.40	93.94 \pm 3.11	4.38 \pm 0.72	2.67 \pm 0.71	21.86 \pm 3.89
RS	92.35 \pm 1.83	96.01 \pm 1.40	4.41 \pm 0.79	1.24 \pm 0.31	22.05 \pm 4.23
RH-ex	92.62 \pm 1.76	96.29 \pm 1.32	4.37 \pm 0.73	3.06 \pm 0.88	21.85 \pm 3.92
RH-sh	88.65 \pm 5.56	93.45 \pm 4.00	4.4 \pm 0.72	2.69 \pm 0.66	21.96 \pm 3.91
RHS	92.07 \pm 1.95	95.8 \pm 1.57	4.4 \pm 0.75	3.36 \pm 1.08	21.99 \pm 4.03
RHS-ex	92.74 \pm 1.78	96.4 \pm 1.31	4.38 \pm 0.73	3.52 \pm 1.02	21.89 \pm 3.92
RHS-ex-sm	93.72 \pm 1.74	96.75 \pm 1.29	4.38 \pm 0.73	3.52 \pm 1.02	21.89 \pm 3.92

Results presented as mean \pm standard deviation.

Table 3.2: Comparitave Significance


Corrections	R	RH	RS	RH-ex	RH-sh	RHS	RHS-ex	RHS-ex-sm	T-test
R		0.000	0.000	0.000	0.000	0.508	0.000	0.000	R
RH	0.000		0.000	0.000	0.031	0.000	0.000	0.000	RH
RS	0.000	0.000		0.000	0.000	0.008	0.000	0.000	RS
RH-ex	0.000	0.000	0.000		0.000	0.000	0.000	0.000	RH-ex
RH-sh	0.000	0.003	0.000	0.000		0.000	0.000	0.000	RH-sh
RHS	0.711	0.000	0.007	0.000	0.000		0.000	0.000	RHS
RHS-ex	0.000	0.000	0.000	0.000	0.000	0.000		0.000	RHS-ex
RHS-ex-sm	0.000	0.000	0.000	0.000	0.000	0.000	0.000		RHS-ex-sm
WSR-test	R	RH	RS	RH-ex	RH-sh	RHS	RHS-ex	RHS-ex-sm	Corrections

Significance as calculated with global gamma passrates by a T-test (top right) and a Wilcoxon signed-rank (WSR) test (bottom left) for each correction permutation vs every other permutation. Highlighted in red are the cells that correspond to a test failure (significance value $> .01$). Significance values below 10^{-3} are truncated for ease of viewing.


3.2 Smoothing Strength

Table 3.3: Smoothing Strength


Smoothing std (mm)	Local Passrate	Global Passrate
0	$89.46 \pm 4.40\%$	$93.94 \pm 3.11\%$
0.5	$92.00 \pm 2.15\%$	$95.77 \pm 1.69\%$
0.75	$92.07 \pm 1.95\%$	$95.8 \pm 1.57\%$
1	$92.05 \pm 1.95\%$	$95.77 \pm 1.60\%$
1.5	$92.10 \pm 1.96\%$	$95.77 \pm 1.62\%$
2	$92.10 \pm 1.95\%$	$95.76 \pm 1.62\%$
2.5	$92.08 \pm 1.97\%$	$95.73 \pm 1.64\%$
3	$92.07 \pm 1.98\%$	$95.71 \pm 1.66\%$



(a) Global passrates



(b) Local passrates



(c) Local and Global Passrates

Figure 3.1: Smoothing Strength

The effects of increased gaussian smoothing on the gamma passrates of the RHS correction method.

3.3 Sensitivity to Variation

The relationship between time since CT simulation and gamma passrate is given in Figure 3.2, and was used as a proxy for sensitivity of the script to patient anatomy variation. The relationship exists, but appears to be very weak.

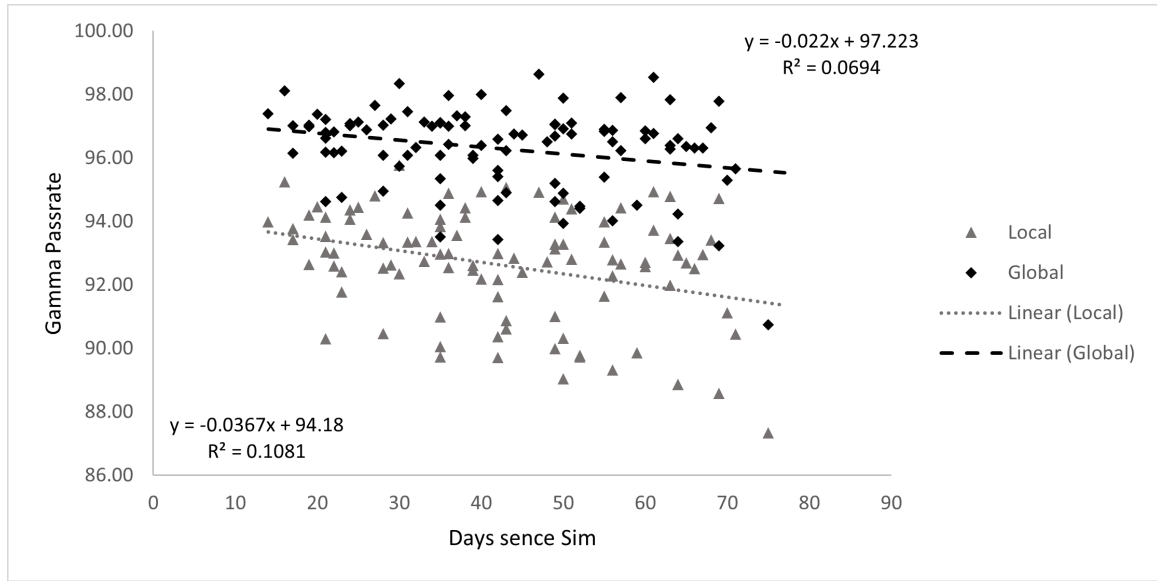


Figure 3.2: Gamma Passrates vs Time Since CT Simulation

Gamma Passrates for each prediction using the RHS-ex-sm method plotted against the number of days since CT simulation. The equations in the top right and bottom left are linear fits for the Global and Local passrates respectively.

3.4 Gamma Evaluations per Field

A single patient was selected to verify the prediction process on a per field basis. Passrates generated using the same gamma parameters (3 %/3 mm), but with a dose cutoff equal to the beam weight times 2 Gy

Table 3.4: Global Gamma Passrates per Field

Beam #	Secondary CBCT				
	A	B	C	D	E
1	97.58%	97.08%	96.52%	93.52%	98.10%
2	95.63%	96.06%	96.14%	95.22%	96.35%
3	96.23%	97.15%	95.48%	96.88%	96.98%
4	96.82%	95.62%	94.70%	97.00%	97.12%
5	97.58%	95.93%	96.13%	97.21%	98.01%
Average	96.77%	96.37%	95.80%	95.97%	97.31%

Global gamma passrates for each predicted field for all 5 SecCBCTs for a single patient using the correction method RHS-ex.

CHAPTER 4

DISCUSSION

4.1 Visualization

Of the methods tested, the RHS-ex-sm method proved the most accurate. In Figure 4.1 one can see one selected slice from one of the patients in the study this illustrates where this method succeeds and fails.

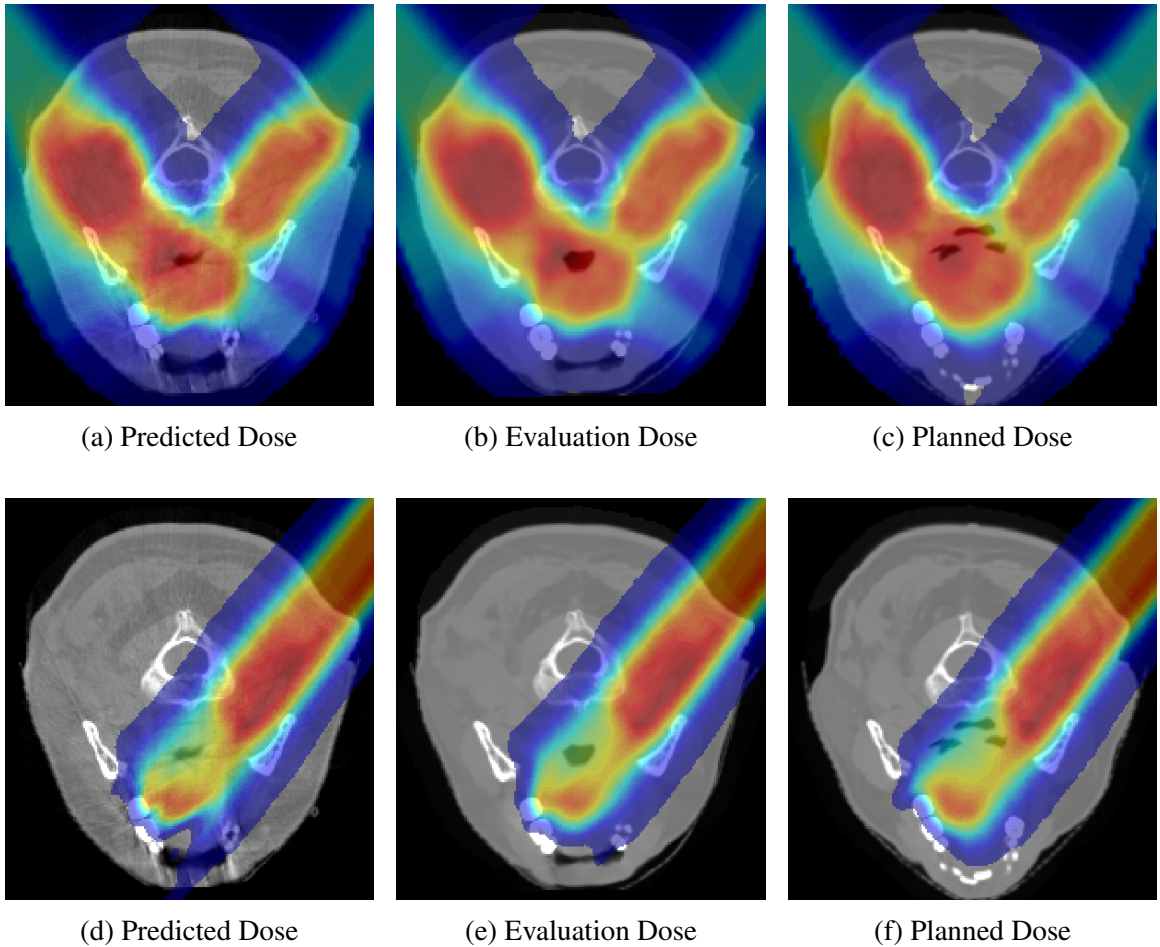


Figure 4.1: Dose Visualization

Slice # 153 from one of the patients in the study. Selected for illustration purposes. Top row (a-c) is the full plan, and the bottom row (d-f) is a single field. Each row has the dose color scale normalized to the maximum point dose among all three images in either the entire distribution or the single field respectively.

When comparing Figure 4.1d and Figure 4.1e the first thing that may pop out is the remarkable similarity in the large scale structure of the high dose region (orange/red) near the end of the beam path. The differences are more pronounced on a smaller scale, and indeed this was to be expected. On a more local scale it should be expected that the shape of the dose distribution would be more dependent on artifacts, but on a larger scale those effects begin to cancel out. The same effect is also expected due to the imperfections in deformable registration. The large scale structure should be less affected by local imperfections due to either process. For this reason, global passrates were deemed more indicative of plan similarity than local passrates.

4.2 CBCT Corrections

4.2.1 Application of the RSP Curve




Figure 4.2: CTV DVH before and after RSP curve application

DVHs displayed are calculated using the CTV from the treatment plan. These structures remained unchanged when applied to the evaluation and prediction doses, and therefore are not representative of actual dose on the structure. This is still useful as a qualitative demonstration of the importance of the RSP curve.

One fairly standard correction for dose calculation is the RSP curve. Since this method is looking at the density differences between images instead of the absolute density one may be forgiven for assuming it is reasonable to omit the HU to RSP conversion. This is not what was found however. The non-linear region seen between $-500HU$ and $500HU$

in Figure 1.3 implies that it's not that simple, and indeed there was a drastic improvement to the prediction accuracy when the RSP curve was applied. This is most easily illustrated by looking at the DVH before and after RSP curve application shown in Figure 4.2. The RSP curve was by far the largest improvement to the prediction model, and because of this it is applied in addition to every other correction explored here.

4.2.2 Histogram Matching Improvements

Scatter and beam hardening artifacts discussed in subsection 1.1.1 can have drastic effects on the actual HU values calculated for a CBCT. It seemed necessary to ensure that the RSP conversion had the most correct values to work with. As discussed in subsection 1.2.2, a histogram matching algorithm was the optimal choice for this, and was used to match the CBCTs to the TPCT.

The naive approach of histogram matching the entirety of the beam paths within each patient actually lead to a decrease in accuracy when compared with just the HU to RSP curve. This was likely due to the inclusion of some of the shoulder region and the heavily artifacted slices containing the jaw. It was decided that the simplest way to correct this problem was to exclude regions of the CBCT containing extreme HU values from the histogram construction. As can be seen in Table 3.1 this lead to a drastic improvement to the histogram matching process.

It was expected that by ignoring the first 55 slices of the CBCTs containing the majority of the shoulders the prediction results could be improved. Not only was this not the case, but the prediction accuracy decreased when ignoring the shoulder region. It is suspected that this is due to a few factors. The first is that the choice of slice #55 was too rigid, and did not account for differences between patients. This should only increase the standard deviation while having only a minimal effect on the average, so something else must be at play. The decrease in average accuracy is suspected to be due to the artifacts elsewhere in the CBCT, like those due to the teeth, that would make up a larger proportion of the

included voxels and therefore hold greater sway over the outcome of the histogram matching process after the removal of the shoulders. This was just speculation however, and the development of a more dynamic process for determining the shoulder boundary and other artifact detections is likely to shed some light on the exact process at play here.

4.2.3 Gaussian Smoothing

Artifact correction in a CBCT is a somewhat involved process taking 4 min in some cases even with graphics processing unit (GPU) acceleration[12]. This is far too slow for an algorithm attempting to make a dose prediction in under 1 min. Artifacts, however may still be a problem especially on a local scale. Luckily all of the CBCTs that will be used are taken on the same patient using very similar if not identical machines. It was suspected that this meant most artifacts would appear near enough to each other that a simple gaussian smoothing of the images would suffice as a rudimentary artifact “corrector.”

A smoothing standard deviation of 0.75 mm was chosen due to the local maximum in Figure 3.1 and its low standard deviation in Table 3.3 This seems to be only a minor choice, as the absolute effect of the smoothing strength is fairly low. The lower spread in the output data was deemed more important than the absolute improvements in local passrates given by higher smoothing standard deviations.

Table 4.1: Smoothing Effects


Label	Global (%)	Stdev	Improvement (s)	T-test	WSR-test
R	95.85	1.39	0.16	0.00	0.00
RS	96.01	1.40			
RH	93.94	3.11	1.86	0.00	0.00
RHS	95.80	1.57			
RH-ex	96.29	1.32	0.11	0.00	0.00
RHS-ex	96.40	1.31			

Selected data from Table 3.1 & Table 3.2 for ease of reference.

Smoothing in general proved to be effective, but only provide a minor improvement. From Table 4.1 we can see that the absolute global accuracy improvement is fairly small, but statistically significant.

4.3 Dose Smoothing

While not applied to every correction method, It is clear that smoothing the predicted dose distributions such that their resolution is more in-line with the calculation resolution of the TPS brings the predictions more in-line with the evaluation doses. Smoothing the dose of the most accurate prediction method (RHS-ex) improved the global passrates from 96.40% to 96.75%, but it is unclear if this would be a useful improvement in a clinical setting. In



(a) Unsmoothed Dose

(b) Smoothed Dose

Figure 4.3: Before and After Dose Smoothing

Slice # 120 selected at random from among the RHS-ex predictions before and after dose smoothing.

Figure 4.3 we can see that the dose smoothing has an effect, but it is unclear if this effect would be aiding in a visual evaluation or just artificially increasing the prediction passrates. Visually the effect of this smoothing is easiest to notice outside the body, and therefore is unlikely to have any effect on the decision to treat.

CHAPTER 5

CONCLUSION

The goal of this work was to produce a method for quick online dose prediction that could be easily implemented in a proton clinic. With prediction times below 30 s, and an expected global accuracy of $96.75\% \pm 1.29\%$, this method is just on the boundary of clinical usefulness.

The speed could easily be improved, as everything was written in Matlab without any parallel computing tools. A non-interpreted language optimized to use multithreaded CPUs or GPU acceleration would hopefully reduce the average prediction time to below 15 s. These improvements would also allow for more complicated CBCT corrections, such as bow-tie filter modeling for beam hardening corrections, or multilayered histogram matching without exceeding 30 s prediction times.

A combination of optimization and the addition of additional correction methods should push this tool/process into the realm of clinical usefulness. The interpolation of dose from one CBCT to another has promise as a useful clinical tool for both therapists and physicists, but the method as presented thus far only lays the foundation for future work.

REFERENCES

- [1] U. Schneider, E. Pedroni, and A. Lomax, “The calibration of CT hounsfield units for radiotherapy treatment planning,” *Physics in Medicine and Biology*, vol. 41, no. 1, pp. 111–124, Jan. 1996.
- [2] A. Sudhyadhom, “On the molecular relationship between hounsfield unit (hu), mass density, and electron density in computed tomography (ct),” *PLOS ONE*, vol. 15, no. 12, pp. 1–15, Dec. 2021.
- [3] R. Schulze *et al.*, “Artefacts in cbct: A review,” *Dentomaxillofacial Radiology*, vol. 40, no. 5, pp. 265–273, 2011, PMID: 21697151. eprint: <https://doi.org/10.1259/dmfr/30642039>.
- [4] D. R. Grimes, D. R. Warren, and M. Partridge, “An approximate analytical solution of the bethe equation for charged particles in the radiotherapeutic energy range,” *Scientific Reports*, vol. 7, no. 1, p. 9781, 2017.
- [5] A.-C. Knopf and A. Lomax, “In vivo proton range verification: A review,” *Phys Med Biol*, vol. 58, no. 15, R131–60, Aug. 2013.
- [6] T. Masuda *et al.*, “Measurement of nuclear reaction cross sections by using cherenkov radiation toward high-precision proton therapy,” *Scientific Reports*, vol. 8, no. 1, p. 2570, 2018.
- [7] X. Zhu and G. El Fakhri, “Proton therapy verification with pet imaging,” *Theranostics*, vol. 3, no. 10, pp. 731–740, Sep. 2013.
- [8] K. C. Jones *et al.*, “Acoustic-based proton range verification in heterogeneous tissue: Simulation studies,” *Physics in Medicine & Biology*, vol. 63, no. 2, p. 025 018, Jan. 2018.
- [9] V. Giacometti, A. R. Hounsell, and C. K. McGarry, “A review of dose calculation approaches with cone beam ct in photon and proton therapy,” *Physica Medica*, vol. 76, pp. 243–276, 2020.
- [10] Y. Onozato *et al.*, “Evaluation of on-board kv cone beam computed tomography–based dose calculation with deformable image registration using hounsfield unit modifications,” *International Journal of Radiation Oncology · Biology · Physics*, vol. 89, no. 2, pp. 416–423, 2014.
- [11] K. Arai *et al.*, “Feasibility of cbct-based proton dose calculation using a histogram-matching algorithm in proton beam therapy,” *Physica Medica*, vol. 33, pp. 68–76, 2017.

- [12] A. Sisniega *et al.*, “High-fidelity artifact correction for cone-beam CT imaging of the brain,” *Physics in Medicine and Biology*, vol. 60, no. 4, pp. 1415–1439, Jan. 2015.
- [13] F. H. Attix, *Introduction to Radiological Physics and Radiation Dosimetry*. Wiley-VCH, 2016.