

Incorporation of Biological Factors in Radiation Therapy Treatment Planning



Mark Brooke

Wolfson College
University of Oxford

Supervisors:

Prof Frank Van den Heuvel

Prof Maria Hawkins

Dr Sam Vinko

A thesis submitted for the degree of

Doctor of Philosophy

Michaelmas 2020



CANCER
RESEARCH
UK

OXFORD
CENTRE

Declaration

I, Mark Brooke declare that the work presented in this thesis is my own, with the exception of the parts listed below, unless stated otherwise in-text, and is in accordance with the University of Oxford's regulations for the degree of Doctor of Philosophy.

- **Chapter 5:** the beam line experiment at the Clatterbridge Cancer Centre (CCC) in Wirral, UK was undertaken in collaboration with the Cockcroft Institute at the University of Liverpool, led by PhD student Jacinta Yap, and Advacam (Prague, Czech Republic) on 28 November 2019. In attendance: Jactina Yap (Cockcroft), Mark Brooke (Oxford), Navrit Bal (Nikhef, Amsterdam, Netherlands), Andrzej Kacperek (CCC) and Roland Schneurer (Cockcroft). Cristina Oancea (Advacam) provided initial and necessary processing of data collected with the Advacam detector using proprietary software, and provided MATLAB scripts to assist with post-processing. My roles: experimental design, equipment setup, data recording and analysis.
- **Chapter 5, Figure 5.1 and Figure 5.2:** reproduced figures, with permission, from the PhD thesis (in production) of Jacinta Yap (Cockcroft Institute, University of Liverpool).
- **Chapter 5, Figure 5.3 and Figure 5.8:** reproduced figures (Fig. 1 and Fig. 2, respectively), with permission, from Granja *et al.*, *Nuclear Inst. and Methods in Physics Research, A* **911** (2018) 142–152.

- **Chapter 6:** the inverse planning algorithm was mathematically developed in collaboration with Yair Censor (University of Haifa, Israel) and Aviv Gibali (ORT Braude College, Israel). Equal contributions were made to the writing of the mathematical paper that resulted [Brooke, Censor and Gibali. (In press). *International Transactions in Operational Research*, <https://doi.org/10.1111/itor.12929>]. Advice on medical physics context was provided from a wider collaboration with Reinhard Schulte (Loma Linda University, USA) and Scott Penfold (University of Adelaide, Australia).
- **Chapter 7, Page 204:** Calculations of normal tissue complication probabilities (NTCP) were performed by feeding dose-volume histogram (DVH) data into a *Python* script written by Pratik Samant (Oxford).
- **Appendix C, Figure C.1:** the Landau fitting procedure was performed using a *Python* script provided by Frank Van den Heuvel (Oxford).

Mark Brooke

13 December 2020

Acknowledgements

This DPhil has been an exciting adventure from start to finish, which began with a 16,000 kilometre move from back home in Adelaide, Australia to Oxford, England. I have learned a great deal these past few years, and while much of it has been consolidated in this thesis, a significant portion of it has been in my own self-growth and in building professional networks and personal relationships. There are many people I'd like to thank for their help in forming this body of work and for keeping me sane, motivated and happy throughout the process.

Firstly, I am immensely grateful to my supervisors, Frank Van den Heuvel, Maria Hawkins and Sam Vinko for their encouragement, support and guidance. They have been sources of knowledge and experience while also providing clarity in the direction of this thesis and helping me build the necessary skills to be successful in my future work. I extend many thanks as well to Suliana Teoh, Francesca Fiorini, Anna Vella, Tom Whyntie, Pratik Samant, Ben George, Amy Elliott, Mark Hill and Muhammad Kasim for stimulating discussions and for providing a fun and supportive work environment.

My external collaborators, Yair Censor, Aviv Gibali, Scott Penfold, Reinhard Schulte, and Jacinta Yap, have been instrumental in offering insights into a number of principles in mathematics, accelerator physics and clinical medical physics. Proper acknowledgement of their ongoing support deserves much more than this short note. I am grateful and excited to work together on more projects in the future.

Thank you to my family for always supporting my pursuits. I'm sorry this one had

to be so far from home... Thank you to my friends here, back home and around the world. And thank you, Molly, for being the most loving, supportive and comforting partner anyone could wish for. I couldn't imagine achieving this without you.



This work was supported by a John Monash Scholarship, a Clarendon Scholarship and Cancer Research UK, grant number C2195/A25197, through a CRUK Oxford Centre DPhil Prize Studentship.

Publications and presentations

Journal Articles

Dynamic string-averaging CQ-methods for the split feasibility problem with percentage violation constraints arising in radiation therapy treatment planning.

M. Brooke, Y. Censor, & A. Gibali. (In press). International Transactions in Operational Research. <https://doi.org/10.1111/itor.12929>

Incorporating oxygenation levels in analytical DNA-damage models – Quantifying the oxygen fixation mechanism.

F. Van den Heuvel, A. Vella, F. Fiorini, M. Brooke, M.A. Hill, T. Maughan. In submission.

Conference Presentations

Quantification of DNA damage and variable RBE in hadron therapy

M. Brooke & F. Van den Heuvel

2018 IEEE Nuclear Science Symposium and Medical Imaging Conference, International Convention Centre, Sydney, Australia, 13 November 2018

On quenching of GafChromic film response in proton therapy

F. Van den Heuvel, M. Brooke & F. Fiorini

PTCOG 58, Manchester Central Convention Complex, Manchester, United Kingdom, 10 June 2019

Investigating the mechanism behind FLASH radiation effects

F. Van den Heuvel & M. Brooke

First International Biophysics Collaboration Meeting, GSI, Darmstadt, Germany, 21
May 2019

A probabilistic Poisson model to determine RBE changes facilitates LET-based optimization in Intensity Modulated Proton Treatment

M. Brooke, F. Van den Heuvel & A. Vella

Blue Ribbon Award ePoster presentation

American Association of Physicists in Medicine (AAPM) Virtual Meeting, Online, 12
July 2020

A flexible projection-based non-convex inverse planning algorithm for intensity-modulated proton therapy

M. Brooke, Y. Censor, A. Gibali, S. Penfold, R. Schulte & F. Van den Heuvel

Blue Ribbon Award ePoster presentation

American Association of Physicists in Medicine (AAPM) Virtual Meeting, Online, 12
July 2020

Oxygen Related Factors in FLASH Radiotherapy

F. Van Den Heuvel, K. Petersson, B. Vojnovic, M. Hill, A. Vella, A. Ryan, M. Brooke,
T.S. Maughan & A.J. Giaccia

American Society for Radiation Oncology (ASTRO) 2020, 26 October 2020

*International Journal of Radiation Oncology*Biography*Physics*, 108(3, Supplement),
e534. <https://doi.org/10.1016/j.ijrobp.2020.07.1668>

FLEPAC: A FLASH therapy treatment planning system for electrons, protons, Helium and Carbon ions

M. Brooke, F. Fiorini, M. Hill, A. Kiltie, J-L. Ruan, A. Ryan, F. Van den Heuvel, A. Vella & B. Vojnovic

European Society for Radiation Oncology (ESTRO) 2020, Online, 30 November 2020

Abstract

Radiation therapy plays a crucial role in the treatment of cancer. Major developments in the field, including 3D tomographic imaging, intensity modulation and novel particle modalities, have all contributed to iterative improvements in precise dose delivery. These advances have enabled highly conformal tumouricidal dose deposition with improved healthy tissue sparing. Proton therapy, in particular, has become an increasingly adopted modality due to its favourable depth-dose profile, allowing for much of the dose to be localised within the target region, with virtually no dose leakage to tissues beyond the target. However, physical dose does not tell the full story; radiobiological mechanisms and their intermodality differences must be well-understood, as these explain the processes by which the energy deposited in cells leads to cell death and, ultimately, to clinical outcomes.

Densely ionising particle tracks offer an enhanced cell killing efficiency over x-rays, which may be quantified through the relative biological effectiveness (RBE). The RBE is a multiplicative factor converting the proton dose to an x-ray dose equivalent that yields the same biological endpoint. In standard clinical practice a constant RBE of 1.1 is used, assuming protons are 10% more effective than x-rays for the same dose, though recently there has been an increased focus on more sophisticated, variable RBE models. Many decades of clinical experience with x-rays enables RBE-weighted proton dose planning under a familiar framework. The dose is shaped according to both absolute and volume-based tolerance doses. Although, in many modern treatment planning systems, dose-volume constraints (DVCs) are approximated as

their exact formulation poses a non-convex optimisation problem.

This thesis proposes novel methods and recommendations for the inclusion of radiobiological factors in treatment planning through (1) a variable but pragmatic RBE model based on DNA double strand break induction, and (2) a flexible, projection-based inverse planning algorithm, suited to non-convex settings, that comprehensively addresses dose-volume effects through the exact modeling of DVCs. This author hypothesises that inclusion of these strategies will enable confident and effective treatment planning not only in terms of physical dose deposition but, crucially, in biological effect.

List of Acronyms

ASIC application-specific integrated circuit

BER base excision repair

BP base pair

CAD computer aided design

CCC Clatterbridge Cancer Centre

CHEF contour-clamped homogeneous electric field

CHO Chinese hamster ovary

CPU central processing unit

CRT conformal radiotherapy

CSDA continuous slowing down approximation

CT computed tomography

CTV clinical target volume

CVaR conditional value-at-risk

DDR DNA damage response

DNA deoxyribonucleic acid

DSA dynamic string averaged

DSB double strand break

DVC dose volume constraint

DVH dose-volume histogram

EQD equivalent uniform dose

ESTRO European Society for Radiation Oncology

EUD equivalent uniform dose

FDG fluorodeoxyglucose

FMO fluence map optimisation

FPS frames per second

FSU functional subunit

GPU graphics processing unit

GTV gross tumour volume

HDC hard dose constraint

HR homologous recombination

ICRP International Commission on Radiological Protection

ICRU International Commission on Radiation Units and Measures

IMFP inelastic mean free path

IMPT intensity modulated proton therapy

IMRT intensity modulated radiation therapy

IPOPT interior point optimisation

IR irradiation

JIT just-in-time (compilation)

LEM local effect model

LET linear energy transfer

LINAC linear accelerator

LQ linear-quadratic

MC Monte Carlo

MCDS Monte Carlo Damage Simulation

MCO multi-criteria optimisation

MCS multiple Coulomb scattering

MIP mixed integer programming

MKM microdosimetric kinetic model

MLC multi-leaf collimator

MOSCFPP multiple operator split common fixed point problem

MPV most probable value

MRI magnetic resonance imaging

MT multi-threading

NER nucleotide excision repair

NHEJ non-homologous end-joining	QUANTEC Quantitative Analyses of Normal Tissue Effects in the Clinic
NHS National Health Service (UK)	
NIST National Institute of Standards and Technology	
NSD nominal standard dose	RBE relative biological effect
NTCP normal tissue complication probability	RMF repair-misrepair-fixation
	ROI region of interest
OAR organ-at-risk	
OER oxygen enhancement ratio	SARP self-adaptive relaxation parameter (method)
OFH oxygen fixation hypothesis	SF surviving fraction
	SOBP spread-out Bragg peak
PBP pristine Bragg peak	SSB single strand break
PBT proton beam therapy	STL stereolithography
PDF probability density function	
PDL Poisson-distributed lesions	TCP tumour control probability
PENTEC Pediatric Normal Tissue Effects in the Clinic	TDRA theory of dual radiation action
PET positron emission tomography	TNM Tumour, node and metastasis
PFGE pulsed-field gel electrophoresis	TOPAS Tool for Particle Simulation
PIDE Particle Irradiation Data Ensemble	TV total variation
PLL potentially lethal lesion	
PMMA polymethyl methacrylate	VMAT volumetric modulated arc therapy
PTV planning target volume	
PVC percentage violation constraint	WET water equivalent thickness

Contents

1	Introduction	1
1.1	Modern radiation therapy: from x-rays to protons	1
1.2	Proton therapy physics	5
1.2.1	Photon interactions with matter	5
1.2.2	Proton interactions with matter	8
1.2.2.1	Inelastic Coulomb interaction	9
1.2.2.2	Elastic Coulomb scattering	10
1.2.2.3	Non-elastic nuclear collisions	11
1.2.3	Multiple scattering	12
1.2.4	Stopping power	14
1.2.5	Depth-dose characteristics	17
1.2.6	Linear energy transfer	19
1.3	Proton radiobiology	23
1.3.1	Temporal action of ionising radiation	23
1.3.2	DNA damage and repair	24
1.3.3	The linear quadratic model for cell survival	29
1.3.4	Fractionated radiotherapy	31
1.3.5	Relative biological effectiveness	34
1.4	Intensity modulated radiation therapy	37
1.4.1	Proton pencil beam scanning	40
1.5	Radiobiological considerations in treatment planning	41

1.5.1	Dose-volume effects	42
1.5.2	Levels of radiobiological optimisation	44
1.6	Scope of this thesis	46
2	Overview of methods	48
2.1	A double strand break induction model for RBE prediction	48
2.1.1	Theoretical development	48
2.1.2	Monte Carlo simulations	49
2.1.3	Implementation in a treatment planning system	49
2.1.4	Comparisons with in-vitro data	49
2.1.5	Beam line experiment at the Clatterbridge Cancer Centre	50
2.2	A flexible projection-based inverse planning algorithm for radiation therapy	50
2.2.1	Theoretical development	50
2.2.2	Practical implementation	51
2.2.3	Testing on synthetic and clinical treatment plans	51
3	A DNA double strand break induction model for RBE prediction	52
3.1	Introduction	52
3.1.1	Importance of the RBE endpoint	53
3.2	Overview of existing RBE models	56
3.2.1	Wilkens and Oelfke (2004) phenomenological model	58
3.2.2	Carabe <i>et al.</i> (2007) phenomenological model	59
3.2.3	Chen and Ahmad (2011) phenomenological model	60
3.2.4	Wedenberg <i>at al.</i> (2013) phenomenological model	61
3.2.5	McNamara <i>et al.</i> (2015) phenomenological model	62
3.2.6	Mairani <i>et al.</i> (2017) phenomenological model	62
3.2.7	Microdosimetric-kinetic model	63

3.2.8	Local effect model	64
3.2.9	Repair-misrepair-fixation model	66
3.2.10	Van den Heuvel (2014) DNA damage model	68
3.3	A new DNA double strand break induction model	69
3.3.1	Inelastic mean free path calculation	71
3.3.2	Parameterising the damage model	74
3.3.3	Relative biological effectiveness	77
3.4	Validation of DSB-based RBE models	78
3.4.1	Detection of double strand breaks	79
3.4.1.1	Immunofluorescent staining	80
3.4.1.2	Pulsed-field gel electrophoresis	81
3.4.1.3	Neutral comet assay	83
3.4.2	Clonogenic survival assay	83
3.4.2.1	Degeneracy of the linear quadratic parameters	88
3.5	Discussion	90
4	Monte Carlo simulation studies and comparisons with <i>in vitro</i> data	100
4.1	Introduction	100
4.2	Materials and Methods	101
4.2.1	Monte Carlo studies	101
4.2.1.1	Creating a TOPAS extension for scoring DSB damage	101
4.2.1.2	RBE calculations in pristine proton pencil beams	103
4.2.1.3	RBE calculations in the SOBP region	104
4.2.2	RBE calculation in a treatment planning system	105
4.2.3	Comparisons with <i>in vitro</i> experimental data	107
4.2.3.1	The Particle Irradiation Data Ensemble	107
4.2.3.2	Calculation of RBE	108
4.3	Results	110

4.3.1	Monte Carlo studies	110
4.3.2	RBE calculation in a treatment planning system	114
4.3.3	Comparisons with <i>in vitro</i> experimental data	117
4.4	Discussion	120
5	Clinical beam line experiment at the Clatterbridge Cancer Centre	130
5.1	Introduction	130
5.2	Methods and materials	132
5.2.1	The Advacam MiniPIX-Timepix detector	132
5.2.2	Beam line measurements	134
5.2.3	Calculation of LET spectra	138
5.2.4	Simulated Monte Carlo spectra	140
5.3	Results	146
5.4	Discussion	152
6	A flexible projection-based inverse planning algorithm for radiation therapy	161
6.1	Introduction	161
6.1.1	Mathematical challenges with dose-volume criteria	162
6.2	Methods	169
6.2.1	Hard dose constraints (HDCs)	170
6.2.2	Dose volume constraints (DVCs)	171
6.2.3	Problem statement: the non-convex split feasibility problem (SFP)	171
6.2.4	Split-feasibility seeking methods for solving the SFP	173
6.2.5	Further insight into Byrne's CQ formula	177
6.2.6	Dynamic self-adaptive relaxation parameter methods	178
6.3	Discussion	181

7	Implementation and testing of a novel inverse planning algorithm	185
7.1	Introduction	185
7.2	Methods	186
7.2.1	Two-dimensional numerical examples	188
7.2.1.1	Test Case 1	188
7.2.1.2	Test Case 2	189
7.2.2	Simplified base-of-skull chordoma	190
7.2.3	The TROTS proton therapy data set	192
7.2.4	Code acceleration for clinical efficiency	193
7.3	Results	194
7.4	Discussion	199
8	Concluding remarks	209
A	Derivation of the interaction length density function	212
B	C++ scripts for TOPAS extensions	217
C	Stable distribution fits to MiniPIX LET measurements	225
D	Note on convergence for the DSACQ algorithm using SARP	227
E	Dose difference forest plots for all TROTS plans	229
F	Extensions to the SARP method	234
F.1	Threshold condition	234
F.2	Prioritising structures	234
	References	238

List of Figures

1.1	Schematics showing the three major mechanisms by which photons interact with matter in the context of radiotherapy; (a) photoelectric interaction, (b) Compton scattering, and (c) pair production. e^- represents an electron while e^+ represents a positron.	6
1.2	Illustration of proton interaction mechanisms pertinent to radiation therapy: (a) Inelastic Coulomb scattering (if the impact parameter, b , is of the order of the atomic radius, a , then “knock-on” electrons can result), (b) elastic Coulomb scattering with atomic nuclei, and (c) non-elastic collisions. e , p , n and γ represent an electron, proton, neutron and photon, respectively.	9
1.3	Classical scattering in three dimensions. Charged particles with mass m , kinetic energy E , speed v_0 and initial momentum $p_0 = mv_0$ are deviated by a spherical potential whose origin is on the z -axis. Particles within an infinitesimal patch of area $d\sigma = b db d\phi$ at azimuthal angle ϕ with impact parameter b scatter into a corresponding infinitesimal solid angle $d\Omega = \sin(\theta)d\theta d\phi$ where θ is the polar scatter angle.	12
1.4	Proton stopping power contributions (left axis) and range (right axis) in water as functions of kinetic energy. S_{el} : electronic stopping power; S_{nuc} : nuclear stopping power; R_{CSDA} : CSDA range; $R = \alpha E^p$: range approximation by Bragg-Kleeman rule [29] as in (1.12). Data sourced from the NIST pSTAR database [54].	16

1.5	Illustrative comparison of a proton depth-dose curve to that of a conventional megavoltage (MV) photon beam. Individual pristine proton pencil beams produce the Bragg peaks shown which together contribute to a homogeneous spread out Bragg peak (SOBP) covering the target region.	18
1.6	Illustration of low and high LET radiation at the scale of the DNA double-helix, and the corresponding ionisation density pattern. Black lines represent particle tracks. Each ionisation event is marked in red.	20
1.7	Dose, LET_t and LET_d depth profiles along the central axis of a 100 MeV monoenergetic proton pencil beam.	22
1.8	Illustration of DNA strand breaks (in orange) caused by indirect and direct damage events, with each ionisation event marked in red. A primary particle is incident from the left, giving rise to a secondary electron. Indirect damage occurs by ionisation of nearby water molecules which results in hydroxyl radicals ($OH\cdot$) then ionising the DNA. Direct damage occurs by ionisation of the DNA by the incident particle itself. Multiple strand breaks within 10 base pairs are labelled double strand breaks (DSBs), which may be due to a single particle track (bottom track shown) or multiple.	26
1.9	(a) Examples of a linear quadratic (LQ) cell survival curves for early and late responding tissues. The radiosensitivity of the tissue is characterised by α/β , measured in units of Gy, which is the dose at which the linear and quadratic components of the curve are equal. (b) Examples of cell survival trends for low and high LET radiation, and a fractionated low LET treatment regimen. The linear component, or initial shoulder, to the curve is shown as a dashed line of the same colour. . .	32

1.10	The number of publications discussing variable RBE has noticeably increased over the past two decades. Based on search of data available on sciencedirect.com on 17 September 2020 with the following search strings: ("relative biological effectiveness" OR "RBE") AND (proton OR protons) AND ("variable" OR "biophysical" OR "nonconstant" OR "non-constant" OR "changing").	36
1.11	Simplified illustration of active pencil beam scanning for IMPT dose delivery. Proton pencil beams of a nominal energy pass through a range shifter before being modulated by orthogonal electromagnets, such that the resulting Bragg peak is positioned at the required location in the target.	41
3.1	Left: geometric setup for possible interaction lengths of a linear particle trajectory through a cylindrical section representing the sensitive region containing the DNA, in which an ionisation event may cause lasting DNA damage. Right: calculated and Monte Carlo generated interaction length distribution for $d = 2.37$ nm and $2H = 6.8$ nm.	70
3.2	IMFP in DNA as calculated by (3.33) with parameters listed in Table 3.2 for protons and electrons, both with and without a plasmonic correction. Parameters were fitted to experimental data by Tan <i>et al.</i> [235] (protons) and NIST [236, 237] (electrons).	72
3.3	Proton collisional stopping power in liquid water, shown as a function of kinetic energy between 1 keV and 10 GeV. Stopping power is maximised at approximately 90 keV, which is also the energy at which the IMFP in Figure 3.2(a) is minimised. Plotted data was collected from the <i>NIST pSTAR</i> database [54].	73

3.4	The DSB induction model, (3.31), fit through a single “conversion efficiency” parameter, ε (listed in Table 3.3), to relative Monte-Carlo-generated DSB counts using the MCDS damage simulation software [159, 160]. Fits are shown for (a) protons and (b) electrons. The colour scale corresponds to the interaction length probability density, f_L	75
3.5	DSB yield, as calculated by (3.34), shown with and without the plasmonic correction term in (3.33) as a function of kinetic energy for (a) protons and (b) electrons. Monte-Carlo-generated DSB counts using MCDS 3.1 [159, 160] are also shown, and were used to fit the parameters listed in Table 3.3.	76
3.6	Degeneracy in LQ parameters, α and β , for curves that predict cell survival within 10% ($\delta = 0.1$, black points), 5% ($\delta = 0.05$, blue points) and 2% ($\delta = 0.02$, red points) agreement. Data points were generated using a Monte Carlo sampling method. $\Delta\alpha$ and $\Delta\beta$ represent the change in α and β , respectively, between two equally valid survival curves.	89
3.7	Qualitative representation of the addition of two pencil beam energy spectra in overlapping voxels. $\phi_1(E)$ is the fluence spectrum due to pencil beam 1 in the highlighted voxel and $\phi_2(E)$ is the spectrum for pencil beam 2. The total spectrum in the voxel of interest is $\Phi = \phi_1 + \phi_2$. Gray shaded region represents an arbitrary target volume.	93
3.8	Left: predicted DSB yield using MCDS [159, 160] for different oxygen concentrations (0%, 5%, 10%, 20%, 40% and 60%). Right: the oxygen enhancement ratio (OER) as a function of proton energy, calculated as the ratio of the predicted DSB yield at a specified oxygen concentration to that at zero oxygen partial pressure. .	95

3.9	(a) Relative SSB induction likelihood, according to (3.50), with fitting parameter $\epsilon = 1.128$ and 95% confidence interval (C.I.) of (1.068, 1.188), fit to Monte-Carlo-generated SSB counts using MCDS version 3.1 [159, 160]. (b) Most likely SSB induction response shown as a function of proton energy, against the simulated MCDS data.	98
4.1	RBE _{cd} , calculated using (3.35), along the central axis of pristine proton pencil beams in water between 60 MeV and 120 MeV, shown alongside the relative dose deposited by each beam.	110
4.2	Visualised dose and RBE _{cd} -weighted dose scored in TOPAS [161] for a pristine 60 MeV proton pencil beam in water.	111
4.3	Summary of RBE _{cd} calculations shown in Figure 4.1, showing a decrease in RBE, both at the Bragg peak and at the 95% dose fall-off, as the beam energy is increased.	112
4.4	RBE _{cd} , calculated using (3.35), shown along with the dose-averaged and track-averaged LET, as scored in TOPAS along the central axis, for pristine proton pencil beams between 60 MeV and 120 MeV in water.	113
4.5	Relationship between RBE _{cd} , calculated using (3.35), and LET from Monte Carlo (MC) simulations. Also shown is the fit of (4.11) to the 550 data points with both dose-averaged and track-averaged LET. See Table 4.2 for the fitting parameters.	114
4.6	(a, b, c): RBE _{cd} , calculated using (3.35), is shown at 1 mm intervals along the central axis of three SOBPs at 60 MeV, 100 MeV and 180 MeV maximum energy, respectively, for $n = 5, 10$ and 20 energy intervals forming the SOBP according to (4.5). (d, e, f): Variation in the RBE _{cd} calculation in the middle 50% of the SOBP for 60 MeV, 100 MeV and 180 MeV, respectively. The darker pink shaded region represents the 68% confidence interval in the mean RBE _{cd} while the lighter pink region represents the 95% confidence interval.	115

4.7	(a) dose map, at isocentre, for a head-and-neck plan distributed with the <i>matrad</i> treatment planning software [163]. (b) RBE_{cd} map corresponding to the dose map in (a), generated by modifying the LET calculation subroutine. The gross and planning target volumes, GTV and PTVs respectively, and spinal cord contours are shown for reference.	116
4.8	Histogram showing the RBE_{cd} value distribution in the pixels of the dose map slice shown in Figure 4.7.	117
4.9	Percentage difference between predicted RBE and the empirical RBE for surviving fractions between 5% and 95%. Predicted RBE is according to the six phenomenological models summarised in Table 3.1. Each point represents the RBE prediction error in a single experiment from the PIDE database [164], while the red line indicates the average error across all experiments. Blue shaded regions (from darkest to lightest) give envelopes of one, two and three standard deviations from the mean.	118
4.10	(a) Comparison of RBE predictions to experimental data provided by the PIDE database [164]. Experimental RBE_{SF} across surviving fractions of 5% to 95% are shown as box plots (black). Predictions from two models in this work, $RBE_{SF}+f(LET_d)$ (yellow) and $RBE_{SF}+PDL$ (green), given by (4.6)-(4.7) and (4.8)-(4.9) respectively, are shown along with the models of McNamara <i>et al.</i> [190] and Mairani <i>et al.</i> [191]. RBE_{cd} predictions from the $f(L_\Delta)$ envelope in (4.11) are overlaid. (b) An increasing trend of α values with LET is shown against the prediction from this work.	119
4.11	Schematic of proton pencil beams impinging on a cube phantom. Examples of local proton energy spectra in the plateau region and at the Bragg peak are shown against the relative DSB induction likelihood. A hashed region in the lower left plot shows the component of the energy spectrum that contributes to complex DNA damage.	123

4.12	RBE _{cd} -weighted dose scored along the central axis in TOPAS [161] for proton pencil beams between 60 MeV and 120 MeV. Two RBE _{cd} scoring techniques are shown, dose-weighted and fluence-weighted.	125
4.13	Comparison of RBE predictions to experimental data provided by the PIDE database [164]. Experimental RBE _{SF} across surviving fractions of 5% to 95% are shown as box plots (black). Predictions from two models in this work, RBE _{SF} + $f(LET_d)$ (yellow) and RBE _{SF} +PDL (green), given by (4.6)-(4.7) and (4.8)-(4.9) respectively, are shown along with the RMF model [201, 202, 203] as in (3.26) and (3.24) with RBE _{DSB} calculated using the PDL model. RBE _{cd} predictions from the $f(L_\Delta)$ envelope in (4.11) are overlaid. (b) Affect of changing the diameter of the cell nucleus or sub-cellular domain on the RMF+ $f(LET_d)$ prediction.	128
5.1	CAD model of the CCC treatment line. Courtesy of J. Yap. Reproduced from [347] with permission.	131
5.2	Schematic of the TOPAS model of the CCC treatment line [165]. Courtesy of J. Yap. Reproduced from [347] with permission.	132
5.3	Schematic (a) and photograph (b) of the active sensor and Timepix ASIC chip of the MiniPIX hybrid semiconductor detector. The full detector (c) is around the size of a large USB data stick. Reproduced from [351, Fig. 1] with permission. . .	133
5.4	The Advacam MiniPIX-Timepix detector placed in a custom-made 3D-printed housing, which is attached to a remotely controlled rotating base.	133
5.5	Setup of the Advacam MiniPIX-Timepix detector at the CCC treatment line. Here, the detector is placed at a 45 degree angle to the central axis of the beam. . . .	134
5.6	Visualisation in Advacam PIXet proprietary software [359] of individual clusters in a single frame acquired over 10 ms. Relative pixel intensities represent the amount of charge deposited. Run R7 (60° detector angle, 10 mm Lucite) from Table 5.1 is shown. Figure adapted from screenshot provided by C. Oancea (Advacam, Prague).	137

5.7	Visualisation in Advacam PIXet proprietary software [359]. Left: all integrated frames. Relative pixel intensities represent the amount of charge deposited. Right: cluster height distribution, showing energy deposited per pixel. Run R7 (60° detector angle, 10 mm Lucite) from Table 5.1 is shown. Screenshot provided by C. Oancea (Advacam, Prague).	138
5.8	Example of particle tracking in the MiniPIX Timepix detector. Reproduced from [351, Fig. 2] with permission.	140
5.9	Spread of the beam from phase space of 10^6 primary protons collected immediately downstream of the nozzle, after collimation. (a) Emittance in the sagittal axis. (b) Emittance in the coronal axis. (c) Position distribution of protons in the transverse plane.	142
5.10	Histogram of proton energy from phase space of 10^6 primary protons collected immediately downstream of the nozzle, after collimation. A Gaussian fit with mean 60.04 MeV and standard deviation 0.48 MeV is shown.	143
5.11	(a-i) Histograms of energy deposition as measured by the MiniPIX sensor (“meas.”) for all runs R1 to R9, and as calculated from simulations S1a to S9a which include aluminium shielding.	148
5.12	(a-i) Histograms of LET contributions calculated from the MiniPIX sensor data (“meas.”) for all runs R1 to R9, compared to LET_d output from simulations S1a to S9a which include aluminium shielding. All LET values are in silicon.	149
5.13	Boxplots showing the spread, median and MPV of LET contributions in silicon calculated from the MiniPIX sensor data beside those retrieved from the simulations S1a to S9a. Mean $LET_d(H_2O)$ values scored from a water cube phantom simulation are also shown. Runs for which particle counting statistics were considerably low, corresponding to $WET = \infty$, (R5, R6 and R9) have been faded out.	150

5.14	All LET results, both measured and simulated, corresponding to finite WET plotted as their MPV against a simulated depth-dose profile in a water phantom. Error bars indicate the 95% confidence interval centred on the mean. LET values in silicon have been converted to water equivalent values, $LET(H_2O)$, for all detector measurements and simulations.	152
5.15	Cross-section of scored LET_d in TOPAS at depths of (a) 10 mm, (b) 30 mm, (c) 32 mm and (d) 33 mm in a water phantom of side lengths 35 mm. Bin sizes of 1 mm^3 were used for scoring. Images have been upscaled by a factor of 10 and interpolated.	153
5.16	(a-i) Histograms of the kinetic energy of protons impinging on the surface of the sensor, as scored in the simulations S1a to S9a (with aluminium shielding) and S1b to S9b (without shielding).	154
5.17	Boxplots showing the spread, median and MPV of DSB counts, and the associated RBE_{cd} by DSB induction scored from the simulations S1a to S9a. Runs for which particle counting statistics were considerably low, corresponding to $WET = \infty$, (R5, R6 and R9) have been faded out.	155
5.18	All simulated RBE_{cd} estimates corresponding to finite WET plotted as their MPV against a simulated depth-dose profile in a water phantom. Error bars indicate the 95% confidence interval centred on the mean.	156
5.19	Stopping power (left axis) and CSDA range (right axis) in Aluminium, showing the energy range in which protons traversing the 1 mm shielding layer may be absorbed and not reach the detector.	157
5.20	Distribution of deposited energy and track lengths from the MiniPIX data (runs R1 to R9, Table 5.1) with level curves of constant LET superimposed for reference.	158
5.21	Most probable values (MPVs) of LET plotted against MPVs of RBE_{cd} collected from all simulations in Table 5.2. The shaded region represents the area between the empirical RBE_{cd} fit, (4.11), to LET_d (bottom) and LET_t (top) presented in Chapter 4.	159

6.1	Each pencil beam gives rise to a 3D dose grid of dimensions $m_x \times m_y \times m_z$ consisting of $m = m_x m_y m_z$ voxels, which is transformed to a column vector of length m . For n pencil beams, columns are concatenated to form the $m \times n$ dose matrix A	164
6.2	An example of convex feasibility-seeking in two dimensions. The feasible region (shaded) arises from the intersection of lower-halfplanes defined by lines H_1 and H_2 . Starting with an arbitrary initial estimate $\mathbf{x}^{(0)}$, the next iterate is computed by projecting $\mathbf{x}^{(0)}$ orthogonally onto H_1 . Sequential projections are made between H_1 and H_2 , with the estimate converging on a feasible solution, \mathbf{x}^*	167
6.3	The possible choices of the vector of beamlet intensities \mathbf{x} is restricted to the set of linear constraints C . The SFP in (6.17) requires that for each structure s with a DVC applied to it, one must also restrict the dose vector $A_s \mathbf{x}$ to the possibilities in the set of non-linear constraints Q_s . Any vector \mathbf{x}^* in the set $\Gamma = C \cap A^{-1}(Q_s)$ will satisfy both the linear and the non-linear constraints and therefore be labelled a feasible solution.	173
6.4	Orthogonal projection $P_\Omega(\mathbf{z})$ of an element \mathbf{z} onto the hyperslab Ω . In the case that \mathbf{z} is above the hyperplane H_1 , which defines a boundary of Ω , the projection is made onto H_1 . The relaxation parameter λ allows the projection to (a) fall on H_1 exactly, (b) overshoot H_1 , or (c) undershoot H_1	176
6.5	In the case of 6 voxels representing an OAR, if we allow the dose in one third of them to overflow, then the operator P_Q takes the lowest violation (voxel 3) and decreases the contained dose to the prescribed maximum.	177
6.6	Examples of (a) fully-sequential and (b) fully-simultaneous projections onto four hyperplanes, H_i for $i \in \{1, 2, 3, 4\}$. The current solution vector $\mathbf{x}^{(k)}$ is projected onto the four hyperplanes in order to determine the next iterate(s), $\mathbf{x}^{(k+1)}$	182

7.1	Flowchart of Algorithm 4. N_{cycles} is the number of algorithmic iterations, N_{struct} is the number of structures, and N_{voxels} is the number of voxels comprising the current structure. Counter increments are implied within each loop, with indices starting at 1 by default.	186
7.2	(a) A single Gaussian “dose” kernel contribution shown at one grid point. (b) Homogeneous flood dose of 50 units formed by superimposing all 1156 contributions. (c) The optimised dose map showing the structures for which the Test Case 1 prescription in Table 7.1 was applied.	189
7.3	(a) Flood dose, with mean of 50 units, formed by superimposing 484 Gaussian dose contributions. (b) The optimised dose map showing the structures for which the Test Case 2 prescription in Table 7.1 was applied.	190
7.4	Cumulative dose-volume histogram (DVH) for Test Case 1 showing the percentage of each structure that has received a certain dose. HDC and DVC prescriptions are shown as filled circles.	194
7.5	Number of total violations as a function of number of the algorithmic cycles for Test Case 1. A decrease indicates improvement in meeting the prescription. . . .	195
7.6	Percentage of violations as a function of the number of algorithmic cycles for Test Case 1, shown separately for HDCs (minimum and maximum doses) and DVCs. An upper DVC is that which is applied to an avoidance structure while a lower DVC is that which is applied to a target structure.	196
7.7	Log-loss curves showing the reduction in the total number of constraint violations as the number of cycles increases, for (a) Test Case 2, (b) Test Case 3 and (c) TROTS Case 1. All show both parameter-update schemes, SARP1 and SARP2, as well as the static method in which the relaxation parameters are set to their mid-range.	197

7.8	Mean number of constraint violations, as a percentage of the total number of voxels, shown against the number of algorithmic cycles for Test Case 2 using SARP2. Violations have been separated into those on DVC constraints, HDC constraints, and the average of these two. The static method are shown for comparison. . . .	198
7.9	Forest plots showing non-zero dose differences in TROTS Case 1 distributed in the voxels of the structures of interest. Higher doses obtained using the DSACQ algorithm, when compared to the included TROTS solution, manifest as positive values. One asterisk beside the data indicates there is a statistically significant difference in the doses; two or three indicate the absolute difference is statistically greater than 0.25 Gy or 0.50 Gy respectively. Significance is defined by a Wilcoxon signed-rank test returning a probability of $p < 0.05$	199
7.10	Transverse (top row) and sagittal (bottom row) CT slices for TROTS Case 1, showing the final dose after 500 iterations of the DSACQ algorithm using SARP2 (a, d), the actual dose included with the TROTS data (b, e), and the dose difference between the former and the latter (c, f). Absolute doses above 0.5 Gy or differences less than 2% are not shown. The treatment isocentre is shown by a black cross, and the chosen slices are indicated by dashed gray lines.	200
7.11	Cumulative dose-volume-histogram (DVH) for TROTS Case 1. Solid lines represent the solution after 500 cycles of the DSACQ algorithm using SARP2. Dashed lines represent the solution included in the TROTS data.	201
7.12	Log-loss curves showing the relative reduction in the total number of constraint violations as the number of cycles increases, for all 20 TROTS proton cases. Results using the SARP2 method are shown against the static method in which the relaxation parameters are set to their mid-range.	202

7.13	Run time comparison (per cycle) between original <i>MATLAB</i> implementation and accelerated <i>Python</i> acceleration of the DSACQ algorithm through multi-threading (MT) and just-in-time (JIT) compilation for TROTS Case 1. Bottom (blue) bars: all HDC and DVC projections. Top (orange): only HDC projections.	203
7.14	Cumulative dose-volume-histogram (DVH) for TROTS Case 1 after code acceleration. Solid lines represent the solution after 50 cycles of the DSACQ algorithm with static relaxation parameters and no DVCs. Dashed lines represent the solution included in the TROTS data.	204
7.15	Share of constraint violations between Target A and Target B in Test Case 2. The outer ring represents the 499th algorithmic cycle while the inner ring represents the 500th cycle, after an update of the relaxation parameters with SARP2. This behaviour occurs each time the relaxation parameters are updated. A compromise between the two solutions is to take the average.	205
7.16	Cumulative dose-volume-histogram (DVH) for Test Case 2. Solid lines represent the solution after 500 cycles of the DSACQ algorithm using SARP2. Dashed lines represent the solution after 499 cycles.	207
A.1	Shaded regions represent the integration region $\mathcal{D} := \{(x, y) \in \mathbb{R}_+^2 \mid xd/y \leq \ell\}$ for examples where (a) $\ell \leq H$, and (b) $\ell > H$	214
A.2	The integration region $\mathcal{D}_{\text{ext}} := \{(x, y) \in \mathbb{R}_+^2 \mid y \leq \sqrt{1 - H^2/\ell^2}\}$ is given by the shaded area.	215
C.1	(a-i) Histograms of LET contributions calculated from the MiniPIX sensor data (purple shaded area) for all runs R1 to R9, with associated Landau (black line) and general stable distribution fits (red line). Fitting procedure performed by <i>Python</i> script provided by F. Van den Heuvel.	226

E.1	Forest plots showing non-zero dose differences for TROTS Cases 1–6 distributed in the voxels of the structures of interest. Higher doses obtained using the DSACQ algorithm, when compared to the included TROTS solution, manifest as positive values.	230
E.2	Forest plots showing non-zero dose differences for TROTS Cases 7–12 distributed in the voxels of the structures of interest. Higher doses obtained using the DSACQ algorithm, when compared to the included TROTS solution, manifest as positive values.	231
E.3	Forest plots showing non-zero dose differences for TROTS Cases 13–18 distributed in the voxels of the structures of interest. Higher doses obtained using the DSACQ algorithm, when compared to the included TROTS solution, manifest as positive values.	232
E.4	Forest plots showing non-zero dose differences for TROTS Cases 19–20 distributed in the voxels of the structures of interest. Higher doses obtained using the DSACQ algorithm, when compared to the included TROTS solution, manifest as positive values.	233
F.1	Log-loss curves showing the reduction in the total number of constraint violations as the number of cycles increases, for (a, d) Test Case 2, (b, e) Test Case 3 and (c, f) TROTS Case 1. All parameter-update methods are shown as well as the static method in which the relaxation parameters are set to their mid-range. Plots (a-c) show SARP2 with a threshold condition and (d-f) show the newly formulated SARP3.	235
F.2	Example of SARP3 implementation. (a) Total violations broken down into HDC and DVC violations for three structures. (b) Parameter ranges are adjusted such that the step-sizes of the projections associated with each structure scale with the number of violations. The SARP2 methodology is then applied within these new ranges.	236

List of Tables

3.1	Summary of phenomenological and mechanistic proton RBE models described in Section 3.2.	57
3.2	Parameters fitted to experimental data, by Tan <i>et al.</i> [235] (protons) and NIST [236, 237] (electrons), using a Levenberg-Marquadt procedure, for the IMFP calculation in (3.33). Parameters are provided for both the Bethe-Barkas formalism and the plasmonic correction term.	73
3.3	Parameter fits to Monte Carlo generated DSBs using MCDS 3.1 [159, 160] along with their 95% confidence interval, for DSB conversion efficiency, ε , low-level (minimal) damage, a , and high-level (maximal) damage, b	76
3.4	Levels of confidence for validating DSB-based RBE models.	79
4.1	Beam parameters used to generate monoenergetic and modulated pencil beams in TOPAS [161]. SOBP width, χ , and weighting parameter, p , are implemented in (4.4) and (4.5) respectively.	106
4.2	Parameter fits to (4.11) using 550 data points for both dose-averaged LET (LET_d) and track-averaged LET (LET_t). 95% confidence intervals are shown in brackets. The coefficient of determination for this fit was $R^2 = 0.999$ against LET_d and $R^2 = 0.984$ against LET_t	113
5.1	Details of each Clatterbridge experiment with the Advacam MiniPIX-Timepix detector.	136

5.2	Details of each Clatterbridge simulation in TOPAS, corresponding to the experiments in Table 5.1.	145
5.3	Calculated and measured LET values from all runs in Table 5.1 and all simulations in Table 5.2. All LET values are in silicon except for the water phantom profile, LET(H ₂ O).	146
5.4	Calculated DSB counts and RBE _{cd} values from all simulations in Table 5.2. . . .	147
7.1	Prescriptions chosen for all test cases.	191

Chapter 1

Introduction

“Therapy should be permanently backed up by scientific research without which no progress is possible. Moreover, the search for pure knowledge is one of the important needs of mankind. . . ”

— Marie Skłodowska Curie, *Remarks at the opening of the Radium Institute, Warsaw, 1932. Found in [1].*

1.1 Modern radiation therapy: from x-rays to protons

It is currently expected that fifty percent of people in the United Kingdom will be diagnosed with cancer during their lifetime [2]. Research into cancer treatments, therefore, has been, and continues to be, an incredibly important and globally consolidated effort. A recent report published by Public Health England in partnership with the NHS and Cancer Research UK compared the three major, but not mutually exclusive, treatment regimes for cancer: chemotherapy, radiation therapy and surgery. The report found that 27% of all cancer treatments in England between 2013 and 2016 involved radiation therapy [3], a technique by which ionising radiation is used to kill tumour cells while sparing, as best as possible, surrounding healthy organs. This figure drastically increases to around 64% on average for head and neck cancer types, including up to 83% for cancers of the oropharynx [4]. Macmillan Cancer Support has reported that by the end of this year there will be 2.9 million people in the UK

living with cancer and that this will increase to four million people by 2030 [5], a 38% increase over the next decade. In 2015 the European Society for Radiation Oncology (ESTRO) performed a health economics study to assess the increased demand for radiation therapy by the year 2025 [6], finding that a 16% increase in demand for radiotherapy services across Europe should be expected, with some countries expecting up to a 30% increase. The study intended to signal to European health policy makers that further investment in radiotherapy was needed along with a reduction in the disparity of current resources across Europe. Looking ahead, as we live longer lives it is clear that cancer prevalence will continue to increase, lest a panacea arrives, and that the need for quality radiotherapy services will only increase.

The advent of radiotherapy dates back to Roentgen's accidental discovery of x-rays¹ in 1895 [7], for which he was awarded the 1901 Nobel Prize in Physics (reviewed by Mould in [8]), and to their almost immediate application to the treatment of skin cancers. Perhaps the first treatment was performed by Grubbe in 1896 for a breast cancer patient [9]. Radiation, and particularly radioactivity of natural sources, was a developing field, rife with interest from researchers including Antoine Henri Becquerel, Marie Skłodowska Curie and Pierre Curie. The invention of the modern x-ray tube by Coolidge [10] two decades after Roentgen's discovery allowed for the production of higher energy, more penetrative x-rays and thus the ability to target non-superficial deep-seated tumours [11]. Initial treatment outcomes were limited in their success and presented a range of side effects as there lacked a mechanistic understanding among proponents of radiotherapy regarding the radiobiological principles at play. It was during the 1920s that radiobiological considerations came to the forefront in both treatment and radiation safety² [12, 13].

External beam radiotherapy as we know it today can be attributed to the pro-

¹The field of radiology used to be referred to as "Roentgenology" after Roentgen's discovery and x-rays were called "cathode rays".

²The International Commission on Radiological Protection (ICRP) was established in 1928. See <http://www.icrp.org>.

duction of high energy megavoltage photons through the use of cobalt-60 as radiation sources and the invention of linear accelerators (LINACs) in the 1950s. Though, an influx of newly mechanically-plausible treatment options was impeded by limitations in patient imaging necessary for accurately selecting appropriate radiation fields. The invention of computed tomography (CT) for imaging by Cormack and Hounsfield in 1973 brought higher density resolution three-dimensional scans, allowing for improved radiological diagnosis and radiotherapy planning. This earned them the 1979 Nobel Prize in Physiology or Medicine³. A key development in radiotherapy, dubbed the most exciting addition since CT imaging [14, 15], was the ability to deliver conformal doses to the tumour through intensity modulated radiation therapy (IMRT), the history of which has been reviewed in detail by Bortfeld [16]. Conceptualised by Brahme in 1982 [17], IMRT posed radiotherapy as an inverse problem that presented alternatives to the impracticality of fabricating individualised passive scatter compensators specific to each treatment. By the end of the decade, inverse planning algorithms began to surface as alternatives to trial-and-error-based treatment planning. An algebraic approach by Altschuler and Censor [18] (see also [19, 20, 21]) was particularly transformative as it replaced an impractical integral inversion problem with one in which both the patient geometry and external radiation field were discretised, allowing for iterative numerical techniques. The first example of modern inverse optimisation for radiotherapy was presented by Webb in 1989 [22]. Only a few years later, Webb presented an inverse optimisation algorithm that allowed for precise two-dimensional x-ray modulation using a large number of programmable tungsten “leaves” [23]. The multi-leaf collimator (MLC), as it was named, has since been a staple of conventional radiotherapy since its first clinical usage in 1996 [24].

Almost in parallel to the developments of x-ray therapy, researchers were exploring the potential uses of protons for radiotherapy, albeit at a slower pace and with more

³Press release, NobelPrize.org, Nobel Media AB 2020, <https://www.nobelprize.org/prizes/medicine/1979/press-release/>, Accessed 28/11/2020.

hesitancy in clinical adoption. Around the same time as Roentgen's discovery of x-rays, Thomson discovered the electron [25]. Modern atomic theory developed over the following two decades as Rutherford investigated the existence of a positively charged nucleus around which electrons orbit [26], which culminated in his discovery of the proton in 1919 [27]. In 1946, Wilson proposed the use of accelerated protons for therapeutic use [28], building on William Henry Bragg's 1903 discovery of the energy loss characteristics of charged particles [29]. Bragg noticed that charged particles have a finite range in matter, depending on their velocity, and that they give up most of their energy to the material just prior to stopping. This sharp *Bragg peak* in dose at the end-of-range is the cornerstone of proton therapy, allowing for highly conformal dose to the tumour volume with minimal dose leakage to healthy tissues. Protons were first used at research accelerators at the Lawrence-Berkeley National Laboratory (1954) and Uppsala University (1967) prior to their first hospital-based use in 1989 for the treatment of ocular tumours at the Clatterbridge Cancer Centre in the United Kingdom [30]. The Loma Linda University Medical Centre in the United States shortly followed, expanding the use of protons to other treatment sites in 1990 [30].

Widespread use of proton therapy, despite its promising dose-shaping characteristics for paediatric cancers and complex treatment sites, was initially hindered due to the high costs associated with the construction of treatment facilities and the maintenance of highly specialised cyclic particle accelerators. However, benefits of improved healthy tissue sparing and, consequently, fewer side effects as well as technological innovations helped elevate the treatment to regular usage. Since the first clinical use of protons in the 1950s over 220,000 patients have been treated with protons globally, as of November 2020, with 47% of these treatments occurring in the past five years alone [31, 32, 33]. Protons currently account for 85% of particle therapy modalities, with the other 15% consisting mostly of carbon ions [33]. As of September 2020, there

are five proton therapy centres currently in use in the United Kingdom with two additional centres under construction [34, 35]. Protons indeed offer therapeutic benefit over conventional x-rays, however radiobiological differences between these modalities have recently been the subject of increased attention. In order to optimally and confidently administer proton treatments, understanding and quantification of these differences must be thoroughly explored.

1.2 Proton therapy physics

1.2.1 Photon interactions with matter

Photons are uncharged, massless quanta of energy and momentum whose interaction with traversed media can only be described stochastically. They are unaffected by the Coulombic fields set up by the nuclei and electrons that comprise the material through which they pass, and instead tend to lose large amounts of their kinetic energy in discrete interactions. Three of these interactions, namely photoelectric absorption, Compton scattering and pair production, are of interest in radiotherapy, as they result in energetic electrons which go on to deposit dose in the material and cause cellular damage by the interaction mechanisms discussed in the next subsection. Although photons are used in conventional radiotherapy, it is their “offspring”, the secondary electrons, that cause damage. While photons can travel large distances and even through the entire patient, the electrons produced may travel only up to a few centimetres from their origin [36, p. 184-185].

In a photoelectric interaction (Figure 1.1(a)) a photon impinges on an atom and gives up all of its energy to a bound electron in an inner shell. If the photon energy, $h\nu$, where h is Planck’s constant and ν is the frequency, is greater than the binding energy then the electron is ejected from the atom with excess kinetic energy. This ejected electron is called a photoelectron. The vacancy left in the energy shell is then




Figure 1.1: Schematics showing the three major mechanisms by which photons interact with matter in the context of radiotherapy; (a) photoelectric interaction, (b) Compton scattering, and (c) pair production. e^- represents an electron while e^+ represents a positron.

filled with an electron from a higher energy band. To conserve energy, a low-energy photon with energy equal to the difference between the two energy levels is emitted. These are labelled *characteristic x-rays* as the cascade, or emission pattern, is unique to the atomic energy structure of the element and can be used to identify the presence of the element in a sample. Sometimes another atomic electron is ejected in preference to characteristic x-rays; these are called Auger electrons. The probability of any of the three listed interaction mechanisms occurring can be linked to the atomic number, Z , of the material or the energy of the incident photon, $h\nu$. The photoelectric cross-section increases in direct proportion to $(Z/E)^3$ where E is the kinetic energy of the incident particle. This sensitive dependence on the atomic number is exploited in diagnostic radiology, in which tissues can be distinguished by contrast [37, p. 24]. In low- Z tissues the emission of Auger electrons can dominate that of characteristic x-rays.

In a Compton interaction (Figure 1.1(b)) a photon interacts with a loosely bound electron, transferring some of its energy to it. The electron recoils with gained kinetic energy while the photon loses an equal amount of energy and changes direction such that momentum in the system is conserved. Compton scatter is most important for photons with energies between approximately 100 keV and 10 MeV, and is the

dominant interaction process for photons ejected from linear accelerators (1 to 20 MV beams) for conventional radiotherapy [38, p. 6]. The likelihood of interaction decreases with increasing energy and is independent of the atomic number of the material.

Photons of energies above 1.022 MeV are able to interact with nuclei to undergo pair production (Figure 1.1(c)). This is when the photon is converted into an electron and positron pair, each with a binding energy of 0.511 MeV. The kinetic energies of the electron and positron must add to $(h\nu - 1.022)$ MeV in order to conserve energy. Positrons are short-lived and travel only a short distance in ordinary matter before being annihilated through interaction with an electron. Two photons of 511 keV are emitted in this process. The likelihood of pair production increases with increasing photon energy, above 1.022 MeV of course, and increases approximately in proportion to Z^2 .

While the trajectory and interactions of an individual photon are not deterministic, the collective behaviour of a large number of them is, thankfully, predictable. A sufficiently sized “bunch” of photons can be described by beam intensity and beam quality, two terms which do not apply to individual photons. The beam intensity refers to the number of photons, which is often measured as particle fluence (the number of particles per unit area), while the beam quality refers to the penetrative power of the beam, which is related not to the number of photons but instead to their energy. As a beam of photons passes through a material it is attenuated by the stochastic action of a proportion of photons interacting with atoms and nuclei. This attenuation leads to a reduction in fluence, Φ , that is mathematically described by

$$\Phi(x) = \Phi_0 \exp\left(-\rho x \left(\frac{\mu}{\rho}\right)\right) \quad (1.1)$$

where Φ_0 is the fluence of photons incident on a material of mass density ρ , $\Phi(x)$ is

the fluence at depth x in the material and (μ/ρ) is known as the mass attenuation coefficient. This coefficient describes the fractional reduction in intensity of a beam of photons over the distance travelled, considering all possible interactions, and has units of $\text{cm}^2 \text{g}^{-1}$. The mass attenuation coefficient, taking into account the density of the absorber, is very useful for radiation shielding calculations. Note that (1.1) assumes a narrow beam of monoenergetic photons. In practice, beams are polyenergetic, consisting of a spectrum of photon energies. It is indeed a broad spectrum that is produced by an x-ray tube for conventional radiotherapy. In this case Φ must be treated as a function of energy and so must (μ/ρ) . There is an interesting interplay between intensity and beam quality which is most noticeable, and problematic, in computed tomography (CT) as it introduces imaging artifacts. As the beam passes through material, lower energy photons exhibit a larger attenuation cross-section than those of higher energy, which shifts the expected value of the energy spectrum to higher values. This effect is known as *beam hardening*.

1.2.2 Proton interactions with matter

Contrary to photons, charged particles are continuously slowed down, and eventually stopped, in material through consistent interaction between the Coulombic fields of the incident particle and atomic nuclei in the material. A proton traversing a material will “feel” every other source of charge present in that material due to the electromagnetic interaction. One may intuit that the possible interaction mechanisms that can occur for protons travelling through a material depend on the distance between the proton and atomic nuclei, as the Coulombic repulsion between them is inversely proportional to the square of this distance, and the amount of time spent in such proximity⁴. This can be further distilled down to the impact parameter, b , which is defined as the shortest distance between the initial (straight) proton trajectory and

⁴Both of these concepts are captured in the impulse of the Coulombic force.

the nucleus (see Figure 1.3), and the speed, v , or equivalently the kinetic energy, E , of the proton. In the context of proton therapy we are interested in the three interaction mechanisms that have the most dosimetric (and by virtue, biological) consequences; (1) inelastic Coulomb interaction with atomic electrons, (2) elastic Coulomb scattering with atomic nuclei, and (3) non-elastic nuclear collisions.

1.2.2.1 Inelastic Coulomb interaction

Consider Bohr's classical view of the atom [39], with atomic radius a . The most likely scenario for a passing proton is that its trajectory is relatively far away from the atomic nucleus ($b \gg a$), either giving enough energy to a valence electron to ionize the atom or exciting an electron to a higher energy level which subsequently returns to its ground state and emits a photon. In either case the proton loses some energy, but it is negligibly deflected from its original trajectory due to its rest mass being over 1800 times greater than that of the electron. These types of interactions in

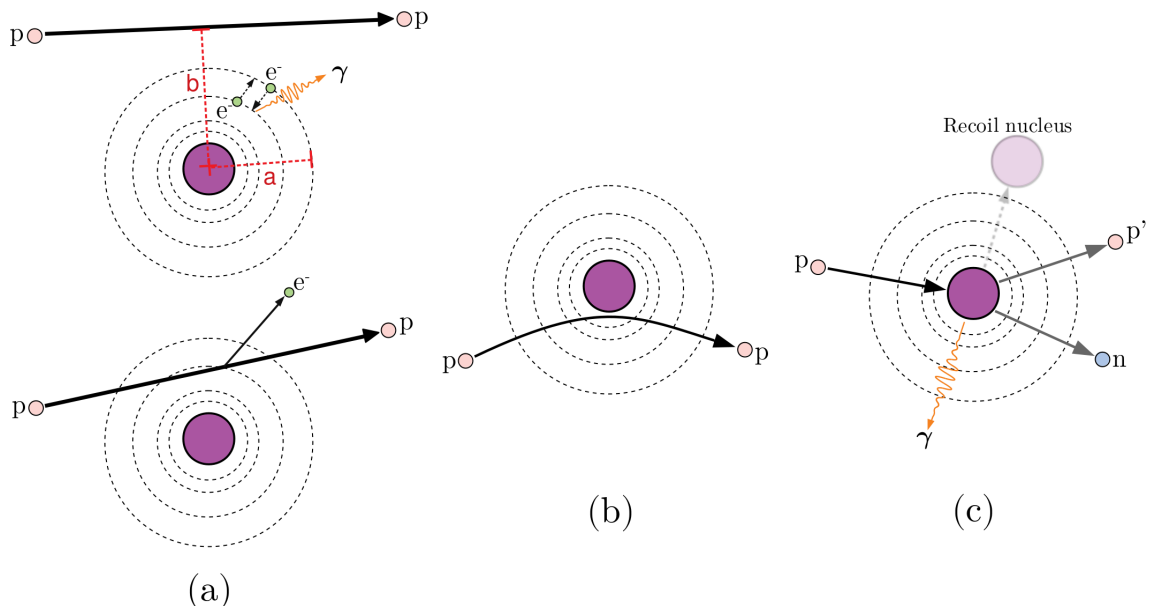


Figure 1.2: Illustration of proton interaction mechanisms pertinent to radiation therapy: (a) Inelastic Coulomb scattering (if the impact parameter, b , is of the order of the atomic radius, a , then “knock-on” electrons can result), (b) elastic Coulomb scattering with atomic nuclei, and (c) non-elastic collisions. e , p , n and γ represent an electron, proton, neutron and photon, respectively.

which $b \gg a$ are often called “soft collisions” and they explain the quasi-continuous loss of proton energy. Section 1.2.4 introduces a statistical measure of this energy loss for a beam of protons, called the *stopping power*. “Knock-on” collisions, shown in the lower portion of Figure 1.2(a), are likely to occur when the impact parameter is of the order of the atomic radius ($b \sim a$). In this scenario the proton interacts with a single atomic electron in a manner more analogous to a billiard ball collision, imparting a large amount of kinetic energy to it. In fact, the energy is large enough that this electron (labelled a δ -ray) can travel a significant distance and warrant its own track of collisions that is of dosimetric interest.

1.2.2.2 Elastic Coulomb scattering

As a proton passes close to an atomic nucleus ($b \ll a$) the strong Coulombic repulsion force dominates. Unlike the interaction with electrons, the large mass of the nucleus causes the proton to deviate from its original path, as shown in Figure 1.2(b). In response there is a small recoil of the nucleus and in most cases a negligible amount of energy lost by the proton. Therefore, the collision is considered elastic. This elastic scattering determines the sharpness of lateral penumbra in the dose deposited within a patient [40] and can be described statistically using multiple Coulomb scattering (MCS) theory (briefly discussed in Section 1.2.3). At this stage it is useful to mention that although our discussion here is focused on protons, all of these physical concepts apply to electrons as well. Although, they are drawn toward the nucleus rather than repelled from it, and their small mass means they experience much more considerable deflections. Electrons may lose energy in this process, which manifests as photons, however elastic scattering with the nucleus is about 30 to 50 times more likely to occur than radiative losses [36, p. 163]. If photons are produced, they can have an energy of up to the incident kinetic energy of the electron. The spectrum of photons produced is known as *bremmstrahlung*, and this is in fact the source of x-rays (produced

from an x-ray tube) for conventional radiotherapy. While proton bremsstrahlung is theoretically possible, it is of insignificant concern in proton therapy.

1.2.2.3 Non-elastic nuclear collisions

Non-elastic nuclear collisions involve the incident (*primary*) proton being captured by an atomic nucleus followed by spallation products (protons, neutrons, deuterons, tritons, alpha particles or heavier ions) directed largely in the incident direction of the primary proton, an example of which is shown in Figure 1.2(c). These nuclear fragments contribute to the *secondary* radiation field, a naming convention which distinguishes them from the primary proton field. In order for the interaction to occur, the proton must have an energy greater than the electrostatic binding energy of the nucleons, which is approximately 8 MeV for anatomical materials. While these non-elastic scattering events with the nucleus are few compared to the elastic and inelastic Coulomb interactions described above, they nonetheless have considerable consequences in dose, biological effect and radiation protection. Secondary protons comprise about 10% of the dose in the patient, however the impact on the spatial dose distribution for therapeutic fields is small [see, e.g., 41]. Heavier secondary particles contribute only about 1% of the dose [42] and most travel only a very short distance from their creation, however their energy losses over this distance can have increased biological significance and lead to increased cell-killing that would not be suggested solely by the physical dose deposited. We are now moving into the realm of linear energy transfer (LET) (Section 1.2.6) and relative biological effectiveness (RBE) (Section 1.3.5), which are of paramount importance throughout this thesis. As such, finer details are left until later on.

A non-elastic interaction caused by a secondary neutron can leave the nucleus in an unstable excited energy state of “induced radioactivity”. Analogous to an electron moving down an energy level to fill a vacancy, the nucleus is returned to its ground




Figure 1.3: Classical scattering in three dimensions. Charged particles with mass m , kinetic energy E , speed v_0 and initial momentum $p_0 = mv_0$ are deviated by a spherical potential whose origin is on the z -axis. Particles within an infinitesimal patch of area $d\sigma = b db d\phi$ at azimuthal angle ϕ with impact parameter b scatter into a corresponding infinitesimal solid angle $d\Omega = \sin(\theta)d\theta d\phi$ where θ is the polar scatter angle.

state by a nucleon (or several nucleons) dropping to lower energy levels. However the energy gaps between levels are much larger and give rise to high-energy γ -ray photons. It may take several minutes for this decay and emission to occur. For this reason one must be acutely aware of persisting radioactivity in the treatment room, and especially in the linear accelerator itself, after the beam has been turned off.

1.2.3 Multiple scattering

In the case of elastic scatter of a single proton due to a positively charged atomic nucleus, classical scattering theory allows us to determine the distribution of possible deflection angles. Extending this to a multitude of successive small-angle scattering events in a beam of protons requires a probabilistic approach to model the nett angular deflection through some finite thickness. This series of small-angle deflections is referred to as multiple Coulomb scattering (MCS). We shall begin with a brief recount of Rutherford scattering [26] before providing a short, practical overview of MCS theory.

Consider a proton incident on a spherical potential as shown in Figure 1.3. Protons within an infinitesimal patch of area $d\sigma = b db d\phi$ with impact parameter b and az-

imuthal angle ϕ scatter into a corresponding infinitesimal solid angle $d\Omega = \sin(\theta)d\theta d\phi$ where θ is the polar scatter angle. In Rutherford's calculations [26], the differential scattering cross-section, or probability of scatter, $d\sigma$, for scattering into the solid angle $d\Omega$ is given by

$$\frac{d\sigma}{d\Omega} = \left(\frac{Ze^2}{8\pi\epsilon_0 m v_0^2} \right)^2 \text{cosec}^4(\theta/2) \quad (1.2)$$

where Z is the atomic number of the nucleus, e is the electronic charge, ϵ_0 is the vacuum permittivity, m is the mass of the proton, v_0 is its initial speed, and θ is the projected scattering angle as shown in Figure 1.3. We omit derivations here for brevity but it should be noted that (1.2) can also be derived quantum mechanically by considering the first Born approximation of a plane wave incident on a screened Yukawa potential, and it can also be recovered from the Mott cross-section for a spin- $\frac{1}{2}$ particle in the presence of a nucleus [43].

The amount of spread, or nett deflection, of a proton beam due to elastic scattering is characterised by the *scattering power*, defined as the rate of increase, with depth, x , in a material, of the mean square of the projected scattering angle, θ . That is,

$$T(x) := \frac{d\langle\theta^2\rangle}{dx}. \quad (1.3)$$

The scattering power is an important concept in modelling MCS. The most accurate model of MCS is provided by Molière [44, 45], however a simplified Gaussian model of the probabilistic proton trajectory is provided by Fermi-Eyges theory [46, 47], which is sufficiently accurate for small deflection angles. In Fermi-Eyges theory the variance of the scattering angle, $\text{var}(\theta)$, is given by the zero-th moment of the scattering power,

$$\text{var}(\theta) = \int_0^x T(\xi)d\xi, \quad (1.4)$$

and in assuming a Gaussian scattering process one may infer the scattering power

from

$$d\langle\theta^2\rangle = dx \int \theta^2 d\sigma. \quad (1.5)$$

Here, $d\sigma$ is the usual differential scattering probability, for which we could substitute the Rutherford cross-section in (1.2). However, the Mott and Rutherford formulae both assume that the nucleus sets up an electric field equivalent to that of a point positive charge. A more accurate formula that takes into account the finite size of the nucleus and the screening of its electric field by orbital electrons is provided by Goudsmit and Saunderson [48, 49]. Gottschalk presented an explicit expression for the scattering power [50] which has been widely applied in proton CT reconstruction⁵.

1.2.4 Stopping power

As discussed in Section 1.2.2, inelastic Coulomb interactions between projectile protons and atomic electrons are responsible for a quasi-continuous loss of proton energy as they traverse media. The expected rate of energy loss, dE , over a distance dx is known as the *stopping power*,

$$S(E) := \left\langle \frac{dE}{dx} \right\rangle, \quad (1.6)$$

and it is a concept that applies not only to proton beams, but to any charged particle beam. Note that the $\langle\cdot\rangle$ notation is often omitted as it is well-understood that stopping power is an expected quantity over numerous energy loss events and should not be applied to a single particle. It is also common to report the *mass stopping power*, S/ρ , given the particles traverse a material of mass density ρ , in units of $\text{MeV cm}^2\text{g}^{-1}$. The Bethe-Bloch formula [52, 53] is most often used to determine the stopping power in therapeutic energy ranges, of about 3 MeV to 300 MeV [50]. In its full form, it is

⁵This includes work published by the author of this thesis in February 2020, in which heterogeneous moments of the scattering power are used to develop a most likely path formalism for proton CT imaging [51].

given by

$$\frac{S}{\rho} := -\frac{1}{\rho} \left\langle \frac{dE}{dx} \right\rangle = 4\pi N_A r_e^2 m_e c^2 \frac{Z}{A} \frac{z^2}{\beta^2} \left[\log \left(\frac{2m_e c^2 \gamma^2 \beta^2}{I} \right) - \beta^2 - \frac{\delta}{2} - \frac{C}{Z} \right] \quad (1.7)$$

where N_A is Avogadro's number, r_e is the classical electron radius, $m_e c^2$ is the rest mass of the electron, c is the speed of light in a vacuum, z is the charge of the particle, Z is the atomic number of the material traversed, A is its mass number, and I is its mean excitation potential. $\beta := v/c$ where v is the velocity of the particle. The last two terms represent relativistic and quantum mechanical corrections that must be considered at very high energy (shielding correction, δ) or very low energy (shell correction, C). In practice, for proton therapy applications, the following simplification can be used with negligible loss of accuracy:

$$S = \frac{4\pi}{m_e c^2} \frac{\rho_e}{\beta^2} \left(\frac{e^2}{4\pi\epsilon_0} \right)^2 \left[\log \left(\frac{2m_e c^2 \beta^2}{I(1-\beta^2)} \right) - \beta^2 \right]. \quad (1.8)$$

Here, ρ_e is the electron number density of the material, e is the electronic charge and ϵ_0 is the electric permittivity of free space. The stopping power in composite materials of n elements can be estimated using known mean excitation potentials of individual elements and the Bragg additive rule;

$$\ln I = \sum_{k=1}^n w_k \frac{Z_k}{A_k} \ln I_k \left\langle \frac{Z}{A} \right\rangle^{-1} \quad (1.9)$$

where

$$\left\langle \frac{Z}{A} \right\rangle = \sum_{k=1}^n w_k \frac{Z_k}{A_k} \quad (1.10)$$

and w_k is the fractional weight per volume of the k -th element.

The total stopping power can be broken down into three independent terms; the electronic stopping power, S_{el} , due to the loss of energy to atomic electrons; the nuclear stopping power, S_{nuc} , due to elastic collisions with atomic nuclei; and the




Figure 1.4: Proton stopping power contributions (left axis) and range (right axis) in water as functions of kinetic energy. S_{el} : electronic stopping power; S_{nuc} : nuclear stopping power; R_{CSDA} : CSDA range; $R = \alpha E^p$: range approximation by Bragg-Kleeman rule [29] as in (1.12). Data sourced from the NIST pSTAR database [54].

radiative stopping power, S_{rad} , due to Bremsstrahlung from nonelastic nuclear interactions. As mentioned previously, radiative losses are negligible in proton therapy [40] and so calculation of S_{rad} is rarely necessary. Figure 1.4 shows the relative contributions of S_{el} and S_{nuc} to the total stopping power in liquid water for proton energies between 1 keV and 300 MeV, using the NIST pSTAR database [54]. The S_{nuc} contribution is more pronounced at lower energies, however it is only at energies below 20 keV at which the contribution becomes larger than 1% of the total stopping power. This contribution quickly jumps to approximately 25% at 1 keV.

The range that protons travel in a material is dependent on their energy loss. Therefore, the range is also a stochastic quantity, defined as the depth at which half of the protons have come to rest in the material. The average range may be calculated

in a continuous slowing down approximation (CSDA) as follows:

$$R = \int_0^R dx = \int_{E_0}^0 \left(\frac{dx}{dE} \right) dE = \int_{E_0}^0 \frac{1}{S} dE, \quad (1.11)$$

where E_0 is the initial kinetic energy of the proton. The empirically-derived Bragg-Kleeman rule [29],

$$R = \alpha E^{p_0} \quad (1.12)$$

for $\alpha = 0.0022$ and $p_0 = 1.77$ (according to Bortfeld and Schlegel [55, 56]) may be used to approximate the CSDA range for proton energies above approximately 1 MeV, as demonstrated in Figure 1.4.

1.2.5 Depth-dose characteristics

The cornerstone of proton therapy, the *Bragg peak* in the depth-dose distribution, comes as a result of two main phenomena; (1) how energy is transferred in electromagnetic interactions, calculated through the stopping power formula in (1.7); and (2) a fundamental difference between electrons, which are the dose-inducing particles in conventional photon therapy, and protons. It can be seen in (1.7) that the stopping power is directly proportional to the inverse square of the particle velocity, $v = \beta c$. Therefore, as the particle slows down, more energy is lost per unit length, until it gives up all available energy and stops. Stopping power is independent of the mass of the particle, so one may, rightly, question why the electron dose deposition is so different to that of protons (as illustrated in Figure 1.5). The answer does in fact relate to the relative masses of the two particles, even though the stopping power is described identically. The proton, having over 1800 times the mass of the electron, suffers smaller deflections in its interaction with electrons in the target material. Electrons, on the other hand, suffer large angle deflections. This scatter results in a diffuse dose deposition in which there is no defined Bragg peak for the beam as a whole.

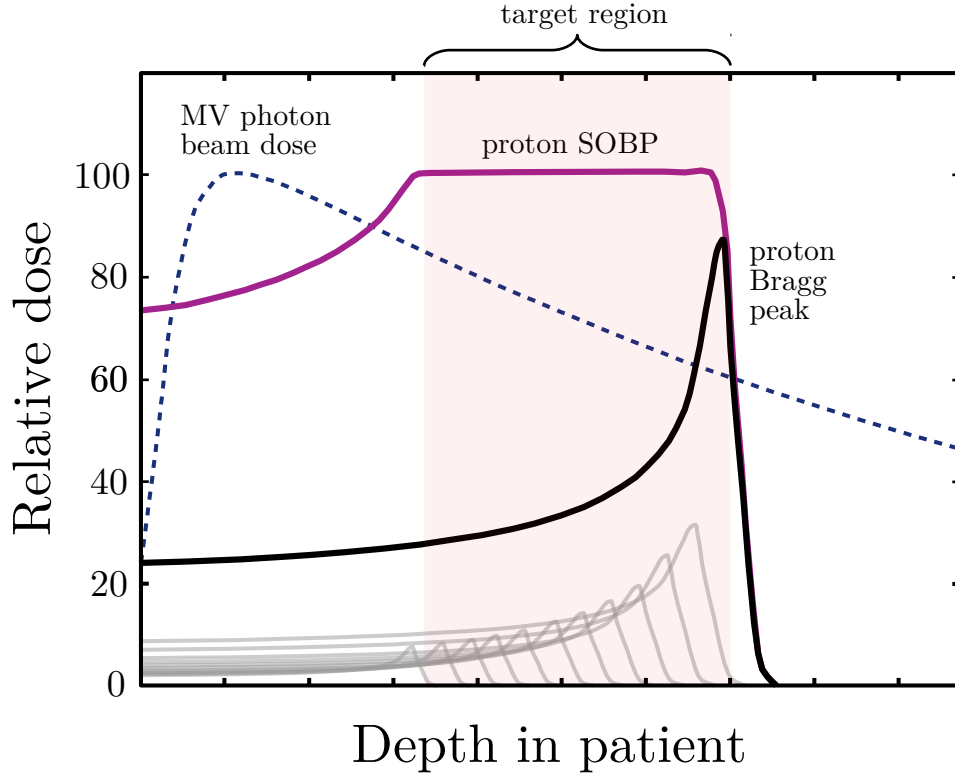


Figure 1.5: Illustrative comparison of a proton depth-dose curve to that of a conventional megavoltage (MV) photon beam. Individual pristine proton pencil beams produce the Bragg peaks shown which together contribute to a homogeneous spread out Bragg peak (SOBP) covering the target region.

The dose, \mathcal{D} , deposited in a small volume of material is given by the energy, ϵ , transferred to the material by all particles passing through it, divided by the mass, m , of that volume. That is, $\mathcal{D} = \epsilon/m$, which has units of J kg^{-1} , otherwise known as *gray (Gy)*. If the particle fluence in a small volume of material with mass density ρ is known for the entire spectrum of kinetic energies then the dose in that volume can be calculated using

$$\mathcal{D} = \sum_{\text{particle type}} \int_0^{E_{\max}} \Phi(E) \frac{S_{\text{el}}}{\rho}(E) dE \quad (1.13)$$

where E_{\max} is the maximum particle energy in that volume, $\Phi(E)$ is the particle fluence, or number of particles, passing through that volume with kinetic energy E , and S_{el} is the material-dependent electronic (or *collisional*) stopping power. A

summation must be applied over all types of ionising particles that pass through the volume.

The pronounced proton Bragg peak allows for highly localised dose deposition at the target, low dose on entry into the patient, and virtually no dose distal of the peak. In comparison, the dose due to a photon beam is distributed more broadly and all the way through to the exit boundary, with the highest dose being deposited only a few centimetres past the skin surface. There are therefore clear benefits to using protons over photons in terms of physical dose conformality; tumoural tissue can be precisely targeted while reducing dose leakage to surrounding healthy tissues. In practice, a monoenergetic proton pencil beam is not a wise choice for treatment as the dose at the Bragg peak is, by definition, highly heterogeneous. Instead, the incident, or nominal, energies of many individual proton pencil beams are modulated in order to produce a homogeneous dose that covers the entire target. (Finer details of how this is achieved in three dimensions are provided in Section 1.4.1.) The resulting dose distribution is known as a *spread out Bragg peak (SOBP)* and is shown in Figure 1.5 alongside its constituent monoenergetic pencil beam doses.

1.2.6 Linear energy transfer

The linear energy transfer (LET) is defined by the International Commission on Radiation Units and Measures (ICRU) as the restricted linear electronic stopping power for a given material [57]. That is, it is equal to the differential mean energy transferred to the material through electronic interactions per unit path length, including the energy carried away from the primary track by secondary electrons but only up to some threshold Δ ;

$$\text{LET}_\Delta := S_{\text{el}} - \left\langle \frac{dE}{dx} \right\rangle_{\text{KE} > \Delta}. \quad (1.14)$$

The second term in (1.14) is the mean energy transferred to the material by secondary electrons with initial kinetic energies larger than Δ . The unrestricted LET, denoted LET_{∞} or with the subscript dropped entirely, is often used for simplicity. In this case the LET is assumed to be equal to S_{e1} and it includes all secondary electrons in the calculation.

The LET provides a microdosimetric measure of the radiation quality, which provides further insight into the density of ionisations along a proton track, as illustrated in Figure 1.6 at the scale of the DNA double-helix. Therefore, it is commonly used in a first approximation to explain any observed differences in the biological effect of radiation, which is discussed in finer detail in the next Section and in Chapter 3. Photons produce a sparse ionisation pattern in comparison to protons or heavier ions and exhibit a correspondingly low LET.

If one were to investigate each proton track individually, even when confined to

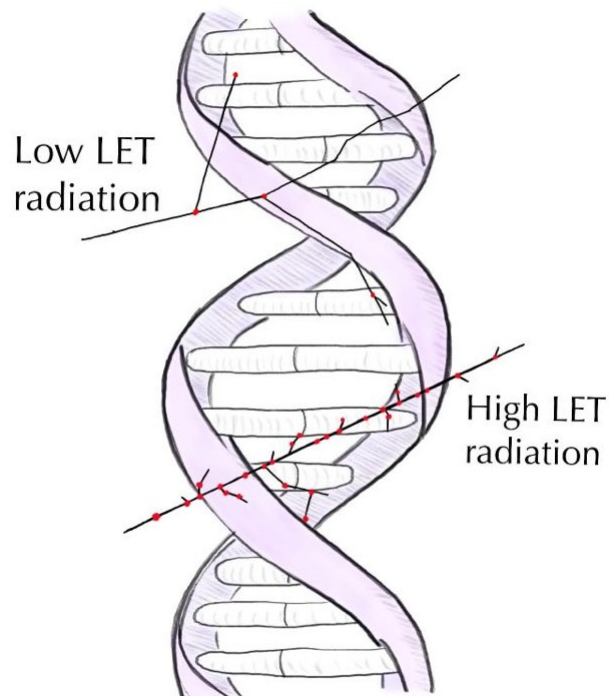


Figure 1.6: Illustration of low and high LET radiation at the scale of the DNA double-helix, and the corresponding ionisation density pattern. Black lines represent particle tracks. Each ionisation event is marked in red.

a very small volume, it would be highly unlikely to observe a consistent amount of energy being transferred per track length. Instead, possible values for the LET follow a distribution, which is discussed further in Chapter 5. In order for the LET to be useful as a voxelised quantity and be mapped to patient CT coordinates, as is the case for dose, the distribution is replaced in practice by the average value. There are two main methods by which the LET is averaged; track-averaging (LET_t) and dose-averaging (LET_d). The track-averaged (equivalently, “fluence-averaged”) LET is defined as the average value of the electronic stopping power weighted by the energy spectrum of the protons present in the voxel of interest. That is,

$$\text{LET}_t = \frac{\int_0^\infty S_{\text{el}}(E)\Phi(E)dE}{\int_0^\infty \Phi(E)dE} \quad (1.15)$$

where $\Phi(E)$ gives the number of protons entering the voxel with a kinetic energy between E and $E + dE$. Dose-averaged LET is instead calculated by weighting each track by its contribution to the total dose due to all protons entering the voxel;

$$\text{LET}_d = \frac{\int_0^\infty S_{\text{el}}(E)\mathcal{D}(E)dE}{\int_0^\infty \mathcal{D}(E)dE} = \frac{\int_0^\infty S_{\text{el}}^2(E)\Phi(E)dE}{\int_0^\infty S_{\text{el}}(E)\Phi(E)dE}. \quad (1.16)$$

Even though the dose contribution from secondary protons is relatively small compared to primary protons, their contribution to LET_d is significant [58, 59]. These contributions have often been omitted in analytical calculations of the LET, but are included in Monte Carlo approaches [60, 61, 62, 59, 63]. The LET_d is more practically calculable and is therefore used in biological treatment planning, however factors such as track structure heterogeneity and energy straggling, particularly toward the proton end-of-range, cause considerable disagreement between the two calculation methods. Dose-averaging will always lead to a larger LET due to the uncertainty in the measurement of electronic stopping power. This is universally recognised, though, by this author’s understanding, has rarely been explained formulaically as




Figure 1.7: Dose, LET_t and LET_d depth profiles along the central axis of a 100 MeV monoenergetic proton pencil beam.

follows. The variance in a random variable X can, in general, be calculated using

$$\text{var}(X) = \langle X^2 \rangle - \langle X \rangle^2. \quad (1.17)$$

Applying this relationship to (1.15) and (1.16) yields

$$LET_d = LET_t \left(1 + \frac{\text{var}(S_{el})}{LET_t^2} \right). \quad (1.18)$$

Here, the LET_d has a clear and explicit dependence on the uncertainty in the measurement of the electronic stopping power and it is always true that $LET_d \geq LET_t$. Figure 1.7 demonstrates this disagreement between the two quantities and the tendency for LET to increase exponentially with depth toward the Bragg peak and distal dose fall-off. Monte Carlo simulation results for a 100 MeV monoenergetic proton pencil beam are shown. The methods by which this data was produced are described in detail in Chapter 4.

1.3 Proton radiobiology

1.3.1 Temporal action of ionising radiation

In previous sections we explored the physical interaction mechanisms by which protons, or indeed any charged particles, are slowed down and deflected as they traverse materials. The trajectories of these particles and the energy loss events they encounter along these paths are responsible for the spatial distribution of dose in the material. This physical dose is simply a measure of the amount of localised energy transferred to the material. Many decades of experience in conventional x-ray therapy has enabled physical dose values to be circumstantially linked to clinical outcomes, without necessarily needing to understand the exact biological mechanisms that caused these outcomes. Factors leading to biologic injury must be understood in order to resolve differences in clinical outcomes between photons and protons, or indeed any other novel radiation treatment. Understanding all the mechanisms at play, from physics, to chemistry, and to biology, allows clinicians to plan better treatments and improve patient outcomes.

The affects of ionising radiation on cells can be broken into a timeline of events, from the initial physical interaction causing ionisation of molecules, to chemical reactions caused by the breakage of molecular bonds, and eventually to biologic responses that follow from chemical changes. Physical interactions between the incoming radiation and cellular targets occur in a very small time frame, on the order of 10^{-15} seconds. In the likely event that a particle ionises a water molecule, as a cell consists mostly of water, the resulting H_2O^+ ion is short-lived ($\sim 10^{-10}$ s). It's chemical instability is a result of an unpaired valence electron. This makes this ion a *radical*⁶. The ion goes on to react with ordinary water to form a hydronium ion (H_3O^+) and

⁶Note that, in general, an ion need not also be a radical.

a hydroxyl radical ($\text{OH}\cdot$). These two chemical processes are summarised below:



where e^- is an ejected valence-shell electron. Additionally, the induction of an excited state of the water molecule, H_2O^* , can also result in hydroxyl production;



Hydroxyl radicals are highly reactive, with a mean lifetime of around 10^{-9} seconds. They diffuse through the cell only a short distance but this can be enough to reach, and interact with, critical cellular targets such as the DNA. The time between chemical effects and expression of biologic damage can vary significantly. This could be hours, days, years, or even generations in the case of hereditary defects.

1.3.2 DNA damage and repair

The deposition of energy within a cell is stochastic and thus so is the spatial distribution of initial ionisations. These initial ionisation events and subsequent ionisations caused by diffused radicals are able to induce damage in any cellular target. However, many of these targets, such as ribosomes, mitochondria, or cytoplasm, for example, have redundancies (multiple copies or rapid turnover) which mitigate the affects of molecular damage. By contrast, there are only two copies of DNA and turnover is relatively very slow. This lack of redundancy paired with the DNA's command of cellular functions makes DNA the main target for lasting or lethal damage [64, p. 9, 23] [65, p. 11] [66, 67].

The DNA double-helix consists of two opposing strands of a sugar-phosphate

backbone connected at regular intervals to purine (adenine or guanine) or pyrimidine (cytosine or thymine) bases. The strands are linked to each other by hydrogen bonds between the bases, however adenine must always pair with thymine and guanine must pair with cytosine. There is seemingly endless depth to which we could discuss DNA biology, from gene expression to oncogenesis, however for the purposes of this thesis we restrict our discussion to only a few key ideas.

DNA damages may occur from *direct* or *indirect* ionisation or excitation of the DNA. Direct damage entails the particulate radiation interacting with the DNA along its trajectory. Indirect damage occurs when water surrounding the DNA (within a radius about twice that of the DNA itself [64, p. 10]) is first ionised and resulting chemically reactive radicals diffuse to, and interact with, the DNA. Hydroxyl radicals, produced through the reactions shown in (1.19), (1.20) and (1.21), are the most reactive with the DNA compared to other radical species, accounting for up to two-thirds of damage in conventional x-ray therapy [68]. The consequences of these indirect damages are, in general, highly dependent on the amount of “fixation” by molecular oxygen, which renders the lesion permanent. If there is an absence of molecular oxygen near the damage site in the microseconds that follow initial ionisation, DNA radicals formed by indirect events are more readily repaired. However, if the DNA radical reacts with oxygen to form a peroxide then the DNA is not able to be repaired at this site. This is known as the *oxygen fixation hypothesis*. The difference between hypoxic (very low to null oxygen concentration) cells and aerobic, or normally oxygenated, cells has been shown to be significant in the case of x-rays, with up to a three-fold difference in the eventual damage yield [see, e.g., 69, 70, 71, 72, 73]. This multiplicative factor by which the presence of oxygen tends to increase the severity of injury is known as the *oxygen enhancement ratio (OER)*. Explicitly, it is most often quantified as the ratio of doses $\mathcal{D}([\text{pO}_2]^a)/\mathcal{D}([\text{pO}_2]^b)$ that produce the same biological or clinical outcome for different molecular oxygen percentage concentrations




Figure 1.8: Illustration of DNA strand breaks (in orange) caused by indirect and direct damage events, with each ionisation event marked in red. A primary particle is incident from the left, giving rise to a secondary electron. Indirect damage occurs by ionisation of nearby water molecules which results in hydroxyl radicals ($\text{OH}\cdot$) then ionising the DNA. Direct damage occurs by ionisation of the DNA by the incident particle itself. Multiple strand breaks within 10 base pairs are labelled double strand breaks (DSBs), which may be due to a single particle track (bottom track shown) or multiple.

$[\text{pO}_2]^a \leq [\text{pO}_2]^b$. The OER is typically more significant for low LET radiation such as x-rays, for which indirect damage events dominate, compared to higher LET radiation such as low-energy protons. We explore this in some more detail in Chapter 3, including in the context of a novel treatment option known as *FLASH* radiation therapy, which is increasingly gaining traction in the radiotherapy community.

Whether by direct or indirect ionisations, the resultant DNA radicals may ultimately lead to lasting DNA damage if not repaired by cellular mechanisms. Types of damage fall under a number of classifications however only some are pertinent in radiation therapy. For example, low energy ultraviolet photons (approximately 4 to 30 eV) are able to cause adjacent thymine and cytosine bases to covalently bond together causing a “kink” in the DNA. These kinks, properly referred to as pyrimi-

dine dimers, are readily repaired through nucleotide excision repair (NER) and are of negligible consequence relative to clustered ionisations at the deoxyribose (sugar) or phosphate groups comprising the DNA backbone. In general, isolated base damages from ionising radiation are repaired successfully through base excision repair (BER). Conversely, strand breakages can have significant consequences, the severity of which is related to their spatial proximity. Single strand breaks (SSBs) do not often lead to cell kill because the opposite strand is able to act as a template for repair [66]. Double strand breaks (DSBs) may be formed by single or multiple particle tracks, or multiple indirect damages, within ten base pairs, or one helical turn of the DNA [74]. Although, at therapeutic doses the probability of a DSB being induced by two or more independent particle tracks can be considered negligible; a dose of larger than 10^4 Gy would be required for such an occurrence to become likely [75]. Contrary to isolated SSBs, these DSBs can cleave chromatin into distinct pieces at the damage sites, making it much more difficult to repair. Figure 1.8 illustrates strand breaks caused by indirect and direct damage events. Multiple strand breaks within ten base pairs, as shown in the figure, are labelled DSBs. Subcategories of DSBs are often used to distinguish the density of clustered lesions. Using the classification system of Michalik and Frankenberg [76], a *simple* DSB is labelled as consisting of two opposing SSBs with at most one additional base damage while a *complex* DSB is one that consists of multiple SSBs and base damages within a close vicinity. Complex DSBs are commonly associated with high LET radiation as the density of ionisation events is greater than for low LET radiation. It should be noted that labels of damage complexity have differed historically, with some being strictly guided by the number of strand breaks [see, e.g., 77, 78, 74] while others include base damages and whether the damages were directly or indirectly induced [see, e.g., 76, 79]. In fact, Schipler and Iliakis [80] describe in great detail six different complexity labels for DSBs, though our two-level classification system is all that is needed as context for this thesis.

A dose of 1 Gy due to x-ray radiation typically results in a relatively small number of DSBs (up to about 40), twenty-five times fewer than the number of SSBs [64, p. 13], however it is primarily due to the induction of these DSBs in nuclear DNA that approximately one-third of exposed mammalian cells are killed [65, p. 13]. There are two major DSB repair pathways; *homologous recombination (HR)* and *non-homologous end-joining (NHEJ)* [81, 82]. HR is an error-free repair process that occurs in late-interphase (S or G2 cell cycle phases). It avoids errors by requiring physical contact with the undamaged sister chromatid, which acts as a template for repair. This process cannot occur in early-interphase (G1), prior to replication, as there is no template chromatid present. Instead, early-interphase NHEJ repair involves rejoining the broken DNA strands. However, NHEJ is prone to errors as there is no “ground-truth” template for repair, and small base sequence variations can occur near the site of the DSB. It is interesting to note that the presence of a sister chromatid after replication does not guarantee the error-free HR process. In fact, both HR and NHEJ compete. The regulation of the necessary repair proteins is in part related to the recruitment of 53BP1 proteins [83], which is explained in some further detail in Chapter 3, Section 3.4. In each case, unpaired bases at the breaks tend to join together when in proximity, causing chromosome aberrations. Lethal aberrations usually fall under *dicentric*s, by which two pre-replication chromosomes join at their breaks, *rings*, by which a single pre-replication chromosome joins to itself at two different break sites, or *anaphase bridges*, by which two chromatids post-replication join at their breaks.

While our discussion here is not intended as a comprehensive take on the biological processes at play, it is nonetheless important to highlight the links between the damages to DNA and the eventual chromosomal effects that then contribute to cell death or dysfunction. DNA, with a diameter of less than 2.5 nm, is tightly wound around histones to form a densely packed chromatin structure. The width of the

entire chromosome in metaphase is more than five hundred times larger than the DNA itself [64, p. 24]. DSB damages in the DNA can be the ultimate cause of chromosome aberrations, which interrupt regular cell division. Complex DSB damage may occur early in interphase, causing unrepaired breaks to be replicated when an identical strand is produced, or they may occur in one or both of the chromatids in late-interphase.

The possible DNA damage response (DDR) pathways are highly interrelated and complex, comprising *sensor* proteins that survey the genome for damages and subsequent *effector* pathways that dictate the action that the cell takes as a response to radiation-induced damage [65, p. 14]. The cell may repair itself, become blocked in its progression through the cell cycle, or undergo a programmed (*apoptotic*) death. Sufficient trauma to the cell may even cause a less orderly death known as *necrosis*.⁷

1.3.3 The linear quadratic model for cell survival

Cell survival curves relate the radiation dose received by cells to the fraction of those that survive damage caused by irradiation. The definition of cell death, however, is actually rather subjective, so in the context of radiobiological studies *reproductive cell death* is often the endpoint used. A cell that has lost its ability to proliferate sustainably and form a large colony over at least several mitotic cycles is not *clonogenic*, has lost its reproductive integrity, and is considered dead. This is despite some cellular mechanisms remaining functional [84], [64, p. 35]. *In-vitro* radiobiological studies, meaning those which are conducted in the lab, outside of the natural biological environment, rather than in living beings (*in-vivo*), often use reproductive cell death as the measured endpoint. Through the preparation of cell lines that vary in type, function, and host (e.g. human or animal) and the harvesting of specimens from normal tissues or tumours, one may test a range of hypotheses related to the cellular

⁷The cell may also die by way of senescence or autophagy, however these pathways are arguably much less relevant in a discussion of injury as a result of radiotherapy.

response of radiation. These studies involve irradiating a cell culture and comparing it to a control through *clonogenic survival assay*, which is discussed in Chapter 3 (Section 3.4).

The linear-quadratic (LQ) model is one of the simplest, yet most ubiquitous, models used in radiation biology. In this model, the surviving fraction, $0 \leq S \leq 1$, of cells exposed to a single dose, \mathcal{D} , of radiation is given by

$$S = \exp(-\alpha\mathcal{D} - \beta\mathcal{D}^2) \quad (1.22)$$

where α and β are two constant fitting parameters that indicate radiosensitivity. The linear response term, α , describes the initial slope of the shouldered curve and represents cell death due to “single-hit” events while β describes the quadratic component and represents cell death due to second-order “multiple-hit” events [see, e.g., 85]. A single-hit event is one in which lethal damage to the cell is caused by a single event. The higher order $\mathcal{O}(\mathcal{D}^2)$ term then may refer to the combination of multiple sub-lethal and potentially lethal damage eventuating in cell death [86], which on their own would not have been lethal. In the case of a proton beam, one could consider intratrack damage (induced by a single proton track) to scale linearly with dose and intertrack damage (caused by multiple tracks) to scale with polynomial dose terms. Intratrack contributions can be intuitively understood by taking the limit of (1.22) as $\mathcal{D} \rightarrow 0^+$, in which the probability of two or more damage sites being co-located is minimal. Pragmatically, truncating the model to the leading linear and quadratic terms sufficiently describes empirical data. While some studies in the literature have had an empirical curve-fitting focus [see, e.g., 87], others have applied more detailed mechanistic approaches (see [88], [89] and a recent review by McMahon [85]).

The concept that eventually led to the LQ model can be traced back to the 1920s, when an assumption was made that lethal irradiation-induced damage events were

Poisson-distributed in proportion to the dose received by the cell [90]. Under this assumption, one could hypothesise that the probability of a cell surviving a single delivered dose of \mathcal{D} is proportional to $e^{-\mathcal{D}}$. This is because survival is interpreted as the cell having experienced zero lethal events. After the advent of modern clonogenic assay techniques in the 1950s [91] however, more and more empirical data showed a shouldered response curve [92] and led to doubts in the purely log-linear relationship. Formulations similar to what is now known as the LQ model were subsequently proposed, and although it is difficult to trace the model back to a single source [85], one could perhaps attribute the pioneering work to Neary [93] and Kellerer and Rossi [88].

The α/β ratio, in units of gray (Gy), is the dose at which the linear and quadratic contributions to cell death are equal. This gives an indication of the curvature and can therefore also be used as an indicator of the radiosensitivity of different tissues. Early responding tissues, such as skin (epidermis) and intestine, have a large α/β ratio of around 10 Gy while late responding tissues, such as the spinal cord and lung, have a smaller α/β ratio of around 2 or 3 Gy.

1.3.4 Fractionated radiotherapy

Having just discussed survival curves that describe the cellular response to a single radiation dose, we must address the fact that almost all treatments are *fractionated*. That is, the total dose is spread out over time into discrete treatment sessions. Typical cell survival curves for fractionated and single-dose treatments are shown in Figure 1.9. The guiding principle of fractionation is that differences between the radiosensitivity of the irradiated tissues, or even between the cells within a tumour, can be exploited by allowing sufficient time for sublethal damage sites to be repaired between treatment sessions. Cancer cells are typically more radioresistant than the surrounding healthy tissue so the benefit of repair is biased toward those healthy

tissues without compromising the overall tumour control. The benefits of this treatment regime were first noticed by Regaud in 1911; delivering small dose irradiations to ram testes 15 days apart achieved sterilisation without burns to the skin, which were otherwise observed in single-dose delivery [94, 95]. However, it took until the late 1920s for the radiotherapy community to come to a consensus on the benefits of fractionation; some likened small dose fractions to “weak” radiation and that efficacy would be decreased [96]. Early proponents of fractionated treatment speculated that discrepancies in cellular metabolic rates could explain differences in radiosensitivity. A number of skin-sparing observations, similar to those of Regaud, began to surface over the following decade [12, 97, 98]. Total treatment time, in addition to the number of fractions, was later found to affect cellular response (again, experimenting with skin) [99, 100] which led to attempts at biological iso-effect modelling [101, 102, 103, 104, 105]. Such models were used to determine *tolerance doses*

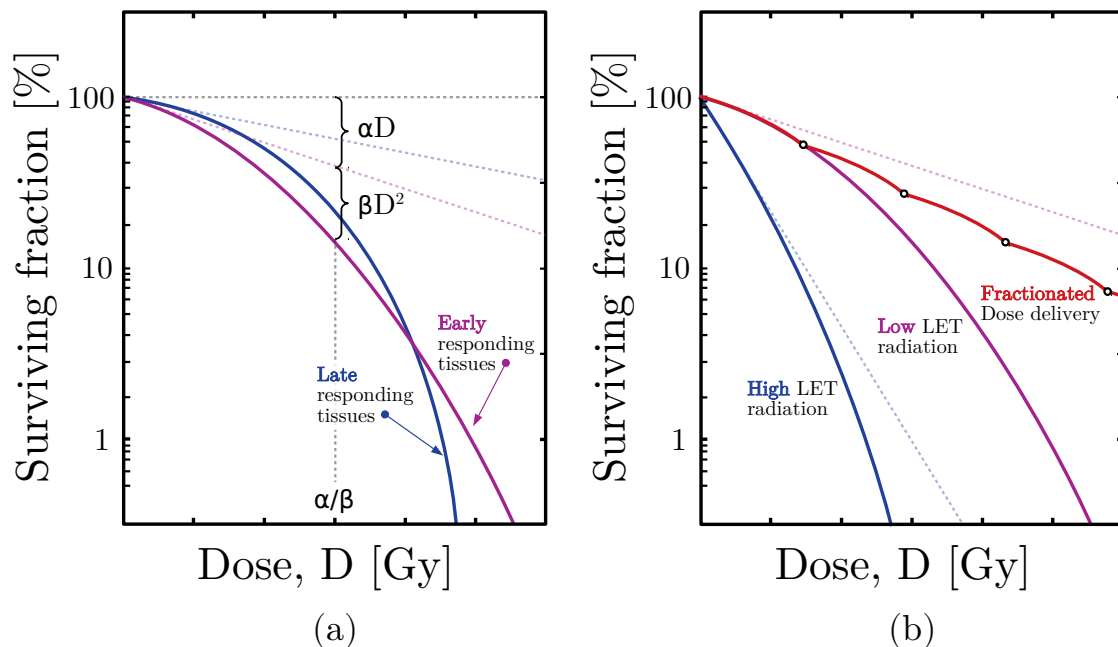


Figure 1.9: (a) Examples of a linear quadratic (LQ) cell survival curves for early and late responding tissues. The radiosensitivity of the tissue is characterised by α/β , measured in units of Gy, which is the dose at which the linear and quadratic components of the curve are equal. (b) Examples of cell survival trends for low and high LET radiation, and a fractionated low LET treatment regimen. The linear component, or initial shoulder, to the curve is shown as a dashed line of the same colour.

for healthy tissues, or the dose beyond which would likely cause organ failure. Ellis [103, 104] introduced the nominal standard dose (NSD) model, building upon a power law relationship between tolerance dose and total irradiation time, which was inspired by photochemistry [101, 102]. The NSD model, however, was based largely on empirical data from skin experiments and could not be easily generalised to other normal (non-cancerous) tissues (see [105] and references therein). It was eventually replaced with a thorough LQ-based formalism by Barendsen [105].

If the total dose, \mathcal{D} , of a treatment is to be delivered in N equal fractions then the dose per fraction is $d = \mathcal{D}/N$ and the final surviving fraction is predicted by the product of the LQ predictions for each individual fraction. That is,

$$S(\mathcal{D}) = [S(d)]^N = \exp(-\alpha dN - \beta d^2 N) = \exp \left\{ -\alpha \mathcal{D} \left(1 + \frac{d}{\alpha/\beta} \right) \right\}. \quad (1.23)$$

More generally, for fractions of different doses, which is uncommon in practice, if the n -th fraction ($n = 1, 2, \dots, N$) delivers a dose of d_n then

$$S(\mathcal{D}) = \prod_{n=1}^N \exp(-\alpha d_n - \beta d_n^2) = \exp \left\{ -\sum_{n=1}^N \alpha d_n \left(1 + \frac{d_n}{\alpha/\beta} \right) \right\}. \quad (1.24)$$

The influence of the α/β ratio on the shape of the survival curve can be appreciated in (1.23) (or (1.24)).

Withers [106] and Steel [107] summarised five cellular processes and characteristics that govern fractionation effects, popularised as the “5 Rs of radiobiology”⁸, which apply both to normal tissues and to tumours. They are:

1. **radiosensitivity.** Cells and tissues vary in their response to radiation. They can be early- or late-responding. Tissues in which cellular damage manifests after long delays are more sensitive to fraction size than most tumours or acutely responding normal tissues [108].

⁸Not to be confused with “5 R’s” (with an apostrophe), to which Steel was strongly opposed.

2. **Repair.** Fractionated dose delivery allows for sublethal damage repair between treatment sessions. Capacity for such repair is generally greater for healthy tissues than for tumour cells.
3. **Repopulation.** Clonogenic tumour cells that survived the previous fraction of radiation can increase their rate of proliferation and give rise to more cells that must be killed in the following fraction. Increasing the time between treatments accentuates this undesirable effect.
4. **Redistribution.** Irradiation of cells can activate DNA damage checkpoints, defined points in the cell cycle at which cells may be blocked or delayed in their progression to the G1, S and G2 phases. As radiosensitivity is dependent on the phase of the cell cycle, fractionated dose delivery allows for redistribution of phases in the cell populations, thereby mitigating the affects of radioresistant phases.
5. **Reoxygenation.** Tumours are mostly hypoxic, meaning they contain a very small amount of molecular oxygen, due to their poor vascular structure. In the absence of oxygen, fewer damage sites are molecularly fixed, resulting in a smaller number of permanent damage sites. Efficacy of treatment is therefore reduced. Fractionation allows for reoxygenation of the tumour as aerobic cells are killed and the gross volume decreases in size.

Potential relationships between these “5 Rs” and the Hallmarks of Cancer [109] have been discussed by Harrington [110].

1.3.5 Relative biological effectiveness

While protons were initially proposed in 1946 for their favourable depth-dose distribution [28], their enhanced radiobiological effectiveness has since been the subject of

further consideration. Physical dose deposition does not on its own adequately provide a generalised description of tumour control or complications arising in normal tissues. In fact, protons are more effective than x-rays at cell kill per unit of dose deposited [111]. As we have explored in this Section, the translation of initial physical damage to subsequent chemical and biological processes that eventuate in cellular injury is complex and depends on a number of factors. Densely ionising particle tracks offer an increased cell-killing efficiency over sparsely ionising x-rays, which may be quantified through various biological endpoints, such as the LET. A data-rich history of conventional radiation therapy has made it convenient to summarise inter-modality radiobiological differences as a weighted x-ray dose equivalent, referred to as the *relative biological effectiveness (RBE)*. The RBE is defined as the ratio of doses required by a reference radiation, such as a megavoltage x-ray beam, and a test radiation, such as 200 MeV protons, to yield the same biological effect. That is,

$$\text{RBE}_{\langle\text{endpoint}\rangle} = \frac{\text{Dose of reference radiation}}{\text{Dose of test radiation}}. \quad (1.25)$$

Values obtained for RBE are restricted to the biological *endpoint* chosen. For example, a commonly used endpoint is the number of surviving cells in a culture by clonogenic survival assay following irradiation. However, values inferred from cell culture experiments in laboratories can differ significantly from those inferred from animal or human subjects. An in-depth discussion on the consequences of the chosen RBE endpoint is provided in Chapter 3.

The RBE depends not only on radiation type and biological endpoint, but on dose, LET, fractionation regimen and tissue type. Combining this complexity with that of accurate measurement and experimental reproducibility leads to a large RBE uncertainty. In 2007 the ICRU published guidelines for dose prescription in proton therapy [112], in which they stated that “*clinical experience has not been interpreted*

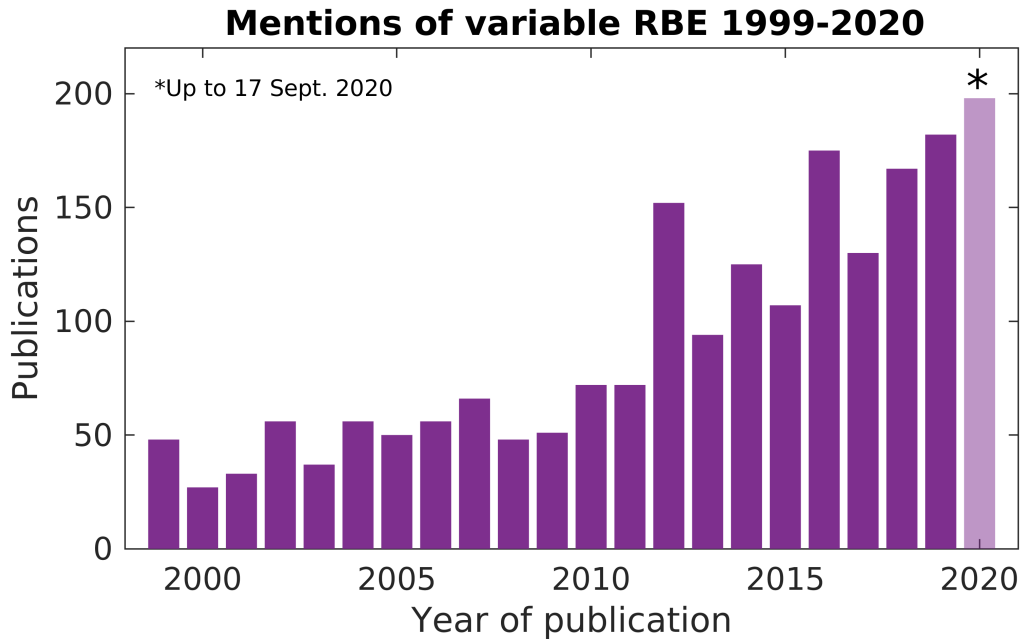


Figure 1.10: The number of publications discussing variable RBE has noticeably increased over the past two decades. Based on search of data available on [sciencedirect.com](https://www.sciencedirect.com) on 17 September 2020 with the following search strings: ("relative biological effectiveness" OR "RBE") AND (proton OR protons) AND ("variable" OR "biophysical" OR "nonconstant" OR "non-constant" OR "changing").

as indicating that RBE is different from 1.10". This recommendation of uniformly reducing the dose to 90% of the x-ray prescription has since been the clinical norm, and experimental data loosely supports that, on average, proton RBE is approximately 1.1 [113]. However, there has been a push toward variable RBE modelling for treatment planning as the RBE may be significantly greater than 1.1 toward the distal edge of the Bragg peak [see, e.g., 114, 115, 116, 117, 118, 119]. In fact, the number of publications discussing variable RBE has noticeably increased over the past two decades (see Figure 1.10).

The surviving fraction of irradiated cells in laboratory experiments was mentioned above as the most commonly used RBE endpoint. Following our earlier discussions of cell survival curves, we are now well-placed to quantify the RBE by surviving fraction, denoted RBE_{SF} :

$$RBE_{SF} = \frac{D_x}{D} \quad (1.26)$$

where the proton dose, \mathcal{D} , is given by

$$\mathcal{D} = \begin{cases} -\frac{1}{2} \left(\frac{\alpha}{\beta} \right) + \frac{1}{2} \sqrt{\left(\frac{\alpha}{\beta} \right)^2 - \frac{4}{\beta} \log S}, & \beta > 0, \\ -\frac{1}{\alpha} \log S, & \beta = 0, \end{cases} \quad (1.27)$$

and the equivalent x-ray dose, \mathcal{D}_x , is determined by the same formulae, substituting for α_x , β_x and S_x . The surviving fraction, S , is given by (1.22). It can be shown that

$$\text{RBE} = \frac{1}{2\mathcal{D}} \left[-\left(\frac{\alpha}{\beta} \right)_x + \sqrt{\left(\frac{\alpha}{\beta} \right)_x^2 + 4\mathcal{D} \left(\frac{\alpha}{\beta} \right)_x \frac{\alpha}{\alpha_x} + 4\mathcal{D}^2 \frac{\beta}{\beta_x}} \right]. \quad (1.28)$$

In Chapter 3 and Chapter 4 we explore the RBE concept in detail and propose a biophysical model for quantifying the RBE in treatment planning systems.

1.4 Intensity modulated radiation therapy

Prior to a patient receiving external beam radiotherapy, a patient-specific treatment plan must be produced using computer software. Geometrical structures critical to the success of the treatment are first delineated (contoured) over CT images of the patient, which may be registered with other modalities such as magnetic resonance imaging (MRI) or positron emission tomography (PET). The most relevant regions of interest (ROIs) that must be delineated for every plan, according to the ICRU in Reports 78 [112] and 83 [120], are:

- **Gross tumour volume (GTV):** the visible or directly detectable volume of the primary tumour and any metastases;
- **Clinical target volume (CTV):** the tissue volume relevant for therapy, containing the tumour and any immediate microscopic extension (e.g. to lymph nodes) of the malignant disease;

- **Planning target volume (PTV):** a geometrical construction containing the CTV. Margins are applied in accordance with the ICRU guidelines [121, 122] to ensure the CTV receives at least 95% of the prescribed dose in 90% of patients when accounting for setup errors and position uncertainty;
- **Organ at risk (OAR):** normal tissues susceptible to radiation damage which may require alteration of the treatment plan or prescription.

The planning objective is to make the best possible compromise between delivering adequate lethal dose coverage to the PTV and minimising the dose to any OARs. A case may be forward-planned, by which the planner alters factors such as the beam angles, energy and relative intensity before calculating the dose and assessing the overall treatment quality, or inverse-planned, by which a set of desired constraints and targets for the dose distribution is defined, and the number of fields (beams) and their orientations are chosen. The former manual optimisation technique is used in conventional 3D conformal radiotherapy (CRT) while the latter is used in intensity modulated radiotherapy (IMRT). CRT has historically been the dominant technique in photon treatment delivery, with IMRT coming later to the commercial market in the 1980s to 1990s [123, 124]. Conceptualised in the 1960s, IMRT offered more precise beam-shaping than the static field-shaping methods of CRT, thereby enabling highly conformal dose distributions to be achieved. However, its potential would not be fully recognised until decades later due to limitations in computing power; complex inverse planning algorithms were simply too slow and demanding for the majority of computer systems at the time. It has since become the standard delivery technique for many tumour sites [125, 126, 127, 128, 129, 130].

In IMRT planning, whether photon- or proton-based, the central concept is to divide the radiation field defined for treatment into individual beamlets whose relative fluence can be altered to more accurately shape the dose pattern. For photon IMRT, the field aperture of the treatment beam is discretised into individual beamlet

contributions to the field, often called “bixels”, which are modulated through the use of compensators and a multileaf collimator (MLC). In the case of intensity modulated proton therapy (IMPT), the intensities of individual proton pencil beams are modulated. In either modality there are at least two distinct optimisations involved in the planning procedure. The first is to determine suitable beam angles and, for protons, also the beamlet energies as these determine the Bragg peak locations. The second step is a *fluence map optimisation* (FMO), a focus of this thesis (Chapter 6 and Chapter 7), in which the beamlet intensities are modulated using inverse planning algorithms. Photon IMRT involves a third optimisation step to determine both spatial and temporal MLC leaf-sequencing required to realise the chosen intensity pattern.

The use of IMRT over 3D CRT has enabled a shift in clinical paradigms. The ability to achieve highly conformal dose distributions has naturally been met with higher standards of healthy tissue sparing, albeit while still maintaining tumour control. Dose-escalation and hypofractionation have also been topics of interest, in which the approach to planning is based more on keeping toxicity levels acceptable, rather than as low as possible, while delivering a higher dose to the tumour compared to that which would usually be prescribed [124].

We highlight, as an example, an argument for the use of IMRT techniques for the treatment of head-and-neck cancer. This treatment site involves a number of OARs in close proximity to the target volume and, therefore, demands precise shaping of the beam. The severity of dysphagia and xerostomia due to irradiation of the salivary glands is strongly attributed to long term quality of life, with adverse affects arising in swallowing ability, taste and speech [131, 124]. Clinical research has demonstrated improved preservation of the salivary glands and fewer instances of these adverse affects when using IMRT over conventional CRT methods [132, 131, 133].

1.4.1 Proton pencil beam scanning

Conventional IMRT involves source megavoltage photons through a linear accelerator (LINAC). In contrast, high energy protons are produced through cyclic particle accelerators such as cyclotrons or synchrotrons prior to being directed through a beam transport system. In IMPT a three-dimensional dose distribution is obtained dynamically via superposition of many individual pencil beams, by a method known as *active scanning*, also called *pencil beam scanning*. Through the acceleration of charged particles to a desired energy and their deflection by use of electromagnets, each pencil beam can be tuned such that the target volume is irradiated slice-by-slice according to a pre-calculated path. This concept is illustrated in Figure 1.11. Proton therapy clinics have increasingly opted for the active scanning technique over older *passive scattering* techniques, which shape the dose distribution via physical filters, modulators and scatterers. Active scanning IMPT has been demonstrated to achieve both improved sparing in healthy tissues and more accurate high-dose delivery to target volumes when compared to passively scattered protons [134, 135, 136, 137, 138]. However, both delivery methods bring their own merits. Passive delivery systems are less complex, require fewer adjustments to the beam source and can reduce the total treatment time, but they also produce increased neutron contamination in the beam, are less efficient and are not as suitable for larger or more geometrically complex target volumes [139]. Pencil beam scanning is able to achieve more conformal target coverage, however the time-component of the scanning method can amplify *interplay effects*, in which random positioning errors due to organ motion contribute to localised regions of excessive, or insufficient, dose. Techniques have been developed to reduce interplay effects, including averaged layer re-scanning, patient immobilisation and breathing management [see, e.g., 140, 141]. Despite the clearly different dose delivery methods between photon and proton IMRT, inverse plan optimisation for

both modalities is of roughly the same order of complexity [142].

1.5 Radiobiological considerations in treatment planning

In order to best achieve therapeutic effectiveness and mitigate undesired effects, biological factors must be considered during treatment planning. As more clinics around the world adopt conformal non-conventional therapies such as external beam proton and heavy-ion IMRT, biological equivalence metrics such as RBE (Section 1.3.5) become vital in accounting for radiobiological differences between radiation modalities. However, the RBE is not all-encompassing; it simply allows for the physical dose due to non-conventional modalities to be assessed under the same framework that has been used for x-ray treatments for many decades, over which a wealth of clinical experience has been accumulated. This is a crucial first step that provides the photon-equivalent dose map, in units of Gy(RBE), but it does not necessarily inform us of how to best prescribe the dose, shape it or fractionate its delivery.

The *equivalent dose (EQD)* is an isoeffect metric, based on the LQ model, which converts the total dose delivered in a prescribed fractionation regimen to a common




Figure 1.11: Simplified illustration of active pencil beam scanning for IMPT dose delivery. Proton pencil beams of a nominal energy pass through a range shifter before being modulated by orthogonal electromagnets, such that the resulting Bragg peak is positioned at the required location in the target.

reference (2 Gy equal fractions, for example) by equating (1.23) and (1.24). In general, the equivalent dose, for the same amount of cell kill, deposited over equal fractions of d_{ref} is given by

$$\text{EQD} = \frac{\sum_{n=1}^N d_n \left(1 + \frac{d_n}{\alpha/\beta}\right)}{1 + \frac{d_{\text{ref}}}{\alpha/\beta}} \quad (1.29)$$

where d_n is the actual dose deposited in the n -th fraction and N is the total number of fractions.

The irradiation volume of an organ and the spatial heterogeneity of the dose it receives have important biological consequences. These are known as *dose-volume effects*.

1.5.1 Dose-volume effects

A naive approach to treatment planning would entail a simple regime; the dose in each organ at risk (OAR) would be minimised while maximising tumouricidal dose. This approach, however, is unsatisfactory for two main reasons. First, excessive radiation dose can lead to large amounts of cell “debris” as necrosis dominates apoptotic, or programmed, cell death; necrotic injury does not signal phagocytosis in neighbouring cells, limiting the natural healing process of the body. Second, cancer cells intersperse themselves among healthy tissue, so increased cellular injury to the tumour is met with increased toxicity to the healthy tissue [143]. A balance must be struck between risking tumour recurrence with too low of a dose and causing excess injury with too high of a dose. An appropriate regime for maximising therapeutic effectiveness is therefore to deliver a homogeneous, or uniform, dose to the target region [144], which was formalised by the ICRU in Report 50 [121]. The treatment planning process would then involve setting what we shall refer to as *hard dose constraints* (HDCs), comprised of upper dose bounds on OARs and both lower and upper dose bounds on target regions. In practice, radiation fields are heterogeneous, particularly within healthy

tissues surrounding the target volume that receive dose “spillage”. Consequently, it is important to understand how partial irradiation of an organ influences pathogenesis [145, 146, 147] and to incorporate such dose-volume effects in treatment planning.

Withers *et al.* [148] introduced the concept of functional subunits (FSUs) to relate dose-volume effects to tissue architecture. An FSU is the largest subvolume of a tissue that can be restored from a single surviving stem or clonogenic cell. They may be arranged in *parallel* or in *series* (or in a combination of the two), depending on the organ, with starkly different clinical consequences.⁹ The spinal cord, oesophagus and intestines are pertinent examples of structures with serial FSU organisation, and are sometimes referred to as “tubular” or “chain” organs [147, 143]. Uniform distribution of dose throughout the organ can, in general, be tolerated however localised “hot spots” of excessive dose that destroy one FSU can compromise the functionality of the organ as a whole. In contrast, parallel (or “rope”) organs such as kidneys, lungs or the liver may retain their overall functionality even when some individual FSUs lose their function. However, relatively low dose distributed through the entire volume can cause the organ to lose its overall functionality [143]. In reality the distinction between serial and parallel organs is not well-defined; some biological functions are affected in a parallel manner while others are affected in a serial manner.

Dose-volume considerations have made their way prominently into the clinic. In fact, irradiation volume is of such importance in determining normal tissue complications that it has been paired with the “5 Rs of radiation biology”¹⁰ [106, 107] described in Chapter 1 [145]. The cumulative *dose-volume histogram* (DVH) is widely used to assess plan quality. The DVH shows the proportion of a defined structure that receives at least a given dose level. Though it is not a metric for the spatial distribution of dose, it does allow the planner to consider both hard dose bounds and

⁹The “series” (or “serial”) and “parallel” naming convention is analogous to that used to describe the interdependence of components in electrical circuits as a current is applied.

¹⁰Unfortunately “irradiation volume” does not begin with an “R” but this hasn’t stopped the radiation biology community from co-opting it as an addition to the “5 Rs”.

dose-volume effects using a single graphic. Clinical evidence accumulated over many years has enabled planners to set desired control points to help shape DVH curves. These control points consist of the usual HDCs as well as *dose-volume constraints (DVCs)*, which specify an amount of permitted *overflow* or *underflow* from a prescription dose. A DVC applied to an OAR may be denoted, in the usual notation, $D_{X\%} < Y$, which indicates that the dose to $X\%$ of the structure must not exceed Y Gy. Alternatively, a DVC applied to a target region may be denoted $D_{X\%} > Y$, which indicates that the dose to $X\%$ of the structure must be at least Y Gy. The inclusion of DVCs into planning has shown demonstrable benefits for a number of treatment sites. For example, evidence suggests that late rectal toxicity is reduced through the use of DVCs for prostate cancer treatment with photon IMRT [149] and, more recently, dose-volume parameters for the small bowel have been found to be significantly predictive of toxicity in photon IMRT treatment of rectal cancer [150]. In 2010 the QUANTEC (Quantitative Analysis of Normal Tissue Effects in the Clinic) initiative published guidance on normal tissue dose tolerances with a focus on volume effects [151] (see also [152]), and in 2019 the PENTEC (Pediatric Normal Tissue Effects in the Clinic) initiative proposed further dose-volume reporting standards for childhood cancer outcomes [153].

Although imposing DVCs has therapeutic advantages, their implementation in inverse plan optimisation is complex and optimal solutions are not necessarily guaranteed. This is explored in detail in Chapter 7 in which a new inverse planning algorithm is presented for handling DVCs accurately and appropriately.

1.5.2 Levels of radiobiological optimisation

Nahum and Uzan [154] categorised radiobiological optimisation into five distinct levels of increasing sophistication:

- (I) **Level I.** Isotoxic individualisation of the prescribed dose at a fixed number of

fractions;

(II) **Level II.** Incorporation of a fractionation scheme;

(III) **Level III.** Use of radiobiological functions in the inverse planning framework;

(IV) **Level IV.** Inclusion of functional imaging (excluding for the purposes of delineating volumes).

Almost all current clinical protocols will fall under a Level II optimisation, including organ tolerances, fractionated dose delivery and dose-volume considerations. In proton therapy where RBE plays a role, as we have discussed in Section 1.3.5, a constant multiplier of the dose (increasing it by 10% everywhere) is often applied. This could, by definition, be considered a Level III optimisation but in practical terms it is a small improvement from Level II. However, variable RBE models, which we explore in Chapter 3, certainly belong to Level III if they are used during inverse planning. These higher level radiobiological frameworks are not exclusive to particle therapies. In fact, in standard photon IMRT a Level III plan may be achieved through inclusion of EQD ((1.29)), for fractionation effects, and a quantity known as the (*generalised*) *equivalent uniform dose (EUD)* [155, 156], for dose-volume effects. The EUD is sometimes used as a mathematically simple, albeit less accurate, substitute for DVCs (Level II). It is calculated using

$$\text{EUD} = \left(\sum_k v_k \mathcal{D}_k^a \right)^{\frac{1}{a}} \quad (1.30)$$

where a is a tissue-specific parameter reported in the literature based on observational dose-volume effect data [see, e.g., 157, 152], and v_k is a defined fractional organ volume that receives a dose of \mathcal{D}_k . The summation in (1.30) is over all fractional volumes. In the simplest, and most common, scenario these fractional volumes are the voxels defined in the dose grid or planning CT and all are of the same size. The

EUD, and other radiobiological quantities, can be used to calculate tumour control probability (TCP) and normal tissue complication probability (NTCP) [see, e.g., 158]. All treatment planning systems implicitly seek to balance therapeutic effectiveness and toxicity, though some do so through explicit maximisation of the therapeutic ratio, TCP/NTCP.

A plan can be raised to a Level IV optimisation through additional considerations of, for example, clonogen location or labelling of hypoxic regions. For example, pharmacokinetics of radionuclides imaged with PET can be used to map levels of molecular oxygen in the tumour, which can then be used to apply boosted doses to hypoxic regions that are typically more radioresistant than well-oxygenated regions. This concept is explored in more detail in Chapter 3.

1.6 Scope of this thesis

This thesis investigates the importance of radiobiological factors in radiation therapy planning, with a focus on proton therapy, and proposes novel methods and recommendations for the inclusion of such factors in treatment planning.

The structure of this thesis is segmented into two distinct investigations; Chapter 3 to Chapter 5 detail the development and testing of a new biophysical model for quantifying the relative biological effect (RBE) of dose in proton therapy planning. Chapter 6 and Chapter 7 formulate and test an inverse optimisation algorithm for treatment planning of intensity modulated radiation therapy (IMRT), agnostic to the chosen radiation modality, which provides comprehensive handling of dose-volume effects in the spatial distribution of dose in the patient. A brief overview of the methods undertaken is provided in Chapter 2.

This author hypothesises that inclusion of DNA damage modelling and comprehensive plan optimisation strategies, as a result of the methods presented in this

thesis, will enable confident and effective plan creation not only in terms of physical dose deposition but, crucially, in biological effect.

Chapter 2

Overview of methods

2.1 A double strand break induction model for RBE prediction

2.1.1 Theoretical development

In Chapter 1 a case was outlined for the use of variable RBE modelling in proton radiation therapy, in order to accurately weight proton doses to x-ray dose equivalents, for which patient outcomes are much better understood due to a wealth of clinical experience. Further, the mechanisms by which such an enhanced cell-killing efficiency may result from protons were discussed. Efficiencies in DSB induction in nuclear DNA were identified as the major source of an increase in the RBE toward the proton end-of-range. Therefore, in Chapter 3 a simple biophysical RBE model based on DSB induction is developed for use in proton therapy and its possible uses in radiobiologically constrained inverse plan optimisation are explored. The model evaluates DNA damage clustering through a probabilistic Poisson-arrival process of ionisation events, based on the local energy spectrum of particles traversing a voxel. Parametric fitting is performed to DSB counts simulated using the MCDS microdosimetric Monte Carlo software [159, 160].

2.1.2 Monte Carlo simulations

Implementation of the RBE model proposed in Chapter 3 is described in Chapter 4, which is tested for proton pencil beams of varying energies and a reference 6 MV photon beam using Monte Carlo methods. Simulated depth-RBE profiles are assessed against dose and LET for both monoenergetic (pristine Bragg peak) and modulated (SOBP) beams in a water phantom using TOPAS [161], which wraps and extends the Geant4 [162] Monte Carlo toolkit.

2.1.3 Implementation in a treatment planning system

Following Monte Carlo simulations of proton beams in water, Chapter 4 goes on to calculate RBE in the voxels of an archived head-and-neck cancer treatment plan using an empirical relationship from the generated Monte Carlo data. An open source treatment planning suite, *matRad* (DKFZ, Heidelberg, 2020) [163], containing an LET-calculation subroutine, is modified in order to obtain approximate RBE values in each voxel. The visualised RBE map is assessed both qualitatively and quantitatively to determine the variability of predicted RBE throughout the patient. This is compared to the standard clinical recommendation of assuming a constant RBE of 1.1 in all voxels.

2.1.4 Comparisons with in-vitro data

Chapter 4 concludes with an analysis of the proposed RBE model against empirical cell survival data collected *in vitro*, using the Particle Irradiation Data Ensemble (PIDE) (GSI Helmholtzzentrum für Schwerionenforschung GmbH, 2020) [164]. The DSB induction model presented in this thesis is compared to a range of phenomenological RBE models based on cell survival as well as alternate biophysical models. The model is re-imagined for a number of practical implementations suitable for both

laboratory experiments and computerised treatment planning.

2.1.5 Beam line experiment at the Clatterbridge Cancer Centre

Chapter 5 presents a clinical proton beam line experiment undertaken at the Clatterbridge Cancer Centre (CCC) in Wirral, United Kingdom, home to the world's first hospital-based proton therapy facility. Following simulations of idealised beam dynamics and measurements in the previous Chapter, real-world measurements of the radiation field using a MiniPIX-Timepix silicon semiconductor detector (Advacam s.r.o., Czech Republic, 2020) were compared to detailed simulations of the CCC beam line developed by J. Yap at the University of Liverpool [165]. Quantities such as LET, expected DSB yield and RBE were calculated from raw and simulated data and analysed in context of experimental findings from previous radiobiological studies in the literature.

2.2 A flexible projection-based inverse planning algorithm for radiation therapy

2.2.1 Theoretical development

The aforementioned methods address concerns about differences in biological and therefore clinical effects between doses delivered by protons or by conventional megavoltage photons. These aim to build confidence in the representation of dose in the patient using standard protocols that have developed in clinics over many decades. However, we also discussed in Chapter 1 (Section 1.5) the sensitivity of different tissues to heterogeneity in the spatial dose delivered. Precise shaping of the dose distribution is necessary in order to reap the full benefits of conformal IMRT, re-

ardless of the radiation modality used. Chapter 6 presents a novel inverse planning algorithm that shapes dose with dose-volume effects explicitly included in the problem formulation. Any number of DVCs may be modelled exactly, with no convex approximations necessary. A split-feasibility approach is taken as opposed to conventional cost-minimisation, by which a feasible dose solution is found through iteratively projecting the solution vector onto sets representing dose constraints.

2.2.2 Practical implementation

The inverse planning algorithm introduced in Chapter 6 is revisited in Chapter 7 with a practical and pragmatic approach. Having introduced the necessary mathematical theory, this Chapter focuses instead on the implementation, including examples of how one may exploit the flexibility of the algorithm in choosing the weighting and order in which projections are executed onto constraint sets. This flexibility allows for certain constraints to be automatically prioritised depending on levels of violation in the plan. Possibilities for hardware acceleration, for example through parallelisation, of the algorithm are discussed.

2.2.3 Testing on synthetic and clinical treatment plans

Finally, Chapter 7 then tests the performance of the inverse planning algorithm on synthetically generated test cases and archived head-and-neck proton treatment plans with many avoidance structures and constraints. The algorithm is tested in *MATLAB* (The MathWorks, Inc.) and an accelerated implementation is explored using processor multithreading and just-in-time compilation techniques with the *Python* programming language (Python Software Foundation, 2020). Dose solutions are assessed for plan quality using dose map visualisation, dose-volume histograms, Wilcoxon signed-rank tests and constraint violation counts. Overall suitability of the algorithm and possibilities for its deployment in treatment planning systems are discussed.

Chapter 3

A DNA double strand break induction model for RBE prediction

3.1 Introduction

In Chapter 1 we discussed a range of radiobiological considerations in proton therapy and, in particular, the consensus that protons and other potentially high-LET radiation modalities are more effective than x-rays at cell kill per unit of dose deposited [111]. We introduced the concept of relative biological effectiveness (RBE), which quantifies the efficacy of protons relative to a reference x-ray treatment for a given biological endpoint, and we outlined the need for widespread implementation of variable RBE models in opposition to a constant scaling factor [see, e.g., 114, 115, 166, 116, 117, 119, 118, 167]. Such adoption of variable RBE models in treatment planning systems would enable optimised dose delivery based on both physical and biological objectives, with potential to increase tumour control and reduce toxicities to healthy organs.

In this chapter we present a simple and pragmatic biophysical DNA double strand break (DSB) induction model for predicting RBE in proton beams and explore its possible uses in radiobiologically constrained inverse plan optimisation. In this formalism, DSB induction is modelled as a Poisson arrival process with respect to the ionisation cross-section (equivalently, the mean free path) of the incident particle and

the length of interaction with the DNA. The RBE, restricted to complex clustered damages to the DNA in the form of induced DSBs, is defined locally by the ratio of DSB yield at the same dose in the proton beam and the reference radiation. That is,

$$\text{RBE}_{\text{cd}} := \frac{\text{Proton DSB yield (cell}^{-1} \text{ Gy}^{-1})}{\text{Reference DSB yield (cell}^{-1} \text{ Gy}^{-1})} \Big|_{\text{isodose}}. \quad (3.1)$$

The subscript “cd” represents *clustered damage*. The model presented in this chapter scales linearly with dose, which enables RBE-weighted dose pencil beam kernels to be calculated only once, prior to optimisation of the pencil beam intensities.

We begin with a brief commentary on the choice of RBE endpoint and make the case that DSB induction serves as a practical primary endpoint for biophysical models, upon which additional physicochemical and biological considerations can be added. This is followed by a review of some of the most commonly used empirical and mechanistic RBE models in research and practice (Section 3.2), before delving into the details of the proposed model (Section 3.3). Techniques for validating DSB-based RBE models, and their limitations, are discussed in detail in Section 3.4.

3.1.1 Importance of the RBE endpoint

The RBE is heavily dependent on the biological endpoint chosen for comparison between the proton and x-ray modalities. On the surface the most clinically relevant endpoints may appear to be those derived from patient data statistics, such as long-term survival, lung fibrosis, or early reactions including, for example, skin erythema (redness, similar to a sunburn), pharyngeal inflammation (sore throat or difficulty swallowing) or xerostomia (dry mouth). These are all tangible endpoints, meaning they come directly from clinical observations rather than from extrapolated experimental findings, however in practice they do not provide robust RBE measures for three main reasons. Firstly, complication probabilities for normal tissues are highly

non-linear with dose and depend not only on the integral dose to the organ but also on its spatial distribution [168, 151]. Interpatient variability therefore renders it difficult to generalise clinical outcomes. Secondly, outcomes directly attributed to irradiation may be conflated by external factors. This is particularly problematic for adjuvant therapies such as chemotherapy or surgery. Comprehensive and standardised reporting of patient data helps to identify and remove possible sources of conflation, though the decimated data set is often not large enough to determine an RBE value with high confidence. Finally, diagnosis and assessment of some clinical outcomes are inherently subjective (e.g. loss of appetite, nausea or fatigue), which is another source of large variability in the reporting of these outcomes. Although, it should be noted that strides made in the areas of structural and functional imaging, for example, have allowed for metrics that better quantify some endpoints, such as radiation-induced necrosis [169].

The vast majority of empirical RBE data in the literature to-date has come from *in vitro* experiments using reproductive cell survival as the endpoint, though alternate endpoints such as gene expression patterns specific to radiation type have also been explored [see, e.g., 170, 171]. *In vivo* experiments are able to offer more insight into the biological response in humans, in which the tissue microenvironment plays an important role that is unable to be emulated in cell cultures. It is also typical of *in vitro* studies to use Chinese hamster ovary (CHO) or fibroblast (V79) cell lines, which have lower α/β values (see Chapter 1, Section 1.3.4) than early reacting tissues studied *in vivo* [172]. This being said, measuring RBE *in vivo* poses practical and ethical challenges of its own; Paganetti and colleagues have reported that measuring just a 5% RBE difference can require several hundred animals [113, 172].

Despite reasonably large databases of *in vitro* results now being readily available (for example, the *Particle Irradiation Data Ensemble*¹ (PIDE) from GSI [164]), it is

¹Available at: <http://gsi.de/bio-pide>. Free registration required for access.

still claimed that there is insufficient biological data, or too much variability in the data, to justify clinical use of variable RBE models for proton therapy [173]. However, heavier ion therapies, such as carbon ion therapy, currently incorporate variable RBE into the treatment planning process, with different models being used across Asia and Europe.

The major physical mechanisms by which protons and heavy ions are able to induce higher levels of cellular injury compared to photons is well-understood by examining track structure [see, e.g., 174]. Further, there is sufficient evidence of the association between high-LET protons and enhanced biological end-of-range effects [175, 176, 177, 178, 179] to challenge the recommendation of using a constant RBE of 1.1 clinically [113, 112]. Pragmatic approaches have been suggested to avoid elevated LET in normal tissues, including inverse optimisation with constrained LET-weighted dose [180, 181, 182]. In practice, even without such algorithms, clinicians often construct geometrical margins to avoid placing the Bragg peak too close to critical structures.

The major mechanism for cell kill is believed to be the induction of DSBs in nuclear DNA [66, 67]. At the Bragg peak, ionisation events are densely clustered and therefore able to cause more complex DNA damage which enzymatic DNA damage response (DDR) mechanisms cannot repair as easily as lesions caused by sparse ionisation patterns [167]. If DSBs induced by radiation do persist through repair attempts then cell death occurs most often through mitotic catastrophe [183, 184]. Biophysical models (discussed in the next section) have been derived to connect the concepts of ionisation density, DSB formation and subsequent cell death, however they vary in their predictions and in how efficiently and effectively they can be implemented in treatment planning systems.

In much the same way that in a first approximation LET can be used in the planning process to mitigate biological uncertainties, the ratio of initial DSBs based

on local energy spectra offers a pragmatic quantification of first order biological effects. It also invites more complex modeling through modularity. Physicochemical and biological process that follow the initial DSB pattern are able to be modelled independently, if one wishes to include this detail. It should be noted however that the accuracy of these models can easily be overshadowed by uncertainties introduced by factors such as the spatial dose distribution, patient positioning and movement during treatment, and inter-patient tissue response. One must also consider how effectively these models can be implemented algorithmically into inverse plan optimisation. A highly sensitive and highly non-linear objective function may be impractical where a simpler linear approximation could be used in its place.

3.2 Overview of existing RBE models

There have been a number of proton RBE models proposed over the past 25 years, the majority of which are phenomenological; they set out to find empirical relationships between LET and tissue specific properties derived from the LQ model (Chapter 1, Section 1.3.3) [185, 186, 187, 188, 189, 190, 191]. Others have taken a more mechanistic approach, developing biophysical first-principles (and semi-empirical) models which more closely reflect the biological formation and processing of DNA damage, and subsequent likelihood of cell death [192, 193, 194, 195, 198, 199, 200, 201, 202, 203]. Here we examine a selection of the most widely used models in research and treatment and compare them, where possible, under a common framework. This is not intended as an exhaustive review and so for finer details on the models themselves and their implementation, the reader is referred to the original papers cited throughout this section and the comprehensive reviews by Rørvik *et al.* [204] (phenomenological) and Stewart *et al.* [205] (mechanistic).

Considering the endpoint of reproductive cell survival, and following Dale and

Table 3.1: Summary of phenomenological and mechanistic proton RBE models described in Section 3.2.

Model	Category	RBE _{max}	RBE _{min}	Fitting data	Limitations
Wilkens & Oelfke (2004) [185]	Phen.	$\frac{1}{\alpha_x} (0.1 + 0.02 \text{ LET}_d)$	1	V79	Only one cell type used for fitting. Only 19 data points used. No $(\alpha/\beta)_x$ dependence.
Carabe <i>et al.</i> (2007) [186, 187]	Phen.	$0.843 + 0.154 \frac{2.686}{(\alpha/\beta)_x} \text{ LET}_d$	$1.09 + 0.006 \frac{2.686}{(\alpha/\beta)_x} \text{ LET}_d$	V79	Only one cell type used for fitting. Differences in reference radiation used.
Chen & Ahmad (2012) [188]	Phen.	$0.1 \frac{1 - \exp(-0.0013 \text{ LET}_d^2)}{\alpha_x} + \frac{0.045 \alpha_x \text{ LET}_d}{\alpha_x}$	1	V79	Only one cell type used for fitting. Only 14 data points used. No $(\alpha/\beta)_x$ dependence.
Wedenberg <i>et al.</i> (2013) [189]	Phen.	$1 + \frac{0.434 \text{ LET}_d}{(\alpha/\beta)_x}$	1	V79-379A, V79-753B, DLD1, SQ20B, C1-1, C3H10T1/2, SCC25, HCT116, M/10, HF19	Only 24 data points used in total across 10 cell lines. No data with LET < 6 keV/μm.
McNamara <i>et al.</i> (2015) [190]	Phen.	$0.999 + \frac{0.356 \text{ LET}_d}{(\alpha/\beta)_x}$	$1.101 - 0.004 \sqrt{(\alpha/\beta)_x} \text{ LET}_d$	Multiple cell lines from 76 studies, listed in [115].	RBE _{min} < 1.0 can be predicted.
Mairani <i>et al.</i> (2017) [191]	Phen.	$1 + \frac{0.377 \text{ LET}_d}{(\alpha/\beta)_x}$	$\exp \left[-\frac{1}{2} \left(\frac{\text{LET}_d - 3.28}{27.90} \right)^2 \right]$	V79, C3H10T1/2, M/10	RBE _{min} = 1 can be assumed. Large or undefined $(\alpha/\beta)_x$ for C3H10T1/2, M/10 cell lines.
MKM (Hawkins <i>et al.</i> (1994)) [192, 193, 194, 195, 196, 197]	Mech.†	$\frac{\alpha_0}{\alpha_x} + \frac{1}{(\alpha/\beta)_x} \frac{1}{\rho \pi r_d^2} \bar{y}_D$	1	Dependent on study	α_0 varies between experiments. DSB formation not explicitly modelled.
LEM (Scholz <i>et al.</i> (1996)) [198, 199, 200]	Mech.	No explicit expression.		Dependent on study	Large computational burden. Limited accuracy in earlier versions (before LEM IV).
RMF (Carlson <i>et al.</i> (2008)) [201, 202, 203]	Mech.‡	$\text{RBE}_{\text{DSB}} \left(1 + \frac{2\bar{z}_F \text{RBE}_{\text{DSB}}}{(\alpha/\beta)_x} \right)$	RBE _{DSB}	N/A	Continuous energy loss and straight path assumptions break down at proton end-of-range.

† ρ is the mass density of the spherical domain of radius r_d . \bar{y}_D is the dose-averaged lineal energy in that domain.

‡ \bar{z}_F is the frequency-mean specific energy per radiation event in a cell nucleus.

Jones [206], one may rewrite (1.28), presented in Chapter 1, as

$$\text{RBE} = \frac{1}{2\mathcal{D}} \left[- \left(\frac{\alpha}{\beta} \right)_x + \sqrt{\left(\frac{\alpha}{\beta} \right)_x^2 + 4\mathcal{D} \left(\frac{\alpha}{\beta} \right)_x \text{RBE}_{\max} + 4\mathcal{D}^2 \text{RBE}_{\min}^2} \right] \quad (3.2)$$

where \mathcal{D} is the dose delivered by the proton modality, $(\alpha/\beta)_x$ is the photon reference (x-ray or gamma-ray) tissue radiosensitivity parameter and

$$\text{RBE}_{\max} := \lim_{\mathcal{D} \rightarrow 0^+} \text{RBE} = \frac{\alpha}{\alpha_x}, \quad \text{RBE}_{\min} := \lim_{\mathcal{D} \rightarrow \infty} \text{RBE} = \sqrt{\frac{\beta}{\beta_x}}. \quad (3.3)$$

Through distilling down the calculation of cell survival RBE to finding expressions for RBE_{\max} and RBE_{\min} , a range of models can be compared under a common framework. Phenomenological models have been typically derived through a two-parameter linear regression to *in vitro* data, over one or multiple cell lines, to a function that is first order in $\text{LET}_d/(\alpha/\beta)_x$. Most models provide explicit expressions for the ratio of the linear dose response between protons and x-rays (or gamma-rays), α/α_x , and assume a negligible difference in the quadratic response. That is, $\beta = \beta_x$. The main features of six phenomenological and three mechanistic models are outlined below and summarised in Table 3.1. Each model is presented with optimal fitting parameters, if reported, that have been extracted from the literature.

3.2.1 Wilkens and Oelfke (2004) phenomenological model

Wilkens and Oelfke [185] proposed a simple linear relationship between the linear proton dose response, α , and dose-averaged unrestricted LET_d for the purposes of

fast RBE estimation in the iterative inverse planning process;²

$$\text{RBE}_{\max} = \frac{1}{\alpha_x}(0.1 + 0.02 \text{ LET}_d), \quad (3.4)$$

$$\text{RBE}_{\min} = 1. \quad (3.5)$$

The model was intended for use in conjunction with an analytical LET approximation [207], and was applied to treatment planning in a subsequent paper [208]. Only V79 cell lines were used to determine the two best-fit parameters in (3.4), from 19 data points [204] with $\text{LET}_d > 3 \text{ keV}/\mu\text{m}$. The response of α to increasing LET was found to deviated from a linear model beyond LET_d values of $30 \text{ keV}/\mu\text{m}$, but the linear model was kept for $3 \text{ keV}/\mu\text{m} < \text{LET}_d < 30 \text{ keV}/\mu\text{m}$ as larger values are rarely observed in proton therapy. The reference radiation used to calculate the RBE was ^{60}Co γ -rays and the V79 radiosensitivity parameter was calculated at $(\alpha/\beta)_x = 3.758 \text{ Gy}$ ($\alpha_x = 0.112 \text{ Gy}^{-1}$, $\beta_x = 0.0298 \text{ Gy}^{-2}$). As only one cell line was used, the empirical formula for α was explicitly determined and it does not generalise to other cell lines. Furthermore, mixed radiation fields due to secondary particles produced in non-elastic nuclear interactions are not accounted for. However, in a first approximation the model was useful as a heuristic approach to handling end-of-range effects due to its simplicity.

3.2.2 Carabe *et al.* (2007) phenomenological model

In contrast to the model proposed by Wilkens and Oelfke [185], the Carabe *et al.* model [186, 187] accounts for tissue-specific reference radiosensitivity, $(\alpha/\beta)_x$ in addition to LET_d in both the linear and quadratic terms, corresponding to RBE_{\max} and

²In all empirical expressions for RBE_{\max} and RBE_{\min} presented in this work, the dimensions of constants can be determined according to $[\alpha] = \text{Gy}^{-1}$, $[\beta] = \text{Gy}^{-2}$ and $[\text{LET}] = \text{keV}/\mu\text{m}$.

RBE_{min} respectively. The model gives the following:

$$\text{RBE}_{\max} = 0.843 + 0.154 \frac{2.686}{(\alpha/\beta)_x} \text{LET}_d, \quad (3.6)$$

$$\text{RBE}_{\min} = 1.09 + 0.006 \frac{2.686}{(\alpha/\beta)_x} \text{LET}_d. \quad (3.7)$$

The model was fit to V79 survival data from seven separate publications (found in [187]), comprising 44 data points in total. The additional constant of 2.686 Gy in the numerators of (3.6) and (3.7) is the average value of the measured $(\alpha/\beta)_x$ values across the reference modalities, though the spread in these measurements could not be found in the literature. In taking the average, one must be wary of differences in the tissue radiosensitivity to different reference radiations and also in other experimental conditions which may have an impact on the results. For example, one of the data sets used by Carabe *et al.* used 240 kVp x-rays [209] while another used ⁶⁰Co γ -rays [210].

3.2.3 Chen and Ahmad (2011) phenomenological model

In 2011 Chen and Ahmad [188] introduced a nonlinear model to describe RBE_{max} while assuming no difference in the quadratic dose response;

$$\text{RBE}_{\max} = \frac{0.1}{\alpha_x} + \frac{1 - \exp(-0.0013 \text{LET}_d^2)}{0.045\alpha_x \text{LET}_d}, \quad (3.8)$$

$$\text{RBE}_{\min} = 1. \quad (3.9)$$

Unlike most of the other models, there is no dependence on $(\alpha/\beta)_x$ factored into the model. V79 survival data from four separate studies and LET values collected from Monte Carlo simulations in Geant4 [162] were used for the fitting procedure, comprising a total of 14 data points. The reference radiation is reported in [188] to be ⁶⁰Co, with an endpoint of 10% surviving fraction, giving $(\alpha/\beta)_x = 0.3$ Gy

($\alpha_x = 0.1 \text{ Gy}^{-1}$, $\beta_x = 0.3 \text{ Gy}^{-2}$.) As was the case for the previously mentioned models, secondary particles heavier than protons, such as ^4He , are not included in the model. However, Chen and Ahmad make a note of the fact that these particles contribute less than 1% of the total dose. The majority of that dose is also in the entrance region, where nuclear fragmentation is most likely to occur, away from the Bragg peak.

3.2.4 Wedenberg *et al.* (2013) phenomenological model

The Wedenberg *et al.* model [189] is a simplified variation of the model by Carabe *et al.* [186, 187], with the quadratic dose response, β , assumed to be the same as the photon response, β_x , and only one fitting parameter needed for α ;

$$\text{RBE}_{\max} = 1 + \frac{0.434 \text{ LET}_d}{(\alpha/\beta)_x}, \quad (3.10)$$

$$\text{RBE}_{\min} = 1. \quad (3.11)$$

The linear relationship between RBE_{\max} and LET_d is again only considered valid for $\text{LET}_d < 30 \text{ keV}/\mu\text{m}$. Unlike the models already presented, ten different cell lines (V79-379A, V79-753B, DLD1, SQ20B, C1-1, C3H10T1/2, SCC25, HCT116, M/10, HF19) were used in the fitting procedure, however the minimum LET value in the data set was $6 \text{ keV}/\mu\text{m}$, which is large enough to cause a clinically significant radiobiological effect [211]. Ideally, lower values of LET should included in the fitting procedure. Experimental data was sourced from six studies, comprising 24 data points in total. Efforts were made to select studies in which the reference $(\alpha/\beta)_x$ values were consistent for each cell line. Overall, the fitting procedure and statistical analysis by Wedenberg *et al.* is presented with more rigour than in the previously discussed models thus far.

3.2.5 McNamara *et al.* (2015) phenomenological model

The McNamara *et al.* model, published in 2015, built on earlier works, especially those by Wedenberg *et al.* [189] and Jones [212]. Tissue radiosensitivity is incorporated into the quadratic response term, β , while the linear response term, α , is described by the same relationship as (3.10);

$$\text{RBE}_{\max} = 0.999 + \frac{0.356 \text{ LET}_d}{(\alpha/\beta)_x}, \quad (3.12)$$

$$\text{RBE}_{\min} = 1.101 - 0.004\sqrt{(\alpha/\beta)_x} \text{ LET}_d. \quad (3.13)$$

It should be noted that the constant of 0.999 in (3.12) can sometimes be misreported as 0.991 due to an error in the original paper by McNamara *et al.* [190]; it is assumed the true value should be taken from eq(7) there and not from the preceding text. The main advantage of this model, compared to other phenomenological models, is that a large and diverse database of experimental results was used in the fitting procedure. A range of cell lines, and therefore $(\alpha/\beta)_x$ values, were used, albeit this also presented a source of large uncertainty in the regression due to the spread of the data. Data points were given a statistical weight in the fitting procedure according to uncertainties reported by Paganetti [115]. A total of 285 data points from a subset of the 76 different studies documented by Paganetti [115] were used, with the restrictions that $\text{LET}_d < 20 \text{ keV}/\mu\text{m}$ and $(\alpha/\beta)_x < 30 \text{ Gy}$.

3.2.6 Mairani *et al.* (2017) phenomenological model

In 2017 Mairani *et al.* [191] proposed a model that explores higher-order terms in LET_d and can be applied to secondary fragments such as ^4He . It followed previous work specifically for helium ion beams [213, 214]. The prediction of RBE_{\max} is of the same form as (3.10) and (3.12) while that of RBE_{\min} is given by an exponential

relationship;

$$\text{RBE}_{\max} = 1 + \frac{0.377 \text{ LET}_d}{(\alpha/\beta)_x}, \quad (3.14)$$

$$\text{RBE}_{\min} = \exp \left[-\frac{1}{2} \left(\frac{\text{LET}_d - 3.28}{27.90} \right)^2 \right]. \quad (3.15)$$

The fitting procedure used the same data as the Wedenberg *et al.* model [189] with an additional seven data points from studies on V79-753B (200 kVp x-ray reference), C3H10T1/2 and M/10 (^{60}Co γ -ray reference) cell lines. A total of 31 data points were used with LET values up to 37.8 keV/ μm . Every data point had an associated uncertainty which was used in the χ^2 minimisation procedure. A range of parameterizations for RBE_{\max} were investigated, including quadratic, cubic, and exponential factors of LET_d , however the linear model shown in (3.14) provided the best fit to the data. Mairani *et al.* found that $\text{RBE}_{\min} = 1$ is a valid assumption due to the exponential term in (3.15) providing a negligible deviation from unity over the LET range [191].

3.2.7 Microdosimetric-kinetic model

The microdosimetric-kinetic model (MKM, or “MK model”) [192, 193, 194, 195, 196, 197] builds upon the theory of dual radiation action (TDRA) [215, 216] to relate the radiation field in small sensitive subvolumes (or “domains”) of the cell nucleus to the fate of potentially lethal lesions (PLLs) in DNA. These PLLs may be repaired or converted into lethal chromosomal aberrations. For example, an acentric fragment could form as the result of a single PLL, or two PLLs in close proximity may combine to form a dicentric chromosome. It is assumed in the TDRA that the formation of PLLs (or “sublesions”) inside spherical domains is proportional to the energy imparted to the site, which can be quantified with the probability distribution of the specific energy, z . In the MK model the radiation field within a domain is characterised by

the lineal energy, $y := \epsilon/\bar{\ell}$, where ϵ is the energy deposited and $\bar{\ell}$ is the mean chord length of the radiation track through that volume. The quadratic component of the cell survival curve is assumed to be no different to that for photon radiation ($\beta = \beta_x$) while the linear response, α , is assumed to increase with the dose-averaged lineal energy \bar{y}_D , which can be approximated by LET_d . The model assumes

$$\text{RBE}_{\max} = \frac{\alpha_0}{\alpha_x} + \frac{1}{(\alpha/\beta)_x} \frac{1}{\rho\pi r_d^2} \bar{y}_D, \quad (3.16)$$

$$\text{RBE}_{\min} = 1, \quad (3.17)$$

where ρ is the density of the spherical domain, r_d is its radius and α_0 is a constant representing the value of α in the limit that the lineal energy approaches zero. The formulation of the MK model in terms of microdosimetry and the formation of PLLs differs to the explicit DSB induction modelling of the two models that immediately follow, the local effect model and the repair-misrepair-fixation model. However, PLLs manifest primarily as DSBs [195, 205].

3.2.8 Local effect model

The local effect model (LEM) [198], specifically the latest version, LEM IV [199, 200], relates radiation-induced DSBs to reproductive cell survival by predicting DSB clustering in chromatin loops. DSBs are classified as “isolated” (iDSB), which are considered simpler for cellular mechanisms to repair, or “clustered” (cDSB), referring to multiple DSBs in close proximity within a chromatin loop, which are less likely to be successfully repaired. The clustering of DSBs is quantified through a cluster index,

$$C := \frac{N_{\text{cDSB}}}{N_{\text{iDSB}} + N_{\text{cDSB}}} \quad (3.18)$$

where N_{iDSB} and N_{cDSB} represent the number of loops with isolated or clustered DSBs, respectively. The model, which applies not only to protons but to any radiation type,

originally (LEM I) assumed that equal doses in a target region lead to equal local effects, thus the quality of the radiation does not matter. LEM IV revised this concept to assume that equal local DSB densities, instead of physical dose, lead to equal biological effects. The RBE is governed by the microscopic distribution of energy deposition and therefore the track structure of the radiation. A homogeneous distribution of DNA is assumed within the cell nucleus and the cell volume is broken down into a large number of disjoint subvolumes. Assuming a homogeneous dose \mathcal{D} throughout the cell, the probability of zero lethal events, $\mathcal{P}(0)$, and therefore the surviving fraction, S , may be calculated using Poisson statistics,

$$S(\mathcal{D}) = \mathcal{P}(0) = \exp(-\langle N(\mathcal{D}) \rangle) \quad (3.19)$$

where $\langle N(\mathcal{D}) \rangle$ is the expected number of lethal events, or DSBs, within the cell. In the limit that the number of subvolumes becomes infinitely large and each one has a volume of dV , the local dose, $d_V(x, y, z)$, deposited (homogeneously) throughout the subvolume can be used to determine the expected number of lethal events over the whole cell. Given that $\mathcal{D} = \int d_V(x, y, z)dV$, one may write

$$\langle N(\mathcal{D}) \rangle = \int \frac{-\log S(d_V(x, y, z))}{V} dV. \quad (3.20)$$

In practice the volume of the cell is discretised and the local dose in each voxel can be computed using amorphous track structure simulations, although this is only really necessary for particle beams and not for the photon reference radiation, which can be assumed homogeneous. One computation of the local dose does not suffice, as it is a stochastic quantity. Instead, it, or equivalently the surviving fraction calculated in the subvolume, should be averaged over a large number of independent simulations (or cells) [217]. The dose-response curves, given by (3.19), for protons and for the reference modality can then be used to calculate the RBE. The full method behind

RBE calculation using the LEM IV is beyond the scope of this thesis, but is described well by, for example, Karger and Peschke [218] and Stewart *et al.* [205]. While we cannot explicitly write RBE_{\max} or RBE_{\min} for the LEM, it is of importance to note that up to eleven default parameters and five input parameters can be required to implement the model [205], which relate to track structure (if using analytic methods), cell geometry, typical photon DSB and single strand break (SSB) yields, DNA content of cell nuclei, and reference radiation response. In its full implementation in treatment planning, it can present a large computational burden. A first order approximation of the dose response from only intra-track damages (relating to the α term) may be implemented rapidly, though methods have been developed to approximate the combination of damage sites by multiple particle tracks (relating to the β term) [see, e.g., 219]. Recent benchmarking to survival data for HSG, V79 and CHO-K1 cell lines with various ions (protons, helium, carbon, oxygen and neon) showed limited accuracy and variability between the first three iterations of the LEM model, while LEM IV agreed well with the data [220]. However, the LEM IV is yet to be widely implemented in clinical treatment planning systems [218].

3.2.9 Repair-misrepair-fixation model

The repair-misrepair-fixation (RMF) model [201, 202, 203] explicitly links cell survival RBE to the induction of initial DSBs and subsequent lethal chromosome aberrations. The number of induced DSBs per gray of dose deposited per 10^9 basepairs of DNA is denoted Σ ($\text{Gy}^{-1} \text{Gbp}^{-1}$) and the frequency-weighted mean specific energy per radiation event in a cell nucleus $\langle z \rangle_F$ (Gy). In the limit that dose is small in comparison to the radiosensitivity parameter, α/β (in whichever modality), the linear and quadratic

dose responses are parameterised by tissue-specific constants θ and κ by

$$\alpha = \theta\Sigma + \kappa\langle z \rangle_F \Sigma^2, \quad (3.21)$$

$$\beta = \frac{\kappa}{2}\Sigma^2. \quad (3.22)$$

The linear term in (3.21) represents unrepaired or misrepaired DSBs while the quadratic term represents intra-track chromosomal aberrations. The expression in (3.22) represents inter-track aberrations; that is, those from multiple particle tracks on the same target. In the framework we have been using thus far,

$$\text{RBE}_{\max} = \text{RBE}_{\text{DSB}} \left(1 + \frac{2\langle z \rangle_F \text{RBE}_{\text{DSB}}}{(\alpha/\beta)_x} \right), \quad (3.23)$$

$$\text{RBE}_{\min} = \text{RBE}_{\text{DSB}} \quad (3.24)$$

where $\text{RBE}_{\text{DSB}} := \Sigma/\Sigma_x$, the ratio of initial DSBs from a proton beam to those from a ^{60}Co γ -ray reference beam. In a first approximation, under the assumptions that protons are not deviated from their path by multiple scattering, energy losses can be described adequately by the continuous slowing down approximation (CSDA), the target is spherical with diameter d and composed of water, the mean specific energy may be written as [201, 221, 222]

$$\langle z \rangle_F \approx \frac{0.204 \text{ LET}_d}{d^2}. \quad (3.25)$$

Further, $d \approx 4 \mu\text{m}$ usually suffices for cell survival data relating to human and mammalian cells [205]. Therefore, a simpler variation of (3.23) may be approximated as

$$\text{RBE}_{\max} \approx \text{RBE}_{\text{DSB}} \left(1 + \frac{0.0256 \text{ LET}_d}{(\alpha/\beta)_x} \text{RBE}_{\text{DSB}} \right). \quad (3.26)$$

At the proton end-of-range the assumptions of continuous energy loss and straight paths can break down, and energy and path length straggling may decrease the localised energy deposition [223]. The mechanistic formulation of the RMF model allows it to account for changing radiation quality, energy and biological target (e.g. cell type). An auxiliary model must be used to determine the DSB-induction RBE, RBE_{DSB} . The Monte Carlo Damage Simulation (MCDS) software by Stewart *et al.* [159, 160] has commonly been used to estimate the number of DSBs for a continuous range of particle energies in the therapeutic range (and beyond). MCDS is also able to simulate the interplay between particle LET and the effects of oxygen concentration in the cell [159, 160, 224, 225, 226].

3.2.10 Van den Heuvel (2014) DNA damage model

Van den Heuvel [227] published a geometrical approach to DNA damage by charged particles using a single-particle model. The DNA is approximated as a cylinder with appropriate dimensions and the energy-dependent mean free path $\lambda(E)$ between successive ionisations is used to infer the probability of DSB generation. A larger angle of incidence implies a longer path length through the DNA and therefore higher LET. This was linked to a greater likelihood of clustered damage. The problem was shown to reduce to a Cauchy distribution in the high-LET regimen, with the distribution of DSBs reformulated as a function of energy, explicitly. The derived complex damage response function, F_{cd} , was stated as

$$F_{\text{cd}}(E) = (a - b) \frac{2}{\pi} \left[\tan^{-1} \left(\frac{E - E_0}{\Gamma/2} \right) \right] + b, \quad (3.27)$$

which gives the yield of DSBs per 10^9 basepairs of DNA per 1 gray of absorbed dose. The parameters a , b , Γ , and E_0 were fitted to DSB counts from particle tracks simulated in MCDS [228] through a two-stage χ^2 minimisation process, using the

differential distribution, dF_{cd}/dE , and the cumulative distribution, F_{cd} . Parameters were fit to simulated data for electrons, protons, helium ions and carbon ions.

Van den Heuvel argued that for implementation in a treatment planning system, if the energy spectrum, $\Phi(E)$, for each type of particle is known in a voxel then the mean number of DSBs, M_{cd} , can be calculated using the response function ((3.27)) and the dose in that voxel, D [227, 229, 230];

$$M_{cd} = D \times \frac{\int_0^\infty \Phi(E)F_{cd}(E)dE}{\int_0^\infty \Phi(E)dE}. \quad (3.28)$$

The RBE by DSB-induction could then be calculated by dividing the value of M_{cd} for the proton treatment by that for the reference treatment.

3.3 A new DNA double strand break induction model

Here a method is proposed for quantifying clustered DNA damage induced by proton irradiation. The energy-dependent mean free path $\lambda(E)$ of a particle infers that in a track segment of length L the mean number of interactions is given by $\bar{n} = L/\lambda$. Given that a particle has travelled through nuclear DNA, using the inelastic mean free path (IMFP) we may refer to \bar{n} as the average number of lesions induced within the DNA. The distance L shall be known as the interaction length. The conditional probability, $\mathcal{P}(\lambda|L)$, that two or more lesions occur within the segment, contributing to a DSB, is given by a Poisson distribution,

$$\mathcal{P}(\lambda|L) = \sum_{k=2}^{\infty} \frac{1}{k!} \left(\frac{\varepsilon L}{\lambda}\right)^k \exp\left(-\varepsilon \frac{L}{\lambda}\right) = 1 - \left(1 + \varepsilon \frac{L}{\lambda}\right) \exp\left(-\varepsilon \frac{L}{\lambda}\right), \quad (3.29)$$

where a factor $\varepsilon > 0$ has been included to represent the efficiency with which an interaction is converted to DNA damage. This factor linearly scales the mean number

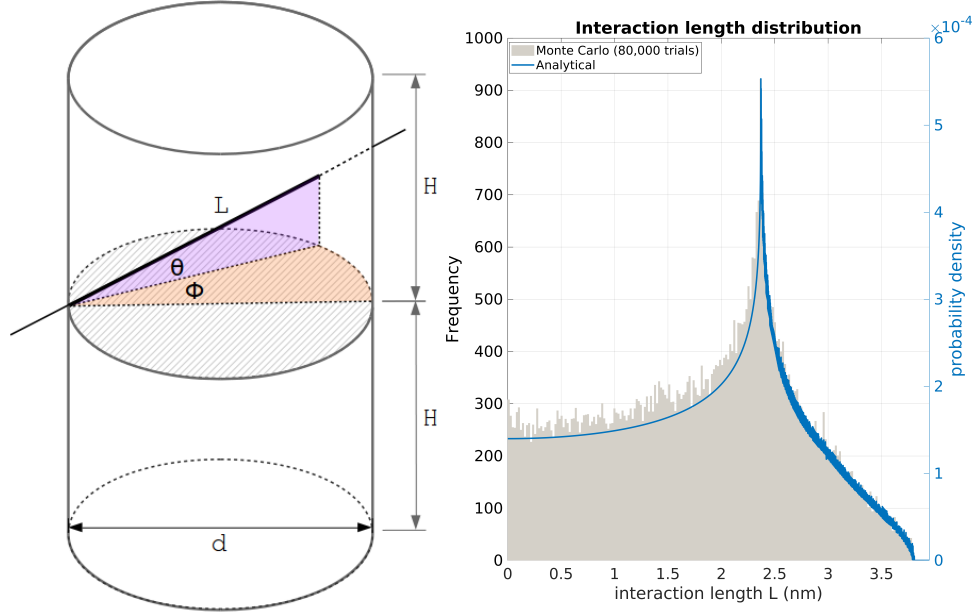


Figure 3.1: Left: geometric setup for possible interaction lengths of a linear particle trajectory through a cylindrical section representing the sensitive region containing the DNA, in which an ionisation event may cause lasting DNA damage. Right: calculated and Monte Carlo generated interaction length distribution for $d = 2.37$ nm and $2H = 6.8$ nm.

of interactions in a given segment. To classify a track as contributing toward complex damage, all lesions along the track must be restricted to the dimensions of the DNA double-helix and within ten base-pairs (bp). The choice of ten base-pairs is in keeping with previous studies detailed in the literature [74, 231, 232, 159, 160, 225, 233, 234]. Therefore, L must take values from the distribution of permitted interaction lengths.

We begin by approximating 20 bp in a physiological B-DNA double-helix as a cylinder of diameter $d = 2.37$ nm and height $2H = 6.8$ nm. Without loss of generality, the incident particle may assume entry at the midpoint of the cylinder with uniform random variables $\phi \in (-\frac{\pi}{2}, \frac{\pi}{2})$ and $\theta \in (-\frac{\pi}{2}, \frac{\pi}{2})$ defining the azimuthal and polar components, respectively, of the trajectory as shown in Figure 3.1. Monte Carlo methods may be used to generate the probability density function (PDF) of possible interaction lengths L , also shown in Figure 3.1. Further, it can be shown analytically

(derivation provided in Appendix A) that the PDF is given by

$$f_L(\ell) = \begin{cases} \frac{4}{\pi^2 d} \int_0^1 \frac{y}{\sqrt{1-y^2} \sqrt{1-(\ell y/d)^2}} dy, & \ell \leq H, & (3.30a) \\ \frac{4H}{\pi^2 \ell^2} \frac{\arcsin\left(\frac{\ell}{d} \sqrt{1-(H/\ell^2)}\right)}{\sqrt{1-(H/\ell^2)}} & \ell > H. & (3.30b) \\ \quad + \frac{4}{\pi^2} \frac{\partial}{\partial \ell} \int_{\sqrt{1-(H/\ell^2)}}^1 \frac{\arcsin\left(\frac{\ell y}{d}\right)}{\sqrt{1-y^2}} dy, & & \end{cases}$$

The PDF, f_L , may be used as weighting coefficients in a linear combination of infinitely many Poisson-distributed random variables. The resulting likelihood measure, \mathcal{L}_{cd} , is proportional to the probability of a single proton inducing two or more lesions on opposing DNA strands, and may be interpreted as a relative likelihood measure for the induction of complex DNA damage;

$$\mathcal{L}_{cd}(\lambda; \ell) := f_L(L) \mathcal{P}(\lambda|L). \quad (3.31)$$

3.3.1 Inelastic mean free path calculation

Our proposed Poisson model for damage induction is based on the inelastic mean free path (IMFP) of a particle incident on a DNA segment. The IMFP, λ , is dependent on the kinetic energy of the particle, E , and is inversely proportional to the cross-section of inelastic interaction with the target material, σ_{in} . That is,

$$\lambda(E) \sim 1/\sigma_{in}(E). \quad (3.32)$$

As the IMFP decreases, meaning a greater number ionisation events may occur within a segment, the probability of inducing a DSB increases accordingly. The accuracy of our damage estimate in (3.31) therefore heavily depends on the accuracy of the IMFP calculation.

To calculate the IMFP we follow methods employed by Van den Heuvel [227]. The traditional Bethe-Barkas formalism [52, 238] is widely used to describe the IMFP for particles of relatively large kinetic energies, above the threshold energy for producing core ionisations in the target nucleus. At low energies, however, Ziaja *et al.* [239] argue that the band structure of the target material is essential when considering the behaviour of ionisations. They incorporate knowledge of the band gap, which is the minimal energy required to generate an electron-hole pair, and the behaviour of plasmon excitations at low energies, to more accurately describe the IMFP;

$$\lambda(E) = \frac{E - E_0 \exp(-B/A)}{A \log(E/E_0) + B} + \frac{\sqrt{E}}{A_1(E - E_{th})^{B_1}} \quad (3.33)$$

where E_0 , E_{th} , A , B , A_1 , and B_1 are fitting parameters. The first term comes from the traditional Bethe-Barkas formalism and the second term is the plasmonic correction [239, eq.(7)]. Using a Levenberg-Marquadt fitting procedure, the parameters in (3.33) were fitted to IMFP data in DNA for electrons [236, 237] and protons [235], as shown in Table 3.2 and Figure 3.2. The first term, if used exclusively, erroneously describes

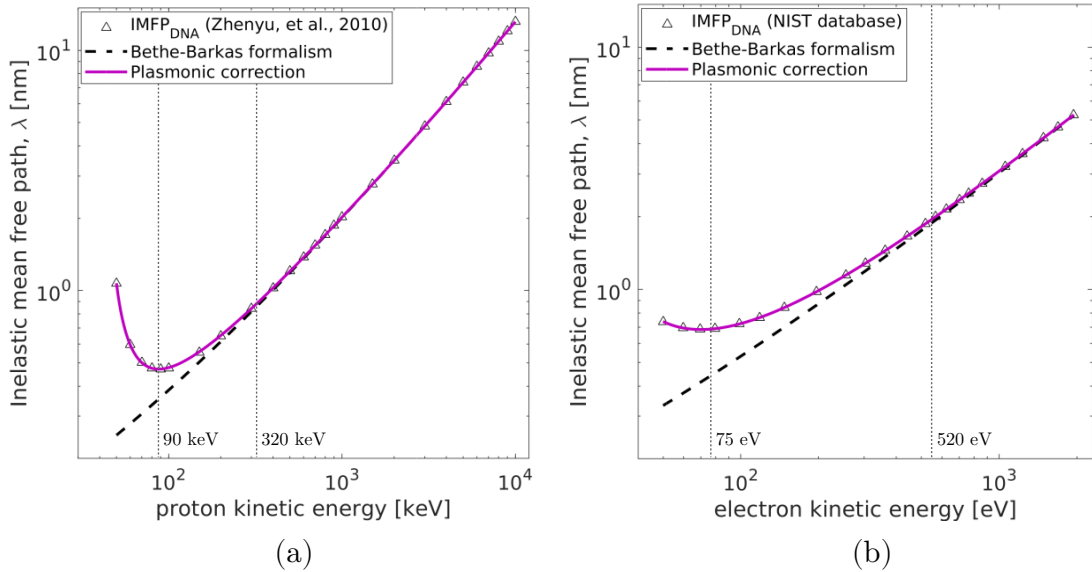


Figure 3.2: IMFP in DNA as calculated by (3.33) with parameters listed in Table 3.2 for protons and electrons, both with and without a plasmonic correction. Parameters were fitted to experimental data by Tan *et al.* [235] (protons) and NIST [236, 237] (electrons).




Figure 3.3: Proton collisional stopping power in liquid water, shown as a function of kinetic energy between 1 keV and 10 GeV. Stopping power is maximised at approximately 90 keV, which is also the energy at which the IMFP in Figure 3.2(a) is minimised. Plotted data was collected from the *NIST pSTAR* database [54].

a monotonic decrease in λ as E decreases. In DNA, we see the plasmonic correction begin to take effect at energies lower than 320 keV (protons) and 520 eV (electrons). The IMFP reaches a minimum at 90 keV (protons) and 75 eV (electrons), which indicates that these are the approximate energies at which DSB induction is most likely to occur.

Notably, proton stopping power, which describes the expected rate of energy loss of protons through a material, also peaks around 90 keV in liquid water, as shown in Figure 3.3. In liquid water we see a maximal stopping power of approximately 830 MeV cm²/g, which, in a naive approximation that accounts only for the density

Table 3.2: Parameters fitted to experimental data, by Tan *et al.* [235] (protons) and NIST [236, 237] (electrons), using a Levenberg-Marquadt procedure, for the IMFP calculation in (3.33). Parameters are provided for both the Bethe-Barkas formalism and the plasmonic correction term.

	Bethe-Barkas formalism			plasmonic correction		
	A [keV/nm]	B [keV/nm]	E_0 [keV]	A_1	B_1	E_{th} [keV]
protons	117.007	-318.694	1.00	0.681	1.249	42.396
electrons	0.0659	-0.1282	0.001	22.228	1.010	0.0243

of DNA being approximately 1.3 times that of water [240], translates to an LET of around 108 keV/ μm . This density scaling is in line with the methods of Champion and colleagues [241, 242], who applied it to cellular media comprising hydrated DNA. Thorough investigations, including experimental and simulation studies, of stopping power in DNA have yielded maximal values between 110 and 140 keV/ μm , all occurring for a proton energy of approximately 100 keV [243, 244, 245, 246, 247, 248].

It is of interest to mention that an LET of 100 keV/ μm is generally considered the optimal ionisation density for generating a DSB, with a common explanation being that the mean separation between ionisations at this LET coincides with the diameter of the DNA [see, eg, 64, p. 108-10]. We offer an alternate explanation; for protons with a kinetic energy of approximately 90 keV, the mean separation between ionisation events is minimised which, in turn, maximises the probability of DSB induction.

3.3.2 Parameterising the damage model

In a clinical proton beam, primary and secondary protons are responsible for the great majority (approximately 97%) of deposited dose. We must assume at this stage, however, that the damage contribution by secondary electrons is non-negligible. Regardless, we must model these contributions as this provides the necessary reference for calculating the RBE, as secondary electrons are responsible for the vast majority of DNA damage in conventional x-ray beam therapy. We therefore tune the efficiency parameter, ε , in (3.31) for protons and electrons separately such that our predictive model fits simulated DSBs via the Monte Carlo method.

The microdosimetric Monte Carlo damage simulation software by Semenenko and Stewart, MCDS (version 3.1) [159, 160], was used to obtain approximate DSB yields in cells irradiated by monoenergetic protons and electrons for a range of input energies. In the MCDS software, clustered elementary damages to DNA are generated through detailed track-structure simulations and categorised into levels of complexity

described by Nikjoo *et al.* in [78] and [74]. A DSB is defined as two strand breaks on opposite strands within 10 base pairs (bp) of each other. The DNA content of the cell was assumed to be 1 Gbp (giga base pair) on average, the radius of the cell nucleus was set to 5 μm and 10,000 cells were used in the damage simulation. Hypoxic environmental conditions (0 mmHg O_2 pressure) were chosen, in a first approximation, in order to limit the effects of chemical repair and oxygen fixation, which are not included in our physical damage induction model. DSB yields per cell per unit dose (Gy), induced by an individual particle track, were extracted for 294 protons with kinetic energies between 10 eV and 110 MeV, and 66 electrons with kinetic energies between 0.1 eV and 10 MeV.

As MCDS returns absolute DSB counts, we convert these to relative induction probabilities by scaling the data between 0 and 1 such that the largest count becomes 1, indicating maximal damage, and the smallest count returns 0, indicating minimal damage. In this form, the efficiency parameter, ε in (3.29) was allowed to vary in a Levenberg-Marquadt method to best fit the data to (3.31) with f_L normalised such that the most likely interaction length has a density of unity. This was performed for

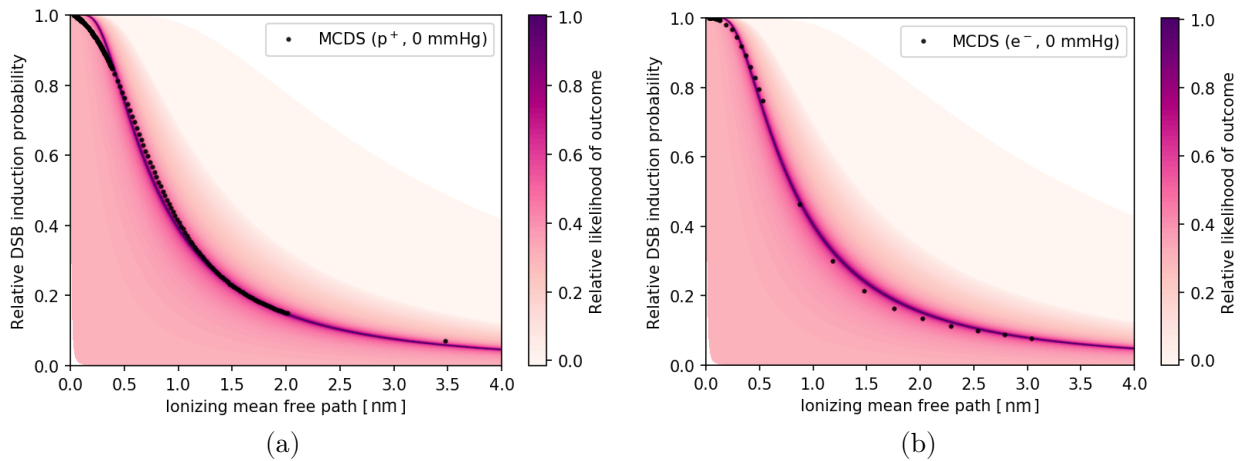


Figure 3.4: The DSB induction model, (3.31), fit through a single “conversion efficiency” parameter, ε (listed in Table 3.3), to relative Monte-Carlo-generated DSB counts using the MCDS damage simulation software [159, 160]. Fits are shown for (a) protons and (b) electrons. The colour scale corresponds to the interaction length probability density, f_L .

Table 3.3: Parameter fits to Monte Carlo generated DSBs using MCDS 3.1 [159, 160] along with their 95% confidence interval, for DSB conversion efficiency, ε , low-level (minimal) damage, a , and high-level (maximal) damage, b .

	ε	a [$\text{cell}^{-1}\text{Gy}^{-1}$]	b [$\text{cell}^{-1}\text{Gy}^{-1}$]
protons	0.5837 (0.5739, 0.5935)	2.8725 (2.6783, 3.0667)	27.9034 (27.7876, 28.0192)
electrons	0.5847 (0.5755, 0.5939)	2.8362 (2.7488, 2.9236)	28.2680 (28.1704, 28.3656)




Figure 3.5: DSB yield, as calculated by (3.34), shown with and without the plasmonic correction term in (3.33) as a function of kinetic energy for (a) protons and (b) electrons. Monte-Carlo-generated DSB counts using MCDS 3.1 [159, 160] are also shown, and were used to fit the parameters listed in Table 3.3.

both protons and electrons, separately. The resulting likelihood measures for DSB induction, shown in Figure 3.4(a) and Figure 3.4(b), are given by (3.31) with the fitted parameter, ε , provided in Table 3.3.

In addition to modelling the relative likelihood of DSB induction, we may predict the yield, Y_{cd} , of DSBs using the unscaled data generated using the MCDS software. That is,

$$Y_{\text{cd}}(\lambda; L) = (b - a)\mathcal{L}_{\text{cd}}(\lambda; L) + a \quad (3.34)$$

where a and b are parameters representing the minimum and maximum DSB counts, respectively. Fits for protons and electrons are provided in Table 3.3.

Kinetic energies were converted in the fitting procedure to the inelastic mean free path, $\lambda(E)$, using (3.33) with the parameters listed in Table 3.2. As the MCDS software (version 3.1) used to obtain the simulated DSB counts does not take into account plasmonic effects, even very low energy particles are assumed to generate DSBs with high likelihood. The model was therefore parameterised using the standard Bethe-Barkas formalism [52, 238] for calculating the IMFP. The final damage model includes the plasmonic correction. The associated DSB yields for the most likely interaction length of $\ell_0 = d$, the diameter of the DNA, are shown in Figure 3.5(a) and Figure 3.5(b).

3.3.3 Relative biological effectiveness

We may extend the damage model to calculate the relative biological effectiveness (RBE), restricted to DSB induction, between a radiation field set up by a proton beam and that set up by a reference modality. The reference in almost every case is a photon field generated for conventional x-ray therapy in which the vast majority of dose contributions are attributed to secondary electrons. In a specified volume of interest, which is most suitably set commensurate with the voxel size in the patient dose matrix, the restricted relative biological effectiveness, RBE_{cd} , is given by the ratio of the expected DSB yield in the modality of interest, $\langle Y_{\text{cd}} \rangle$, to that in the reference modality, $\langle Y_{\text{cd}} \rangle_{\text{ref}}$. Note that this requires an underlying assumption that initial DSB formation is linear with dose, which is considered valid for standard therapeutic doses [205]. That is,

$$\text{RBE}_{\text{cd}} = \frac{\langle Y_{\text{cd}} \rangle}{\langle Y_{\text{cd}} \rangle_{\text{ref}}}. \quad (3.35)$$

The (dose-averaged) expected DSB yield may be calculated using

$$\langle Y_{\text{cd}} \rangle = \frac{1}{\mathcal{D}} \sum_{\substack{\text{particle} \\ \text{type}}} \int_0^{E_{\text{max}}} \Phi(E) \frac{S_{\text{col}}}{\rho}(E) Y_{\text{cd}}(\lambda(E), \ell_0) dE \quad (3.36)$$

where $\Phi(E)$ is the particle fluence at kinetic energy E , up to a maximum of E_{\max} , through the volume of interest, ρ is the mass density of the medium, $\frac{S_{\text{col}}}{\rho}(E)$ is the unrestricted mass collision stopping power for this energy and medium, and

$$\mathcal{D} = \sum_{\substack{\text{particle} \\ \text{type}}} \int_0^{E_{\max}} \Phi(E) \frac{S_{\text{col}}}{\rho}(E) dE \quad (3.37)$$

is the dose to the medium in the volume of interest. Note that the summation in (3.36) and (3.37), for our purposes, separates the electron energy spectrum from the proton energy spectrum.

3.4 Validation of DSB-based RBE models

While it is understood that initial DSB generation in the cell nucleus, and the subsequent threat of chromosome aberrations, plays a crucial role in cell death, there are many contributing factors in determining the biological outcome of cell irradiation (IR). As discussed previously, these include, but are not limited to, environmental conditions, beam dynamics, spatio-temporal track structure, cell type, tissue vasculature, and repair kinetics. Though despite such complexity, a mixture of phenomenological and first-principles modelling is able to provide useful tools that account for variability in outcomes between different treatment modalities and methodologies. However, these biophysical models must have demonstrable predictive capabilities validated by real-world experimentation and clinical evidence, and the uncertainty in these predictions must always be recognised.

This author proposes that the accuracy of RBE models may be assessed through three distinct levels of validation, in order of decreasing confidence; (I) validation of direct processes and mechanisms modelled (e.g. number of DSBs generated), (II) endpoint validation (e.g. reproductive cell survival), and (III) comparison between

Table 3.4: Levels of confidence for validating DSB-based RBE models.

Validation level	Description	Relative confidence	Example
I	Validation of direct processes and mechanisms modelled.	●●●●●	Absolute number of DSBs measured through immunostaining, PFGE or comet assay, compared to predicted DSB yield.
II	Endpoint validation.	●●●●○	Surviving fraction ratio between proton and reference modality from clonogenic assay compared to cell survival RBE from mechanistic DSB-based model.
III	Cross-endpoint comparison.	●●○○○	<i>In vivo</i> animal survival rate compared to ratio of DSBs in proton and reference modalities.

alternate endpoints. These are summarised, with examples, in Table 3.4. In the context of modelling RBE based on initial DSB induction, predicted DSB counts can be compared directly to experimentally measured DSB counts (Level I validation), or the predicted likelihood of cell death due to initial DSBs could be compared to results from clonogenic survival assays (Level II). In order to understand the benefits and limitations of these validation methods, the most common experimental techniques used to detect DSBs and to measure cell survival post-IR are reviewed below.

3.4.1 Detection of double strand breaks

A range of methodologies exist for quantifying DSBs in prepared cell samples, though the vast majority can be segmented into either immunocytological techniques, such as immunofluorescent staining with γ H2AX [249], or physical fragmentation, such as pulsed-field gel electrophoresis (PFGE) [250] or single-cell electrophoresis, also known as the neutral comet assay [251]. The former involves incubating the cells with antibodies that encourage the recruitment of fluorescent proteins to damage sites, thus showing these sites as visual foci under immunofluorescence microscopy, while the latter involves physically separating fragments that have resulted from double

stranded cleavage of the DNA.

3.4.1.1 Immunofluorescent staining

Immunostaining, for the purposes of DSB detection, refers to techniques by which DSBs may be quantified through the intensity and location of fluorescent foci that reveal themselves under fluoroscopic imaging. Flow cytometry can alternatively be used in the analysis, however there are additional logistic requirements for this method due to the use of specialist equipment. Using both an appropriate primary antigen that will seek out a specific molecule and a secondary antibody with a fluorescent tag, binding can occur between the antigens at the site of the target molecule and appear as visual foci. The histone H2AX, which plays a critical role in surveillance of genomic integrity, becomes phosphorylated in response to IR-induced DSBs to form γ H2AX [249, 252]. γ H2AX is therefore widely used today as a biomarker to signal DNA damage in mammalian cells. Foci can be detected almost immediately following irradiation of the cells, with the yield stabilising within 10 to 30 minutes [249, 253].

Care must be taken in verifying, as best as practically achievable, that observed foci indeed correspond to sites of DSB damage. It is known that γ H2AX foci are not limited to DSB sites, and can in fact appear at SSB sites as well [254]. Detectability of proteins at IR-induced foci is dependent on, among other factors, the amount of those proteins recruited to the damage site. As mentioned in Chapter 1, the NHEJ pathway dominates over the HR pathway in DSB repair for mammalian cells, by a factor of about 5:1 [255], however the components of NHEJ exhibit high expression and no more than a few molecules are necessary for repair [256, 257]. As a result, foci at these NHEJ-associated damage sites may be conflated with foci due to weak binding of the repair processing proteins to undamaged chromatin [258], which are also visible under fluorescence at high sensitivity. It is possible to discern DSB- and non-DSB-induced foci associated with NHEJ proteins using a cellular fractionation

assay [259].

Another commonly used biomarker is 53BP1, which, unlike γ H2AX, does not tend to accumulate at sites of replicative stress [260]. It has therefore been suggested that using both 53BP1 and γ H2AX in combination can discern IR-induced DSBs from those induced by replicative stress [261, 262]. The reader is referred to a recent review by Mavragani *et al.* [258] for an in-depth discussion on the suitability of other molecular DSB biomarkers, which is beyond the scope of this work. Additionally, a 2015 review by Rothkam *et al.* [263] provides a thorough discussion of the biological significance of DNA damage foci in general, including comparisons of the most common biomarkers.

Strides are continually being made to improve quantification of DSB induction. For example, a recently developed technique using confocal microscopy and 3D reconstruction in place of two-dimensional immunofluorescent imaging enables improved detection of DSB-induced γ H2AX foci [264].

3.4.1.2 Pulsed-field gel electrophoresis

In contrast to immunostaining techniques, pulsed-field gel electrophoresis (PFGE) [250], which superseded older techniques of neutral sucrose gradient sedimentation [265] and neutral elution [266], relies on physical separation of broken DNA fragments in order to quantify DSBs. In PFGE, cells are embedded in agarose, a polysaccharide matrix, due to its stability with the surrounding environment and its low reactive capabilities with DNA. Agarose consists of pore sizes that are of sufficient volume to separate portions of DNA molecules. Electrophoretic methods are then applied to separate DNA fragments based on their size across the agarose matrix. An alternating electric field is applied to the gel which causes DNA radicals to orientate themselves to align with the field. In this process, smaller fragments move quicker than larger ones, therefore allowing them to be sorted by size into the pores. The size of the

fragments can then be detected with a staining agent, such as ethidium bromide, which binds to the DNA and fluoresces under ultraviolet light. An advantage of PFGE when compared to other methods is that the DNA can be transferred for Western or Southern blot following the process.

PFGE generally requires that cells be exposed to larger doses of radiation (≥ 1 Gy) than for γ H2AX immunostaining [267], which has been used down to 1 mGy (for low LET photon radiation) [183]. Over-expression of a protein may generate artifacts [268], and it has been shown that a proportion of DSBs counted through PFGE may manifest as heat-induced artifacts during lysis when post-irradiation cell incubation is carried out at temperatures of around 37°C or higher [269, 270, 271]. The apparent DSB count can even be double the IR-induced amount between incubation at room temperature and 37°C [269]. The pH level has also been shown to affect DSB counts; alkaline conditions can promote more DSBs than neutral conditions [272]. Factors influencing assay sensitivity, including lysis conditions, pulse time and gradient of the electric field, temperature, and concentrations of both the cell sample and of agarose, have been explored by Blöcher and Kunhi [273]. Their studies were based on a particular type of PFGE called contour-clamped homogeneous electric field (CHEF) electrophoresis, which quantifies migrating DNA fractions without physical mass-based separation, although the results were expected to apply broadly.

It is understood that PFGE can only provide approximate DSB levels in mammalian cells. The process exhibits variability in DNA fragmentation, and therefore apparent DSB yield, throughout the cell cycle which does not necessarily reflect DSB-related mechanisms, but may be explained by DNA replication upon entry into S phase [274]. In some circumstances DNA replication is able to lead to DSBs that are formed physiologically rather than induced by IR [275, 276, 277, 278]. For example, DNA replication forks may progress over sites of SSB damage to form a DSB [278]. A recent study by Kawashima *et al.* [279] proposed a method for detecting DSBs at

DNA replication sites by combining PFGE with a Western blot assay to search for DNA replication markers. It must also be noted that the absolute number of DSBs counted will change in proportion to the DNA content of the cell so, for example, a higher DSB yield can be expected if the cells are irradiated during the G2 phase than during G1.

3.4.1.3 Neutral comet assay

Single-cell gel electrophoresis, better known as the comet assay [251], involves a similar process to that of PFGE, however differences in DNA damage and repair can be detected at the level of a single cell, which makes it particularly useful for examining biopsy specimens. Cells are embedded in agarose and lysed, and, upon application of an electric field, damaged DNA migrates across the gel, taking the appearance of a comet. (Undamaged DNA also migrates across the gel, but it remains in tact and does not produce a comet-like streak.) The comet assay has not been as useful as a quantitative DSB detection method, however there are improvements in this space with the advent of machine learning models. See, for example, recent work by Atila *et al.* [280] in classification of DSB damage using convolutional neural networks. In a 2012 study by Forchharmer and colleagues [281], it was shown that even when implementing a standard procedure across different laboratories, there were substantial disagreements between the number of measured strand breaks revealed by comet assay. The study concluded that comet assay validation trials should take further steps to evaluate the implementation of standard procedures.

3.4.2 Clonogenic survival assay

Cell survival following irradiation can be quantified by studying the surviving fraction of cultured cells in controlled experimental conditions. The technique, known as clonogenic survival assay, was introduced in the 1950s [91] to assess the reproductive

capabilities of single HeLa cells derived from a cervical carcinoma³. Modeling the surviving fraction as a function of dose is achieved through a two-parameter linear regression to the linear quadratic (LQ) model, covered in Chapter 1 (Section 1.3.3), against stratified sampling data. This model assumes that lethal IR-induced damage events are Poisson-distributed in proportion to the dose received by the cell [90]. The determination of RBE between radiation modalities using cell survival as the endpoint is common, and this is typically considered the endpoint that is most clinically relevant [205]. At this point, though, we must be clear on what is meant by “cell survival”; lethal events are understood to ultimately lead to a loss of reproductive integrity of the cell, at which point the death of the cell is marked as the loss of its ability to reproduce indefinitely. Interestingly, despite a cell being labelled dead by this definition, it may retain other capabilities, such as synthesising proteins or even replicating its DNA through mitosis, as long as it does not produce a significant number of progeny [84]. The RBE by cell survival is the ratio of doses required in the interest (e.g. protons) and reference (e.g. 6 MV x-rays) modalities to leave a prescribed fraction remaining⁴. In practice, the surviving fraction is actually calculated by taking the ratio of plating efficiency between the irradiated sample and the control. This plating efficiency refers to the proportion of cells seeded into the dish whose progeny go on to form a colony. Incorporating this measure into the calculation serves to normalise the surviving fraction to the control sample, and to filter out unwanted factors unrelated to IR-induced damage. It is vital that these factors that are external to the influence of the radiation are controlled as best as possible as they are significant sources of error.

One may refer to protocols detailed by Munshi *et al.* [84], Franken *et al.* [283] or other works in the existing literature for detailed protocols on clonogenic assay for

³The HeLa cell line is the most well-known and prominently used human cell line in cancer biology research, having originated from a single sample taken from a cervical biopsy during a woman’s cancer treatment in Maryland in the early 1950s. The fascinating history of HeLa cells is covered in a review by Masters [282].

⁴The mathematical derivation of RBE by cell survival is provided in Chapter 1, Section 1.3.5.

the purposes of radiation studies. Here we highlight only a pertinent selection of error sources. Samples must be kept in a stable environment of constant temperature and carbon dioxide concentration. Once removed from the incubator to be irradiated (or stained for counting) these environmental conditions are disturbed, which may confound those disturbances induced by radiation effects. In any laboratory that shares incubators for multiple experiments there is also the possibility of inadvertently disturbing the environment, or the cells themselves, as the incubator is accessed regularly during the incubation period, which can often last over a week. When the dish is bumped, colonies may break apart to artificially increase the calculated plating efficiency [84]. Initial preparation involves preparation of a solution containing a number of the cells of interest and a growth medium. Healthy cells will tend to adhere to the surface of the flask, so an enzyme called trypsin is introduced to loosen the cell membrane and help form a single-cell suspension [64, p. 36]. The concentration of cells in the suspension can be estimated using various methods, however a hemocytometer is often preferred due to its low cost and ease of use. This device is essentially a glass slide with etched grids. Cells are counted manually under a microscope within separate grid areas and the results are averaged, although human error can still introduce a large variance in the calculated concentration. Further, the suspension is assumed homogeneous when, in practice, gravity causes the cells to drift to the base of the flask. This problem is addressed by delicately swirling and “shaking” the flask prior to extraction. The number of cells seeded into the solution, or the dilution factor used to reduce this number, is carefully chosen such that on the day of the experiment the number of cells in the suspension is of the order of 10^6 and those plated onto the dish have a confluency of at least 70% [284]. Confluency is a measure of how much surface area is occupied by the cells. If the number of cells is very small then the statistics are not suitable to reliably assess plating efficiency and surviving fraction, but if the number is too large then colonies either tend to merge together and make the task

of counting them difficult [64, p. 37] or signalling networks may be altered [see, e.g., 285]. The desired number of seeded cells during the initial preparation is calculated using knowledge of the expected surviving fraction (from theory, previous studies or otherwise) and the doubling time of the cell line. This is approximately 12 to 14 hours for V79 Chinese hamster lung fibroblast cells, a commonly used cell line (discussed in more detail below) in *in vitro* irradiation experiments. After cells are irradiated they are put back in the incubator for another week or more, allowing for at least six cell divisions [284], before they are counted. A dye such as methylene blue is added to the dish. This is absorbed by healthy colonies and the excess can be washed away, which allows for simple identification of colonies by the human eye. However, some cells that have not lost their proliferation potential but are nonetheless damaged in other ways may not readily absorb the dye and may be missed in the count. Some cells may also fail to stick to the dish and will be washed away with the excess dye.

V79 Chinese hamster lung fibroblast cells are often used in radiobiological experiments due to a number of suitable characteristics including rapid growth rate (reduces overall time needed in the lab), high plating efficiency (limits influence of external factors) and quick recovery from being stored frozen in liquid nitrogen, which allows cells of similar characteristics to be easily stocked in large quantities [286]. V79-derived cell lines with different characteristics created through mutagenic agents are often quite stable, so experimentation can be carried out on variants (mutants) of standard V79 cells in which certain DDR pathways are inhibited and repair kinetics are interrupted. In this way it is possible to isolate, to a degree, the effects of initial IR-induced cellular damage from those following attempts by the cell at repair. Caution must be taken, however, in generalising DDR observations from V79 studies across distinct mutations [287] and, of course, in further generalising results to human cells or *in vivo* studies. For example, it has been shown that V79 cells do not express functional p53 [287], a protein that is crucial for tumour suppression in humans [288].

In a recent study, Mara *et al.* [289] investigated the accuracy of RBE prediction models in low $(\alpha/\beta)_x$ cell lines, concluding that these models are best suited to cell lines with large $(\alpha/\beta)_x$ values and insufficiently accurate for those with low values. This limits our ability to test these models on some common *in vitro* cell lines, including V79 cells, though it is promising for use on early-reacting healthy tissues *in vivo*.

The above commentary has been intended to inform the reader of the many aspects of a typical clonogenic survival assay procedure that may introduce substantial errors to the measured surviving fraction. These include human errors associated with preparation and measurement, subjectivity of measurements such as cell confluency, and variation that is inherent in the biological processes in response to non-IR-induced factors, including trypsinisation, absorption of staining dye, and environmental dependence (e.g. CO₂ concentration and temperature dependence). All this being said, biologists are well aware of the large variance in experimental results, which is often much larger than those encountered in experimental physics, and so a heavy emphasis is put on the reproducibility of results. Well-documented procedures are carefully repeated multiple times and several cell cultures are used in each experiment. It is also common to search for overall trends in the empirical data over many samples rather than relying on absolute measurements. Some of the validation techniques discussed above may be best used in combination. For instance, clonogenic cell survival trends have been successfully predicted in human squamous cell carcinoma by measuring residual DSBs, using γ H2AX immunostaining, in parallel to clonogenic assay 24 hours post-IR [290].

Due to the shouldered dose response curve, according to the LQ model the RBE can vary substantially with the surviving fraction that is chosen as the endpoint. A typical experiment may report, for example, RBE₁₀ or RBE₃₇, corresponding to surviving fractions of 10% or 37%, respectively. The former is arbitrary while the latter captures the natural *e*-folding (reduction of progeny by a factor of $\exp(-1)$) of

the Poisson arrival process governing lethal IR-induced events. As discussed in Section 3.2, it is most useful to compare cell survival trends to single- and multiple-track damage through the behaviour of survival at very low dose and very high dose. In this framework the focus shifts from the surviving fraction endpoint to the predictions of $\text{RBE}_{\max} = \alpha/\alpha_x$ and $\text{RBE}_{\min} = \sqrt{\beta/\beta_x}$, first defined in (3.3). In the case of the RMF model [201, 202, 203] (Section 3.2.9) there is an explicit relationship between the RBE for cell survival and the RBE by DSB induction, which can be determined by an auxiliary model such as the probabilistic Poisson model proposed in this Chapter. In general, sufficiently high-LET radiation accentuates the linear component of the cell survival curve to the point that $\mathcal{O}(\mathcal{D}^2)$ corrections can be neglected and the surviving fraction approximates $S = \exp(-\alpha\mathcal{D})$ [65, p. 69-70] [205]. Additionally, even up to very high doses of hundreds of gray, initial DSB formation scales linearly with dose [205]. In a first approximation, considering intra-track damages, the ratio of the linear LQ response parameters, α/α_x , between protons and the reference radiation follow the trend of the cell survival RBE and provide a useful maximum RBE estimate.

3.4.2.1 Degeneracy of the linear quadratic parameters

An anti-correlation exists between the α and β terms of the LQ model, meaning an increase in one may be compensated by a decrease in the other in such a way as to produce the same cell survival curve, within some degree of uncertainty. This has been discussed in some detail by Friedrich *et al.* [164], highlighting that “best-fit” parameters are not necessarily unique. Slight differences in either parameter can translate to differences in the calculated RBE, and this applies to both the proton data and the reference radiation data. A supplementary theoretical analysis of this degeneracy in the values of the LQ parameters is provided here.

Suppose there are two cell survival curves given by $S_1(\mathcal{D}) = e^{-\alpha_1\mathcal{D}-\beta_1\mathcal{D}^2}$ and $S_2(\mathcal{D}) = e^{-\alpha_2\mathcal{D}-\beta_2\mathcal{D}^2}$. Let us say these two curves describe the data equally well if the

predicted surviving fractions do not differ by more than $100\delta\%$ for any dose \mathcal{D} in the experimental range. Then we require, for all \mathcal{D} ,

$$(1 - \delta)S_1(\mathcal{D}) \leq S_2(\mathcal{D}) \leq (1 + \delta)S_1(\mathcal{D}) \quad (3.38)$$

$$\implies 1 - \delta \leq \frac{S_2(\mathcal{D})}{S_1(\mathcal{D})} \leq 1 + \delta. \quad (3.39)$$

Now, for small values of δ , truncating the Taylor series for $\log(1 + \delta)$ at the first term




Figure 3.6: Degeneracy in LQ parameters, α and β , for curves that predict cell survival within 10% ($\delta = 0.1$, black points), 5% ($\delta = 0.05$, blue points) and 2% ($\delta = 0.02$, red points) agreement. Data points were generated using a Monte Carlo sampling method. $\Delta\alpha$ and $\Delta\beta$ represent the change in α and β , respectively, between two equally valid survival curves.

gives

$$-\delta \leq \log \frac{S_2(\mathcal{D})}{S_1(\mathcal{D})} \leq \delta \quad (3.40)$$

$$\implies |\log S_2(\mathcal{D}) - \log S_1(\mathcal{D})| \leq \delta \quad (3.41)$$

$$\implies |\alpha_1 \mathcal{D} + \beta_1 \mathcal{D}^2 - \alpha_2 \mathcal{D} - \beta_2 \mathcal{D}^2| \leq \delta \quad (3.42)$$

$$\implies |\Delta\alpha \mathcal{D} + \Delta\beta \mathcal{D}^2| \leq \delta. \quad (3.43)$$

Figure 3.6 shows the degeneracy in the LQ parameters for curves that predict survival within 10% ($\delta = 0.1$), 5% ($\delta = 0.05$) and 2% ($\delta = 0.02$) agreement. These data points were generated using a simple Monte Carlo method; 100,000 $\Delta\alpha$ and $\Delta\beta$ values were generated from a random uniform distribution in the interval $(-0.5, 0.5)$ and were plotted on the axes if (3.43) was satisfied for all doses from 0 Gy up to 2 Gy in steps of 1 cGy.

3.5 Discussion

A simple biophysical RBE model based on DNA double strand break (DSB) induction has been developed for proton therapy. The model excludes physicochemical processes such as oxygen fixation and DNA repair, but captures the fundamental physical principles behind the enhanced cell-killing effect of high-LET radiation. As discussed above, the model may be used as an input into already-established biophysical cell survival models such as the RMF model [201, 202, 203], or it may be used on its own in a first approximation to increased cell-killing efficiency. The latter approach allows a pragmatic formalism for radiobiological optimisation of proton therapy treatment plans in which biological uncertainties can be reduced in normal tissues. This is similar to the approach of LET-weighted dose optimisation [see, e.g., 180, 181, 182]. The aim in such optimisation techniques is to shift the proton end-of-range toward

the central region of the target volume and away from the margins where zones of high-LET may overlap with critical structures [291]. Apart from LET-weighted dose planning, another heuristic approach to biological planning along the same lines is proton track-end penalisation [292]. In this methodology, regular dose-based inverse planning is carried out with additional objectives which penalise protons stopping in organs at risk.

The DNA damage model proposed in this Chapter is a single particle interaction model. A useful consequence is that the model scales linearly with dose and damage-weighted (or RBE-weighted) dose contributions due to overlapping proton pencil beams simply add together. Suppose for a given voxel in the dose grid of a patient plan there are N proton pencil beams, each contributing fluence spectra $\phi_k(E)$, $k = 1, 2, \dots, N$, to the final proton field $\Phi(E)$, as shown qualitatively in Figure 3.7. That is,

$$\Phi(E) = \sum_{k=1}^N \phi_k(E). \quad (3.44)$$

We find, from (3.36), that the dose-averaged DSB yield due to the proton field is given by

$$\langle Y_{\text{cd}} \rangle = \frac{1}{\mathcal{D}} \int_0^{E_{\text{max}}} \left(\sum_{k=1}^N \phi_k \right) \frac{S_{\text{col}}}{\rho} Y_{\text{cd}}(\lambda, \ell_0) dE \quad (3.45)$$

$$= \frac{1}{\mathcal{D}} \sum_{k=1}^N \int_0^{E_{\text{max}}} \phi_k \frac{S_{\text{col}}}{\rho} Y_{\text{cd}}(\lambda, \ell_0) dE \quad (3.46)$$

$$\implies \mathcal{D} \langle Y_{\text{cd}} \rangle = \sum_{k=1}^N \mathcal{D}_k \langle Y_{\text{cd}} \rangle_k \quad (3.47)$$

where \mathcal{D}_k is the dose deposited in the voxel due to the k -th pencil beam. It is clear from (3.47) that DSB yield, as predicted by the damage model, due to multiple pencil

beams scales linearly with dose. This naturally extends to the RBE-weighted dose;

$$\mathcal{D}_{\text{RBE}} = \sum_{k=1}^N (\mathcal{D}_{\text{RBE}})_k. \quad (3.48)$$

In the context of iterative inverse planning algorithms, which is explored in Chapter 6, this linearity enables the pencil beam dose matrix to be calculated just once, prior to optimisation. In standard inverse planning of intensity modulated proton therapy (IMPT) one may refer to dose *kernels* as the building blocks of a treatment plan. These kernels describe the dose deposition due to a single proton pencil beam chosen for the plan, taking into account the positioning of the beam and the heterogeneity of patient anatomy. Using the RBE model discussed here, RBE-weighted dose kernels can be similarly defined. Analytical proton dose kernels have been proposed [293, 294], although simulating the beam by Monte Carlo methods is more accurate due to the stochastic nature of dose deposition. In either method, only the local energy spectrum, Φ , in a voxel is needed to calculate the RBE-weighted dose.

It must be stressed that linearity of the RBE-weighted dose applies only under the assumption that second-order dose effects, such as the formation of clustered lesions due to multiple particle tracks in close proximity, are negligible. If the effect of combining pencil beams were non-linear then the matrix would require recalculation each time the relative pencil beam weights were changed. There is a well-known saturation, or “overkill”, effect in particle therapy in which the RBE tends to reach an inflection point and decrease above a certain LET. This overkill zone is in the LET range of 100 keV/ μm to 150 keV/ μm . Compared to heavier ions, the size and charge of the proton, and associated energy straggling [223], keeps the maximum proton LET down to well below around 30 keV/ μm in clinical beams. Therefore, no overkill effect is expected and, instead, the RBE rises monotonically as proton LET increases. Additionally, in proton therapy (and conventional radiotherapy) the prescribed dose is

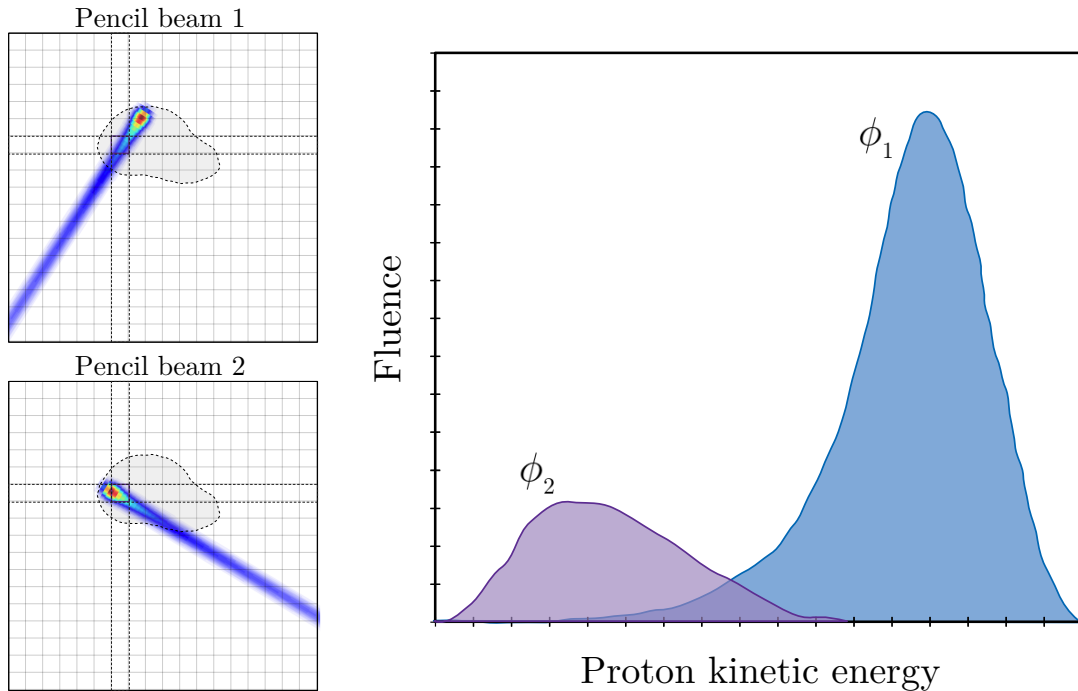


Figure 3.7: Qualitative representation of the addition of two pencil beam energy spectra in overlapping voxels. $\phi_1(E)$ is the fluence spectrum due to pencil beam 1 in the highlighted voxel and $\phi_2(E)$ is the spectrum for pencil beam 2. The total spectrum in the voxel of interest is $\Phi = \phi_1 + \phi_2$. Gray shaded region represents an arbitrary target volume.

usually delivered in fractions of less than 3 Gy, so non-linear dose effects are mitigated. This is not the case for hypofractionated radiotherapy, in which fewer fractions are used to deliver a higher dose per fraction, and this should be taken into consideration when implementing the DNA damage model in these circumstances.

The proposed DNA damage model has been formulated under an assumption of a hypoxic environment in order to limit the effects of the oxygen fixation and chemical repair. As introduced briefly in Chapter 1 (Section 1.3.2), hypoxic tissues exhibit up to a threefold reduction in radiosensitivity compared to their normoxic counterparts when irradiated by photons [295, 296, 297, 298, 299]. The oxygen fixation hypothesis (OFH) [300] postulates that this oxygen enhancement effect is due to radical-induced indirect DNA damage becoming “fixed” by molecular oxygen, inhibiting chemical repair processes from restoring the damage. In contrast, without oxygen, radiosensitising peroxy radicals are unable to form and chemical repair is mostly successful.

Mechanistic models based on extensions of the OFH have been proposed for explaining the oxygen enhancement effect [301, 302]. In the high-LET regime of heavy ions oxygen enhancement is less pronounced due to the number of direct ionisations of the DNA being sufficiently large and dominating the contribution to complex damage. In the DNA damage model proposed in this work, oxygen enhancement effects were avoided so that direct damage from ionisations within the DNA would dominate over indirect damage and so that the model would not overestimate the lethality of radiation to the tumour. In the majority of cases, aberrant vascularisation and compromised blood supply in solid tumours induces hypoxic regions [see, e.g., 303]. Future work will address the inclusion of oxygen concentration into the model, which will avoid underestimation of damage to healthy, well-oxygenated tissues.

The modularity of the model allows for the initial DSB yield to be calculated independently of the oxygen enhancement. It has been demonstrated that the rate of DSB rejoining does not change with oxygen concentration [304], however the relative yield of lethal DNA radicals from indirect damage is related to the amount of oxygen binding that is able to occur. This exhibits a saturation behaviour governed by a second order differential equation, the solution to which is a rational function dependent on the oxygen concentration [305, 306]. This rational form has been used by Stewart *et al.* in modelling the oxygen enhancement effect in the MCDS microdosimetric Monte Carlo software [159, 160, 224, 225, 226]. In [224] initial DSB formation is assumed at the upper limit, in normoxic conditions, with the yield reduced by a multiplicative factor representing the fraction of DNA radicals removed through chemical repair. The impact of oxygen concentration on DSB yield was investigated for protons using MCDS, the results of which are shown in Figure 3.8. The oxygen enhancement ratio (OER) was calculated as the ratio of the predicted DSB yield at a specified oxygen concentration to that at zero oxygen partial pressure. (Note that the plasmonic correction to the proton mean free path used in this work is not used

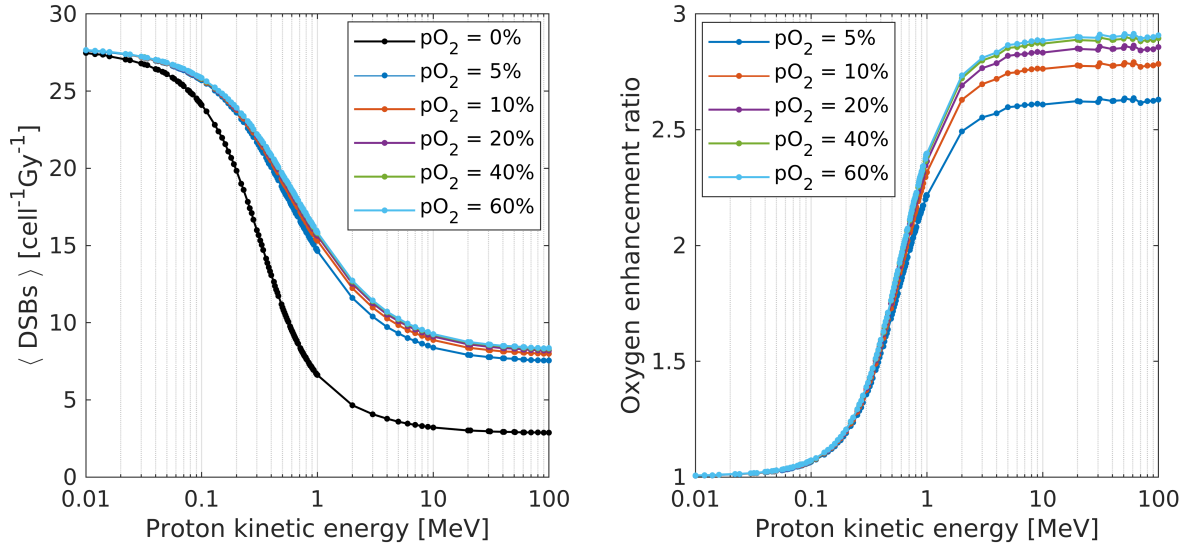


Figure 3.8: Left: predicted DSB yield using MCDS [159, 160] for different oxygen concentrations (0%, 5%, 10%, 20%, 40% and 60%). Right: the oxygen enhancement ratio (OER) as a function of proton energy, calculated as the ratio of the predicted DSB yield at a specified oxygen concentration to that at zero oxygen partial pressure.

in MCDS calculations.) A saturation behaviour is evident; for proton energies up to 1 MeV the OER for all non-zero oxygen concentrations follows the same trend, increasing up to approximately 2.0. There are discrepancies of up to 10% between the OER values for different oxygen levels at energies above 1 MeV, however the change in the absolute contribution to damage is only about 1 DSB cell⁻¹Gy⁻¹. The consideration of indirect DNA damage may warrant an increase in the diameter of the sensitive cylindrical volume, d , in (3.30a) and (3.30b). This also serves to best retain the physical interpretation of the interaction-to-DSB conversion efficiency parameter, ε , in (3.29).

Implementation of the proposed RBE_{cd} model in IMPT inverse planning fits under the label of a Level III biological optimisation, using the framework of Nahum and Uzan [154] outlined in Chapter 1 (Section 1.5). Through the inclusion of oxygen concentration this can be extended to Level IV. In practice the voxelised dose grid can be registered to functional positron emission tomography (PET) images

that contain local oxygenation information [305]. This information can be gathered, for example, from differences in the uptake of radiopharmaceutical agents such as ^{18}F -fluorodeoxyglucose (FDG). Changes in tumour metabolism are an indicator of hypoxia levels and can be signified through the uptake of FDG for aerobic glycolysis, known as the Warburg effect [307]. This is explored in detail by Rejandran *et al.* [308]. The proposed use of both radiological and functional (or molecular) images to inform a “radiobiological” target volume for planning led to dose- and LET-painting techniques [309, 310], with success demonstrated in IMPT [311, 312]. Regions of interest identified as hypoxic could be prescribed a dose boost through either delineation of new sub-volumes within the target region or by a “paint-by-numbers” approach on a voxel-by-voxel basis [310]. Models and algorithms for biological optimisation have since been proposed for particle therapy based on continuous measures of local oxygenation [313, 314, 315].

The incorporation of oxygen concentration into the DSB induction model also opens the possibility of building an optimisation framework that accounts for the enhanced tissue-sparing effects of ultra-high dose rate radiotherapy, referred to as *FLASH* therapy [316, 317, 318]. This is currently being developed. A potential mechanism of the tissue-sparing effect in *FLASH* therapy is that of transient hypoxia induced by rapid oxygen depletion. At dose rates of tens of gray per second or higher, a depletion model can explain the apparent decrease in radiosensitivity for normoxic tissues of low or intermediate oxygen tension [319].

The modularity of the damage model has been mentioned above for the inclusion of an independent oxygen enhancement model. In general one may introduce modular additions of independent chemical and biological processes, which depend on the initial DSB yield. For example, in addition to the mechanistic RBE models discussed in Section 3.2, a 2018 study by Wang *et al.* [320] modelled initial IR-induced DSB formation and subsequent repair through the NHEJ pathway. A Poisson-based formalism

was used to model the surviving fraction of irradiated cells, which was reduced to the standard LQ model through a truncated Taylor expansion. The linear response, α , includes separable terms related to DSB yield, enzymatic repair success, and overkill effects (which for therapeutic protons are negligible). While the phenomenological models discussed in Section 3.2 do not explicitly model repair, mechanistic biophysical models such as the one presented in this work may be formulated such that independently designed repair models may be added.

Throughout this Chapter the DSB, or complex damage, induction likelihood has referred to the likelihood of a single particle track giving rise to a cluster of two or more ionisations within ten basepairs. To further generalise the model and take full advantage of its Poisson nature, different ionisation cluster size possibilities can be considered separately. This enables narrower classification of damages. For completeness, the single strand breaks (SSBs) may also be considered; we may alter (3.29) to model the relative SSB induction probability by replacing the summation with only a $k = 1$ term, which represents a cluster size of exactly one. That is,

$$\mathcal{P}_{\text{SSB}}(\lambda|L) = \varepsilon \frac{L}{\lambda} \exp\left(1 - \varepsilon \frac{L}{\lambda}\right). \quad (3.49)$$

Note that the exponential term differs from the standard Poisson definition due to a normalisation constant of $1/e$. As in (3.31), the likelihood measure for SSB induction can similarly be expressed as

$$\mathcal{L}_{\text{SSB}}(\lambda; \ell) := f_L(L) \mathcal{P}_{\text{SSB}}(\lambda|L). \quad (3.50)$$

Similar to the methods in Section 3.3.2, the SSB induction model for protons was parameterised using simulated SSB counts collected from MCDS [159, 160]. Figure 3.9 shows that a reasonable fit ($\varepsilon = 1.128$, 95% C.I. (1.068, 1.188)) can be obtained for the low-energy, high-LET part of the proton spectrum, however the prediction

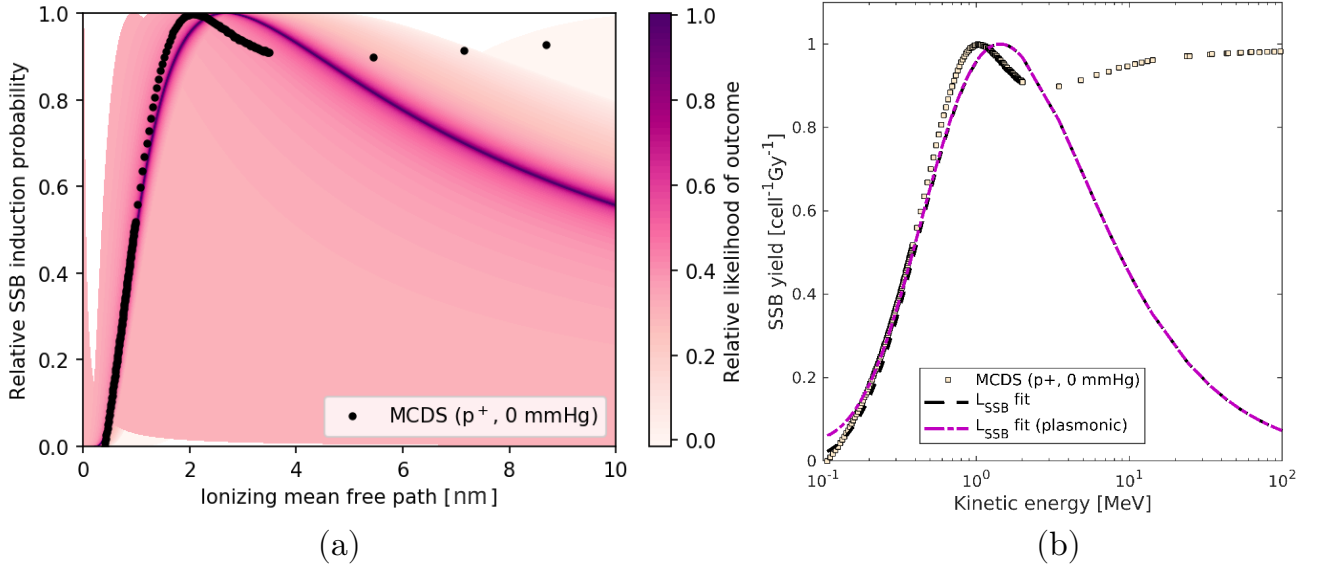


Figure 3.9: (a) Relative SSB induction likelihood, according to (3.50), with fitting parameter $\epsilon = 1.128$ and 95% confidence interval (C.I.) of (1.068, 1.188), fit to Monte-Carlo-generated SSB counts using MCDS version 3.1 [159, 160]. (b) Most likely SSB induction response shown as a function of proton energy, against the simulated MCDS data.

deviates from the data around the peak of SSB production and beyond the inflection point. The apparent increase in SSB production for proton energies of around 2 MeV or higher cannot be explained by a single-particle Poisson model, and is likely the result of simulated secondary particles being included in the damage contribution.

The extension of the model presented in this work to heavier ion therapies is beyond the scope of this thesis. It is, however, important to note that doing so requires knowledge of the ionising mean free path of these particles, which is not as extensively covered in the current literature as it is for electrons and protons [see, e.g., 321]. The Van den Heuvel DNA damage model [227] outlined in Section 3.2, which inspired the damage model presented in this work, has a formulation explicitly in terms of particle energy. This makes it a suitable candidate for modelling damage due to particles other than protons and electrons, however a consequence of this approach is that plasmonic effects can no longer be considered.

In summary, a simple biophysical model has been proposed based on the relative induction likelihood of complex clustered DNA damage and its applicability to IMPT

treatment planning has been explored as a pragmatic approach to handling enhanced cell-killing effects at the proton end-of-range. Techniques for validating this model, or any DSB-based RBE model, and their limitations were explored in detail, and possible extensions to the model have been discussed, such as the inclusion of oxygenation and chemical repair modelling, and narrower classification of damage based on ionisation cluster sizes. In the Chapter that follows, we explore the implementation of the model in a series of Monte Carlo studies and compare its predictions to *in vitro* data from literature against a selection of RBE models discussed in Section 3.2.

Chapter 4

Monte Carlo simulation studies and comparisons with *in vitro* data

4.1 Introduction

The previous Chapter outlined the development of a single-particle DSB induction model for predicting RBE in proton therapy. Hereon we refer to this as the Poisson-Distributed Lesions (PDL) model for the sake of convenience. In this Chapter we explore its practical implementation in a series of Monte Carlo (MC) simulations, test its predictions against published cell survival data, explore its use in a treatment planning system, and suggest both simplified analytical and MC techniques for its widespread use.

We begin by outlining the methods by which the PDL model may be directly incorporated into Monte Carlo software and provide relevant programmatic scripts in Appendix B. RBE_{cd} predictions are investigated in comparison to dose and LET for a range of proton pencil beam energies impinging on a water phantom. This is followed by a simulation study of modulated proton beams as they are used to form SOBPs. Corresponding RBE_{cd} and LET values are used to fit an empirical RBE_{cd} -LET relationship, which is then used to estimate the RBE distribution in an IMPT treatment plan using an open source treatment planning system. Finally, the PDL model is discussed in the context of cell survival RBE and the predictions of the

model are compared to a catalogue of 30 independent *in vitro* studies comprising 161 data points.

4.2 Materials and Methods

4.2.1 Monte Carlo studies

TOPAS [161] (version 3.1.2), which wraps and extends the Geant4 simulation toolkit [162], was used to investigate the damage model proposed in Chapter 3 through a series of MC simulations. The particle tracking step-size and included physics processes were unaltered from the default settings in TOPAS, which have been optimised for proton therapy applications [322].

We explore, with the parameterised model, the behaviours of DSB production, the associated RBE_{cd} , and the relationship between RBE_{cd} and LET, with depth for proton beams of varying energies. This investigation is conducted with monoenergetic proton pencil beams, which produce a pristine Bragg peak, as well as modulated polyenergetic beams, which produce a spread out Bragg peak (SOBP). We compare our results, where possible, both with experimental findings in the existing literature and with a selection of RBE models discussed in Chapter 3, Section 3.2.

4.2.1.1 Creating a TOPAS extension for scoring DSB damage

TOPAS enables the compilation of user-defined extensions, written as C++ scripts which call standard Geant4 classes. In order to directly score the expected number of DSBs in a voxelised geometry, a binned scorer, which inherits from a TOPAS base class of `TsVBinnedScorer`, was created to accumulate damage contributions due to each event. An event is defined as a single particle having entered a voxel of interest and is handled through a `ProcessHits` method. The aggregation and output of data is handled by the base class.

The calculation of (3.36) consists of first checking the type of particle that triggered the event. If the particle is not an electron nor a proton, or if its kinetic energy is less than the plasmonic threshold, E_{th} (see Table 3.2), then there is no contribution to damage and the accumulator remains unchanged. If the particle is an electron or a proton with large enough kinetic energy, the inelastic mean free path is calculated using (3.33), followed by Y_{cd} , using (3.34). The energy deposited to the medium, ϵ , and the voxel volume, ν , can be computed using a standard MC dose calculation method. The dose deposited by the k -th particle entering a voxel, d_k , is given by

$$d_k = \frac{\epsilon_k}{\rho\nu} \quad (4.1)$$

where ρ is the mass density of the medium in this voxel, assuming this is homogeneous within the voxel. The final step is to multiply the damage yield, Y_{cd} , by the dose and the statistical weight of the Monte Carlo track, w_k , to give the final contribution. Over all particles traversing the voxel, the dose-averaged DSB yield is given by

$$\langle Y_{\text{cd}} \rangle = \frac{1}{\mathcal{D}} \left[\sum_{k=1}^{N_{p^+}} (Y_{\text{cd}}^{p^+})_k d_k^{p^+} w_k^{p^+} + \sum_{k=1}^{N_{e^-}} (Y_{\text{cd}}^{e^-})_k d_k^{e^-} w_k^{e^-} \right] \quad (4.2)$$

where the first summation is over all N_{p^+} protons, the second summation is over all N_{e^-} electrons, and $\mathcal{D} = \sum_{k=1}^{N_{p^+} + N_{e^-}} d_k$ is the total dose scored in the voxel by proton and electron contributions. The quantity inside the square brackets is what is accumulated by this custom TOPAS scorer. The divisor, the total dose, \mathcal{D} , can be collected using the default TOPAS dose scorer. The C++ scripts compiled as TOPAS extensions, which score proton and electron induced damage, are located in Appendix B.

4.2.1.2 RBE calculations in pristine proton pencil beams

Proton pencil beams with nominal beam energies of 60 MeV to 120 MeV, in intervals of 10 MeV, were created in TOPAS and fired from an ideal (theoretical) beam source toward a cylindrical water phantom of radius 5 cm and length 42 cm. Each simulation generated 5×10^7 primary proton histories. The beam source was placed coaxial to the cylinder and 4 cm from the entry surface. Details of the beam parameters can be found in Table 4.1. A vacuum wire-frame parallel world scorer was defined for each beam such that the expected DSB yield, LET, particle fluence, dose and other microdosimetric quantities could be extracted in 30 radial bins of 1 mm and between 40 and 130 depth bins of 1 mm, depending on the beam energy. The number of depth bins was chosen to score up to 20% beyond the proton range, as calculated by the continuous slowing down approximation (CSDA), for each energy. These query depths should not be confused with the step-size in the MC simulation itself. A reference 6 MV x-ray beam comprised of 5×10^8 primary photon histories was defined in TOPAS to generate a 10 cm \times 10 cm square field. The same scoring geometry as detailed above was used to calculate the DSB yield due to electrons in each voxel.

Additionally, kinetic energy spectra were collected at the upstream entry surfaces of each voxel in an effort to visualise damage contributions in terms of the energy spectra in selected voxels. Filters were applied within the TOPAS implementation to separately score electrons, primary protons and secondary protons. Phase space scorers allow for raw data on surface scorers to be recorded, such as absolute position, energy, direction and particle type. However, this approach was computationally demanding both in the collection and post-processing of data. Instead, binned scorers were set up to dynamically increment the particle count in each voxel and energy bin as a particle with a certain energy was registered. Energies were scored from 10 eV to 10 keV in bin sizes of 10 eV for electrons and from 1 keV to the maximum proton energy in 1 keV bins for protons.

We saw in Chapter 3 that a number of studies have ventured to find analytical correlations between RBE and LET, often postulating that the RBE increases linearly with LET [see, e.g., 185, 186, 189, 190, 191]. Therefore, corresponding pairs of RBE_{cd} and LET data in each voxel were used to fit an approximate empirical relationship, $RBE_{cd} \approx f(L_{\Delta})$, between these quantities for track- or dose-averaged LET, L_{Δ} .

4.2.1.3 RBE calculations in the SOBP region

In contrast to a single monoenergetic proton pencil beam, which produces a sharp peak in the depth-dose distribution, an appropriately weighted assortment of pencil beams at different energies can produce a spread out Bragg peak (SOBP). An SOBP spreads a more uniform dose along the depth of the target region, leading to more predictable tumour control. Three separate SOBPs were created in TOPAS for three maximum proton energies, 60 MeV, 100 MeV and 180 MeV, using the methods of Jette and Chen [323] and Bortfeld and Schlegel [55] (see also [56]), as outlined below.

We wish to create a proton SOBP which covers a fraction, $0 < \chi < 1$, of the full range, R_0 , of the largest energy pencil beam. The largest energy pencil beam has a kinetic energy of E_0 , and the range R_0 can be calculated using (1.12). If n pencil beams are chosen to form the SOBP, then the first step is to determine the energies necessary to place the proton end-of-range at the end of $n - 1$ equal intervals within the SOBP region. The range of the k -th pencil beam for $k = 0, 1, \dots, n - 1$ is given by

$$r_k = \left[1 - \left(1 - \frac{k}{n-1} \right) \chi \right] R_0 \quad (4.3)$$

and the corresponding energies are given by

$$e_k = \left(\frac{r_k}{\alpha} \right)^{1/p_0} \quad (4.4)$$

where $\alpha = 0.0022$ and $p_0 = 1.77$ as in (1.12). The final step is to set the relative

(normalised) weights, w_k , of each pencil beam so that a flat SOBP is produced. Jette and Chen [323] gave these weights in terms of a fitting parameter p which is not necessarily the same as the p_0 value used by Bortfeld and Schlegel [55]. Optimal values can be found for a limited number of maximum proton energies and SOBP widths in [323, Table 1]. The weights are given by

$$w_k = \begin{cases} 1 - \left(1 - \frac{1}{2(n-1)}\right)^{1-1/p}, & k = 0, \\ \left(1 - \frac{1}{n-1} \left(k - \frac{1}{2}\right)\right)^{1-1/p} - \left(1 - \frac{1}{n-1} \left(k + \frac{1}{2}\right)\right)^{1-1/p}, & k = 1, 2, \dots, n-2, \\ \left(\frac{1}{2(n-1)}\right)^{1-1/p}, & k = n-1. \end{cases} \quad (4.5)$$

Details of the beam parameters used in TOPAS for each of the SOBPs can be found in Table 4.1. Dose, LET and damage were scored along the central axis with the methods described in the previous subsection. Statistical variation in the calculated RBE_{cd} was analysed after halving and doubling the number of beam energies used to produce the SOBP. For the 60 MeV and 100 MeV SOBPs, 5, 10 and 20 energy intervals were used. For the 180 MeV SOBP, 10, 20 and 40 intervals were used.

It is hypothesised that if the number of Bragg peaks forming the overall SOBP is increased then the variation in the RBE_{cd} will decrease. The reasoning behind this is that many closely positioned Bragg peaks should allow for a smooth and monotonic increase in LET, whereas, for example, two Bragg peaks separated by some distance allow for a drop in LET between the peaks.

4.2.2 RBE calculation in a treatment planning system

Thus far, RBE_{cd} calculations have been made using MC simulation data from pristine pencil beams or SOBPs in a homogeneous water phantom. We now explore the RBE_{cd} spectra in the context of a radiotherapy treatment plan, with potentially

Table 4.1: Beam parameters used to generate monoenergetic and modulated pencil beams in TOPAS [161]. SOBP width, χ , and weighting parameter, p , are implemented in (4.4) and (4.5) respectively.

	Monoenergetic proton pencil beams	Modulated proton pencil beams		
Energy spread (% max energy)	0.758	0.758		
Position distribution	Gaussian	Flat		
Position cutoff shape	Ellipse	Rectangle		
Position spread in X direction (cm)	0.65	1.0		
Position spread in Y direction (cm)	0.65	1.0		
Angular distribution	Gaussian	Gaussian		
Angular spread in X direction (rad)	0.0032	0.0032		
Angular spread in Y direction (rad)	0.0032	0.0032		
Energy spectrum type	N/A	Continuous		
SOBP beam energy (MeV)	N/A	60	80	100
SOBP width, χ	N/A	0.25	0.25	0.25
Weighting parameter, p	N/A	1.21	1.42	1.69

over a thousand individual pencil beams of different energies, incident from multiple directions on a heterogeneous target.

The open source treatment planning software *matRad*¹ [163] (DKFZ, Heidelberg, 2020) was used to run an inverse optimisation procedure, based on physical dose only, on a head and neck case provided as an example with the *matRad* distribution. The plan included two planning target volumes (PTVs) of prescription doses 63 Gy and 70 Gy, and 21 avoidance structures. Delineated structures were defined on a planning CT with voxels of side lengths of 0.94 mm and depths of 3 mm. A three-field IMPT treatment was to be delivered in 30 fractions. Dose calculations in *matRad* are based on the methods of Hong *et al.* [293] and the analytical proton pencil beam model of Bortfeld [56], while multiple scattering is approximated with the Highland formula [324]. Appropriate dose spot locations (see Chapter 1, Section 1.4.1) are determined

¹Available at: <https://e0404.github.io/matRad/>.

based on the radiological depth necessary to place the Bragg peak at regular spacing within the target regions. Tabulated depth dose profiles and lateral Gaussian spread parameters are interpolated to estimate the dose deposited in each voxel. A regular lateral spot distance (referred to as “bixel width” in the software) of three voxels was used. An interior point optimisation (IPOPT) method [325] is used to perform inverse optimisation in matRad.

The matRad software includes an LET calculation subroutine which employs similar methods to the dose calculation. The empirical RBE_{cd} -LET relationship, $f(L_{\Delta})$, derived from the aforementioned MC studies allowed for this subroutine to be altered such that the predicted RBE_{cd} was returned instead of the LET. The optimised dose distribution due to a single fraction was compared to the corresponding RBE_{cd} distribution.

4.2.3 Comparisons with *in vitro* experimental data

4.2.3.1 The Particle Irradiation Data Ensemble

The *Particle Irradiation Data Ensemble*² (PIDE) is a database released by GSI (GSI Helmholtzzentrum für Schwerionenforschung GmbH, 2020) [164] which features compiled results of *in vitro* cell survival experiments gathered from a literature survey. Each included experiment has reported dose response curves for both ion and photon (reference) irradiation, allowing for RBE to be calculated using the surviving fraction (SF) endpoint. The database collates information such as the cell line used, cell cycle phase, genomic length of diploid cells, ion species used, LET and energy of ions at the target, LQ dose response parameters, α and β , for both the ion and reference radiation, and whether measurements were taken at a pristine Bragg peak (PBP) or in an SOBP.

It should be noted that this collation is of many different studies performed under

²Available at: <http://gsi.de/bio-pide>. Free registration required for access.

different conditions, with differing methods, spanning over five decades. Caution must therefore be taken in comparing the results of these experiments under the same framework. As noted by Friedrich *et al.* [164], individual experiments often published either the averaged LET of the ions or their kinetic energy, but not both. In such cases the missing data was approximated using in-house software at the GSI biophysics department. Further, different studies calculated LET by different methods, which were not always explicitly stated. Of those that did report the calculation method, some used track averaged LET while others used dose averaged LET; we know from Chapter 1, Section 1.2.6 that these quantities can be significantly different from each other. Different studies also use different photon reference modalities. Orthovoltage x-rays, for example, are known to exhibit their own RBE with respect to megavoltage reference beams or ^{60}Co γ -rays. Friedrich and his colleagues also cite uncertainties in plating efficiencies (see Chapter 3, Section 3.4) as being large sources of error that were not able to be quantified for most studies. This affects the integrity of the cell survival data and consequently the LQ parameters. PIDE version 2.0 was used from the beginning of this study, however a recent release (version 3.1) has since been obtained, detailing 161 experimental proton data sets from 30 separate studies. Version 3.1 includes both the LQ parameters reported in the original paper and those calculated by GSI using the raw experimental data. It also fixed errors in the previous database and added new studies to it. The data used in this thesis has since been checked against PIDE version 3.1 and updated where necessary.

4.2.3.2 Calculation of RBE

The RBE by surviving fraction (SF), RBE_{SF} was calculated using (1.26) and (1.27) for all proton data in the PIDE database. Surviving fractions from 5% to 95%, in intervals of 5%, were assessed. $\text{RBE}_{\text{max}} = \alpha/\alpha_x$ and $\text{RBE}_{\text{min}} = \sqrt{\beta/\beta_x}$ values were calculated according to the phenomenological models described in Chapter 3, Section

3.2 (summarised in Table 3.1). These were used to calculate the predicted RBE at all SF levels, using (3.2), and compare these results to the empirical data from the PIDE database.

In addition to the PDL model, the experimental data was also compared to two related models using the RBE_{SF} framework in (3.2). Under the assumptions of the PDL model, we set

$$RBE_{\max} = f(L_{\Delta}), \quad (4.6)$$

$$RBE_{\min} = 1, \quad (4.7)$$

where $f(L_{\Delta})$ is the empirical RBE_{cd} -LET relationship derived from the MC data, mentioned in Section 4.2.1.2. Additionally, the PDL model was used directly with the mean proton energy at the target, given in the PIDE database;

$$RBE_{\max} = \frac{Y_{cd}^{p+}(\lambda(E), \ell_0)}{\langle Y_{cd}^{e-} \rangle_{\text{ref}}}, \quad (4.8)$$

$$RBE_{\min} = 1. \quad (4.9)$$

In (4.8), $Y_{cd}^{p+}(\lambda(E), \ell_0)$ is the calculated DSB yield due to protons with kinetic energy E and interaction length ℓ_0 with the DNA, as in (3.34). Note that, unlike in the MC simulations, the local energy spectrum at the cell target is unknown for the experiments in the PIDE data. We must therefore use the singular energy value given in the dataset, or estimate this quantity from the LET. $\langle Y_{cd}^{e-} \rangle_{\text{ref}}$ is the mean DSB yield over the entire electron energy spectrum due to a reference beam. Here, for an x-ray reference, it is assumed the DSB yield per unit dose is approximately constant and can therefore be reduced to a single value that is independent of the local energy spectrum. We explore this assumption further in Section 4.3.1.




Figure 4.1: RBE_{cd} , calculated using (3.35), along the central axis of pristine proton pencil beams in water between 60 MeV and 120 MeV, shown alongside the relative dose deposited by each beam.

4.3 Results

4.3.1 Monte Carlo studies

The RBE_{cd} , calculated using (3.35) alongside the custom TOPAS scorer detailed in Section 4.2.1.1, is shown against the relative dose as a function of depth for each of the proton pencil beams in Figure 4.1. RBE_{cd} values were calculated by dividing the proton modality DSB yield by the reference modality DSB yield in corresponding voxels. The RBE_{cd} for all beams is close to 1.0 in the entry region and rises steeply prior to the Bragg peak. Both the maximum RBE achieved, at the end-of-range, and the steepness of the RBE_{cd} -depth profile are largest at the lowest beam energy of 60 MeV. Figure 4.3 demonstrates a general decrease of the RBE_{cd} with increasing beam energy, both at the Bragg peak and at the distal edge (defined at 5% of the maximum dose). A maximum RBE_{cd} of approximately 2.0 was calculated for the 60 MeV pencil beam while 1.6 was calculated for the 120 MeV beam. The RBE_{cd} -depth profile is

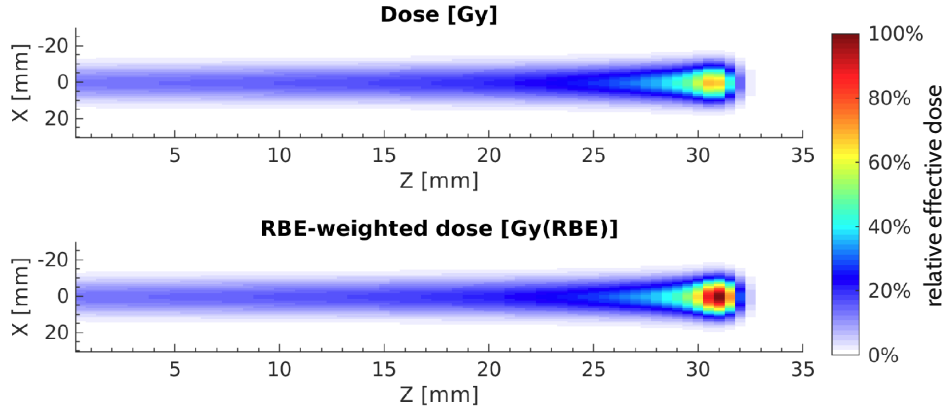


Figure 4.2: Visualised dose and RBE_{cd} -weighted dose scored in TOPAS [161] for a pristine 60 MeV proton pencil beam in water.

seen to closely follow the dose-averaged LET profile toward the distal dose fall-off in Figure 4.4. The LET increase, however, is more gradual while the RBE_{cd} increase is sharper.

A negligible difference in electron-induced DSB yield was found across all voxels for the 6 MV x-ray reference beam, as is expected from low LET photon irradiation [205]. The mean DSB yield was $2.839 \text{ cell}^{-1}\text{Gy}^{-1}$ with only a $0.001 \text{ cell}^{-1}\text{Gy}^{-1}$ difference expected within three standard deviations. Further, it was found in all proton beam simulations that the electron contribution to dose was less than 2% of that from protons. Using these two pieces of information, we may define a simplified version of the damage model in (3.35) that requires only the local proton energy spectrum;

$$\text{RBE}_{\text{cd}} = \frac{\langle Y_{\text{cd}}^{p+} \rangle}{2.839}. \quad (4.10)$$

Here, $\langle Y_{\text{cd}}^{p+} \rangle$ is the expected DSB yield per unit dose due exclusively to protons. Note that this is lower than a standard low-LET experimental DSB yield due to (1) the lower estimate of cellular DNA content (1 Gbp) in the MCDS simulations and (2) the fact that only intratrack direct damage is considered. However, the effect on proton DSB yield is proportional, so the ratio of these quantities is unaffected.

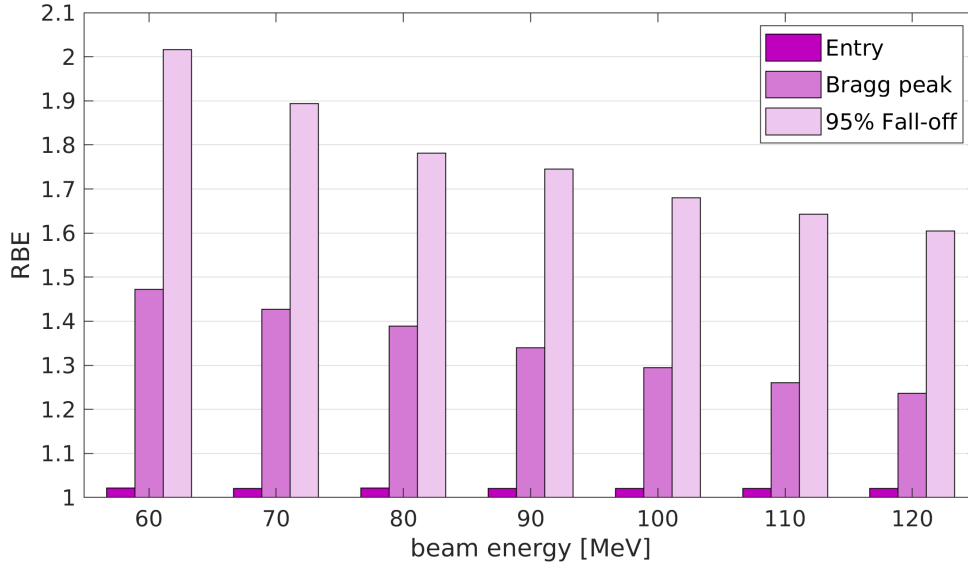


Figure 4.3: Summary of RBE_{cd} calculations shown in Figure 4.1, showing a decrease in RBE, both at the Bragg peak and at the 95% dose fall-off, as the beam energy is increased.

LET and RBE_{cd} values were compared for 550 data points, which accounted for each beam up to the 95% dose fall-off. A clear relationship is shown between these quantities in Figure 4.5, especially between the RBE_{cd} and the dose-averaged LET. The RBE_{cd} remains flat around 1.0 for LET values of less than $3 \text{ keV}/\mu\text{m}$ before rising monotonically in an approximately linear fashion toward the highest LET value scored ($\sim 23 \text{ keV}/\mu\text{m}$). A rational function, f , was fit to the data;

$$\text{RBE}_{\text{cd}}(L_{\Delta}) = f(L_{\Delta}) = \frac{p_1 L_{\Delta}^2 + p_2 L_{\Delta} + p_3}{L_{\Delta} + q_1}, \quad (4.11)$$

where L_{Δ} is the LET (dose-averaged or track-averaged) in $\text{keV}/\mu\text{m}$ and p_1 , p_2 , p_3 and q_1 are the fitting parameters, shown in Table 4.2. The coefficient of determination for this fit was $R^2 = 0.999$ against LET_d and $R^2 = 0.984$ against LET_t .

Nine SOBPs were generated in TOPAS, consisting of three maximum beam energies, each with three different pencil beam compositions. The results of these simulations can be seen in Figure 4.6. All setups outlined in Table 4.1 produced reasonably flat SOBPs on the central axis with less than a 5% difference between dose extrema,

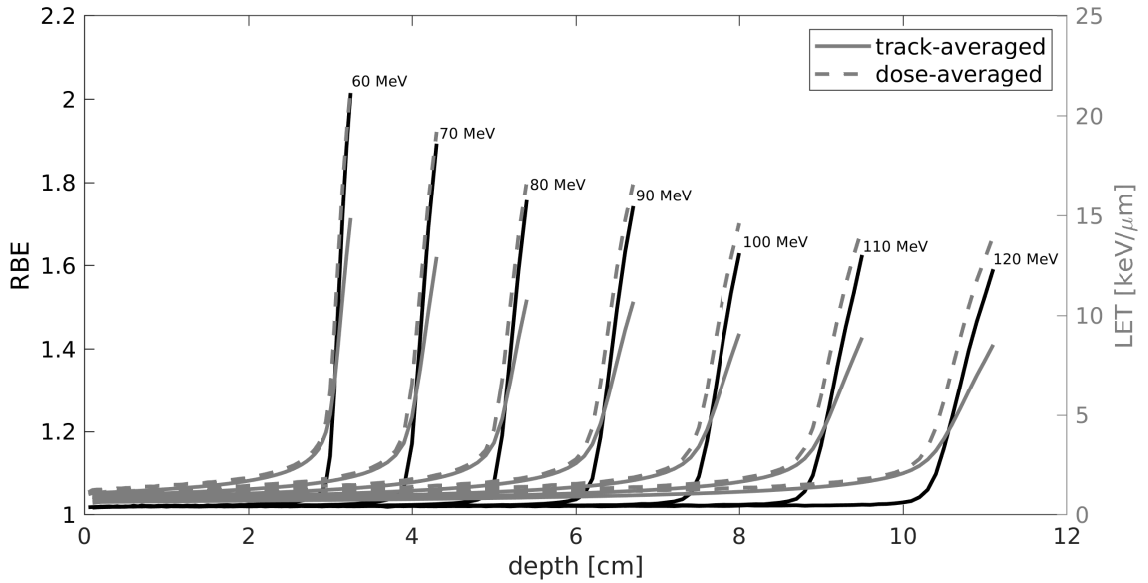


Figure 4.4: RBE_{cd} , calculated using (3.35), shown along with the dose-averaged and track-averaged LET, as scored in TOPAS along the central axis, for pristine proton pencil beams between 60 MeV and 120 MeV in water.

Table 4.2: Parameter fits to (4.11) using 550 data points for both dose-averaged LET (LET_d) and track-averaged LET (LET_t). 95% confidence intervals are shown in brackets. The coefficient of determination for this fit was $R^2 = 0.999$ against LET_d and $R^2 = 0.984$ against LET_t .

	$L_{\Delta} = LET_d$	$L_{\Delta} = LET_t$
p_1 [(keV/ μm) $^{-1}$]	0.06480 (0.06388, 0.06573)	0.09913 (0.09553, 0.1027)
p_2 [unitless]	0.8606 (0.8548, 0.8664)	0.7881 (0.7748, 0.8013)
p_3 [keV/ μm]	4.844 (4.467, 5.222)	2.097 (1.627, 2.568)
q_1 [keV/ μm]	4.653 (4.278, 5.027)	1.934 (1.470, 2.397)

as shown in Figure 4.6(a, b, c). All SOBPs cover roughly 25% of the depth-dose profile. The RBE_{cd} is shown to stay at 1.0 throughout the plateau region before stepping up to around 1.1 in the SOBP region. Toward the distal edge of the SOBP the RBE_{cd} rises sharply. Figure 4.6(d, e, f) shows the statistical variation in the RBE_{cd} measured through the middle 50% of each SOBP. The mean RBE_{cd} is highest for the 60 MeV SOBP, at 1.16, and decreases with increasing beam energy (1.10 for 100 MeV and 1.06 for 180 MeV). The spread in RBE_{cd} through the SOBP decreases only very slightly as the number, n , of pencil beams that make up the SOBP are increased,

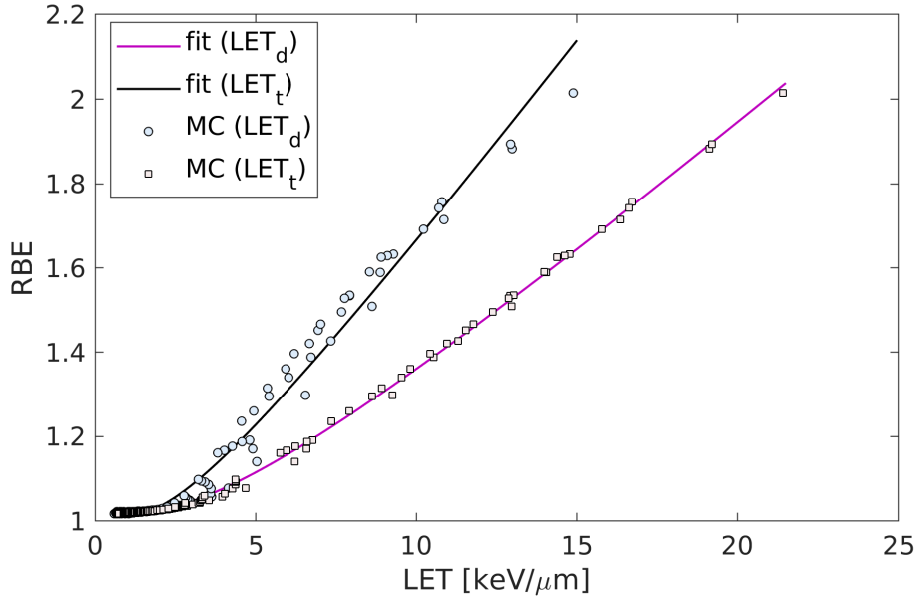


Figure 4.5: Relationship between RBE_{cd} , calculated using (3.35), and LET from Monte Carlo (MC) simulations. Also shown is the fit of (4.11) to the 550 data points with both dose-averaged and track-averaged LET. See Table 4.2 for the fitting parameters.

however there is also a very small decrease in the RBE_{cd} as n is increased. This could be due to an increased number of lower LET protons from the other pencil beams dominating the spectrum and thus reducing the average damage contribution. The 68% and 95% confidence intervals in the RBE_{cd} , over all values of n , are shown in Figure 4.6(d, e, f). The widths of these confidence intervals are shown to decrease as the SOBP energy is increased. This is most likely due to two factors; (1) as the SOBP energy is increased, so is the absolute width of the SOBP and as a consequence the build up in both dose and LET is more gradual; (2) a larger number of RBE_{cd} measurements were taken in wider SOBPs, which naturally leads to smaller statistical variation.

4.3.2 RBE calculation in a treatment planning system

The RBE_{cd} map produced for a head-and-neck plan distributed with the matRad treatment planning software is shown alongside the dose map for the isocentre slice

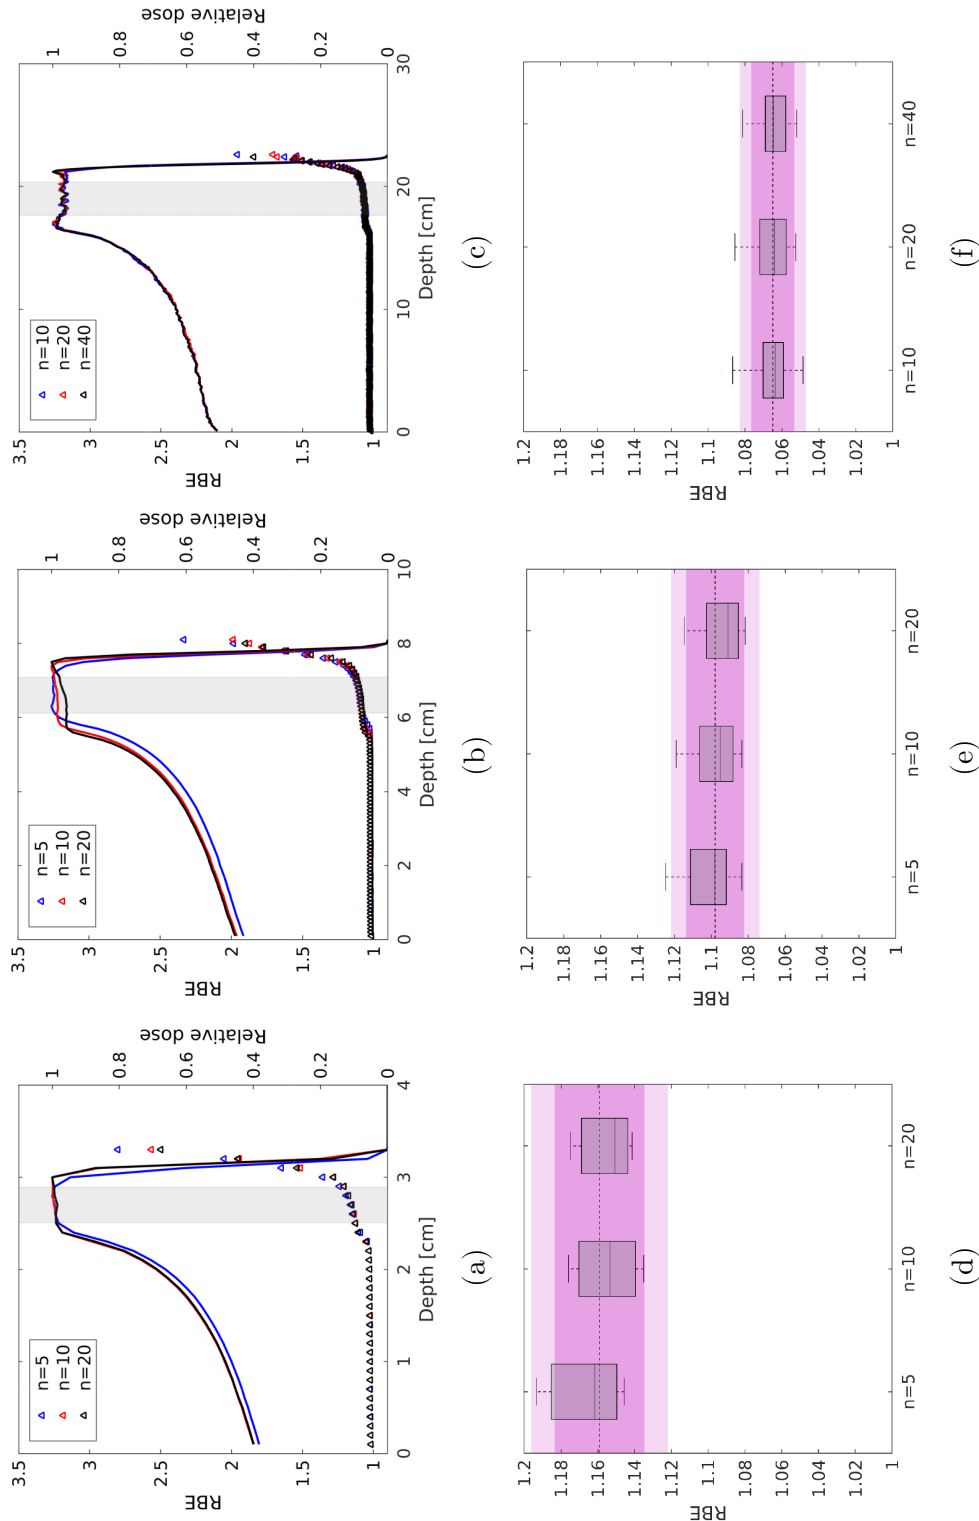


Figure 4.6: (a, b, c): RBE_{cd} , calculated using (3.35), is shown at 1 mm intervals along the central axis of three SOBPs at 60 MeV, 100 MeV and 180 MeV maximum energy, respectively, for $n = 5, 10$ and 20 energy intervals forming the SOBPs according to (4.5). (d, e, f): Variation in the RBE_{cd} calculation in the middle 50% of the SOBPs for 60 MeV, 100 MeV and 180 MeV, respectively. The darker pink shaded region represents the 68% confidence interval in the mean RBE_{cd} while the lighter pink region represents the 95% confidence interval.

in Figure 4.7. Two planning target regions were delivered uniform doses of approximately 2.0 Gy and 2.5 Gy per fraction. RBE_{cd} is approximately 1.0 in the medial-posterior low-dose region, however there is swathe of elevated RBE_{cd} in the target region, between around 1.06 to 1.10, and a further increase up to 1.2 at the edges of the dose fields. A histogram of the RBE_{cd} values in the isocentre slice voxels is displayed in Figure 4.8. It is clear in both the visualised RBE_{cd} map and the histogram that the proton dose is predicted to be between about 6% and 10% more effective across the middle of the target regions. This result is expected, given the SOBPs simulation results in Section 4.3.1; many proton pencil beams were modulated to create SOBPs and thus produce the homogeneous dose fields in the treatment plan. Of the patient voxels, 41% gave an RBE_{cd} of 1.05 or greater, with 15% greater than 1.10 and 3% greater than 1.5.

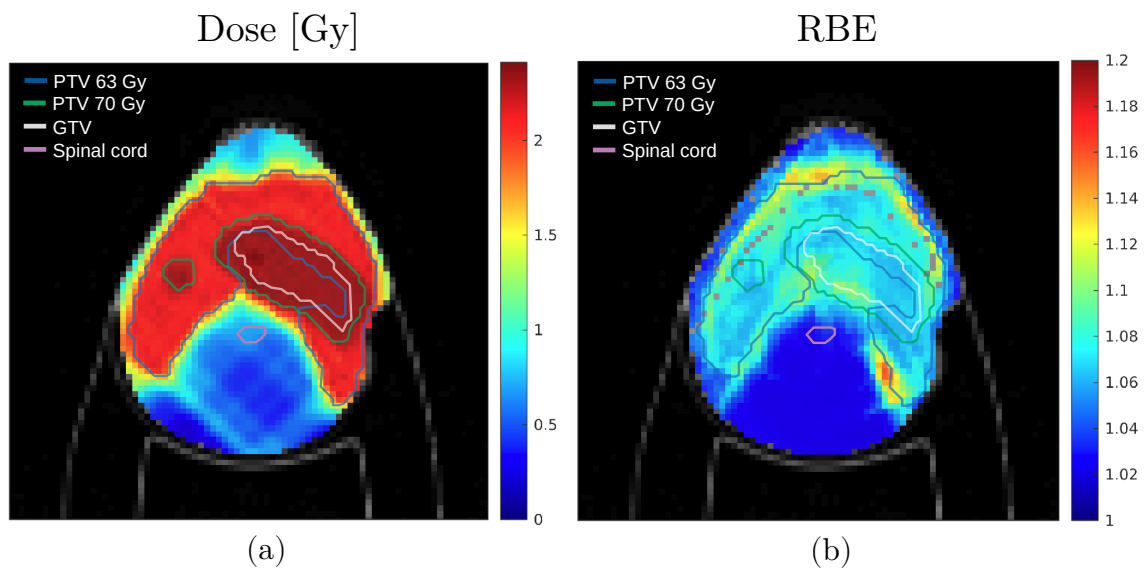


Figure 4.7: (a) dose map, at isocentre, for a head-and-neck plan distributed with the *matrad* treatment planning software [163]. (b) RBE_{cd} map corresponding to the dose map in (a), generated by modifying the LET calculation subroutine. The gross and planning target volumes, GTV and PTVs respectively, and spinal cord contours are shown for reference.




Figure 4.8: Histogram showing the RBE_{cd} value distribution in the pixels of the dose map slice shown in Figure 4.7.

4.3.3 Comparisons with *in vitro* experimental data

The error in the RBE prediction using the phenomenological models of Wilkens and Oelfke [185], Carabe *et al.* [186], Chen and Ahamd [188], Wedenberg *et al.* [189], McNamara *et al.* [190] and Mairani *et al.* [191], was assessed for surviving fraction endpoints between 5% and 95% survival. Figure 4.9 shows the percentage difference between the predictions and the RBE_{SF} calculated using the empirical LQ model data from the PIDE database. Each point represents the prediction error for an individual experiment while the red line indicates the average error across all experiments. Blue shaded regions (from darkest to lightest) give envelopes of one, two and three standard deviations from the mean. It can be seen that the two most recent models by McNamara *et al.* [190] and Mairani *et al.* [191] seem to be the most accurate in terms of mean error and spread. However, while the mean error is below 5% for both models, errors in most individual experiments are much larger and generally increase with surviving fraction.

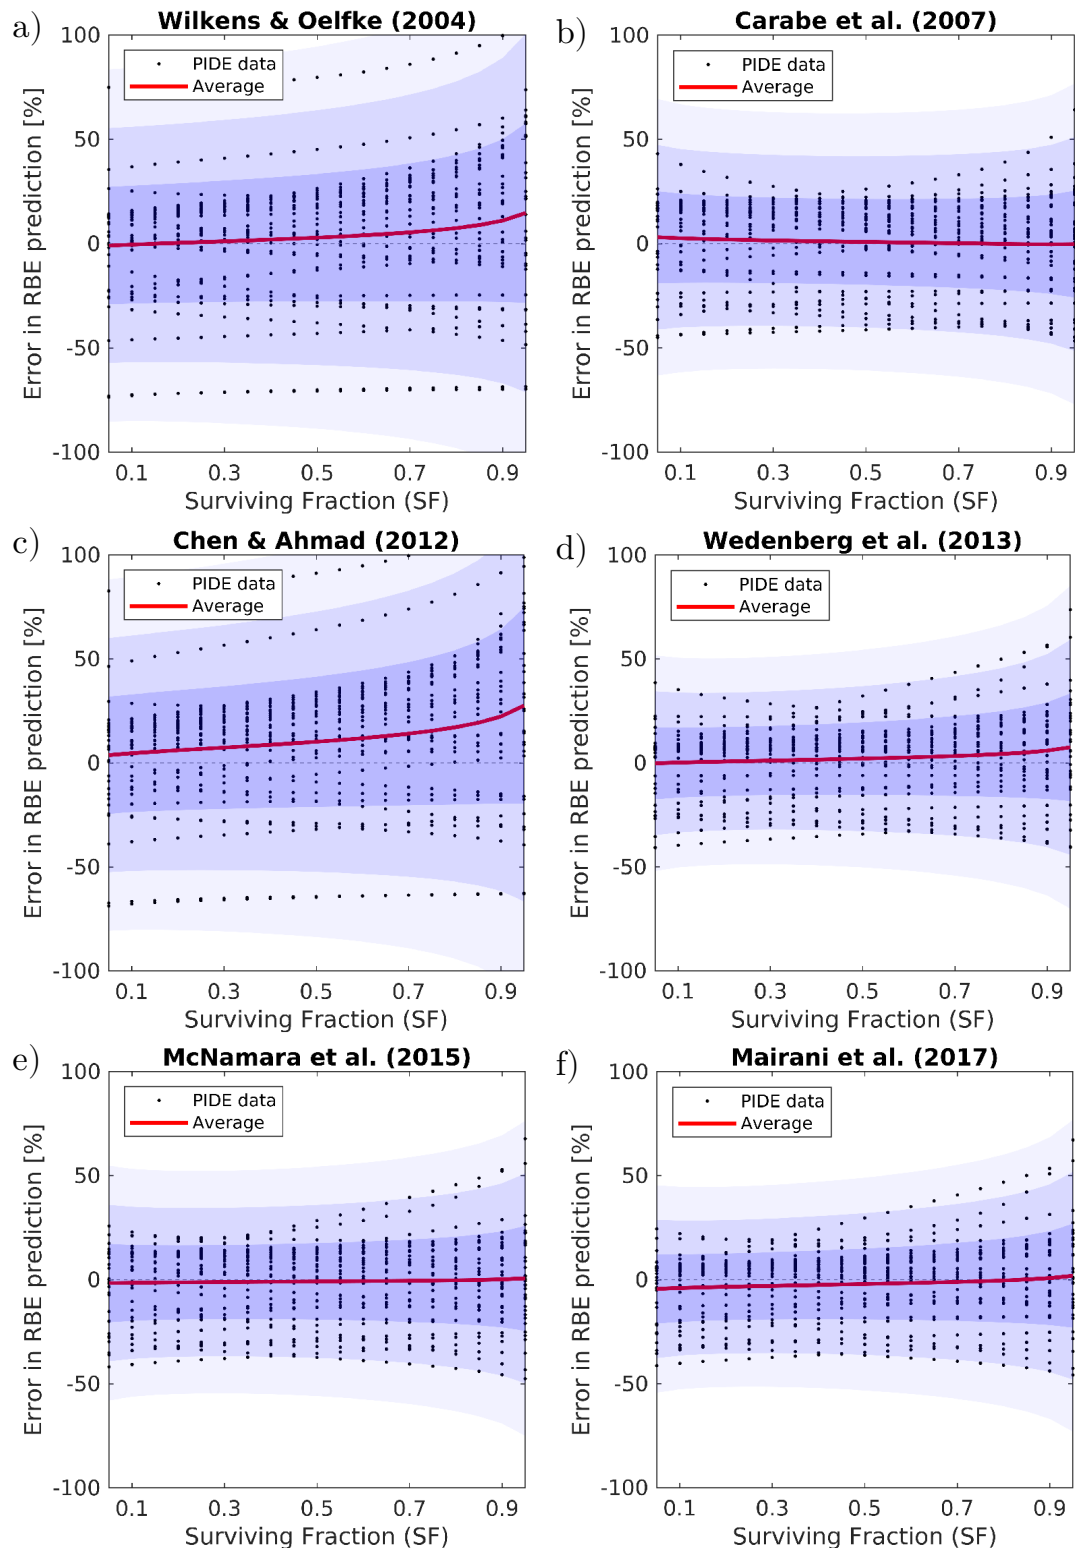


Figure 4.9: Percentage difference between predicted RBE and the empirical RBE for surviving fractions between 5% and 95%. Predicted RBE is according to the six phenomenological models summarised in Table 3.1. Each point represents the RBE prediction error in a single experiment from the PIDE database [164], while the red line indicates the average error across all experiments. Blue shaded regions (from darkest to lightest) give envelopes of one, two and three standard deviations from the mean.

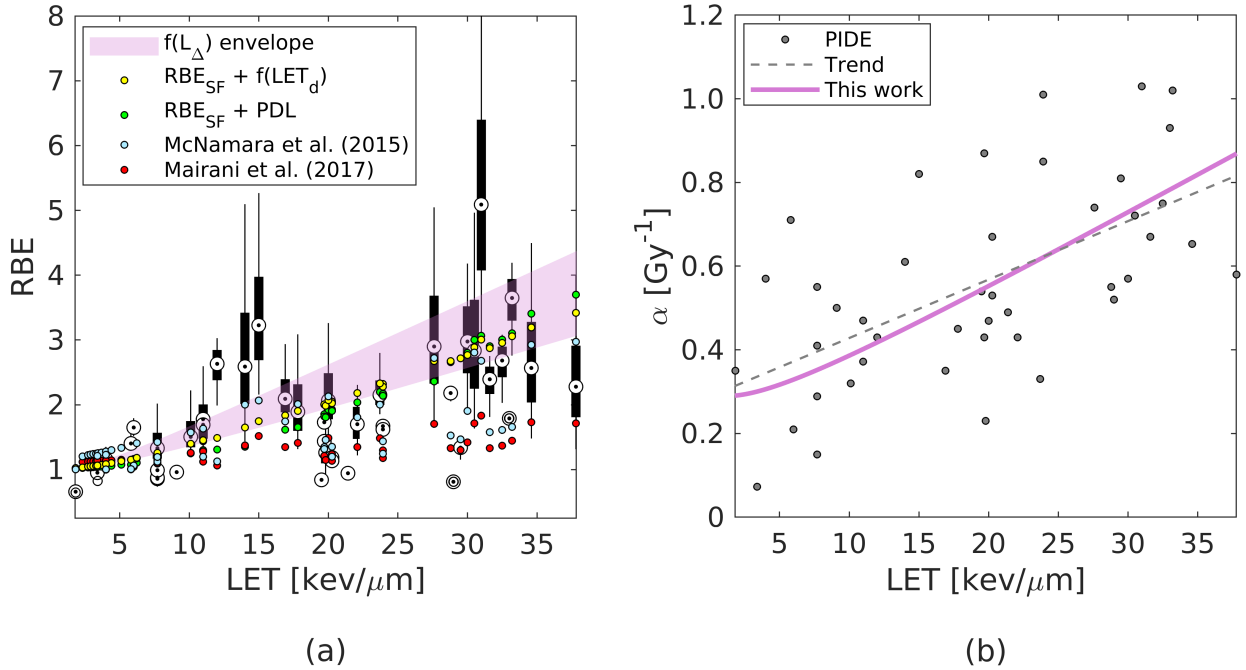


Figure 4.10: (a) Comparison of RBE predictions to experimental data provided by the PIDE database [164]. Experimental RBE_{SF} across surviving fractions of 5% to 95% are shown as box plots (black). Predictions from two models in this work, RBE_{SF}+ $f(\text{LET}_d)$ (yellow) and RBE_{SF}+PDL (green), given by (4.6)-(4.7) and (4.8)-(4.9) respectively, are shown along with the models of McNamara *et al.* [190] and Mairani *et al.* [191]. RBE_{cd} predictions from the $f(L_\Delta)$ envelope in (4.11) are overlaid. (b) An increasing trend of α values with LET is shown against the prediction from this work.

Figure 4.10(a) shows the empirical RBE_{SF} across all surviving fractions as standard box plots. Predictions from two models in this work, RBE_{SF}+ $f(\text{LET}_d)$ and RBE_{SF}+PDL, given by (4.6)-(4.7) and (4.8)-(4.9) respectively, are shown along with those of McNamara *et al.* [190] and Mairani *et al.* [191]. The empirical RBE-LET relationship, $f(L_\Delta)$, for LET values between those obtained via dose-averaging and track-averaging, is overlaid. There is both large intraexperimental variation in the RBE_{SF}, due to the surviving fraction endpoint choices, and interexperimental variation between results for protons at similar LET. The latter could be due to a multitude of factors such as those described in Chapter 3, Section 3.4, including differences in cell type, cell preparation, plating efficiencies and LET measurement technique. A decrease in RBE is seen in two studies beyond approximately 33 keV/μm, however

for therapeutic proton beams we shall restrict our study to LET values less than 30 keV/ μm . There are also some studies that show excessively large RBE for their corresponding LET, however it is beyond the scope of this thesis to explore the causes of such discrepancies in the PIDE database. Despite these uncertainties, a general increase in RBE is observed as the LET is increased. The predictions of this work describe this trend well and exhibit greater precision than both the McNamara and Mairani models. The predictions by the Mairani model are shown to underestimate the RBE, most noticeably for LET values above 20 keV/ μm . This underestimation is also evident in the McNamara model predictions, though to a lesser extent.

Linear LQ response parameter values, α , are shown in Figure 4.10(b) for each experiment from the PIDE database. Similar to the RBE discussion above, a large spread in the α values between experiments is noticeable, however they clearly increase as the LET is increased. This general trend is shown by a qualitative linear fit to the data. Using (4.8) and the simplification of the reference DSB yield in (4.10), the α parameter is directly proportional to the proton-induced DSB yield and to the RBE. In the limit as LET approaches zero, α can be expected to approach α_x , the linear response for the reference radiation, giving a maximum RBE of unity. The mean value of α_x for the experiments shown was 0.304 Gy⁻¹. Therefore, $f(\text{LET}_d)$ was scaled to achieve this value in the low-LET limit. The result, shown in Figure 4.10(b), describes the empirical trend well.

4.4 Discussion

Following the development of a biophysical RBE model based on DSB induction in nuclear DNA, presented in Chapter 3, a series of Monte Carlo simulations were carried out in order to evaluate the implementation of the model and its predictions. The model, referred to here as the Poisson Distributed Lesions (PDL) model, was used

to score RBE_{cd} in a voxelised water phantom for both monoenergetic proton pencil beams and SOBPs of varying nominal energies. An empirical relationship between RBE_{cd} and LET was derived from the Monte Carlo results, which was used to modify an LET calculation subroutine in the matRad open source treatment planning system [163] such that voxelised RBE_{cd} distributions could be calculated without the need for Monte Carlo methods. This modified functionality was tested on a head-and-neck plan optimised according to physical dose constraints. Predictions of the PDL model at varying proton energy (or LET) were compared to experimental results compiled in the PIDE database [164] and to competing predictions by six published RBE models.

Monte Carlo studies showed a tendency for RBE_{cd} at the Bragg peak and distal edge to be larger for lower energy beams than for higher energy beams. Maximum RBE_{cd} was estimated at 2.0 for a 60 MeV proton pencil beam but was 1.6 for double the beam energy, at 120 MeV. The trend closely followed that of the LET, particularly dose-averaged LET, however the increase in RBE_{cd} proximal to the Bragg peak was sharper than the more gradual increase in LET. The tendency for RBE_{cd} to decrease with increasing pencil beam energy could be considered in IMPT planning when selecting the direction and energy of treatment fields. In this way, uncertainty in the biological effect at the proton end-of-range could be partially mitigated. For example, according to the simulation results in Section 4.3.1, increasing pencil beam energy from 60 MeV to 90 MeV corresponds to a reduction in RBE_{cd} by up to 13%, while the proton range is extended by approximately 33 mm. This range extension can be accommodated by altering the location and angle of incidence of the beam such that it initially passes through more tissue. However any advantage in reducing the RBE_{cd} must be weighed against the implications of depositing more dose in healthy tissues, and particularly in critical structures where radiation could have otherwise been avoided. Passive scatter of high energy beams prior to entering the patient could also theoretically reduce the biological end-of-range effects however the

increased scatter can reduce sharpness of the lateral dose penumbra and would require sophisticated dose modelling such as Monte Carlo methods to be used in the inverse planning procedure.

Average RBE_{cd} values of approximately 1.1 were calculated in the middle of simulated SOBPs, ranging from 1.07 for a 180 MeV beam to 1.16 for a 60 MeV SOBP. The mean RBE_{cd} for a 100 MeV SOBP was exactly 1.10. This finding falls in line with much of the experimental evidence in the existing literature [see, e.g., 113], on which the constant RBE recommendation of the ICRU in Report 78 [112] was based. It is interesting to note that while the RBE_{cd} is elevated to around 1.1 in the SOBP, there is negligible excess biological effect predicted in the entrance and plateau regions leading up to it. In all simulations, RBE_{cd} was approximately 1.0 for depths up to about 90% of the expected range of the lowest energy primary proton in the beam. This could allow for dose escalation, as the target region would receive an enhanced biological effective dose while proximal healthy tissue would not. Further evidence is needed from independent studies to support or refute this claim.

Further evidence of elevated RBE_{cd} throughout target regions was demonstrated through analysis of the head-and-neck treatment plan using matRad. Predicted RBE_{cd} in these regions was consistently between 1.06 and 1.10, which further increases up to 1.20 at the target boundaries. Conversely, the RBE_{cd} remained around 1.0 in low-dose voxels that were not close to proton Bragg peak locations. Voxelised RBE_{cd} was able to be calculated in matRad by modifying an existing LET calculation subroutine to instead use the empirical RBE_{cd} -LET model, $f(L_{\Delta})$, in (4.11). While this is a pragmatic method for treatment planning systems that already implement an optimised LET calculation technique, micro-scale track structure and knowledge of the stopping power are not actually required for the PDL model. Only the local energy spectrum within a voxel is needed. This is already obtained either analytically or via a Monte Carlo method for the dose calculation. Thus, the PDL should be able

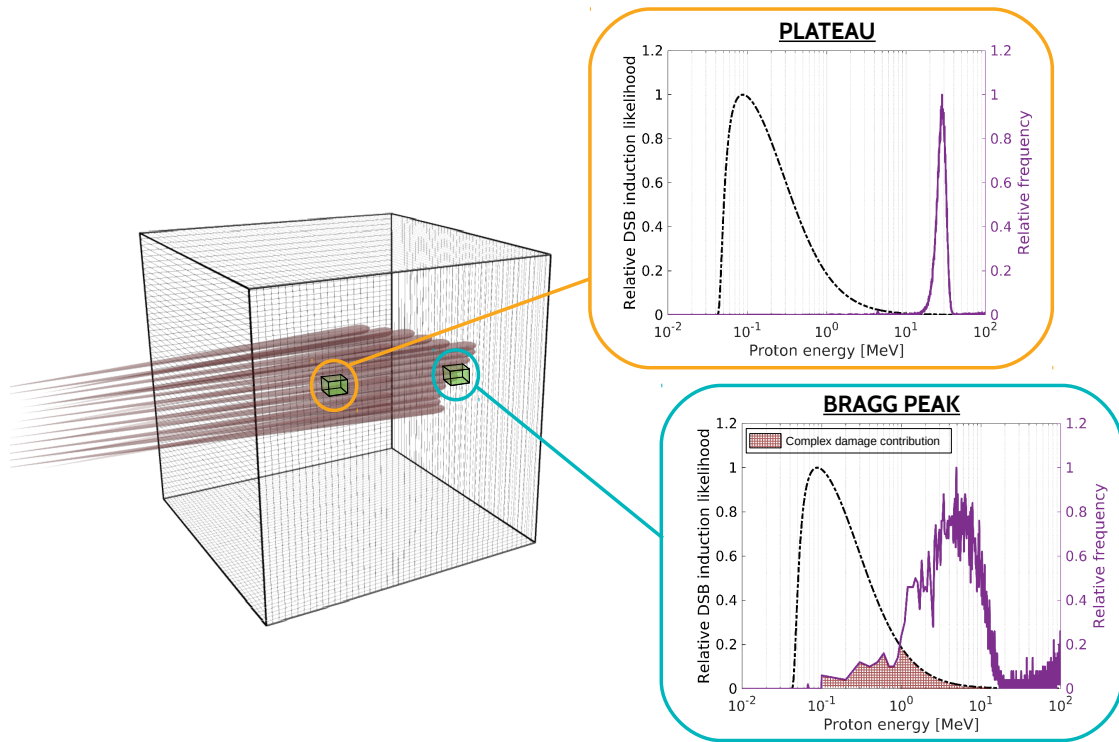


Figure 4.11: Schematic of proton pencil beams impinging on a cube phantom. Examples of local proton energy spectra in the plateau region and at the Bragg peak are shown against the relative DSB induction likelihood. A hashed region in the lower left plot shows the component of the energy spectrum that contributes to complex DNA damage.

to be easily implemented in any modern treatment planning system. The model is currently being implemented³ in a research version of *RayStation*⁴ (RaySearch Laboratories, 2020) to be tested on a range of archived clinical cases and compared to other RBE models. Figure 4.11 illustrates how voxelised energy spectra can describe the associated RBE_{cd} . The schematic shows proton pencil beams entering a phantom, with example energy spectra from Monte Carlo simulations shown for two voxels at different depths. In the plateau region most protons exhibit high kinetic energies compared to those that could induce a complex DNA damage response, but in the Bragg peak the spectrum broadens and shifts to lower energies. A larger proportion of protons therefore lead to complex damage and the RBE_{cd} increases. One advan-

³With thanks to Erik Traneus for helpful discussions and technical assistance.

⁴Product information available at: <https://www.raysearchlabs.com/raystation/>

tage of the PDL model, as discussed in Chapter 3, is that the predicted DSB yield, and therefore the RBE_{cd} , scales linearly with dose; this is not necessarily the case for other endpoints, such as surviving fraction, which show large variability between cell types. This enables the RBE_{cd} -weighted dose due to overlapping pencil beams to simply add. By contrast, first attempts to extend TOPAS for radiobiological scoring used a number of models, including those listed in Table 3.1, and it was found that non-linearity in the often-used LQ parameters α and β required post-processing for accurate calculation [326].

It was shown that the electron contribution to DSB yield due to a 6 MV reference x-ray beam could be assumed constant. While the same is not true for electron contributions to damage in a proton beam, the proton contributions outweigh the dose-averaged yield by around 30 times. These two considerations enabled a simplified version of the PDL model, (4.10), to be used in practice. In examining the kinetic energy spectra of protons from the MC simulations, an observation was made that secondary protons offered only a very small contribution to the shape of the spectrum and that the RBE_{cd} did not noticeably change when secondary protons were neglected from the calculation. However, a thorough quantitative analysis on this subject should be undertaken before further simplifying the PDL model. A dose-weighted average of the DSB yield was used in (3.36) and (4.2) to calculate the RBE_{cd} . An alternate fluence-weighted technique was investigated for simpler implementation. Under this technique and the exclusion of the electron energy spectrum (3.36) becomes

$$\langle Y_{\text{cd}} \rangle = \frac{\int_0^{E_{\text{max}}} \Phi(E) Y_{\text{cd}}^{p+}(\lambda(E), \ell_0) dE}{\int_0^{E_{\text{max}}} \Phi(E) dE} \quad (4.12)$$

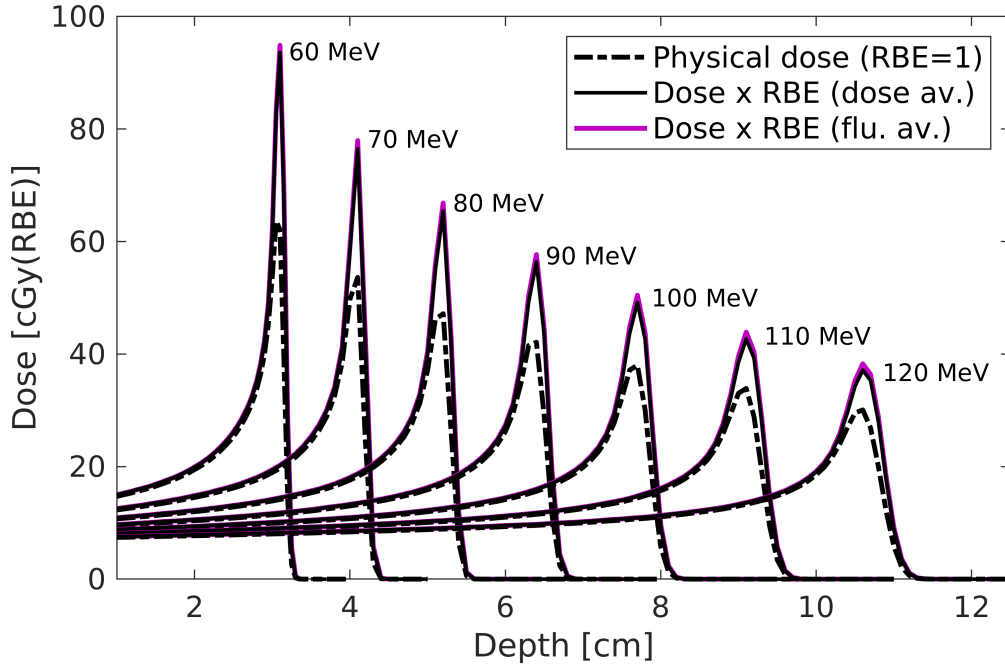


Figure 4.12: RBE_{cd} -weighted dose scored along the central axis in TOPAS [161] for proton pencil beams between 60 MeV and 120 MeV. Two RBE_{cd} scoring techniques are shown, dose-weighted and fluence-weighted.

and the corresponding Monte Carlo calculation in (4.2) becomes

$$\langle Y_{cd} \rangle = \frac{\sum_{k=1}^{N_{p^+}} (Y_{cd}^{p^+})_k w_k^{p^+}}{\sum_{k=1}^{N_{p^+}} w_k^{p^+}}. \quad (4.13)$$

Figure 4.12 shows the RBE_{cd} -weighted dose for the pencil beam simulations previously used in this work, implementing both the dose-averaged and the fluence-averaged techniques. The fluence-averaged technique performed similarly up to the Bragg peak, at which point it overestimated the maximum RBE_{cd} by 2.3% on average with a standard deviation of 0.6%. Given the intrinsic uncertainty in the biological effect itself, it may be reasonable to account for this small overestimation and use the fluence-averaged technique in practice due to its favourable implementation; it does not require a dose calculation or knowledge of the stopping power.

It would be of interest to compare the electron-induced damage yield for other clinical reference beams, depending on which reference modality is most appropriate. It is widely accepted that there is a difference in biological effectiveness between orthovoltage and megavoltage x-rays [327]. Further, photon energies above the binding energy of encountered nuclei (around 8 MeV) may result in neutron contamination of the radiation field. This is also a possible source of increased RBE. The reference for the RBE_{cd} calculation should therefore be specifically chosen based on circumstance rather than a single value being used universally.

The empirical RBE_{cd} -LET relationship in this work was derived from the results of simulations in water as a first approximation. Material heterogeneity in a patient leads to variations in local energy spectra and dose deposition which may have an affect on the calculated RBE. Grzanka *et al.* [59] investigated the influence of material composition on LET, a similar microdosimetric quantity, finding less than a 1% difference between values in water and brain tissue. Materials such as bone, which are expected to have a greater influence on the spectrum, were not included in the study. Tissue heterogeneity also presents biological variations due to the radiosensitivities, or α/β ratios, of different tissues, which should be accounted for through the addition of chemical repair modelling to complement initial DSB induction predictions. There are a number of recent publications in DSB repair modelling following proton irradiation which could be utilised [328, 329, 320, 330, 331, 332]. The combination of both physical and chemical models of cellular damage may be tested using more recently developed nanoscale Monte Carlo software. For example, the *Geant4-DNA* and *TOPAS-nBio* projects have been used to investigate track structure on sub-cellular targets, such as DNA and mitochondria, as well as water radiolysis and DNA repair kinetics [333, 334, 245, 335, 246, 336, 337, 338, 339, 340, 341, 342].

The predictions of the PDL model were compared to 161 data points from 30 independent *in vitro* proton irradiation experiments. Two additional models, based

on the PDL formalism, were introduced to relate the DSB-induction RBE to cell survival RBE through the surviving fraction endpoint, namely, the $\text{RBE}_{\text{SF}} + f(\text{LET}_d)$ and $\text{RBE}_{\text{SF}} + \text{PDL}$ models. Both were shown to well describe the empirical relationship between RBE and LET, while the two most accurate phenomenological RBE models for this data, the McNamara [190] and Mairani [191] models, underestimated the RBE for LET values larger than $20 \text{ keV}/\mu\text{m}$. The Mairani model consistently underestimated the RBE for the entire LET range. The PDL model also predicted an upward trend in the linear LQ response term α that agreed well with the trend in the empirical data, although the spread in the data was large. Care must be taken in extrapolating the trend through higher LET values due to saturation effects, discussed in Chapter 3, however for clinical proton beams the LET is limited to about $30 \text{ keV}/\mu\text{m}$. Using the PDL model directly with only one energy value as the input, instead of using the full local spectrum, proved to be surprisingly precise when compared to the use of the empirical RBE-LET formula, $f(\text{LET}_d)$. This may enable the model to be used more widely, particularly in *in vitro* experiments, where either direct measurement or modelling of the local energy spectrum on the cellular target is lacking.

The repair-misrepair-fixation (RMF) model [201, 202, 203] was discussed in Chapter 3, Section 3.2. The assumptions of the RMF differ to those of the $\text{RBE}_{\text{SF}} + \text{PDL}$ used in this work in that the maximum RBE (low dose RBE_{SF}) is greater than the RBE by DSB induction and minimum RBE (high dose RBE_{SF}) is exactly equal to it. The RMF assumes an auxiliary RBE_{DSB} model that accounts for inter-track chromosomal aberrations whereas the PDL is purely a single-particle, intra-track damage model. Nevertheless, the use of the PDL as an auxiliary model to the RMF was briefly explored. The RMF was used according to (3.26) and (3.24). Figure 4.13(a) shows the predictions of this model against the two other RBE_{SF} models proposed in this work. Only a small discrepancy between the three models is apparent. Note that

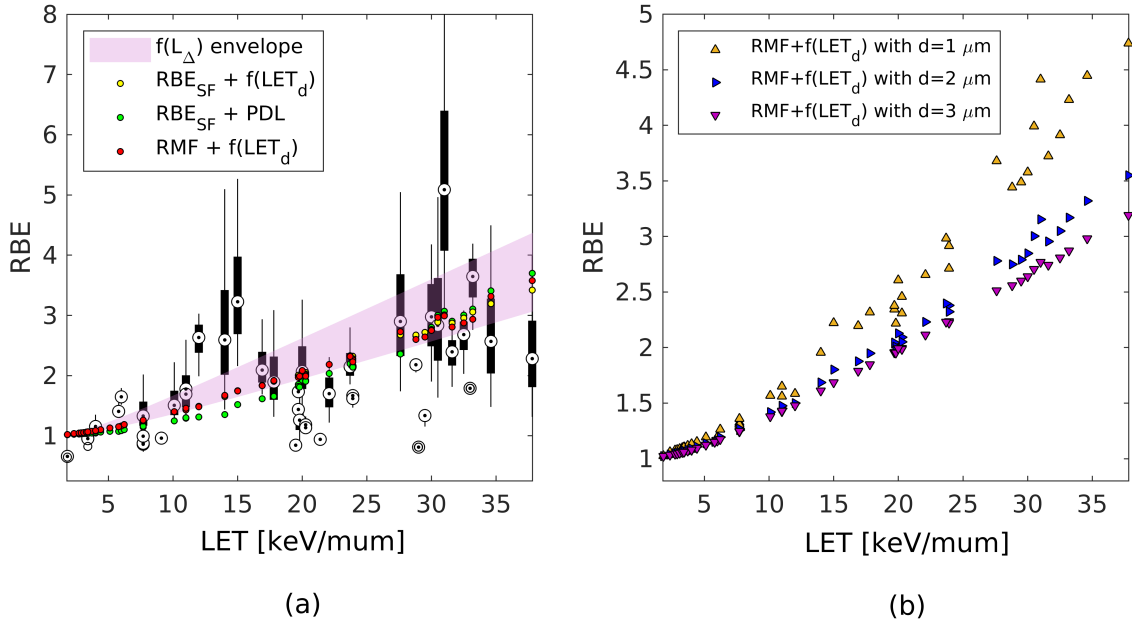


Figure 4.13: Comparison of RBE predictions to experimental data provided by the PIDE database [164]. Experimental RBE_{SF} across surviving fractions of 5% to 95% are shown as box plots (black). Predictions from two models in this work, $RBE_{SF} + f(LET_d)$ (yellow) and $RBE_{SF} + PDL$ (green), given by (4.6)-(4.7) and (4.8)-(4.9) respectively, are shown along with the RMF model [201, 202, 203] as in (3.26) and (3.24) with RBE_{DSB} calculated using the PDL model. RBE_{cd} predictions from the $f(L_{\Delta})$ envelope in (4.11) are overlaid. (b) Affect of changing the diameter of the cell nucleus or sub-cellular domain on the $RMF + f(LET_d)$ prediction.

RBE_{max} in the RMF model has been simplified assuming that the specific energy is calculated for a cell target of radius $4 \mu\text{m}$, as suggested by Stewart *et al.* [205]. The microdosimetric-kinetic (MK) model [192, 193, 194, 195] similarly uses the specific energy to relate the local energy spectrum inside sub-cellular domains to chromosomal aberrations caused by the combination of two potentially lethal lesions. However, the targets for binary misrepair in the MK model have a diameter of approximately 0.5 to $0.9 \mu\text{m}$ [205]. Changes in the diameter of the cell or the sub-cellular target can have a significant affect on the RBE prediction with the RMF model, as shown in Figure 4.13(b).

Possible sources of variation in the experimental data provided in the PIDE database [164] was discussed in Section 4.2.3. Any large discrepancies in the data for which a cause was not immediately obvious were addressed by referring to the original papers detailing the experiment. For instance, data by Wouters *et al.* [343]

seemed to overestimate RBE at low LET, producing values of around 1.4 for LET less than 5 keV/ μm . The authors, in the original paper, explained that their use of asynchronous V79 cells for the ^{60}Co reference experiment required a two-population LQ model. This method, described in [344], calculated the surviving fraction using $S = \Omega e^{-\alpha_s \mathcal{D} - \beta_s \mathcal{D}^2} + (1 - \Omega) e^{-\alpha_r \mathcal{D} - \beta_r \mathcal{D}^2}$ where Ω and $1 - \Omega$ are the fractions of sensitive and resistant cells whose response parameters are α_s, β_s and α_r, β_r respectively. By using this model in place of the standard LQ formalism, the calculated RBE was reduced from 1.4 to between 1.0 and 1.1 at low LET. This is just one example of how knowledge of the individual experiments is sometimes required in addition to simply having access to parameters collected *en masse*. Future comparisons with the PIDE database, or any experimental data for that matter, should consider such discrepancies between experiments and justify their inclusion in the study.

Chapter 5

Clinical beam line experiment at the Clatterbridge Cancer Centre

5.1 Introduction

In Chapter 4 we investigated the predictions of induced initial DSBs and RBE_{cd} in pristine proton pencil beams and homogeneous fields (SOBPs) using the model described in Chapter 3. Monte Carlo simulations were carried out using ideal beam sources, for which the initial position, momentum direction and energy of the source particles were controlled via parameters set in TOPAS. Each generated particle, called a *history*, would have an associated position, momentum direction and energy sampled from a distribution, which could be described as either uniform, for which any value in the allowed interval is equally probable, or Gaussian, for which the likelihood of a value being chosen is related to user-defined mean, variance and cutoff extrema. While these parameters allow for more realistic, rather than superficial, beam modelling, it is unable to closely emulate real-world beam conditions observed at the nozzle of a treatment machine. For this reason, experiments were undertaken at a clinical proton beam line¹ to measure the energy spectra of protons after passing through varying absorber thicknesses, using a semiconductor detector. Measured spectra were used

¹This work was carried out in collaboration with the Cockcroft Institute at the University of Liverpool, being led by Jacinta Yap, and Advacam (Prague, Czech Republic).

to calculate approximate LET and RBE_{cd} depth profiles, which were compared to results from Monte Carlo simulations using an advanced Geant4 [162] model of the beam line [345, 346] which has since been extended in TOPAS [161, 165].

The Clatterbridge Cancer Centre (CCC) in Wirral, United Kingdom is home to the world's first hospital-based proton beam therapy (PBT) facility, having treated patients with ocular melanoma for over 30 years². The PBT beam line produces 60 MeV protons at the treatment isocentre, with a maximum penetration depth of approximately 31 mm in water, a sub-millimetre distal fall-off and penumbra of 1.5 mm [348]. The source of protons is a Scanditronix MC-60 PF isochronous cyclotron, initially commissioned in 1984 for neutron therapy. Scattering foils placed in the vacuum tube of the beam line help shape and homogenise the proton beam while an upstream modulator and range shifter control the lateral penumbra. Further technical details of the beam line, and of clinical PBT operations at CCC, can be found in [349].

Previous radiobiological studies at the CCC PBT beam line [176, 350] have investigated both LET and RBE in cell cultures along pristine and spread-out Bragg peaks, including analyses of double strand breaks (DSBs) through immunofluorescence [176] and comet assay [350] techniques, which were discussed in Chapter 3. Applying the RBE_{cd} model presented in Chapter 3 to empirical data collected at CCC as well as simulated data is therefore a natural extension to the Monte Carlo studies in Chapter




Figure 5.1: CAD model of the CCC treatment line. Courtesy of J. Yap. Reproduced from [347] with permission.

²The first patients at the Clatterbridge Cancer Centre were treated in June of 1989.




Figure 5.2: Schematic of the TOPAS model of the CCC treatment line [165]. Courtesy of J. Yap. Reproduced from [347] with permission.

4, as predicted LET and RBE values can be compared to published data from the same site in order to validate the model.

5.2 Methods and materials

5.2.1 The Advacam MiniPIX-Timepix detector

The Advacam MiniPIX-Timepix detector is a portable, USB-operated silicon semiconductor detector with an application-specific integrated circuit (ASIC) capable of tracking individual particles. It has been described as a “*miniaturised, low-power and easily deployable particle micro-tracker*” [351]. The Timepix chip [352, 353], developed by the Medipix Collaboration, based at CERN, has been used for a wide variety of applications. These range from radiation dosimetry (see, e.g., [354]) to radiation detection on-board the International Space Station [355, 356, 357].

The detector consists of a $300\ \mu\text{m}$ thick silicon sensor comprised of an active area of $14 \times 14\ \text{mm}^2$, divided into a 256×256 pixel array. Each pixel has dimensions of $55 \times 55\ \mu\text{m}^2$ and offers an individual channel for readout up to 45 frames per second. In Time-Over-Threshold (TOT) mode, the Timepix ASIC is able to track individual particles, making it an ideal candidate for this study. The detector, shown in Figure 5.3, features a 1 mm thick aluminium sliding window which provides shielding to the sensor when necessary [351].

The detector was held in place throughout the experiment by a plastic housing




Figure 5.3: Schematic (a) and photograph (b) of the active sensor and Timepix ASIC chip of the MiniPIX hybrid semiconductor detector. The full detector (c) is around the size of a large USB data stick. Reproduced from [351, Fig. 1] with permission.

that was designed and 3D-printed by members of the Quasar Group at the University of Liverpool³. The plastic housing was secured to a remote-controlled rotating plate and the sensor was connected via USB cable to a laptop that could be accessed from the control room via remote desktop connection.




Figure 5.4: The Advacam MiniPIX-Timepix detector placed in a custom-made 3D-printed housing, which is attached to a remotely controlled rotating base.

³Quasar group (Liverpool): <https://www.liverpool.ac.uk/quasar/>

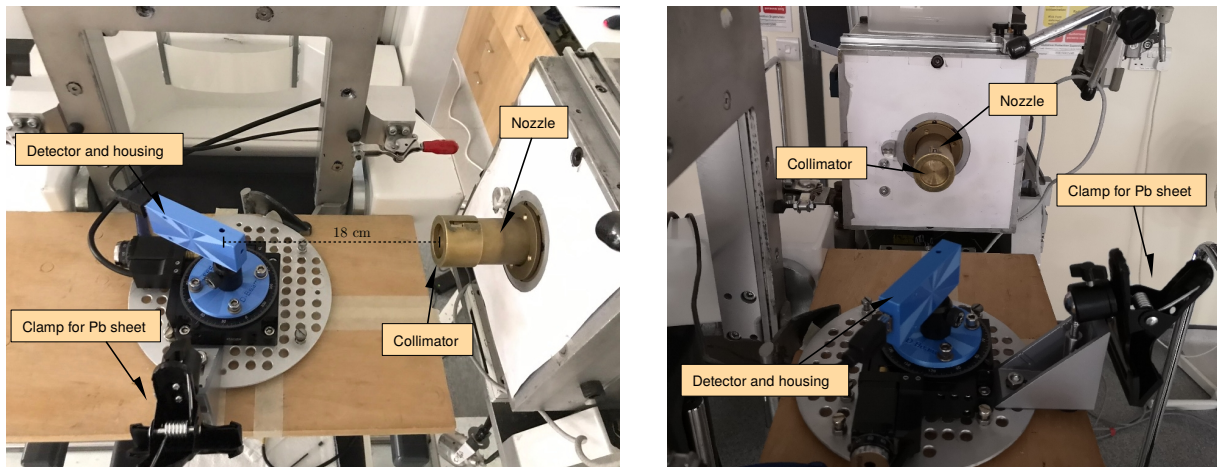


Figure 5.5: Setup of the Advacam MiniPIX-Timepix detector at the CCC treatment line. Here, the detector is placed at a 45 degree angle to the central axis of the beam.

5.2.2 Beam line measurements

The detector housing was secured to a bench top approximately 18 cm from the nozzle such that the centre of the sensor was in line with the central axis of the beam. This was achieved by adjusting the height of the treatment chair and checking alignment with a light source positioned behind a 1 mm diameter brass collimator at the nozzle. This collimator was kept in place for the entire experiment. It has been demonstrated that placing the MiniPIX detector at an angle ($> 20^\circ$) to the incident particles can improve particle tracking resolution [358]. The remotely-controlled rotating stand was therefore initially set to an angle of 45 degrees to the central axis.

The CCC beam line typically operates at a beam current in the order of 10^{10} protons per second. A current of this magnitude would over-saturate the MiniPIX detector and possibly cause damage to the electronics. The lowest possible beam current setting was therefore used, which was estimated between 10^3 and 10^5 protons per second for the experiments. Prior to the beam being turned on for the first time, several background readings were taken with the detector and, in order to protect the sensor from potentially high flux rates, a lead sheet with a 200 μm diameter pinhole aligned with the central axis was secured upstream of the detector using a vice and

the detector's 1 mm thick aluminium sliding window was closed.

A remote desktop client ensured that all data could be processed and visualised in real-time, using the Advacam *PIXet* software [359], from the safety of the control room. With the beam turned on, live visualisation of energy deposition allowed for the beam current to be varied to a suitable level (10^3 to 10^4 protons per second) and confirmed appropriate positioning of the sensor with respect to the central axis of the beam. Following a general rule of thumb [358], the beam current was set such that no more than 10% of the pixels in the sensor were populated with particle tracking information, as this helps optimise the resolving capabilities of the detector by avoiding over-saturation. A bias voltage of 30 V was applied with an acquisition rate of approximately 33 frames per second (FPS), which ensured that within each frame there was sufficient time for individual particles to induce a proportional charge response. Decreasing the frame rate would risk miscapture of the full particle track while increasing the frame rate would risk saturating pixels to the point that individual tracks could not be identified. A pragmatic approach was taken with the bias voltage, which affects the sensitivity of the charge deposition to energy conversion; setting this to 30 V appeared to achieve suitable sensitivity in registering single particle tracks. It should be noted, as well, that the sensor did not collect data at a constant rate. Open shutter time was approximately 10 ms followed by an additional 20 ms of dead time for the ASIC to process and record the data. A new "frame" was written to the cluster file every 30 ms. With the beam turned off, the lead sheet was subsequently removed and the experiment could proceed. However, the aluminium sliding cover remained in the closed position as the flux rate was still too high at the lowest possible beam settings. Placing increasingly thicker absorbers in the path of the beam would decrease the flux rate observed at the sensor, however the shielding was kept in place throughout the remainder of the experiment for consistency.

Six blocks of Lucite material, also known as Polymethyl methacrylate (PMMA),

Table 5.1: Details of each Clatterbridge experiment with the Advacam MiniPIX-Timepix detector.

Run ID	Detector angle ($^{\circ}$)	Lucite thickness (mm)	Aluminium shielding (mm)	WET (mm)	Acquisition time (s)	Number of clusters	Cluster rate (s^{-1})
BG1	N/A	N/A	N/A	N/A	131.25	63	59.26
BG2	N/A	N/A	N/A	N/A	95.26	48	56.60
BG3	N/A	N/A	N/A	N/A	43.56	21	55.46
BG4	N/A	N/A	N/A	N/A	254.58	130	52.24
R1	45	10.00	1.4	14.42	131.35	9440	693.9
R2	45	20.10	1.4	26.00	68.07	7277	909.3
R3	45	24.40	1.4	30.87	58.44	3445	581.1
R4	45	25.48	1.4	31.96	58.44	2270	460.3
R5	45	25.87	1.4	∞^*	58.44	569	294.3
R6	45	26.40	1.4	∞^*	58.44	179	203.9
R7	60	10.00	2.0	15.64	58.44	21610	1341
R8	60	24.40	2.0	32.00	58.44	5325	693.9
R9	60	25.48	2.0	∞^*	58.44	344	259.2

* A WET of ∞ indicates that the vast majority of protons have stopped within the material and therefore do not reach the detector.

varied in their thickness, were placed in the path of the beam to emulate radiation field measurements at different depths in a homogeneous phantom. These depths, listed in Table 5.1, covered the entry region, the Bragg peak and the distal dose fall-off. Each block of Lucite was manually placed approximately 6 cm upstream from the detector, perpendicular to the central axis. In order to obtain suitable particle counting statistics, acquisition times were varied to account for reduced proton transmission as the Lucite thickness was increased (refer to Table 5.1).

Following six runs at 45 degrees, the angle of the detector to the central axis was adjusted to 60 degrees and the experimental procedure was repeated for one block of Lucite corresponding to the entrance region (10.00 mm) and two blocks corresponding roughly to the Bragg peak (24.40 mm and 25.48 mm). Examples of visualised data in the PIXet software are displayed in Figure 5.6 (individual tracks) and Figure 5.7 (integrated frames and cluster height distribution). The water equivalent thickness (WET) through which the beam had passed once it reached the detector was calculated for each setup, accounting for the thickness and angle of the aluminium sliding

cover.

An iterative numerical method similar to that presented in [360] was carried out to determine the WET. Stopping power tables for water, PMMA (Lucite) and aluminium were extracted from the NIST pSTAR database [54] and interpolated to approximate the stopping power at energies between 0.5 MeV and 62 MeV for all three materials. The expected energy of protons exiting a material of thickness t_m cm, given that they entered the material with a mean energy of E_0 MeV, can be calculated using an Euler method,

$$E_{n+1} = E_n - \frac{S_m(E_n)t_m}{N + 1}, \quad (5.1)$$

where the material thickness has been subdivided $N + 1$ times, $n = 0, 1, \dots, N$, and S_m is the stopping power in the material of interest. The WET is then calculated by

$$\text{WET} = t_m \frac{\bar{S}_m}{S_m} \quad (5.2)$$




Figure 5.6: Visualisation in Advacam PIXet proprietary software [359] of individual clusters in a single frame acquired over 10 ms. Relative pixel intensities represent the amount of charge deposited. Run R7 (60° detector angle, 10 mm Lucite) from Table 5.1 is shown. Figure adapted from screenshot provided by C. Oancea (Advacam, Prague).

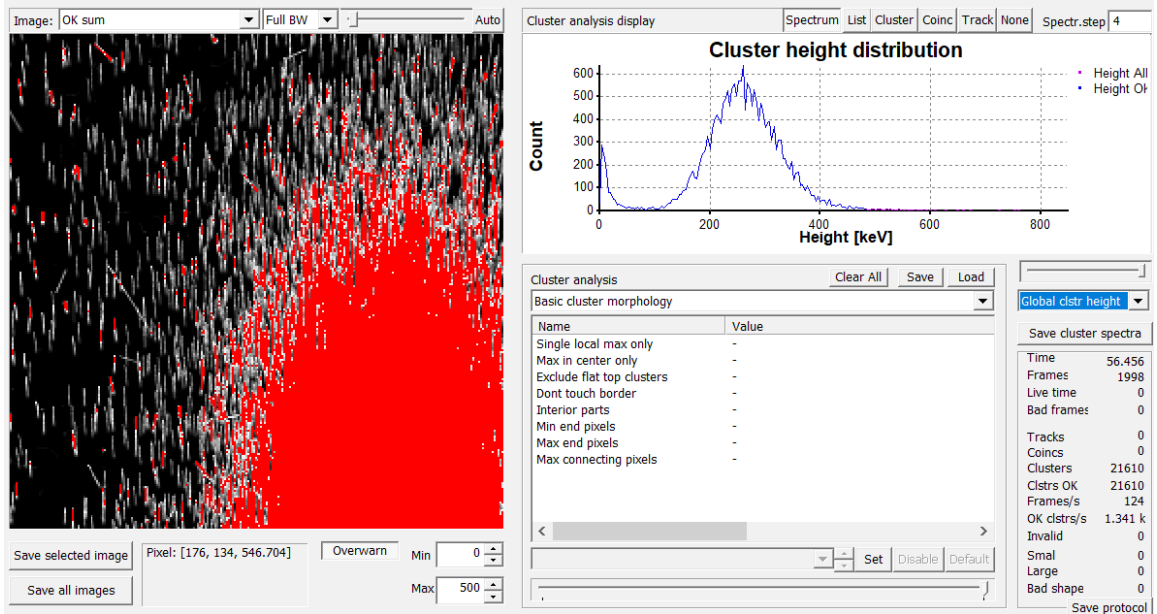


Figure 5.7: Visualisation in Advacam PIXet proprietary software [359]. Left: all integrated frames. Relative pixel intensities represent the amount of charge deposited. Right: cluster height distribution, showing energy deposited per pixel. Run R7 (60° detector angle, 10 mm Lucite) from Table 5.1 is shown. Screenshot provided by C. Oancea (Advacam, Prague).

where

$$\bar{S} = \frac{\int_{E_f}^{E_0} S dE}{\int_{E_f}^{E_0} dE} \quad (5.3)$$

and $E_f := E_{N+1}$ from (5.1). The output energy, E_f , from the Lucite block was used as the input energy to the 1 mm aluminium sliding cover, whose effective thickness, t_m , was either 1.4 mm (placed at 45°) or 2.0 mm (60°), due to the trigonometry. The total WET (listed in Table 5.1) was calculated as the sum of those for the Lucite block and the aluminium shielding layer.

5.2.3 Calculation of LET spectra

Each particle that reaches the sensor of the detector manifests as a cluster of pixels that produce electronic signals in response to a cascade of charge depositions. Events were captured in log files, which were sent to Advacam and converted to cluster list files using proprietary software. These cluster lists contained information for indi-

vidual tracks resolved by the PIXet software [359], including deposited energy per pixel, cluster area, incidence angles and the projected track length onto the sensor plane. Additionally, parameters related to track morphology were provided, such as linearity and roundness, which assist in the identification of particle type. For example, electrons often create curly (non-linear) and thin tracks while ions create wide and straight tracks. In the case of protons, electrons and other energetic charged particles, a large angle of incidence with the sensor provides clearer track structure and accentuates differences in morphology between clusters caused by different particle types. Hence detector angles of 45 and 60 degrees were chosen to more easily separate proton- and electron-induced events. The reader is referred to the works of Opalka, Granja and colleagues [361, 362, 358] for detailed methods of track pattern recognition.

The deposited energy per pixel was integrated over pixels corresponding to a single event cluster and divided by the full length of the track to retrieve the LET contribution in silicon. The full length of the 3D track, ℓ_{3D} , was calculated using the elevation angle, β , and the projected track length, $\ell_{proj.}$, with $\ell_{3D} = \ell_{proj.} / \sin \beta$ (refer to Figure 5.8). Contributions from all clusters were binned according to the Freedman-Diaconis rule [363], which is suitable for data with heavy-tailed distributions and has low sensitivity to outliers. The Landau distribution, a stable distribution with stability and skewness parameters $\alpha = 1$ and $\beta = 1$, respectively, well describes the energy deposition of a charged particle in an absorber of finite (and usually thin) thickness⁴ [364, 365], and therefore can be fitted to the binned LET data. The moments of the Landau distribution are undefined, and as such the concepts of mean or variance are unhelpful. Instead, the most probable value (MPV), at which the peak of the density function is located, is most appropriate for reducing the estimated LET

⁴The term “thickness” here is an abstraction from the absolute depth of the material, which is related to the strength of attenuation offered by the material and the energy of the particles that impinge on it.




Figure 5.8: Example of particle tracking in the MiniPIX Timepix detector. Reproduced from [351, Fig. 2] with permission.

to a single value [366]. The MPV was therefore determined, both with a Landau fit and with the α and β stable distribution parameters allowed to vary freely within their permitted intervals of $0 < \alpha \leq 1$ and $-1 \leq \beta \leq 1$. This method provided a measure of how well the Landau distribution fits to the data.

5.2.4 Simulated Monte Carlo spectra

Monte Carlo simulations were conducted in which geometrical setup and beam dynamics associated with each experimental setup were accurately modelled. This required accurate modelling of the beam delivery system, all six Lucite blocks and the detector.

An advanced Geant4 [162] model of the CCC beam line [345, 346, 165] developed by J. Yap (University of Liverpool) and colleagues at University College London has recently been repurposed for use in TOPAS [161]. The model includes all components of the beam delivery system, including scattering foils, modulators and dose monitors placed between the initial vacuum tube and the final treatment nozzle. All components are accurately represented using computer-aided design (CAD) based on on-site measurements, which were imported to TOPAS as stereolithography (STL)

files. An artificial 50 μm gap between any two components is enforced to avoid geometry overlap errors encountered by TOPAS during the simulation. The reader is referred to [347] for a detailed discussion of this and other technical aspects of the model setup.

As mentioned in Chapter 4, TOPAS features default physics processes and particle tracking settings that have been optimised for proton therapy applications [322]. However, TOPAS also allows modular physics lists in which the user may select modules that contain, for example, transport models or cross-section data that better suit their application. Additionally, cuts and thresholds can be adjusted to control the speed of the Monte Carlo simulation or the precision to which individual particles are tracked through a volume. The CCC model was benchmarked to quality assurance measurements obtained at the treatment facility, from which optimal performance was obtained through the inclusion of the electromagnetic physics module `g4em-standard_opt3` and the high energy physics module `g4h-phy_QGSP_BIC_HP`. Production cuts were set to limit the mean free path of secondary particles to a minimum of 10 μm for electrons and protons, and 100 μm for photons and positrons.

The source of primary protons in the CCC model, set to emulate the output of the cyclotron, is sampled from a Gaussian distribution with a mean energy of 62.20 MeV and a spread (standard deviation) of 0.53 MeV. The passage of protons through the 16.5 m transport line results in a mean output energy of approximately 60 MeV. Initial simulations suggested a transmission rate of around 3% to 5% between the source and the nozzle, likely due to scatter and a noticed tendency for TOPAS to encounter a significant number of stuck tracks. Exploiting multi-threading on 23 CPU threads, it took on average 8.5 hours to yield 1.5 million protons exiting the nozzle. This posed a heavy computational demand on the simulations, as suitable particle counting statistics are vital for Monte Carlo studies. Both Lucite block thickness and the 1 mm aluminium shielding upstream of the sensor contribute to attenuation of

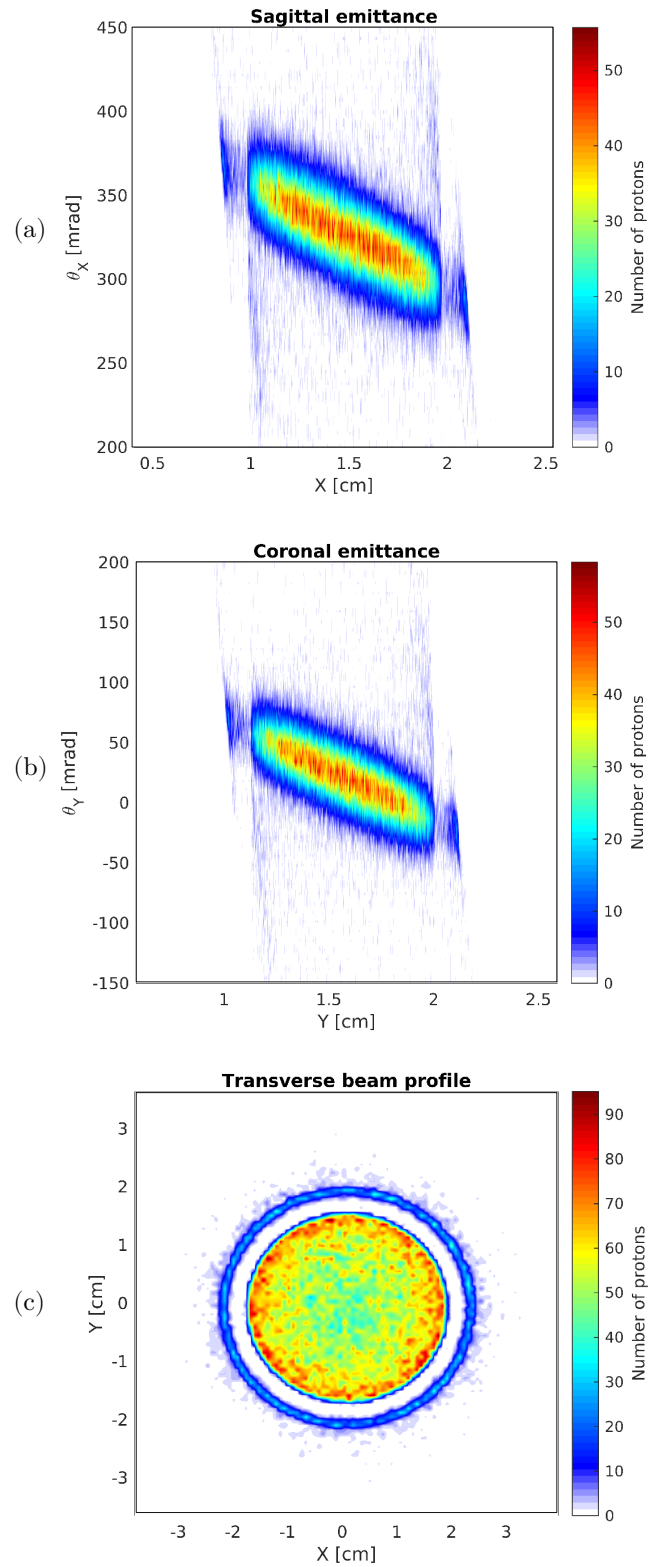


Figure 5.9: Spread of the beam from phase space of 10^6 primary protons collected immediately downstream of the nozzle, after collimation. (a) Emittance in the sagittal axis. (b) Emittance in the coronal axis. (c) Position distribution of protons in the transverse plane.




Figure 5.10: Histogram of proton energy from phase space of 10^6 primary protons collected immediately downstream of the nozzle, after collimation. A Gaussian fit with mean 60.04 MeV and standard deviation 0.48 MeV is shown.

the beam, which can result in only a small fraction of the original primary protons being recorded within the detector. To counter this, a very large number (10^6) of primary protons must exit the nozzle. This number, of course, depends on a number of factors in addition to attenuation. For example, increasing the number of the voxels in the scoring geometry – in this case, the number of pixels comprising the sensor – also demands an increase in the number of primaries, so that suitable statistics are obtained in each voxel. In order to considerably speed up the simulations, a phase space scorer was placed perpendicular to the central beam axis at the nozzle in TOPAS, and one million protons were recorded. Primary particles were then generated using sampling from this phase space file, removing the need to model the entire treatment line in subsequent simulations. The phase space data was also used to analyse the beam emittance (Figure 5.9) and energy spread (Figure 5.10) at the nozzle. A mean energy of 60.04 MeV was observed with a standard deviation of 0.48 MeV. Some minor leakage can be seen in the beam cross-section in Figure 5.9(c) owing to the aforementioned $50 \mu\text{m}$ gap between geometrical components. This was not expected to have any appreciable effect on the simulation results.

The MiniPIX Timepix sensor was modelled in TOPAS with an active area of $14 \times 14 \text{ mm}^2$ and a depth of $300 \mu\text{m}$. It was assigned the predefined Geant4 material `G4_Si`, representing silicon, and rotated by an angle of 45 or 60 degrees as in the experimental setup. Simulations were carried out for six different blocks of Lucite, defined using the TOPAS material definition of Lucite, with a mean excitation energy of 74.0 eV and a mass density of 1.19 g/cm^3 , placed 6 cm upstream of the detector. Two simulations were run for each block; one with a 1 mm thick aluminium (TOPAS material with mass density 2.70 g/cm^3) sheet placed approximately 5 mm from, and parallel to, the sensor plane, and one without. This allowed for the differences in WET to be accounted for in the simulations. A process of trial and error was used to determine suitable numbers of primary protons in each simulation, shown in Table 5.2, in order to obtain suitable statistics when scoring particles in the detector geometry.

The pixel array of the sensor was emulated by dividing the silicon geometry into a binned scorer in which each voxel independently recorded the dose-averaged LET (LET_d) of protons traversing their volume. Originally a grid of 6×6 bins was chosen to represent the 256×256 array, as the computational demand was relatively low and the detector would often recruit no more than 40 pixels per cluster, so each bin could correspond roughly to a single cluster. However, after preliminary simulations it was decided that a 28×28 grid of 784 bins provided more reliable statistics for creating the LET histograms to compare to the experimental data. Increasing the number of bins further would have drastically increased the computational demand. Measures for the number of induced DSBs and the RBE_{cd} (by DSB induction) were also recorded in these bins using the custom TOPAS scorer and associated methods discussed in Chapter 4. Additionally, the kinetic energy spectra of the incident protons between 40 keV and 60 MeV were collected at the surface of entry into the detector using a bin width of 5 keV.

The empirical energy deposition data from the detector was compared to the

simulation data by interpolating the stopping power in silicon, using the NIST pSTAR database [54], for the spectrum of kinetic energies, and then multiplying by sampled 3D track path lengths recorded by the detector during the experiment.

Finally, another simulation was set up with a water cube phantom of side lengths 35 mm, in which the LET in water, $LET_d(H_2O)$, and dose were scored at 1 mm depth intervals with no detector modelling. This provided further comparison of the different LET estimates against the depth-dose curve.

Table 5.2: Details of each Clatterbridge simulation in TOPAS, corresponding to the experiments in Table 5.1.

Run ID	Detector angle ($^\circ$)	Lucite thickness (mm)	Aluminium shielding (mm)	WET (mm)	Number of primaries	Number detected	CPU time [†] (s)
S1a	45	10.00	1.4	14.42	1.0×10^6	151640 (15.16%)	44208
S1b	45	10.00	0	11.49	1.0×10^6	152270 (15.23%)	43920
S2a	45	20.10	1.4	26.00	1.5×10^6	220931 (14.73%)	100800
S2b	45	20.10	0	23.09	1.5×10^6	221799 (14.79%)	104256
S3a	45	24.40	1.4	30.87	4.0×10^6	113177 (2.83%)	132134
S3b	45	24.40	0	28.04	4.0×10^6	585268 (14.63%)	149760
S4a	45	25.48	1.4	31.96	4.0×10^7	8037 (0.02%)	872640
S4b	45	25.48	0	29.29	4.0×10^7	5724690 (14.31%)	978912
S5a	45	25.87	1.4	∞^*	6.0×10^7	4253 (<0.01%)	1004856
S5b	45	25.87	0	29.74	6.0×10^7	7689854 (12.82%)	1009352
S6a	45	26.40	1.4	∞^*	6.0×10^7	1714 (<0.01%)	509856
S6b	45	26.40	0	30.35	6.0×10^7	4351454 (7.25%)	537288
S7a	60	10.00	2.0	15.64	1.0×10^6	135600 (13.56%)	44976
S7b	60	10.00	0	11.49	1.0×10^6	136749 (13.67%)	43128
S8a	60	24.40	2.0	32.00	4.0×10^6	4675 (0.12%)	134020
S8b	60	24.40	0	28.04	4.0×10^6	527294 (13.18%)	148680
S9a	60	25.48	2.0	∞^*	4.0×10^7	7832 (0.02%)	893040
S9b	60	25.48	0	29.29	4.0×10^7	4955708 (12.39%)	969600

[†] Simulations were multi-threaded across 24 CPU threads. Thus the real-time computation was approximately 24 times faster than values shown here.

* A WET of ∞ indicates that the vast majority of protons have stopped within the material and therefore do not reach the detector.

Table 5.3: Calculated and measured LET values from all runs in Table 5.1 and all simulations in Table 5.2. All LET values are in silicon except for the water phantom profile, LET(H₂O).

Run ID	WET (mm)	LET (keV/ μ m)				LET _d (H ₂ O) profile (keV/ μ m)	
		MPV	Q1	Median	Q3	Mean	Std. dev
R1	14.42	3.3	2.5	3.4	4.1	1.9	0.2
R2	26.00	4.3	3.8	4.6	5.6	2.7	0.2
R3	30.87	12.5	8.3	11.5	14.4	10.8	1.9
R4	31.96	15.5	6.5	11.4	15.4	18.2	1.5
R5	∞^*	7.0	4.2	7.6	11.3	N/A	N/A
R6	∞^*	1.5	1.0	1.4	3.0	N/A	N/A
R7	15.64	2.6	2.1	2.7	3.4	1.9	0.2
R8	32.00	13.5	6.3	10.8	14.2	18.5	1.4
R9	∞^*	13.5	8.0	11.4	14.2	N/A	N/A
S1a	14.42	2.2	1.8	2.0	2.3	1.9	0.2
S1b	11.49	1.6	1.5	1.6	1.9	1.9	0.2
S2a	26.00	2.5	2.5	2.6	2.7	2.7	0.2
S2b	23.09	2.1	2.0	2.1	2.3	2.3	0.2
S3a	30.87	13.6	13.3	13.5	13.7	10.8	1.9
S3b	28.04	3.1	3.0	3.1	3.1	3.2	0.3
S4a	31.96	16.5	14.9	16.9	19.0	18.2	1.5
S4b	29.29	6.1	6.0	6.1	6.1	4.1	0.5
S5a	∞^*	13.5	13.8	15.7	20.3	N/A	N/A
S5b	29.74	8.5	8.5	8.5	8.5	4.8	0.7
S6a	∞^*	15.0	14.6	17.2	22.1	N/A	N/A
S6b	30.35	11.6	11.5	11.6	11.6	7.5	1.3
S7a	15.64	1.9	1.7	1.9	2.2	1.9	0.2
S7b	11.49	1.4	1.4	1.5	1.7	1.9	0.2
S8a	32.00	14.5	12.6	15.2	18.2	18.5	1.4
S8b	28.04	3.2	3.2	3.2	3.3	3.2	0.3
S9a	∞^*	12.5	11.4	13.1	15.6	N/A	N/A
S9b	29.29	7.0	6.9	7.0	7.0	4.1	0.5

* A WET of ∞ indicates that the vast majority of protons have stopped within the material and therefore do not reach the detector.

5.3 Results

Raw data from log files and cluster lists, summarised in Table 5.1, show that the number of clusters recognised by the PIXet software was adequate for statistical analysis of the results for the thinnest three absorbers for the 45° measurements (R1, R2, R3 and R4) and for the two thinnest absorbers for the 60° measurements

Table 5.4: Calculated DSB counts and RBE_{cd} values from all simulations in Table 5.2.

Run ID	WET (mm)	$\langle \text{DSBs} \rangle$ ($\text{cell}^{-1}\text{Gy}^{-1}$)				RBE_{cd}			
		MPV	Q1	Median	Q3	MPV	Q1	Median	Q3
S1a	14.42	2.9	2.9	2.9	2.9	1.0	1.0	1.0	1.0
S1b	11.49	2.9	2.9	2.9	2.9	1.0	1.0	1.0	1.0
S2a	26.00	2.9	2.9	2.9	2.9	1.0	1.0	1.0	1.0
S2b	23.09	2.9	2.9	2.9	2.9	1.0	1.0	1.0	1.0
S3a	30.87	5.8	5.4	5.6	5.8	2.0	1.9	2.0	2.0
S3b	28.04	3.0	3.0	3.0	3.0	1.1	1.1	1.1	1.1
S4a	31.96	5.5	5.2	6.2	7.8	1.9	1.8	2.2	2.7
S4b	29.29	3.6	3.6	3.6	3.6	1.3	1.3	1.3	1.3
S5a	∞^*	5.0	5.0	6.0	7.6	1.8	1.8	2.1	2.7
S5b	29.74	4.2	4.2	4.2	4.2	1.5	1.5	1.5	1.5
S6a	∞^*	7.0	5.2	6.3	8.2	2.5	1.8	2.2	2.9
S6b	30.35	5.1	5.1	5.1	5.1	1.8	1.8	1.8	1.8
S7a	15.64	2.9	2.9	2.9	2.9	1.0	1.0	1.0	1.0
S7b	11.49	2.9	2.9	2.9	2.9	1.0	1.0	1.0	1.0
S8a	32.00	4.5	4.7	5.7	7.5	1.6	1.7	2.0	2.6
S8b	28.04	3.0	3.0	3.0	3.0	1.1	1.1	1.1	1.1
S9a	∞^*	5.5	4.6	5.5	6.4	1.9	1.6	1.9	2.3
S9b	29.29	3.8	3.8	3.8	3.8	1.3	1.3	1.3	1.3

* A WET of ∞ indicates that the vast majority of protons have stopped within the material and therefore do not reach the detector.

(R7 and R8). Conversely, runs R5, R6 and R9 yielded clusters not too dissimilar to the background readings (BG1 to BG4). This was due to the vast majority of protons having been stopped by the aluminium layer preceding the sensor (WET = ∞). Indeed, the simulation summary, Table 5.2, shows that only 0.02% or fewer of the primary protons at the nozzle reached the detector in the simulations of these runs (S5a, S6a and S9a). Meanwhile, simulations with the aluminium sliding window removed (S5b, S6b and S9b) showed much higher transmission, ranging from around 7% to 13%. The thicknesses of Lucite chosen prior to the experiment were based on their WET without the aluminium shielding. With the shielding, the Bragg peak was shifted back so that R3, with a WET of 30.87 mm, corresponds to the Bragg peak while R4 and R8 (WET values of 31.96 mm and 32.00 mm, respectively) correspond to the distal dose fall-off.

The deposited dose within the detector matches well with that calculated from the

simulated data (Figure 5.11) for the runs with suitable statistics mentioned above, however a prominent right-skew to the simulated data is present. This is believed to be, in part, due to the increased spread caused by the number of scoring bins (784) defined in TOPAS being significantly smaller than the number of channels (65,536) in the sensor. The empirical data also shows increased noise at low energies, which

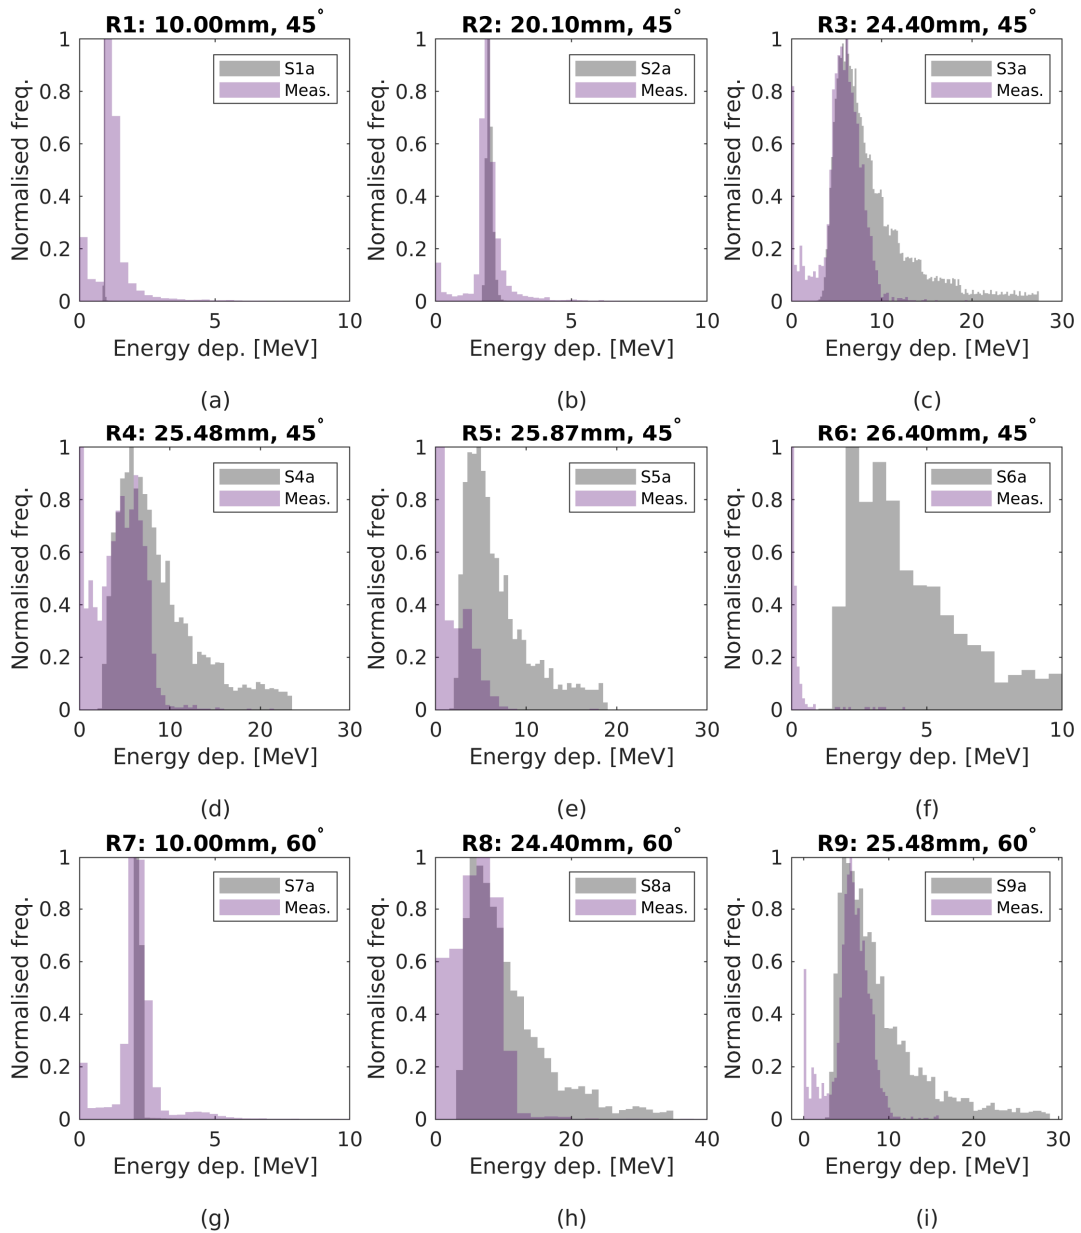


Figure 5.11: (a-i) Histograms of energy deposition as measured by the MiniPIX sensor (“meas.”) for all runs R1 to R9, and as calculated from simulations S1a to S9a which include aluminium shielding.

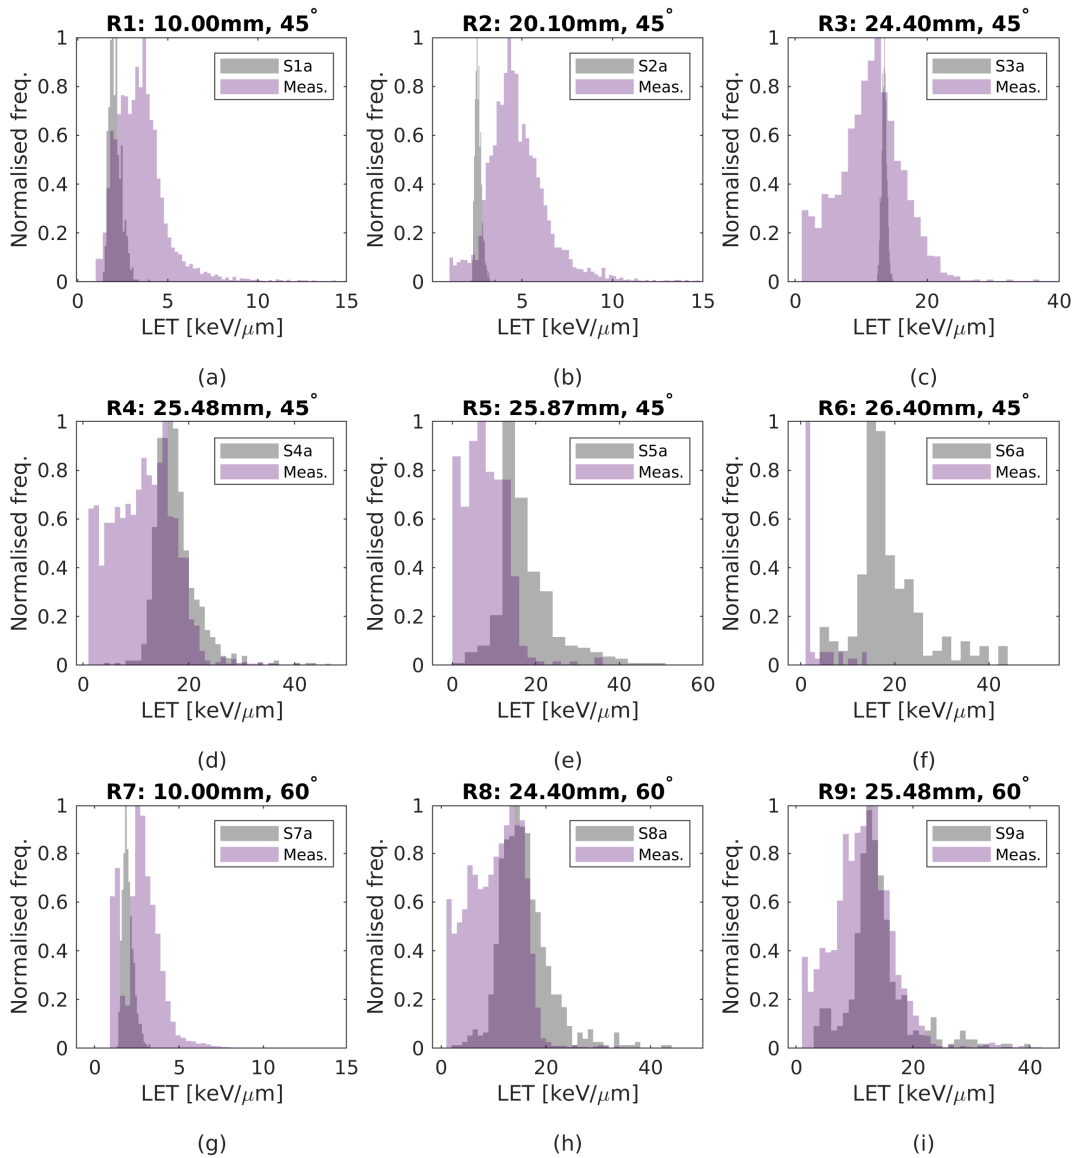


Figure 5.12: (a-i) Histograms of LET contributions calculated from the MiniPIX sensor data (“meas.”) for all runs R1 to R9, compared to LET_d output from simulations S1a to S9a which include aluminium shielding. All LET values are in silicon.

is not present in the simulated data. This could be caused by the misclassification of tracks by particle type based on their morphology, allowing electron-induced clusters to be included in the recorded data.

Histograms of the measured and simulated LET contributions are shown for all runs in Figure 5.12, and are summarised in Figure 5.13 and Table 5.3. Landau and other stable distribution fits to the LET data can be found in Appendix C. In the case that no suitable fit to the data was found, the MPV was taken as the mode of the




Figure 5.13: Boxplots showing the spread, median and MPV of LET contributions in silicon calculated from the MiniPIX sensor data beside those retrieved from the simulations S1a to S9a. Mean $LET_d(H_2O)$ values scored from a water cube phantom simulation are also shown. Runs for which particle counting statistics were considerably low, corresponding to $WET = \infty$, (R5, R6 and R9) have been faded out.

data. A general increase in the LET is observed as the Lucite thickness, and therefore the WET, is increased. The LET ranges from around $3 \text{ keV}/\mu\text{m}$ in the entrance region to between $15 \text{ keV}/\mu\text{m}$ at the distal edge, with a value of approximately $12 \text{ keV}/\mu\text{m}$ at the Bragg peak. The spread, namely the interquartile range, of the data also increases with WET, due to the lower particle count in the detector. All LET values from the empirical measurements and detector simulations were converted to water equivalent values, $LET(H_2O)$, using the empirical formula of Benton *et al.* [367] and plotted alongside the water phantom simulation LET as a function of WET in Figure 5.14. The figure shows a clear trend of increasing LET toward and beyond the Bragg peak. Values in the entrance and build up regions did not exceed $5 \text{ keV}/\mu\text{m}$ and therefore would not be expected to cause any clinically significant radiobiological effect [211] when compared to megavoltage x-ray treatment. This is also reflected in the calculated RBE_{cd} estimates, which were between 1.0 and 1.1 for WET values

below around 28 mm. In general the increasing trend in LET, both by simulation and through detector measurements, shown in Figure 5.14, resembles expectations. Figure 5.13 shows that the interquartile range of the simulated LET generally overlaps with that of the measured LET, and that the spread in the LET spectra increased as the Lucite thickness was increased. This was expected as there were fewer particles reaching the detector for larger WET values, resulting in decreased statistics for the LET calculation.

Edge-of-field effects were briefly explored in the water phantom simulation by scoring the LET away from the central axis in 1 mm^3 bins. Figure 5.15 shows a cross-section of scored LET_d in the entrance region and around the Bragg peak and distal dose fall-off. Lateral penumbra LET_d was almost twice as high as at the central axis at a depth of 10 mm but this was still below $4 \text{ keV}/\mu\text{m}$ and would not be expected to lead to any substantial difference in biological outcome. At the Bragg peak and beyond the difference was much smaller, with LET_d varying by less than 5% between the central axis and lateral penumbra.

The incident proton kinetic energy spectra, displayed in Figure 5.16, show the substantial affect of the inclusion of thin aluminium shielding placed immediately upstream of the sensor. The distribution of proton energies was systematically shifted toward lower energies, leading to higher LET and RBE_{cd} estimates. Simulated DSB counts in the detector, summarised in Table 5.4 and Figure 5.17, were shown to double toward the Bragg peak and distal edge when compared to the entrance region. The associated RBE_{cd} ranged from 1.0 in the entrance and build up regions to approximately 2.0 at the Bragg peak. A sharp increase in the estimated RBE_{cd} is observed leading up to the Bragg peak (Figure 5.18). A decrease in RBE_{cd} is then seen at the distal edge, however this observation is not as prominent in the interquartile range (Figure 5.17) as it is in the calculated MPV.




Figure 5.14: All LET results, both measured and simulated, corresponding to finite WET plotted as their MPV against a simulated depth-dose profile in a water phantom. Error bars indicate the 95% confidence interval centred on the mean. LET values in silicon have been converted to water equivalent values, $\text{LET}(\text{H}_2\text{O})$, for all detector measurements and simulations.

5.4 Discussion

The MiniPIX-Timepix semiconductor detector was used to measure the energy spectra of individual protons passing through Lucite absorbers of varying thickness at the CCC clinical proton beam line. This enabled calculation of LET spectra at varying WET along the depth-dose curve, which were compared to simulations using an advanced TOPAS model of the treatment line. The calculated LET was in general agreement with simulated dose-averaged LET and within expectations, with values around or below $3 \text{ keV}/\mu\text{m}$ in the entrance region before a sharp increase up to between 13 and $18 \text{ keV}/\mu\text{m}$ at the distal dose fall-off. Those values prior to the build-up in the last 10% of the proton range are therefore of little biological consequence. The spread in the calculated LET increased substantially toward larger WET values, a trend that was also present in the simulation data. This was a consequence of a sig-




Figure 5.15: Cross-section of scored LET_d in TOPAS at depths of (a) 10 mm, (b) 30 mm, (c) 32 mm and (d) 33 mm in a water phantom of side lengths 35 mm. Bin sizes of 1 mm^3 were used for scoring. Images have been upscaled by a factor of 10 and interpolated.

nificant number of protons stopping in the absorber, or the aluminium shielding layer, thus not reaching the detector. If the experiment were repeated, longer acquisition times would be allowed for each run to achieve suitable statistics. However, this was challenging on the day of the experiment as the beam line was needed for treatment and for other experiments, limiting the time available. This was also the reason why more runs were not performed with finer increments in the Lucite thickness. Given the limited beam time, there were plans to return to CCC in early-to-mid 2020 after analysing the first batch of data and being able to revise and improve the experimental procedure. Unfortunately a number of logistical factors prevented this from going ahead⁵. The revision to the procedure that would have had the most influence was the

⁵Note: returning to CCC for subsequent measurements was not possible due to the lockdown measures imposed in England during the Covid-19 pandemic.

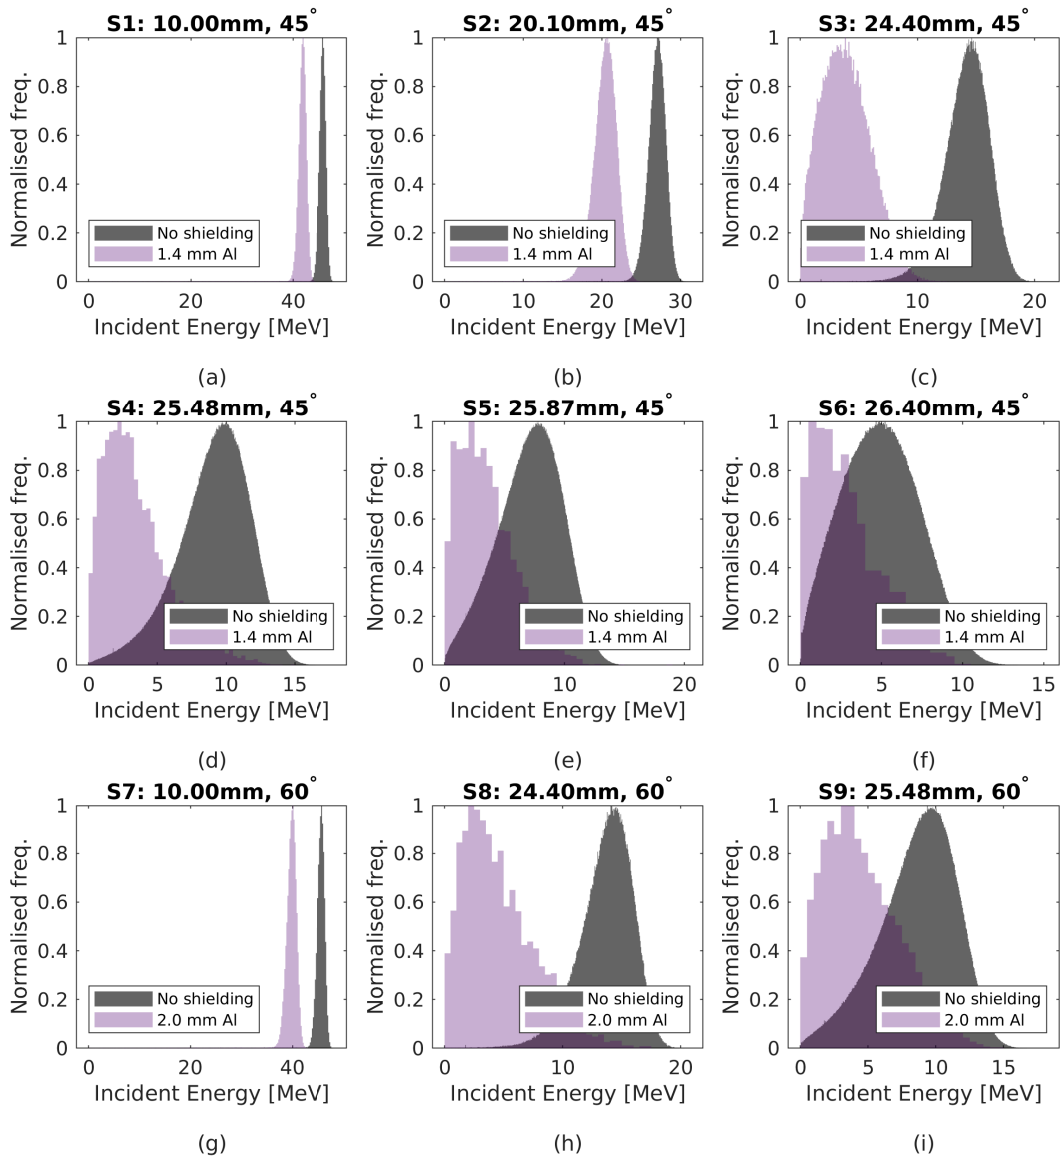


Figure 5.16: (a-i) Histograms of the kinetic energy of protons impinging on the surface of the sensor, as scored in the simulations S1a to S9a (with aluminium shielding) and S1b to S9b (without shielding).

removal of the aluminium shielding layer. This was kept in place at first to prevent damage to sensor components caused by a high dose rate from the cyclotron, however as Lucite thickness was increased this became less of an issue and the shielding could have been removed.

Aluminium is a very effective absorber compared to water or Lucite. Figure 5.19 shows the range of protons in the material under the continuous slowing down approximation (CSDA). Notice that incident protons below approximately 13 MeV

are expected to stop within 1 mm, and would therefore not reach the detector. Not only does this effectively apply a high-pass filter to the energy spectra, but for protons of energies larger than this threshold of 13 MeV, it artificially provides a source of low-energy, high-LET protons to the sensor. This effect was partially accounted for by calculating the WET due to both the Lucite block and the aluminium layer, however factors such as the rate of transmission and scatter still have a significant affect on the spread of the calculated LET in each pixel of the sensor.

Some discrepancies between the mean LET_d values collected in TOPAS from the water phantom and those LET values calculated in the detector, both actual and simulated, could be partly explained by differences in the methods by which LET is calculated. The ProtonLET scorer provided in TOPAS is based on general Monte Carlo methods outlined by Cortés-Giraldo, Granville and colleagues [60, 61]. In general, variation between Monte Carlo scoring techniques is larger for LET_d than for LET_t [61]. While LET_d was scored in the phantom, the LET calculation from




Figure 5.17: Boxplots showing the spread, median and MPV of DSB counts, and the associated RBE_{cd} by DSB induction scored from the simulations S1a to S9a. Runs for which particle counting statistics were considerably low, corresponding to $WET = \infty$, (R5, R6 and R9) have been faded out.




Figure 5.18: All simulated RBE_{cd} estimates corresponding to finite WET plotted as their MPV against a simulated depth-dose profile in a water phantom. Error bars indicate the 95% confidence interval centred on the mean.

the measured MiniPIX data was neither track-averaged nor dose-averaged, as it was taken to be the MPV of a spectrum of voxelised LET contributions. Additionally, while WET allows comparison between depth profiles across different materials, it is not a substitute for material properties, which influence the particle energy spectrum.

For small bin sizes in a scoring geometry, there is a possibility of artificially over-estimating LET contributions due to the voxel size being smaller than the mean path length between events [368]. Measurements of the LET are confined to the scoring bin, rather than the track being traced beyond these limits. The scoring geometry defined for the detector simulations consisted of voxels with dimensions $0.5 \text{ mm} \times 0.5 \text{ mm} \times 0.3 \text{ mm}$. The largest possible “track length” within a voxel was therefore $768 \mu\text{m}$, the longest corner-to-corner distance. In addition, the angle of the detector in the simulations could have allowed a portion of particle tracks to traverse one voxel but then leak into another. This second voxel would then record a higher LET track that has already lost energy through its traversal of the first voxel. This can also




Figure 5.19: Stopping power (left axis) and CSDA range (right axis) in Aluminium, showing the energy range in which protons traversing the 1 mm shielding layer may be absorbed and not reach the detector.

be a source of overestimation. Figure 5.20 provides an alternate visualisation of the MiniPIX data which, to this author’s knowledge, has not been presented before in the existing literature. It shows the distribution of deposited energy and track lengths that contribute to the LET calculation, from all runs, with level curves of constant LET superimposed for reference. It captures energy, track length and LET spectra in a single plot. With this graphic we see that imposing an artificial cutoff track length of $768 \mu\text{m}$ is unlikely to overestimate the LET by any considerable amount within the range of deposited energies that were actually recorded in the MiniPIX detector.

The extent to which MiniPIX cluster files were contaminated with electron tracks was unknown, as filtration techniques were mostly contained within the PIXet software and in supplementary scripts provided by Advacam⁶. As such, a thorough analysis of filtration techniques was beyond the scope of this thesis. Some degree

⁶Scripts provided by C. Oancea (Advacam, Prague).




Figure 5.20: Distribution of deposited energy and track lengths from the MiniPIX data (runs R1 to R9, Table 5.1) with level curves of constant LET superimposed for reference.

of contamination is to be expected as distinguishing particle types through track morphology alone is error-prone, particularly for short, low-energy tracks. This was partly accounted for by placing the detector on an angle (45° or 60°) to the incident particle beam, however any electrons that were erroneously included in the cluster files would have contributed to the lower part of the energy deposition spectrum and influenced the LET calculations. An example of this, prior to filtration, can be seen in the cluster height spectrum for run R7, shown in Figure 5.7. The secondary low energy peak is most likely due to electron tracks, and it can be seen in Figure 5.11(g) that at least some of these tracks have remained post-filtration. Contamination of this sort together with low particle counts in the detector for some runs may explain the deviation of the LET spectra from a Landau distribution.

DSB counts and the associated RBE_{cd} (by DSB-induction) were scored in all

detector simulations listed in Table 5.2 according to the methods outlined in Chapter 3 and Chapter 4. The trend of estimated RBE_{cd} , starting at 1.0 in the entrance region and then sharply increasing in the last 10% of the depth-dose curve corroborated the trend in LET. As was the case for the LET, the spread in the data increased dramatically with the thickness of the Lucite block used, due to lower particle counts in the detector. The arguments presented above for the affect of the aluminium shielding on the particle energy spectrum equally apply here to the accuracy of the RBE_{cd} estimates.

As mentioned in the Section 5.1, previous radiobiological experiments undertaken at CCC involving RBE determination and DSB analysis include those of Chaudhary *et al.* [176] and Carter *et al.* [350]. In [176] the DNA damage response of human skin fibroblast (AG01522) cells was studied using a 53BP1 foci formation assay following irradiation at several depths in the 60 MeV monoenergetic beam generated with the




Figure 5.21: Most probable values (MPVs) of LET plotted against MPVs of RBE_{cd} collected from all simulations in Table 5.2. The shaded region represents the area between the empirical RBE_{cd} fit, (4.11), to LET_d (bottom) and LET_t (top) presented in Chapter 4.

CCC treatment line. A statistically significant difference was found between the number of foci counted 24 hours post-irradiation at the Bragg peak due to protons and a reference irradiation (1 Gy of 225-kVp x-rays). The ratio of foci can be calculated from [176, Fig.2]. Assuming this ratio represents only irradiation-induced DSBs then it corresponds to an RBE estimate of 2.0, which exactly matches the prediction of this current study. It should be noted, however, that no significant difference was found in foci formation immediately following irradiation; the discrepancy was only present in sustained damage. The study by Carter *et al.* [350] investigated mechanistic processing of DNA damage following irradiations of HeLa and oropharyngeal squamous cell carcinoma (UMSCC74A and UMSCC6) cells at the CCC beam line. Comet assay and clonogenic survival assay techniques were used to assess damage to cells placed downstream of the same 24.40 mm Lucite absorber used in this study. RBE was estimated between 1.67 and 1.85 across all cell lines. This does not match with the findings of the 24.40 mm simulations, S3b and S8b (without shielding), however Carter also reported an effective beam energy of 58 MeV, rather than the 60.02 MeV value used in this study.

Previous work by Chaudhary *et al.* [175] at another facility, running a 62 MeV proton beam line, examined the relationship between LET and RBE. In that study RBE was measured with a surviving fraction endpoint by clonogenic assay using AG01522 and glioma (U87) cells. The RBE_{cd} -LET relationship inferred from simulations in this study (Table 5.2) is shown in Figure 5.21. The shaded region represents the area between the empirical RBE_{cd} fit, (4.11), to LET_d (bottom) and LET_t (top) presented in Chapter 4. The simulation results from this work fall mostly within this range, however they best resemble the relationship between RBE_{cd} and track-averaged LET. The general trend and range of RBE_{cd} values resembles those of [175, Fig.5], considering the variation between cell lines and surviving fraction endpoint (50%, 10% or 1%).

Chapter 6

A flexible projection-based inverse planning algorithm for radiation therapy

6.1 Introduction

In Chapter 1 we discussed the increasingly adopted technique of intensity modulated radiation therapy (IMRT) for dose delivery. IMRT, whether by photons or protons, offers improved dose conformality to the tumour region when compared to conventional 3D conformal radiotherapy (CRT), while also limiting the dose deposited in surrounding healthy tissues. We have explored, in previous Chapters, the added therapeutic benefits of proton beams, or indeed heavier charged particle modalities, over photons in terms of both their physical dose deposition characteristics and radiobiological consequences. These benefits do not come without scrutiny or compromise; highly localised dose at Bragg peak positions demands highly accurate and precise dose delivery, while biological range uncertainties call for scrupulous beam selection criteria. In order to fully realise the potential of proton therapy, sophisticated and clinically implementable treatment planning algorithms are needed.

A new and novel algorithm for radiation therapy inverse planning is presented in this Chapter in the context of fluence map optimisation (FMO) for intensity modulated proton therapy (IMPT). Though, it should be noted that the techniques discussed here are agnostic to the type of radiation and how it is delivered; all that is

required is that the final dose pattern can be expressed as a linear combination of discrete contributions and that, in practice, the relative contributions of these discrete sources can be modulated. As such, the algorithm may be equally applied to the FMO formulations of photon or heavy-ion IMRT. The algorithm, presented here as the *dynamic string-averaging CQ-method (DSACQ)*, is based on a *split-feasibility* formulation of the IMRT inverse problem, contrary to conventional gradient-descent-based minimisation of an objective function. The DSACQ method iteratively applies metric projections onto sets that represent the imposed dose constraints, while also allowing flexibility in the weighting and order in which projections are executed onto the individual constraint sets. Dose-volume effects are handled through the inclusion of dose volume constraints (DVCs). Representing these DVCs leads to a non-convex optimisation problem, so they are often approximated with additional linear constraints or handled with impractical heuristics. In contrast, the DSACQ algorithm is based on iterative projections onto both convex and non-convex constraints, and is able to generate feasible dose solutions, if such solutions exist, with any number of linear or non-linear constraints imposed. DVCs are represented exactly, with no linear approximation necessary.

6.1.1 Mathematical challenges with dose-volume criteria

FMO techniques applied to the selection of beamlet weights for radiotherapy have been largely based on weighted least-squares models, drawn from analogies with tomographical image reconstruction [369]. In a least-squares approach, a dose-dependent objective (or “cost”) function is defined and linear programming methods are used to select a dose map that minimises this cost function. This is the technique used most widely by commercial treatment planning systems [370], most often achieved through gradient descent methods. The cost function, f , is usually a weighted combination of structure-specific piecewise-quadratic functions, f_s , which penalise any deviation

from the dose prescription. That is, for a given spatial dose distribution, $\mathcal{D}(x, y, z)$, in the volume of interest,

$$f(\mathcal{D}(x, y, z)) = \sum_{s=1}^{N_s} w_s f_s(\mathcal{D}(x, y, z)), \quad (6.1)$$

where N_s is the number of defined structures to which prescriptions are applied, and w_s , for $s \in \{1, 2, \dots, N_s\}$, are weighting factors that set the relative priority of these prescriptions. In practice each penalisation function, f_s , is a sum of penalty terms, each of which applies to an individual voxel in a discretised dose array. Under this fully-discretised approach the problem statement for the inverse planning of IMPT may be formulated using systems of inequalities (see, e.g., [21, 371] and references therein). In this approach a discretised three-dimensional dose array is created to be coincident with voxels in the patient CT or co-registered planning images. Each beamlet can be viewed as giving rise to a separate, independent dose grid, and by rearranging all voxels in a known order one may record in a single column the dose in each voxel due to a single beamlet. Appending all beamlet contributions forms a *dose-influence* matrix, A , in which the element in the i -th row and j -th column gives the dose in voxel i due to beamlet j . This is demonstrated graphically in Figure 6.1 for the case of proton pencil beams in IMPT. The dose distribution can be changed by adjusting the relative weights (intensities) of each pencil beam. Multiplying A by the column vector of all relative beamlet intensities, \mathbf{x} , gives the total dose in each voxel. The goal of the planner is then to find a suitable vector, \mathbf{x} , with non-negative entries, such that the dose values obtained in the vector $A\mathbf{x}$ satisfy, as best as possible, the prescribed constraints for each structure.

Suppose the i -th voxel in the dose array is within an OAR and receives a dose of $d_i(\mathbf{x})$, dependent on the beamlet intensity vector, \mathbf{x} . An upper dose bound, d_{\max} , imposed on this voxel sets up the prescription $d_i(\mathbf{x}) \leq d_{\max}$. Applying this prescription

to every voxel within the OAR would yield a corresponding penalty term [370],

$$f_{\text{OAR}} = \sum_{\substack{\text{all voxels } i \\ \text{within OAR}}} \max\{0, (d_i(\mathbf{x}) - d_{\text{max}})^2\}. \quad (6.2)$$

This penalty function renders the optimisation problem well-behaved from a mathematical standpoint, as it is both convex and differentiable, allowing for linear programming methods to be implemented. Though, in deviating from HDCs to incorporate additional DVC prescriptions, the FMO inverse problem becomes non-convex and the cost function non-differentiable. This significantly increases the complexity of the problem and gives rise to multiple local minima in the cost function [372, 373]. While an estimate of the solution can be verified against the plan prescription with efficiency (in polynomial time), the pursuit of a global optimum is classified as a non-deterministic polynomial time hard (NP-hard)¹ problem [372, 374]. In the presence of DVCs, conventional gradient descent algorithms can easily fall into “traps” in which a suboptimal solution results. An illustrative analogy is that of finding the lowest point

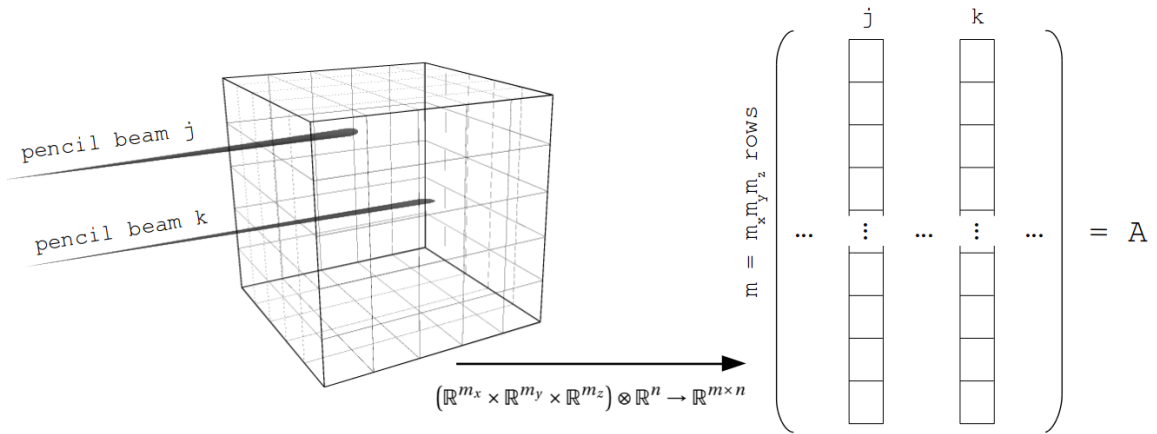


Figure 6.1: Each pencil beam gives rise to a 3D dose grid of dimensions $m_x \times m_y \times m_z$ consisting of $m = m_x m_y m_z$ voxels, which is transformed to a column vector of length m . For n pencil beams, columns are concatenated to form the $m \times n$ dose matrix A .

¹The NP-hard problem classification is far beyond the scope of this thesis, so more detail is not provided here. In fact, it is part of a famous question in computer science posed almost 50 years ago. The Clay Mathematics Institute offers a million dollar prize to whomever eventually finds a solution. See: <https://claymath.org/millennium-problems/p-vs-np-problem>

on a mountainous landscape. As far as the algorithm is concerned, a shallow valley holds as much merit as a deep gorge; the first encountered is the accepted solution. Despite this glaring problem, FMO algorithms have traditionally consisted of simple least-squares formulations with a nonlinear volume-sensitive penalty term appended to the cost function [375, 376]. Metaheuristics such as simulated annealing, used as local search methods, have been demonstrated to alleviate the suboptimality problem [22, 23, 377].

Issues of non-convexity arising from DVCs are often skirted entirely by imposing gradations of linear dose constraints on a number of subvolumes, inserted manually in the structure of interest [378]. Advanced interior point methods are able to provide a solution very efficiently [379, 380, 381, 382, 383], however, this approach amplifies an intrinsic limitation of IMRT inverse planning: multi-objectiveness [see, e.g., 370]. A large number of conflicting dose objectives are applied to structures in close proximity, which may never be fulfilled due to physical limitations on dose conformality. Censor *et al.* [384] addressed infeasibility in the solution space by introducing an $\{\alpha, \beta\}$ -relaxation², meaning no more than $100\alpha\%$ of the dose bound inequalities can be violated by no more than $100\beta\%$. This, in fact, is mathematically equivalent to imposing a DVC; a proportion of voxels within a structure are permitted to violate their HDCs, but with an upper bound imposed on the magnitude of the violation. For given values of α and β there is degeneracy in the choice of which specific voxels may violate their constraints. The IMRT inverse problem therefore becomes a combinatorial problem which can be solved using a variety of mixed integer programming (MIP) methods [385, 386, 387, 388, 389, 390, 391, 392], including later formulations specifically for IMPT [393, 394], though such methods can be prohibitively impractical and time-consuming at the scale of clinical plan complexity [372, 384, 370, 394]. MIP for-

²The term “relaxation” is used carefully. Throughout this Chapter it is (unavoidably) mentioned by three unique definitions. Here it refers to allowing a certain amount of overflow or underflow compared to the dose prescription. It is later used in reference to projection onto sets, and to approximating non-convex sets with a convex hull.

mulations hard-code and strictly enforce the prescribed DVCs, which differ from the more commonly used least-squares formulations in which DVCs are only encouraged through additional terms in the objective function. The former are sometimes referred to as “hard DVCs” and the latter “soft DVCs”. In aiming for fast plan production in clinical settings, soft DVCs are generally favoured [395, 376, 396, 379, 397, 398]. Mukherjee *et al.* [399] recently proposed a two-phase optimisation algorithm that makes use of both soft and hard DVC formulations; a convex heuristic method first identifies voxels likely to receive doses violating the DVC thresholds, before applying hard dose limits on those voxels.

Convex relaxations of DVC prescriptions are attractive alternatives to hard-DVC-based MIP algorithms for flexible and computationally tractable inverse planning [400], however the inverse problem is merely an approximation under such a relaxation and, in the case of least-squares linear programs, “optimality” of the solution remains governed by a largely subjective cost function. Plan quality is subject to expert interpretation and it is therefore not uncommon for physicians to disagree at the stage of plan approval [401, 402]. Due to this subjectivity and the multi-objectiveness of IMRT inverse problems, the weighting factors, w_s , in (6.1) for each structure in the cost function must be judiciously adjusted numerous times by a trial-and-error process before a plan is considered clinically acceptable [403]. This issue has been somewhat mitigated through the adoption of multi-criteria optimisation (MCO) techniques in modern treatment planning systems, whereby a number of *pareto-optimal* solutions may be presented to the planner that are noticeably different, but still equally valid choices based on the value of the cost function [404, 405, 406, 407, 408]. For example, one plan may place a higher priority on achieving tumour dose coverage while another may sacrifice some coverage for reduced leakage to an adjacent critical structure. Both plans may be considered optimal by some metric, but perhaps not by others. While MCO enables more possible treatment scenarios to be explored,

they remain computationally inefficient in practice and expert clinical judgment is required to assess overall plan quality [400]. Some attention has been given to the direct use of dose-response metrics in place of DVC-relaxation. Convex problems may be formulated with quasi-biological models such as tumour control probabilities (TCP) and normal tissue complication probabilities (NTCP) based on, for example, equivalent uniform dose (EUD) [155, 156, 396, 409]. Others have abstracted concepts from areas of operational research, including rectifier (or sigmoidal) activation [406], conditional value-at-risk (CVaR) [379, 380, 410] and stochastic dominance [411]. However, there are some drawbacks with this dose-response approach. Quantities such as TCP, NTCP and EUD are highly degenerate functions of dose-volume combinations, and can therefore be insensitive to dose hot spots and dose heterogeneity in general [396, 409]. While it is certainly convenient to reduce DVH data down to a single value, there is significant loss of information in this process and no universally accepted metric to substitute case-by-case clinical judgement.




Figure 6.2: An example of convex feasibility-seeking in two dimensions. The feasible region (shaded) arises from the intersection of lower-halfplanes defined by lines H_1 and H_2 . Starting with an arbitrary initial estimate $\mathbf{x}^{(0)}$, the next iterate is computed by projecting $\mathbf{x}^{(0)}$ orthogonally onto H_1 . Sequential projections are made between H_1 and H_2 , with the estimate converging on a feasible solution, \mathbf{x}^* .

The discussions above have expressed a number of challenges in the incorpora-

tion of dose-volume effects in radiation therapy treatment planning. Conventional cost-minimisation approaches are subject to multiple local minima and therefore sub-optimal plan quality, while convex reformulations of the inverse problem are approximate and subjective. Addressing dose-volume effects through DVC prescriptions allows for the robust consideration of well-defined complication probabilities accumulated over decades of documented clinical experience [385, 372]. Further, these constraints are physically well-understood and tangible, unlike “black-box” models used in the calculation of metrics such as TCP and NTCP [372]. In defining multiple DVCs on a structure there is also flexibility in the amount of control placed on the desired DVH curves. Exact modelling of DVC prescriptions implemented through MIP methods can produce an optimal solution, if one exists, however such methods can be prohibitively impractical. Multi-objectiveness of planning criteria and subjectiveness in the assessment of overall plan quality shift the problem of IMRT inverse planning from one of optimisation to one of *compromisation*. Here, we define *compromisation* as the process of finding any *feasible* dose solution that results from unavoidable trade-offs between meeting the imposed constraints. *Feasibility-seeking*, as opposed to cost-minimisation, algorithms are therefore well suited to IMRT inverse problems, and have been a topic of development over the past four decades [18, 21, 412, 413, 414, 384, 371, 415, 416, 378, 417]. See, in particular, the commentary by Censor [412] and references therein. In a fully-discretised approach to IMRT inverse planning, as described earlier, HDCs may be represented by half-spaces in the beamlet intensity vector space. Feasible dose solutions for IMPT plans have been shown to result from iterative projections of the beamlet intensity vector onto such half-spaces [416]. Figure 6.2 demonstrates how orthogonal projections may be used to converge on an a feasible solution to a convex feasibility-seeking problem. DVCs manifest as additional $\{\alpha, \beta\}$ -relaxations [384] which give rise to non-convex sets in the dose vector space. The result is a *split* feasibility problem between these two mathe-

mathematical spaces of different dimensions; the space of beamlet intensity vectors, \mathbf{x} , and the space of dose vectors, $A\mathbf{x}$. Split feasibility problems have been examined with mathematical rigour in cases where all sets are convex [418, 419, 420, 421, 422, 423], although there have been a limited number of recent studies which explored the behaviour of iterative projections onto non-convex sets [424, 425, 426, 427]. In previous work, $\{\alpha, \beta\}$ -relaxations were formalised as percentage violation constraints (PVCs) specifically for the purpose of modelling DVCs in radiation therapy [378]. Non-convex DVCs were imposed in a continuous manner, as opposed to integer methods, through the use of the ℓ_0 sparsity-norm, which counts the number of non-zero entries in a vector. This Chapter builds upon those methods, offering improved flexibility in its implementation and suggestions for clinical use, including an automatic parameter-selection protocol that requires no intervention from the user.

6.2 Methods

Section 6.1 introduced the fully-discretised IMRT inverse problem and offered insight into its reformulation as a split-feasibility problem, which is briefly summarised as follows:

1. each beamlet gives rise to a column of a matrix, A , which gives the dose deposited in each voxel due to one beamlet;
2. the relative number of fired particles in each beamlet is given in a column, \mathbf{x} ;
3. $A\mathbf{x}$ is then the weighted average of the beamlet doses, and therefore gives the total dose in each voxel;
4. mathematical feasibility-seeking methods are then used to choose suitable beamlet intensities, \mathbf{x} , such that the prescribed dose constraints are met.

Suppose we have S_a closed volumes representing avoidance structures (OARs) and S_t closed volumes representing target structures, such as the PTV. Suppose a structure with identification index $s \in \{1, \dots, S_a, S_a + 1, \dots, S_a + S_t\}$ contains m_s voxels. The matrix $A_s = (a_i^j) \in \mathbb{R}_+^{m_s \times n}$ contains the forward-calculated dose due to n pencil beams with unit radiation intensity. Variable relative pencil beam intensities (which must be nonnegative) are stored in the column vector $\mathbf{x} = (x_j)_{1 \leq j \leq n}$. That is, $(a_s)_j^i x_j$ is the dose deposited in voxel i due to pencil beam j . The total dose in voxel i , which shall be denoted d_s^i , due to all beamlets is given by the superposition of all beamlet contributions,

$$d_s^i := \sum_{j=1}^n (a_s)_j^i x_j = \langle \mathbf{a}_s^i, \mathbf{x} \rangle \quad (6.3)$$

where $\langle \cdot, \cdot \rangle$ is the Euclidean inner product and \mathbf{a}_s^i is the i -th row vector of A_s . The column vector containing the dose in each voxel of structure s is given by

$$\mathbf{d}_s = A_s \mathbf{x}. \quad (6.4)$$

6.2.1 Hard dose constraints (HDCs)

For each avoidance structure $s \in \{1, \dots, S_a\}$,

$$\mathbf{0} \leq \mathbf{d}_s \leq \mathbf{u}_s, \quad (6.5)$$

where the i -th element of $\mathbf{u}_s = (u_s^i)_{1 \leq i \leq m_s}$ is the maximum dose permitted for voxel i of structure s . For each target structure $s \in \{S_a + 1, \dots, S_a + S_t\}$,

$$\mathbf{l}_s \leq \mathbf{d}_s \leq \mathbf{u}_s, \quad (6.6)$$

where the i -th element of $\mathbf{l}_s = (l_s^i)_{1 \leq i \leq m_s}$ is the minimum dose permitted for voxel i of structure s .

6.2.2 Dose volume constraints (DVCs)

For each avoidance structure $s \in \{1, \dots, S_a\}$,

$$\mathbf{0} \leq \mathbf{d}_s \leq (1 + \beta_s)\mathbf{u}_s, \quad (6.7)$$

$$\|(\mathbf{d}_s - \mathbf{u}_s)_+\|_0 \leq \alpha_s m_s \quad (6.8)$$

where the i -th element of $\mathbf{u}_s = (u_s^i)_{1 \leq i \leq m_s}$ is the maximum dose permitted for voxel i of structure s , but overflow of $(100\beta_s)\%$ is permitted in at most $(100\alpha_s)\%$ of the voxels. $\beta_s \in (0, 1)$ and $\alpha_s \in [0, 1]$. The sparsity norm operation $\|\cdot\|_0$ counts the number of positive, non-zero entries in a vector. For each target structure $s \in \{S_a + 1, \dots, S_a + S_t\}$,

$$(1 - \beta_s)\mathbf{l}_s \leq \mathbf{d}_s \leq \mathbf{u}_s, \quad (6.9)$$

$$\|(\mathbf{l}_s - \mathbf{d}_s)_+\|_0 \leq \alpha_s m_s \quad (6.10)$$

where the i -th element of $\mathbf{l}_s = (l_s^i)_{1 \leq i \leq m_s}$ is the minimum dose permitted for voxel i of structure s , but underflow of $(100\beta_s)\%$ is permitted in at most $(100\alpha_s)\%$ of the voxels. Here, $\beta_s \in (0, 1)$ and $\alpha_s \in [0, 1]$.

6.2.3 Problem statement: the non-convex split feasibility problem (SFP)

In the previous subsection we defined constraints which may be imposed on the absorbed dose matrix A_s for some structure s . These give rise to a system of linear inequalities in the case of (6.5)-(6.9) and a non-linear system in the case of (6.8) and (6.10). All structures and their linear constraints may be combined using block-matrices as follows. Let $A \in \mathbb{R}_+^{M \times n}$ be the matrix formed by blocks of all matrices A_s

where $M := \sum_{s=1}^{S_a+S_t} m_s$. Vectors $b, c \in \mathbb{R}^M$ are similarly defined such that each block-vector corresponds to linear constraints imposed on structure s . Written explicitly,

$$A := \begin{pmatrix} A_1 \\ A_2 \\ \vdots \\ A_{S_a} \\ -A_{S_a+1} \\ -A_{S_a+2} \\ \vdots \\ -A_{S_a+S_t} \end{pmatrix}, \quad b := \begin{pmatrix} \mathbf{v}_1 \\ \mathbf{v}_2 \\ \vdots \\ \mathbf{v}_{S_a} \\ -\mathbf{v}_{S_a+1} \\ -\mathbf{v}_{S_a+2} \\ \vdots \\ -\mathbf{v}_{S_a+S_t} \end{pmatrix}, \quad c := \begin{pmatrix} \mathbf{0} \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \\ -\mathbf{u}_{S_a+1} \\ -\mathbf{u}_{S_a+2} \\ \vdots \\ -\mathbf{u}_{S_a+S_t} \end{pmatrix}, \quad (6.11)$$

where for all avoidance structures, $s \in \{1, 2, \dots, S_a\}$,

$$\mathbf{v}_s := \begin{cases} \mathbf{u}_s, & \text{if an HDC is applied to } s, \\ (1 + \beta_s)\mathbf{u}_s, & \text{if a DVC is applied to } s, \end{cases} \quad (6.12a)$$

$$(6.12b)$$

and for all target structures, $s \in \{S_a + 1, S_a + 2, \dots, S_a + S_t\}$,

$$\mathbf{v}_s := \begin{cases} \mathbf{l}_s, & \text{if an HDC is applied to } s, \\ (1 - \beta_s)\mathbf{l}_s, & \text{if a DVC is applied to } s. \end{cases} \quad (6.13a)$$

$$(6.13b)$$

The linear HDC constraints may now be expressed as the set

$$C := \{\mathbf{x} \in \mathbb{R}_+^n \mid \mathbf{c} \leq \mathbf{A}\mathbf{x} \leq \mathbf{b}\}. \quad (6.14)$$

We may instead define M hyperslabs in \mathbb{R}^n , namely

$$C_i := \{\mathbf{x} \in \mathbb{R}_+^n \mid c^i \leq \langle \mathbf{a}^i, \mathbf{x} \rangle \leq b^i\}, \quad i = 1, 2, \dots, M, \quad (6.15)$$

where \mathbf{a}^i is the i -th row of A , and b^i and c^i are the i -th components of \mathbf{b} and \mathbf{c} , respectively. We must now also address the non-linear constraints arising from DVC constraints in (6.8) and (6.10). Let $\Delta \subseteq \{1, 2, \dots, S_a, S_a + 1, \dots, S_t\}$ be the set




Figure 6.3: The possible choices of the vector of beamlet intensities \mathbf{x} is restricted to the set of linear constraints C . The SFP in (6.17) requires that for each structure s with a DVC applied to it, one must also restrict the dose vector $A_s \mathbf{x}$ to the possibilities in the set of non-linear constraints Q_s . Any vector \mathbf{x}^* in the set $\Gamma = C \cap A^{-1}(Q_s)$ will satisfy both the linear and the non-linear constraints and therefore be labelled a feasible solution.

containing the identification numbers of all structures to which a DVC is applied.

For each structure $s \in \Delta$ we define the non-convex set

$$Q_s := \begin{cases} \{\mathbf{d}_s \in \mathbb{R}^{m_s} \mid \|(\mathbf{d}_s - \mathbf{u}_s)_+\|_0 \leq \alpha_s m_s\}, & s \leq S_a, \\ \{\mathbf{d}_s \in \mathbb{R}^{m_s} \mid \|(\mathbf{1}_s - \mathbf{d}_s)_+\|_0 \leq \alpha_s m_s\}, & s > S_a. \end{cases} \quad (6.16a)$$

$$(6.16b)$$

The IMPT (or more generally, IMRT) inverse problem now becomes a multiple-set non-convex *split-feasibility problem* (SFP) in which we seek a vector of beamlet intensities, \mathbf{x}^* , that satisfies the set of linear constraints, C , while $A\mathbf{x}^*$ simultaneously satisfies all sets, Q_s , for $s \in \Delta$, of non-linear dose-space constraints. The problem statement may be written as follows:

$$\text{Find } \mathbf{x}^* \in \bigcap_{i=1}^M C_i \text{ such that } A_s \mathbf{x}^* \in Q_s \forall s \in \Delta. \quad (6.17)$$

6.2.4 Split-feasibility seeking methods for solving the SFP

Let us briefly regress from the context of treatment planning and phrase the SFP mathematically on only two nonempty closed convex sets C and Q . If A represents

a bounded linear transformation from one Hilbert space to another, $A : \mathcal{H}_1 \rightarrow \mathcal{H}_2$, and $C \subseteq \mathcal{H}_1$ and $Q \subseteq \mathcal{H}_2$ then the SFP is to find the element $\mathbf{x}^* \in C$ such that $A\mathbf{x}^* \in Q$ [428]. The space of all possible solutions is given by $\Gamma = C \cap A^{-1}(Q)$. In most purposes we restrict our domain to Euclidean space, where $C \subseteq \mathbb{R}^N$, $Q \subseteq \mathbb{R}^M$ and A is a real-valued matrix with dimensions $M \times N$. Censor and Elfving in their 1994 paper [429] first proposed iterative methods for solving the SFP which required computing matrix inverses at each step. The well-known CQ algorithm by Byrne [430, 431], which does not require matrix inverses, is now commonly used in signal processing and image reconstruction.

Byrne's CQ algorithm starts with an arbitrary initial estimate of the solution, $\mathbf{x}^{(0)}$, and for each successive iteration, $k \in \mathbb{N}$, the solution is updated via

$$\mathbf{x}^{(k+1)} = P_C \left(\mathbf{x}^{(k)} + \gamma A^T (P_Q - I) A \mathbf{x}^{(k)} \right), \quad (6.18)$$

where I is the identity operator ($I(\mathbf{z}) = \mathbf{z}$ for any element \mathbf{z}), A^T is the matrix transpose of A and $\gamma \in (0, 2/L)$. L is the largest eigenvalue of the matrix $A^T A$. The operator P_Ω denotes the orthogonal projection of a point onto a set Ω .

In the case of treatment planning with a single DVC, Penfold *et al.* [378] proposed an algorithmic framework using both Byrne's CQ algorithm and an automatic relaxation method that uses cyclic projections. In order to include multiple DVCs one may follow the work of Wang and Xu [432] which details a cyclic algorithm for solving the multiple-set split feasibility problem as defined by Censor and Elfving [429].

In the context of our treatment planning problem with multiple DVCs we may use Algorithm 1, based on the work of Wang and Xu. Alternatively, making use a cost function, the SFP in (6.17) could be rephrased as a minimisation problem as follows:

Find $\mathbf{x}^* \in \mathbb{R}_+^n$ which minimises

$$f(\mathbf{x}) = \frac{1}{2} \left\{ \sum_{i=1}^M \|\mathbf{x} - P_{C_i}(\mathbf{x})\|^2 + \sum_{s \in \Delta} \|A_s \mathbf{x} - P_{Q_s}(A_s \mathbf{x})\|^2 \right\}. \quad (6.21)$$

In either methodology, one must compute the projections onto the hyperslabs C_i , defined in (6.15), and onto the non-convex sets Q_s , defined in (6.16a) and (6.16b). If an element $\mathbf{z} \in \mathbb{R}^n$ is not within C_i then it is projected orthogonally onto the nearest

Algorithm 1 Fully sequential dynamic string-averaging CQ-method

```

1: Set an arbitrary  $\mathbf{x}^{(0)} \in \mathbb{R}^n$ .
2: Set maximum number of allowed iterations  $k_{\max}$ .
3: Set voxel iteration index  $k = 0$ .
4: while  $k < k_{\max}$  do
5:    $i(k) \leftarrow (k \bmod M) + 1$  ▷ Cyclic control over rows of  $A$ .
6:   Get structure index  $s(i(k))$ .
7:   ▷  $s(i(k))$  from lookup table.
8:   if  $s(i(k)) \in \Delta$  then
9:     
$$\mathbf{x}^{(k+1)} \leftarrow P_{C_{i(k)}} \left( \mathbf{x}^{(k)} - \gamma A_{s(i(k))}^T (I - P_{Q_{s(i(k))}}) A_{s(i(k))} \mathbf{x}^{(k)} \right), \quad (6.19)$$

10:    ▷  $\gamma$  may take any value in  $(0, 2/\|A_{s(i(k))}\|^2)$ .
11:    ▷  $\|A_{s(i(k))}\|^2$  is the Frobenius matrix norm.
12:   end if
13:   
$$\mathbf{x}^{(k+1)} \leftarrow P_{C_{i(k)}}(\mathbf{x}^{(k)}). \quad (6.20)$$

14:   for all elements  $(x^{(k+1)})_j$  in  $\mathbf{x}^{(k+1)}$  do
15:     if  $(x^{(k+1)})_j < 0$  then
16:       Set  $(x^{(k+1)})_j \leftarrow 0$ .
17:     end if
18:   end for
19:    $k \leftarrow k + 1$ 
20: end while

```




Figure 6.4: Orthogonal projection $P_{\Omega}(\mathbf{z})$ of an element \mathbf{z} onto the hyperplane H_1 , which defines a boundary of Ω . In the case that \mathbf{z} is above the hyperplane H_1 , the projection is made onto H_1 . The relaxation parameter λ allows the projection to (a) fall on H_1 exactly, (b) overshoot H_1 , or (c) undershoot H_1 .

hyperplane which defines a boundary of C_i . Otherwise, no action is taken. That is,

$$P_{C_i}(\mathbf{z}) = \begin{cases} \mathbf{z}, & \mathbf{c}^i \leq \langle \mathbf{a}^i, \mathbf{z} \rangle \leq \mathbf{b}^i, & (6.22a) \\ \mathbf{z} + \lambda \frac{\mathbf{c}^i - \langle \mathbf{a}^i, \mathbf{z} \rangle}{\langle \mathbf{a}^i, \mathbf{a}^i \rangle} \mathbf{a}^i, & \langle \mathbf{a}^i, \mathbf{z} \rangle < \mathbf{c}^i, & (6.22b) \\ \mathbf{z} + \lambda \frac{\mathbf{b}^i - \langle \mathbf{a}^i, \mathbf{z} \rangle}{\langle \mathbf{a}^i, \mathbf{a}^i \rangle} \mathbf{a}^i, & \langle \mathbf{a}^i, \mathbf{z} \rangle > \mathbf{b}^i, & (6.22c) \end{cases}$$

where $\lambda \in (0, 2)$ is a user-chosen relaxation parameter allowing the projection to over- or under-shoot the hyperplane instead of falling exactly onto it. This is demonstrated in Figure 6.4.

The projection onto the set Q_s follows a different recipe. First, a translation is made to simplify the projection. For a generic element, $\mathbf{w} \in \mathbb{R}^{m_s}$,

$$P_{Q_s}(\mathbf{w}) := \begin{cases} P_{\bar{Q}_s}(\mathbf{w} - \mathbf{u}_s) + \mathbf{u}_s, & s \leq S_a, & (6.23a) \\ -P_{\bar{Q}_s}(\mathbf{l}_s - \mathbf{w}) + \mathbf{l}_s, & s > S_a, & (6.23b) \end{cases}$$

where

$$\bar{Q}_s := \{\mathbf{y} \in \mathbb{R}^{m_s} \mid \|\mathbf{y}_+\|_0 \leq \alpha_s m_s\}. \quad (6.24)$$

We have translated our problem from projecting $\mathbf{w} \in \mathbb{R}^{m_s}$ onto Q_s to projecting the shifted point $\mathbf{w}' \in \mathbb{R}^{m_s}$ onto \bar{Q}_s . First we count the number of positive entries in the vector \mathbf{w}' , denoted $l = \|\mathbf{w}'_+\|_0$. If $l \leq \alpha_s m_s$ then \mathbf{w}' is already in the set \bar{Q}_s and no action is taken. However, if $l > \alpha_s m_s$ then $P_{\bar{Q}_s}(\mathbf{w}')$ replaces the $[l - \alpha_s m_s]$ smallest positive components of \mathbf{w}' with zeros and leaves the other components unchanged.

6.2.5 Further insight into Byrne’s CQ formula

In the context of treatment planning, Byrne’s formula ((6.18)) is used to iteratively update the vector of beamlet intensities, \mathbf{x} , by using projections onto the constraint sets. Here, we simplify the IMRT problem by considering only one DVC set, Q , and one structure, whose dose-influence matrix is denoted A . Byrne has proven convergence in the general case³ [430], but why should one expect it to converge? We explain here three major components of the procedure:

$$\mathbf{x}^{(k+1)} = P_C \left(\underbrace{\mathbf{x}^{(k)} + \gamma \underbrace{A^T \left(\underbrace{P_Q(A\mathbf{x}^{(k)}) - A\mathbf{x}^{(k)}}_{\textcircled{1}} \right)}_{\textcircled{2}}}_{\textcircled{3}} \right), \quad (6.25)$$

- ① The operation P_Q involves comparing the dose vector, $A\mathbf{x}$, to the HDC (minimum or maximum permitted dose) imposed on the structure represented by A . Any voxel containing dose below the minimum or above the maximum is violating the HDC. The projection operator, (6.23a) or (6.23b), “removes” the smallest violations by setting the dose to the HDC value. This process can be




Figure 6.5: In the case of 6 voxels representing an OAR, if we allow the dose in one third of them to overflow, then the operator P_Q takes the lowest violation (voxel 3) and decreases the contained dose to the prescribed maximum.

³Note that Byrne’s proof requires static, pre-chosen relaxation parameters, λ and γ .

seen more clearly in the example in Figure 6.5.

- ② The difference between the new dose distribution and the previous is given by $P_Q(A\mathbf{x}) - A\mathbf{x}$. We wish to infer from this the difference between the intensity vectors needed to achieve these dose distributions. We could use the matrix inverse, A^{-1} , to transform back to the intensities space, but computing the inverse is inefficient. Instead, we may use the fact that $A^T A$ is an orthogonal matrix, so the transformation $\mathbf{x} \mapsto A^T A\mathbf{x}$ preserves lengths. That is, if $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ and A has n rows then⁴

$$(A^T A\mathbf{x} - A^T A\mathbf{y})^T (A^T A\mathbf{x} - A^T A\mathbf{y}) = (\mathbf{x} - \mathbf{y})^T (\mathbf{x} - \mathbf{y}). \quad (6.26)$$

- ③ The bracketed term in (6.25) is the new estimate of beamlet intensities (weighted by γ) after trying to satisfy the non-linear constraints in Q . This estimate is now projected onto the linear constraint set C , giving a compromise between the two sets of constraints. This entire process is repeated over many iterations.

6.2.6 Dynamic self-adaptive relaxation parameter methods

In the previous subsection we explored how we may seek a solution to the SFP through the use of iterative projection methods, namely (6.19), and (6.20) in Algorithm 1, the generalised extension of Byrne's CQ algorithm (6.18) [430, 431] based on the work of Wang and Xu [432]. Recall that these iterative methods involve projections onto the sets C_i for $i = 1, 2, \dots, M$, and Q_s for $s \in \Delta$. The constraints imposed by these sets can be categorised as follows.

- The sets C_i arise from constraining the maximum and minimum permitted dose in voxel i . Note that this includes any permitted over- and under-flow. These sets encode the HDCs.
- The sets Q_s arise from limiting the number of voxels in structure s that are

⁴The proof of (6.26) is straightforward using $(A^T A)^T (A^T A) = I$ where I is the identity.

permitted to have dose larger than the specified maximum or smaller than the specified minimum. These sets encode the DVCs.

Projections onto C_i are “relaxed” with $\lambda \in (0, 2)$ and projections onto Q_s are relaxed with $\gamma \in (0, 2/\|A_{s(i(k))}\|^2)$. These values are called *relaxation parameters* as they allow a projection to fall within close proximity of the set without necessitating that it fall exactly on the boundary. It is common practice to choose an arbitrary λ and γ prior to the first iteration and retain the same values through all iterations. However the possibility of using dynamic parameters $\lambda^{(k)}$ and $\gamma^{(k)}$, which may be updated on each iteration, k , has been explored [433, 434].

A method is proposed here which monitors the percentage of voxels violating their constraints and accordingly updates structure specific $\lambda_s^{(k)}$ and $\gamma_s^{(k)}$ parameters to shift the attention of the algorithm to satisfying the constraints which are most violated. We shall refer to this as the method of *Self-Adaptive Relaxation Parameters* (SARP). Using SARP, (6.19) is replaced by

$$\mathbf{x}^{(k+1)} = P_{C_{i(k)}} \left(\mathbf{x}^{(k)} - \gamma_{s(i(k))}^{(k)} A_{s(i(k))}^T (I - P_{Q_{s(i(k))}}) A_{s(i(k))} \mathbf{x}^{(k)} \right), \quad (6.31)$$

where the $P_{C_{i(k)}}$ projections use $\lambda_{s(i(k))}^{(k)}$ in place of λ in (6.22a)-(6.22c).

We propose here two variants of the SARP method. In both versions, the relaxation parameters, λ_s and γ_s , are initially set to their mid-range. After each algorithmic cycle of the Dynamic String-Averaging CQ-Method (DSACQ), meaning one pass through all imposed constraints, the number of voxels violating each type of constraint are counted. After some number of consecutive cycles (e.g. 10) it is assessed whether any of the constraint violations are starting to stabilise; that is, the number of violations does not significantly decrease from one iteration to the next. If this is the case then the relaxation parameters are updated.

One version of the algorithm, SARP1 (Algorithm 2), uses a “bisection-like” method to update the parameters while another, SARP2 (Algorithm 3), assigns exact values based on the ratio of the number of violations between the HDCs and the DVCs. Both of these variations of the SARP method encourage over-relaxation in the projections

onto the sets representing the least satisfied constraints and under-relaxation in those onto the constraint sets that are most satisfied.

The ultimate purpose of any parameter-update scheme would be to shift the attention of the algorithm onto the objectives⁵ that need the most improvement.

Algorithm 2 Self-adaptive Relaxation Parameter method: version 1 (SARP1)

Usage: If called at the k -th iterate of the Dynamic String-Averaging CQ-Method (Algorithm 1), SARP1 executes the following:

```

1: for each structure  $s$  with a DVC do
2:   Count elements of  $A_s x^{(k)}$  for which  $x^{(k)}$  violates HDC constraints
3:   Set as  $\eta_C^s$ 
4:   Count elements of  $A_s x^{(k)}$  which violate DVC constraints
5:   Set as  $\eta_Q^s$ 
6:   if  $\eta_C^s > \eta_Q^s$  then
7:      $\triangleright$  There are more HDC violations than DVC violations
8:      $\triangleright$  Place more emphasis on HDC projections
9:

```

$$\gamma_s^{(k+1)} \leftarrow \frac{1}{2} \gamma_s^{(k)} \quad (6.27)$$

$$\lambda_s^{(k+1)} \leftarrow \frac{1}{2} \lambda_s^{(k)} + \frac{1}{2} \lambda_{s,\max} \quad (6.28)$$

```

10:  else if  $\eta_Q^s > \eta_C^s$  then
11:     $\triangleright$  There are more DVC violations than HDC violations
12:     $\triangleright$  Place more emphasis on DVC projections

```

$$\lambda_s^{(k+1)} \leftarrow \frac{1}{2} \lambda_s^{(k)} \quad (6.29)$$

$$\gamma_s^{(k+1)} \leftarrow \frac{1}{2} \gamma_s^{(k)} + \frac{1}{2} \gamma_{s,\max} \quad (6.30)$$

```

13:  else
14:    Do nothing.
15:  end if
16: end for

```

⁵The objectives here are to minimise the number of voxels that violate a certain type of constraint.

Algorithm 3 Self-adaptive Relaxation Parameter method: version 2 (SARP2)

Usage: If called at the k -th iterate of the Dynamic String-Averaging CQ-Method (Algorithm 1), SARP2 executes the following:

- 1: **for** each structure s with a DVC **do**
- 2: Count elements of $A_s x^{(k)}$ for which $x^{(k)}$ violates HDC constraints
- 3: Set as η_C^s
- 4: Count elements of $A_s x^{(k)}$ which violate DVC constraints
- 5: Set as η_Q^s

$$\lambda_s^{(k+1)} \leftarrow \frac{\eta_C^s}{\eta_C^s + \eta_Q^s} \lambda_{s,\max} \quad (6.32)$$

$$\gamma_s^{(k+1)} \leftarrow \frac{\eta_Q^s}{\eta_C^s + \eta_Q^s} \gamma_{s,\max} \quad (6.33)$$

- 6: **end for**
-

6.3 Discussion

This work provides a proof of concept for a projection-based IMRT inverse planning algorithm, called the *dynamic string-averaging CQ-method (DSACQ)*, that is mathematically robust to the order and weighting with which projections are executed, even in non-convex settings imposed by multiple DVCs. A self-adaptive relaxation parameter (SARP) method has been proposed to continually monitor constraint violations and shift the attention of the algorithm to better meet those constraints that are most violated. This is achieved by dynamically adjusting the amount of over- or under-relaxation of projections onto individual constraint sets. Unlike cost-minimisation algorithms, the DSACQ method uses split-feasibility seeking techniques to obtain a *feasible* dose solution, based on the imposed constraints, rather than claiming optimality of the solution, which is both subject to clinical interpretation and unable to be mathematically guaranteed using a non-differentiable and non-convex cost function. The DSACQ method is able to exactly model DVC prescriptions, without the need for linear approximations, and avoid impractical combinatorial optimisation methods through the use of an ℓ_0 -sparsity constraint applied to sets representing the DVCs.

The DSACQ algorithm has been presented here in its practical medical physics

context. However, its more generalised formulation has been outlined in a mathematical paper [417] in which convergence properties have been rigorously analysed⁶. Though beyond the scope of this thesis, it may be of interest to the reader to recognise that the IMRT inverse problem there is formulated in terms of nonempty fixed-point sets defined on *Landweber-type* operators (see, e.g., [435, 436, 437, 438]). The applicability of convergence properties discussed in [417] to the SARP methods detailed here (Algorithm 2 and Algorithm 3) is briefly visited in Appendix D. In this Chapter the DSACQ is presented, for the sake of simplicity, as a fully-sequential algorithm with unit weights. That is, the solution vector, \mathbf{x} , is successively updated in a sequence according to individual projections onto each structure and each voxel. Any prioritisation between structures or constraints occurs only through changing the relaxation parameters, and not through premultiplying weights to the update formulae in (6.19) or (6.20). Though, in practice, additional flexibility is gained by varying weights and implementing string-averaged projections (see, e.g., [439, 440, 441, 442]). Figure 6.6 illustrates an example in which fully-simultaneous *block-iterative* (or “*Cimmino*”) projections differ from fully-sequential (or “*Kaczmarz*”) projections, performed with four hyperplanes. The shaded region in Figure 6.6(b) represents the possible solu-




Figure 6.6: Examples of (a) fully-sequential and (b) fully-simultaneous projections onto four hyperplanes, H_i for $i \in \{1, 2, 3, 4\}$. The current solution vector $\mathbf{x}^{(k)}$ is projected onto the four hyperplanes in order to determine the next iterate(s), $\mathbf{x}^{(k+1)}$.

⁶With many thanks and credit to co-authors Yair Censor and Aviv Gibali.

tion space for the next iterate of the solution, $\mathbf{x}^{(k+1)}$, which depends on the weights assigned to each individual projection. Block-iterative methods are advantageous for code parallelisation and allow for shorter run times, however one must be careful in the length of strings averaged and the weights used, as the rate of convergence may generally worsen.

Reducing computational time in the feasibility-seeking procedure, possibly through use of SARP methods (explored in Chapter 7) or running on graphics processing units (GPUs), for example, opens up the possibility of flooding target volumes with a large number of pencil beams. Therefore, the task of the clinician in setting up the treatment plan prior to the FMO step becomes simpler. However, this could lead to an inefficient over-parameterisation of the beamlet intensities. In this case, an ℓ_1 -regularised objective function [443] could be defined such that the linear HDCs are imposed while keeping the number of beamlets used to a minimum. From here, the usual feasibility-seeking procedure could be used on only the portion of the initial beamlets with non-zero weights. In photon IMRT the dose resulting from a continuously rotating beam with a variable exposure rate may be approximated by a set of fixed beams oriented at small angular gradations [444]. This, in theory, opens possibilities for fully-automated planning of volumetric modulated arc therapy (VMAT).

Tractability of the algorithm could be improved by replacing ℓ_0 -norm constraints with an ℓ_1 -norm convex envelope. In such a setting, one could redefine the DVC projections of the dose. For example, for a DVC on an OAR with maximum permitted dose vector, $\mathbf{u} \in \mathbb{R}^{m_s}$, one could set

$$P_{\bar{Q}_s}(A_s \mathbf{x}) := \inf\{\mathbf{0}, A_s \mathbf{x} - \mathbf{u}\} + P_{\bar{Q}_s}(\sup\{\mathbf{0}, A_s \mathbf{x} - \mathbf{u}\}) + \mathbf{u} \quad (6.34)$$

where

$$P_{\bar{Q}_s}(\mathbf{y}) := \arg \min_{\bar{q} \in \bar{Q}_s} \|\mathbf{y} - \bar{q}\|_2^2 \quad (6.35)$$

and

$$\bar{Q}_s := \{\mathbf{y} \in \mathbb{R}_+^{m_s} \mid \|\mathbf{y}\|_1 \leq \alpha m_s\}. \quad (6.36)$$

Note that Q_s remains non-convex and \bar{Q}_s no longer represents a strict dose volume constraint. Instead of counting the number of voxels violating their constraints with $\|\cdot\|_0$, we now measure the integral amount of dose above the threshold using $\|\cdot\|_1$. To address this, one could apply a weighting factor to the upper bound, αm_s , to better reflect the integral dose (see, e.g., [445]), or one could apply a “sparsifying” basis. An example of such a procedure for ℓ_1 -minimisation is provided below:

$$\min_{\mathbf{x} \in \mathbb{R}^n} \|W^{(k)} y^{(k)}\|_1 = \sum_j w_j^{(k)} \cdot |y_j^{(k)}| \quad (6.37)$$

subject to $y_j \geq 0 \forall j$. The weighting factors w_j are set to unity on the first iteration, and then updated for the k -th iterate according to

$$w_j^{(k)} = \frac{1}{|y_j^{(k)}| + \epsilon^{(k)}}, \quad (6.38)$$

where $\epsilon^{(k)} > 0$ is small. This method has been explored for improving the efficiency of dose delivery in IMRT by maximising sparsity in the vector of beamlet intensities [446]. The authors also explored total variation (TV) minimisation to promote sparsity in fluence-map variations. This concept could be promising for implementing dose-homogeneity constraints in the feasibility-seeking procedure described in this Chapter. For instance, it may be possible to simultaneously impose constraints on the Fourier space transform of the dose vector for a given structure. In this way, peaks corresponding to steep dose gradients within the structure (away from the boundary) could be reduced, thereby enforcing dose homogeneity within the structure.

In the following Chapter, we shift focus from theoretical development of the DSACQ algorithm to its practical implementation and analyse its performance in a variety of synthetically-generated test cases as well as archived clinical plans.

Chapter 7

Implementation and testing of a novel inverse planning algorithm

7.1 Introduction

In Chapter 6 we examined the inverse problem of radiation therapy treatment planning in terms of matrix algebra and developed a formalism in which dose volume constraints (DVCs) can be exactly represented by a sparsity norm constraint, with no convex relaxation necessary. The problem statement for inverse planning was presented as a non-convex split feasibility problem, to which we sought a solution based on iterative projection methods inspired by Byrne’s well-known CQ method [430, 431]. A new algorithm was proposed, called the *dynamic string-averaging CQ-method (DSACQ)*, which allows flexibility in the weighting and order in which projections onto constraint sets are executed, and in the number of non-convex DVCs imposed. The former gives rise to a variety of methods that have the potential to improve convergence to a feasible dose solution, while the latter is essential for fulfilling modern treatment planning objectives. Dynamic parameter-update methods were posed as an extension to the algorithm. These methods dynamically adjust projection-related parameters to shift the activity of the algorithm to the constraints that are most violated. Two possible automated implementations were detailed under the title of *self-adaptive relaxation parameter (SARP)* methods.

In this Chapter we focus on the implementation of the DSACQ algorithm, including its practicality, performance and translation into clinical use. Convergence

behaviour of the DSACQ algorithm is evaluated in a variety of settings and performance of the SARP dynamic parameter-update methods are compared. This analysis begins with two rudimentary synthetic plans created on two-dimensional pseudo-dose grids, progresses to a more realistic three-dimensional plan with only two structures, and then to full archived head-and-neck treatment plans containing up to 23 structures. Finally, code acceleration is investigated to ascertain whether the DSACQ algorithm can be quickly implemented in clinical time limitations.

7.2 Methods

We detail here a practical implementation of the fully sequential DSACQ method (Algorithm 1), which was applied in *MATLAB* version R2019a (The MathWorks, Inc.) for all test cases that follow. First, or in the zero-th iteration of the algorithm, beamlet weights are set to unit intensity. That is, $\mathbf{x}^{(0)} = (1, 1, \dots, 1)^T$, before running through an iterative scheme described step-by-step by the pseudo-code in Algorithm 4 (also summarised as a flow chart in Figure 7.1). Note that in the theoretical formulation of the inverse problem in Chapter 6 it was most suitable to work in the framework of block-matrices, where the full dose-influence matrix A was constructed of structure-specific submatrices, A_s , for S_a avoidance structures and S_t target structures. In this

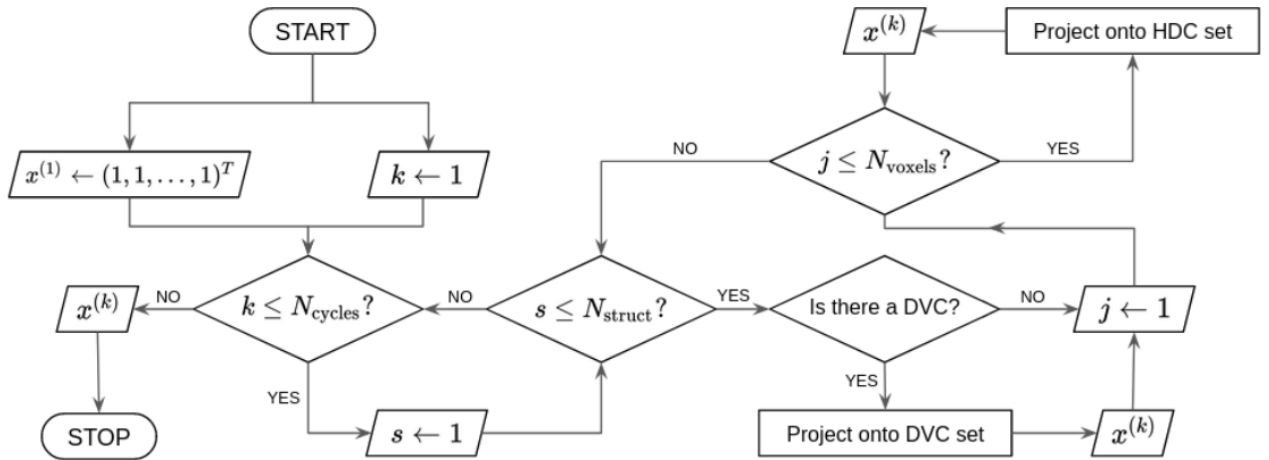


Figure 7.1: Flowchart of Algorithm 4. N_{cycles} is the number of algorithmic iterations, N_{struct} is the number of structures, and N_{voxels} is the number of voxels comprising the current structure. Counter increments are implied within each loop, with indices starting at 1 by default.

Algorithm 4 Pseudo-code example for implementation of the DSACQ algorithm

```

1: Initialise  $\mathbf{x}^{(0)} = (1, 1, \dots, 1)^T$ .
2: Initialise cycle number  $k = 1$ .
3: Choose max cycles  $N_{\text{cycles}}$ .
4: while  $k < N_{\text{cycles}}$  do
5:   for  $s \in \{1, 2, \dots, S_a, S_a + 1, \dots, S_a + S_t\}$  do
6:     if  $s \in \Delta$  then
7:       Choose some  $0 < \gamma_s^{(k)} < 2/\|A_s\|^2$ .
8:        $\mathbf{x}^{(k)} \leftarrow \mathbf{x}^{(k)} - \gamma_s^{(k)} A_s^T (A_s \mathbf{x}^{(k)} - P_{Q_s}(A_s \mathbf{x}^{(k)}))$ 
9:     end if
10:    for  $i \in \{1, 2, \dots, m_s\}$  do
11:       $\mathbf{x}^{(k)} \leftarrow P_{C_s^i}(\mathbf{x}^{(k)})$   $\triangleright$  Free choice of  $0 < \lambda_s^{(k)} < 2$  within  $P_{C_s^i}$ .
12:    end for
13:     $\mathbf{x}^{(k+1)} \leftarrow \mathbf{x}_+^{(k)}$   $\triangleright$  Enforce nonnegativity constraint.
14:     $k \leftarrow k + 1$ 
15:  end for
16: end while

```

practical implementation the matrix A need not be numerically evaluated; each A_s matrix may be treated independently. This is, programmatically, both simpler and more efficient than storing a very large matrix in computer memory. Now, in order to accommodate this method, a small change to the definition of the hyperslabs C_i in (7.1) is necessary. Without loss of generality, we define, for each structure s ,

$$C_s^i := \{\mathbf{x} \in \mathbb{R}_+^n \mid c_s^i \leq \langle \mathbf{a}_s^i, \mathbf{x} \rangle \leq b_s^i\}, \quad i = 1, 2, \dots, m_s, \quad (7.1)$$

where \mathbf{b}_s and \mathbf{c}_s are the s -th block-vectors of the matrices \mathbf{b} and \mathbf{c} originally defined in (6.11), and m_s is the number of voxels contained in structure $s \in \{1, 2, \dots, S_a, S_a + 1, \dots, S_a + S_t\}$. With this adjusted definition of the hyperslabs representing the hard dose constraints (HDCs), the orthogonal (metric) projections onto them, $P_{C_s^i}$, are defined just as they were in (6.22a)–(6.22c) and the problem statement of (6.17) may be rewritten as

$$\text{Find } \mathbf{x}^* \in \bigcap_{s=1}^{S_a+S_t} \bigcap_{i=1}^{m_s} C_s^i \text{ such that } A_s \mathbf{x}^* \in Q_s \quad \forall s \in \Delta. \quad (7.2)$$

In Algorithm 4 there are two “for” loop control cycles beginning on Line 5 and

Line 10. These are written such that the structures may be chosen in any order, without replacement, and so may the voxels within each structure. The exact control sequences may be freely chosen by the user to suit their application, however in this work each cycle of the algorithm invokes a randomly generated ordering. In this way, possible bias, introduced by always considering one structure before another, can be avoided. Here, a preset number of cycles, N_{cycles} , are performed before terminating the iterative method and assessing the solution. The user may choose to replace this with a tolerance-based stopping criterion.

7.2.1 Two-dimensional numerical examples

7.2.1.1 Test Case 1

A two-dimensional 512×512 pixel grid, representing 262,144 pixels, was set up as a region of dosimetric interest. In a clinical treatment plan this would represent a patient geometry and these pixels would be replaced by a large number of three-dimensional voxels. Two spatial dimensions are assumed for this plan for the sake of simplicity, though the problem may be easily generalised to three dimensions. A total of 1,156 uniformly distributed Gaussian “dose” kernel contributions were created to emulate radiation dose and cover the entire grid. Each kernel was given a standard deviation of 20 pixels and an amplitude such that a homogeneous flood dose was achieved with a mean value of 50 units, as shown in Figure 7.2(a, b). From this point forward it shall be assumed that pixel values correspond to “dose”.

Thus far, 1,156 different matrices of dimensions 512×512 have been introduced, each corresponding to a separate dose kernel. In order to form a *dose-influence* matrix, A , for use in inverse planning, each matrix must be collapsed to a single column vector with 262,144 entries while keeping track of which indices corresponded to which spatial positions in the dose grid. The matrix A is formed by all column vectors and therefore has 262,144 rows and 1,156 columns. A prescription comprised of four HDCs and three DVCs¹, shown in Table 7.1, was applied to three arbitrarily

¹DVCs in Table 7.1 are written in the standard notation, $D_{V\%}$, which is the dose that is received by exactly $V\%$ of the structure.

defined square regions representing two avoidance structures, “Avoidance A” and “Avoidance B”, and a target structure, “Target”. Pixel indices corresponding to each of these structures were used to extract the structure-specific block-matrices A_1 , A_2 and A_3 , respectively. The vector of beamlet intensities was initialised to $\mathbf{x}^{(0)} = (1, 1, \dots, 1)^T$ with 1,156 entries.

Forty cycles ($N_{\text{cycles}} = 40$) of Algorithm 4 were applied to the problem described above in order to reduce the dose in the avoidance structures and elevate it in the target structure, according to the prescription in Table 7.1. In this initial test of the DSACQ algorithm, the relaxation parameters, λ_s and γ_s , for each structure were set to their midrange values, 1 and $\|A_s\|^{-2}$, respectively. That is, projections were executed exactly onto the hyperplanes defined by the constraints, with no under- or over-relaxation. Explicitly, $\lambda_1 = \lambda_2 = \lambda_3 = 1$, $\gamma_1 = 1.546 \times 10^{-6}$, $\gamma_2 = 1.545 \times 10^{-6}$, and $\gamma_3 = 1.030 \times 10^{-6}$.

7.2.1.2 Test Case 2

A similar problem to Test Case 1, described above, was set up on a 100×100 pixel grid with 484 Gaussian dose contributions, each with a standard deviation of 12 pixels, giving a mean dose of 50 units (see Figure 7.3). A coarser grid than that used in Test Case 1 encourages a larger degree of overlapping dose contributions. Additionally, the




Figure 7.2: (a) A single Gaussian “dose” kernel contribution shown at one grid point. (b) Homogeneous flood dose of 50 units formed by superimposing all 1156 contributions. (c) The optimised dose map showing the structures for which the Test Case 1 prescription in Table 7.1 was applied.




Figure 7.3: (a) Flood dose, with mean of 50 units, formed by superimposing 484 Gaussian dose contributions. (b) The optimised dose map showing the structures for which the Test Case 2 prescription in Table 7.1 was applied.

target prescription was split over two square regions with more challenging constraints (see Table 7.1). The DSACQ algorithm was then applied using both variations of the dynamic parameter-update method, SARP1 (Algorithm 2) and SARP2 (Algorithm 3), as well a static parameter method in which the relaxation parameters are set to their mid-range for all projections.

7.2.2 Simplified base-of-skull chordoma

The past two test cases have been simplified to two spatial dimensions with Gaussian pseudo-dose contributions, which have no resemblance to pencil beam doses encountered in proton therapy. In this next case, we progress to real proton dose contributions from a simplified version of an archived head-and-neck radiotherapy plan. We have used the 3D base-of-skull chordoma treatment plan originally used in the 2017 paper by Penfold *et al.* [378] which introduced the *Dose Volume Split Feasibility* (DVSF) algorithm inspiring the current work. Only two structures are present; an organ-at-risk (OAR) and a planning target volume (PTV), to be avoided and covered, respectively, as best as practically achievable. The only OAR in this plan is the brainstem. The DSACQ algorithm was run for 2000 cycles using both example SARP methods and a static parameter method.

Table 7.1: Prescriptions chosen for all test cases.

Plan	Structure	HDCs	DVCs
Test Case 1	Avoidance A	$D_{\max} = 25$	$D_{10\%} \leq 20$
	Avoidance B	$D_{\max} = 40$	$D_{25\%} \leq 30$
	Target	$D_{\min} = 60$	$D_{90\%} \geq 65$
		$D_{\max} = 70$	
Test Case 2	Avoidance A	$D_{\max} = 54$	$D_{7\%} \leq 50$
	Avoidance B	$D_{\max} = 48$	$D_{5\%} \leq 45$
	Target A	$D_{\min} = 62$	$D_{90\%} \geq 65$
		$D_{\max} = 70$	
	Target B	$D_{\min} = 60$	$D_{85\%} \geq 62$
		$D_{\max} = 65$	
Test Case 3	Brainstem	$D_{\max} = 54$	$D_{5\%} \leq 50$
	PTV	$D_{\min} = 74.9$	$D_{95\%} \geq 70$
		$D_{\max} = 66.5$	
TROTS Case 1	Parotid (left)	$D_{\max} = 6.20$	$D_{50\%} \leq 0.05$
	Parotid (right)	$D_{\max} = 56.60$	$D_{50\%} \leq 3.05$
	SMG (left)	$D_{\max} = 24.60$	$D_{50\%} \leq 3.05$
	SMG (right)	$D_{\max} = 68.90$	$D_{50\%} \leq 28.65$
	Spinal Cord	$D_{\max} = 20.80$	$D_{50\%} \leq 0.05$
	Brainstem	$D_{\max} = 19.70$	$D_{50\%} \leq 0.05$
	SCM	$D_{\max} = 68.40$	$D_{50\%} \leq 33.55$
	MCM	$D_{\max} = 54.40$	$D_{50\%} \leq 2.25$
	MCI	$D_{\max} = 48.80$	$D_{50\%} \leq 0.25$
	MCRico	$D_{\max} = 8.90$	$D_{50\%} \leq 0.45$
	Oesophagus	$D_{\max} = 6.70$	$D_{50\%} \leq 0.55$
	Larynx	$D_{\max} = 55.00$	$D_{50\%} \leq 1.15$
	Oral Cavity	$D_{\max} = 69.90$	$D_{50\%} \leq 4.35$
	Patient	$D_{\max} = 69.80$	$D_{50\%} \leq 0.05$
	CTV High	$D_{\min} = 64.70$	$D_{50\%} \geq 67.45$
		$D_{\max} = 70.10$	
	CTV Intermediate	$D_{\min} = 54.50$	$D_{50\%} \geq 59.45$
		$D_{\max} = 70.00$	
	CTV Low Shrunk	$D_{\min} = 53.00$	$D_{50\%} \geq 55.05$
$D_{\max} = 57.30$			

As detailed in Section 2.5.2 of the work by Penfold *et al.* [378], the *Pinnacle* treatment planning system (Philips Healthcare, Koninklijke Philips N.V.) was originally used to contour the PTV and brainstem (OAR) before the DICOM RT structure files were imported into a *MATLAB* script. Contours were mapped over the CT coordinates and a dose grid was created in *MATLAB* to match that defined in *Pinnacle*. The dose grid dimensions were $42 \times 43 \times 9$ voxels with resolutions in the x , y and z

dimensions of 2 mm, 2 mm and 3 mm, respectively. The x - y pixels of the grid were twice as large as those in the CT images. Two IMPT fields were defined in a research release of *Pinnacle* for proton pencil beam scanning, containing 574 and 564 pencil beams, at angles of 80° and 280° , respectively. Proximal PTV coverage was ensured using a range shifter of 7.5 cm thickness for both fields. Pencil beams were set to 80% layer overlap, 0.6 cm lateral spot resolution, 0.4 cm lateral target margin, and 3 standard deviation dose spread during dose calculation. The dose grids corresponding to unit intensity of each proton pencil beam were exported, after being calculated using the *Pinnacle* analytical pencil beam scanning algorithm.

7.2.3 The TROTS proton therapy data set

The Radiotherapy Optimisation Test Set (TROTS²) contains 20 proton beam therapy (PBT) treatment plans for patients with head and neck cancer that were treated with 3-beam intensity-modulated proton therapy (IMPT) [447]. These PBT plans were created from re-purposed IMRT plans for oropharyngeal cancer patients [448]. The data includes the agreed-upon final dose solution for comparison. One intention of the data set is to measure the performance and quality of mathematical solvers.

Here, we provide an in-depth analysis of the DSACQ algorithm applied to the first case, which we shall refer to as TROTS Case 1. Although the original patient data was unable to be traced³, the IMPT plans in the TROTS database were based on a published “wish list” for oropharyngeal cancer treatment using photon IMRT [448]. TROTS Case 1 most closely resembled the wish list for Stage III cancer of the hypopharynx, with TNM (tumour, node and metastasis) staging T4aN2bM0. This means: (1) that the tumour has invaded at least one of the larynx, tongue, muscles in the jaw, roof of the mouth, or jawbone (T4a); (2) the cancer has spread to more than one lymph node on the same side as the primary tumour volume, but none measure larger than 6 cm (N2b); and (3) the cancer has not metastasised (M0). The original

²Available at: <https://sebastiaanbreedveld.nl/trots/>

³Helpful information was provided by correspondence with Dr. Sebastiaan Breedveld, the creator of the TROTS data set, however due to the anonymisation of patient data and the age of the data set, specific case information was unable to be retrieved.

prescription, on which TROTS Case 1 appears to be based, consisted of multiple target regions between 54 Gy and 66 Gy [449]. The plan was generated from three coplanar beams with gantry angles of 50°, 180° and 310°.

The TROTS plans contain multi-criteria and robust optimisation objectives, however there are no DVC prescriptions. Therefore, dose-volume histogram (DVH) curves for each structure were reproduced using the archived dose solution included with the TROTS plan, in order to extract approximate $D_{50\%}$ DVCs as well as minimum and maximum dose constraints. The purpose of this was to seek similar DVH curves using the DSACQ algorithm. A number of “ring” structures are present in the original plan as a means to better shape the target dose. In this implementation these structures were neglected and, instead, only the main target definition was used. Resulting dose solutions were compared via their cumulative DVH curves and spatial dose maps. Wilcoxon signed rank tests were performed at the 95% confidence level to determine statistically significant variations in the dose differences between the DSACQ solution and the accepted solution.

Finally, a convergence analysis, using 500 cycles of the DSACQ algorithm, was performed for all remaining 19 TROTS proton plans using the same methods as those described above, using both a static parameter method and the SARP2 method.

7.2.4 Code acceleration for clinical efficiency

The DSACQ algorithm was first implemented as an object-oriented program in *MATLAB* so that it would be easily communicable with *matRad*⁴ [163], an open-source dose calculation and optimisation toolkit, which also runs in *MATLAB*, used previously in Chapter 4. The TROTS data set is also formatted as *MATLAB* data files. In an effort to increase computational efficiency, the program was rewritten in *Python*, exploiting just-in-time (JIT) compilation and CPU multi-threading (MT). Run times were compared between the original and accelerated implementations.

⁴Available at: <http://e0404.github.io/matRad/>

7.3 Results

The performance of the DSACQ algorithm has been assessed both visually, using spatial dose maps and DVH curves, and quantitatively through metrics such as the number of voxels violating their constraints and Wilcoxon signed rank tests.

Each cycle of the DSACQ algorithm was shown to incrementally improve the dose solution and converge on the prescription in all cases, neglecting finer clinical considerations. Figure 7.2(c) shows a visualisation of the dose solution for Test Case 1 following the algorithmic procedure. It is common in the clinic to evaluate plans using their DVH, which shows the percentage of each structure that has received a certain dose. Figure 7.4 shows a suitable DVH for this plan, with the prescriptions in Table 7.1 being approximately met. General convergence to the solution is indicated by a decrease in the total number of pixels violating the constraint imposed upon them, shown in the log-loss plot in Figure 7.5. Further, log-loss plots for all four types of constraints (minimum dose, maximum dose, lower DVC and upper DVC) are displayed in Figure 7.6. Again, these all show a general decrease in the number of




Figure 7.4: Cumulative dose-volume histogram (DVH) for Test Case 1 showing the percentage of each structure that has received a certain dose. HDC and DVC prescriptions are shown as filled circles.




Figure 7.5: Number of total violations as a function of number of the algorithmic cycles for Test Case 1. A decrease indicates improvement in meeting the prescription.

violations and, therefore, indicate that the solution gradually improves as the number of cycles increases.

Log-loss plots show the total number of constraint violations to decrease with the number of algorithmic cycles for all subsequent cases (Figure 7.7). Oscillatory behaviour is observed in both SARP implementations; the number of violations sharply increases or decreases at each update of the relaxation parameters. This reflects shifts in the activity of the algorithm to the constraints that are most violated. This is most evident in Test Case 2, where the prescription is split over two disjoint target regions. An explanation of this behaviour is left to the Discussion. In general, improvement in the solution at each iteration is shown not only by a decrease in the number of violations, but also by smaller fluctuations in this number. Figure 7.8 breaks down the total violations, using SARP2, from Figure 7.7(a) into the mean percentages of DVC violations and HDC violations. It can be seen that the average of these violations follows a downward trend, and that the algorithm's attempt to meet the DVCs induces more violations in the HDCs, and vice versa. Setting static relaxation parameters, and therefore static step-sizes in the projections, was able to reduce the total number of violations further than using either SARP method within 5000 algorithmic




Figure 7.6: Percentage of violations as a function of the number of algorithmic cycles for Test Case 1, shown separately for HDCs (minimum and maximum doses) and DVCs. An upper DVC is that which is applied to an avoidance structure while a lower DVC is that which is applied to a target structure.

cycles. Figure 7.7(b, c) both indicate an increased rate of improvement in the final dose solution through the use of SARP2.

TROTS Case 1, the most complex plan, achieved target coverage within 0.5 Gy of the accepted clinical dose solution and offered the same or better sparing in 11 out of 13 organs at risk, according to Figure 7.9 and Wilcoxon signed-rank tests. The left submandibular gland (SMG) received a median dose of 0.29 Gy greater than that in the accepted plan and an overall dose difference that is statistically greater than 0.5 Gy. (Significance here is defined by a Wilcoxon signed-rank test returning a probability of $p < 0.05$.) The median dose of the right SMG, however, was lowered by 2.7 Gy. Excess dose in the brainstem did not statistically exceed 0.25 Gy and there was no increase in median dose. Figure 7.10 shows the spatial dose distribution




Figure 7.7: Log-loss curves showing the reduction in the total number of constraint violations as the number of cycles increases, for (a) Test Case 2, (b) Test Case 3 and (c) TROTS Case 1. All show both parameter-update schemes, SARP1 and SARP2, as well as the static method in which the relaxation parameters are set to their mid-range.

for one transaxial and one saggital slice, with some signs of localised hot spots in the dose difference. In terms of the DVH (Figure 7.11), minimum and maximum dose constraints are shown to be well met. The majority of $D_{50\%}$ doses are the same or lower than those in the original TROTS plan for organs at risk, and the same or higher for target regions, with the exception of two ring structures, “CTV High Ring 0-10 mm Outside” and “CTV Combined Ring 0-10 mm”, both of which received less dose than was originally prescribed. Recall that these ring structures were excluded from the DSACQ inverse planning criteria, but were originally included in TROTS as a means to more precisely shape the dose distribution.

Log-loss curves in Figure 7.12 show the relative reduction in the total number of constraint violations as the number of cycles increases, for all 20 TROTS proton cases. Results using the SARP2 method are shown against the static method in which the relaxation parameters are set to their mid-range. It can be seen that plan quality, by this metric, gradually improves with the number of algorithmic cycles, in all cases. The mean violation reduction is over 99% for both the static method and the SARP2 method. SARP2 deviates from the static method after approximately 200 cycles, once the violations begin to stabilise, offering only a slight improvement on the static method. Dose differences between the DSACQ solution and the accepted

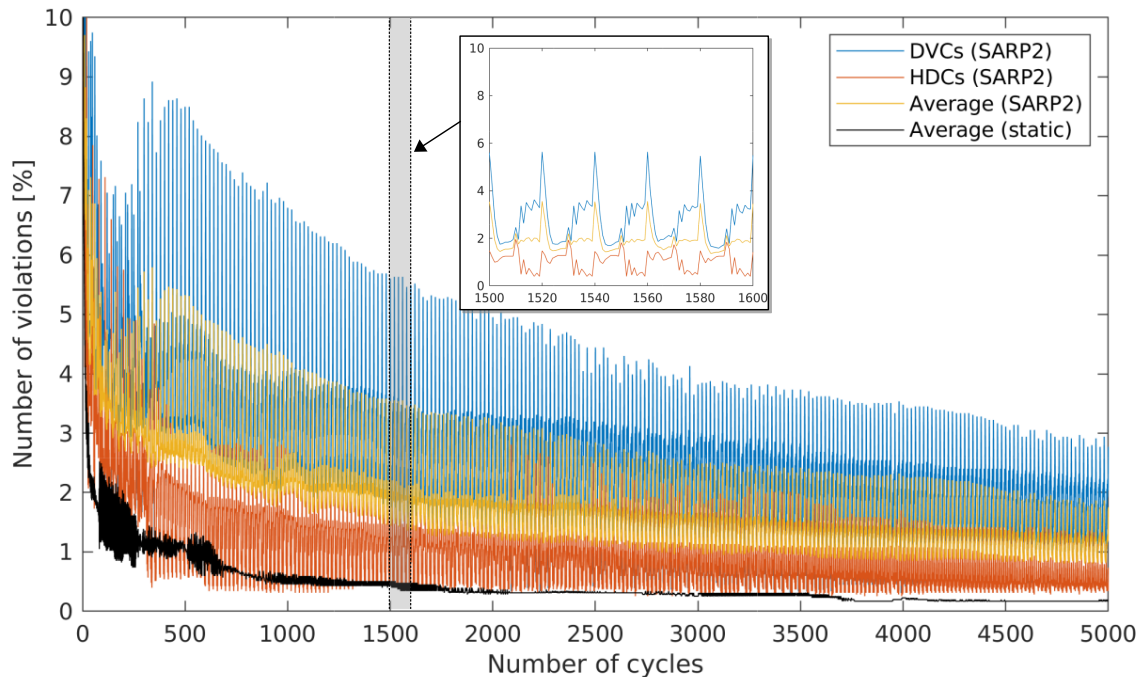


Figure 7.8: Mean number of constraint violations, as a percentage of the total number of voxels, shown against the number of algorithmic cycles for Test Case 2 using SARP2. Violations have been separated into those on DVC constraints, HDC constraints, and the average of these two. The static method are shown for comparison.

dose solution are shown, for reference, as forest plots in Appendix E, Figures E.1–E.4. Target coverage is within 1 Gy of the accepted dose solution across all plans. Healthy tissue dosage is generally equal or better except for some overdosage (within 2 Gy) of the constrictor muscles, in Plans 3, 6, 10, 13 and 15, and the submandibular gland in Plan 14.

Accelerated implementation in the *Python* programming language (version 2) using CPU multi-threading (MT) and just-in-time (JIT) compilation on up to 20 CPU threads reduced run time by a factor of 15.4, down to an average of 0.145 seconds per cycle for the most complex clinical case. A suitable solution was achieved after approximately 300 cycles (43.5 seconds). Optimising only on linear constraints (minimum and maximum dose limits) reduced this time to only 15.6 seconds with an average run time of 0.052 seconds per cycle. The DVH shown in Figure 7.14 was achieved in only 2.6 seconds, after 50 iterations of the DSACQ algorithm with no DVCs applied. Figure 7.13 shows the significant speed-up offered by the accelerated




Figure 7.9: Forest plots showing non-zero dose differences in TROTS Case 1 distributed in the voxels of the structures of interest. Higher doses obtained using the DSACQ algorithm, when compared to the included TROTS solution, manifest as positive values. One asterisk beside the data indicates there is a statistically significant difference in the doses; two or three indicate the absolute difference is statistically greater than 0.25 Gy or 0.50 Gy respectively. Significance is defined by a Wilcoxon signed-rank test returning a probability of $p < 0.05$.

code, compared to the standard *MATLAB* implementation.

7.4 Discussion

The DSACQ algorithm achieved feasible dose solutions for all test cases, showing a general trend of decreasing constraint violations as the number of algorithmic iterations was increased. Following from the rigorous mathematical analysis of convergence properties presented in the mathematical paper detailing the DSACQ algorithm [417], these results give increased confidence that the algorithm performs as required in the context of radiotherapy inverse planning. In its simplest form, the DSACQ

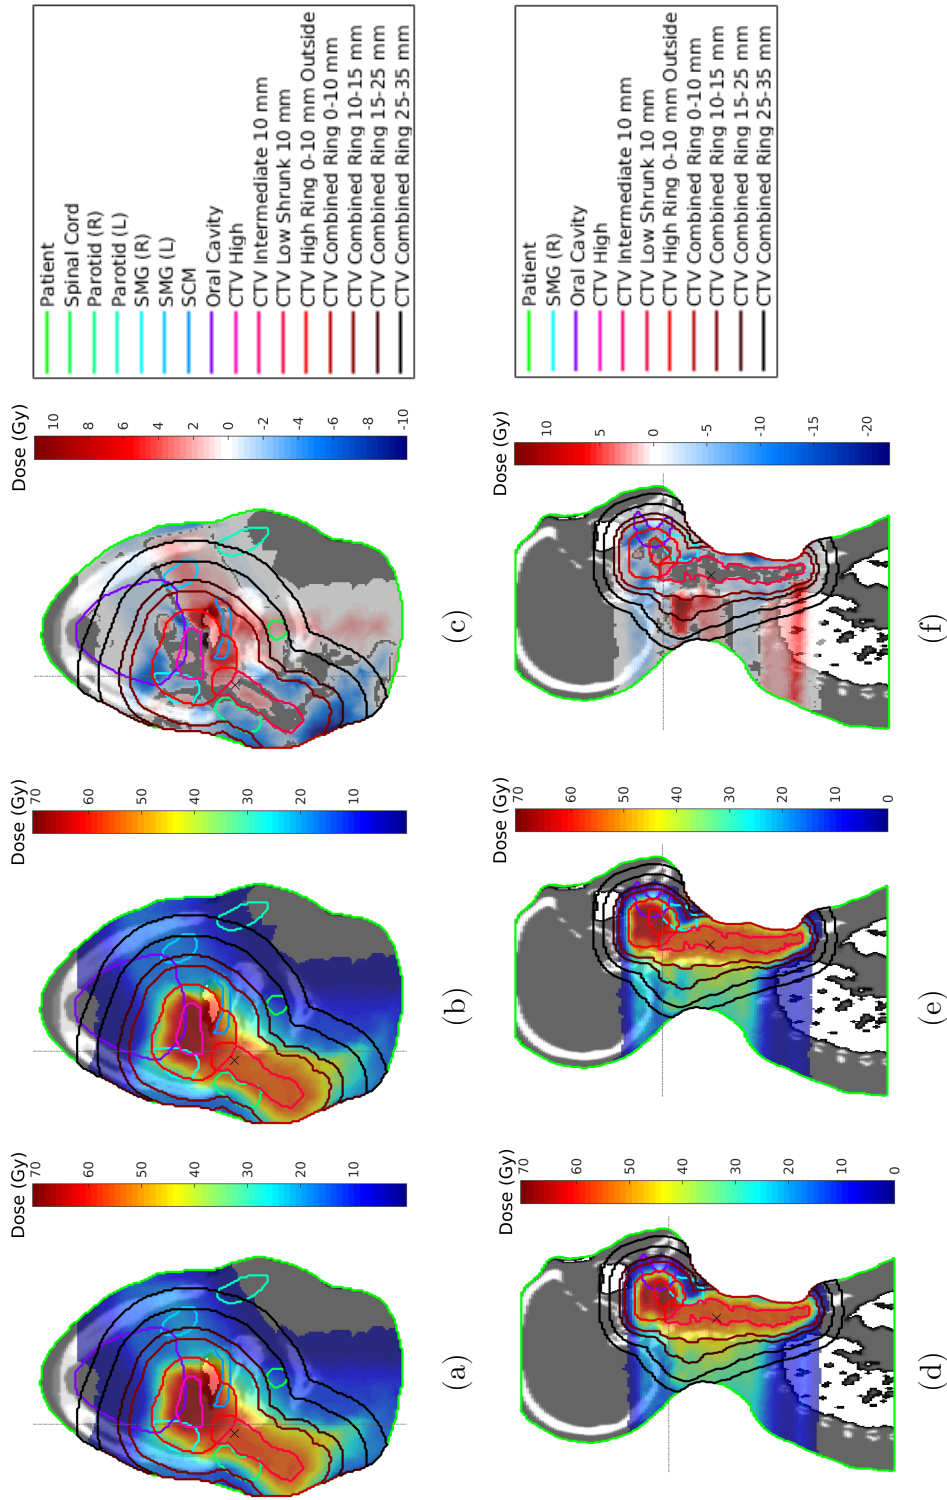


Figure 7.10: Transverse (top row) and sagittal (bottom row) CT slices for TROTS Case 1, showing the final dose after 500 iterations of the DSACQ algorithm using SARP2 (a, d), the actual dose included with the TROTS data (b, e), and the dose difference between the former and the latter (c, f). Absolute doses above 0.5 Gy or differences less than 2% are not shown. The treatment isocentre is shown by a black cross, and the chosen slices are indicated by dashed gray lines.




Figure 7.11: Cumulative dose-volume-histogram (DVH) for TROTS Case 1. Solid lines represent the solution after 500 cycles of the DSACQ algorithm using SARP2. Dashed lines represent the solution included in the TROTS data.

algorithm operates with static relaxation parameters that are set to their mid-range, meaning each iterate is projected exactly onto the constraint sets. However, dynamic parameter-update methods can be applied to continually re-weight these projections in order to satisfy the constraints that are most violated. Two possible candidates for a self-adaptive relaxation parameter (SARP) method were introduced in Chapter 6, namely SARP1 and SARP2, though the reader is encouraged to explore other possible implementations. Two additional implementations are provided in Appendix F; a threshold condition dictating when to adjust the relaxation parameters, and an extension to shift the focus of the algorithm not only between the HDCs and DVCs, but between structures as well.

Oscillatory behaviour is observed when implementing both SARP1 and SARP2 for Test Case 2. This is due to interplay between the two target prescriptions, which are split over two disjoint regions. The algorithm cannot find a solution that satisfies both target prescriptions, so it tends to oscillate between meeting the requirements




Figure 7.12: Log-loss curves showing the relative reduction in the total number of constraint violations as the number of cycles increases, for all 20 TROTS proton cases. Results using the SARP2 method are shown against the static method in which the relaxation parameters are set to their mid-range.

of Target A and meeting those of Target B. Figure 7.15 serves to explain this phenomenon; it shows the share of constraint violations between both target regions, before and after an update of the relaxation parameters. In this case, a compromise could be to take the average of the two intensity vector solutions, or to switch to the static parameter method after a sustained period of highly oscillatory behaviour. It is worth noting that, in terms of the cumulative DVH, the difference between the two solutions is minuscule (Figure 7.16). As both target prescriptions create steep dose gradients, the number of constraint violations is very sensitive to small changes in the beamlet intensity vector. The “bisection-like” method of SARP1 encourages smaller adjustments at each iteration except for when the majority of violations switches between the HDCs and the DVCs. This is why some oscillatory behaviour can also be seen in Figure 7.7(b).

In the two-dimensional problems of Test Case 1 and Test Case 2, Gaussian pseudo-




Figure 7.13: Run time comparison (per cycle) between original *MATLAB* implementation and accelerated *Python* acceleration of the DSACQ algorithm through multi-threading (MT) and just-in-time (JIT) compilation for TROTS Case 1. Bottom (blue) bars: all HDC and DVC projections. Top (orange): only HDC projections.

dose contributions were created in place of realistic proton pencil beam doses for simplicity. In reality the spatial interaction of pencil beam contributions is non-uniform and spans three dimensions. The DSACQ algorithm is a fluence-based inverse planning algorithm that seeks a feasible dose solution by adjusting the weights of pre-curved pencil beams. Prior to this step, pencil beam angles and energies must be selected through a separate process. This is similar in principle to the choice of equally spaced Gaussian kernels in the two-dimensional test cases, such that a homogeneous dose flooded the target regions. It should be noted that the choice of pencil beam contributions in the TROTS proton cases are beyond the scope of this work.

It can be concluded, from a combination of Wilcoxon signed-rank tests, cumulative DVH curves and visualised spatial dose maps, that the final dose solution achieved by the DSACQ algorithm for TROTS Case 1 did not degrade overall plan quality compared to the accepted TROTS solution. In fact, most organs received less dose. However the dose difference alone is not sufficient to claim better healthy tissue sparing or tumour control. For this we may refer to models that calculate normal tissue complication probabilities (NTCP) and tumour control probabilities (TCP). This is currently being investigated using methods detailed by Gay and Niemierko [158] and




Figure 7.14: Cumulative dose-volume-histogram (DVH) for TROTS Case 1 after code acceleration. Solid lines represent the solution after 50 cycles of the DSACQ algorithm with static relaxation parameters and no DVCs. Dashed lines represent the solution included in the TROTS data.

Warkentin *et al.* [450]. For example, while Figure 7.9 could indicate better sparing of the ipsilateral submandibular gland (“SMG (R)”) using the DSACQ algorithm, preliminary NTCP calculations⁵ do not necessarily support this; both dose solutions give a 100% complication probability for the endpoint of decreased salivary flow. However, inputting the ipsilateral parotid gland DVH data into the NTCP model indicated that the DSACQ solution decreased the likelihood of xerostomia by as much as 12%, from 35% down to 23%. Even though xerostomia is a subjective endpoint, these predictions are contradictory, as both the SMG and the parotid glands are responsible for salivary production. Further clarification and analysis of NTCP and TCP modelling techniques, which is beyond the scope of this thesis, should take place in future studies. It is of interest to note that the models predicted full tumour control (100%) and

⁵NTCP calculations were performed by feeding DVH data into a *Python* script written by Dr. Pratik Samant (Oxford, UK), which is based on the methods contained in [158, 450].




Figure 7.15: Share of constraint violations between Target A and Target B in Test Case 2. The outer ring represents the 499th algorithmic cycle while the inner ring represents the 500th cycle, after an update of the relaxation parameters with SARP2. This behaviour occurs each time the relaxation parameters are updated. A compromise between the two solutions is to take the average.

less than a 0.1% chance of brainstem or spinal cord necrosis, for both dose solutions. The DSACQ solution also decreased the likelihood of laryngeal cartilage necrosis from 51% to 10.5%.

The dose difference maps in Figure 7.10(c, f) show localised regions of significant dose difference, known as “hot spots” (overdose) and “cold spots” (underdose), as large as 10 Gy. Their presence is likely due to the decision to remove target ring structures from the inverse planning process, as these help to shape the dose. While these regions are not immediately recognisable in the DSACQ solution (Figure 7.10(a, d)), care should be taken in practice to redistribute any hot or cold spots that may arise in the final solution. This is standard practice in the treatment planning workflow. While the algorithm includes volumetric constraints in the form of DVCs, it is possible for these to be satisfied even when a small subvolume of the structure is significantly overdosed or underdosed. In addition to subtle corrections being performed post-optimisation, hot spot effects can also be somewhat mitigated by defining a new structure as the disjunctive union of the patient with all other structures and applying

to it a maximum dose constraint. This structure represents the voxels of the patient that otherwise have no constraints applied to them.

Applying 500 cycles of the DSACQ algorithm to all 20 TROTS proton therapy plans led to a 99% reduction, on average, in the number of constraint violations when compared to the initial dose distribution given by setting all beamlets to unit intensity. Given this performance, it is not surprising that the static- and dynamic-parameter (SARP2) methods only differed after approximately 200 cycles and there was little gain from using SARP2. This particular automatic parameter-update method was designed to activate when the number of violations begin to stabilise. In cases where the static method provides a suitable rate of convergence, there is little to no expected gain from the SARP methods. Testing different SARP implementations in a large number of clinical plans, covering different treatment sites and plan complexities, will shed more light on the overall performance of these methods in comparison to conventionally used static relaxation parameter methods.

The purpose in using the TROTS data set in this work has not been to improve plan quality; rather, it has been to investigate whether the DSACQ algorithm is capable of producing similar plan quality to the original solutions. Using DVH curves of the original solutions to set tolerance dose constraints and $D_{50\%}$ DVC prescriptions has indeed enabled the DSACQ algorithm to produce feasible dose solutions. Moving forward, alternative prescriptions and other clinical plans may be trialed with further consultation from clinicians.

The feasibility of using the DSACQ algorithm in clinical settings relies not only on the quality of the solution it obtains, but on the time it takes to reach that solution. Standard implementation in *MATLAB* took approximately 11 minutes to arrive at a suitable solution while the accelerated implementation in *Python* took under one minute, which allows for efficient clinical usage. Furthermore, the solution obtained in only 2.6 seconds using no DVCs was not too dissimilar from the accepted solution. One cycle of linear projections could be completed in an average time of 50 ms. It is therefore suggested that some number of cycles of linear projections be performed prior to “turning on” the DVCs. This allows for the algorithm to quickly




Figure 7.16: Cumulative dose-volume-histogram (DVH) for Test Case 2. Solid lines represent the solution after 500 cycles of the DSACQ algorithm using SARP2. Dashed lines represent the solution after 499 cycles.

shift pencil beam weights closer to the space of feasible solutions before introducing the more computationally intensive DVC projections. Additional speed-up can be achieved through the choice of hardware, as the computational time is expected to scale linearly with the number of CPU cores used during multi-threading. Porting certain subroutines, such as matrix multiplication, to a GPU was investigated briefly, however the computational overhead involved typically led to decreased performance. This is expected for relatively small-scale problems. It could prove useful to investigate GPU implementation for plans with higher resolution dose maps, and therefore larger matrices, or with larger dose fields, and therefore less sparse matrices. Further decreasing the time taken to complete one cycle opens the possibility of real-time user interaction. This could involve dynamically adjusting the relative weighting of projections or the relaxation parameters through on-screen tools and immediately seeing the result. There are a myriad of GUI designs that could enable this. The graphical component of displaying the resulting DVH or dose map would most likely involve GPU-based calculation methods to avoid computational bottlenecks.

In summary, the DSACQ algorithm introduced in Chapter 6 has been tested on a variety of cases, including archived head-and-neck cases containing up to 23 delineated structures with both linear and nonlinear constraints imposed. All test

cases returned feasible dose solutions. An accelerated version of the algorithm which makes use of multi-threading and just-in-time compilation dramatically reduced the run time. Numerical convergence can be achieved well within clinically imposed time restrictions, however finer spatial considerations such as hot and cold dose spots may still need judicious correction post-optimisation.

Chapter 8

Concluding remarks

Radiation therapy is a crucial pillar in the treatment of cancer, tangential to chemotherapy and surgical intervention. Advances in IMRT and proton dose delivery have enabled highly conformal tumouricidal dose deposition, with proton therapy in particular able to drive down dose leakage to healthy tissues. However, as has been vehemently expressed throughout this thesis, physical dose deposition in the patient does not tell the full story; radiobiological mechanisms and their intermodality differences must be well-understood, as these describe the processes by which the energy deposited in cells eventuates in cell death and ultimately clinical outcomes. It has become increasingly clear that proton RBE is dependent on a number of factors, the most convincing of which is arguably the LET, or corresponding ionisation density, as protons are slowed through cellular targets. Despite the evidence of there existing a variable, depth-dependent proton RBE, there are few examples of radiobiological optimisation strategies employed in treatment planning systems beyond the simplified use of a constant multiplicative factor to express the equivalent x-ray dosage.

X-ray and proton treatments alike are planned based on empirical dose-volume thresholds and extreme (upper or lower) dose tolerances observed for a range of tissues over many years of clinical experience in conventional radiotherapy. In order to plan with confidence, one must be able to firstly represent the proton dose, if applicable, as an accurate x-ray equivalent in units of Gy(RBE), and then shape the dose to meet all the required HDCs and DVCs. This thesis has proposed novel methods and recommendations for the inclusion of radiobiological factors in treatment planning

through (1) a variable but pragmatic RBE model based on the relative likelihood of DNA double-strand-break induction, and (2) a flexible projection-based inverse planning algorithm that comprehensively addresses dose-volume effects through the inclusion of exact DVC modelling on any number of target or avoidance structures.

The biophysical RBE model presented in this thesis captures the fundamental physical process behind the phenomenon of increased cell-killing efficiency toward the distal edge of the proton Bragg peak. As a single particle interaction model, predictions in DSB induction scale linearly with dose, allowing predictions in overlapping voxels due to separate pencil beams to simply add. Consequently, standard fluence-based inverse optimisation may be applied to the RBE-weighted dose distribution. This RBE model may be used in conjunction with any FMO-based inverse planning algorithm, which of course includes the algorithm presented in this thesis. This algorithm takes a feasibility-seeking approach, in contrast to conventional cost-minimisation techniques, and was demonstrated to generate feasible dose solutions in non-convex settings, with any number of linear or non-linear constraints imposed. A flexible implementation is offered, allowing projection parameters to vary freely within their permitted interval, whether by the self-adaptive methods proposed or by manual user intervention.

This thesis set out to explore an hypothesis:

Inclusion of DNA damage modelling and comprehensive plan optimisation strategies, as a result of the methods presented in this thesis, will enable confident and effective plan creation not only in terms of physical dose deposition but, crucially, in biological effect.

Evidence for variable RBE and for suitability of the proposed RBE model in practice was presented through a range of Monte Carlo simulations, implementation strategies, experimental data comparisons and analysis of a clinical beam line experiment, while the proposed inverse planning algorithm produced suitable dose solutions to a number of synthetically generated, as well as complex archived, treatment plans. In future work, both the RBE model and planning algorithm should be tested on a large variety of treatment plans and compared to alternate models and algorithms in order to

build confidence in the approach among specialists and to identify and strengthen any shortcomings. This author suggests that multi-disciplinary improvements to the RBE model be explored in more detail, such as the inclusion of DNA damage repair and tissue-specific response modelling, and that alternate parameter-update methods, GPU acceleration and user interaction possibilities be proposed and trialled in the inverse planning algorithm. With these focuses in mind, iterations upon the initial strategies explored in this thesis will continually improve the confidence with which radiobiological factors are considered in radiation therapy treatment planning.

Appendix A

Derivation of the interaction length density function

If we assume a particle enters the cylinder as shown in Figure 3.1 with azimuthal angle $\phi \in (-\frac{\pi}{2}, \frac{\pi}{2})$ and polar angle $\theta \in (-\frac{\pi}{2}, \frac{\pi}{2})$ then the interaction length ℓ is given by

$$\ell(\theta, \phi) = \begin{cases} d \frac{\cos \phi}{\cos \theta}, & |\tan \theta| \leq \frac{H}{d} \frac{1}{\cos \phi}, \\ \frac{H}{\sin \theta}, & |\tan \theta| > \frac{H}{d} \frac{1}{\cos \phi}, \end{cases} \quad (\text{A.1a})$$

$$(\text{A.1b})$$

where d is the diameter of the cylindrical volume. If the particle exits through the curved surface of the cylinder then (A.1a) applies. If it exits through the top or bottom of the cylinder then (A.1b) applies. Reformulating the problem in terms of $x := \cos \phi$ and $y := \cos \theta$ gives

$$\ell(x, y) = \begin{cases} \frac{xd}{y}, & x \leq \frac{H}{d} \frac{y}{\sqrt{1-y^2}}, \\ \frac{H}{\sqrt{1-y^2}}, & x \geq \frac{H}{d} \frac{y}{\sqrt{1-y^2}}. \end{cases} \quad (\text{A.2a})$$

$$(\text{A.2b})$$

Both angles are independently chosen from uniform distributions, $\Phi \sim U(-\frac{\pi}{2}, \frac{\pi}{2})$ and $\Theta \sim U(-\frac{\pi}{2}, \frac{\pi}{2})$, with corresponding density functions f_Φ and f_Θ given by

$$f_\Phi(\phi) = \begin{cases} \frac{1}{\pi}, & \phi \in (-\frac{\pi}{2}, \frac{\pi}{2}), \\ 0, & \text{otherwise,} \end{cases} \quad (\text{A.3a})$$

$$(\text{A.3b})$$

and

$$f_{\Theta}(\theta) = \begin{cases} \frac{1}{\pi}, & \theta \in (-\frac{\pi}{2}, \frac{\pi}{2}), \\ 0, & \text{otherwise.} \end{cases} \quad (\text{A.4a})$$

$$\quad (\text{A.4b})$$

We now define the random variables $X \sim \cos \Phi$ and $Y \sim \cos \Theta$. The density function for X is given by

$$f_X(x) = \sum_k f_{\Phi}(\phi_k) \left| \frac{d\phi}{dx} \right|_{\phi=\phi_k} \quad (\text{A.5})$$

where ϕ_k are the k solutions to $x = \cos \phi$. We have one solution in the first quadrant of the unit circle, $\phi_1 \in [0, \frac{\pi}{2})$, and the other in the fourth quadrant, $\phi_2 = -\phi_1 \in (-\frac{\pi}{2}, 0]$.

Therefore,

$$f_X(x) = \frac{1}{\pi} \left| \frac{-1}{\sin \phi_1} \right| + \frac{1}{\pi} \left| \frac{1}{\sin \phi_1} \right| = \frac{2}{\pi} \frac{1}{\sqrt{1-x^2}}. \quad (\text{A.6})$$

Similarly,

$$f_Y(y) = \frac{2}{\pi} \frac{1}{\sqrt{1-y^2}}. \quad (\text{A.7})$$

The joint density of X and Y , denoted f_{XY} , is simply the product of these independent densities. That is, $f_{XY} = f_X f_Y$.

Let us now define $L \sim Xd/Y$ to be a random variable representing the interaction length in the case that the proton exits through the curved surface of the cylinder. The cumulative density function for L , or the probability of a particle having an interaction length of at most ℓ , is given by

$$F_L(\ell) := \mathcal{P}(L \leq \ell) = \iint_{(x,y) \in \mathcal{D}} f_{XY}(x,y) dx dy \quad (\text{A.8})$$

where

$$\mathcal{D} := \left\{ (x,y) \in \mathbb{R}_+^2 \mid \frac{xd}{y} \leq \ell \right\}. \quad (\text{A.9})$$

The set \mathcal{D} is represented by the shaded region in Figure A.1. Note that the boundaries of the region for $\ell \leq H$ are different to those for $\ell > H$. The integration in (A.8)

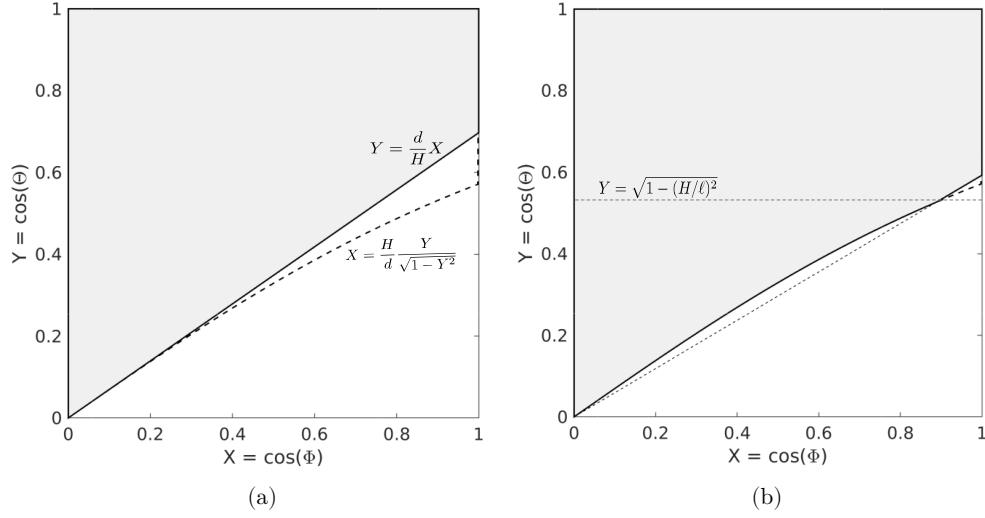


Figure A.1: Shaded regions represent the integration region $\mathcal{D} := \{(x, y) \in \mathbb{R}_+^2 \mid xd/y \leq \ell\}$ for examples where (a) $\ell \leq H$, and (b) $\ell > H$.

becomes

$$F_L(\ell) = \begin{cases} \int_0^1 \int_0^{y\ell/d} f_{XY} dx dy, & \ell \leq H, \quad (\text{A.10a}) \\ \int_0^{\sqrt{1-(H/\ell)^2}} \int_0^{\frac{H}{d} \frac{y}{\sqrt{1-y^2}}} f_{XY} dx dy & \ell > H. \quad (\text{A.10b}) \\ + \int_{\sqrt{1-(H/\ell)^2}}^1 \int_0^{y\ell/d} f_{XY} dx dy, & \end{cases}$$

Finally, the density, f_L , is given by

$$f_L(\ell) = \frac{\partial}{\partial \ell} F_L(\ell). \quad (\text{A.11})$$

It can be shown with a few substitutions that the analytical solution is

$$f_L(\ell) = \begin{cases} \frac{4}{\pi^2} \frac{\partial}{\partial \ell} \int_0^1 \frac{\arcsin\left(\frac{ly}{d}\right)}{\sqrt{1-y^2}} dy, & \ell \leq H, \quad (\text{A.12a}) \\ \frac{4}{\pi^2} \frac{\partial}{\partial \ell} \int_0^{\sqrt{1-(H/\ell)^2}} \frac{\arcsin\left(\frac{Hy}{d\sqrt{1-y^2}}\right)}{\sqrt{1-y^2}} dy & \ell > H, \quad (\text{A.12b}) \\ + \frac{4}{\pi^2} \frac{\partial}{\partial \ell} \int_{\sqrt{1-(H/\ell)^2}}^1 \frac{\arcsin\left(\frac{ly}{d}\right)}{\sqrt{1-y^2}} dy, & \end{cases}$$




Figure A.2: The integration region $\mathcal{D}_{\text{ext}} := \{(x, y) \in \mathbb{R}_+^2 | y \leq \sqrt{1 - H^2/\ell^2}\}$ is given by the shaded area.

which can be further simplified to

$$f_L(\ell) = \begin{cases} \frac{4}{\pi^2 d} \int_0^1 \frac{y}{\sqrt{1-y^2} \sqrt{1-(\ell y/d)^2}} dy, & \ell \leq H, & (\text{A.13a}) \\ \frac{4H}{\pi^2 \ell^2} \frac{\arcsin\left(\frac{\ell}{d} \sqrt{1-(H/\ell^2)}\right)}{\sqrt{1-(H/\ell^2)}} & \ell > H. & (\text{A.13b}) \\ \quad + \frac{4}{\pi^2} \frac{\partial}{\partial \ell} \int_{\sqrt{1-(H/\ell)^2}}^1 \frac{\arcsin\left(\frac{\ell y}{d}\right)}{\sqrt{1-y^2}} dy, & & \end{cases}$$

We have derived the probability density function, f_L , for the interaction length distribution of a particle that passes through a cylindrical volume as shown in Figure 3.1 with diameter d and height $2H$, given the particle exits through the curved surface of the cylinder and not through the top or bottom surfaces. For the purposes of this work, we may take (A.13a) and (A.13b) to be the final solution, as a DNA double-strand break is most likely to result from a particle passing through the full diameter of the DNA double helix within 10 base pairs. Nevertheless, we will now derive the full solution for completeness.

In the case that the particle exits through the top or bottom surfaces, we define

$L_{\text{ext}} \sim H\sqrt{1 - Y^2}$ where the subscript ‘ext’ is given to signify that this is an extension to the formulation above. The integration region, \mathcal{D}_{ext} , becomes

$$\mathcal{D}_{\text{ext}} := \left\{ (x, y) \in \mathbb{R}_+^2 \mid y \leq \sqrt{1 - H^2/\ell^2} \right\}, \quad (\text{A.14})$$

which is represented by the shaded region in Figure A.2, and the resulting density function is

$$f_{L_{\text{ext}}}(\ell) = \frac{4}{\pi^2} \left[\int_0^{\frac{\ell}{d}\sqrt{1-(H/\ell)^2}} \frac{1}{\sqrt{1-x^2}} \arcsin\left(\frac{xd}{\sqrt{H^2+x^2d^2}}\right) dx + \int_{\frac{\ell}{d}\sqrt{1-(H/\ell)^2}}^1 \frac{1}{\sqrt{1-x^2}} \arcsin\left(\sqrt{1-(H/\ell)^2}\right) dx \right]. \quad (\text{A.15})$$

The full solution, including all possible surfaces through which the particle can exit, is therefore given by (A.13a) for $\ell \leq H$ and (A.13b)+(A.15) for $\ell > H$.

Appendix B

C++ scripts for TOPAS extensions

Main script for damage inflicted by protons

```
1 // * *
2 // * Custom TOPAS extension for scoring accumulated double strand *
3 // * break (DSB) counts per cell. *
4 // * *
5 // * Author: Mark Brooke (mark.brooke@oncology.ox.ac.uk) *
6 // * (markdanielbrooke@gmail.com) *
7 // * September 2020 *
8 // *****
9 //
10
11 #include "ProtonTotalDamagePlasmonic.hh"
12 #include "G4ParticleTable.hh"
13 #include "G4ParticleDefinition.hh"
14
15 #include "G4VSolid.hh"
16 #include "G4VPhysicalVolume.hh"
17 #include "G4VPVParameterisation.hh"
18 #include "G4UnitsTable.hh"
19
20
21 ProtonTotalDamagePlasmonic::ProtonTotalDamagePlasmonic(TsParameterManager* pM,
22 TsMaterialManager* mM, TsGeometryManager* gM, TsScoringManager* scM,
23 TsExtensionManager* eM, G4String scorerName, G4String quantity,
24 G4String outFileName, G4bool isSubScorer)
25 : TsVBinnedScorer(pM, mM, gM, scM, eM, scorerName, quantity,
26 outFileName, isSubScorer)
27 {
28 SetUnit("mole");
29 fProtonDefinition = G4ParticleTable::GetParticleTable()
```

```

30         ->FindParticle("proton");
31     }
32
33
34     ProtonTotalDamagePlasmonic::~ProtonTotalDamagePlasmonic() {}
35
36
37     G4bool ProtonTotalDamagePlasmonic::ProcessHits(G4Step* aStep, G4TouchableHistory*)
38     {
39         if (!fIsActive) {
40             fSkippedWhileInactive++;
41             return false;
42         }
43
44         G4double ke = aStep->GetTrack()->GetKineticEnergy()/MeV;
45         G4double edep = aStep->GetTotalEnergyDeposit()/joule;
46         if ( ke > 42.396/1000 && edep > 0. ) {
47
48             if (aStep->GetTrack()->GetParticleDefinition()==fProtonDefinition){
49
50                 G4double T2 = 1;
51                 G4double A1 = 117.0067884;
52                 G4double B1 = -318.6944555;
53
54                 G4double mfp = (ke*1000-T2*exp(-B1/A1))/(A1*log(ke*1000/T2)+B1);
55                 // Plasmonic effects:
56                 mfp += sqrt(ke*1000)/(0.681*pow((ke*1000-42.396),1.249));
57
58                 // Calculate the number of DSBs produced per cell per Gy
59                 // due to this single particle track
60                 G4double damage = 25.031*(1-(1+1.3519/mfp)*exp(-1.3519/mfp))+2.8725;
61
62                 // Multiply by the dose deposited by this particle
63                 G4int idx = ((G4TouchableHistory*)(aStep->GetPreStepPoint()
64                     ->GetTouchable()))->GetReplicaNumber(indexDepth);
65                 G4double cubicVolume = ComputeVolume(aStep, idx);
66                 G4double density = aStep->GetTrack()->GetStep()->GetPreStepPoint()
67                     ->GetMaterial()->GetDensity()/(g/cm3);
68                 G4double dose = edep / ( density * cubicVolume ) / 1000;
69                 // factor of 1000 above needed to convert g to kg
70                 // Multiply by the statistical weight of the track
71                 dose *= aStep->GetPreStepPoint()->GetWeight();
72
73                 AccumulateHit(aStep, damage*dose*mole );
74                 // The unit "mole" is used as a placeholder. It carries no significance.

```

```
75     // The accumulated result is the total number of DSBs per cell
76     // in the voxel of interest.
77     } else {
78         AccumulateHit(aStep, 0*mole );
79     }
80     return true;
81 }
82 return false;
83 };
84
85 G4double ProtonTotalDamagePlasmonic::ComputeVolume(G4Step* aStep, G4int idx){
86     G4VPhysicalVolume* physVol = aStep->GetPreStepPoint()->GetPhysicalVolume();
87     G4VPVParameterisation* physParam = physVol->GetParameterisation();
88     G4VSolid* solid = 0;
89     if(physParam)
90     { // for parameterized volume
91         if(idx<0)
92         {
93             G4ExceptionDescription ED;
94             ED << "Incorrect replica number --- GetReplicaNumber: "
95                 << idx << G4endl;
96             G4Exception("G4PSDoseDeposit::ComputeVolume",
97                 "DetPS0004", JustWarning, ED);
98         }
99         solid = physParam->ComputeSolid(idx, physVol);
100         solid->ComputeDimensions(physParam, idx, physVol);
101     }
102     else
103     { // for ordinary volume
104         solid = physVol->GetLogicalVolume()->GetSolid();
105     }
106
107     return solid->GetCubicVolume()/cm3;
108 };
```

Header file for damage inflicted by protons

```
1 //
2 // *****
3 // *
4 // * Header file for the custom TOPAS extension for scoring *
5 // * accumulated double strand break (DSB) counts per cell. *
6 // *
7 // * Author: Mark Brooke (mark.brooke@oncology.ox.ac.uk) *
8 // * (markdanielbrooke@gmail.com) *
9 // * September 2020 *
10 // *****
11 //
12
13 #ifndef ProtonTotalDamagePlasmonic_hh
14 #define ProtonTotalDamagePlasmonic_hh
15
16 #include "TsVBinnedScorer.hh"
17
18 class ProtonTotalDamagePlasmonic : public TsVBinnedScorer
19 {
20 public:
21     ProtonTotalDamagePlasmonic(TsParameterManager* pM, TsMaterialManager* mM,
22         TsGeometryManager* gM, TsScoringManager* scM, TsExtensionManager* eM,
23         G4String scorerName, G4String quantity, G4String outFileNames,
24         G4bool isSubScorer);
25
26     virtual ~ProtonTotalDamagePlasmonic();
27
28     G4bool ProcessHits(G4Step*, G4TouchableHistory*);
29
30     G4double ComputeVolume(G4Step*, G4int);
31
32 protected:
33     G4ParticleDefinition* fProtonDefinition;
34 };
35 #endif
```

Main script for damage inflicted by electrons

```
1 // Scorer for ElectronTotalDamagePlasmonic
2 //
3 // *****
4 // *
5 // * Custom TOPAS extension for scoring accumulated double strand *
6 // * break (DSB) counts per cell. *
7 // *
8 // * Author: Mark Brooke (mark.brooke@oncology.ox.ac.uk) *
9 // * (markdanielbrooke@gmail.com) *
10 // * September 2020
11 // *****
12 //
13
14 #include "ElectronTotalDamagePlasmonic.hh"
15 #include "G4ParticleTable.hh"
16 #include "G4ParticleDefinition.hh"
17
18 #include "G4VSolid.hh"
19 #include "G4VPhysicalVolume.hh"
20 #include "G4VPVParameterisation.hh"
21 #include "G4UnitsTable.hh"
22
23
24 ElectronTotalDamagePlasmonic::ElectronTotalDamagePlasmonic(TsParameterManager* pM,
25 TsMaterialManager* mM, TsGeometryManager* gM, TsScoringManager* scM,
26 TsExtensionManager* eM, G4String scorerName, G4String quantity,
27 G4String outFileFileName, G4bool isSubScorer)
28 : TsVBinnedScorer(pM, mM, gM, scM, eM, scorerName, quantity,
29 outFileFileName, isSubScorer)
30 {
31 SetUnit("mole");
32 fElectronDefinition = G4ParticleTable::GetParticleTable()
33 ->FindParticle("e-");
34 }
35
36
37 ElectronTotalDamagePlasmonic::~ElectronTotalDamagePlasmonic() {}
38
39
40 G4bool ElectronTotalDamagePlasmonic::ProcessHits(G4Step* aStep, G4TouchableHistory*)
41 {
42 if (!fIsActive) {
43 fSkippedWhileInactive++;
```

```

44     return false;
45 }
46
47 G4double ke = aStep->GetTrack()->GetKineticEnergy()/MeV;
48 G4double edep = aStep->GetTotalEnergyDeposit()/joule;
49 if ( ke > 0.0243/1000 && edep > 0. ) {
50
51     if (aStep->GetTrack()->GetParticleDefinition()==fElectronDefinition){
52
53         G4double T2 = 1*0.001;
54         G4double A1 = 65.898*0.001;
55         G4double B1 = -128.23*0.001;
56
57         G4double mfp = (ke*1000-T2*exp(-B1/A1))/(A1*log(ke*1000/T2)+B1);
58         // Plasmonic effects:
59         mfp += sqrt(ke*1000)/(22.228*pow((ke*1000-0.0243),1.01));
60
61         // Calculate the number of DSBs produced per cell per Gy
62         // due to this single particle track
63         G4double damage = 25.432*(1-(1+1.3858/mfp)*exp(-1.3858/mfp))+2.836;
64
65         // Multiply by the dose deposited by this particle
66         G4int idx = ((G4TouchableHistory*)(aStep->GetPreStepPoint()
67             ->GetTouchable()))->GetReplicaNumber(indexDepth);
68         G4double cubicVolume = ComputeVolume(aStep, idx);
69         G4double density = aStep->GetTrack()->GetStep()->GetPreStepPoint()
70             ->GetMaterial()->GetDensity()/(g/cm3);
71         G4double dose = edep / ( density * cubicVolume ) / 1000;
72         // factor of 1000 above needed to convert g to kg
73         // Multiply by the statistical weight of the track
74         dose *= aStep->GetPreStepPoint()->GetWeight();
75
76         AccumulateHit(aStep, damage*dose*mole );
77         // The unit "mole" is used as a placeholder. It carries no significance.
78         // The accumulated result is the total number of DSBs per cell
79         // in the voxel of interest.
80     } else {
81         AccumulateHit(aStep, 0*mole );
82     }
83     return true;
84 }
85 return false;
86 };
87
88 G4double ElectronTotalDamagePlasmonic::ComputeVolume(G4Step* aStep, G4int idx){

```

```
89   G4VPhysicalVolume* physVol = aStep->GetPreStepPoint()->GetPhysicalVolume();
90   G4VPVParameterisation* physParam = physVol->GetParameterisation();
91   G4VSolid* solid = 0;
92   if(physParam)
93   { // for parameterized volume
94     if(idx<0)
95     {
96       G4ExceptionDescription ED;
97       ED << "Incorrect replica number --- GetReplicaNumber: "
98         << idx << G4endl;
99       G4Exception("G4PSDoseDeposit::ComputeVolume",
100                 "DetPS0004", JustWarning, ED);
101     }
102     solid = physParam->ComputeSolid(idx, physVol);
103     solid->ComputeDimensions(physParam, idx, physVol);
104   }
105   else
106   { // for ordinary volume
107     solid = physVol->GetLogicalVolume()->GetSolid();
108   }
109
110   return solid->GetCubicVolume()/cm3;
111 };
```

Header file for damage inflicted by electrons

```
1 //
2 // *****
3 // *
4 // * Header file for the custom TOPAS extension for scoring *
5 // * accumulated double strand break (DSB) counts per cell. *
6 // *
7 // * Author: Mark Brooke (mark.brooke@oncology.ox.ac.uk) *
8 // * (markdanielbrooke@gmail.com) *
9 // * September 2020 *
10 // *****
11 //
12
13 #ifndef ElectronTotalDamagePlasmonic_hh
14 #define ElectronTotalDamagePlasmonic_hh
15
16 #include "TsVBinnedScorer.hh"
17
18 class ElectronTotalDamagePlasmonic : public TsVBinnedScorer
19 {
20 public:
21     ElectronTotalDamagePlasmonic(TsParameterManager* pM,
22         TsMaterialManager* mM, TsGeometryManager* gM, TsScoringManager* scM,
23         TsExtensionManager* eM, G4String scorerName, G4String quantity,
24         G4String outFileName, G4bool isSubScorer);
25
26     virtual ~ElectronTotalDamagePlasmonic();
27
28     G4bool ProcessHits(G4Step*, G4TouchableHistory*);
29
30     G4double ComputeVolume(G4Step*, G4int);
31
32 protected:
33     G4ParticleDefinition* fElectronDefinition;
34 };
35 #endif
```

Appendix C

Stable distribution fits to MiniPIX LET measurements

The Landau distribution [364] is a special case of stable distribution, defined below [see, e.g., 451], with stability parameter $\alpha = 1$ and skewness $\beta = 1$. Its moments are undefined, so it has no mean or variance, but it can be useful for finding a most probable value (MPV) of a distribution.

Definition C.1 (Stable distribution) (i) Let X_1 and X_2 be independent copies of a random variable X . Then X is said to be stable if for any constants $a > 0$ and $b > 0$ the random variable $aX_1 + bX_2$ has the same distribution as $cX + d$ for some constants $c > 0$ and d . (ii) A random variable X is called stable if its characteristic function, ψ , can be written as

$$\psi(t; \alpha, \beta, c, \mu) = \exp(it\mu - |ct|^\alpha(1 - i\beta\text{sgn}(t)\Phi)) \quad (\text{C.1})$$

where $\text{sgn}(t)$ is $+1$ if t is positive or -1 if t is negative, and $\Phi = \tan(\pi\alpha/2)$ if $\alpha \neq 1$ or $\Phi = -2 \log |t|/\pi$ if $\alpha = 1$.

The probability density function (PDF) for the Landau distribution can be written as

$$f(x) = \frac{1}{\pi c} \int_0^\infty e^{-t} \cos \left[t \left(\frac{x - \mu}{c} \right) + \frac{2t}{\pi} \log \left(\frac{t}{c} \right) \right] dt. \quad (\text{C.2})$$

Landau fits, as well as stable distribution fits where parameters are permitted to vary, have been made to the measured LET data from the MiniPIX detector in

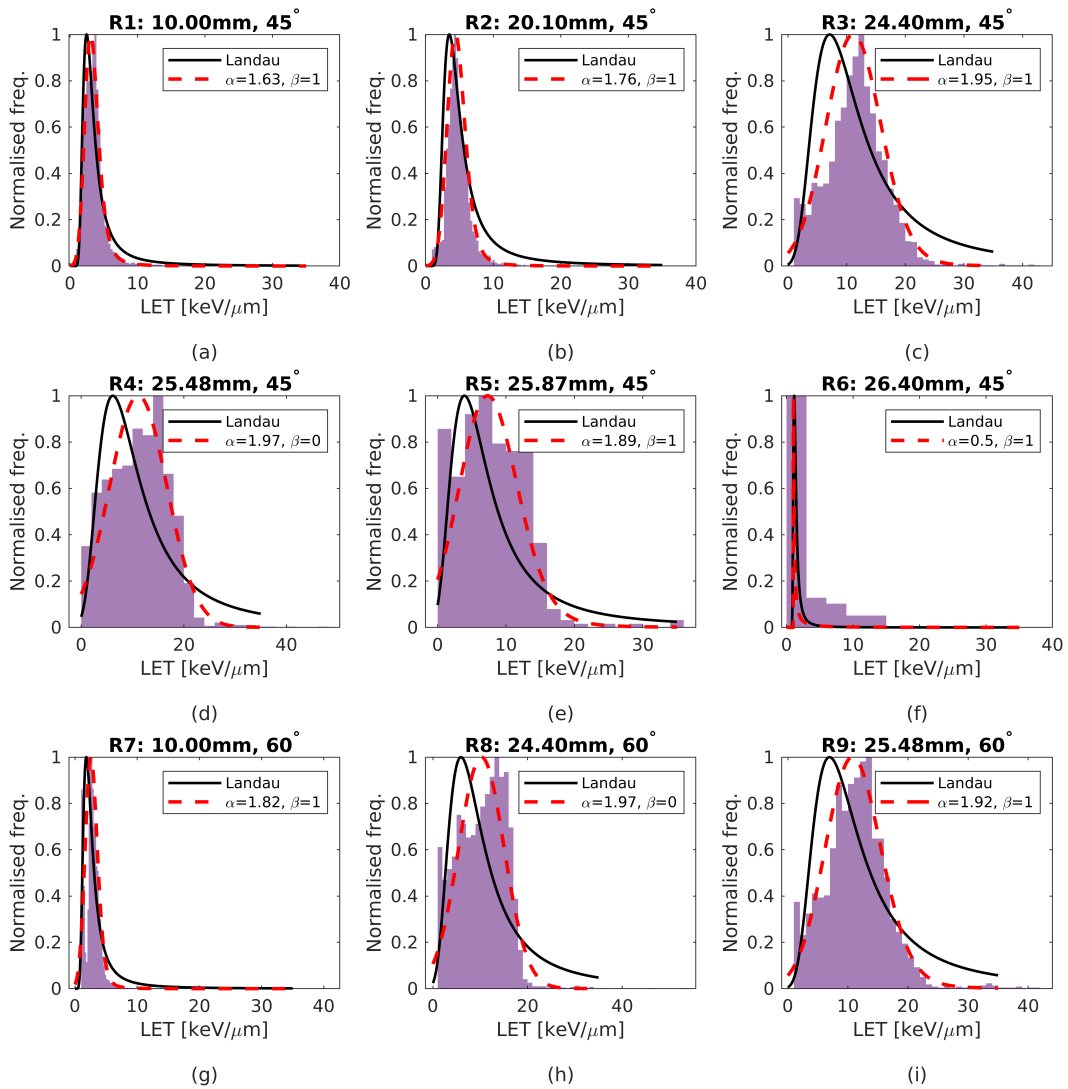


Figure C.1: (a-i) Histograms of LET contributions calculated from the MiniPIX sensor data (purple shaded area) for all runs R1 to R9, with associated Landau (black line) and general stable distribution fits (red line). Fitting procedure performed by *Python* script provided by F. Van den Heuvel.

Chapter 5. Figure C.1 shows discrepancies between the fits for some measurements, indicating a Landau fit is not always appropriate. Where a Landau was not appropriate, the MPV was taken to be the mode of the data.

Appendix D

Note on convergence for the DSACQ algorithm using SARP

In Equation 4.1 of the dynamic string-averaging CQ-method (DSACQ) paper preprint [?] we defined a Landweber-type operator $V_j : \mathcal{H} \rightarrow \mathcal{H}$ for $j = 1, 2, \dots, r$ such that

$$V_j := I - \gamma_j A_j^*(I - T_j)A_j \quad (\text{D.1})$$

for quasi-nonexpansive operators T_j and relaxation parameters $\gamma_j \in (0, \|A_j\|_F^{-2})$. Using another selection of operators $U_j : \mathcal{H} \rightarrow \mathcal{H}$, we defined $R_j : \mathcal{H} \rightarrow \mathcal{H}$ by $R_j := U_j V_j$. For our purposes we set $U_j := P_{C_j}$ and $T_j := P_{Q_j}$ where P_{C_j} represents orthogonal projection onto the set of all HDCs for structure j and P_{Q_j} represents orthogonal projection onto the set of DVCs for structure j . (U_j or T_j may alternatively be set to the identity operator if there is no constraint on the structure.) The algorithmic scheme then applies to the current iterate of the solution vector $x^{(k)}$ the operators R_j successively (the details of which are unimportant for this note).

For every operator V_j let us define N_j variations, V_{j,n_j} for $n_j = 1, 2, \dots, N_j$, which differ only in their relaxation γ_{j,n_j} . For a sufficiently large N_j -sized sub-sample of the parameter space of γ_j , choosing an operator V_{j,n_j} is equivalent to choosing V_j with the desired relaxation. This vastly increases the number of operators we are working with, but we can always choose a suitable control sequence to emulate the desired relaxation choices. This is equivalent to changing the value of γ_j on each iteration.

There is a caveat to the argument above. The number of variations we may define

for each operator is finite or, at best, countably infinite. However the parameter space for γ_j is uncountable. The argument is therefore only valid for a countable subset of all possible parameter values, which of course causes no concern in the discretised approach. Notice, regardless, that for both variations of the SARP method presented in Chapter 6, the parameter space is countable. For example, the parameter space using SARP2 is $\mathbb{Q} \cap (0, \|A_j\|_F^{-2})$.

Finally, the *multiple-operator split common fixed point problem (MOSCFPP)* (Problem 1 in [?]) is now formulated such that a feasible solution must satisfy

$$x^* \in \text{Fix}(V_{j,n_j}) = A_j^{-1}\text{Fix}(T_j) \quad \forall n_j \quad \forall j = 1, 2, \dots, r. \quad (\text{D.2})$$

This is automatically satisfied if the original problem is satisfied, namely $x^* \in \text{Fix}(V_j)$ for $j = 1, \dots, r$, because a fixed point remains a fixed point under any relaxation within the allowed range, $(0, \|A_j\|_F^{-2})$. A similar argument may be applied to the λ_j -relaxation of the operators U_j . Thus, the convergence properties that apply to the DSACQ algorithm as described in [417] also apply under the inclusion of SARP1 and SARP2 (Algorithm 2 and Algorithm 3, respectively, in Chapter 6).

Appendix E

Dose difference forest plots for all TROTS plans

Supplementary figures are presented overleaf, showing dose differences between the DSACQ solution and the archived solution obtained for all 20 TROTS plans discussed in Chapter 7.

E | Dose difference forest plots for all TROTS plans

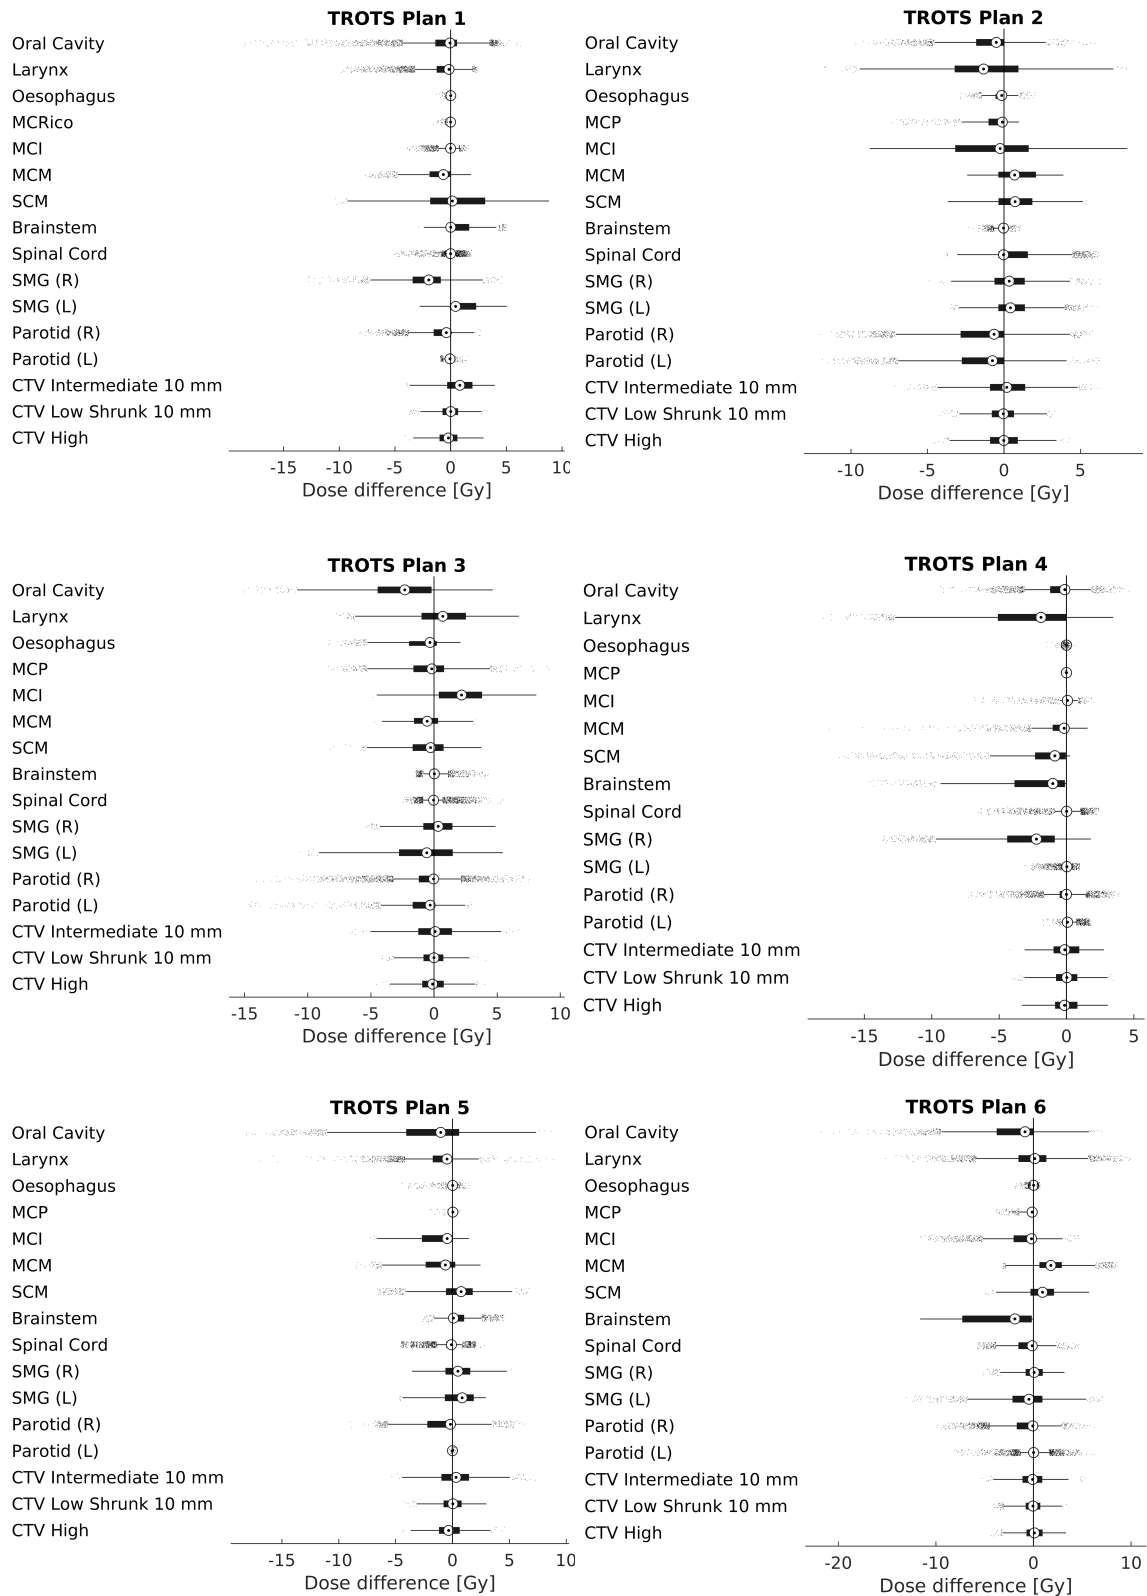


Figure E.1: Forest plots showing non-zero dose differences for TROTS Cases 1–6 distributed in the voxels of the structures of interest. Higher doses obtained using the DSACQ algorithm, when compared to the included TROTS solution, manifest as positive values.

E | Dose difference forest plots for all TROTS plans

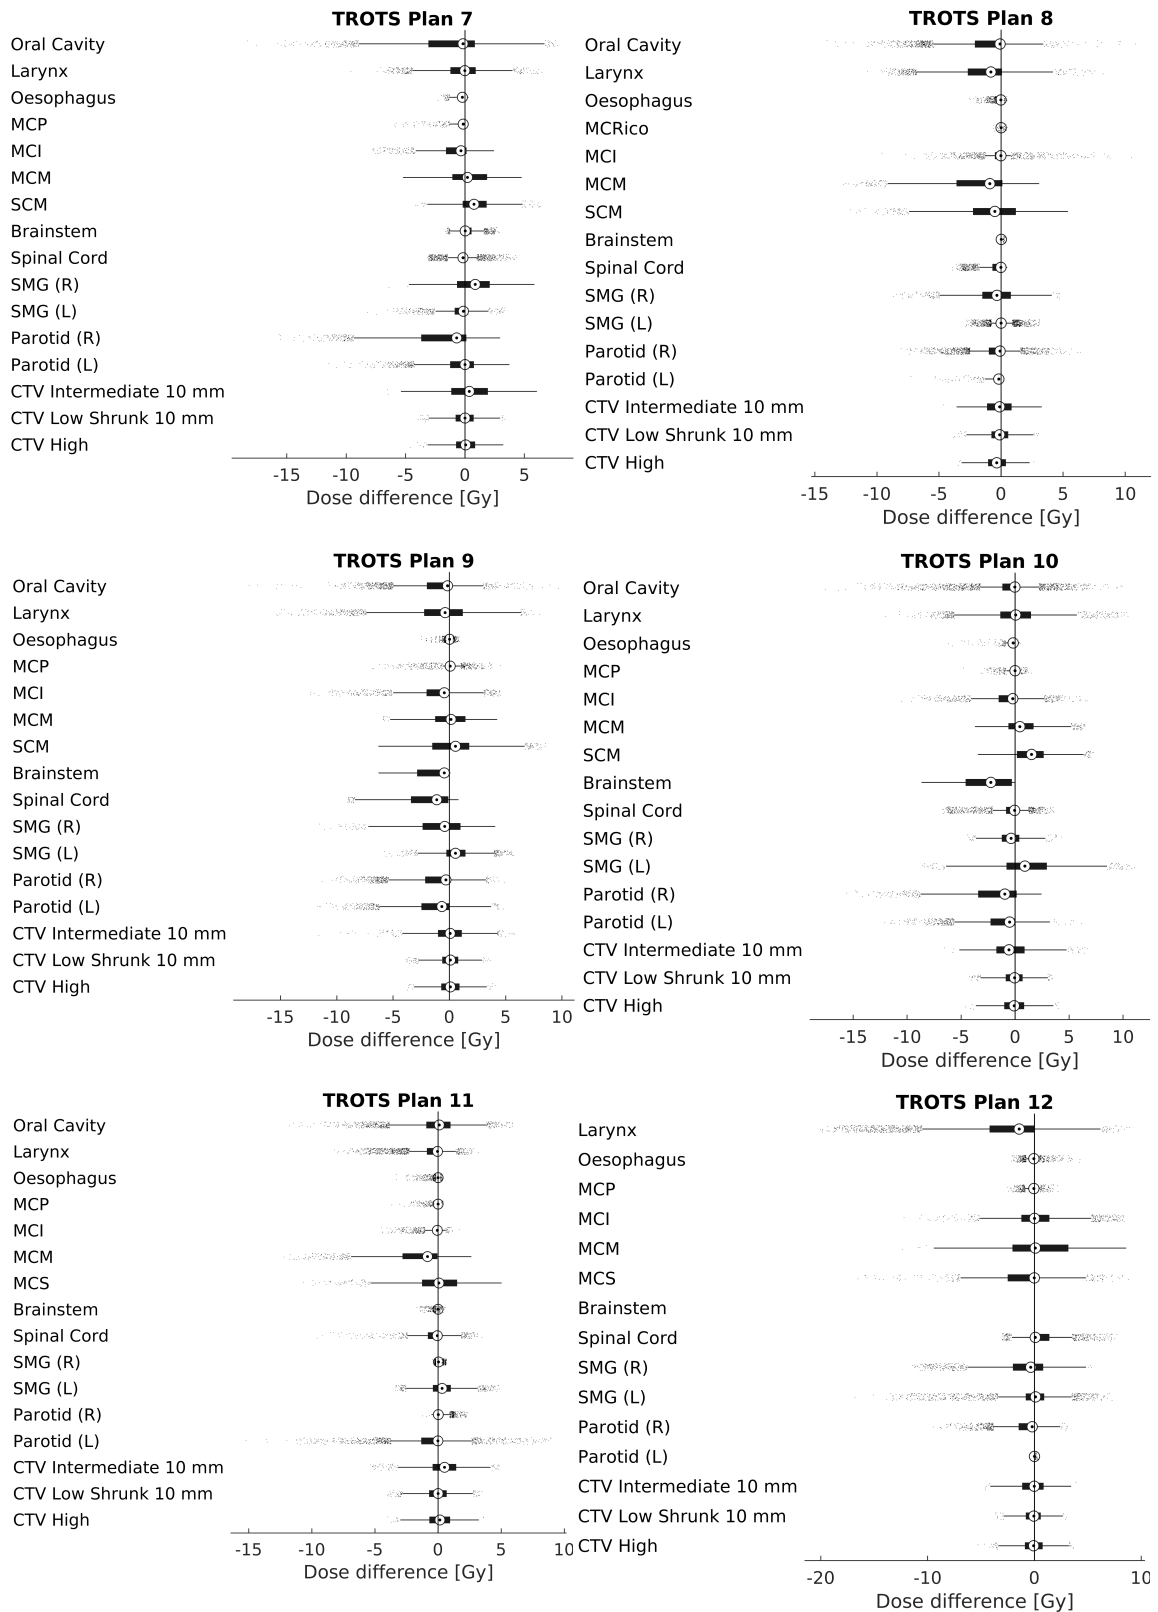


Figure E.2: Forest plots showing non-zero dose differences for TROTS Cases 7–12 distributed in the voxels of the structures of interest. Higher doses obtained using the DSACQ algorithm, when compared to the included TROTS solution, manifest as positive values.

E | Dose difference forest plots for all TROTS plans




Figure E.3: Forest plots showing non-zero dose differences for TROTS Cases 13–18 distributed in the voxels of the structures of interest. Higher doses obtained using the DSACQ algorithm, when compared to the included TROTS solution, manifest as positive values.




Figure E.4: Forest plots showing non-zero dose differences for TROTS Cases 19–20 distributed in the voxels of the structures of interest. Higher doses obtained using the DSACQ algorithm, when compared to the included TROTS solution, manifest as positive values.

Appendix F

Extensions to the SARP method

F.1 Threshold condition

In the SARP methodology presented in Chapter 6 and tested in Chapter 7, the number of voxels that violate their constraints are monitored at each algorithmic cycle. The results shown in Figure 7.7 arose from updating the relaxation parameters according to SARP1 and SARP2 every 10 cycles. Experiments were also undertaken using a threshold condition, which only updates the parameters if the total number of violations has not decreased over the past 10 cycles. This reduced most of the oscillatory behaviour in Test Case 2 and led to a lower number of violations in both Test Case 2 and Test Case 3 (Figure F.1(a, b)), however the violation count was larger for TROTS Case 1 than when no threshold was applied (Figure F.1(c)).

Enforcing the threshold condition forces the algorithm to continue with the same relaxation parameter values until there is no longer any improvement in the solution. The rate of decrease of the violations, and therefore the rate of convergence, may be slower or faster than with some other choice of parameters, so the threshold condition does not guarantee superior convergence.

F.2 Prioritising structures

In SARP1 and SARP2 the attention of the algorithm is shifted between the HDCs and the DVCs for each structure. A third suggestion, SARP3 (Algorithm 5), is made

here, which also shifts attention to the *structures* that have the most violations. In this method the mid-range of the permitted interval of the relaxation parameters for each structure is determined by the total percentage of violations relative to the mean over all structures, as demonstrated by the example in Figure F.2. The SARP2 methodology is then applied to this new parameter range. In this way, an attempt is made to simultaneously prioritise both the structures that need the most attention and the most violated constraints imposed on those structures.

Figure F.1(d-f) shows a lower number of violations reached for Test Case 2 and Test Case 3, and a similar trajectory for TROTS Case 1, when comparing SARP3 to both SARP1 and SARP2. The rate of convergence to a feasible solution is faster than the static parameter method in Test Case 2 and TROTS Case 1.

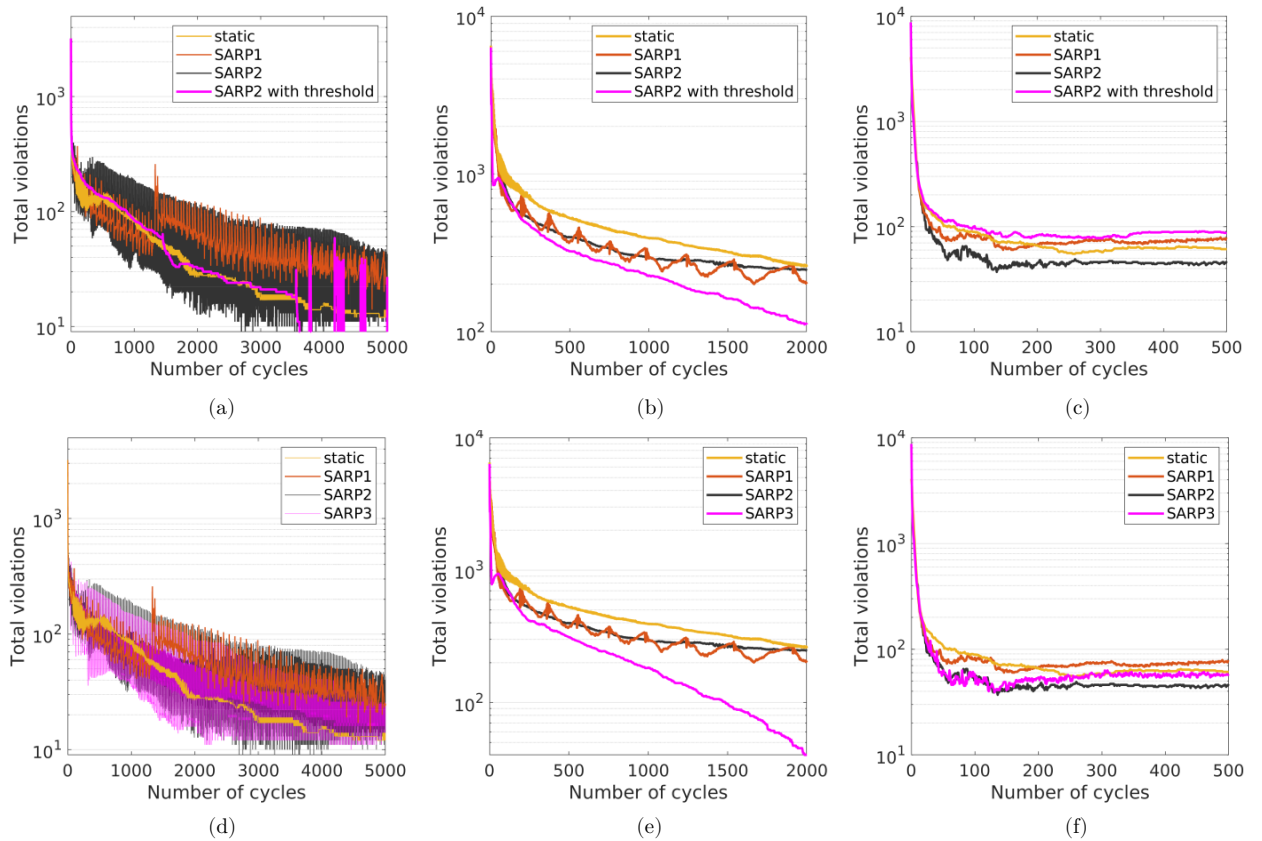


Figure F.1: Log-loss curves showing the reduction in the total number of constraint violations as the number of cycles increases, for (a, d) Test Case 2, (b, e) Test Case 3 and (c, f) TROTS Case 1. All parameter-update methods are shown as well as the static method in which the relaxation parameters are set to their mid-range. Plots (a-c) show SARP2 with a threshold condition and (d-f) show the newly formulated SARP3.

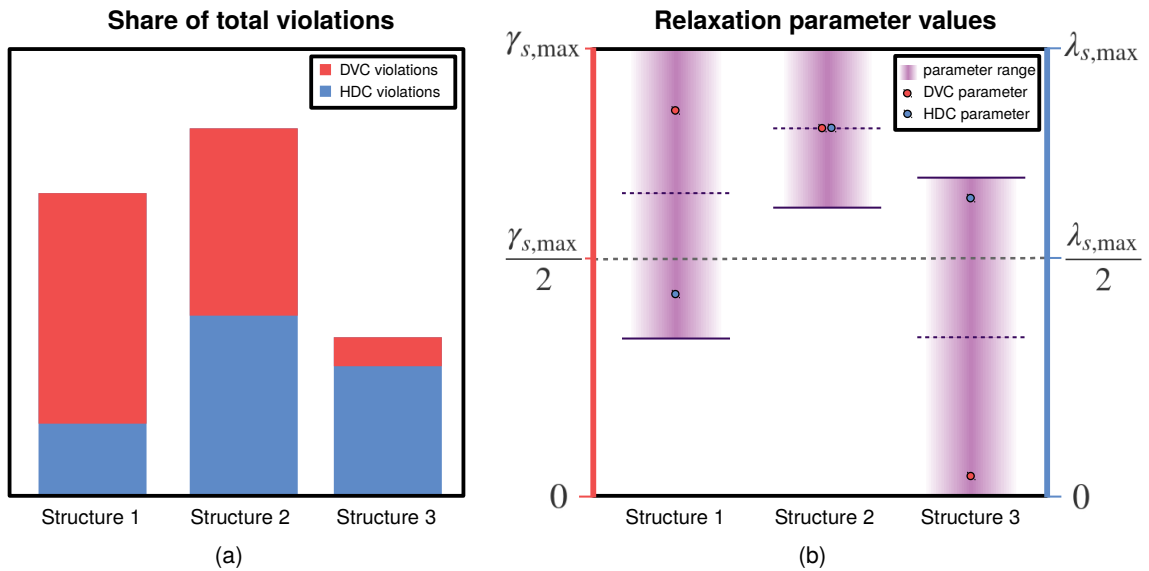


Figure F.2: Example of SARP3 implementation. (a) Total violations broken down into HDC and DVC violations for three structures. (b) Parameter ranges are adjusted such that the step-sizes of the projections associated with each structure scale with the number of violations. The SARP2 methodology is then applied within these new ranges.

Algorithm 5 Self-adaptive Relaxation Parameter method: version 3 (SARP3)

Usage: If called at the k -th iterate of the Dynamic String-Averaging CQ-Method (Algorithm 1), SARP3 executes the following:

- 1: **for each structure s do**
- 2: Count elements of $A_s x^{(k)}$ for which $x^{(k)}$ violates HDC constraints.
- 3: Set as η_C^s
- 4: Count elements of $A_s x^{(k)}$ which violate DVC constraints.
- 5: Set as η_Q^s . ▷ If s has no DVCs then $\eta_Q^s = 0$
- 6: Calculate total proportion, v_s , of voxels violating constraints:

$$v_s := \frac{1}{m_s} (\eta_C^s + \eta_Q^s) \quad (\text{F.1})$$

- 7: **end for**
- 8: Calculate the mean proportion, \bar{v} of voxels violating constraints, over all N_s structures:

$$\bar{v} := \frac{1}{N_s} \sum_{s'} v_{s'} \quad (\text{F.2})$$

- 9: **for each structure s do**
- 10: Calculate the midpoint, M :

$$M := \frac{1}{2} \left(1 + \frac{v_s - \bar{v}}{\sum_{s'} |v_{s'} - \bar{v}|} \right) \quad (\text{F.3})$$

- 11: **if s has a DVC then**
- 12: Calculate the range, R :

$$R := 2 \min(M, 1 - M) \quad (\text{F.4})$$

- 13: Update the parameters:

$$\lambda_s^{(k+1)} \leftarrow \left[M + \left(\frac{\eta_C^s}{\eta_C^s + \eta_Q^s} - \frac{1}{2} \right) R \right] \lambda_{s,\max} \quad (\text{F.5})$$

$$(\text{F.6})$$

$$\gamma_s^{(k+1)} \leftarrow \left[M + \left(\frac{\eta_Q^s}{\eta_C^s + \eta_Q^s} - \frac{1}{2} \right) R \right] \gamma_{s,\max} \quad (\text{F.7})$$

- 14: **else**
- 15: Update the parameters:

$$\lambda_s^{(k+1)} \leftarrow M \lambda_{s,\max} \quad (\text{F.8})$$

- 16: **end if**
 - 17: **end for**
-

References

- [1] Anna Gasinska. The contribution of women to radiobiology: Marie Curie and beyond. *Reports of Practical Oncology & Radiotherapy*, 21(3):250 – 258, 2016.
- [2] Cancer Research UK. Cancer risk statistics. Available at: <https://www.cancerresearchuk.org/health-professional/cancer-statistics/risk>, 2020. Accessed: 25-11-2020.
- [3] Jess Fraser, Emma Hope, Jack Anderson, Wouter Verstraete, Ravneet Sandhu, and Sean McPhail. *Chemotherapy, Radiotherapy and Surgical Tumour Resections in England*. Public Health England, 2020.
- [4] Public Health England National Cancer Registration & Analysis Service. Chemotherapy, radiotherapy and surgical tumour resections in england. Available at: <https://www.cancerdata.nhs.uk/treatments>, 2020. Accessed: 25-11-2020.
- [5] Macmillan Cancer Support. Statistics fact sheet. Available at: https://www.macmillan.org.uk/_images/cancer-statistics-factsheet_tcm9-260514.pdf, 2019. Accessed: 25-11-2020.
- [6] Josep M. Borrás, Yolande Lievens, Michael Barton, Julieta Corral, Jacques Ferlay, Freddie Bray, and Cai Grau. How many new cancer patients in europe will require radiotherapy by 2025? An ESTRO-HERO analysis. *Radiotherapy and Oncology*, 119(1):5–11, 2016.
- [7] WK Röntgen. Über eine neue art von strahlen: vorläufige mitteilung. *Sitzungsber. Phys. Med. Gesell.*, 1895.
- [8] R F Mould. Röntgen and the discovery of x-rays. *The British Journal of Radiology*, 68(815):1145–1176, 1995. PMID: 8542220.
- [9] Emil H Grubbé. Priority in the therapeutic use of x-rays. *Radiology*, 21(2):156–162, 1933.
- [10] William David Coolidge. A powerful röntgen ray tube with a pure electron discharge. *Physical Review*, 2(6):409, 1913.

- [11] Ernest O Lawrence and M Stanley Livingston. The production of high speed light ions without the use of high voltages. *Physical Review*, 40(1):19, 1932.
- [12] Henri Coutard. Roentgen therapy of epitheliomas of the tonsillar region, hypopharynx and larynx from 1920 to 1926. *Am J Roentgenol*, 28:313–331, 1932.
- [13] Henri Coutard. Principles of x ray therapy of malignant diseases. *The Lancet*, 224(5784):1–8, 1934.
- [14] C Nutting. Intensity-modulated radiotherapy (IMRT): the most important advance in radiotherapy since the linear accelerator? *The British Journal of Radiology*, 76(910):673–673, 2003. PMID: 14512325.
- [15] Søren M. Bentzen. Radiation therapy: Intensity modulated, image guided, biologically optimized and evidence based. *Radiotherapy and Oncology*, 77(3):227–230, 2005.
- [16] Thomas Bortfeld. IMRT: a review and preview. *Physics in Medicine and Biology*, 51(13):R363–R379, 2006.
- [17] A Brahme, J E Roos, and I Lax. Solution of an integral equation encountered in rotation therapy. *Physics in Medicine and Biology*, 27(10):1221–1229, 1982.
- [18] Martin D Altschuler and Yair Censor. Feasibility solutions in radiation therapy treatment planning. In *Proceedings of the eighth international conference on the use of computers in radiation therapy*, pages 220–224. Silver Spring, MD, USA: IEEE Computer Society Press, 1984.
- [19] Martin D Altschuler, William D Powlis, and Yair Censor. Teletherapy treatment planning with physician requirements included in the calculation: I. concepts and methodology. *Optimization of cancer radiotherapy*, pages 443–452, 1985.
- [20] Y Censor, WD Powlis, and MD Altschuler. On the fully discretized model for the inverse problem of radiation therapy treatment planning. In *Proceedings of the thirteenth annual northeast bioengineering conference*, volume 1, pages 211–214. Institute of Electrical and Electronics Engineers (IEEE), Inc New York, 1987.
- [21] Yair Censor, Martin D. Altschuler, and William D. Powlis. A computational solution of the inverse problem in radiation-therapy treatment planning. *Applied Mathematics and Computation*, 25(1):57–87, 1988.
- [22] S Webb. Optimisation of conformal radiotherapy dose distribution by simulated annealing. *Physics in Medicine and Biology*, 34(10):1349–1370, 1989.

- [23] S Webb. Optimization by simulated annealing of three-dimensional, conformal treatment planning for radiation fields defined by a multileaf collimator: II. inclusion of two-dimensional modulation of the x-ray intensity. *Physics in Medicine and Biology*, 37(8):1689–1704, 1992.
- [24] C.Clifton Ling, Chandra Burman, Chen S. Chui, Gerald J. Kutcher, Stephen A. Leibel, Thomas LoSasso, Radhe Mohan, Thomas Bortfeld, Larry Reinstein, Spiridon Spirou, X.H. Wang, Qiuwen Wu, Michael Zelefsky, and Zvi Fuks. Conformal radiation treatment of prostate cancer using inversely-planned intensity-modulated photon beams produced with dynamic multileaf collimation. *International Journal of Radiation Oncology • Biology, • Physics*, 35(4):721 – 730, 1996.
- [25] Joseph John Thomson. XI. cathode rays. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 44(269):293–316, 1897.
- [26] E. Rutherford. LXXIX. the scattering of α and β particles by matter and the structure of the atom. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 21(125):669–688, 1911.
- [27] Ernest Rutherford. LIV. collision of α particles with light atoms. IV. an anomalous effect in nitrogen. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 37(222):581–587, 1919.
- [28] R R Wilson. Radiological use of fast protons. *Radiology*, ;47(5):487–91, 1946.
- [29] William Henry Bragg and Richard Kleeman. XXXIX. on the α particles of radium, and their loss of range in passing through various atoms and molecules. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 10(57):318–340, 1905.
- [30] Xiufang Tian, Kun Liu, Yong Hou, Jian Cheng, and Jiandong Zhang. The evolution of proton beam therapy: Current and future status. *Molecular and clinical oncology*, 8(1):15–21, 2018.
- [31] Martin Jermann. Particle therapy statistics in 2014. *International Journal of Particle Therapy*, 2(1):50–54, 2015.
- [32] Martin Jermann. Particle therapy patient statistics (per end of 2016). Available at: https://www.ptcog.ch/archive/patient_statistics/Patientstatistics-updateDec2016.pdf, Dec 2016. Accessed: 25-11-2020
- [33] Martin Jermann. Particle therapy patient statistics (per end of 2019). Available at: <https://www.ptcog.ch/index.php/ptcog-patient-statistics>, Nov 2020. Accessed: 25-11-2020

- [34] Particle therapy facilities in clinical operation (last update: Sep 2020).
<https://www.ptcog.ch/index.php/facilities-in-operation>, 2020. Accessed: 25-11-2020.
- [35] Particle therapy facilities under construction (update nov 2020).
<https://www.ptcog.ch/index.php/facilities-under-construction>, 2020. Accessed: 25-11-2020.
- [36] Frank Herbert Attix. Introduction to Radiological Physics and Radiation Dosimetry. John Wiley & Sons, 1986.
- [37] A. Sibtain, A. Morgan, and N. MacDougall. *Physics for Clinical Oncology*. Radiotherapy in Practice. OUP Oxford, 2012.
- [38] P. Hoskin. *External Beam Therapy*. Radiotherapy in Practice. OUP Oxford, 2012.
- [39] N. Bohr. I. on the constitution of atoms and molecules. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 26(151):1–25, 1913.
- [40] Wayne D Newhauser and Rui Zhang. The physics of proton therapy. *Physics in Medicine and Biology*, 60(8):R155–R209, 2015.
- [41] Andrew J. Wroe, Iwan M. Cornelius, and Anatoly B. Rosenfeld. The role of nonelastic reactions in absorbed dose distributions from therapeutic proton beams in different medium. *Medical Physics*, 32(1):37–41, 2005.
- [42] L. Verhey, H. Blattman, P. M. Deluca, and D. Miller. Report 59. *Journal of the International Commission on Radiation Units and Measurements*, os30(2), 2016.
- [43] NF Mott. The scattering of fast electrons by atomic nuclei. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 124(794):425–442, 1929.
- [44] G. Moliere. Theorie der streuung schneller geladener teilchen ii. *Z. Naturforsch.*, Ba:78, 1948.
- [45] H. A. Bethe. Molière’s theory of multiple scattering. *Phys. Rev.*, 89:1256–1266, 1953.
- [46] Bruno Rossi and Kenneth Greisen. Cosmic-ray theory. *Rev. Mod. Phys.*, 13:240–309, 1941.
- [47] Leonard Eyges. Multiple scattering with energy loss. *Phys. Rev.*, 74:1534–1535, 1948.
- [48] S. Goudsmit and J. L. Saunderson. Multiple scattering of electrons. *Phys. Rev.*, 57:24–29, 1940.

- [49] S. Goudsmit and J. L. Saunderson. Multiple scattering of electrons. ii. *Phys. Rev.*, 58:36–42, 1940.
- [50] Bernard Gottschalk. On the scattering power of radiotherapy protons. *Medical Physics*, 37(1):352–367, 2010.
- [51] Mark D. Brooke and Scott N. Penfold. An inhomogeneous most likely path formalism for proton computed tomography. *Physica Medica*, 70:184 – 195, 2020.
- [52] H. Bethe. Zur theorie des durchgangs schneller korpuskularstrahlen durch materie. *Annalen der Physik*, 397(3):325–400, 1930.
- [53] F. Bloch. Zur bremsung rasch bewegter teilchen beim durchgang durch materie. *Annalen der Physik*, 408(3):285–320, 1933.
- [54] M.J. Berger, J.S. Coursey, M.A. Zucker and J. Chang. Available at: <http://physics.nist.gov/Star>. National Institute of Standards and Technology, Gaithersburg, MD., 2005. Accessed: 03-09-2020.
- [55] Thomas Bortfeld and Wolfgang Schlegel. An analytical approximation of depth - dose distributions for therapeutic proton beams. *Physics in Medicine and Biology*, 41(8):1331–1339, 1996.
- [56] Thomas Bortfeld. An analytical approximation of the bragg curve for therapeutic proton beams. *Medical physics*, 24(12):2024–2033, 1997.
- [57] David J. Thomas. ICRU report 85: fundamental quantities and units for ionizing radiation. *Radiation Protection Dosimetry*, 150(4):550–552, 2012.
- [58] C Grassberger and H Paganetti. Elevated LET components in clinical proton beams. *Physics in Medicine and Biology*, 56(20):6677–6691, 2011.
- [59] Leszek Grzanka, Oscar Ardenfors, and Niels Bassler. Monte Carlo simulations of spatial LET distributions in clinical proton beams. *Radiation Protection Dosimetry*, 180(1-4):296–299, 2018.
- [60] M A Cortés-Giraldo and A Carabe. A critical study of different Monte Carlo scoring methods of dose average linear-energy-transfer maps calculated in voxelized geometries irradiated with clinical proton beams. *Physics in Medicine and Biology*, 60(7):2645–2669, 2015.

- [61] Dal A Granville and Gabriel O Sawakuchi. Comparison of linear energy transfer scoring techniques in Monte Carlo simulations of proton beams. *Physics in Medicine and Biology*, 60(14):N283–N291, 2015.
- [62] Fada Guan, Christopher Peeler, Lawrence Bronk, Changran Geng, Reza Taleei, Sharmalee Randeniya, Shuaiping Ge, Dragan Mirkovic, David Grosshans, Radhe Mohan, and Uwe Titt. Analysis of the track- and dose-averaged LET and LET spectra in proton therapy using the Geant4 Monte Carlo code. *Medical Physics*, 42(11):6234–6247, 2015.
- [63] Dirk Wagenaar, Linh T Tran, Arturs Meijers, Gabriel Guterres Marmitt, Kevin Souris, David Bolst, Benjamin James, Giordano Biasi, Marco Povoli, Angela Kok, Erik Traneus, Marc-Jan van Goethem, Johannes A Langendijk, Anatoly B Rosenfeld, and Stefan Both. Validation of linear energy transfer computed in a Monte Carlo dose engine of a commercial treatment planning system. *Physics in Medicine & Biology*, 65(2):025006, 2020.
- [64] Amato J Giaccia Eric J Hall. *Radiobiology for the Radiologist*. Wolters Kluwer Health, 7 edition, 2011.
- [65] M. Joiner and A. van der Kogel. *Basic Clinical Radiobiology Fourth Edition*. Hodder Education, 2009.
- [66] Keith W. Caldecott. Single-strand break repair and genetic disease. *Nature Reviews Genetics*, 9:619–631, 2008.
- [67] J. F. Ward. Biochemistry of DNA lesions. 104(2):S103–S111, 1985.
- [68] Jean Cadet, Kelvin JA Davies, Marisa HG Medeiros, Paolo Di Mascio, and J Richard Wagner. Formation and repair of oxidatively generated damage in cellular DNA. *Free Radical Biology and Medicine*, 107:13–34, 2017.
- [69] George DD Jones and Peter O’Neill. Kinetics of radiation-induced strand break formation in single-stranded pyrimidine polynucleotides in the presence and absence of oxygen; a time-resolved light-scattering study. *International Journal of Radiation Biology*, 59(5):1127–1145, 1991.
- [70] Wendy Knapp Pogozelski and Thomas D Tullius. Oxidative strand scission of nucleic acids: routes initiated by hydrogen abstraction from the sugar moiety. *Chemical Reviews*, 98(3):1089–1108, 1998.
- [71] P Wardman. The importance of radiation chemistry to radiation and free radical biology (the 2008 silvanus thompson memorial lecture). *The British Journal of Radiology*, 82(974):89–104, 2009.

- [72] Peter Wardman. Time as a variable in radiation biology: the oxygen effect. *Radiation Research*, 185(1):1–3, 2016.
- [73] Jake C Forster, Michael JJ Douglass, Wendy M Phillips, and Eva Bezak. Monte Carlo simulation of the oxygen effect in DNA damage induction by ionizing radiation. *Radiation research*, 190(3):248–261, 2018.
- [74] H. Nikjoo, P. O’Neill, W. E. Wilson, and D. T. Goodhead. Computational Approach for Determining the Spectrum of DNA Damage Induced by Ionizing Radiation. *Radiation Research*, 156(5):577 – 583, 2001.
- [75] H. Nikjoo and D. T. Goodhead. Track structure analysis illustrating the prominent role of low-energy electrons in radiobiological effects of low-LET radiations. *Physics in Medicine and Biology*, 36(2):229 – 238, 1991.
- [76] V. Michalik and D. Frankenberg. Simple and complex double-strand breaks induced by electrons. *International Journal of Radiation Biology*, 66(5):467–470, 1994.
- [77] D.E. Charlton, H. Nikjoo J.L. Humm. Calculation of Initial Yields of Single- and Double-strand Breaks in Cell Nuclei from Electrons, Protons and Alpha Particles. *International Journal of Radiation Biology*, 56(1):1-19, 1989.
- [78] D. T. Goodhead H. Nikjoo, P. O’Neill and M. Terrissol. Computational modelling of low-energy electron-induced DNA damage by early physical and chemical events. *International Journal of Radiation Biology*, 71(5):467–483, 1997.
- [79] H. Nikjoo, P. O’Neill, M. Terrissol and D.T. Goodhead. Quantitative modelling of DNA damage using Monte Carlo track structure method. *Radiation and Environmental Biophysics*, 38:31–38, 1999.
- [80] Agnes Schipler and George Iliakis. DNA double-strand-break complexity levels and their possible contributions to the probability for error-prone processing and repair pathway choice. *Nucleic acids research*, 41(16):7589–7605, 2013.
- [81] Lepakshi Ranjha, Sean M Howard, and Petr Cejka. Main steps in DNA double-strand break repair: an introduction to homologous recombination and related processes. *Chromosoma*, 127(2):187–214, 2018.
- [82] Nicholas R Pannunzio, Go Watanabe, and Michael R Lieber. Nonhomologous DNA end-joining for repair of DNA double-strand breaks. *Journal of Biological Chemistry*, 293(27):10512–10523, 2018.

- [83] Atsushi Shibata and Penny A. Jeggo. Roles for 53BP1 in the repair of radiation-induced DNA double strand breaks. *DNA Repair*, 93:102915, 2020. Tribute to Samuel H. Wilson: Shining Light on Base Excision DNA Repair.
- [84] A. Munshi, M. Hobbs, and R. Meyn. Clonogenic cell survival assay. *Methods in Molecular Medicine*, 110:21–8, 2005.
- [85] Stephen Joseph McMahon. The linear quadratic model: usage, interpretation and challenges. *Physics in Medicine & Biology*, 64(1):01TR01, 2018.
- [86] Joel S. Bedford. Sublethal damage, potentially lethal damage, and chromosomal aberrations in mammalian cells exposed to ionizing radiations. *International Journal of Radiation Oncology • Biology, • Physics*, 21(6):1457 – 1469, 1991.
- [87] W.K. Sinclair. The shape of radiation survival curves of mammalian cells cultured in vitro. *Biophysical Aspects of Radiation Quality*, 21–43, 1966.
- [88] Albrecht M. Kellerer and Harald H. Rossi. RBE and the primary mechanism of radiation action. *Radiation Research*, 47(1):15–34, 1971.
- [89] K H Chadwick and H P Leenhouts. A molecular theory of cell survival. *Physics in Medicine and Biology*, 18(1):78–87, 1973.
- [90] J. A. Crowther and William Bate Hardy. Some considerations relative to the action of x-rays on tissue cells. *Proceedings of the Royal Society of London. Series B, Containing Papers of a Biological Character*, 96(674):207–211, 1924.
- [91] Theodore T. Puck and Philip I. Marcus. Action of x-rays on mammalian cells. *Journal of Experimental Medicine*, 103(5):653–666, 1956.
- [92] Tikvah Alper, N. E. Gillies, and M. M. Elkind. The sigmoid survival curve in radiobiology. *Nature*, 186(4730):1062–1063, 1960.
- [93] G.J. Neary. Chromosome aberrations and the theory of rbe. *International Journal of Radiation Biology and Related Studies in Physics, Chemistry and Medicine*, 9(5):477–502, 1965.
- [94] C Regaud and T Nogier. Sterilization rontgenienne totale et definitive, sans radiodermite, des testicules du belier adulte: conditions de sa realisation. *Compt Rend Soc de Biol*, 70:202–203, 1911.

- [95] Philip P. Connell and Samuel Hellman. Advances in radiotherapy and implications for the next century: A historical perspective. *Cancer Research*, 69(2):383–392, 2009.
- [96] Howard D. Thames. On the origin of dose fractionation regimens in radiotherapy. *Seminars in Radiation Oncology*, 2(1):3 – 9, 1992.
- [97] W. S. MacComb and Edith H. Quimby. The rate of recovery of human skin from the effects of hard or soft roentgen rays or gamma rays. *Radiology*, 27(2):196–207, 1936.
- [98] Edith H. Quimby and W. S. MacComb. Further studies on the rate of recovery of human skin from the effects of roentgen- or gamma-ray irradiation. *Radiology*, 29(3):305–312, 1937.
- [99] J. F. Fowler, R. L. Morgan, J. A. Silvester, D. K. Bewley, and B. A. Turner. Experiments with fractionated x-ray treatment of the skin of pigs. i—fractionation up to 28 days. *The British Journal of Radiology*, 36(423):188–196, 1963.
- [100] J. F. Fowler, D. K. Bewley, R. L. Morgan, and J. A. Silvester. Experiments with fractionated x-irradiation of the skin of pigs II. – fractionation up to five days. *The British Journal of Radiology*, 38(448):278–284, 1965.
- [101] M. Strandqvist. Studien über die kumulative wirkung der röntgenstrahlen bei fraktionierung. *Acta Radiol*, 55(SUPPL.):1–300, 1944.
- [102] William Bond. Radium beam dose-time relations. *Journal of the Faculty of Radiologists*, 3(2):111 – 124, 1951.
- [103] F. Ellis. The relationship of biological effect to dose-time fractionation factors in radiotherapy. *Current Topics in Radiation Research*, 4:357–397, 1968.
- [104] Frank Ellis. Dose, time and fractionation: A clinical hypothesis. *Clinical Radiology*, 20(1):1 – 7, 1969.
- [105] G. W. Barendsen. Dose fractionation, dose rate and iso-effect relationships for normal tissue responses. *International Journal of Radiation Oncology • Biology, • Physics*, 8(11):1981–1997, 1982.
- [106] H. Withers. The four R's of radiotherapy. Advances in radiation biology. Vol. 5. Elsevier, 241–271, 1975.
- [107] G. Gordon Steel, T.J. McMillan, and J.H. Peacock. The 5 Rs of radiobiology. *International Journal of Radiation Biology*, 56(6):1045–1048, 1989.

- [108] Howard D Thames Jr., H. Rodney Withers, Lester J. Peters, and Gilbert H. Fletcher. Changes in early and late radiation responses with altered dose fractionation: Implications for dose-survival relationships. *International Journal of Radiation Oncology • Biology, • Physics*, 8(2):219–226, 1982.
- [109] Douglas Hanahan and Robert A. Weinberg. The hallmarks of cancer. *Cell*, 100(1):57–70, 2000.
- [110] K. Harrington, P. Jankowska, and M. Hingorani. Molecular biology for the radiation oncologist: the 5rs of radiobiology meet the hallmarks of cancer. *Clinical Oncology*, 19(8):561–571, 2007.
- [111] André Wambersie. Rbe, reference RBE and clinical rbe: Applications of these concepts in hadron therapy. *Strahlentherapie und Onkologie*, 175(2):39, 1999.
- [112] Zuofeng Li. Prescribing, recording, and reporting proton-beam therapy: International commission on radiation units and measurements report 78. *J ICRU*, 7:210, 2007.
- [113] Harald Paganetti, Andrzej Niemierko, Marek Ancukiewicz, Leo E Gerweck, Michael Goitein, Jay S Loeffler, and Herman D Suit. Relative biological effectiveness (RBE) values for proton beam therapy. *International Journal of Radiation Oncology • Biology, • Physics*, 53(2):407–421, 2002.
- [114] Harald Paganetti and Peter Van Luijk. Biological Considerations When Comparing Proton Therapy With Photon Therapy. *Seminars in Radiation Oncology*, 23(2):77–87, 2013.
- [115] Harald Paganetti. Relative biological effectiveness (RBE) values for proton beam therapy. Variations as a function of biological endpoint, dose, and linear energy transfer. *Physics in Medicine and Biology*, 59(22):R419–R472, 2014.
- [116] Francesco Tommasino and Marco Durante. Proton radiobiology. *Cancers*, 7(1):353–381, 2015.
- [117] Tracy Underwood and Harald Paganetti. Variable Proton Relative Biological Effectiveness: How Do We Move Forward?. *International Journal of Radiation Oncology • Biology, • Physics* 95(1):56–58, 2016.
- [118] Radhe Mohan, Christopher R. Peeler, Fada Guan, Lawrence Bronk, Wenhua Cao, and David R. Grosshans. Radiobiological issues in proton therapy. *Acta Oncologica*, 56(11):1367–1373, 2017.

- [119] Bleddyn Jones. Proton radiobiology and its clinical implications. *Ecancermedicalscience*, 11:1–11, 2017.
- [120] ICRU Prescribing, recording, and reporting photon-beam intensity-modulated radiation therapy (IMRT). *ICRU report*, 83(10):27–40, 2010.
- [121] Douglas Jones. ICRU report 50—prescribing, recording and reporting photon beam therapy. *Medical Physics*, 21(6):833–834.
- [122] T. Landberg, J. Chavaudra, J. Dobbs, J. P. Gerard, G. Hanks, J. C. Horiot, K. A. Johansson, T. Möller, J. Purdy, N. Suntharalingam, and H. Svensson. Report 62. *Journal of the International Commission on Radiation Units and Measurements*, os32(1), 1999.
- [123] Intensity Modulated Radiation Therapy Collaborative Working Group. Intensity-modulated radiotherapy: current status and issues of interest. *International Journal of Radiation Oncology • Biology • Physics*, 51(4):880–914, 2001.
- [124] T. S. Hong, M. A. Ritter, W. A. Tomé, and P. M. Harari. Intensity-modulated radiation therapy: emerging cancer treatment technology. *British Journal of Cancer*, 92(10):1819–1824, 2005.
- [125] R. Mohan, Q. Wu, M. Manning, and R. Schmidt-Ullrich. Radiobiological considerations in the design of fractionation strategies for intensity-modulated radiation therapy of head and neck cancers. *International Journal of Radiation Oncology • Biology • Physics*, 46:619–630, 2000.
- [126] C. M. Nutting, D. J. Convery, V. P. Cosgrove, C. Rowbottom, A. R. Padhani, S. Webb, and D. P. Dearnaley. Reduction of small and large bowel irradiation using an optimized intensity-modulated pelvic radiotherapy technique in patients with prostate cancer. *International Journal of Radiation Oncology • Biology • Physics*, 48:649–656, 2000.
- [127] C. E. Coles, A. M. Moody, C. B. Wilson, and N. G. Burnet. Reduction of radiotherapy-induced late complications in early breast cancer: the role of intensity-modulated radiation therapy and partial breast irradiation: Part I - normal tissue complications. *Clinical Oncology*, 17:16–24, 2005.
- [128] A. Narayana, J. Yamada, S. Berry, P. Shah, M. Hunt, P. H. Gutin, and S. A. Leibel. Intensity-modulated radiotherapy in high-grade gliomas: Clinical and dosimetric results. *International Journal of Radiation Oncology • Biology • Physics*, 64:892–897, 2006.

- [129] S. Sura, V. Gupta, E. Yorke, A. Jackson, H. Amols, and K. E. Rosenzweig. Intensity-modulated radiation therapy (IMRT) for inoperable non-small cell lung cancer: The Memorial Sloan-Kettering Cancer center (MSKCC) experience. *Radiotherapy and Oncology*, 87:17–23, 2008.
- [130] Leire Arbea, Luis Isaac Ramos, Rafael Martínez-Monge, Marta Moreno, and Javier Aristu. Intensity-modulated radiation therapy (IMRT) vs. 3D conformal radiotherapy (3DCRT) in locally advanced rectal cancer (LARC): dosimetric comparison and clinical implications. *Radiation Oncology*, 5(1):17, 2010.
- [131] XiaoShen Wang and Avraham Eisbruch. IMRT for head and neck cancer: reducing xerostomia and dysphagia. *Journal of radiation research*, 57 Suppl 1(Suppl 1):i69–i75, 2016.
- [132] Piet Dirix and Sandra Nuyts. Evidence-based organ-sparing radiotherapy in head and neck cancer. *The Lancet Oncology*, 11(1):85–91, 2010.
- [133] Feng Teng, Wenjun Fan, Yanrong Luo, Zhongjian Ju, Hanshun Gong, Ruigang Ge, Fang Tong, Xinxin Zhang, and Lin Ma. Reducing xerostomia by comprehensive protection of salivary glands in intensity-modulated radiation therapy with helical tomotherapy technique for head-and-neck cancer patients: A prospective observational study. *BioMed Research International*, 2019:2401743–2401743, 2019.
- [134] A. J. Lomax, E. Pedroni, H. P. Rutz, and G. Goitein. The clinical potential of intensity modulated proton therapy. *Zeitschrift für Medizinische Physik*, 14:147–152, 2004.
- [135] X. Zhang, Y. Li, X. Pan, L. Xiaoqiang, R. Mohan, R. Komaki, J. D. Cox, and J. Y. Chang. Intensity-modulated proton therapy reduces the dose to normal tissue compared with intensity-modulated radiation therapy or passive scattering proton therapy and enables individualized radical radiotherapy for extensive stage IIIB non-small-cell lung cancer: A virtual clinical study. *International Journal of Radiation Oncology • Biology • Physics*, 77:357–366, 2010.
- [136] J. Welsh, D. Gomez, M. B. Palmer, B. A. Riley, A. V. Mayankkumar, R. Komaki, L. Dong, X. R. Zhu, A. Likhacheva, Z. Liao, W. L. Hofstetter, J. A. Ajani, and J. D. Cox. Intensity-modulated proton therapy further reduces normal tissue exposure during definitive therapy for locally advanced distal esophageal tumors: A dosimetric study. *International Journal of Radiation Oncology • Biology • Physics*, 81:1336–1342, 2011.
- [137] H. M. Kooy and C. Grassberger. Intensity modulated proton therapy. *British Journal of Radiology*, 88:20150195, 2015.

- [138] L. V. van Dijk, R. J. H. M. Steenbakkers, B. ten Haken, H. P. van der Laan, A. A. van't Veld, J. A. Langendijk, and E. W. Korevaar. Robust intensity modulated proton therapy (IMPT) increases estimated clinical benefit in head and neck cancer patients. *PLoS ONE*, 11:1–14, 2016.
- [139] Hui Liu and Joe Y. Chang. Proton therapy in clinical practice. *Chinese Journal of Cancer*, 30(5):315–326, 2011.
- [140] Joao Seco, Daniel Robertson, Alexei Trofimov, and Harald Paganetti. Breathing interplay effects during proton beam scanning: simulation and statistical analysis. *Physics in Medicine and Biology*, 54(14):N283–N294, 2009.
- [141] Laleh Kardar, Yupeng Li, Xiaoqiang Li, Heng Li, Wenhua Cao, Joe Y. Chang, Li Liao, Ronald X. Zhu, Narayan Sahoo, Michael Gillin, Zhongxing Liao, Ritsuko Komaki, James D. Cox, Gino Lim, and Xiaodong Zhang. Evaluation and mitigation of the interplay effects of intensity modulated proton therapy for lung cancer in a clinical setting. *Practical Radiation Oncology*, 4(6):e259–e268, 2014.
- [142] A Lomax. Intensity modulation methods for proton radiotherapy. *Physics in Medicine and Biology*, 44(1):185–205, 1999.
- [143] Allen Holder. Designing radiotherapy plans with elastic constraints and interior point methods. *Health Care Management Science*, 6(1):5–16, 2003.
- [144] David Sonderman and Philip G. Abrahamson. Radiotherapy treatment design using mathematical programming models. *Operations Research*, 33(4):705–725, 1985.
- [145] Wolfgang Dörr and Albert J van der Kogel. The volume effect in radiotherapy. *Basic Clinical Radiobiology. 4th ed. London: Hodder Arnold*, pages 191–206, 2009.
- [146] W. Dörr and M. Schmidt. 7.05 - normal tissue radiobiology. In Anders Brahme, editor, *Comprehensive Biomedical Physics*, pages 75 – 95. Elsevier, Oxford, 2014.
- [147] W. Dörr. Radiobiology of tissue reactions. *Annals of the ICRP*, 44(1_suppl):58–68, 2015.
- [148] H. Rodney Withers, Jeremy M.G. Taylor, and Boguslaw Maciejewski. Treatment volume and tissue tolerance. *International Journal of Radiation Oncology • Biology, • Physics*, 14(4):751 – 759, 1988.
- [149] S. L. Gulliford, K. Foo, R. C. Morgan, E. G. Aird, A. M. Bidmead, H. Critchley, P. M. Evans, S. Gianolini, W. P. Mayles, A. R. Moore, B. Sanchez-Nieto, M. Partridge, M. R.

- Sydes, S. Webb, and D. P. Dearnaley. Dose-volume constraints to reduce rectal side effects from prostate radiotherapy: Evidence from MRC RT01 trial ISRCTN 47772397. *International Journal of Radiation Oncology • Biology • Physics*, 76:747–754, 2010.
- [150] Daniel L.P. Holyoake, Mike Partridge, and Maria A. Hawkins. Systematic review and meta-analysis of small bowel dose–volume and acute toxicity in conventionally-fractionated rectal cancer radiotherapy. *Radiotherapy and Oncology*, 138:38 – 44, 2019.
- [151] Lawrence B. Marks, Ellen D. Yorke, Andrew Jackson, Randall K. Ten Haken, Louis S. Constine, Avraham Eisbruch, Søren M. Bentzen, Jiho Nam, and Joseph O. Deasy. Use of normal tissue complication probability models in the clinic. *International Journal of Radiation Oncology • Biology • Physics*, 76(3):S10–S19, Mar 2010.
- [152] Søren M. Bentzen, Louis S. Constine, Joseph O. Deasy, Avi Eisbruch, Andrew Jackson, Lawrence B. Marks, Randall K. Ten Haken, and Ellen D. Yorke. Quantitative analyses of normal tissue effects in the clinic (quantec): An introduction to the scientific issues. *International Journal of Radiation Oncology • Biology • Physics*, 76(3):S3–S9, Mar 2010.
- [153] L.S. Constine, C.M. Ronckers, C.-H. Hua, A. Olch, L.C.M. Kremer, A. Jackson, and S.M. Bentzen. Pediatric normal tissue effects in the clinic (pentec): An international collaboration to analyse normal tissue radiation dose–volume response relationships for paediatric cancer patients. *Clinical Oncology*, 31(3):199 – 207, 2019.
- [154] Alan E. Nahum and Julien Uzan. (Radio)biological optimization of external-beam radiotherapy. *Computational and Mathematical Methods in Medicine*, 2012:1–13, 2012.
- [155] Andrzej Niemierko. Reporting and analyzing dose distributions: a concept of equivalent uniform dose. *Medical physics*, 24(1):103–110, 1997.
- [156] Andrzej Niemierko. A generalized concept of equivalent uniform dose (EUD). *Med Phys*, 26(6):1100, 1999.
- [157] C. Burman, G.J. Kutcher, B. Emami, and M. Goitein. Fitting of normal tissue tolerance data to an analytic function. *International Journal of Radiation Oncology • Biology • Physics*, 21(1):123 – 135, 1991. Three-Dimensional Photon Treatment Planning Report of the Collaborative Working Group on the Evaluation of Treatment Planning for External Photon Beam Radiotherapy.
- [158] Hiram A. Gay and Andrzej Niemierko. A free program for calculating EUD-based NTCP and TCP in external beam radiotherapy. *Physica Medica*, 23(3-4):115–125, 2007.

- [159] V. A. Semenenko and R. D. Stewart. A fast Monte Carlo algorithm to simulate the spectrum of DNA damages formed by ionizing radiation. *Radiation Research*, 161(4):451 – 457, 2004.
- [160] V. A. Semenenko and R. D. Stewart. Fast Monte Carlo simulation of DNA damage formed by electrons and light ions. *Physics in Medicine and Biology*, 51(7):1693–1706, 2006.
- [161] J. Perl, J. Shin, J. Schumann, B. Faddegon, and H. Paganetti. TOPAS: An innovative proton Monte Carlo platform for research and clinical applications. *Medical Physics*, 39:6818, 2012.
- [162] S. Agostinelli et al. Geant4 - a simulation toolkit. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 506(3):250 – 303, 2003.
- [163] Hans-Peter Wieser, Eduardo Cisternas, Niklas Wahl, Silke Ulrich, Alexander Stadler, Henning Mescher, Lucas-Raphael Müller, Thomas Klinge, Hubert Gabrys, Lucas Burigo, Andrea Mairani, Swantje Ecker, Benjamin Ackermann, Malte Ellerbrock, Katia Parodi, Oliver Jäkel, and Mark Bangert. Development of the open-source dose calculation and optimization toolkit matrad. *Medical Physics*, 44(6):2556–2568, 2017.
- [164] T. Friedrich, U. Scholz, T. Elsaesser, M. Durante, and M. Scholz. Systematic analysis of RBE and related quantities using a database of cell survival experiments with ion beam irradiation. *Journal of Radiation Research*, 54(3):494–514, 2013.
- [165] Jacinta Yap, Javier Resta-López, Andrzej Kacperek, Roland Schnuerer, Simon Jolly, Stewart Boogert, and Carsten Welsch. Beam characterisation studies of the 62 MeV proton therapy beamline at the clatterbridge cancer centre. *Physica Medica*, 77:108 – 120, 2020.
- [166] Minna Wedenberg and Iuliana Toma-Dasu. Disregarding rbe variation in treatment plan comparison may lead to bias in favor of proton plans. *Medical Physics*, 41(9):091706, 2014.
- [167] B. Jones, S.J. McMahon, and K.M. Prise. The radiobiology of proton therapy: Challenges and opportunities around relative biological effectiveness. *Clinical Oncology*, 30(5):285 – 292, 2018.
- [168] B. Emami, J. Lyman, A. Brown, L. Cola, M. Goitein, J. E. Munzenrider, B. Shank, L. J. Solin, and M. Wesson. Tolerance of normal tissue to therapeutic irradiation. *International Journal of Radiation Oncology • Biology, • Physics*, 21(1):109–122, 1991.
- [169] Nishant Verma, Matthew C. Cowperthwaite, Mark G. Burnett, and Mia K. Markey. Differentiating tumor recurrence from treatment necrosis: a review of neuro-oncologic imaging strategies. *Neuro-oncology*, 15(5):515–534, 2013.

- [170] Qi Liu, Tracy SA Underwood, Jong Kung, Meng Wang, Hsiao-Ming Lu, Harald Paganetti, Kathryn D Held, Theodore S Hong, Jason A Efstathiou, and Henning Willers. Disruption of *slx4-mus81* function increases the relative biological effectiveness of proton radiation. *International Journal of Radiation Oncology • Biology, • Physics*, 95(1):78–85, 2016.
- [171] Steffen Nielsen, Niels Bassler, Leszek Grzanka, Jan Swakon, Pawel Olko, Christian Nicolaj Andreassen, Jens Overgaard, Jan Alsner, and Brita Singers Sørensen. Differential gene expression in primary fibroblasts induced by proton and cobalt-60 beam irradiation. *Acta Oncologica*, 56(11):1406–1412, 2017.
- [172] Harald Paganetti. Proton Relative Biological Effectiveness – Uncertainties and Opportunities. *International Journal of Particle Therapy*, 5(1):2–14, 2018.
- [173] Henning Willers, Antino Allen, David Grosshans, Stephen J. McMahon, Cläre von Neubeck, Claudia Wiese, and Bhadrasain Vikram. Toward a variable RBE for proton beam therapy. *Radiotherapy and Oncology*, 128(1):68 – 75, 2018.
- [174] Dudley T. Goodhead. Energy deposition stochastics and track structure: what about the target? *Radiation Protection Dosimetry*, 122(1-4):3–15, 2007.
- [175] Pankaj Chaudhary, Thomas I. Marshall, Francesca M. Perozziello, Lorenzo Manti, Frederick J. Currell, Fiona Hanton, Stephen J. McMahon, Joy N. Kavanagh, Giuseppe Antonio Pablo Cirrone, Francesco Romano, Kevin M. Prise, and Giuseppe Schettino. Relative biological effectiveness variation along monoenergetic and modulated bragg peaks of a 62-MeV therapeutic proton beam: A preclinical assessment. *International Journal of Radiation Oncology • Biology, • Physics*, 90(1):27 – 35, 2014.
- [176] Pankaj Chaudhary, Thomas I. Marshall, Frederick J. Currell, Andrzej Kacperek, Giuseppe Schettino, and Kevin M. Prise. Variations in the processing of DNA double-strand breaks along 60-MeV therapeutic proton beams. *International Journal of Radiation Oncology • Biology, • Physics*, 95(1):86–94, May 2016.
- [177] John J Cuaron, Chang Chang, Michael Lovelock, Daniel S Higginson, Dennis Mah, Oren Cahlon, and Simon Powell. Exponential increase in relative biological effectiveness along distal edge of a proton bragg peak as measured by deoxyribonucleic acid double-strand breaks. *International Journal of Radiation Oncology • Biology, • Physics*, 95(1):62–69, 2016.
- [178] Brita Singers Sørensen, Niels Bassler, Steffen Nielsen, Michael R Horsman, Leszek Grzanka, Harald Spejlborg, Jan Swakoń, Paweł Olko, and Jens Overgaard. Relative biological

- effectiveness (RBE) and distal edge effects of proton radiation on early damage in vivo. *Acta Oncologica*, 56(11):1387–1391, 2017.
- [179] Maria Saager, Peter Peschke, Stephan Brons, Jürgen Debus, and Christian P Karger. Determination of the proton RBE in the rat spinal cord: Is there an increase towards the end of the spread-out bragg peak? *Radiotherapy and Oncology*, 128(1):115–120, 2018.
- [180] Drosoula Giantsoudi, Clemens Grassberger, David Craft, Andrzej Niemierko, Alexei Trofimov, and Harald Paganetti. Linear energy transfer-guided optimization in intensity modulated proton therapy: feasibility study and clinical potential. *International Journal of Radiation Oncology • Biology, • Physics*, 87(1):216–222, Sep 2013. 23790771[pmid].
- [181] Jan Unkelbach, Pablo Botas, Drosoula Giantsoudi, Bram L. Gorissen, and Harald Paganetti. Reoptimization of intensity modulated proton therapy plans based on linear energy transfer. *International Journal of Radiation Oncology Biology Physics*, 96(5):1097–1106, 2016.
- [182] Stephen J McMahon, Harald Paganetti, and Kevin M Prise. Let-weighted doses effectively reduce biological variability in proton radiotherapy planning. *Physics in Medicine & Biology*, 63(22):225009, 2018.
- [183] K. Rothkamm, I. Krüger, L.H. Thompson, and M. Löbrich. Pathways of DNA double-strand break repair during the mammalian cell cycle. *Molecular and Cellular Biology*, 23(16):5706–5715, 2003.
- [184] Thomas Helleday, Justin Lo, Dik C. van Gent, and Bevin P. Engelward. DNA double-strand break repair: From mechanistic understanding to cancer treatment. *DNA Repair*, 6(7):923 – 935, 2007.
- [185] J J Wilkens and U Oelfke. A phenomenological model for the relative biological effectiveness in therapeutic proton beams. *Physics in Medicine and Biology*, 49(13):2811–2825, 2004.
- [186] Alejandro Carabe-Fernandez, Roger G. Dale, and Bleddyn Jones. The incorporation of the concept of minimum RBE (RBE_{\min}) into the linear-quadratic model and the potential for improved radiobiological analysis of high-LET treatments. *International Journal of Radiation Biology*, 83(1):27–39, 2007.
- [187] Alejandro Carabe, Maryam Moteabbed, Nicolas Depauw, Jan Schuemann, and Harald Paganetti. Range uncertainty in proton therapy due to variable biological effectiveness. *Physics in Medicine and Biology*, 57(5):1159–1172, 2012.

- [188] Y. Chen and S. Ahmad. Empirical model estimation of relative biological effectiveness for proton beam therapy. *Radiation Protection Dosimetry*, 149(2):116–123, 2011.
- [189] Minna Wedenberg, Bengt K. Lind, and Björn Hårdemark. A model for the relative biological effectiveness of protons: The tissue specific parameter α/β of photons is a predictor for the sensitivity to LET changes. *Acta Oncologica*, 52(3):580–588, 2013.
- [190] Aimee L McNamara, Jan Schuemann, and Harald Paganetti. A phenomenological relative biological effectiveness (RBE) model for proton therapy based on all published in vitro cell survival data. *Physics in Medicine and Biology*, 60(21):8399–8416, 2015.
- [191] A Mairani, I Dokic, G Magro, T Tessonnier, J Bauer, T T Böhlen, M Ciocca, A Ferrari, P R Sala, O Jäkel, J Debus, T Haberer, A Abdollahi, and K Parodi. A phenomenological relative biological effectiveness approach for proton therapy based on an improved description of the mixed radiation field. *Physics in Medicine and Biology*, 62(4):1378–1395, 2017.
- [192] Roland B Hawkins. A statistical theory of cell killing by radiation of varying linear energy transfer. *Radiation research*, 140(3):366–374, 1994.
- [193] RB Hawkins. A microdosimetric-kinetic model of cell death from exposure to ionizing radiation of any let, with experimental and clinical applications. *International Journal of Radiation Biology*, 69(6):739–755, 1996.
- [194] Roland B. Hawkins. A microdosimetric-kinetic theory of the dependence of the RBE for cell death on LET. *Medical Physics*, 25(7 I):1157–1170, 1998.
- [195] Roland B Hawkins. Mammalian cell killing by ultrasoft x rays and high-energy radiation: an extension of the mk model. *Radiation research*, 166(2):431–442, 2006.
- [196] Yuki Kase, Tatsuaki Kanai, Naruhiro Matsufuji, Yoshiya Furusawa, Thilo Elsässer, and Michael Scholz. Biophysical calculation of cell survival probabilities using amorphous track structure models for heavy-ion irradiation. *Physics in Medicine & Biology*, 53(1):37, 2007.
- [197] T Sato, R Watanabe, Y Kase, C Tsuruoka, M Suzuki, Y Furusawa, and K Niita. Analysis of cell-survival fractions for heavy-ion irradiations based on microdosimetric kinetic model implemented in the particle and heavy ion transport code system. *Radiation protection dosimetry*, 143(2-4):491–496, 2011.
- [198] M. Scholz and G. Kraft. Track structure and the calculation of biological effects of heavy charged particles. *Advances in Space Research*, 18(1-2):5–14, 1996.

- [199] Thilo Elsässer, Wilma K Weyrather, Thomas Friedrich, Marco Durante, Gheorghe Iancu, Michael Krämer, Gabriele Kragl, Stephan Brons, Marcus Winter, Klaus-Josef Weber, et al. Quantification of the relative biological effectiveness for ion beam radiotherapy: direct experimental comparison of proton and carbon ion beams and a novel approach for treatment planning. *International Journal of Radiation Oncology • Biology, • Physics*, 78(4):1177–1183, 2010.
- [200] Thomas Friedrich, Uwe Scholz, Thilo Elsässer, Marco Durante, and Michael Scholz. Calculation of the biological effects of ion beams based on the microscopic spatial damage distribution pattern. *International Journal of Radiation Biology*, 88(1-2):103–107, 2012.
- [201] David J. Carlson, Robert D. Stewart, Vladimir A. Semenenko, and George A. Sandison. Combined Use of Monte Carlo DNA Damage Simulations and Deterministic Repair Models to Examine Putative Mechanisms of Cell Killing. *Radiation Research*, 169(4):447 – 459, 2008.
- [202] Malte C. Frese, Victor K. Yu, Robert D. Stewart, and David J. Carlson. A mechanism-based approach to predict the relative biological effectiveness of protons and carbon ions in radiation therapy. *International Journal of Radiation Oncology • Biology, • Physics*, 83(1):442–450, 2012.
- [203] Florian Kamp, Gonzalo Cabal, Andrea Mairani, Katia Parodi, Jan J. Wilkens, and David J. Carlson. Fast biological modeling for voxel-based heavy ion treatment planning using the mechanistic repair-misrepair-fixation model and nuclear fragment spectra. *International Journal of Radiation Oncology • Biology, • Physics*, 93(3):557–568, 2015.
- [204] Eivind Rørvik, Lars Fredrik Fjæra, Tordis J Dahle, Jon Espen Dale, Grete May Engeseth, Camilla H Stokkevåg, Sara Thörnqvist, and Kristian S Ytre-Hauge. Exploration and application of phenomenological RBE models for proton therapy. *Physics in Medicine & Biology*, 63(18):185013, 2018.
- [205] Robert D. Stewart, David J. Carlson, Michael P. Butkus, Roland Hawkins, Thomas Friedrich, and Michael Scholz. A comparison of mechanism-inspired models for particle relative biological effectiveness (RBE). *Medical Physics*, 45(11):e925–e952, 2018.
- [206] Roger G Dale and Bleddyn Jones. The assessment of RBE effects using the concept of biologically effective dose. *International Journal of Radiation Oncology • Biology, • Physics*, 43(3):639 – 645, 1999.
- [207] Jan J Wilkens and Uwe Oelfke. Analytical linear energy transfer calculations for proton therapy. *Medical physics*, 30(5):806–815, 2003.

- [208] Jan J. Wilkens and Uwe Oelfke. Optimization of radiobiological effects in intensity modulated proton therapy. *Medical Physics*, 32(2):455–465, 2005.
- [209] G. Schettino, M. Folkard, K. M. Prise, B. Vojnovic, A. G. Bowey, and B. D. Michael. Low-dose hypersensitivity in Chinese hamster V79 cells targeted with counted protons using a charged-particle microbeam. *Radiation Research*, 156(5):526 – 534, 2001.
- [210] G. Coutrakon, J. Cortese, A. Ghebremedhin, J. Hubbard, J. Johanning, P. Koss, G. Maudsley, C. R. Slater, C. Zuccarelli, and J. Robertson. Microdosimetry spectra of the loma linda proton beam and relative biological effectiveness comparisons. *Medical Physics*, 24(9):1499–1506, 1997.
- [211] Clemens Grassberger, Alexei Trofimov, Anthony Lomax, and Harald Paganetti. Variations in linear energy transfer within clinical proton therapy fields and the potential for biological treatment planning. *International Journal of Radiation Oncology • Biology, • Physics*, 80(5):1559–1566, 2011.
- [212] Bleddyn Jones. Towards achieving the full clinical potential of proton therapy by inclusion of LET and RBE models. *Cancers*, 7(1):460–480, 2015.
- [213] A. Mairani, I. Dokic, G. Magro, T. Tessonnier, F. Kamp, D. J. Carlson, M. Ciocca, F. Cerutti, P. R. Sala, A. Ferrari, T. T. Böhlen, O. Jäkel, K. Parodi, J. Debus, A. Abdollahi, and T. Haberer. Biologically optimized helium ion plans: Calculation approach and its in vitro validation. *Physics in Medicine and Biology*, 61(11), 2016.
- [214] A Mairani, G Magro, I Dokic, SM Valle, T Tessonnier, R Galm, M Ciocca, K Parodi, A Ferrari, O Jäkel, et al. Data-driven RBE parameterization for helium ion beams. *Physics in Medicine & Biology*, 61(2):888, 2016.
- [215] Albrecht M Kellerer and Harald H Rossi. The theory of dual radiation action. *Current Topics in Radiation Research Quarterly*, 85–158, 1974.
- [216] Albrecht M. Kellerer and Harald H. Rossi. A generalized formulation of dual radiation action. *Radiation Research*, 178(2), 2012.
- [217] M Scholz, AM Kellerer, W Kraft-Weyrather, and G Kraft. Computation of cell survival in heavy ion beams for therapy. *Radiation and Environmental Biophysics*, 36(1):59–66, 1997.
- [218] Christian P Karger and Peter Peschke. RBE and related modeling in carbon-ion therapy. *Physics in Medicine & Biology*, 63(1):01TR02, 2017.

- [219] Michael Krämer and Michael Scholz. Rapid calculation of biological effects in ion radiotherapy. *Physics in Medicine & Biology*, 51(8):1959, 2006.
- [220] Caterina Monini, Gersende Alphonse, Claire Rodriguez-Lafrasse, Étienne Testa, and Michaël Beuve. Comparison of biophysical models with experimental data for three cell lines in response to irradiation with monoenergetic ions. *Physics and Imaging in Radiation Oncology*, 12:17–21, 2019.
- [221] HH Rossi. Microdosimetry. *Biophysical aspects of radiation quality*, 81–94, 1966.
- [222] ICRU Microdosimetry. Report no. 36. *International Commission on Radiation Units and Measurements*, Bethesda, Maryland, 20814, 1983.
- [223] Albrecht M Kellerer and D Chmelevsky. Criteria for the applicability of LET. *Radiation Research*, 63(2):226–234, 1975.
- [224] Robert D Stewart, Victor K Yu, Alexandros G Georgakilas, Constantinos Koumenis, Joo Han Park, and David J Carlson. Effects of radiation quality and oxygen on clustered DNA lesions and cell death. *Radiation Research*, 176(5):587–602, 2011.
- [225] Robert D Stewart, Seth W Streitmatter, David C Argento, Charles Kirkby, John T Goorley, Greg Moffitt, Tatjana Jevremovic, and George A Sandison. Rapid MCNP simulation of DNA double strand break (DSB) relative biological effectiveness (RBE) for photons, neutrons, and light ions. *Physics in Medicine & Biology*, 60(21):8249, 2015.
- [226] Robert D Stewart. Induction of DNA damage by light ions relative to ^{60}Co γ -rays. *International journal of particle therapy*, 5(1):25–39, 2018.
- [227] Frank Van Den Heuvel. A closed parameterization of DNA-damage by charged particles, as a function of energy - A geometrical approach. *PLoS ONE*, 9(10):1–14, 2014.
- [228] V. A. Semenenko and R. D. Stewart. A fast Monte Carlo algorithm to simulate the spectrum of DNA damages formed by ionizing radiation. *Radiation Research*, 161(4):451–457, 2004.
- [229] Y Hsiao and R D Stewart. Monte Carlo simulation of DNA damage induction by X-rays and selected radioisotopes. *Phys Med Biol*, 53(1):233, 2008.
- [230] F Van den Heuvel, Jean-Pierre Locquet, and S Nuyts. Beam energy considerations for gold nano-particle enhanced radiation treatment. *Physics in Medicine and Biology*, 55(16):4509, 2010.

- [231] Michael A Huels, Badia Boudaïffa, Pierre Cloutier, Darel Hunting, and Leon Sanche. Single, double, and multiple double strand breaks induced in DNA by 3– 100 eV electrons. *Journal of the American Chemical Society*, 125(15):4467–4477, 2003.
- [232] YW Huang, CY Pan, YY Hsiao, TC Chao, CC Lee, and CJ Tung. Monte Carlo simulations of the relative biological effectiveness for DNA double strand breaks from 300 MeV u^{-1} carbon-ion beams. *Physics in Medicine & Biology*, 60(15):5995, 2015.
- [233] Sylvain Meylan, Sébastien Incerti, Mathieu Karamitros, Nicolas Tang, Marta Bueno, Isabelle Clairand, and Carmen Villagrasa. Simulation of early DNA damage after the irradiation of a fibroblast cell nucleus using geant4-DNA. *Scientific reports*, 7(1):1–15, 2017.
- [234] Konstantinos P. Chatzipapas, Panagiotis Papadimitroulas, Mohammad Obeidat, Kristen A. McConnell, Neil Kirby, George Loudos, Niko Papanikolaou, and George C. Kagadis. Quantification of DNA double-strand breaks using Geant4-DNA. *Medical Physics*, 46(1):405–413, 2019.
- [235] Zhenyu Tan, Yueyuan Xia, Mingwen Zhao, and Xiangdong Liu. Proton inelastic mean free path in a group of bioorganic compounds and water in 0.05–10MeV range – including higher-order corrections. *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 268(14):2337 – 2342, 2010.
- [236] C. J. Powell and A. Jablonski. NIST standard reference database 71. Available at: <http://dx.doi.org/10.18434/T48C78>, 1999. Accessed: 03-09-2020.
- [237] C. J. Powell and A. Jablonski. Evaluation of calculated and measured electron inelastic mean free paths near solid surfaces. *Journal of Physical and Chemical Reference Data*, 28(1):19–62, 1999.
- [238] Walter H. Barkas, John N. Dyer, and Harry H. Heckman. Resolution of the Σ^{-} -mass anomaly. *Phys. Rev. Lett.*, 11:26–28, Jul 1963.
- [239] Beata Ziaja, Richard A. London, and Janos Hajdu. Ionization by impact electrons in solids: Electron mean free path fitted over a wide energy range. *Journal of Applied Physics*, 99(3):033514, 2006.
- [240] G.D. Birnie, D. Rickwood and A. Hell. Buoyant densities and hydration of nucleic acids, proteins and nucleoprotein complexes in metrizamide. *Biochimica et Biophysica Acta - Nucleic Acids and Protein Synthesis*, 331(2):283-294, 1973.

- [241] C. Champion, M.A. Quinto, J.M. Monti, M.E. Galassi, P.F. Weck, O.A. Fojón, J. Hanssen and R.D. Rivarola. Water versus DNA: new insights into proton track-structure modelling in radiobiology and radiotherapy. *Physics in Medicine and Biology*, 60:7805–7828, 2015.
- [242] M.E. Alcocer-Ávila, M.A. Quinto, J.M. Monti, R.D. Rivarola and C. Champion. Proton transport modeling in a realistic biological environment by using TILDA-V. *Scientific Reports*, 9:14030, 2019.
- [243] Tan Z., Xia Y., Zhao M. and Liu X. Proton stopping power in a group of bioorganic compounds over the energy range of 0.05–10 MeV. *Nucl. Instrum. Methods Phys. Res.*, 248:1, 2006.
- [244] Ziegler J. F., Ziegler M. D., and Biersack J. P. SRIM: The Stopping and Range of Ions in Matter. *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 268(11-12):1818–1823, 2010.
- [245] S. Incerti, G. Baldacchino, M. Bernal, R. Capra, C. Champion, Z. Francis, P. Guèye, A. Mantero, B. Mascialino, P. Moretto, P. Nieminen, C. Villagrasa, and C. Zacharatou. The Geant4-DNA project. *International Journal of Modeling, Simulation, and Scientific Computing*, 01(02):157–178, 2010.
- [246] Francis, Z., Incerti, S., Karamitros, M., Tran, H. N., and Villagrasa, C. Stopping power and ranges of electrons, protons and alpha particles in liquid water using the Geant4-DNA package. *Nucl. Instrum. Methods Phys. Res.*, 269:2307, 2011.
- [247] Bernal, M.A., Bordage, M.C., Brown, J.M.C., Davidková, M., Delage, E., El Bitar, Z., Enger, S.A., Francis, Z., Guatelli, S., Ivanchenko, V.N. and Karamitros, M. Track structure modeling in liquid water: A review of the Geant4-DNA very low energy extension of the Geant4 Monte Carlo simulation toolkit. *Phys. Med.*, 31:861, 2015.
- [248] Mustapha Krim, Jamal Inchaouh, Noura Harakat, Abdenbi Khouaja, Meriem Fiak, Mohammed Reda Mesradi, Ahmed Fathi, Omar Jdair, and Redouane El baydaoui. Cross sections, stopping power and bragg peak range calculation of proton collisions with the DNA base adenine. *Japanese Journal of Applied Physics*, 58(9):096001, 2019.
- [249] E.P. Rogakou, D.R. Pilch, A.H. Orr, V.S. Ivanova, and W.M. Bonner. DNA double-stranded breaks induce histone H2AX phosphorylation on serine 139. *Journal of Biological Chemistry*, 273(10):5858–5868, 1998. cited By 3582.

- [250] D. D. Ager, W. C. Dewey, K. Gardiner, W. Harvey, R. T. Johnson, and C. A. Waldren. Measurement of radiation-induced DNA double-strand breaks by pulsed-field gel electrophoresis. *Radiation Research*, 122(2):181–187, 1990.
- [251] O. Ostling and K.J. Johanson. Microelectrophoretic study of radiation-induced DNA damages in individual mammalian cells. *Biochemical and Biophysical Research Communications*, 123(1):291 – 298, 1984.
- [252] Olga A. Sedelnikova, Emmy P. Rogakou, Igor G. Panyutin, and William M. Bonner. Quantitative detection of 125IdU-Induced DNA double-strand breaks with γ -H2AX antibody. *Radiation Research*, 158(4):486–492, 2002.
- [253] Emmy P. Rogakou, Chye Boon, Christophe Redon, and William M. Bonner. Megabase chromatin domains involved in DNA double-strand breaks in vivo. *Journal of Cell Biology*, 146(5):905–916, 1999.
- [254] Mischa G. Vrouwe, Alex Pines, Rene M. Overmeer, Katsuhiko Hanada, and Leon H. F. Mullenders. UV-induced photolesions elicit ATR-kinase-dependent signaling in non-cycling cells through nucleotide excision repair-dependent and -independent pathways. *Journal of Cell Science*, 124(3):435–446, 2011.
- [255] Andrea Beucher, Julie Birraux, Leopoldine Tchouandong, Olivia Barton, Atsushi Shibata, Sandro Conrad, Aaron A Goodarzi, Andrea Krempler, Penny A Jeggo, and Markus Löbrich. Atm and artemis promote homologous recombination of radiation-induced DNA double-strand breaks in G2. *The EMBO Journal*, 28(21):3413–3427, 2009.
- [256] Simon Bekker-Jensen, Claudia Lukas, Risa Kitagawa, Fredrik Melander, Michael B. Kastan, Jiri Bartek, and Jiri Lukas. Spatial organization of the mammalian genome surveillance machinery in response to DNA strand breaks. *Journal of Cell Biology*, 173(2):195–206, 2006.
- [257] Doug W. Chan, Benjamin Ping-Chi Chen, Sheela Prithivirajasingh, Akihiro Kurimasa, Michael D. Story, Jun Qin, and David J. Chen. Autophosphorylation of the DNA-dependent protein kinase catalytic subunit is required for rejoining of DNA double-strand breaks. *Genes Development*, 16(18):2333–2338, 2002.
- [258] Ifigeneia V. Mavragani, Zacharenia Nikitaki, Spyridon A. Kalospyros, and Alexandros G. Georgakilas. Ionizing radiation and complex DNA damage: From prediction to detection challenges and biological significance. *Cancers*, 11(11):1789, 2019.

- [259] J. Drouet, C. Delteil, J. Lefrançois, P. Concannon, B. Salles, and P. Calsou. DNA-dependent protein kinase and XRCC4-DNA ligase IV mobilization in the cell in response to DNA double strand breaks. *Journal of Biological Chemistry*, 280(8):7060–7069, 2005.
- [260] Aaron A. Goodarzi and Penny A. Jeggo. Irradiation induced foci (IRIF) as a biomarker for radiosensitivity. *Mutation Research/Fundamental and Molecular Mechanisms of Mutagenesis*, 736(1):39 – 47, 2012.
- [261] Sebastien de Feraudy, Ingrid Revet, Vladimir Bezrookove, Luzviminda Feeney, and James E. Cleaver. A minority of foci or pan-nuclear apoptotic staining of γ H2AX in the S phase after UV damage contain DNA double-strand breaks. *Proceedings of the National Academy of Sciences*, 107(15):6870–6875, 2010.
- [262] Kai J. Neelsen, Isabella M.Y. Zanini, Raquel Herrador, and Massimo Lopes. Oncogenes induce genotoxic stress by mitotic processing of unusual replication intermediates. *Journal of Cell Biology*, 200(6):699–708, 2013.
- [263] Kai Rothkamm, Stephen Barnard, Jayne Moquet, Michele Ellender, Zohaib Rana, and Susanne Burdak-Rothkamm. DNA damage foci: Meaning and significance. *Environmental and Molecular Mutagenesis*, 56(6):491–504, 2015.
- [264] Nico Ruprecht, Martin N. Hungerbühler, Ingrid B. Böhm, and Johannes T. Heverhagen. Improved identification of DNA double strand breaks: γ -h2ax-epitope visualization by confocal microscopy and 3d reconstructed images. *Radiation and Environmental Biophysics*, 58(2):295–302, May 2019.
- [265] P.V. Hariharan and Franklin Hutchinson. Neutral sucrose gradient sedimentation of very large DNA from bacillus subtilis: Ii. double-strand breaks formed by gamma ray irradiation of the cells. *Journal of Molecular Biology*, 75(3):479 – 494, 1973.
- [266] Matthews O. Bradley and Kurt W. Kohn. X-ray induced DNA double strand break production and repair in mammalian cells as measured by neutral filter elution. *Nucleic Acids Research*, 7(3):793–804, 10 1979.
- [267] Emma L. Leatherbarrow, Jane V. Harper, Francis A. Cucinotta, and Peter O’Neill. Induction and quantification of γ -H2AX foci following low and high let-irradiation. *International Journal of Radiation Biology*, 82(2):111–118, 2006.
- [268] Julien Vignard, Gladys Mirey, and Bernard Salles. Ionizing-radiation induced DNA double-strand breaks: A direct and indirect lighting up. *Radiotherapy and Oncology*, 108(3):362 – 369, 2013.

- [269] George D. D. Jones, Tracy V. Boswell, and John F. Ward. Effects of postirradiation temperature on the yields of radiation-induced single- and double-strand breakage in sv40 DNA. *Radiation Research*, 138(2):291–296, 1994.
- [270] E. S. Henle, R. Roots, W. R. Holley, and A. Chatterjee. DNA strand breakage is correlated with unaltered base release after gamma irradiation. *Radiation Research*, 143(2):144–150, 1995.
- [271] B. Rydberg. Radiation-induced heat-labile sites that convert into DNA double-strand breaks. *Radiation Research*, 153(6):805–812, 2000. cited By 87.
- [272] Maryann B. Flick, Raymond L. Warters, Linda S. Yasui, and Robert E. Krisch. Measurement of radiation-induced DNA damage using gel electrophoresis or neutral filter elution shows an increased frequency of DNA strand breaks after exposure to ph 9.6. *Radiation Research*, 119(3):452–465, 1989.
- [273] D. Blöcher and M. Kunhi. DNA double-strand break analysis by chef electrophoresis. *International Journal of Radiation Biology*, 58(1):23–34, 1990.
- [274] G.E. Iliakis, L. Metzger, N. Denko, and T.D. Stamato. Detection of DNA double-strand breaks in synchronous cultures of cho cells by means of asymmetric field inversion gel electrophoresis. *International Journal of Radiation Biology*, 59(2):321–341, 1991.
- [275] Katsuhiko Hanada, Magda Budzowska, Mauro Modesti, Alex Maas, Claire Wyman, Jeroen Essers, and Roland Kanaar. The structure-specific endonuclease mus81–eme1 promotes conversion of interstrand DNA crosslinks into double-strand breaks. *The EMBO Journal*, 25(20):4921–4932, 2006.
- [276] Rachael Elizabeth Hawtin, David Elliot Stockett, Oi Kwan Wong, Thomas Helleday, and Judith Ann Fox. Homologous recombination repair is essential for repair of vosaroxin-induced DNA double-strand breaks. *Oncotarget*, 1(7):606–619, 2010.
- [277] Katsuhiko Hanada, Tomohisa Uchida, Yoshiyuki Tsukamoto, Masahide Watada, Nahomi Yamaguchi, Kaoru Yamamoto, Seiji Shiota, Masatsugu Moriyama, David Y. Graham, and Yoshio Yamaoka. Helicobacter pylori infection introduces DNA double-strand breaks in host cells. *Infection and Immunity*, 82(10):4182–4189, 2014.
- [278] Mohamed E. Ashour, Reham Atteya, and Sherif F. El-Khamisy. Topoisomerase-mediated chromosomal break repair: an emerging player in many games. *Nature Reviews Cancer*, 15(3):137–151, 2015.

- [279] Yuri Kawashima, Nahomi Yamaguchi, Rie Teshima, Hisashi Narahara, Yoshio Yamaoka, Hirofumi Anai, Yoshihiro Nishida, and Katsuhiro Hanada. Detection of DNA double-strand breaks by pulsed-field gel electrophoresis. *Genes to Cells*, 22(1):84–93, 2017.
- [280] Ümit Atila, Yusuf Yargı Baydilli, Eftal Sehirli, and Muhammed Kamil Turan. Classification of DNA damages on segmented comet assay images using convolutional neural network. *Computer Methods and Programs in Biomedicine*, 186:105192, 2020.
- [281] Lykke Forchhammer, Clara Ersson, Steffen Loft, Lennart Möller, Roger W. L. Godschalk, et al. Inter-laboratory variation in DNA damage using a standard comet assay protocol. *Mutagenesis*, 27(6):665–672, 07 2012.
- [282] John R. Masters. HeLa cells 50 years on: the good, the bad and the ugly. *Nature Reviews Cancer*, 2(4):315–319, 2002.
- [283] Nicolaas A. P. Franken, Hans M. Rodermond, Jan Stap, Jaap Haveman, and Chris van Bree. Clonogenic assay of cells in vitro. *Nature Protocols*, 1(5):2315–2319, 2006.
- [284] Haloom Rafehi, Christian Orłowski, George T. Georgiadis, Katherine Ververis, Assam El-Osta, and Tom C. Karagiannis. Clonogenic assay: adherent cells. *Journal of Visualized Experiments: JoVE*, (49):2573, 2011.
- [285] Árpád Szöör, László Ujlaky-Nagy, Gábor Tóth, János Szöllösi, and György Vereb. Cell confluence induces switching from proliferation to migratory signaling by site-selective phosphorylation of pdgf receptors on lipid raft platforms. *Cellular Signalling*, 28(2):81 – 93, 2016.
- [286] Matthews O. Bradley, Bijoy Bhuyan, Mary C. Francis, Robert Langenbach, Andrew Peterson, and Eliezer Huberman. Mutagenesis by chemical agents in v79 chinese hamster cells: A review and analysis of the literature: A report of the gene-tox program. *Mutation Research/Reviews in Genetic Toxicology*, 87(2):81 – 142, 1981.
- [287] Wenren Chaung, Li-Jun Mi, and Robert J. Boorstein. The p53 Status of Chinese hamster V79 cells frequently used for studies on DNA damage and DNA repair. *Nucleic Acids Research*, 25(5):992–994, 03 1997.
- [288] Colleen A. Brady and Laura D. Attardi. p53 at a glance. *Journal of Cell Science*, 123(15):2527–2532, 2010.
- [289] Elisabeth Mara, Monika Clausen, Suphalak Khachonkham, Simon Deycmar, Clara Pessy, Wolfgang Dörr, Peter Kuess, Dietmar Georg, and Sylvia Gruber. Investigating the impact of

- alpha/beta and LET_d on relative biological effectiveness in scanned proton beams: An in vitro study based on human cell lines. *Medical Physics*, 47(8):3691–3702, 2020.
- [290] Apostolos Menegakis, Ala Yaromina, Wolfgang Eicheler, Annegret Dörfler, Bettina Beuthien-Baumann, Howard D. Thames, Michael Baumann, and Mechthild Krause. Prediction of clonogenic cell survival curves based on the number of residual DNA double strand breaks measured by γ H2AX staining. *International Journal of Radiation Biology*, 85(11):1032–1041, 2009.
- [291] Armin Lühr, Cläre von Neubeck, Mechthild Krause, and Esther G. C. Troost. Relative biological effectiveness in proton beam therapy - current knowledge and future challenges. *Clinical and Translational Radiation Oncology*, 9:35–41, 2018.
- [292] Erik Traneus and Jakob Ödén. Introducing proton track-end objectives in intensity modulated proton therapy optimization to reduce linear energy transfer and relative biological effectiveness in critical structures. *International Journal of Radiation Oncology • Biology, • Physics*, 103(3):747–757, 2019.
- [293] Linda Hong, Michael Goitein, Marta Bucciolini, Robert Comiskey, Bernard Gottschalk, Skip Rosenthal, Chris Serago, and Marcia Urie. A pencil beam algorithm for proton dose calculations. *Physics in Medicine & Biology*, 41(8):1305, 1996.
- [294] B Gottschalk. Deterministic proton dose calculation from first principles. *Physics in Medicine & Biology*, 63(13):135016, jul 2018.
- [295] L. H. Gray, A. D. Conger, M. Ebert, S. Hornsey, and O. C. A. Scott. The concentration of oxygen dissolved in tissues at the time of irradiation as a factor in radiotherapy. *The British Journal of Radiology*, 26(312):638–648, 1953.
- [296] D. Ewing and E. L. Powers. Oxygen-dependent sensitization of irradiated cells. Technical report, United States, 1979. CONF-790229-4.
- [297] C. Clifton Ling, Howard B. Michaels, Leo E. Gerweck, Edward R. Epp, and Eleanor C. Peterson. Oxygen sensitization of mammalian cells under different irradiation conditions. *Radiation Research*, 86(2):325–340, 1981.
- [298] T. T. Kwok and R. M. Sutherland. The radiation response of cells recovering after chronic hypoxia. *Radiation Research*, 119(2):261–270, 1989.
- [299] Paul Okunieff, Michael Hoeckel, Eamonn P. Dunphy, Karlheinz Schlenger, Claudia Knoop, and Peter Vaupel. Oxygen tension distributions are sufficient to explain the local response of

- human breast tumors treated with radiation alone. *International Journal of Radiation Oncology • Biology, • Physics*, 26(4):631–636, 1993.
- [300] Peter Alexander and Arthur Charlesby. Physico-chemical methods of protection against ionizing radiations. In *Radiobiology Symposium*, volume 49. Butterworth London, 1954.
- [301] David Robert Grimes and Mike Partridge. A mechanistic investigation of the oxygen fixation hypothesis and oxygen enhancement ratio. *Biomedical physics & engineering express*, 1(4):045209–045209, 2015.
- [302] Hans Liew, Carmen Klein, Frank T. Zenke, Amir Abdollahi, Jürgen Debus, Ivana Dokic, and Andrea Mairani. Modeling the effect of hypoxia and DNA repair inhibition on cell survival after photon irradiation. *International journal of molecular sciences*, 20(23):6054, 2019.
- [303] William R. Wilson and Michael P. Hay. Targeting hypoxia in cancer therapy. *Nature Reviews Cancer*, 11(6):393–410, 2011.
- [304] David J Carlson, Robert D Stewart, and Vladimir A Semenenko. Effects of oxygen on intrinsic radiation sensitivity: A test of the relationship between aerobic and hypoxic linear-quadratic (lq) model parameters a. *Medical physics*, 33(9):3105–3115, 2006.
- [305] F. Van Den Heuvel, D. De Ruyscher, W. Van Elmpt, and S. Nuyts. Using a closed analytical expression to determine biological effects depending on radiation spectrum and oxygen level. *International Journal of Radiation Oncology • Biology, • Physics*, 87(2):S97, Oct 2013.
- [306] Gordon R. Kepner. Saturation behavior: a general relationship described by a simple second-order differential equation. *Theoretical Biology and Medical Modelling*, 7(1):11, 2010.
- [307] Otto Warburg. On the origin of cancer cells. *Science*, 123(3191):309–314, 1956.
- [308] Joseph G. Rajendran, David A. Mankoff, Finbarr O’Sullivan, Lanell M. Peterson, David L. Schwartz, Ernest U. Conrad, Alexander M. Spence, Mark Muzi, D. Greg Farwell, and Kenneth A. Krohn. Hypoxia and glucose metabolism in malignant tumors. *Clinical Cancer Research*, 10(7):2245–2252, 2004.
- [309] C.Clifton Ling, John Humm, Steven Larson, Howard Amols, Zvi Fuks, Steven Leibel, and Jason A. Koutcher. Towards multidimensional radiotherapy (md-crt): biological imaging and biological conformality. *International Journal of Radiation Oncology • Biology, • Physics*, 47(3):551–560, 2000.
- [310] Søren M. Bentzen. Theragnostic imaging for radiation oncology: dose-painting by numbers. *The Lancet Oncology*, 6(2):112–117, 2005.

- [311] Indira Madani, Antony J. Lomax, Francesca Albertini, Petra Trnková, and Damien C. Weber. Dose-painting intensity-modulated proton therapy for intermediate- and high-risk meningioma. *Radiation Oncology*, 10:72–72, 2015.
- [312] Katrien Konings, Charlot Vandevoorde, Bjorn Baselet, Sarah Baatout, and Marjan Moreels. Combination therapy with charged particles and molecular targeting: A promising avenue to overcome radioresistance. *Frontiers in Oncology*, 10:128–128, 2020.
- [313] E Scifoni, W Tinganelli, W K Weyrather, M Durante, A Maier, and M Krämer. Including oxygen enhancement ratio in ion beam treatment planning: model implementation and experimental verification. *Physics in Medicine and Biology*, 58(11):3871–3895, 2013.
- [314] Walter Tinganelli, Marco Durante, Ryoichi Hirayama, Michael Krämer, Andreas Maier, Wilma Kraft-Weyrather, Yoshiya Furusawa, Thomas Friedrich, and Emanuele Scifoni. Kill-painting of hypoxic tumours in charged particle therapy. *Scientific Reports*, 5(1):17016, 2015.
- [315] O Sokol, M Krämer, S Hild, M Durante, and E Scifoni. Kill painting of hypoxic tumors with multiple ion beams. *Physics in Medicine & Biology*, 64(4):045008, 2019.
- [316] C. D. Town. Effect of high dose rates on survival of mammalian cells. *Nature*, 215(5103):847–848, 1967.
- [317] Vincent Favaudon, Laura Caplier, Virginie Monceau, Frédéric Pouzoulet, Mano Sayarath, Charles Fouillade, Marie-France Poupon, Isabel Brito, Philippe Hupé, Jean Bourhis, et al. Ultrahigh dose-rate flash irradiation increases the differential response between normal and tumor tissue in mice. *Science Translational Medicine*, 6(245):245ra93–245ra93, 2014.
- [318] Pierre Montay-Gruel, Kristoffer Petersson, Maud Jaccard, Gaël Boivin, Jean-François Germond, Benoit Petit, Raphaël Doenlen, Vincent Favaudon, François Bochud, Claude Bailat, et al. Irradiation in a flash: Unique sparing of memory in mice after whole brain irradiation with dose rates above 100 gy/s. *Radiotherapy and Oncology*, 124(3):365–369, 2017.
- [319] Kristoffer Petersson, Gabriel Adrian, Karl Butterworth, and Stephen J. McMahon. A quantitative analysis of the role of oxygen tension in flash radiation therapy. *International Journal of Radiation Oncology • Biology, • Physics*, 107(3):539 – 547, 2020.
- [320] Wenjing Wang, Chunyan Li, Rui Qiu, Yizheng Chen, Zhen Wu, Hui Zhang, and Junli Li. Modelling of cellular survival following radiation-induced DNA double-strand breaks. *Scientific Reports*, 8(1):16202, 2018.

- [321] Larry H Toburen. Ionization and charge-transfer: basic data for track structure calculations. *Radiation and Environmental Biophysics*, 37(4):221–233, 1998.
- [322] C. Zacharatou Jarlskog and H. Paganetti. Physics settings for using the Geant4 toolkit in proton therapy. *IEEE Transactions on Nuclear Science*, 55(3):1018–1025, 2008.
- [323] David Jette and Weimin Chen. Creating a spread-out Bragg peak in proton beams. *Physics in Medicine and Biology*, 56(11):N131–N138, 2011.
- [324] B. Gottschalk, A.M. Koehler, R.J. Schneider, J.M. Sisterson, and M.S. Wagner. Multiple coulomb scattering of 160 MeV protons. *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 74(4):467 – 490, 1993.
- [325] Andreas Wächter and Lorenz T. Biegler. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical Programming*, 106(1):25–57, 2006.
- [326] Lisa Polster, Jan Schuemann, Ilaria Rinaldi, Lucas Burigo, Aimee L McNamara, Robert D Stewart, Andrea Attili, David J Carlson, Tatsuhiko Sato, José Ramos Méndez, Bruce Faddegon, Joseph Perl, and Harald Paganetti. Extension of TOPAS for the simulation of proton radiation effects considering molecular and cellular endpoints. *Physics in Medicine and Biology*, 60(13):5053–5070, 2015.
- [327] H. I. Amols, B. Lagueux, and D. Cagna. Radiobiological effectiveness (RBE) of megavoltage x-ray and electron beams in radiotherapy. *Radiation Research*, 105(1):58–67, 1986.
- [328] Stephen J. McMahon, Jan Schuemann, Harald Paganetti, and Kevin M. Prise. Mechanistic modelling of DNA repair and cellular survival following radiation-induced DNA damage. *Scientific Reports*, 6(1):33290, 2016.
- [329] NT Henthorn, JW Warmenhoven, M Sotiropoulos, RI Mackay, NF Kirkby, KJ Kirkby, and MJ Merchant. In silico non-homologous end joining following ion induced DNA double strand breaks predicts that repair fidelity depends on break density. *Scientific reports*, 8(1):2654, 2018.
- [330] Reza Taleei. Modelling DSB repair kinetics for DNA damage induced by proton and carbon ions. *Radiation Protection Dosimetry*, 183(1-2):75–78, 01 2019.
- [331] John W. Warmenhoven, Nicholas T. Henthorn, Samuel P. Ingram, Amy L. Chadwick, Marios Sotiropoulos, Nickolay Korabel, Sergei Fedotov, Ranald I. Mackay, Karen J. Kirkby, and Michael J. Merchant. Insights into the non-homologous end joining pathway and double

- strand break end mobility provided by mechanistic in silico modelling. *DNA Repair*, 85:102743, 2020.
- [332] Fazeleh S. Mohseni-Salehi, Fatemeh Zare-Mirakabad, Mehdi Sadeghi, and Soudeh Ghafouri-Fard. A stochastic model of DNA double-strand breaks repair throughout the cell cycle. *Bulletin of Mathematical Biology*, 82(1):11, 2020.
- [333] S Chauvie, Z Francis, Susanna Guatelli, S Incerti, B Mascialino, Gerard Montarou, Ph Moretto, P Nieminen, and Maria G Pia. Monte carlo simulation of interactions of radiation with biological systems at the cellular and DNA levels: the Geant4-DNA project. 2006.
- [334] S Chauvie, Zx Francis, S Guatelli, S Incerti, B Mascialino, Gerard Montarou, P Moretto, P Nieminen, and MG Pia. Monte Carlo simulation of electromagnetic interactions of radiation with liquid water in the framework of the Geant4-DNA project. In *2006 IEEE Nuclear Science Symposium Conference Record*, volume 2, pages 796–802, 2006.
- [335] S. Incerti, G. Baldacchino, M. Bernal, R. Capra, C. Champion, Z. Francis, P. Guève, A. Mantero, B. Mascialino, P. Moretto, P. Nieminen, C. Villagrasa, and C. Zacharatou. The Geant4-DNA project. *International Journal of Modeling, Simulation, and Scientific Computing*, 01(02):157–178, 2010.
- [336] S. Incerti, M. Douglass, S. Penfold, S. Guatelli, and E. Bezak. Review of Geant4-DNA applications for micro and nanoscale simulations. *Physica Medica*, 32(10):1187 – 1200, 2016.
- [337] J. Schuemann, A. L. McNamara, J. Ramos-Méndez, J. Perl, K. D. Held, H. Paganetti, S. Incerti, and B. Faddegon. TOPAS-nBio: An extension to the TOPAS simulation toolkit for cellular and sub-cellular radiobiology. *Radiation Research*, 191(2):125 – 138, 2018.
- [338] J Ramos-Méndez, J Perl, J Schuemann, A McNamara, H Paganetti, and B Faddegon. Monte Carlo simulation of chemistry following radiolysis with TOPAS-nBio. *Physics in Medicine & Biology*, 63(10):105014, 2018.
- [339] José Ramos-Méndez, Lucas N Burigo, Reinhard Schulte, Cynthia Chuang, and Bruce Faddegon. Fast calculation of nanodosimetric quantities in treatment planning of proton and ion therapy. *Physics in Medicine & Biology*, 63(23):235015, 2018.
- [340] Aimee L McNamara, José Ramos-Méndez, Joseph Perl, Kathryn Held, Naoki Dominguez, Eduardo Moreno, Nicholas T Henthorn, Karen J Kirkby, Sylvain Meylan, Carmen Villagrasa, Sebastien Incerti, Bruce Faddegon, Harald Paganetti, and Jan Schuemann.

- Geometrical structures for radiation biology research as implemented in the TOPAS-nBio toolkit. *Physics in Medicine & Biology*, 63(17):175018, 2018.
- [341] Hongyu Zhu, Yizheng Chen, Wonmo Sung, Aimee L McNamara, Linh T Tran, Lucas N Burigo, Anatoly B Rosenfeld, Junli Li, Bruce Faddegon, Jan Schuemann, and Harald Paganetti. The microdosimetric extension in TOPAS: development and comparison with published data. *Physics in Medicine & Biology*, 64(14):145004, 2019.
- [342] Hongyu Zhu, Aimee L. McNamara, Stephen J. McMahon, Jose Ramos-Mendez, Nicholas T. Henthorn, Bruce Faddegon, Kathryn D. Held, Joseph Perl, Junli Li, Harald Paganetti, and Jan Schuemann. Cellular response to proton irradiation: a simulation study with TOPAS-nBio. *Radiation Research*, 194(1):9–21, 2020.
- [343] B. G. Wouters, G. K. Y. Lam, U. Oelfke, K. Gardey, R. E. Durand, and L. D. Skarsgard. Measurements of relative biological effectiveness of the 70 MeV proton beam at TRIUMF using Chinese hamster V79 cells and the high-precision cell sorter assay. *Radiation Research*, 146(2):159–170, 1996.
- [344] L. D. Skarsgard, M. W. Skwarchuk, and B. G. Wouters. The survival of asynchronous v79 cells at low radiation doses: Modeling the response of mixed cell populations. *Radiation Research*, 138(1):S72–S75, 1994.
- [345] UCL High Energy Physics Group Proton Beam Therapy. Clatterbridge simulation model. Available at: <http://www.hep.ucl.ac.uk/pbt/wiki/Clatterbridge>, 2020. Accessed: 11-08-2020
- [346] J S L Yap, R Schnuerer, C P Welsch, N J S Bal, M Fransen, F Linde, J L Parsons, and A Kacpererek. Beam characterisation using Medipix3 and EBT3 film at the Clatterbridge proton therapy beamline. In *Proceedings of IPAC2019*, 8–11, 2019.
- [347] Jacinta Yap. *Characterisation Studies of Proton Beamlines for Medical Applications and Beam Diagnostics Integration*. PhD thesis, The University of Liverpool, 2020.
- [348] The Clatterbridge Cancer Centre. Eye proton therapy. Available at: <https://www.clatterbridgecc.nhs.uk/professionals/physics-department/cyclotron>, 2015. Accessed: 17-08-2020.
- [349] A. Kacpererek. Protontherapy of eye tumours in the UK: A review of treatment at Clatterbridge. *Applied Radiation and Isotopes*, 67(3):378–386, 2009.

- [350] Rachel J. Carter, Catherine M. Nickson, James M. Thompson, Andrzej Kacperek, Mark A. Hill, and Jason L. Parsons. Complex DNA damage induced by high linear energy transfer alpha-particles and protons triggers a specific cellular DNA damage response. *International Journal of Radiation Oncology • Biology, • Physics*, 100(3):776–784, 2018.
- [351] Carlos Granja, Karel Kudela, Jan Jakubek, Pavel Krist, David Chvatil, Jan Stursa, and Stepan Polansky. Directional detection of charged particles and cosmic rays with the miniaturized radiation camera minipix timepix. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 911:142 – 152, 2018.
- [352] X. Llopart, R. Ballabriga, M. Campbell, L. Tlustos, and W. Wong. Timepix, a 65k programmable pixel readout chip for arrival time, energy and/or photon counting measurements. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 581(1):485 – 494, 2007.
- [353] Rafael Ballabriga, Michael Campbell, and Xavier Llopart. Asic developments for radiation imaging applications: The medipix and timepix family. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 878:10 – 23, 2018.
- [354] Peter Rubovič, Benedikt Bergmann, Daniela Ekendahl, Jiří Hůlka, Libor Judas, Zdeněk Kohout, Stanislav Pospíšil, and Ivan Štekl. Timepix detector as a tool for x-ray and gamma dosimetry. *Radiation Measurements*, 107:39 – 42, 2017.
- [355] D Turecek, L Pinsky, J Jakubek, Z Vykydal, N Stoffle, and S Pospisil. Small dosimeter based on timepix device for international space station. *Journal of Instrumentation*, 6(12):C12037–C12037, 2011.
- [356] Nicholas Stoffle, Lawrence Pinsky, Martin Kroupa, Son Hoang, John Idarraga, Clif Amberboy, Ryan Rios, Jessica Hauss, John Keller, Amir Bahadori, Edward Semones, Daniel Turecek, Jan Jakubek, Zdenek Vykydal, and Stanislav Pospisil. Timepix-based radiation environment monitor measurements aboard the international space station. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 782:143 – 148, 2015.
- [357] Martin Kroupa, Amir Bahadori, Thomas Campbell-Ricketts, Anton Empl, Son Minh Hoang, John Idarraga-Munoz, Ryan Rios, Edward Semones, Nicholas Stoffle, Lukas Tlustos, Daniel Turecek, and Lawrence Pinsky. A semiconductor radiation imaging pixel detector for space radiation dosimetry. *Life Sciences in Space Research*, 6:69 – 78, 2015.

- [358] Carlos Granja, Jan Jakubek, Stepan Polansky, Vaclav Zach, Pavel Krist, David Chvatil, Jan Stursa, Marek Sommer, Ondrej Ploc, Satoshi Kodaira, and Maria Martisikova. Resolving power of pixel detector timepix for wide-range electron, proton and ion detection. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 908:60 – 71, 2018.
- [359] D. Turecek, J. Jakubek, and P. Soukup. USB 3.0 readout and time-walk correction method for timepix3 detector. *Journal of Instrumentation*, 11(12):C12065–C12065, 2016.
- [360] Rui Zhang and Wayne D Newhauser. Calculation of water equivalent thickness of materials of arbitrary density, elemental composition and thickness in proton beam irradiation. *Physics in Medicine and Biology*, 54(6):1383–1395, 2009.
- [361] L Opalka, C Granja, B Hartmann, J Jakubek, O Jaekel, M Martisikova, S Pospisil, and J Solc. 3D measurement of the radiation distribution in a water phantom in a hadron therapy beam. *Journal of Instrumentation*, 7(01):C01085–C01085, 2012.
- [362] L Opalka, C Granja, B Hartmann, J Jakubek, O Jaekel, M Martisikova, S Pospisil, and J Solc. Linear energy transfer and track pattern recognition of secondary radiation generated in hadron therapy beam in a PMMA target. *Journal of Instrumentation*, 8(02):C02047–C02047, 2013.
- [363] David Freedman and Persi Diaconis. On the histogram as a density estimator:l2 theory. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 57(4):453–476, 1981.
- [364] L. Landau. On the energy loss of fast particles by ionization. *J. Phys. (USSR)*, 8:201–205, 1944.
- [365] D.H. Wilkinson. Ionization energy loss by charged particles part i. the landau distribution. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 383(2):513 – 515, 1996.
- [366] Paulina Stasica, Jakub Baran, Carlos Granja, Nils Krah, Grzegorz Korcyl, Cristina Oancea, Monika Pawlik-Niedźwiecka, Szymon Niedźwiecki, Marzena Rydygier, Angelo Schiavi, Antoni Rucinski, and Jan Gajewski. A simple approach for experimental characterization and validation of proton pencil beam profiles. *Frontiers in Physics*, 8:346, 2020.
- [367] E.R. Benton, E.V. Benton and A.L. Frank. Conversion between different forms of LET. *Radiation Measurements*, 45(8):957-959, 2010.

- [368] Shin J. Schuemann J. Faddegon B. Paganetti H. Ramos J. McNamara A. Hall A. Perl, J. TOPAS MC, Inc. Available at: <http://www.topasmc.org/user-guide>, 2020. Accessed: 04-09-2020.
- [369] T Bortfeld, J Bürkelbach, R Boesecke, and W Schlegel. Methods of image reconstruction from projections applied to conformation radiotherapy. *Physics in Medicine and Biology*, 35(10):1423–1434, 1990.
- [370] Yin Zhang and Michael Merritt. Dose-volume-based IMRT fluence optimization: A fast least-squares approach with differentiability. *Linear Algebra and its Applications*, 428(5):1365 – 1387, 2008. Special Issue on Linear and Nonlinear Models and Algorithms in Intensity-Modulated Radiation Therapy.
- [371] Yair Censor and Jan Unkelbach. From analytic inversion to contemporary IMRT optimization: Radiation therapy planning revisited from a mathematical perspective. *Physica Medica*, 28(2):109–118, 2012.
- [372] J. O. Deasy. Multiple local minima in radiotherapy optimization problems with dose–volume constraints. *Medical Physics*, 24(7):1157–1161, 1997.
- [373] Qiuwen Wu and Radhe Mohan. Multiple local minima in IMRT optimization based on dose–volume criteria. *Medical Physics*, 29(7):1514–1527, 2002.
- [374] Ali T Tuncel, Felisa Preciado, Ronald L Rardin, Mark Langer, and Jean-Philippe P Richard. Strong valid inequalities for fluence map optimization problem under dose-volume restrictions. *Annals of Operations Research*, 196(1):819–840, 2012.
- [375] Paul S. Cho, Shinhak Lee, Robert J. Marks II, Seho Oh, Steve G. Sutlief, and Mark H. Phillips. Optimization of intensity modulated beams with volume constraints using two methods: Cost function minimization and projections onto convex sets. *Medical Physics*, 25(4):435–443, 1998.
- [376] Spiridon V. Spirou and Chen-Shou Chui. A gradient inverse planning algorithm with dose-volume constraints. *Medical Physics*, 25(3):321–333, 1998.
- [377] G. S. Mageras and R. Mohan. Application of fast simulated annealing to optimization of conformal radiation treatments. *Medical Physics*, 20(3):639–647, 1993.
- [378] Scott Penfold, Rafał Zalas, Margherita Casiraghi, Mark Brooke, Yair Censor, and Reinhard Schulte. Sparsity constrained split feasibility for dose-volume constraints in inverse planning

- of intensity-modulated photon or proton therapy. *Physics in Medicine & Biology*, 62(9):3599, 2017.
- [379] H Edwin Romeijn, Ravindra K Ahuja, James F Dempsey, Arvind Kumar, and Jonathan G Li. A novel linear programming approach to fluence map optimization for intensity modulated radiation therapy treatment planning. *Physics in Medicine & Biology*, 48(21):3521, 2003.
- [380] H Edwin Romeijn, Ravindra K Ahuja, James F Dempsey, and Arvind Kumar. A new linear programming approach to radiation therapy treatment planning problems. *Operations Research*, 54(2):201–216, 2006.
- [381] Dionne M Aleman, Daniel Glaser, H Edwin Romeijn, and James F Dempsey. Interior point algorithms: guaranteed optimality for fluence map optimization in IMRT. *Physics in Medicine and Biology*, 55(18):5467–5482, 2010.
- [382] D.M. Aleman, V.V. Mišić, and M.B. Sharpe. Computational enhancements to fluence map optimization for total marrow irradiation using IMRT. *Computers Operations Research*, 40(9):2167 – 2177, 2013.
- [383] Sebastiaan Breedveld, Bas van den Berg, and Ben Heijmen. An interior-point implementation developed and tuned for radiation therapy treatment planning. *Computational Optimization and Applications*, 68(2):209–242, 2017.
- [384] Yair Censor, Adi Ben-Israel, Ying Xiao, and James M. Galvin. On linear infeasibility arising in intensity-modulated radiation therapy inverse planning. *Linear Algebra and Its Applications*, 428(5-6):1406–1420, 2008.
- [385] Mark Langer and Joseph Leong. Optimization of beam weights under dose-volume restrictions. *International Journal of Radiation Oncology • Biology, • Physics*, 13(8):1255–1260, 1987.
- [386] Mark Langer, Richard Brown, Marsha Urie, Joseph Leong, Michael Stracher, and Jeremy Shapiro. Large scale optimization of beam weights under dose-volume restrictions. *International Journal of Radiation Oncology • Biology, • Physics*, 18(4):887–893, 1990.
- [387] Mark Langer, Richard Brown, Peter Kijewski, and Chul Ha. The reliability of optimization under dose-volume limits. *International Journal of Radiation Oncology • Biology, • Physics*, 26(3):529–538, 1993.

- [388] Eva K Lee, Tim Fox, and Ian Crocker. Integer programming applied to intensity-modulated radiation therapy treatment planning. *Annals of Operations Research*, 119(1-4):165–181, 2003.
- [389] Felisa Preciado-Walters, Ronald Rardin, Mark Langer, and Van Thai. A coupled column generation, mixed integer approach to optimal planning of intensity modulated radiation therapy for cancer. *Mathematical Programming*, 101(2):319–338, 2004.
- [390] Michael C Ferris, Robert R Meyer, and Warren D’Souza. Radiation treatment planning: Mixed integer programming formulations and approaches. In *Handbook on modelling for discrete optimization*, pages 317–340. Springer, 2006.
- [391] H. Rocha, J.M. Dias, B.C. Ferreira, and M.C. Lopes. Discretization of optimal beamlet intensities in IMRT: A binary integer programming approach. *Mathematical and Computer Modelling*, 55(7):1969 – 1980, 2012.
- [392] Yihua Lan, Cunhua Li, Haozheng Ren, Yong Zhang, and Zhifang Min. Fluence map optimization (fmo) with dose–volume constraints in IMRT using the geometric distance sorting method. *Physics in Medicine & Biology*, 57(20):6407, 2012.
- [393] Maryam Zaghian, Gino Lim, Wei Liu, and Radhe Mohan. An automatic approach for satisfying dose-volume constraints in linear fluence map optimization for impt. *Journal of cancer therapy*, 5(2):198, 2014.
- [394] Pengfei Zhang, Neng Fan, Jie Shan, Steven E. Schild, Martin Bues, and Wei Liu. Mixed integer programming with dose-volume constraints in intensity-modulated proton therapy. *Journal of Applied Clinical Medical Physics*, 18(5):29–35, 2017.
- [395] Thomas Bortfeld. Clinically relevant intensity modulation optimization using physical criteria. In *XII International Conference on the Use of Computers in Radiation Therapy, 1997*. Medical Physics Publishing, 1997.
- [396] Qiuwen Wu and Radhe Mohan. Algorithms and functionality of an intensity modulated radiotherapy optimization system. *Medical physics*, 27(4):701–711, 2000.
- [397] Sebastiaan Breedveld, Pascal RM Storchi, Marleen Keijzer, Arnold W Heemink, and Ben JM Heijmen. A novel approach to multi-criteria inverse planning for IMRT. *Physics in Medicine & Biology*, 52(20):6339, 2007.
- [398] Marita Falkinger, Stefan Schell, Johannes Müller, and Jan J Wilkens. Prioritized optimization in intensity modulated proton therapy. *Zeitschrift für Medizinische Physik*, 22(1):21–28, 2012.

- [399] Sovanlal Mukherjee, Linda Hong, Joseph O. Deasy, and Masoud Zarepisheh. Integrating soft and hard dose-volume constraints into hierarchical constrained IMRT optimization. *Medical Physics*, 47(2):414–421, 2020.
- [400] Anqi Fu, Bari Ungun, Lei Xing, and Stephen Boyd. A convex optimization approach to radiation treatment planning with dose constraints. *Optimization and Engineering*, 20(1):277–300, 2019.
- [401] Mark Langer, Eva K Lee, Joseph O Deasy, Ronald L Rardin, and James A Deye. Operations research applied to radiotherapy, an nci-nsf-sponsored workshop february 7–9, 2002. *International Journal of Radiation Oncology • Biology • Physics*, 57(3):762–768, 2003.
- [402] Benjamin E Nelms, Greg Robinson, Jay Markham, Kyle Velasco, Steve Boyd, Sharath Narayan, James Wheeler, and Mark L Sobczak. Variation in external beam treatment plan quality: An inter-institutional study of planners and planning systems. *Practical Radiation Oncology*, 2:296–305, 2012.
- [403] Margie A Hunt, Ching-Yeh Hsiung, Spirodon V Spirou, Chen-Shou Chui, Howard I Amols, and Clifton C Ling. Evaluation of concave dose distributions created using an inverse planning system. *International Journal of Radiation Oncology • Biology, • Physics*, 54(3):953–962, 2002.
- [404] M Monz, K. H. Küfer, T. R. Bortfeld, and C. Thieke. Pareto navigation-algorithmic foundation of interactive multi-criteria IMRT planning. *Physics in medicine and biology*, 53:985–998, 2008.
- [405] H.W. Hamacher and K.-H. Küfer. Inverse radiation therapy planning — a multiple objective optimization approach. *Discrete Applied Mathematics*, 118(1):145 – 161, 2002. Special Issue devoted to the ALIO-EURO Workshop on Applied Combinatorial Optimization.
- [406] Tarek Halabi, David Craft, and Thomas Bortfeld. Dose-volume objectives in multi-criteria optimization. *Physics in Medicine and Biology*, 51(15):3809–3818, 2006.
- [407] C. Thieke, K. H. Küfer, M. Monz, A. Scherrer, F. Alonso, U. Oelfke, P. E. Huber, J. Debus, and T. Bortfeld. A new concept for interactive radiotherapy planning with multicriteria optimization: First clinical evaluation. *Radiotherapy and Oncology*, 85:292–298, 2007.
- [408] R Bokrantz and K Miettinen. Improved plan quality in multicriteria radiation therapy optimization by projections onto the Pareto surface. 42(10):5862–5870, 2012.

- [409] Qiuwen Wu, Radhe Mohan, Andrzej Niemierko, and Rupert Schmidt-Ullrich. Optimization of intensity-modulated radiotherapy plans based on the equivalent uniform dose. *International Journal of Radiation Oncology • Biology, • Physics*, 52(1):224 – 235, 2002.
- [410] Timothy C.Y. Chan, Houra Mahmoudzadeh, and Thomas G. Purdie. A robust-cvar optimization approach with application to breast cancer therapy. *European Journal of Operational Research*, 238(3):876 – 885, 2014.
- [411] Nilay Noyan, Gábor Rudolf, and Andrzej Ruszczyński. Relaxations of linear programming problems with first order stochastic dominance constraints. *Operations Research Letters*, 34(6):653 – 659, 2006.
- [412] Yair Censor. *Mathematical Aspects of Radiation Therapy Treatment Planning: Continuous Inversion Versus Full Discretization and Optimization Versus Feasibility*, pages 101–112. Springer New York, New York, NY, 1999.
- [413] Y. Xiao, D. Michalski, J.M. Galvin, and Y. Censor. The least-intensity feasible solution for aperture-based inverse planning in radiation therapy. *Annals of Operations Research*, 119(1):183–203, 2003.
- [414] Wei Chen, G Herman, and Yair Censor. Algorithms for satisfying dose volume constraints in intensity-modulated radiation therapy. *Mathematical methods in biomedical imaging and intensity-modulated radiation therapy (IMRT)*, pages 97–106, 2008.
- [415] Ran Davidi, Yair Censor, Reinhard W Schulte, Sarah Geneser, and Lei Xing. Feasibility-seeking and superiorization algorithms applied to inverse treatment planning in radiation therapy. *Contemp. Math.*, 636:83–92, 2015.
- [416] S. Penfold, M. Casiraghi, T. Dou, R. Schulte, and Y. Censor. A feasibility-seeking algorithm applied to planning of intensity modulated proton therapy: A proof of principle study. *Medical Physics*, 42:3338, 2015.
- [417] Mark Brooke, Yair Censor, and Aviv Gibali. (In press). Dynamic string-averaging CQ-methods for the split feasibility problem with percentage violation constraints arising in radiation therapy treatment planning. *International Transactions in Operational Research*, <https://doi.org/10.1111/itor.12929>.
- [418] Yair Censor and Tommy Elfving. A multiprojection algorithm using bregman projections in a product space. *Numerical Algorithms*, 8(2):221–239, 1994.

- [419] Qingzhi Yang. The relaxed cq algorithm solving the split feasibility problem. *Inverse problems*, 20(4):1261, 2004.
- [420] Eyal Masad and Simeon Reich. A note on the multiple-set split convex feasibility problem in hilbert space. *Journal of Nonlinear and Convex Analysis*, 8(3):367, 2007.
- [421] Genaro López, Victoria Martín-Márquez, Fenghui Wang, and Hong-Kun Xu. Solving the split feasibility problem without prior knowledge of matrix norms. *Inverse Problems*, 28(8):085004, 2012.
- [422] A Latif, J Vahidi, and M Eslamian. Strong convergence for generalized multiple-set split feasibility problem. *Filomat*, 30(2):459–467, 2016.
- [423] Aviv Gibali, Li-Wei Liu, and Yu-Chao Tang. Note on the modified relaxation cq algorithm for the split feasibility problem. *Optimization Letters*, 12(4):817–830, 2018.
- [424] Hedy Attouch, Jérôme Bolte, and Benar Fux Svaiter. Convergence of descent methods for semi-algebraic and tame problems: proximal algorithms, forward–backward splitting, and regularized gauss–seidel methods. *Mathematical Programming*, 137(1-2):91–129, 2013.
- [425] Robert Hesse, D Russell Luke, and Patrick Neumann. Alternating projections and douglas-rachford for sparse affine feasibility. *IEEE Transactions on Signal Processing*, 62(18):4868–4881, 2014.
- [426] Heinz H Bauschke, Hung M Phan, and Xianfu Wang. The method of alternating relaxed projections for two nonconvex sets. *Vietnam Journal of Mathematics*, 42(4):421–450, 2014.
- [427] Aviv Gibali, Shoham Sabach, and Sergey Voldman. Non-convex split feasibility problems: Models, algorithms and theory. arXiv:2010.04504, 2020.
- [428] Peiyuan Wang, Jianjun Zhou, Risheng Wang, and Jie Chen. New generalized variable stepsizes of the CQ algorithm for solving the split feasibility problem. *Journal of inequalities and applications*, 2017(1):135, 2017.
- [429] Yair Censor and Tommy Elfving. A multiprojection algorithm using Bregman projections in a product space. *Numerical Algorithms*, 8(2):221–239, 1994.
- [430] Charles Byrne. Iterative oblique projection onto convex sets and the split feasibility problem. *Inverse Problems*, 18(2):441–453, 2002.
- [431] Charles Byrne. A unified treatment of some iterative algorithms in signal processing and image reconstruction. *Inverse Problems*, 20(1):103–120, 2004.

- [432] Fenghui Wang and Hong Kun Xu. Cyclic algorithms for split feasibility problems in Hilbert spaces. *Nonlinear Analysis, Theory, Methods and Applications*, 74(12):4105–4111, 2011.
- [433] Peiyuan Wang, Jianjun Zhou, Risheng Wang, and Jie Chen. New generalized variable stepsizes of the CQ algorithm for solving the split feasibility problem. *Journal of Inequalities and Applications*, 2017(1):135, 2017.
- [434] Qingzhi Yang. On variable-step relaxed projection algorithm for variational inequalities. *Journal of Mathematical Analysis and Applications*, 302(1):166–179, 2005.
- [435] Yair Censor and Alexander Segal. On the string averaging method for sparse common fixed-point problems. *International Transactions in Operational Research*, 16(4):481–494, 2009.
- [436] Andrzej Cegielski. *Iterative methods for fixed point problems in Hilbert spaces*, volume 2057. Springer, 2012.
- [437] Andrzej Cegielski. General Method for Solving the Split Common Fixed Point Problem. *Journal of Optimization Theory and Applications*, 165(2):385–404, 2015.
- [438] Andrzej Cegielski and Fadhel Al-Musallam. Strong convergence of a hybrid steepest descent method for the split common fixed point problem. *Optimization*, 65(7):1463–1476, 2016.
- [439] Yair Censor, Martin D Altschuler, and William D Powlis. On the use of cimmino’s simultaneous projections method for computing a solution of the inverse problem in radiation therapy treatment planning. *Inverse Problems*, 4(3):607, 1988.
- [440] Ron Aharoni and Yair Censor. Block-iterative projection methods for parallel computation of solutions to convex feasibility problems. *Linear Algebra and Its Applications*, 120:165–175, 1989.
- [441] S. N. Penfold, R. W. Schulte, Y. Censor, V. Bashkirov, S. McAllister, K. E. Schubert, and A. B. Rosenfeld. Block-iterative and string-averaging projection algorithms in proton computed tomography image reconstruction. In Y. Censor, M. Jiang, and G. Wang, editors, *Biomedical Mathematics: Promising Directions in Imaging, Therapy Planning and Inverse Problems*, pages 347–367. Medical Physics Publishing, Madison, WI, USA, 2010.
- [442] Scott Penfold and Yair Censor. Techniques in Iterative Proton CT Image Reconstruction. *Sensing and Imaging*, 16(1):19, 2015.

- [443] Emmanuel J Candes, Michael B Wakin, and Stephen P Boyd. Enhancing sparsity by reweighted ℓ_1 minimization. *Journal of Fourier Analysis and Applications*, 14(5-6):877–905, 2008.
- [444] Joseph Mantel, Harold Perry, and James J Weinkam. Automatic variation of field size and dose rate in rotation therapy. *International Journal of Radiation Oncology • Biology, • Physics*, 2(7-8):697–704, 1977.
- [445] Paul S. Cho, Shinhak Lee, Robert J. Marks II, Seho Oh, Steve G. Sutlief, and Mark H. Phillips. Optimization of intensity modulated beams with volume constraints using two methods: Cost function minimization and projections onto convex sets. *Medical Physics*, 25(4):435–443.
- [446] Hojin Kim, Stephen Becker, Rena Lee, Soonhyouk Lee, Sukyoung Shin, Emmanuel Candès, Lei Xing, and Ruijiang Li. Improving IMRT delivery efficiency with reweighted ℓ_1 -minimization for inverse planning. *Medical Physics*, 40(7):071719.
- [447] Sebastiaan Breedveld and Ben Heijmen. Data for trots – the radiotherapy optimisation test set. *Data in Brief*, 12:143 – 149, 2017.
- [448] Peter W.J. Voet, Maarten L.P. Dirkx, Sebastiaan Breedveld, Dennie Fransen, Peter C. Levendag, and Ben J.M. Heijmen. Toward fully automated multicriterial plan generation: A prospective clinical study. *International Journal of Radiation Oncology • Biology, • Physics*, 85(3):866 – 872, 2013.
- [449] S van de Water, A C Kraan, S Breedveld, W Schillemans, D N Teguh, H M Kooy, T M Madden, B J M Heijmen, and M S Hoogeman. Improved efficiency of multi-criteria IMPT treatment planning using iterative resampling of randomly placed pencil beams. *Physics in Medicine and Biology*, 58(19):6969–6983, 2013.
- [450] Brad Warkentin, Pavel Stavrev, Nadia Stavreva, Colin Field, and B. Gino Fallone. A TCP-NTCP estimation module using DVHs and known radiobiological models and parameter sets. *Journal of Applied Clinical Medical Physics*, 5(1):50–63, 2004.
- [451] J.P. Nolan. *Stable Distributions: Models for Heavy-Tailed Data*. Springer New York, 2016.