



Université
de Toulouse

THÈSE

En vue de l'obtention du

DOCTORAT DE L'UNIVERSITÉ DE TOULOUSE

Délivré par : *l'Université Toulouse 3 Paul Sabatier (UT3 Paul Sabatier)*

Présentée et soutenue le *04/10/2017* par :

Richard BROWN

Microbrachytherapy treatment planning

JURY

LYDIA MAIGNE
DAVID SARRUT
LUDOVIC FERRER
ISABELLE BERRY
XAVIER FRANCERIES
MANUEL BARDIÈS

Maître de Conférences, HDR
Directeur de Recherche, HDR
Physicien Médical
Professeur d'Université, HDR
Maître de Conférences, HDR
Directeur de Recherche, HDR

Rapporteur
Rapporteur
Examinateur
Examinatrice
Directeur
Co-encadrant

École doctorale et spécialité :

GEET : Radio-physique et Imagerie Médicale

Unité de Recherche :

Centre de Recherche en Cancérologie de Toulouse (INSERM UMR 1037, Équipe 15)

Directeur de Thèse :

Xavier FRANCERIES

Rapporteurs :

Lydia MAIGNE et David SARRUT

REMERCIEMENTS

Avant de me lancer dans ce manuscrit, j'aimerais prendre un moment pour remercier les personnes, sans qui ces travaux n'auraient pas été possibles.

Premièrement, je tiens à remercier Xavier, mon directeur de thèse, pour m'avoir guidé pendant ces trois ans. Grâce à ta supervision, j'ai eu l'occasion d'approfondir énormément mes connaissances dans ce vaste domaine.

Un grand merci aux autres statutaires de l'équipe – Manuel (el jefe), Marie-Claude, Luc, Emmanuelle et Laure – merci pour tous les conseils indispensables que vous m'avez prodigué, surtout vers la fin quand j'en avais le plus besoin.

Sara B, Daphné, Sara M, Alice et Alan, vous faisiez partie de l'équipe 15 au début de ma thèse. A ce moment-là, j'étais moins à l'aise avec la langue française (pour être gentil), merci d'avoir été patients avec moi.

Julien, nous avons commencé nos thèses le même jour et donc nous avons vécu la plupart des moments de la thèse ensemble. Et honnêtement, je n'aurai pas espéré avoir partagé ces moments avec quelqu'un d'aussi gentil que toi. Il faudra que je trouve un autre expert en vin, en rhum (peut-être pas en viande) pour me conseiller dans les prochaines étapes de ma vie. Bon courage pour la suite !

Jonathan, nous avons beaucoup travaillé ensemble dans le cadre de TheraneaM. Tu m'as corrigé des milliards de mails écrits dans un français d'un enfant âgé de cinq ans. Encore, je vais te solliciter pour corriger ces remerciements ; mais je te promets, ceux-ci seront les derniers mots que tu vas corriger... Merci, tu es vraiment quelqu'un de super.

L'équipe latino – Gustavo, Erick – merci pour les cafés, sans lesquels je n'aurais sûrement jamais pu finir ma thèse... Mais encore plus important que le café, merci surtout pour tous les moments de partage : de musique, de nourriture, d'escalade, de danse. Je vous promets que je vous rendrai visite au Brésil et au Costa Rica aussi vite que possible !

Les autres garçons de l'équipe 15 – Tony, Fred « le dompteur de kangourous », Maxime et Jérémy – bon courage pour vos travaux à venir. Je ne m'inquiète pas pour vous, vous êtes ensemble pour vous entraider et êtes tous vraiment top.

Louise et Jérémy, récemment partis du CRCT. Vous êtes tous les deux des personnes que j'apprécie énormément. Profitez de ces moments de liberté en me rendant visite à Londres ! On se verra très bientôt.

À mes coloc, Gaétan et Corto (et des squatteurs occasionnels, Hugo et Flore), merci pour TOUT ! Pour les repas partagés, pour les cours de français (même s'ils étaient parfois intentionnellement erronés), pour les week-ends à l'étranger, pour les week-ends de ski/surf/wake, pour les concerts de jazz... Bref, merci mille fois, vous avez vraiment forgé mes trois ans en France !

Lastly, the biggest thank you is reserved for my parents. Thank you for the support, both academic and personal, the whole way through my life. Thank you even more for the last three years. I would have never have thought that even at the age of twenty-five, my dad would still be my chief-editor for all written work. Thank you, thank you, I love you both to pieces.

CONTENTS

LIST OF TABLES.....	VI
LIST OF FIGURES	VIII
LIST OF ABBREVIATIONS AND ACRONYMS.....	XVIII
RESUME (FR).....	1
INTRODUCTION (FR)	4
INTRODUCTION.....	13
1 GENERALITIES	16
1.1 CANCER AND RADIOTHERAPY	16
1.2 MODALITIES OF RADIOTHERAPY	16
1.2.1 <i>EXTERNAL BEAM RADIOTHERAPY</i>	17
1.2.2 <i>BRACHYTHERAPY</i>	17
1.2.3 <i>NUCLEAR MEDICINE</i>	18
1.3 MICROBRACHYTHERAPY	18
1.3.1 <i>THE TREATMENT</i>	19
1.4 TREATMENT PLANNING.....	19
1.5 THE CHOICE OF RADIOELEMENT	20
1.6 GLIOBLASTOMAS	20
1.7 TREATMENT REQUIREMENTS	21
1.7.1 <i>IRRADIATION OF THE TUMOUR</i>	22
1.7.2 <i>ORGANS AT RISK</i>	23
1.8 CONCLUSION	23
2 TREATMENT PLANNING.....	24
2.1 INTRODUCTION	24
2.2 THE FORWARD PROBLEM	25
2.3 THE COST FUNCTION	26

2.3.1	<i>TYPES OF COST FUNCTIONS</i>	27
2.3.2	<i>PHYSICAL COST FUNCTIONS</i>	30
2.3.3	<i>RADIOBIOLOGICAL COST FUNCTIONS</i>	31
2.3.4	<i>CHOSEN COST FUNCTIONS</i>	32
2.4	OPTIMISATION ALGORITHMS.....	33
2.4.1	<i>INTERACTIVE ALGORITHMS</i>	34
2.4.2	<i>SINGLE-OBJECTIVE ALGORITHMS</i>	34
2.4.3	<i>MULTI-OBJECTIVE ALGORITHMS</i>	35
2.5	EXPLANATION OF VARIOUS ALGORITHMS.....	37
2.5.1	<i>CONJUGATE GRADIENT AND SIMILAR ALGORITHMS</i>	38
2.5.2	<i>NELDER-MEAD SIMPLEX</i>	39
2.5.3	<i>BEES ALGORITHM</i>	44
2.5.4	<i>NON-DOMINATED SORTING GENETIC ALGORITHM-II</i>	48
2.6	STOPPING CRITERIA.....	54
2.6.1	<i>CONVERGENCE</i>	55
2.6.2	<i>NUMBER OF ITERATIONS</i>	55
2.6.3	<i>TREATMENT PLAN UPDATES</i>	55
2.7	CONCLUSION.....	56
3	MICROBRACHYTHERAPY DOSIMETRY.....	58
3.1	INTRODUCTION.....	58
3.2	KERNEL GENERATION.....	60
3.2.1	<i>S-VALUES</i>	61
3.2.2	<i>GATE</i>	61
3.2.3	<i>MODELLING THE INJECTION</i>	61
3.2.4	<i>GATE PARAMETERS</i>	62
3.2.5	<i>SIMULATION TIMES</i>	63
3.3	PATIENT ABSORBED DOSE DISTRIBUTIONS.....	64
3.4	STRUCTURE-SPECIFIC DOSIMETRY.....	65
3.5	CONSIDERATIONS.....	66
3.5.1	<i>ASSUMPTION OF HOMOGENEITY</i>	66
3.5.2	<i>RADIOELEMENT DISTRIBUTION AND MICROSPHERE CASING</i>	67
3.5.3	<i>CALCULATION TIMES</i>	68
3.6	A WORKED EXAMPLE.....	71
3.7	CONCLUSION.....	75
4	COMPARISON OF OPTIMISATION ALGORITHMS.....	77
4.1	INTRODUCTION.....	77
4.2	INJECTION PARAMETERS.....	79

4.3	TEST CASE 1: FIXED NUMBER OF INJECTIONS AND NO OAR.....	80
4.3.1	<i>NMS</i>	80
4.3.2	<i>BA</i>	83
4.3.3	<i>NSGA-II</i>	84
4.3.4	<i>COMPARISON</i>	86
4.3.5	<i>DISCUSSION</i>	87
4.4	TEST CASE 2: VARIABLE NUMBER OF INJECTIONS AND NO OAR.....	88
4.4.1	<i>NMS</i>	89
4.4.2	<i>BA</i>	90
4.4.3	<i>NSGA-II</i>	90
4.4.4	<i>COMPARISON</i>	94
4.4.5	<i>DISCUSSION</i>	95
4.5	TEST CASE 3: VARIABLE NUMBER OF INJECTIONS AND OAR	95
4.5.1	<i>NSGA-II</i>	96
4.5.2	<i>COMPARISON</i>	97
4.5.3	<i>DISCUSSION</i>	98
4.6	DISCUSSION.....	99
4.6.1	<i>COSTLET WEIGHTINGS</i>	99
4.6.2	<i>THE PARETO FRONT IN 4D (OR MORE)</i>	100
4.6.3	<i>MAXIMUM NUMBER OF UPDATES AS A STOPPING CRITERION</i>	101
4.6.4	<i>IMPROVING NMS</i>	101
4.6.5	<i>OPTIMISATION TIMES</i>	102
4.6.6	<i>THE CHOSEN ALGORITHM</i>	103
4.7	VOXEL SIZES	104
4.7.1	<i>OPTIMISATION</i>	105
4.7.2	<i>DVHS</i>	106
4.7.3	<i>DISCUSSION</i>	107
4.7.4	<i>CONCLUSION</i>	108
4.8	OPTIMISATION OF A REAL TUMOUR	108
4.8.1	<i>TUMOUR SEGMENTATION</i>	109
4.8.2	<i>TREATMENT PLANNING</i>	110
4.8.3	<i>RESULTS FROM THE CHOSEN TREATMENT PLAN</i>	111
4.8.4	<i>DISCUSSION</i>	113
4.9	CONCLUSION.....	114
5	COMPARISON OF INJECTIONS.....	116
5.1	INTRODUCTION	116
5.2	THE CHOICE OF RADIOELEMENT	117

5.2.1	<i>KERNEL GENERATION</i>	118
5.2.2	<i>PARETO OPTIMISATION</i>	119
5.2.3	<i>DVH COMPARISON</i>	120
5.2.4	<i>DISCUSSION</i>	121
5.2.5	<i>CONCLUSION</i>	122
5.3	INITIAL ACTIVITY AND INJECTED VOLUME	122
5.3.1	<i>INITIAL ACTIVITY</i>	122
5.3.2	<i>INJECTION VOLUME</i>	124
5.3.3	<i>DISCUSSION</i>	126
5.4	LARGER TUMOURS	127
5.4.1	<i>CONCLUSION</i>	129
5.5	CYLINDRICAL INJECTIONS	129
5.5.1	<i>THE INJECTOR AND THE INJECTIONS</i>	130
5.5.2	<i>ABSORBED DOSE DISTRIBUTION OF CYLINDRICAL INJECTIONS</i>	131
5.5.3	<i>COMPARISON WITH SPHERICAL INJECTIONS</i>	134
5.5.4	<i>DISCUSSION</i>	135
5.6	CONCLUSION.....	136
	CONCLUSIONS AND FUTURE PERSPECTIVES	138
	CONCLUSIONS ET PERSPECTIVES FUTURES (FR)	143
	BIBLIOGRAPHY	153
6	APPENDIX: CREATION OF THE TPS	156
6.1	INTRODUCTION	156
6.2	DOSITOUL STRUCTURE	157
6.3	USING DOSITOUL.....	159
6.3.1	<i>DATABASE</i>	159
6.3.2	<i>STRUCTURE EDITOR</i>	161
6.3.3	<i>OPTIMISATION BROWSER</i>	163
6.3.4	<i>SETTING UP A NEW OPTIMISATION</i>	163
6.3.5	<i>PERFORMING THE OPTIMISATION</i>	168
6.3.6	<i>VIEWING OPTIMISATION RESULTS</i>	169
6.4	C++	171
6.4.1	<i>QT</i>	172
6.4.2	<i>VTK</i>	172
6.4.3	<i>GDCM</i>	173
6.5	DOSITOUL CLASSES	173
6.5.1	<i>DT2_DATABASE, DT2_PATIENT AND DT2_STUDY</i>	174

6.5.2	<i>DT2_SERIES</i>	174
6.5.3	<i>DT2_DICOM</i>	174
6.5.4	<i>DT2_OPTIMISATION</i>	174
6.5.5	<i>DT2_ARRAY3D</i>	175
6.5.6	<i>DT2_TREATMENTPLAN</i>	176
6.5.7	<i>DT2_STRUCTURE</i>	176
6.5.8	<i>DT2_DVH</i>	177
6.5.9	<i>DT2_COSTLET</i>	177
6.5.10	<i>DT2_LINE, DT2_MATRIX3X3, DT2_POINT AND DT2_VEC3D</i>	177
6.5.11	<i>DT2_INJECTION</i>	178
6.5.12	<i>DT2_VISUALISATION</i>	179
6.6	CONCLUSION.....	180
SCIENTIFIC PRODUCTIONS		182

LIST OF TABLES

TABLE 2.1: NON-EXHAUSTIVE LIST OF PHYSICAL COST FUNCTIONS USED IN RADIOTHERAPY. VOI IS A VOLUME OF INTEREST. VOXELS ARE 3D PIXELS.	28
TABLE 2.2: NON-EXHAUSTIVE LIST OF RADIOBIOLOGICAL COST FUNCTIONS USED IN RADIOTHERAPY. SEE SECTION 2.3.3 FOR A DISCUSSION OF A , Γ_{50} , TD50 AND TCD50.	29
TABLE 2.3: SUMMARY OF THREE ALGORITHMS THAT HAVE BEEN PRESENTED IN THIS CHAPTER AND WILL BE COMPARED IN CHAPTER 4.	57
TABLE 3.1: SIMULATION TIMES FOR KERNELS USED THROUGHOUT THIS WORK. SIMULATION TIMES ARE SHOWN FOR A SINGLE CORE OF A 2.5 GHz MACBOOK PRO. THE 0.1^3 mm^3 KERNEL OF ^{166}Ho WAS SIMULATED WITH 10^{10} PARTICLES BECAUSE THE DECREASED KERNEL SIZE REQUIRES AN INCREASED NUMBER OF PRIMARY PARTICLES, AS DESCRIBED IN SECTION 3.2.4.	64
TABLE 3.2: THE REDUCTION OF THE NUMBER OF VOXELS INCLUDED IN THE CONVOLUTION OF THE KERNEL AS A FUNCTION OF THE USER-DEFINED THRESHOLD THAT DETERMINES IF VOXELS CONTAIN NON-NEGLIGIBLE ABSORBED DOSES.	70
TABLE 4.1: PARAMETERS USED BY THE ALGORITHMS DURING THE OPTIMISATION OF THE ALGORITHMS THROUGHOUT THIS CHAPTER. PARAMETER EXPLANATIONS ARE GIVEN IN SECTIONS 2.5.2 THROUGH 2.5.4.	78
TABLE 4.2: MINIMUM, MAXIMUM, MEAN AND STANDARD DEVIATION OF THE FINAL COST, PERCENTAGE OF ACCEPTABLE PLANS AND THE MEAN OPTIMISATION TIME FOR THE THREE ALGORITHMS. EACH ALGORITHM WAS LAUNCHED 300 TIMES WITH 20 INJECTIONS.	87

TABLE 4.3: MINIMUM, MAXIMUM, MEAN AND STANDARD DEVIATION OF THE FINAL COST, PERCENTAGE OF ACCEPTABLE PLANS AND THE MEAN OPTIMISATION TIME FOR THE THREE ALGORITHMS, EACH LAUNCHED 300 TIMES WITH BETWEEN 1 AND 20 INJECTIONS.	94
TABLE 4.4: MINIMUM, MAXIMUM, MEAN AND STANDARD DEVIATION OF THE FINAL COST, PERCENTAGE OF ACCEPTABLE PLANS AND THE MEAN OPTIMISATION TIME FOR THE THREE ALGORITHMS, EACH LAUNCHED 300 TIMES WITH BETWEEN 1 AND 20 INJECTIONS.	98
TABLE 4.5: OPTIMISATION TIMES AS A FUNCTION OF VOXEL SPACING. THE OPTIMISATIONS WERE PERFORMED OVER 1,000 GENERATIONS.....	106
TABLE 5.1: NON-EXHAUSTIVE LIST OF B ⁻ -EMITTING RADIONUCLIDES. EBMEAN FROM (BAUM, 2014). RANGE CALCULATED IN WATER, WITH INTERPOLATION FROM (ICRU, 1984). HALF-LIVES, ENERGIES AND RANGES ARE EXPRESSED IN DAYS, KEV AND MM, RESPECTIVELY..	117
TABLE 5.2: OPTIMISATION TIMES AS A FUNCTION OF THE SIZE OF THE PTV.....	128

LIST OF FIGURES

FIGURE 2.1: ITERATIVE PROCESS USED TO CREATE OPTIMAL TREATMENT PLANS. THE COST FUNCTION IS USED TO EVALUATE THE QUALITY OF THE SOLUTIONS, AND IF THE STOPPING CRITERIA ARE NOT MET, THE OPTIMISATION ALGORITHM ALTERS THE SOLUTION. THE FORWARD PROBLEM THEN DETERMINES THE OUTCOME OF THE NEWLY PROPOSED SOLUTIONS.	25
FIGURE 2.2: FICTIONAL EBRT DVH OF A PTV. ALSO DEPICTED ARE D98%PTV (50 GY), D2%PTV (60 GY) AND V55 GyPTV (88%).....	30
FIGURE 2.3: PARETO FRONT (COMPOSED OF PARETO-OPTIMAL SOLUTIONS) AND OTHER, NON-PARETO SOLUTIONS. SOLUTIONS ARE CONSIDERED PARETO-OPTIMAL AS LONG AS NO SOLUTION HAS BEEN FOUND THAT IS BETTER THAN IT WITH RESPECT TO ALL COST FUNCTIONS. H, FOR EXAMPLE, IS NOT PARETO-OPTIMAL BECAUSE D IS BETTER BOTH IN REGARDS TO COST1 AND COST2.....	36
FIGURE 2.4: TOTAL COSTS OF THE 7 PARETO-OPTIMAL TREATMENT PLANS SHOWN IN FIGURE 2.3, DISPLAYED FOR DIFFERENT RELATIVE WEIGHTINGS: (A) $w_1=0.1$, $w_2=0.9$; (B) $w_1=0.5$, $w_2=0.5$; AND (C) $w_1=0.9$, $w_2=0.1$. IT CAN BE SEEN THAT FOR AN SO ALGORITHM, THE “BEST” SOLUTION (LOWEST TOTAL COST) IS HEAVILY DEPENDENT ON THE CHOSEN COSTLET WEIGHTINGS.	36

-
- FIGURE 2.5: SCHEMATIC OF NMS. “IS $R < \text{BEST?}$ ” IS AN ABBREVIATION OF “IS $\text{COST}R < \text{COSTBEST?}$ ”. IT CAN BE SEEN THAT THE ALGORITHM TRIES TO EXPAND THE SIMPLEX USING REFLECTIONS AND EXPANSIONS, BUT IF THESE DO NOT PRODUCE GOOD QUALITY RESULTS, THE SIMPLEX SHRINKS WITH CONTRACTIONS AND REDUCTIONS. 40
- FIGURE 2.6: NMS MOVEMENTS (EXCLUDING THE REDUCTION OF ALL SOLUTIONS) FOR A 1D EXAMPLE. 41
- FIGURE 2.7: NMS MOVEMENTS FOR A 2D EXAMPLE. x_w , x_b , x_m , x_r , x_e , x_{cw} AND x_{cr} , ARE THE WORST, BEST, MEAN (FROM ALL BUT THE WORST), REFLECTION, EXPANSION, CONTRACTION OF WORST AND CONTRACTION OF REFLECTION SOLUTIONS, RESPECTIVELY. ALL OTHER SOLUTIONS THAT ARE NEITHER THE BEST NOR THE WORST ARE DENOTED WITH x 42
- FIGURE 2.8: CONTOUR MAP WITH TWO SIMPLEXES. SIMPLEX A RISKS CONVERGING TO THE LOCAL MINIMUM JUST NEXT TO IT, WHEREAS SIMPLEX B IS MUCH MORE LIKELY (BUT STILL NOT GUARANTEED) TO CONVERGE TO THE GLOBAL MINIMUM. 43
- FIGURE 2.9: SCHEMATIC FOR BA. THE SCOUT BEES ARE RANDOMLY DISTRIBUTED AND RANKED. THE TOP-RANKED SCOUT BEES BECOME THE CENTRE OF “ELITE” AND “BEST” FLOWER PATCHES, AROUND WHICH LOCAL SEARCHES ARE PERFORMED. THE BOX, TITLED “SEARCH EACH OF THE FLOWER PATCHES”, IS EXPANDED UPON IN FIGURE 2.10. 45
- FIGURE 2.10: SCHEMATIC OF A FLOWER PATCH SEARCH. THE FORAGER BEES ARE RANDOMLY LOCATED ABOUT A FLOWER PATCH CENTRE AND THEN RANKED. IF ANY OF THE FORAGERS ARE BETTER THAN THE FLOWER PATCH CENTRE, THEY REPLACE IT AS THE CENTRE AND THEN THE PROCESS IS REPEATED. IF NONE OF THE FORAGERS ARE BETTER THAN THE FLOWER PATCH CENTRE, THE SIZE OF THE FLOWER PATCH IS DECREASED AND THEN THE FORAGERS ARE REDISTRIBUTED. ONCE THE FLOWER PATCH SIZE SHRINKS BENEATH A CERTAIN LIMIT, IT IS ASSUMED THAT THE LOCAL MINIMUM HAS BEEN FOUND. THIS LOCAL MINIMUM IS COMPARED TO THE BEST MINIMUM THAT HAS BEEN FOUND SO FAR BY THE ALGORITHM (THE GLOBAL BEST). IF IT IS BETTER THAN THE GLOBAL BEST, IT REPLACES IT. THE ALGORITHM THEN MOVES ON TO SEARCHING THE NEXT FLOWER PATCH. 45
- FIGURE 2.11: RANDOM ALLOCATION OF 20 SCOUT BEES IN A 2D EXAMPLE. THE TOP RANKED SCOUT BEES BECOME THE CENTRE OF ELITE AND BEST FLOWER PATCHES. 46
- FIGURE 2.12: RANDOM DISTRIBUTION OF FORAGER BEES ABOUT FLOWER PATCHES. ELITE AND BEST FLOWER PATCHES HAVE 20 AND 10 FORAGER BEES, RESPECTIVELY (n_{re} AND n_{rb}). . 47

FIGURE 2.13: SCHEMATIC OF NSGA-II. NEW POPULATIONS ARE CREATED BY MIXING SOLUTIONS FROM THE PREVIOUS POPULATION. THE NEW AND PREVIOUS POPULATIONS ARE ADDED TOGETHER, SORTED AND THE BEST HALF ARE KEPT, WHILST THE WORST HALF ARE DISCARDED. 49

FIGURE 2.14: ITERATIVE PROCESS USED BY NSGA-II TO CREATE A NEW POPULATION (BASED ON THE PREVIOUS POPULATION) VIA BINARY TOURNAMENTS, CROSSOVERS AND MUTATIONS. BINARY TOURNAMENT SELECTION CONSISTS OF RANDOMLY SELECTING TWO SOLUTIONS FROM THE PREVIOUS POPULATION AND SELECTING THE BEST OF THE TWO TO BE A PARENT. THIS PRODUCES ONE PARENT, SO THE BINARY TOURNAMENT SELECTION IS REPEATED A SECOND TIME TO GET THE SECOND PARENT. 50

FIGURE 2.15: SOLUTIONS DIVIDED INTO THEIR CONSECUTIVE PARETO FRONTS. THE SECOND PARETO FRONT IS CONSTRUCTED BY REMOVING ALL SOLUTIONS THAT FORM THE FIRST PARETO FRONT AND FINDING ALL THE NON-DOMINATED SOLUTIONS. THESE ARE THEN REMOVED AND THE PROCESS IS REPEATED FOR THE THIRD PARETO FRONT AND SO ON..... 52

FIGURE 3.1: ILLUSTRATION OF THE ABSORBED DOSE DISTRIBUTION FOR (A) AN INJECTION, (B) THE INJECTION PLACED ON THE DOSE GRID, (C) MULTIPLE INJECTIONS CONVOLVED ON THE DOSE GRID AS DICTATED BY A TREATMENT PLAN AND (D) STRUCTURES SUPERPOSED TO CALCULATE STRUCTURE-SPECIFIC ABSORBED DOSE DISTRIBUTIONS..... 59

FIGURE 3.2: ILLUSTRATION OF THE DOSE GRID. FOUR INJECTIONS AND THEIR ASSOCIATED KERNELS ARE SHOWN AT THE EDGES OF THE PTV. IT CAN BE SEEN THAT THE BLUE BOX IS THEREFORE THE SMALLEST POSSIBLE SIZE OF THE DOSE GRID THAT IS SURE TO CONTAIN THE ENTIRETY OF THE RADIATION OF ALL PROPOSED TREATMENT PLANS..... 65

FIGURE 3.3: ILLUSTRATION OF THE RADIOLOGIC DISTANCE. THE RADIOLOGICAL DISTANCE IS THE EQUIVALENT DISTANCE TRAVERSED THROUGH WATER WITH REGARDS TO DENSITY. THE LINE CORRESPONDING TO AIR IS NOT TO SCALE; THE DENSITY OF AIR IS SO LOW THAT THE RADIOLOGICAL DISTANCE THROUGH AIR IS VERY LARGE COMPARED TO WATER AND BONE. 67

- FIGURE 3.4: 2D ILLUSTRATION OF A CUBIC KERNEL CONTAINING A SPHERICAL DISTRIBUTION. **A** IS THE MINIMUM DISTANCE FOR ALL NON-NEGLIGIBLE ABSORBED DOSES TO BE RECORDED. **B** IS THE MAXIMUM DISTANCE FROM THE INJECTION CENTRE TO THE EDGE OF THE KERNEL (2A IN 2D AND 3A IN 3D). HENCE, ALL THE VOXELS BETWEEN **A** AND **B** FROM THE INJECTION CENTRE, SHOWN BY THE AREA **C** (WHICH IS A VOLUME IN 3D), MUST CONTAIN NEGLIGIBLE ABSORBED DOSES..... 69
- FIGURE 3.5: 1D PROJECTION OF THE ABSORBED DOSE (AND S-VALUE) DISTRIBUTION FOR A 5 mL INJECTION OF ^{90}Y WITH AN A0 OF 5 MBQ. THE LIMITS OF THE 5 mL INJECTION ARE ALSO SHOWN. 72
- FIGURE 3.6: SLICES THROUGH THE ABSORBED DOSE DISTRIBUTION OF THE KERNEL OF A 5 mL INJECTION OF ^{90}Y WITH AN A0 OF 5 MBQ. SLICE DISTANCE FROM THE CENTRE OF THE INJECTION IS GIVEN BENEATH THE DISTRIBUTIONS AND ARE EXPRESSED IN MM. ALL ABSORBED DOSES GREATER THAN OR EQUAL TO 100 GY ARE SHOWN AS WHITE. 72
- FIGURE 3.7: 10 INJECTIONS RANDOMLY PLACED WITHIN THE PTV WHICH IS LOCATED NEXT TO AN OAR. THE LIMITS OF THE DOSE GRID CAN BE SEEN AS THE EDGES OF THE WHITE CUBE (A MARGIN OF 12.5 MM ABOUT THE PTV)..... 73
- FIGURE 3.8: SLICES THROUGH THE DOSE GRID. THE EDGES OF THE PTV AND OAR ARE SHOWN IN GREEN AND CYAN, RESPECTIVELY. ABSORBED DOSES GREATER THAN OR EQUAL TO 100 GY ARE SHOWN IN WHITE. 73
- FIGURE 3.9: DVH OF THE PTV AND OAR. THE MAXIMUM ABSORBED DOSE DELIVERED TO THE PTV WAS 11,934 GY, MEANING THAT THE INTERESTING RANGE OF THE DVH FROM 0 TO ~200 GY, CANNOT BE EASILY ANALYSED. FOR THIS REASON, DVHS SHOWN THROUGHOUT THIS WORK WILL BE DISPLAYED UP UNTIL A MAXIMUM OF 200 GY. 74
- FIGURE 3.10: DVH OF THE PTV AND OAR, LIMITED TO A MAXIMUM OF 200 GY. IT CAN BE SEEN THAT THE OAR REQUIREMENT IS RESPECTED BECAUSE THE OAR LINE PASSES UNDERNEATH THE BLACK CIRCLE. CONVERSELY, THE PTV REQUIREMENT IS NOT RESPECTED BECAUSE THE PTV LINE DOES NOT PASS OVER THE BLACK DIAMOND..... 75
- FIGURE 4.1: COST OF NMS AS A FUNCTION OF THE NUMBER OF UPDATES. DURING EACH GENERATION, THE WORST SOLUTION IS ALTERED. A REDUCTION (IN WHICH ALL SOLUTIONS MOVE TOWARDS THE BEST) CAN BE SEEN AT 38 UPDATES. REDUCTIONS DO NOT GUARANTEE LOWERING THE COST FUNCTION OF ALL SOLUTIONS, AND THIS IS WHY SOME OF THE LINES CAN BE SEEN TO INCREASE AT THIS POINT. 81

FIGURE 4.2: EVOLUTION OF THE COST FOR NMS FOR THE BEST AND WORST TREATMENT PLANS FOR AN OPTIMISATION OF 20 INJECTIONS. IT CAN BE SEEN THAT, BECAUSE NMS MOSTLY IMPROVES THE WORST SOLUTION, THE BEST SOLUTION CHANGES LITTLE OVER THE COURSE OF THE OPTIMISATION..... 82

FIGURE 4.3: EVOLUTION OF THE COST AS A FUNCTION OF THE NUMBER OF TREATMENT PLAN UPDATES FOR “ELITE” AND “BEST” FLOWER PATCHES WITH BA. THE START OF EACH LINE IS THE COST AT THE INITIAL CENTRE OF THE FLOWER PATCH (I.E., THE COST OF THE SCOUT THAT WAS USED TO CREATE THE PATCH). AS BETTER SOLUTIONS ARE FOUND BY FORAGER BEES IN THE SURROUNDING SEARCH-SPACE, THEY BECOME THE FLOWER PATCH CENTRE AND THE COST DECREASES ACCORDINGLY. AT THE END OF EACH LINE, THE FLOWER PATCH HAS PASSED BENEATH A USER-DEFINED MINIMUM SIZE AND IT IS ASSUMED THAT THE LOCAL MINIMUM HAS BEEN FOUND. THE BEST LINE CORRESPONDS TO THE BEST SOLUTION FOUND IN ALL FLOWER PATCHES UP UNTIL THAT POINT..... 83

FIGURE 4.4: EVOLUTION OF THE COST AS A FUNCTION OF GENERATION (I.E., ITERATION) FOR A POPULATION OF 10 TREATMENT PLANS IN NSGA-II. SOLUTIONS THAT WERE CREATED BY CROSSEVERS ARE JOINED TO THEIR PARENTS VIA SOLID BLUE LINES, WHEREAS SOLUTIONS SURVIVING FROM PREVIOUS GENERATIONS ARE SHOWN VIA RED DASHED LINES. 85

FIGURE 4.5: COST AS A FUNCTION OF TREATMENT PLAN UPDATES FOR NSGA-II. 86

FIGURE 4.6: FREQUENCY HISTOGRAM FOR THE FINAL COST OF THE THREE ALGORITHMS, EACH LAUNCHED 300 TIMES WITH 20 INJECTIONS. 87

FIGURE 4.7: THE EVOLUTION OF THE TOTAL COST FOR EACH NUMBER OF INJECTIONS THAT WERE OPTIMISED INDIVIDUALLY. IT CAN BE SEEN THAT ONLY BY USING 20 INJECTIONS COULD A TREATMENT PLAN BE FOUND THAT SUFFICIENTLY IRRADIATES THE PTV (AS THIS IS THE ONLY SOLID LINE AMONGST ALL THE DASHED LINES). 89

FIGURE 4.8: EVOLUTION OF BA FOR AN OPTIMISATION WITH A VARIABLE NUMBER OF INJECTIONS. IT CAN BE SEEN THAT AT 626 UPDATES, EFP2 FIND A SOLUTION WITH A LOWER TOTAL COST THAN THE BEST SOLUTION. HOWEVER, THE BEST SOLUTION IS ACCEPTABLE, WHEREAS THE ONE PROPOSED BY EFP2 IS NOT, AND SO THIS SOLUTION DOES NOT BECOME THE BEST. IT IS ONLY AT 668 UPDATES THAT THE SOLUTION PROPOSED BY EFP2 IS BOTH ACCEPTABLE AND LOWER THAN THE BEST, AND AT THIS POINT IT REPLACES THE BEST SOLUTION. 90

FIGURE 4.9: EVOLUTION OF THE PARETO FRONT AS A FUNCTION OF THE NUMBER OF GENERATIONS. COSTINJ IS SHOWN ALONG THE BOTTOM, WHEREAS THE NUMBER OF INJECTIONS ARE SHOWN ALONG THE TOP. THE BEST SO PLAN (LOWEST TOTAL COST WHILST RESPECTING THE TREATMENT CRITERIA) FOR EACH GENERATION IS MARKED WITH A BLACK CROSS.	91
FIGURE 4.10: EVOLUTION OF THE TOTAL COST AS A FUNCTION OF THE NUMBER OF TREATMENT PLAN UPDATES. CLEARLY, WHEN THE USER IS CONCERNED BY THE TOTAL COST (I.E., THE SO COST), IT IS BETTER TO USE NSGA-II AS AN SO ALGORITHM. IF NSGA-II IS USED IN MO MODE, IT CONCENTRATES ITS RESOURCES ON EXPLORING THE ENTIRETY PARETO FRONT, WHICH IS NOT ADVANTAGEOUS WHEN THE USER IS ONLY CONCERNED BY ONE SPECIFIC POINT OF THE PARETO FRONT.	92
FIGURE 4.11: EXAMPLE OF INCREASING DIVERSITY IN NSGA-II DUE TO MUTATIONS AND THE CROWDING DISTANCE. ALL SOLUTIONS WERE INITIALISED IDENTICALLY WITH 10 INJECTIONS, ALL WITH MATCHING COORDINATES. IT CAN BE SEEN THAT, THANKS TO MUTATIONS AND THE CROWDING DISTANCE, THE SOLUTIONS RAPIDLY SPREAD OUT. AFTER JUST 20 GENERATIONS, SOLUTIONS HAVE BEEN PROPOSED THAT COVER THE FULL RANGE OF ALLOWED NUMBERS OF INJECTIONS (1 TO 20).	93
FIGURE 4.12: FREQUENCY HISTOGRAM FOR THE FINAL COST OF THE THREE ALGORITHMS, EACH LAUNCHED 300 TIMES WITH BETWEEN 1 AND 20 INJECTIONS.	94
FIGURE 4.13: 3D PARETO FRONT FOR NSGA-II. THE DATA POINTS ARE COLOURED AS A FUNCTION OF THE NUMBER OF INJECTIONS USED BY THAT PLAN. A COSTINJ OF 70% CORRESPONDS TO 14 INJECTIONS, WHEREAS A COSTINJ OF 100% CORRESPONDS TO 20 INJECTIONS, AS DESCRIBED IN SECTION 2.4.2. THE PLAN WITH THE LOWEST TOTAL COST IS HIGHLIGHTED AS “BEST”.	96
FIGURE 4.14: FREQUENCY HISTOGRAM FOR THE FINAL COST OF THE THREE ALGORITHMS, EACH LAUNCHED 300 TIMES WITH BETWEEN 1 AND 20 INJECTIONS.	97
FIGURE 4.15: 2D PROJECTIONS OF THE 4D PARETO FRONT. FROM THE PROJECTIONS, IT IS POSSIBLE TO ASCERTAIN CORRELATIONS (SUCH AS INCREASING COSTINJ DECREASES COSTPTV), BUT OTHER RELATIONSHIPS ARE HARDER TO ANALYSE (SUCH AS COSTOAR AND COSTCI). FOR THIS REASON, IT WOULD BE BETTER TO DISPLAY ALL PARETO-OPTIMAL RESULTS IN THE FORM OF A TABLE.	100

List of Figures

- FIGURE 4.16: CENTRAL SLICE OF ABSORBED DOSE DISTRIBUTIONS FOR 5 mL INJECTIONS OF ^{166}Ho WITH A_0 'S OF 5 MBQ AND VOXEL SIZES OF (FROM LEFT TO RIGHT) 0.1^3 , 0.5^3 AND 1.0^3 MM^3 104
- FIGURE 4.17: PARETO FRONT THE OPTIMISATION OF THE PTV (10 MM RADIUS) WITH INJECTIONS OF ^{166}Ho FOR KERNELS WITH VOXEL SIZES OF 0.1^3 , 0.5^3 AND 1.0^3 MM^3 105
- FIGURE 4.18: DVHS FOR THE SAME PLAN (THE FIRST THAT PASSES UNDER 5% ON THE Y-AXIS FOR THE VOXEL SIZE OF 0.1^3 MM^3 IN THE PREVIOUS FIGURE) WITH DIFFERENT VOXEL SIZES. . 106
- FIGURE 4.19: DIFFERENCES IN 0.5^3 AND 1.0^3 MM^3 VOXEL SIZE DVHS COMPARED TO THE DVH OF THE 0.1^3 MM^3 VOXEL SIZE. 107
- FIGURE 4.20: 3D VISUALISATION OF THE SKULL OF A MINI-PIG AS SEEN FROM THREE ANGLES. THE GLIOBLASTOMA HAS BEEN SEGMENTED AND IS SHOWN IN PINK. THE LIMITS OF THE DOSE GRID ARE ALSO SHOWN. 110
- FIGURE 4.21: PARETO FRONT FOR THE TREATMENT PLANNING OF A REAL TUMOUR. THE TREATMENT PLAN THAT WAS DEEMED THE BEST ACCORDING THE USER (BECAUSE IT USES THE FEWEST INJECTIONS AMONGST ALL THE ACCEPTABLE PLANS) IS SHOWN WITH THE VERTICAL CROSS, WHEREAS THE BEST SO PLAN IS SHOWN WITH THE DIAGONAL CROSS.. 111
- FIGURE 4.22: 3D DISTRIBUTION OF INJECTIONS THROUGHOUT THE PTV AS SEEN FROM THREE ANGLES. 112
- FIGURE 4.23: SLICES THROUGH THE ABSORBED DOSE DISTRIBUTION. THE CONTOURS OF THE SKULL ARE SHOWN IN WHITE AND THE CONTOURS OF THE PTV ARE SHOWN IN RED. SLICES ARE SHOWN EVERY 2 SLICES (1 MM). 112
- FIGURE 4.24: DVH FOR THE PTV OF THE SELECTED PLAN. THE PTV'S ABSORBED DOSE REQUIREMENT IS ALSO SHOWN WITH THE BLACK DIAMOND. IT CAN BE SEEN THAT THE PLAN CORRESPONDING TO THIS DVH IS ACCEPTABLE BECAUSE IT PASSES OVER THE BLACK DIAMOND. 113
- FIGURE 5.1: PROBABILITY SPECTRA OF B^- -EMISSIONS FOR ^{177}Lu , ^{131}I , ^{166}Ho AND ^{90}Y . DATA FROM MIRD (ECKERMAN AND ENDO, 1989). 118
- FIGURE 5.2: ABSORBED DOSE AS A FUNCTION OF DISTANCE FOR SINGLE 5 mL INJECTIONS WITH A_0 'S OF 5 MBQ FOR THE RADIOELEMENTS ^{177}Lu , ^{131}I , ^{166}Ho AND ^{90}Y . THE PHYSICAL LIMIT OF THE INJECTION (A SPHERE OF 5 mL HAS A RADIUS OF 1.06 MM) IS ALSO SHOWN. 119

FIGURE 5.3: PARETO FRONT FOR OPTIMISATION OF THE PTV WITH 5 mL INJECTIONS OF ^{177}Lu , ^{131}I , ^{166}Ho AND ^{90}Y EACH WITH AN A_0 OF 20 MBQ.	120
FIGURE 5.4: DVHS FOR THE PLANS THAT USE THE FEWEST INJECTIONS WHILST SUFFICIENTLY IRRADIATING THE PTV ($\text{CostPTV} \leq 5\%$). SINCE IT WAS NOT POSSIBLE TO ACHIEVE THIS WITH THE MAXIMUM NUMBER OF INJECTIONS WITH ^{177}Lu , THE DVH FOR 100 INJECTIONS IS SHOWN. FOR THIS REASON, ^{131}I , ^{166}Ho AND ^{90}Y PASS OVER THE DIAMOND AT (95 GY, 95%), WHEREAS ^{177}Lu PASSES UNDERNEATH.....	121
FIGURE 5.5: 1D ABSORBED DOSE DISTRIBUTION FOR A 5 mL INJECTION OF ^{90}Y WITH VARIOUS VALUES OF A_0 . THE PHYSICAL LIMIT OF THE INJECTION (A SPHERE OF 5 mL HAS A RADIUS OF 1.06 MM) IS ALSO SHOWN.	122
FIGURE 5.6: 1D ABSORBED DOSE DISTRIBUTION FOR 5 mL INJECTIONS OF ^{90}Y WITH VARIOUS VALUES OF A_0 . THE Y-AXIS OF THIS FIGURE IS LIMITED TO A MAXIMUM OF 500 GY. THE PHYSICAL LIMIT OF THE INJECTION (A SPHERE OF 5 mL HAS A RADIUS OF 1.06 MM) IS ALSO SHOWN.	123
FIGURE 5.7: PARETO FRONT AS A FUNCTION OF A_0 . CLEARLY INCREASING A_0 REDUCES THE NUMBER OF INJECTIONS THAT ARE REQUIRED TO SUFFICIENTLY IRRADIATE THE PTV, SHIFTING THE PARETO FRONT TOWARDS THE BOTTOM-LEFT OF THE GRAPH.	124
FIGURE 5.8: 1D ABSORBED DOSE PROFILES FOR SINGLE ^{90}Y INJECTIONS OF VARIOUS VOLUMES AND A_0 OF 5 MBQ. THE INJECTION RADII (1.06, 1.34 AND 1.68 MM) ARE SHOWN FOR EACH OF THE VOLUMES.	125
FIGURE 5.9: PARETO FRONT FOR THE OPTIMISATION OF THE PTV WITH BETWEEN 1 AND 20 INJECTIONS OF 5 MBQ ^{90}Y . THE VOLUME OF THESE INJECTIONS VARIED BETWEEN 5 AND 20 mL. IT CAN BE SEEN THAT VARYING THE INJECTION VOLUME HAS A NEGLIGIBLE EFFECT ON THE PARETO FRONT.....	126
FIGURE 5.10: THE PARETO FRONT FOR THE OPTIMISATION OF SPHERICAL PTVS WITH RADII RANGING BETWEEN 10 AND 30 MM. NSGA-II WAS USED WITH 20 mL INJECTIONS OF ^{90}Y WITH A_0 'S OF 20 MBQ. IT CAN BE SEEN THAT TREATMENT PLANS CAN BE FOUND FOR PTVS SMALLER THAN 20 MM. HOWEVER, 30 MM IS TOO VOLUMINOUS TO BE TREATED, EVEN WITH THE MAXIMUM NUMBER OF INJECTIONS, 200.	128
FIGURE 5.11: 3D VISUALISATION OF ALL POSSIBLE INJECTOR TRAJECTORIES (BLUE) WITHIN A SPHERICAL PTV OF RADIUS 10 MM. THE GUIDE NEEDLE IS SHOWN IN GREEN.	130

FIGURE 5.12: S-VALUES THROUGH THE CENTRAL AXIS OF A CYLINDRICAL INJECTION, WITH CROSS-SECTIONAL AREA OF 0.68 mm^2 AND LENGTH OF 20 MM (AS SHOWN BY THE BLACK SOLID LINES). IT CAN BE SEEN THAT THE KERNEL PLATEAUS $\sim 1.5 \text{ mm}$ AFTER THE START AND BEFORE THE END OF THE INJECTION, AS SHOWN BY THE DASHED LINES. THEREFORE, SHORTER INJECTIONS CAN BE MODELLED BY REMOVING PARTS OF THE KERNEL THAT LIE IN THE PLATEAUED AREA. 132

FIGURE 5.13: ILLUSTRATIONS OF THE ABSORBED DOSE PROFILE OF (A) A SPHERICAL INJECTION; (B) A CYLINDRICAL INJECTION, CENTRED ON VOXEL CENTRES; AND (C) THE ABSORBED DOSE DISTRIBUTION REQUIRED FOR A DIAGONALLY PLACED CYLINDRICAL INJECTION. 133

FIGURE 5.14: ABSORBED DOSE DISTRIBUTION FOR A ROTATED AND RESAMPLED CYLINDRICAL INJECTION, BOTH SHOWN TO ITS MAXIMUM ABSORBED DOSE AND CAPPED TO 100 Gy. ... 134

FIGURE 5.15: PARETO FRONTS FOR OPTIMISATIONS OF THE PTV WITH INJECTIONS OF ^{166}Ho IN THE FORM OF SPHERES AND CYLINDERS. IT CAN BE SEEN THAT FOR FEW INJECTIONS (FOR EXAMPLE, $\text{CostInj} = 10$), USING CYLINDRICAL INJECTIONS IS BETTER AT REDUCING THE UNDER-IRRADIATED PORTION OF THE PTV (SMALLER CostPTV). HOWEVER, FOR HIGH VALUES OF CostInj , CostPTV REACHES AN ASYMPTOTE FOR CYLINDRICAL INJECTIONS. THIS IS BECAUSE THERE ARE PORTIONS OF THE PTV THAT ARE OUT OF REACH OF THE CYLINDRICAL INJECTIONS (FOR EXAMPLE, THE UPPER-EXTREMITIES OF FIGURE 5.11). ... 135

FIGURE 6.1: ILLUSTRATION OF THE OBJECT HIERARCHY OF DosiTOUL. 158

FIGURE 6.2: DATABASE WINDOW. AT THE TOP, A LIST OF PATIENTS IS SHOWN. WHEN A PATIENT IS SELECTED, ALL STUDIES FOR THAT PATIENT ARE SHOWN IN THE SECOND LIST. WHEN A STUDY IS SELECTED, ALL SERIES WITHIN THAT STUDY ARE SHOWN IN THE BOTTOM BOX. USING THE BUTTONS ALONG THE BOTTOM, PATIENTS, STUDIES AND SERIES CAN BE ADDED; STRUCTURES CAN BE ADDED TO SERIES; AND OPTIMISATIONS CAN BE PERFORMED. 160

FIGURE 6.3: EDIT STRUCTURES WINDOW. ON THE LEFT, A 3D VISUALISATION OF ALL STRUCTURES IS SHOWN. ON THE RIGHT, THE CORRESPONDING TABLE IS DISPLAYED. STRUCTURES CAN BE ADDED, DELETED AND ALTERED USING THE BUTTONS ALONG THE BOTTOM. 161

FIGURE 6.4: OPTIMISATION BROWSER. OPTIMISATIONS CAN BE CREATED AND DELETED USING “NEW” AND “DELETE”. THEY CAN BE PARAMETERISED USING “EDIT”. THE OPTIMISATION CAN BE EXECUTED USING “START” AND THEN THE RESULTS CAN BE VIEWED WITH “VIEW RESULTS”. 163

FIGURE 6.5: ALGORITHM SETTINGS. THE TOP BOX REQUIRES INFORMATION THAT IS COMMON TO ALL ALGORITHMS (FOR EXAMPLE, THE TIME LIMIT), WHEREAS THE BOTTOM HALF REQUIRES INFORMATION THAT IS SPECIFIC TO A GIVEN ALGORITHM (FOR EXAMPLE, THE POPULATION SIZE FOR NSGA-II).	164
FIGURE 6.6: KERNEL SETTINGS. INFORMATION, SUCH AS THE PATH OF THE GATE SIMULATION AND THE RADIOELEMENT'S HALF-LIFE, IS REQUIRED HERE.	165
FIGURE 6.7: COSTLET SETTINGS. A LIST OF ALL COSTLETS IS SHOWN AND NEW COSTLETS CAN BE ADDED WITH "ADD". CONVERSELY, THEY CAN BE REMOVED WITH "DELETE".....	166
FIGURE 6.8: SETTING UP INDIVIDUAL COSTLETS. HERE, THE PARAMETERS FOR A DOSE-VOLUME COSTLET ARE SHOWN. THE USER CHOOSES THE STRUCTURE TO BE PENALISED, THE INEQUALITY DIRECTION (" $<$ ", " \leq ", " $>$ " OR " \geq ") AND THE REFERENCE ABSORBED DOSE. UPON CLICKING "OK", THE COSTLET IS ADDED TO THE LIST.	167
FIGURE 6.9: INJECTION PARAMETERS. ON THE LEFT, ALL POSSIBLE INJECTIONS ARE SHOWN RELATIVE TO ALL THE STRUCTURES (IN THIS CASE THE SKULL AND THE PTV). ON THE RIGHT, ALL PARAMETERS REQUIRED FOR THE INJECTIONS ARE ENTERED. ONCE ALL PARAMETERS HAVE BEEN ENTERED, ALL POSSIBLE INJECTIONS ARE COMPUTED USING "CALCULATE".	167
FIGURE 6.10: OPTIMISATION PROGRESS. HERE, A REAL-TIME VISUALISATION OF THE PARETO FRONT IS SHOWN. THE USER MAY CANCEL THE OPTIMISATION WITH "STOP" OR CLICK "DONE" ONCE THE OPTIMISATION IS COMPLETE.	169
FIGURE 6.11: OPTIMISATION RESULTS PAGE. AT THE TOP-LEFT, THE FINAL PARETO FRONT IS SHOWN, AT THE BOTTOM-LEFT THE PARETO-OPTIMAL SOLUTIONS ARE SHOWN IN THE TABLE. ONCE A SOLUTION HAS BEEN SELECTED, THE RIGHT-HAND SIDE UPDATES AUTOMATICALLY, SHOWING SLICES OF THE ABSORBED DOSE DISTRIBUTION (TOP-LEFT), THE 3D VISUALISATION OF THE INJECTIONS (TOP-RIGHT), THE DVH (MIDDLE) AND THE COORDINATES OF THE INJECTIONS (BOTTOM). THE SELECTED TREATMENT PLAN CAN BE SAVED USING "SAVE PLAN TO DISK".....	170
FIGURE 6.12: EXAMPLE OF A VTK WIDGET EMBEDDED IN A QT WINDOW VIA A "QVTKWIDGET".	173

LIST OF ABBREVIATIONS AND ACRONYMS

1D	one-dimension, one-dimensional
2D	two-dimension, two-dimensional
3D	three-dimension, three-dimensional
4D	four-dimension, four-dimensional
BA	Bees Algorithm
CAD	Computer-Aided Design
CI	Conformity Index
COIN	COncormal INdex
CPU	Central Processing Unit
CT	Computed Tomography
CTV	Clinical Target Volume
DICOM	Digital Imaging and COmmunications in Medicine
DPK	Dose-Point Kernel
DVH	Dose-Volume Histogram
DVK	Dose-Volume Kernel
EBRT	External Beam Radiation Therapy
EUD	Equivalent Uniform Dose
GATE	Geant4 Application for Tomographic Emission
GDCM	Grassroots DICOM
GPU	Graphics Processing Unit
GTV	Gross Tumour Volume
Gy	Gray
HI	Homogeneity Index
HU	Hounsfield Unit
MIRD	Medical Internal Radiation Dose Committee
MO	Multi-Objective

MRI	Magnetic Resonance Imaging
MRT	Molecular Radiation Therapy
ND	N-dimension, N-dimensional
NMS	Nelder-Mead Simplex
NSGA-II	Non-dominated Sorting Genetic Algorithm-II
NTCP	Normal Tissue Complication Probability
OAR	Organ At Risk
PTV	Planning Target Volume
SIRT	Selective Internal Radiation Therapy
SO	Single-Objective
STL	Standard Triangle Language
TCP	Tumour Control Probability
TPS	Treatment Planning System
TRT	Targeted Radiation Therapy
VOI	Volume Of Interest
VTK	Visualisation ToolKit
WHO	World Health Organisation

RESUME (FR)

Une nouvelle modalité de radiothérapie, la micro-curiethérapie, est en cours de développement. Cette thérapie cible des tumeurs solides inopérables en effectuant des injections de liquide contenant des microsphères radioactives en suspension.

Plusieurs injections sont nécessaires pour suffisamment irradier la zone tumorale et donc, afin d'optimiser le positionnement de ces injections, une méthode de planification de traitement a été développée et validée au cours de cette thèse.

Tout au long de ce travail, trois thèmes principaux seront discutés :

- Comment réaliser une dosimétrie adaptée à cette micro-curiethérapie ?
- Comment effectuer la planification de traitement pour cette modalité ?
- Comment optimiser le plan de traitement afin qu'il soit le plus efficace possible ?

La première question est abordée en discutant comment la distribution de radiation d'une seule injection pourrait être calculée, et comment cette distribution pourrait être convoluée à une distribution 3D réelle dans le patient. Ainsi, les injections proposées par le plan de traitement pourraient être convoluées par la distribution de dose absorbée d'une injection, afin de calculer la dose absorbée dans le patient.

Une fois la distribution de dose absorbée calculée pour le plan de traitement, la prochaine étape concerne l'extraction des informations dosimétriques, tel que les histogrammes dose-volume (HDVs). Ceci permet d'analyser et de comparer rapidement et efficacement les plans de traitement.

Ensuite, les algorithmes d'optimisation sont développés et comparés. Ces algorithmes cherchent des plans de traitement de qualité croissante, en espérant qu'à la fin du processus d'optimisation, le plan de traitement proposé soit de qualité adéquate pour être effectué sur un patient.

Afin de pouvoir produire des plans de traitement de qualité croissante, il est important de pouvoir les classer en les comparant quantitativement. Ces plans de traitement étant assez abstraits – consistant en plusieurs injections distribuées dans un espace 3D et la distribution de dose absorbée résultante – il est nécessaire de les condenser en quelques chiffres qui pourraient être comparés facilement. De plus, ce processus nécessite d'être rapide car cette analyse de plans de traitement doit être réalisée plusieurs fois pendant la planification de traitement, dans un temps cliniquement acceptable. Dans ce contexte, des méthodes de comparaison quantitatives via des « fonctions coûts » sont discutées.

Une fois les méthodologies de dosimétrie pour la micro-curiethérapie et de production de plans de traitement optimaux développées, certains paramètres d'injection sont testés et comparés. Notamment, une étude importante est effectuée sur le radioélément le mieux-adapté aux injections. D'autres études sont effectuées sur le volume injecté idéal, l'activité initiale du radioélément injecté et, sur la forme géométrique des injections (une forme sphérique ou cylindrique).

Avec la dosimétrie, l'algorithme d'optimisation et les paramètres d'injection en place, le développement d'un système de planification de traitement (TPS) complètement opérationnel avec une interface graphique est discuté et les résultats présentés.

Cette nouvelle technologie associée aux développements réalisés au cours de cette thèse démontre la faisabilité, validée sur animal, de pouvoir injecter un liquide contenant des microsphères radioactives en suspension afin de pouvoir traiter efficacement, tout en préservant les tissus sains environnants, dans le cas de tumeurs inopérables aujourd'hui.

La structure de ce travail est donc :

- Chapitre 1 : état de l'art de ce projet. Qu'est-ce que la micro-curiethérapie ? Pourquoi cette thérapie a-t-elle été proposée ? Quelles tumeurs pourraient potentiellement être ciblées ? Qu'est-ce qui est nécessaire pour que cette thérapie soit un succès ?
- Chapitre 2 : revue de bibliographie, introduisant le concept de planification de traitement. Dans ce chapitre, plusieurs concepts, tel que la fonction coût, sont présentés ainsi que l'explication de plusieurs algorithmes d'optimisation qui seront utilisés tout au long de cette recherche.
- Chapitre 3 : méthode de calcul de la dosimétrie de plans de traitement proposés. Ce chapitre se termine avec un exemple travaillé, du début à la fin, de l'analyse dosimétrique d'un plan de traitement.
- Chapitre 4 : comparaison des algorithmes d'optimisation qui ont été introduits dans le Chapitre 2 en utilisant la méthodologie de dosimétrie qui a été présentée dans le Chapitre 3. Plusieurs cas de test sont utilisés pour cette comparaison, chacun ajoutant un niveau de complexité sur le précédent. Une fois les algorithmes comparés et le mieux adapté sélectionné, une étude est effectuée sur l'impact de la taille du voxel sur la dosimétrie ainsi que sur la planification de traitement. Finalement, un plan de traitement est optimisé et analysé pour un vrai glioblastome.
- Chapitre 5 : comparaison des paramètres d'injection qui pourraient être modifiées afin de délivrer le meilleur traitement possible. Ces paramètres incluent le choix du radioélément injecté, le volume de chaque injection et l'activité initiale injectée. Il est aussi proposé dans ce chapitre que les injections cylindriques soient utilisées à la place d'injections sphériques. Les effets sur le calcul de la distribution de dose absorbée et sur la planification de traitement sont alors discutés.
- Chapitre 6 : les méthodologies de tous les chapitres précédents sont rassemblées afin de créer un TPS de micro-curiethérapie. La structure de ce TPS est présentée. L'utilisation des logiciels tiers est aussi discutée dans ce chapitre.
- Ce travail se termine par un chapitre sur les conclusions et les perspectives futures.

INTRODUCTION (FR)

Le cancer et la radiothérapie

Le cancer est actuellement la plus grande cause de mort en France. Les outils principaux dans la lutte contre le cancer sont la chirurgie, la chimiothérapie et la radiothérapie. Ce travail se concentre uniquement sur la radiothérapie, qui pourrait être utilisée, soit en thérapie adjuvante, soit toute seule.

Quand les rayonnements ionisants traversent le corps, de l'énergie est déposée. Indépendamment du fait que l'énergie soit déposée dans les cellules tumorales ou saines, ceci cause la mort cellulaire. La radiothérapie, par conséquent, exploite les rayonnements ionisants en les focalisant sur la tumeur.

Il existe plusieurs formes de radiothérapie, mais l'objectif commun parmi toutes ces modalités reste le même : irradier suffisamment la tumeur tout en réduisant l'irradiation des tissus sains environnants. La sur-irradiation des tissus sains peut induire des effets secondaires à la fois aigus et chroniques, tandis que la sous-irradiation de la tumeur peut ne pas entraîner sa destruction complète voir le développement de tumeurs secondaires.

Les modalités de radiothérapie

Avant de décrire la micro-curiethérapie, une nouvelle forme de radiothérapie, il est nécessaire de discuter trois modalités existantes de radiothérapie et leurs divers avantages et inconvénients. Ces modalités sont : la radiothérapie externe (EBRT), la curiethérapie et la médecine nucléaire, avec un accent particulier sur un type de médecine nucléaire appelé la radioembolisation (SIRT). Une fois les avantages et inconvénients de ces modalités considérés, les raisons du développement de la micro-curiethérapie sont discutées.

La radiothérapie externe

L'EBRT est la modalité la plus répandue de radiothérapie. Celle-ci consiste à utiliser un accélérateur d'électrons qui produit des rayonnements ionisants de haute énergie (des rayons X ou des électrons). La rotation de l'accélérateur autour du patient permet de créer un point focal (ou isocentre) de l'irradiation centré sur la tumeur.

Cependant, l'inconvénient principal de l'EBRT est que le rayonnement ionisant doit traverser tout le patient pour atteindre la tumeur, souvent localisée proche de tissus sains et vulnérables.

Les effets nocifs de la radiation sur des tissus sains sont minimisés :

- En fractionnant le traitement (par exemple, en délivrant peu d'irradiation tous les jours pendant plusieurs semaines), afin que les tissus sains aient le temps de récupérer (car ils récupèrent plus vite que les cellules tumorales) ;
- En tournant l'accélérateur de particules autour du patient, afin que l'irradiation se propage tous les tissus voisins au lieu d'être localisée sur un volume ; et
- En utilisant les collimateurs multilames (MLC) qui adaptent la forme du faisceau, en fonction de l'angle du bras de l'accélérateur, afin de mieux le faire correspondre à la forme de la tumeur.

Malgré la minimisation de l'exposition des tissus sains, d'autres modalités telles que la curiethérapie et la médecine nucléaire ont été développées avec l'objectif de minimiser encore plus l'irradiation des tissus sains (dans certains cas particuliers de tumeurs).

La curiethérapie

La curiethérapie consiste à insérer des matériaux solides radioactifs (par exemple, des grains de métal) dans et autour de la tumeur. L'avantage clé de la curiethérapie est que, en introduisant la source radioactive beaucoup plus proche de la tumeur, d'autres types de radiation pourraient être utilisés. Des sources pourraient être utilisées pour déposer leurs énergies sur des distances beaucoup plus courtes, limitant l'irradiation des tissus sains environnant tout en augmentant si nécessaire la dose absorbée dans la zone tumorale.

Au contraire, l'inconvénient de la curiethérapie est que, parce que les sources doivent être physiquement placées dans la localisation souhaitée, son utilisation est donc limitée aux traitements de tumeurs qui sont facilement accessibles (par exemple, proche de la surface du patient), telles que les tumeurs de la peau, de la prostate et du col de l'utérus.

La médecine nucléaire

La médecine nucléaire peut être divisée en deux sous-catégories, la radiothérapie interne vectorisée (RIV) et la SIRT :

- La RIV consiste à introduire des produits radiopharmaceutiques (i.e., les radioéléments sont attachés aux produits pharmaceutiques) par voie intraveineuse au patient. Les produits pharmaceutiques s'attachent préférentiellement aux récepteurs sur les cellules tumorales ciblées d'où ils émettent leur radiation, en détruisant les cellules avoisinantes.
- La radioembolisation est utilisée pour les tumeurs du foie inopérables. Ceci consiste à injecter des microsphères radioactives dans les artères nourrissant le foie afin que les radioéléments s'accumulent dans la tumeur, émettant leur énergie au voisinage et la détruisent.

L'avantage de la SIRT est qu'elle est beaucoup moins invasive que d'autres modalités de radiothérapie, par exemple la curiethérapie, en raison du fait que les sources radioactives sont introduites par voie intraveineuse plutôt que d'être injectées directement.

L'inconvénient évident de la SIRT est qu'elle manque de précision. Les microsphères ne sont pas introduites directement dans la tumeur, ni attachées aux produits pharmaceutiques qui eux s'accrochent préférentiellement aux récepteurs des cellules tumorales, comme pour la RIV.

La micro-curiethérapie

La recherche continue avec des modalités de radiothérapie plus performantes qui augmentent l'irradiation de la tumeur tout en diminuant l'irradiation des tissus sains. En vue de cet objectif, l'idée de réaliser de traitements par micro-curiethérapie a été proposée.

La micro-curiethérapie peut être considérée comme un hybride entre la curiethérapie et la SIRT. Elle est un hybride entre les deux car les microsphères radioactives (comme pour la SIRT) sont injectées directement dans la tumeur (comme pour la curiethérapie).

En passant de « grains » en curiethérapie aux « microsphères » en micro-curiethérapie, le volume de chaque unité de traitement est réduit. Ceci devrait aider à mieux distribuer la dose absorbée autour de la tumeur, à la fois en augmentant l'irradiation de la tumeur et en diminuant l'irradiation des tissus sains. De plus, en réalisant les injections directement dans la tumeur, il est espéré que l'inconvénient du manque de précision de la SIRT puisse être contourné.

En outre, en utilisant une unité de traitement sous forme liquide (les microsphères sont dans une suspension liquide), les cathéters invasifs qui sont utilisés pour implanter les grains en curiethérapie pourraient être remplacés par une aiguille moins invasive, pour que les tumeurs plus profondes puissent potentiellement être ciblées.

Vu que la distribution de radiation devrait, en théorie, être mieux répartie autour de la tumeur, l'irradiation de cette dernière pourrait donc être augmentée sans endommager les tissus sains environnants. La micro-curiethérapie pourrait donc être un outil clé dans le traitement de cancers radioresistants inopérables.

Le traitement

En micro-curiethérapie, des injections des volumes entre 5 et 20 μL sont utilisées. Chaque injection contient les microsphères suspendues de manière quasi-homogène en solution. Il est supposé que, une fois injecté dans la tumeur, le liquide forme une sphère uniforme autour de l'aiguille, de sorte qu'il y ait une distribution uniforme de microsphères dans la sphère injectée.

Une fois injectés dans la tumeur, il est nécessaire que les radioéléments restent dans leur position et qu'ils émettent l'intégralité de leur énergie. C'est pour cela que les microsphères sont utilisées ; elles sont suffisamment grandes pour que leurs demi-vies biologiques soient quasi-infinies, ce qui fut démontrée expérimentalement (Bult *et al.*, 2012).

La planification de traitement

De toutes les modalités de radiothérapie déjà discutées précédemment, la planification de traitement est utilisée pour tous. La planification est la recherche du meilleur plan de traitement possible qui pourrait être délivré à chaque patient.

Les algorithmes de résolution de problèmes inverses sont souvent exploités afin de trouver les plans de traitement optimaux. Ils modifient itérativement les plans de traitement afin de minimiser des fonctions de pénalisation déterminées par l'utilisateur. Ces fonctions, appelées des « fonctions coûts », pénalisent des caractéristiques non-souhaitées du plan de traitement (par exemple, la sous-irradiation de la tumeur ou bien la sur-irradiation d'un organe sain avoisinant). Par conséquent, lorsque l'algorithme d'optimisation minimise les fonctions coûts, la qualité du plan de traitement augmente.

À la fin de l'optimisation, le plan de traitement proposé peut être validé et éventuellement effectué sur le patient.

Comme il s'agit d'une nouvelle modalité de radiothérapie, une méthode de planification de traitement n'existe pas encore pour la micro-curiethérapie. Une des thématiques clé de ce travail est donc le développement et la validation d'un système de planification de traitement spécifique à la micro-curiethérapie.

Le choix de radioéléments

La curiethérapie va inévitablement exploiter des radioéléments émetteurs β^- . En effet, de ces radioéléments déposent leur énergie dans les quelques millimètres autour du point d'injection. En limitant le dépôt d'énergie dans les quelques millimètres autour des injections, l'irradiation des tissus sains environnant devrait être minimisée.

Cependant, il existe beaucoup de radioéléments émetteurs β^- , chacun avec des caractéristiques différentes. Ces caractéristiques comprennent le spectre d'énergie des particules β^- émis, l'abondance, l'énergie des rayons γ (permettant l'imagerie) et la demi-vie du radioélément. Une partie de ce travail sera donc dédiée à la comparaison des radioéléments possibles afin de chercher celui qui sera le plus adapté à la problématique.

Glioblastome

Il a déjà été mentionné que cette modalité est adaptée aux tumeurs radio-résistantes inopérables. Il est donc possible d'envisager une utilisation potentielle de cette dernière pour les glioblastomes (aussi appelé le glioblastome multiforme) inopérables.

Le glioblastome est la tumeur du cerveau la plus commune chez les adultes, constituant 52% de toutes les tumeurs fonctionnelles du cerveau, 20% de toutes les tumeurs intracrâniennes et affectant environ 2-3 adultes sur 100,000 aux Etats-Unis et à l'UE (Bleeker, Molenaar and Leenstra, 2012). Les glioblastomes sont aussi classifiés niveau 4 par l'organisation mondiale de la santé (la catégorie de tumeur la plus agressive).

Avec une durée de survie médiane de ~14.6 mois (Weller *et al.*, 2005), le pronostic de glioblastomes n'est pas bon. Normalement, ils sont traités via la résection chirurgicale maximale avec la chimiothérapie et l'EBRT adjuvante. Les glioblastomes inopérables ont un pronostic encore pire, typiquement entre 6 et 12 mois. La meilleure façon de les traiter est moins connue, beaucoup d'études n'arrivent pas à augmenter la durée de survie médiane.

Le problème majeur rencontré pendant le traitement de glioblastomes est leur radiorésistance. Avec l'EBRT, il est impossible de suffisamment irradier le glioblastome sans sur-irradier les tissus sains environnants. Cela signifie que la cause de mortalité commune pour un patient avec un glioblastome est la résurgence de la tumeur ou la formation d'une tumeur secondaire.

Dans la section sur la micro-curiethérapie il a été mentionné qu'avec des unités de traitements liquides, cette modalité devrait pouvoir cibler des tumeurs qui sont actuellement inaccessibles pour la curiethérapie, tels que les glioblastomes. De plus, la petite unité de traitement devrait faire en sorte que l'irradiation soit plus confinée à la tumeur et les tissus sains plus épargnés. Ceci signifie que l'irradiation de la tumeur pourrait être augmentée, et donc la probabilité de récurrence de la tumeur ou des tumeurs secondaires devrait être minimisée, et les patients pourraient espérer une augmentation de la durée de survie médiane.

Il est à noter que, même si l'exemple du glioblastome est donné pendant ce travail, ce traitement est potentiellement applicable aux autres types de tumeurs.

Les exigences de traitement

En radiothérapie, l'irradiation est souvent mesurée en dose absorbée avec l'unité gray (Gy), qui est l'énergie déposée par unité de masse (J/kg). La quantité de dose absorbée nécessaire pour détruire des tumeurs dépend non seulement du type de tumeur ciblée mais également de la modalité de radiothérapie qui est utilisée.

Par conséquent, les exigences dosimétriques de micro-curiethérapie doivent être établies expérimentalement (même si elles peuvent être estimées en considérant, par exemple, la dose efficace biologique (Astrahan, 2008)). Sachant que les données expérimentales ne sont pas actuellement disponibles, des estimations ont été réalisées dans ce travail. Toutefois, toutes les méthodologies proposées sont indépendantes des chiffres exacts de dose absorbée choisis, donc ça ne devrait poser aucune difficulté de les changer une fois que les données expérimentales deviendront disponible à une date ultérieure.

Pour une grande partie de ces travaux de recherche, un « planning target volume » (PTV) et un organe à risque seront à la fois utilisés. Tous les deux sont imaginés avec des sphères d'un rayon de 10 mm, placées de sorte que leurs surfaces soient en contact (leurs centres sont séparés par 20 mm). Il y a deux moments dans ce travail où le PTV utilisé est modifié : pendant une étude sur la faisabilité de micro-curiethérapie pour les tumeurs plus volumineuses et pendant la planification de traitement d'un vrai glioblastome.

L'irradiation de la tumeur

D'abord, il faut décider de l'irradiation de la tumeur nécessaire. En radiothérapie, le volume tumoral qui est ciblé est appelé le PTV. Le PTV est composé du « gross tumour volume » (GTV), qui est le volume de la tumeur qui est visible avec l'imagerie telles que la tomodensitométrie (TDM) et l'imagerie par résonance magnétique (IRM). Une marge est ensuite ajoutée autour du GTV, prenant en compte la propagation de la tumeur sous-clinique que ne se voit pas avec l'imagerie, appelé le « clinical target volume » (CTV). Finalement, le PTV est créé en ajoutant une marge autour du CTV, prenant en compte les incertitudes géométriques et le mouvement des organes.

Typiquement, en curiethérapie, la marge entre le CTV et le PTV est zéro. Car le traitement est en contact direct avec la tumeur, et donc le mouvement de la tumeur est négligeable. C'est aussi supposé être le cas avec la micro-curiethérapie. Par conséquent, tout au long de ce travail, le terme PTV – le volume qui est ciblé – est équivalent au CTV.

Pendant le traitement de glioblastomes avec l'EBRT, 60 Gy sont typiquement prescrites au PTV. Bien que 60 Gy soient prescrites pour l'EBRT, cela ne veut pas forcément dire que 60 Gy devrait être prescrite pour le traitement de glioblastomes avec la micro-curiethérapie. Avec l'EBRT, par exemple, la radiation est typiquement délivrée en 30 séances (5 séances par semaine, pendant 6 semaines) de 2 Gy.

La micro-curiethérapie, par contre, est un traitement qui est effectué en une seule séance. Cela est avantageux car le patient doit passer moins de temps à l'hôpital, augmentant la qualité de sa vie. Toutefois, c'est aussi un inconvénient, la dose absorbée prescrite doit être augmentée, dû au fait qu'elle ne soit pas fractionnée. Cet inconvénient ne sera pas forcément un problème car, comme il a été mentionné dans la section de micro-curiethérapie (et sera démontré dans cette recherche), l'irradiation par micro-curiethérapie est plus focalisée sur la tumeur et donc il sera possible d'augmenter la dose absorbée prescrite au PTV sans augmenter l'irradiation des tissus environnant.

Il a donc été supposé que l'équivalent de 60 Gy avec l'EBRT fractionnée en micro-curiethérapie de 100 Gy.

En concordance avec la plupart des environnements cliniques, les plans de traitements ont été considérés convenables si au moins 95% du PTV reçoit au moins 95% de la dose absorbée prescrite au PTV, D_T^{PTV} . Cette exigence est démontrée dans l'inégalité suivante :

$$V^{PTV}(\geq 95\% \cdot D_T^{PTV}) \geq 95\%$$

Les organes à risques

Les OARs sont les organes radiosensibles vitaux qui devraient être protégés pendant les séances de radiothérapie. Les OARs peuvent être divisés en OARs sériels et parallèles.

Les OARs sériels s'arrêtent de fonctionner si une partie de l'OAR s'arrête parce qu'elle est sur-irradiée. Afin de protéger les OARs sériels, il est donc important de considérer la dose absorbée maximale parmi tous les sous-volumes de l'OAR. Dans le cerveau, le chiasma optique, par exemple, est considéré comme un OAR sérial et la dose absorbée maximale permise pour toute partie de cet OAR est de 54 Gy.

Les OARs parallèles, par contre, continuent de fonctionner même si une partie de l'OAR a arrêté de fonctionner. Ils continuent de fonctionner jusqu'à la fraction de l'OAR qui ne fonctionne plus au-delà d'un seuil. Par exemple, la glande pituitaire, continue de fonctionner tant que moins de 10% reçoit plus de 10 Gy. Par conséquent, un plan de traitement convenable doit respecter cette inégalité :

$$V^{OAR}(\geq 10 \text{ Gy}) \leq 10\%$$

Conclusion

Des concepts clés qui seront utilisés tout au long de ce travail ont été introduits.

D'une part, la motivation pour le développement de la micro-curiethérapie a été explorée en discutant des avantages et inconvénients relatifs à d'autres formes existantes de radiothérapie.

L'idée de la planification de traitement sera discutée en détail pendant ce travail et, elle a été introduite. Le concept des fonctions coûts et leur utilisation afin d'évaluer les plans de traitement rapidement a été discutée.

Par la suite, l'importance de sélectionner le radioélément le mieux-adapté a été soulignée.

D'autre part, une utilisation potentielle de la micro-curiethérapie, pour le traitement du glioblastome, a été proposée et devra être validée dans ces travaux de thèse. Avec cette proposition, les exigences dosimétriques qui seront nécessaires pour des traitements de glioblastomes réussis avec la micro-curiethérapie ont été discutées puis validées.

Le TPS spécifique au plan de traitement est détaillé en fin de thèse.

INTRODUCTION

An innovative form of radiotherapy, microbrachytherapy, is currently under development. This therapy targets solid, inoperable tumours by performing injections of a liquid containing radioactive microspheres in suspension.

Many injections are required to sufficiently cover the tumoural volume. A method of treatment planning is required to enable optimal positioning of these injections. The development and validation of such a methodology has been the focus of this research.

Throughout this work, three main questions are addressed:

- How should the dosimetry of microbrachytherapy be performed?
- How should the treatment planning be conducted for this modality?
- What are the optimal injection properties to deliver the most efficient treatment?

The first point is addressed by discussing how the radiation distribution of a single injection can be calculated, and how this distribution can be convolved in the correct position in the patient, such that many injections can be added, and the total absorbed dose distribution can be calculated for a treatment plan with multiple injections.

Once the absorbed dose distribution has been calculated for a given treatment plan, the next topic to be discussed is the extraction of dosimetric information, such as dose-volume histograms (DVHs). This enables treatment plans to be quickly and efficiently analysed and compared.

Following this, optimisation algorithms are developed and compared. These algorithms search for treatment plans of increasing quality, in hoping that at the end of the optimisation process, the proposed treatment plan is of adequate quality to be able to be performed on a patient.

To be able to produce treatment plans of increasing quality, it is important to be able to quantitatively compare and rank plans. This process should condense a fairly abstract treatment plan – consisting of many injections, distributed through 3D space, and its resulting absorbed dose distribution – into several numbers, which can be easily compared. Moreover, this process should be almost instantaneous because the analysis of treatment plans will be performed many times over throughout the treatment planning process. Methods of quantitative comparison via “cost functions” are therefore discussed within this research.

Once methods have been developed for microbrachytherapy dosimetry and the production of optimal treatment plans, certain injection parameters can be tested and compared. Notably, a large study will be carried out into the use of the best-suited radioelement that should be injected. Other studies will be conducted into the ideal injection volume, the ideal initial activity of the injected radioelement and the form the injections should take (spherical or cylindrical).

With the dosimetry, optimisation algorithm and injection parameters all in place, the development of a fully-operational treatment planning system (TPS) with a graphical-user interface is then discussed, and the results displayed.

Microbrachytherapy is a treatment that is still under development. As such, many aspects are protected by intellectual property and cannot be discussed in this work. Notably, the injector and the method of performing the injections will not be addressed. It is hoped that despite the missing information, enough details are given to imagine the pertinence of this treatment, as well as the necessity of the method of treatment planning that is proposed in this work.

The structure of this work is as follows:

- Chapter 1 discusses the generalities of this project – What is microbrachytherapy? Why has it been proposed as a treatment? What tumours could potentially be targeted? What are the dosimetric requirements of this treatment?
- Chapter 2 introduces the concept of treatment planning. In this chapter, many key concepts, such as the cost function, are presented as well as the explanation of several optimisation algorithms that will be used throughout this research.
- Chapter 3 discusses a method of performing the dosimetry of a proposed treatment plan. This chapter finishes with a worked example of the dosimetric analysis of a treatment plan from start to finish.
- Chapter 4 conducts a comparison of the optimisation algorithms that were introduced in Chapter 2 using the methodology of dosimetry that was introduced in Chapter 3. Several test cases are used for the comparison, each test adding a level of complexity on the previous. Once the algorithms have been compared and the best-suited algorithm has been selected, a study is performed into the effect of voxel size on the dosimetry and treatment planning. Finally, a treatment plan is optimised and analysed for a real glioblastoma.
- Chapter 5 compares injection parameters that can be altered in order to deliver the best treatment possible. Such parameters include the choice of injected radioelements, the volume of each injection and their initial activity. It is also proposed in this chapter that cylindrical injections be used instead of spherical injections. The effects on the absorbed dose distribution calculations and the treatment planning process are therefore discussed.
- This work then considers the conclusions and future perspectives of this research.
- Lastly, as Appendix, the methodologies of all the previous chapters are brought together to create a microbrachytherapy TPS. The structure of this TPS is presented, as is the usage of third-party software.

1 GENERALITIES

1.1 Cancer and radiotherapy

Cancer is currently the number one cause of death in France. The principal tools in combatting cancer are surgery, chemotherapy and radiotherapy. This work focuses solely on radiotherapy, which may be used either as an adjuvant therapy or by itself.

When ionising radiation moves through the body, energy is deposited. Regardless of whether the energy is deposited in tumorous or healthy cells, this energy causes cell death. Radiotherapy therefore exploits ionising radiation by focusing it on the tumour.

Many forms of radiotherapy exist, but the common goal amongst the different modalities rests the same: sufficiently irradiate the tumour without over-irradiating surrounding healthy tissue. Over-irradiation of healthy tissue can lead to both acute and chronic secondary effects, whilst under-irradiation of the tumour can result in not completely destroying it or the development of secondary tumours, either at the original location or elsewhere.

1.2 Modalities of radiotherapy

Before describing the new form of radiotherapy, microbrachytherapy, it is necessary to first discuss three existing modalities of radiotherapy and their various advantages and disadvantages. These modalities are external beam radiotherapy (EBRT), brachytherapy and nuclear medicine, with specific emphasis on a type of nuclear medicine called selective internal radiotherapy (SIRT). Once the merits and drawbacks of these modalities have been considered, the reason for the development of microbrachytherapy can be discussed.

1.2.1 External beam radiotherapy

EBRT is the most common form of radiotherapy. It consists of an accelerator that emits high-energy ionising radiation (x-rays or electrons) whilst turning about the patient, the focal point (or isocentre) of the radiation centred on the tumour.

The fundamental drawback of EBRT, however, is that since the accelerator is located outside of the patient, and that the tumour is not always on the surface of the patient, healthy tissue must inevitably be irradiated as the beam passes through to the tumour. The beam also continues on the other side of the tumour (albeit weaker), and healthy tissue is also irradiated at this point.

The harmful effects of the irradiation of healthy tissue are minimised by:

- Fractionating the treatment (for example, delivering a small amount of radiation every day over several weeks), such that healthy tissue has the chance to recover;
- Turning the accelerator around the patient, such that the harmful radiation is shared out by all the surrounding healthy tissue rather than being localised in one area; and
- Using collimators that shape the beam so that it matches the form of the tumour as closely as possible.

Despite minimising exposure to healthy tissue, other modalities of radiotherapy, such as brachytherapy and nuclear medicine, have been developed with the goal of further minimising their irradiation.

1.2.2 Brachytherapy

Brachytherapy consists of inserting solid radioactive materials (for example, in the form of metal grains) in and around the tumour. The key advantage of brachytherapy is that, by introducing the radioactive source much closer to the target, different types of radiation can be used. Sources can be used that deposit their energy over much shorter distances, limiting irradiation of surrounding healthy tissue.

On the contrary, the main disadvantage of brachytherapy is that, because the sources need to be physically placed in the desired location, its usage is therefore limited to treatments of cancers that are easily accessible from the surface of the patient, such as tumours of the skin, prostate, breast and the cervix.

1.2.3 Nuclear medicine

Nuclear medicine can be divided into two subcategories, targeted radiotherapy (TRT) and SIRT:

- TRT, also known as molecular radiotherapy (MRT), consists of introducing radiopharmaceuticals (radioelements attached to pharmaceuticals) intravenously into the patient. The pharmaceuticals bond readily with cell receptors found commonly on the targeted tumorous cells, where they emit their radiation, killing the cells.
- SIRT is used for inoperable liver tumours. It consists of injecting radioactive microspheres into arteries that nourish the liver such that the radioelements collect in the tumour, emitting their radiation and hopefully destroying it.

The advantage of SIRT is that it is much less invasive than modalities such as brachytherapy, owing to the fact that the radioactive sources are introduced intravenously rather than directly.

The obvious disadvantage of SIRT is that it lacks precision. The microspheres are not introduced directly into the tumour, nor are they attached to pharmaceuticals that attach preferentially to tumorous cells, as with TRT.

1.3 Microbrachytherapy

The search continues for ever-better modalities of radiotherapy that increase the irradiation of the tumour whilst simultaneously decreasing the irradiation of healthy tissue. With this goal in mind, the idea of microbrachytherapy was proposed.

Microbrachytherapy can be considered as a hybrid between brachytherapy and SIRT. It is a hybrid between these two modalities because radioactive microspheres (as with SIRT) are injected directly inside of the tumour (as with brachytherapy).

By passing from grains in brachytherapy to microspheres in microbrachytherapy, the volume of each treatment unit is reduced. This should help to better conform the distribution of radiation around the tumour, both increasing the irradiation of the tumour and decreasing the irradiation of surrounding healthy tissue compared to brachytherapy. Furthermore, by performing the injections directly inside of the tumour, it is hoped that the drawback of the lack of precision in SIRT can be countered.

Furthermore, by using a treatment unit in liquid form (the microspheres are in liquid suspension), the invasive catheters that are required to implant the metallic grains in brachytherapy can be replaced by a less invasive needle, such that deeper tumours can potentially be targeted.

Since the distribution of radiation should, in theory, be better conformed about the tumour, the irradiation of the tumour can therefore be increased without putting the surrounding healthy tissue and organs at greater risk. Microbrachytherapy could therefore be a key tool in the treatment of inoperable, radioresistant cancers.

1.3.1 The treatment

In microbrachytherapy, injections of volumes ranging between 5 and 20 μL are used. Each injection contains the microspheres suspended homogeneously in solution. It is assumed that, once injected in the tumour, the liquid forms a uniform sphere about the needle, such that there is an even distribution of microspheres about the needle.

Once they have been injected inside the tumour, it is required that the radioelements rest in their positions and emit the entirety of their energy. It is for this reason that microspheres are used; they are sufficiently large that their biological half-life is quasi-infinite, as has been demonstrated experimentally (Bult *et al.*, 2012).

1.4 Treatment planning

Of all the modalities of radiotherapy already discussed in this introduction, treatment planning is used for all of them. Treatment planning is the search for the best possible treatment that can be delivered to each patient.

Inverse problem-solving algorithms are often exploited to find optimal treatment plans. They iteratively alter treatment plans so as to minimise user-determined penalising functions. These penalising functions, known as cost functions, penalise unwanted characteristics of the treatment plan (for example, under-irradiation of the tumour or over-irradiation of a neighbouring organ). Therefore, as the optimisation algorithm minimises the cost functions, the quality of the treatment plan increases.

At the end of the optimisation, the proposed treatment planning can be validated and eventually performed on the patient.

As it is a new modality of radiotherapy, a method of treatment planning does not currently exist for microbrachytherapy. One of the key thematics of this work is therefore to propose and develop such a methodology.

1.5 The choice of radioelement

Microbrachytherapy will inevitably exploit β^- -emitting radioelements. This is because such radioelements deliver their energy in the few millimetres surrounding the source. By confining the energy deposition to within a few millimetres from the injected sources, irradiation of neighbouring healthy tissue should be minimised.

There exist many β^- -emitting radioelements, however, each with different characteristics. Such characteristics include the energy spectrum of the emitted β^- -particles, the abundance and energy of γ -rays (which enables imaging) and the radioelement's half-life. Part of this work will therefore be dedicated to the comparison of possible radioelements in the search for the one that is best suited to this application.

1.6 Glioblastomas

It was mentioned at the end of Section 1.3 that microbrachytherapy is tailored for inoperable, radioresistant tumours. A potential usage of this modality is inoperable glioblastomas (also known as glioblastoma multiforme).

Glioblastomas are the most common malignant primary brain tumour in adults, constituting 52% of all functional tissue brain tumours, 20% of all intracranial tumours and affecting around 2-3 adults per 100,000 in the USA and EU (Bleeker, Molenaar and Leenstra, 2012). Glioblastomas also fall into the WHO Grade IV category (the most aggressive category of tumours of the central nervous system).

With a median survival time of around 14.6 months (Weller *et al.*, 2005), the prognosis of glioblastomas is poor. They are usually treated via maximal surgical resection with adjuvant chemotherapy and EBRT. Inoperable glioblastomas have an even worse prognosis, typically ranging between 6 and 12 months. Their optimal treatment path, however, is less well-defined, with many studies struggling to significantly improve median survival times.

A major problem encountered when treating glioblastomas is their radio-resistance. With EBRT, it is impossible to sufficiently irradiate the glioblastoma without over-irradiating the surrounding healthy tissue. This means that the common cause of death for a patient suffering from a glioblastoma is tumour reoccurrence or the formation of secondary tumours.

Section 1.3 stated that microbrachytherapy, given its liquid treatment unit, should be able to target tumours that are currently inaccessible for brachytherapy, such as glioblastomas. Furthermore, the small treatment unit volume means that the distribution of radiation should be confined to the tumour, and the surrounding healthy tissue should be better spared. This means that, compared to other modalities of radiotherapy, the irradiation of the tumour can be increased, and hopefully the probability of tumour reoccurrence and secondary tumours should be decreased, serving to improve median survival times of glioblastoma patients.

It should be noted that, although the example of the glioblastoma is given throughout this work, this treatment is potentially applicable for other types of tumours.

1.7 Treatment requirements

In radiotherapy, irradiation is often measured as absorbed dose using the unit gray (Gy), which is the energy deposited per unit mass (J/kg). The amount of absorbed dose required to destroy tumours not only depends on the type of tumour being targeted, but also the modality of radiotherapy that is being employed.

Therefore, the dosimetric requirements of microbrachytherapy need to be determined experimentally (although they can be estimated by considering, for example, the biological effective dose (Astrahan, 2008)). Since experimental data is not currently available, estimates have been used in this research. All methodologies proposed throughout this work, however, are independent of the exact absorbed dose requirements and limits. As such, changing them should not pose any problems when such data becomes available at a later date.

For the majority of this work, both an imaginary planning target volume (PTV) and an imaginary organ at risk (OAR) will be used. They are both assumed to be spheres with radii of 10 mm, placed such that their surfaces touch (their centres are therefore separated by 20 mm). In two instances the PTV is changed: during a study into the feasibility of microbrachytherapy for larger tumours and during the treatment planning of a real glioblastoma.

1.7.1 Irradiation of the tumour

First, it is necessary to decide the required irradiation of the tumour. In radiotherapy, the tumoural volume that is targeted is known as the PTV. The PTV is composed of the gross tumour volume (GTV), which is the volume of the tumour that can be imaged with modalities such as computed tomography (CT) and magnetic resonance imaging (MRI). A margin is then added around the GTV, accounting for subclinical tumour spread that cannot be imaged, called the clinical target volume (CTV). Lastly, the PTV is created by adding a margin around the CTV to account for geometrical uncertainties and organ movement.

Typically, in brachytherapy, the margin between the CTV and the PTV is zero. This is because the treatment is in direct contact of the tumour, and so tumour movement is negligible. This too, is assumed to be the case with microbrachytherapy. Hence, throughout this research, the term PTV – the volume that is targeted – is equivalent to the CTV.

During the treatment of glioblastomas with EBRT, 60 Gy is typically prescribed to the PTV. Although this is the case for EBRT, 60 Gy should not necessarily be prescribed for the treatment of glioblastomas with microbrachytherapy.

EBRT, for example, is often fractionated, typically delivering the radiation in 30 sessions of 2 Gy (5 sessions per week, over 6 weeks). Microbrachytherapy, however, is a treatment that is delivered in a single session. This is advantageous in that the patient passes less time in hospital, increasing their quality of life. However, it is also disadvantageous in that the prescribed absorbed dose must therefore be increased since it is not fractionated.

This disadvantage may not necessarily pose a problem since, as was mentioned in Section 1.3 (and will hopefully be demonstrated throughout this research), the irradiation with microbrachytherapy is more focused on the tumour, such that it is possible to increase the prescribed absorbed dose to the PTV without increasing the irradiation of surrounding tissue.

Throughout this work, it was assumed that 100 Gy for microbrachytherapy was equivalent to 60 Gy with fractionated EBRT.

In accordance with the majority of clinical environments, treatment plans were required to deliver at least 95% of the PTV's prescribed absorbed dose of 100 Gy, D_T^{PTV} , to at least 95% of the PTV. This requirement can be expressed as the following inequality:

$$V^{PTV}(\geq 95\% \cdot D_T^{PTV}) \geq 95\% \tag{1.1}$$

1.7.2 Organs at risk

OARs are vital radiosensitive organs that need to be protected during radiotherapy. OARs can be subdivided into those that are serial and those that are parallel.

Serial OARs will fail if any single part of the OAR fails. This failure occurs when the maximum irradiation of any sub-volume surpasses a given limit. Hence, when protecting a serial OAR, it is important to consider its maximum absorbed dose. In the brain, the optic chiasm, for example, is considered a serial OAR and the maximum absorbed dose permitted for any part of it during EBRT is 54 Gy.

Parallel OARs, on the other hand, will continue to function even if part of the organ has failed. They will continue working until the fraction of the OAR that has failed passes a given threshold. For example, the pituitary gland is considered a parallel OAR, and will continue to function as long the volume receiving greater than or equal to 10 Gy is less than 10%.

Occasionally, an OAR will be used in this work. The OAR will be treated as a parallel OAR which, as with the pituitary gland, will continue to function as long as less than 10% of it receives greater than or equal to 10 Gy. Hence, an acceptable treatment plan respects the following inequality:

$$V^{OAR}(\geq 10 \text{ Gy}) \leq 10\% \tag{1.2}$$

1.8 Conclusion

Throughout this chapter, key concepts have been introduced that will be used throughout this work. Firstly, the motivation for the development of microbrachytherapy was explored by discussing the relative advantages and disadvantages of other existing forms of radiotherapy.

The idea of treatment planning, a concept that will be discussed in great detail throughout this work, was introduced. The idea of cost functions and how they can be used to quickly evaluate treatment plans was also introduced. Following this, the importance of selecting the best-suited radioelement was stressed.

A potential usage of microbrachytherapy, the glioblastoma, was then introduced as well as the potential dosimetric requirements that would have to be put in place for its successful treatment with this modality.

With all of the ideas that have been briefly introduced in this chapter in mind, it is now possible to discuss the first topic of this research – treatment planning.

2 TREATMENT PLANNING

2.1 Introduction

Treatment planning is the search for the plan that delivers the best treatment to a given patient. Modern treatment planning is highly automated, relying on optimisation algorithms to produce plans that satisfy treatment requirements.

Treatment planning is an indispensable tool in modern radiotherapy, regardless of the modality. For a new modality, such as microbrachytherapy, it is therefore imperative to develop a methodology of treatment planning. Since such a methodology does not yet exist for microbrachytherapy, it is explored throughout this research.

The optimisation algorithms exploited throughout this work are known as inverse problem-solving algorithms because they start at the end – with the desired result – and work backwards to calculate the required input that would yield the result that best matches that which was desired.

Such algorithms are typically iterative and a schematic of an inverse problem-solving algorithm is shown below, in Figure 2.1.




Figure 2.1: Iterative process used to create optimal treatment plans. The cost function is used to evaluate the quality of the solutions, and if the stopping criteria are not met, the optimisation algorithm alters the solution. The forward problem then determines the outcome of the newly proposed solutions.

As can be seen in Figure 2.1, inverse problem-solving algorithms function by initially creating treatment plans. Dosimetry is then performed in what is known as the forward problem. Following this, cost functions are used to evaluate the quality of the treatment plans and if the stopping criteria are not met, the optimisation algorithm is used to modify the treatment plans. This loop continues until the stopping criteria are met, at which point the optimisation process is finished and the optimised plan or optimised plans are presented to the user.

The main components of Figure 2.1 are discussed throughout this chapter:

- Section 2.2 introduces the forward problem (but this will be explored in more depth in Chapter 3),
- Section 2.3 discusses the cost function,
- Section 2.4 introduces optimisation algorithms,
- Section 2.5 describes specific optimisation algorithms that will be used throughout this research,
- Section 2.6 discusses the stopping criteria that could be used, and
- Section 2.7 concludes this chapter.

2.2 The forward problem

The forward problem, with regards to inverse problem-solving algorithms, is the part of the optimisation process that follows causality (thus the name). As was mentioned in the introduction to this chapter, inverse problem-solving optimisation algorithms start with the desired outcome, and work backwards to determine the required input. Nevertheless, it is required that part of the process work in the normal direction to determine if the newly proposed solution matches the desired output.

In the specific case of treatment planning, the optimisation algorithm is used to alter the treatment plan's injections, and then the forward problem is used to calculate the plan's resulting irradiation distribution (known as dosimetry).

The method that has been developed to calculate the dosimetry of a given microbrachytherapy treatment plan requires in-depth discussion and analysis of the process. For this reason, the forward problem will not be discussed further in this chapter, but will be explored thoroughly in Chapter 3.

2.3 The cost function

Cost functions are used to quantify the quality of a treatment plan. As input they take a treatment plan's dosimetry (that was calculated by the forward problem), and as output they generate a single number. This is useful because, by condensing somewhat abstract treatment plans into single numbers, the task of comparing treatment plans and eventually ranking them becomes considerably easier. Once the quality of treatment plans can be quantified and the plans can be ranked, the algorithms use this information to propose new treatment plans, hopefully of increased quality.

Cost functions are used to compare the properties of the proposed treatment plan with the properties of the ideal treatment plan. When the difference between the two is decreased, the cost function decreases accordingly. Optimisation algorithms therefore search for treatment plans that minimise cost functions.

Each cost function concentrates on a specific aspect of the treatment plan. One, for example, might be focused on the irradiation of the PTV. Another might be concentrated on the irradiation of healthy tissue. Another might be concentrated on, for example, how long the treatment would take.

All cost functions are designed so that they are minimised as the quality of the treatment plan increases. This small subtlety means that the cost function that controls qualities that should be maximised throughout the optimisation process should be inverted. For example, during the optimisation process, the goal is to *maximise* the irradiation of the tumour. To create a cost function for this, the same objective can be reversed so that the goal is now to *minimise* the under-irradiated part of the tumour. The two are identical, except that the latter is a minimisation and the former a maximisation. This subtlety will be exploited later in this chapter.

2.3.1 Types of cost functions

Many types of cost function exist within the domain of radiotherapy. Cost functions can be divided into those that assess a treatment plan based on physical factors and those that assess radiobiological factors. A simple example of physical factors is the maximum absorbed dose in a region. Conversely, the equivalent uniform dose (EUD) is an example of a radiobiological factor.

Physical factors have the advantage of being intrinsically quantitative; it is easy, for example, to determine the maximum absorbed dose in a structure. Radiobiological factors, on the other hand, are difficult to quantify and are subject to much larger uncertainties. Once quantified, however, radiobiological factors describe the effects of the radiation at a cellular level, which is a much better indication of the efficacy of a treatment plan.

Non-exhaustive lists of physical and radiobiological cost functions are shown below in Tables 2.1 and 2.2, respectively.

Name	Description	Functional form	PTV/OAR
Minimum	Minimum absorbed dose in any voxel, d_i , in a VOI.	$\min(d_i^{VOI})$	PTV or OAR
Maximum	Maximum absorbed dose in any voxel, d_i , in a VOI.	$\max(d_i^{VOI})$	PTV or OAR
Mean	Mean absorbed dose in a VOI.	$\frac{1}{n_{VOI}} \sum_{i \in VOI} d_i$	PTV or OAR
Dose-volume	Absorbed dose for a volume, X %, on a cumulative DVH. Written as $D_{X\%}^{VOI}$.	$D^{VOI}(X\%)$	PTV or OAR
Volume-dose	Volume (%) of a VOI receiving $\geq X$ Gy. Written as V_X^{VOI} .	$V^{VOI}(\geq X \text{ Gy})$	PTV or OAR
Under-irradiated volume-dose	Volume (%) of the PTV receiving $< D_T^{PTV}$. Effectively 100% minus volume-dose.	$V^{PTV}(< D_T^{PTV})$	PTV-
Quadratic dose difference	Sum of the square of the difference between each voxel, d_i , and D_T in the PTV.	$\sum_{i \in PTV} (D_T^{PTV} - d_i)^2$	PTV
Homogeneity index (HI)	The difference between $D_{2\%}^{PTV}$ and $D_{98\%}^{PTV}$ compared to D_T^{PTV} .	$\frac{D_{2\%}^{PTV} - D_{98\%}^{PTV}}{D_T^{PTV}}$	PTV
Conformity index (CI)	Fraction of PTV receiving $\geq D_T^{PTV}$ (V_T^{PTV}/V^{PTV}) multiplied by ratio of PTV receiving $\geq D_T^{PTV}$ compared to total volume (PTV or otherwise) receiving $\geq D_T^{PTV}$ (V_T^{PTV}/V_T^{all}).	$\frac{V_T^{PTV}}{V^{PTV}} \cdot \frac{V_T^{PTV}}{V_T^{all}}$	PTV and healthy tissue
Conformal index (COIN)	All-in-one function that combines the CI and also penalises over-irradiated OARs (where V_{thresh}^i is the fractional volume of the i^{th} OAR receiving $\geq D_{thresh}^{OAR}$).	$1 - CI \cdot \prod_{i=1}^{N_{OARs}} (1 - V_{thresh}^i)$	PTV, healthy tissue and OARs
Number of injections	Number of injections required.	n_{inj}	
Time	Estimated time for given treatment plan.	t	

Table 2.1: Non-exhaustive list of physical cost functions used in radiotherapy. VOI is a volume of interest. Voxels are 3D pixels.

Name	Description	Functional form	PTV/OAR
Equivalent uniform dose (EUD)	Uniform dose to an organ that causes the same effect as the non-uniform dose delivered in reality.	$\left(\frac{1}{N_{VOI}} \sum_{i \in VOI} (d_i)^a \right)^{1/a}$	PTV or OAR
Normal tissue complication probability (NTCP)	The probability of a complication in normal tissue.	$\left(1 + \left(\frac{TD_{50}}{EUD} \right)^{4 \cdot \gamma_{50}} \right)^{-1}$	OAR
Tumour control probability (TCP)	The probability of tumour control.	$\left(1 + \left(\frac{TCD_{50}}{EUD} \right)^{4 \cdot \gamma_{50}} \right)^{-1}$	PTV

Table 2.2: Non-exhaustive list of radiobiological cost functions used in radiotherapy. See Section 2.3.3 for a discussion of a , γ_{50} , TD_{50} and TCD_{50} .

2.3.2 Physical cost functions

2.3.2.1 Homogeneity index

Figure 2.2 shows a fictional DVH that could typically be encountered for a PTV in EBRT. It demonstrates some of the physical cost functions discussed in Table 2.1. The target absorbed dose for the PTV in this example is 55 Gy. Shown in the figure are $D_{98\%}^{PTV}$, $D_{2\%}^{PTV}$ and $V_{55\text{Gy}}^{PTV}$, which are 50 Gy, 60 Gy and 88% of the volume, respectively.

These values can be used to calculate the HI, which is the difference between $D_{98\%}^{PTV}$ and $D_{2\%}^{PTV}$ divided by D_T^{PTV} . This is a useful metric because when the difference between $D_{98\%}^{PTV}$ and $D_{2\%}^{PTV}$ is reduced, then the irradiation of the PTV is homogeneous (represented by a steep gradient in the DVH). Using the $D_{98\%}^{PTV}$ and $D_{2\%}^{PTV}$ that can be seen in Figure 2.2, the DVH in this example has an HI of 0.18.




Figure 2.2: Fictional EBRT DVH of a PTV. Also depicted are $D_{98\%}^{PTV}$ (50 Gy), $D_{2\%}^{PTV}$ (60 Gy) and $V_{55\text{Gy}}^{PTV}$ (88%).

An example of a microbrachytherapy DVH is given after the explanation of microbrachytherapy dosimetry (Chapter 3), in Figure 3.9. It can be seen for this example that maximum absorbed dose in the PTV is extremely high. As such, $D_{2\%}^{PTV}$ is enormous and so therefore is the HI. Hence, it can be seen that HI is a useless metric for microbrachytherapy, which delivers extremely elevated maximum absorbed doses.

2.3.2.2 Quadratic dose difference

The quadratic dose difference cost function is typically used for the PTV in EBRT. In EBRT this is particularly useful because it both penalises over- and under-irradiated voxels in the PTV, promoting a homogeneous coverage of absorbed dose.

As with the HI, this cost function is also inapt for microbrachytherapy. This is because voxels close to the centre of injections are so heavily irradiated that the penalties that they occur are enormous.

For example, for a target absorbed dose of 100 Gy, a single voxel receiving 10,000 Gy contributes $(10,000-100)^2 = 9.8 \times 10^7$ to the cost function, compared to a cost function of 10^4 for a voxel receiving 0 Gy. Using this cost function, the cost is therefore minimised when no injections are used, which is clearly illogical.

2.3.2.3 Under-irradiated volume-dose

HI and the quadratic dose difference cost functions were deemed unsuitable for the PTV with microbrachytherapy because portions of the PTV receive absorbed doses that greatly exceed the target absorbed dose. This is inherent to treatments such microbrachytherapy, where the source is introduced inside the patient.

The ideal microbrachytherapy cost function for the PTV should therefore not penalise over-irradiated voxels. For this reason, of all the physical cost functions, the under-irradiated volume-dose cost function was deemed the most appropriate for the PTV.

As was mentioned in the last paragraph of Section 2.3, although the goal of microbrachytherapy is to maximise the volume that is adequately irradiated (as per Equation 1.1), cost functions should always be minimised. It is for this reason that a cost function that minimises the under-irradiated fraction of the PTV was chosen.

2.3.3 Radiobiological cost functions

As previously mentioned, radiobiological cost functions are subject to larger uncertainties. These come from the parameters (α , γ_{50} , TD_{50} and TCD_{50} are cited in Table 2.2) that need to be experimentally determined.

The EUD, for example, is the theoretical homogeneous radiation distribution that causes the same amount of clonogen death as the inhomogeneous radiation distribution that is delivered in reality (McGary, Grant 3rd. W. and Woo, 2000). The EUD can therefore be used as a cost function, because a complex distribution can be reduced to a single value of absorbed dose that describes the radiobiological effect of the irradiation. Although the EUD was initially designed for just the PTV, it was later adapted for normal tissue (Gay and Niemierko, 2007). The problem is that α is an empirically-determined parameter, dependent on both the modality of radiotherapy and the tissue in question, and as such there is no data available for this new modality.

For mature forms of radiotherapy, such as EBRT, there exists enough patient data to be able to ascertain these parameters. For a new form of radiotherapy, such as microbrachytherapy, however, the use of radiobiological cost functions is currently infeasible.

2.3.4 Chosen cost functions

Having ruled out radiobiological cost functions, the cost functions used throughout this work consist exclusively of physical cost functions selected from Table 2.1.

2.3.4.1 The PTV required cost function

It was mentioned that the under-irradiated volume-dose cost function should be used for the PTV. As shown in Equation 1.1, at least 95% of the PTV needs to receive 95% of D_T^{PTV} , 100 Gy. The required cost function for the PTV therefore simply returns the volume of the PTV receiving less than 95% of D_T^{PTV} :

$$Cost_{PTV} = V^{PTV}(< 95\% \cdot D_T^{PTV}) \quad 2.1$$

In accordance with Section 1.7.1, only treatment plans with $Cost_{PTV} \leq 5\%$ are acceptable, since this corresponds to plans delivering over 95% of D_T^{PTV} to at least 95% of the PTV.

2.3.4.2 The number of injections cost function

The second cost function, $Cost_{inj}$, penalises the number of injections, n_{inj} . This is used because it is clearly advantageous to be able to find a treatment plan that both sufficiently irradiates the PTV, but also does so with a minimal number of injections. The functional form of $Cost_{inj}$ is simply:

$$Cost_{inj} = n_{inj} \quad 2.2$$

2.3.4.3 The OAR cost function

As was mentioned in Section 1.7.2, an OAR will occasionally be used in this research. When this is the case, a corresponding cost function, $Cost_{OAR}$, is therefore also added.

No more than 10% of the OAR can receive over 10 Gy. A dose-volume cost function was therefore chosen, with the following functional form:

$$Cost_{OAR} = V^{OAR}(\geq 10 \text{ Gy}) \quad 2.3$$

2.3.4.4 The conformity index

Very briefly, in Section 4.6.2, the conformity index (CI) will be used. The CI is an interesting metric because, as can be seen in Table 2.1, it quantifies a treatment plan both with regards to the irradiation of the PTV and the sparing of the healthy tissue. The CI is therefore minimised when the volume of tissue receiving greater than or equal to D_T^{PTV} perfectly coincides with the form of the PTV. The CI will be used with the simple following functional form:

$$Cost_{CI} = CI \quad 2.4$$

2.4 Optimisation algorithms

With the cost functions in place, it is now possible to discuss optimisation algorithms, which will use the cost functions to quantify and rank treatment plans and then modify them accordingly, proposing plans of increasing quality.

Optimisation algorithms can be categorised as stochastic and deterministic. Stochastic algorithms exploit randomness to find optimal treatment plans. Deterministic algorithms, conversely, will always converge to the same point if the input conditions are unchanged. Deterministic algorithms, therefore, are highly dependent on their initialisation. Examples of deterministic and stochastic algorithms are Nelder-Mead Simplex (NMS) and the Bees algorithm (BA), respectively.

Optimisation algorithms can also be categorised as interactive, single-objective (SO) and multi-objective (MO) algorithms. These three types of algorithms will be described in the following sections.

It should be noted that the terms “solution” and “treatment” plan will be used interchangeably throughout this chapter. This is because optimisation algorithms search for optimal solutions and in this work, those solutions happen to be treatment plans.

Lastly, it should also be noted that throughout this work, algorithms rank treatment plans first by the number of treatment requirements that are not respected (i.e., $Cost_{PTV} \leq 5\%$ or $Cost_{OAR} \leq 10\%$) and then by increasing cost function. In this way, the top ranked treatment has the smallest cost function of all plans that violate the fewest requirements.

2.4.1 Interactive algorithms

Interactive algorithms will not be used in this research but they are explained briefly here for completeness.

Interactive algorithms require the user to rank a proposed ensemble of solutions at the end of each iteration. This means that no cost functions are required because the user replaces them during the ranking of proposed treatment plans. This is useful for characteristics that should be optimised but are difficultly expressed as functional forms that can be penalised.

It can be seen that interactive algorithms are extremely time-consuming for the user because they require constant surveillance. It is for this reason that they are not considered in this work.

2.4.2 Single-objective algorithms

The majority of optimisation algorithms encountered in literature are SO algorithms. As the name might suggest, SO algorithms contain a single cost function. The cost function is often a weighted sum of multiple costlet functions (cost functions are referred to as costlets when they are combined together). The form of an SO cost function is shown in Equation 2.5:

$$Cost = \sum_i w_i \cdot Costlet_i \quad 2.5$$

where $Costlet_i$ is the i^{th} costlet (for example, any costlet shown in Tables 2.1 or 2.2) and w_i is its relative weighting. The sum of all of the weightings should be equal to unity ($\sum w_i = 1$).

With SO algorithms, each treatment plan has a single associated cost. The treatment plan with the lowest cost function is therefore deemed the best.

The costlets need to be normalised if they are to be combined into a single sum. For example, $Cost_{PTV}$ and $Cost_{OAR}$ are constrained between 0% and 100%. However, because a treatment plan could use any number of injections, $Cost_{inj}$ could take any value. Hence, when $Cost_{inj}$ is combined with $Cost_{PTV}$ and $Cost_{OAR}$, it is expressed as a percentage of the maximum number of injections allowed for that optimisation (a user-defined parameter). This normalises $Cost_{inj}$, allowing it to be used with other costlets. For example, a treatment plan with 15 injections when 20 is the maximum gives a $Cost_{inj}$ of 75%.

SO algorithms are limited in that the user must decide relative weightings prior to starting the optimisation process. If, at the end, the user wishes to alter these weightings, they must restart the optimisation process anew.

Examples of SO algorithms are NMS and BA, which are described in Sections 2.5.2 and 2.5.3, respectively.

2.4.3 Multi-objective algorithms

MO algorithms are advantageous compared to their SO counterparts in that the costlet functions are not joined together. At the end of the optimisation, the user is presented with a set of non-dominated treatment plans that form the Pareto front.

A treatment plan is considered dominated if there exists at least one other treatment plan whose costlet functions are all better. This is depicted with a 2D example in Figure 2.3. Treatment plan D, for example, does not dominate treatment plan C because $C_1 < D_1$. Similarly, treatment plan C does not dominate treatment plan D because $D_2 < C_2$. Neither treatment plan C nor treatment plan D can be deemed better than the other, and so they are referred to as non-dominated, and they form the Pareto front. Treatment plan H, on the other hand, is not part of the Pareto front because it is dominated by treatment plan D ($D_1 < H_1$ and $D_2 < H_2$).

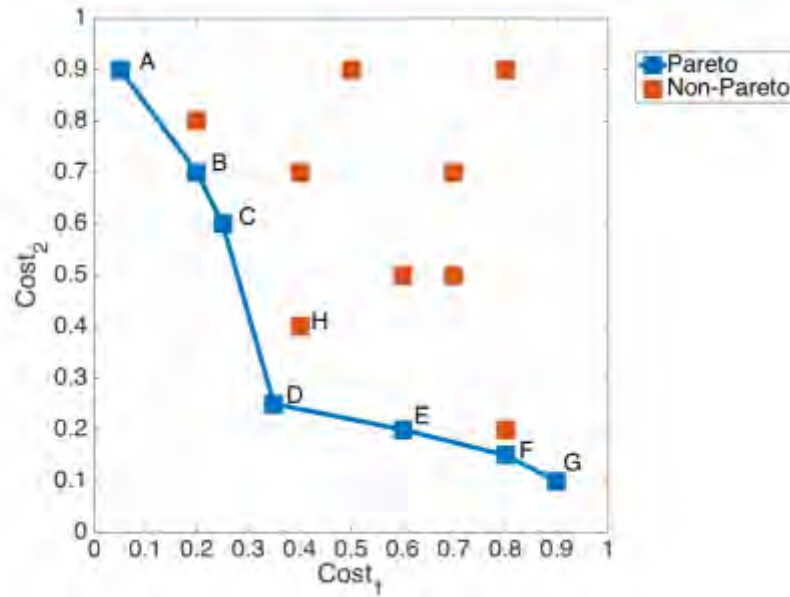


Figure 2.3: Pareto front (composed of Pareto-optimal solutions) and other, non-Pareto solutions. Solutions are considered Pareto-optimal as long as no solution has been found that is better than it with respect to all cost functions. H, for example, is not Pareto-optimal because D is better both in regards to $Cost_1$ and $Cost_2$.

At the end of an MO optimisation, the user is then proposed with all the non-dominated (or Pareto-optimal) solutions, and can select the treatment plan that they prefer. Hence, the crucial difference between SO and MO algorithms is this: with SO algorithms, the importance between the costlet functions must be decided prior to the optimisation, whereas with MO algorithms, this decision takes place at the end. With MO algorithms, the user can change costlet weightings as many times as desired without re-running the optimisation process, whereas SO algorithms need to be restarted each time the weightings change.

This is highlighted in Figure 2.4, where the 7 Pareto-optimal treatment plans of Figure 2.3 are expressed as a weighted sum of their costs for a range of relative weightings.




Figure 2.4: Total costs of the 7 Pareto-optimal treatment plans shown in Figure 2.3, displayed for different relative weightings: (a) $w_1=0.1$, $w_2=0.9$; (b) $w_1=0.5$, $w_2=0.5$; and (c) $w_1=0.9$, $w_2=0.1$. It can be seen that for an SO algorithm, the “best” solution (lowest total cost) is heavily dependent on the chosen costlet weightings.

It can be seen that the “best” treatment plan (lowest total cost) in Figure 2.4a is not the same as in Figure 2.4b or Figure 2.4c. Although an MO algorithm would have proposed all 7 treatment plans, an SO algorithm would have only proposed one. For example, an SO algorithm with weightings of $w_1=0.1$ and $w_2=0.9$ would have only treatment plan A and information relating to the other 6 treatment plans would have been lost. It can therefore be seen that MO algorithms are more powerful, since more information is returned to the user.

Although MO algorithms are more powerful, they are also more complex. This is because instead of minimising a 1D problem (one total cost function) they minimise an ND problem (where N is equal to the number of costlet functions). This means that the computing time of an MO optimisation is often greater than that of an SO optimisation.

The MO algorithm used in this work, as described in Section 2.5.4, is the non-dominated sorting genetic algorithm-II (NSGA-II).

2.5 Explanation of various algorithms

Now that various concepts such as the cost function and stochastic, deterministic, SO and MO algorithms have been introduced, it is possible to discuss in detail the optimisation algorithms that will be compared in the following chapter.

The algorithms that will be presented (and later compared) were chosen to represent the widest range of algorithms possible, consisting of a mix of stochastic and deterministic algorithms and also SO and MO algorithms. All presented algorithms do not require derivatives of the cost function to be calculated, nor do they require the cost function to be continuous, two qualities that are advantageous.

Discussed in Sections 2.5.1 through 2.5.4 are the conjugate gradient (and similar algorithms), NMS, BA and NSGA-II.

2.5.1 Conjugate gradient and similar algorithms

2.5.1.1 The linear problem

The treatment planning process can be considered as a linear problem:

$$AI = D \quad 2.6$$

where I is a binary array, containing an element for each voxel in the PTV. The array element is 1 when there is an injection present in the corresponding voxel and 0 when there is not. D describes the absorbed dose distribution that results from the I injections. D is also an array, containing an element for each voxel of that patient that could potentially be irradiated. Lastly, A is a matrix that describes the absorbed dose in a given voxel for an injection in another given voxel.

Equation 2.6 can therefore be expanded to:

$$\begin{pmatrix} a_{11} & a_{12} & a_{13} & \cdots & a_{1n} \\ a_{21} & a_{22} & a_{23} & & \\ a_{31} & a_{32} & a_{33} & & \\ \vdots & & & \ddots & \\ a_{m1} & & & & a_{mn} \end{pmatrix} \begin{pmatrix} i_1 \\ i_2 \\ i_3 \\ \vdots \\ i_n \end{pmatrix} = \begin{pmatrix} d_1 \\ d_2 \\ d_3 \\ \vdots \\ d_m \end{pmatrix} \quad 2.7$$

where i_j states whether or not the injection is present in the j^{th} voxel of the PTV, d_i is the absorbed dose in the i^{th} voxel of the dose grid and a_{ij} is the absorbed dose in the i^{th} voxel due to an injection in the j^{th} voxel.

2.5.1.2 Solving the linear problem

The linear problem should be solved inversely. The ideal absorbed dose distribution, D^{ideal} , is already known:

$$d_i^{ideal} = \begin{cases} D_T & \text{for } i \in \text{PTV} \\ 0 & \text{otherwise} \end{cases} \quad 2.8$$

and so, the goal is to find the array, I , that delivers a D that best matches D^{ideal} . This is done by rearranging Equation 2.6 such that I is the subject:

$$I = A^{-1}D^{ideal} \quad 2.9$$

where A^{-1} is the inverse of A .

A plethora of algorithms exist that are capable of solving for I , such as the conjugate gradient (Hestenes and Stiefel, 1952), preconditioned conjugate gradient, bi-conjugate gradient, conjugate gradient squared and minimal residual method.

However, the elements of I that are found by these methods are continuous. The solution of this problem, on the contrary, should be binary. For example, each element of I should be either 0 or 1 – either there is an injection or there is not. It is not possible to have 0.673 of an injection. For this reason, all of the algorithms mentioned in the previous paragraph cannot be used for this particular problem.

2.5.2 Nelder-Mead Simplex

NMS (Nelder and Mead, 1965) is a deterministic SO algorithm. The algorithm works by optimising an N -simplex. An N -simplex is an extension of a triangle or a tetrahedron to a shape with $(N + 1)$ vertices. In the case of NMS, N is equal to the number of degrees of liberty of the problem.

Since each injection is free to move in 3 dimensions, $N = 3n_{inj}$. That is to say, if one wishes to optimise a treatment plan that uses 10 injections, the simplex would be composed of 31 treatment plans.

Each vertex of the simplex is a solution (remember that a solution is a treatment plan), and the position of the worst vertex is iteratively changed relative to the mean position of all of the other solutions.

The solutions are altered with a variety of functions: reflection, expansion, contraction and reduction. The schematic of NMS is shown in Figure 2.5.

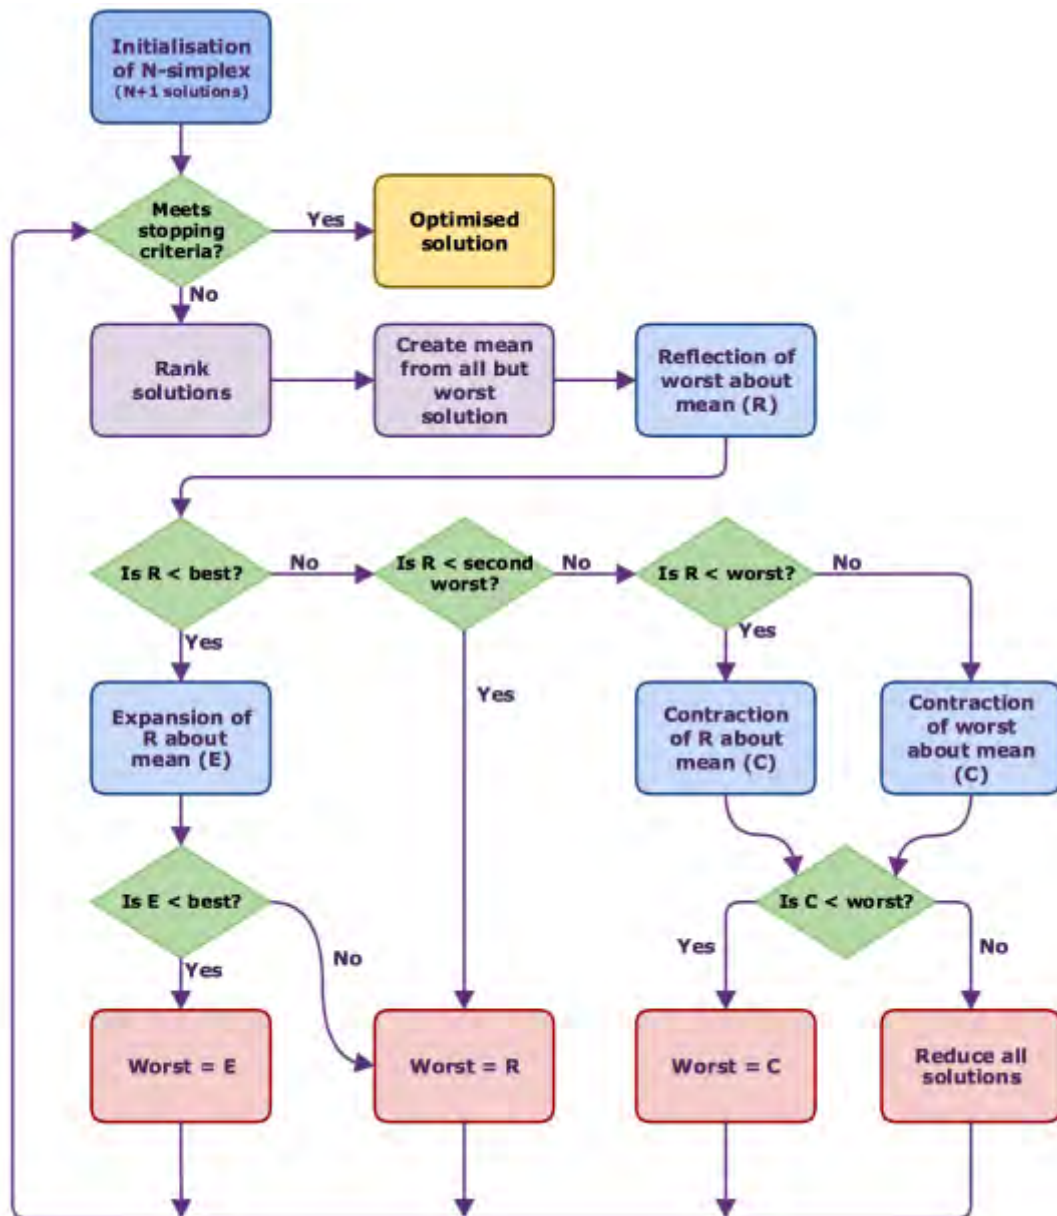


Figure 2.5: Schematic of NMS. “Is $R < best$?” is an abbreviation of “Is $cost(R) < cost(best)$?”. It can be seen that the algorithm tries to expand the simplex using reflections and expansions, but if these do not produce good quality results, the simplex shrinks with contractions and reductions.

The equations governing the reflection, expansion, contraction of the worst, contraction of the reflection and reduction of all solutions are shown below:

$$R = (1 + \alpha)\bar{x} - \alpha x_w \quad 2.10$$

$$E = (1 + \beta)R - \beta\bar{x} \quad 2.11$$

$$C_w = \left(1 - \frac{\gamma}{2}\right)x_w + \frac{\gamma}{2}\bar{x} \quad 2.12$$

$$C_R = \left(1 - \frac{\gamma}{2}\right)R + \frac{\gamma}{2}\bar{x} \quad 2.13$$

$$Red_i = \left(1 - \frac{\delta}{2}\right)x_i + \frac{\delta}{2}gb \quad \text{for all solutions} \quad 2.14$$

where x_w is the worst solution, \bar{x} is the mean of the remaining solutions and gb is the global best (the best solution found throughout the optimisation process so far). The parameters of reflection, expansion, contraction and reduction are α , β , γ , and δ , respectively.

An illustration of the possible movements is shown with a 1D example in Figure 2.6. Since there is only one degree of liberty in this problem, the N -simplex consists of just two vertices. The only movement not shown in Figure 2.6 is reduction, where all points move towards global best. In this example, α , β and γ are all equal to 1 (such that a reflection of 2 about 4 gives 6, an expansion of 6 from 4 gives 8, a contraction of 2 towards 4 gives 3 and a contraction of 6 towards 4 gives 5).




Figure 2.6: NMS movements (excluding the reduction of all solutions) for a 1D example.

Figure 2.7 shows a 2D example the movements encountered in NMS. The original simplex (consisting of 3 vertices) is shown in the top left. The mean, x_m , is calculated using all but the worst solution (x and x_b , but not x_w). Reduction is shown in the last illustration, where all solutions (x and x_w) move towards the best.




Figure 2.7: NMS movements for a 2D example. x_w , x_b , x_m , x_r , x_e , x_{cw} and x_{cr} , are the worst, best, mean (from all but the worst), reflection, expansion, contraction of worst and contraction of reflection solutions, respectively. All other solutions that are neither the best nor the worst are denoted with x .

From the two previous examples, it can be seen that the algorithm preferentially tries to increase the size of the simplex (via reflection and expansion). This means that the search-space (the area in which all solutions exist, i.e., all possible treatment plans) that is explored is as large as possible.

If the reflection and expansion of the worst solution do not produce a better solution, the algorithm shrinks the simplex (via contraction and reduction). This enables the algorithm to converge towards a minimum. Once the size of the simplex is sufficiently small, the algorithm stops. At this point, the solution is a minimum, but the minimum is not guaranteed to be a global minimum.

2.5.2.1 Initialisation

As mentioned in Section 2.3.4.3, deterministic algorithms will always return the same solution if the conditions of initialisation of the problem remain unchanged. Hence, the initialisation of a deterministic algorithm is extremely important.

One can imagine the simplex used in NMS as a triangle flip-flapping its way down a contour map (where height changes as a function of x and y), as shown in Figure 2.8. If the initial triangle is too small, it would get stuck in the first dip encountered (i.e., the solution converges prematurely into a local minimum). It is therefore imperative that the initial simplex covers the largest area possible to increase the chances of convergence towards a better solution. Even for the largest simplex possible, convergence to the global minimum is still not guaranteed (the solutions might straddle the minimum, not knowing that global minimum is between them).

Initialisation was done by randomly generating treatment plans. That is to say, the optimisation of a treatment plan that uses 10 injections was initialised by creating 31 treatment plans, each treatment plan containing the 10 injections randomly distributed within the PTV. This is not a very efficient method, since randomly generating treatment plans does not ensure a large coverage of the search-space (Simplex B in Figure 2.8).




Figure 2.8: Contour map with two simplexes. Simplex A risks converging to the local minimum just next to it, whereas Simplex B is much more likely (but still not guaranteed) to converge to the global minimum.

2.5.2.2 Drawbacks of NMS

There are three main drawbacks of NMS: its dependence on its initialisation, the form the function should have for NMS to work efficiently, and the fact that NMS can only optimise a fixed number of injections at any given time. The first point has already been discussed and the other two are discussed below.

NMS, unlike the other algorithms, is guaranteed to find the minimum of a convex cost function. Convex functions contain one minimum and no local minima. The cost function throughout this problem, however, is extremely unlikely to be convex. This is because there will be many treatment plans of similar quality, each one a local minimum. Spherical PTVs, for example, will have rotational symmetry, so many treatment plans will yield identical, optimal results.

The third drawback of NMS is that, unlike the other algorithms presented in this work, the number of injections is fixed for a given optimisation. This is because the N -simplex can be imagined as an N -dimension ball rolling down an N -dimensional slope. It is therefore not possible to dynamically change N throughout the optimisation. If for example, the number of injections was limited between 1 and 20 and the employed cost function was:

$$Cost = 0.8 \cdot Costlet_{PTV} + 0.2 \cdot Costlet_{inj} \quad 2.15$$

then one would have to repeat the optimisation process 20 times, once for each of the number of injections, to be able to find the optimal treatment plan. This renders NMS considerably slower.

2.5.3 Bees algorithm

BA (Pham *et al.*, 2006) is a stochastic SO algorithm. The algorithm is inspired by the food foraging behaviour of bees.

At any time, a bee colony has a given number of bees. It must therefore decide how to distribute those bees so as to maximise the quantity of foraged food. Scout bees search previously unsearched areas for new food sources. They then report back to the beehive and indicate to the forager bees (via the waggle dance) the quality of the food source that they have just discovered. The more intense the dance, the better the food source.

The forager bees will then search the flower patch immediately surrounding the scout's discovery. More forager bees are sent for more promising food sources. The bees continue to exploit the flower patch until the food runs out, at which point the forager bees will be reallocated to newly discovered flower patches.

The analogy was formed for a computer searching for optimal solutions to a problem. The computer has a finite processing speed, so the computer's resources need to be divided into discovering new areas of the search-space (exploratory behaviour) and searching the immediate vicinity of promising solutions (exploitive behaviour). The scout bees exhibit exploratory behaviour, whereas searching promising flower patches with forager bees demonstrates exploitive behaviour.

The schematics of BA are given in Figure 2.9 and Figure 2.10, followed by detailed explanations in the following sections.

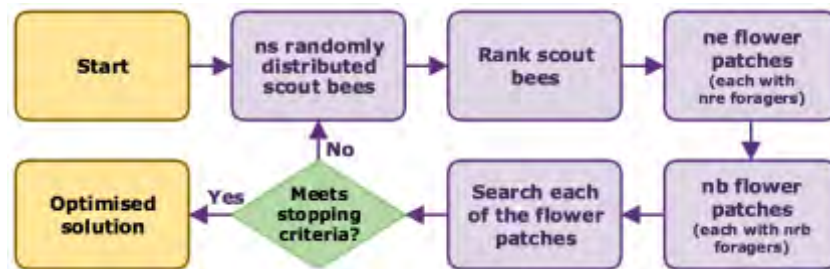


Figure 2.9: Schematic for BA. The scout bees are randomly distributed and ranked. The top-ranked scout bees become the centre of “elite” and “best” flower patches, around which local searches are performed. The box, titled “Search each of the flower patches”, is expanded upon in Figure 2.10.



Figure 2.10: Schematic of a flower patch search. The forager bees are randomly located about a flower patch centre and then ranked. If any of the foragers are better than the flower patch centre, they replace it as the centre and then the process is repeated. If none of the foragers are better than the flower patch centre, the size of the flower patch is decreased and then the foragers are redistributed. Once the flower patch size shrinks beneath a certain limit, it is assumed that the local minimum has been found. This local minimum is compared to the best minimum that has been found so far by the algorithm (the global best). If it is better than the global best, it replaces it. The algorithm then moves on to searching the next flower patch.

2.5.3.1 Initialisation

Initially, ns scout bees are distributed randomly throughout the search-space (ns treatment plans are created, each with a random number of randomly placed injections). The cost functions of these scout bees are evaluated, and the plans are ranked from best to worst.

2.5.3.2 Elite and best flower patches

Once the scout bees are ranked, the top-ranked ne scout bees are designated “elite” scout bees, and the subsequently top-ranked nb scout bees are designated “best” scout bees. The number of elite and best scout bees is chosen such that $ne + nb \leq ns$.

Figure 2.11 shows an example of the random allocation of 20 scout bees within a 2D search space, and the subsequent designation of 3 elite and 2 best scout bees.

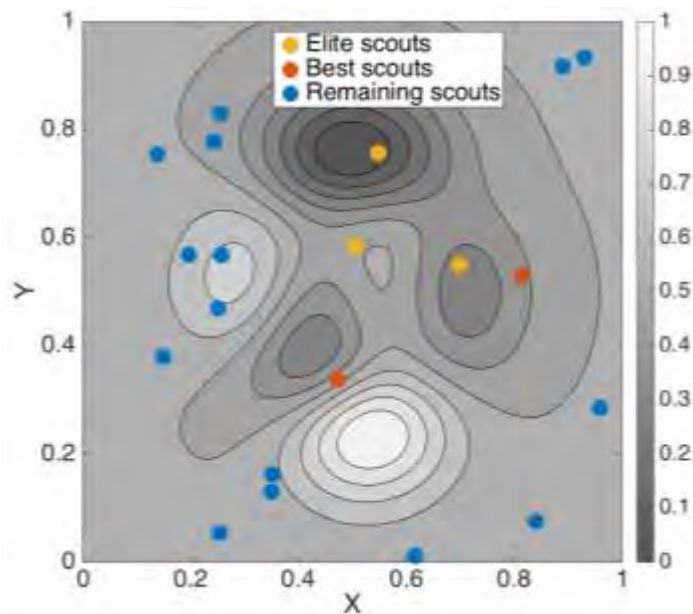


Figure 2.11: Random allocation of 20 scout bees in a 2D example. The top ranked scout bees become the centre of elite and best flower patches.

Flower patches are then created. The elite and best scouts become the centres of elite and best flower patches, respectively.

Each elite flower patch is assigned nre forager bees and each best flower patch is assigned nrb forager bees. The total population of the beehive is therefore:

$$population = ns + ne \cdot nre + nb \cdot nrb \quad 2.16$$

2.5.3.3 Flower patch search

The forager bees are randomly distributed within a user-defined distance, μ , from the flower patch centre, as shown in Figure 2.12. In the given example, nre and nrb are 20 and 10, respectively.




Figure 2.12: Random distribution of forager bees about flower patches. Elite and best flower patches have 20 and 10 forager bees, respectively (nre and nrb).

The fitness of each of the forager bees in a flower patch is evaluated. The forager bees are ranked, and if the cost function of the best forager bee is better than the flower patch centre (i.e., the cost function of the forager bee is lower than the flower patch), this forager bee becomes the centre of the flower patch. If not, the flower patch shrinks by a user-defined factor, $\delta\mu$.

The forager bees are then redistributed around the flower patch centre (which may or may not have changed) and this process continues iteratively until the size of the flower patch passes below a threshold. One can imagine a circle of decreasing radius, centring in on a local minimum. It is assumed that once the flower patch is small enough, then the local minimum has been found.

Once the local minimum has been found, this minimum is compared with the global best (the best treatment plan found by the algorithm in that optimisation run), and if it is better, then the minimum is saved as the global best. The global best treatment plan is the plan that will be proposed to the user at the end of the optimisation process.

The flower patch is then abandoned, and the process is repeated for the other flower patches. Once all flower patches have been explored (and the local minima have been found), the scout bees are redistributed.

2.5.3.4 Drawback of BA

A potential shortfall of BA is the number of user-defined parameters that are required before it can be used. On top of the parameters required for other algorithms (cost function weightings, etc.), users are also required to decide ns , ne , nb , nre , nrb , μ and $\delta\mu$.

A relatively large ns compared to ne , nb , nre and nrb results in more exploratory behaviour because more searches are performed in unknown areas. Similarly, a relatively large ne and nb results in exploratory behaviour because more computer resources are given to areas of the search-space even though they may not seem initially interesting. Conversely, a relatively large nre and nrb results in an increase of exploitive behaviour because more resources are spent on searching areas known to contain promising solutions. Lastly, a large μ and small $\delta\mu$ result in an increase of exploitive behaviour because more resources are spent on finding the flower patch's local minimum.

This drawback, however, can also be presented as an advantage; the liberty in selecting all the parameters means that it is possible to fine-tune BA as much as necessary.

2.5.4 Non-dominated sorting genetic algorithm-II

NSGA-II (Srinivas and Deb, 1994) is an evolutionary algorithm. It is an analogy, inspired by the idea of survival of the fittest, where the best animals survive, passing on their superior genes to future generations, whilst the genes of weaker individuals are phased out. Over many generations, the quality of the population therefore increases.

The schematic of NSGA-II is shown below, in Figure 2.13, and then explained in greater detail afterwards.



Figure 2.13: Schematic of NSGA-II. New populations are created by mixing solutions from the previous population. The new and previous populations are added together, sorted and the best half are kept, whilst the worst half are discarded.

Between animals, new offspring are created by mixing the genes of two parents (referred to in this work as crossovers). Mutations can also occur, encouraging diversity in the population. NSGA-II employs crossovers and mutations to iteratively produce generations of increasing quality.

In nature, the probability that certain characteristics survive to be passed onto future generations is a complicated mix of many factors (strength, speed, intelligence, etc.). In NSGA-II, this is mimicked by ranking the members of the population by taking into account multiple costlet functions simultaneously, through a process known as non-dominated fitness ranking.

Once all the treatment plans are ranked from best to worst, a new population is created using the previous generation's "genes". The new population and the old population are joined together, ranked and the worst half is discarded. This process, known as truncation replacement, ensures that the quality of the new generation is at least as good as the previous generation.

The concepts of the generation of the new population, the non-dominated fitness ranking and truncation replacement are developed further in the following sections.

2.5.4.1 Generation of the new population

The generation of the new population is done by selecting two "parents", mixing their "genes" to create two "children" and then potentially performing mutations on these children. Since this creates two children, the process is repeated $N/2$ to create a population of N treatment plans. The schematic for the production of a new population is shown in Figure 2.14.




Figure 2.14: Iterative process used by NSGA-II to create a new population (based on the previous population) via binary tournaments, crossovers and mutations. Binary tournament selection consists of randomly selecting two solutions from the previous population and selecting the best of the two to be a parent. This produces one parent, so the binary tournament selection is repeated a second time to get the second parent.

Tournament selection is used to choose parent solutions. The first parent is chosen by randomly choosing two solutions from the previous generation and keeping the better ranked solution as the parent. This is then repeated to obtain the second parent.

Crossovers are used to create two children from the two parents. The parents' genes (their injection coordinates) are split in half, and the first child inherits the first half of the first parent's genes and the second half of the second parent's genes. The second child is created with the second half of the first parent's genes and the first half of the second parent's genes. Hence, if the first parent uses 10 injections and the second parent uses 6 injections, the two children will have 8 injections. If the i^{th} injection of Parent 1 and Parent 2 is denoted as x_i and y_i , respectively, then their injections and the injections of their resulting children are expressed as:

$$Parent_1 = (x_1, x_2, x_3, x_4, x_5, x_6, x_7, x_8, x_9, x_{10}) \quad 2.17$$

$$Parent_2 = (y_1, y_2, y_3, y_4, y_5, y_6) \quad 2.18$$

$$Child_1 = (x_1, x_2, x_3, x_4, x_5, y_4, y_5, y_6) \quad 2.19$$

$$Child_2 = (y_1, y_2, y_3, x_6, x_7, x_8, x_9, x_{10}) \quad 2.20$$

Once two children have been created, random mutations are potentially applied, subject to a user-determined mutation probability. In the case of treatment planning, a mutation will either add a randomly generated injection, delete a randomly chosen injection or randomly alter any of the previously existing injections. This process increases diversity in the population, aiding the algorithm better explore the search-space.

The children, whether they have been mutated or not, are then added to the new population and the process is repeated until the new population is complete.

2.5.4.2 *Non-dominated fitness ranking*

As explained in Section 2.4.3, a treatment plan is considered non-dominated as long as no other treatment plan is at least as good it with regards to *all* costlet functions. A population can therefore be sorted into Pareto fronts.

The first Pareto front consists of all the non-dominated solutions. These solutions are then removed from the comparison, and the second Pareto front is formed from all of the newly non-dominated solutions. The second Pareto front will contain solutions that were only dominated by solutions in the first Pareto front. This process is repeated until all of the solutions have been assigned to a Pareto front.

For a population with N treatment plans, there can be 1 Pareto front containing N plans; N Pareto fronts, each containing 1 plan; or anywhere in between. The example that was shown in Figure 2.3 is shown again in Figure 2.15, with all solutions sorted into their respective Pareto fronts.




Figure 2.15: Solutions divided into their consecutive Pareto fronts. The second Pareto front is constructed by removing all solutions that form the first Pareto front and finding all the non-dominated solutions. These are then removed and the process is repeated for the third Pareto front and so on.

Within a Pareto front, no single solution can be said to be better than any other; they are a set of compromises. In the case of treatment planning, the compromise might between the coverage of the PTV and the number of injections – increasing the number of injections (which could be $Cost_1$) would decrease the under-irradiated portion of the PTV (decreasing $Cost_2$) and vice versa.

It is necessary, however, to not only sort the treatment plans into Pareto fronts, but also to rank each of the treatment plans within each Pareto front. In the process of non-dominated fitness ranking, this is done with crowd distances.

Crowding distances are used to favour more unique solutions. By doing so, solutions in sparsely populated areas of the Pareto front are better classed than those in crowded areas, and the future generation will be more evenly spread out.

In the example shown in Figure 2.15, for example, solutions A and G are given the largest crowding distances. They lie on the extremities of the Pareto front, so favouring them will encourage the Pareto front to spread out as far as possible. The crowding distances of the remaining j solutions are calculated using Equation 2.21:

$$Crowd(j) = \sum_i^n \frac{Cost_i^{j+1} - Cost_i^{j-1}}{Cost_i^{max} - Cost_i^{min}} \quad 2.21$$

where n is the number of costlets (2 for the 2D example shown in Figure 2.15), $Cost_i^{j+1}$ and $Cost_i^{j-1}$ are the costs for i^{th} costlets of the neighbouring $(j + 1)^{\text{th}}$ and $(j - 1)^{\text{th}}$ solutions, respectively, and $Cost_i^{max}$ and $Cost_i^{min}$ are the maximum and minimum costs of the i^{th} costlet, respectively. The crowding distance of treatment plan C, for example, is shown in Equation 2.22:

$$Crowd(C) = \frac{D_1 - B_1}{G_1 - A_1} + \frac{B_2 - D_2}{A_2 - G_2} \quad 2.22$$

Once all of the solutions have been sorted into Pareto fronts and the crowding distances have been calculated within each Pareto front, the whole of the population is ranked such that the best solutions have the smallest Pareto front and, for matching Pareto fronts, the largest crowding distances.

Performing the ranking in this manner means that the algorithm can optimise multiple costlet functions without having to combine them using weightings, as is the case for SO algorithms. Once the population has been sorted, the binary selection process can be used to select parents based on their ranking.

2.5.4.3 Truncation replacement

Good optimising algorithms should exhibit a quality known as elitism. Elitism is the idea that a good solution should never be lost. It is possible, for example, that all of the children in a new population are worse than the best solution of a previous population. To avoid losing the best solutions of the previous population, NSGA-II uses a method called truncation replacement, which combines the old and new populations together. The solutions are then ranked, and only the best N treatment plans are kept.

In this fashion, the size of the population remains constant, but the best solutions always survive. This ensures that the quality of the new population is always at least as good as, if not better than, the previous population.

2.5.4.4 *NSGA-II as an SO algorithm*

Although NSGA-II is an MO algorithm, it can also be used as an SO algorithm. By simply applying a single cost function instead of multiple cost functions, the Pareto front becomes 1D. This means that the treatment plan with the lowest cost will dominate all the others, and so will be the only plan in the first Pareto front. Subsequently, the plan with the second lowest cost will dominate all remaining plans, and so will be in the second Pareto front.

In this sense, for a population of N treatment plans there will be N Pareto fronts, each containing just 1 solution. When using NSGA-II as an SO algorithm, calculating the crowding distance becomes unnecessary since the crowding distance was used to rank solutions within the same Pareto front. The solutions are simply ranked by their Pareto front (which is linked directly to their single cost) and the parent selection is then performed as per normal.

2.5.4.5 *Drawback of NSGA-II*

Since NSGA-II mimics evolution, no explicit searching is performed in the immediate vicinity of promising solutions (unlike NMS and BA).

NMS performs a local search by reducing all solutions towards the best. Similarly, BA performs local searches by using forager bees to search around flower patches, moving towards the local minima.

NSGA-II does not perform a local-area search and for this reason, although it performs global searches very efficiently, it is known to converge slowly when in the neighbourhood of promising solutions.

2.6 Stopping criteria

Now that the forward problem, cost function and optimisation algorithms have been introduced, the only box in Figure 2.1 that remains to be discussed is the stopping criteria. Each time an iteration is completed, the algorithm checks to see if any of the criteria have been met, signalling the end of the optimisation process. Various examples of stopping criteria are given in the following subsections.

2.6.1 Convergence

NMS will always converge to a point. When no better solutions have been found, all the treatment plans move towards the best solution (via a reduction) and the size of the N -simplex decreases. Eventually, all of the vertices of the simplex converge to the same point. Since all vertices now match, NMS cannot do any more movements, and so the optimisation process stops. This was therefore used as a stopping criterion for NMS.

Convergence could not, however, be used as a stopping criterion for BA or NSGA-II because their stochastic nature means that they do not converge. For example, BA can continually produce new scout bees and NSGA-II can continually produce new children (although this does not mean that they will continually find better solutions). Hence, other stopping criteria must be found for these methods.

2.6.2 Number of iterations

After considering convergence, designating a maximum number of iterations might seem to be a logical choice of stopping criterion. However, each algorithm (even each configuration of a given algorithm) alters a different number of treatment plans per iteration. NSGA-II with a population of 1,000, for example, will alter 10 times more treatment plans than NSGA-II with a population of 100 per iteration.

Therefore, to use the number of iterations as a stopping criterion does not seem a very apt choice, especially when a comparison is to be performed between the various algorithms later in this work.

2.6.3 Treatment plan updates

Instead of imposing a maximum number of iterations, a maximum number of treatment plan updates was used. That is to say, each time a treatment plan was altered, a counter was incremented. When the counter reached the user-designated maximum, the algorithm stopped. This method was therefore used for all three algorithms throughout this research.

2.7 Conclusion

In this chapter, many key-concepts were presented. In order to find optimal treatment plans, inverse-problem solving algorithms incorporate:

- The forward problem (the calculation of the absorbed dose distribution from a given configuration of injections),
- The cost function (a method of quantifying treatment plans), and
- The optimisation algorithm (a method of altering the treatment plans so as to minimise the cost function).

Both physical and radiobiological cost functions were presented, the latter being ruled out as the relevant data is not yet available for microbrachytherapy. Physical cost functions that are typically used in EBRT were also ruled due to the inherently inhomogeneous radiation distribution encountered in microbrachytherapy.

Finally, costlet functions were created for the PTV, the number of injections, the OAR and the CI. $Cost_{PTV}$ penalised the fraction of the PTV receiving less than 95% of D_T and $Cost_{OAR}$ penalised the fraction of the OAR receiving over 10 Gy. $Cost_{Inj}$ and $Cost_{CI}$ are simply equal to the number of injections and the conformity index, respectively. When the cost functions have been minimised, the treatment plan is optimised.

SO and MO algorithms were also introduced. SO algorithms combine multiple costlet functions into a single cost function with relative weightings between each costlet. MO algorithms, on the other hand, optimise the costlets individually, and the results are presented in the form of a Pareto front.

SO algorithms are advantageous in that the dimensionality of the search-space is reduced from ND to 1D (where N is the number of costlets). This makes the optimisation process much faster. However, if the user then wishes to alter the costlet weightings, they must restart the optimisation process anew. At this point, the speed advantage of SO algorithms would be lost. MO algorithms are slower but more powerful, proposing all optimal solutions, and allowing the user to make the most informed decision possible.

Three algorithms were detailed: NMS, BA and NSGA-II. A summary of these three algorithms is shown in Table 2.3. These three were chosen as they present a wide range of optimisation algorithms, allowing for an interesting comparison in Chapter 4, when the algorithm that is the best-suited for microbrachytherapy treatment planning will be selected. Before that comparison can be performed, however, the forward problem – microbrachytherapy dosimetry – must be discussed in greater detail. This is therefore the subject of the following chapter.

Algorithm	Stochastic/deterministic	SO/MO
NMS	Deterministic	SO
BA	Stochastic	SO
NSGA-II	Stochastic	MO

Table 2.3: Summary of three algorithms that have been presented in this chapter and will be compared in Chapter 4.

3 MICROBRACHYTHERAPY DOSIMETRY

3.1 Introduction

Microbrachytherapy is a new form of radiotherapy. As such, there is no currently established method of dosimetry (for calculating and assessing the irradiation of a patient) for this modality.

In this chapter, the dosimetric methodology that has been developed, and will be used throughout the rest of this research, is detailed.

A treatment plan contains information pertaining to the proposed treatment. For microbrachytherapy, treatment plans will state how many injections will be used and the coordinates of those injections.

From a given treatment plan, it is necessary to then calculate the resulting absorbed dose distribution to be able to evaluate the dosimetry of the PTV, OARs and other healthy tissue in order to evaluate the efficacy of the treatment plan.

The workflow for calculating the dosimetry of a microbrachytherapy treatment plan is shown in Figure 3.1 and detailed below:

- Calculate the absorbed dose distribution for a single injection (known as a kernel),
- Calculate the plan's absorbed dose distribution (also referred to as the dose grid) by convolving the kernel in each of the desired positions corresponding to the treatment plan's coordinates, and
- Superpose binary masks on the dose grid to calculate the distribution in each structure (PTV or OAR).

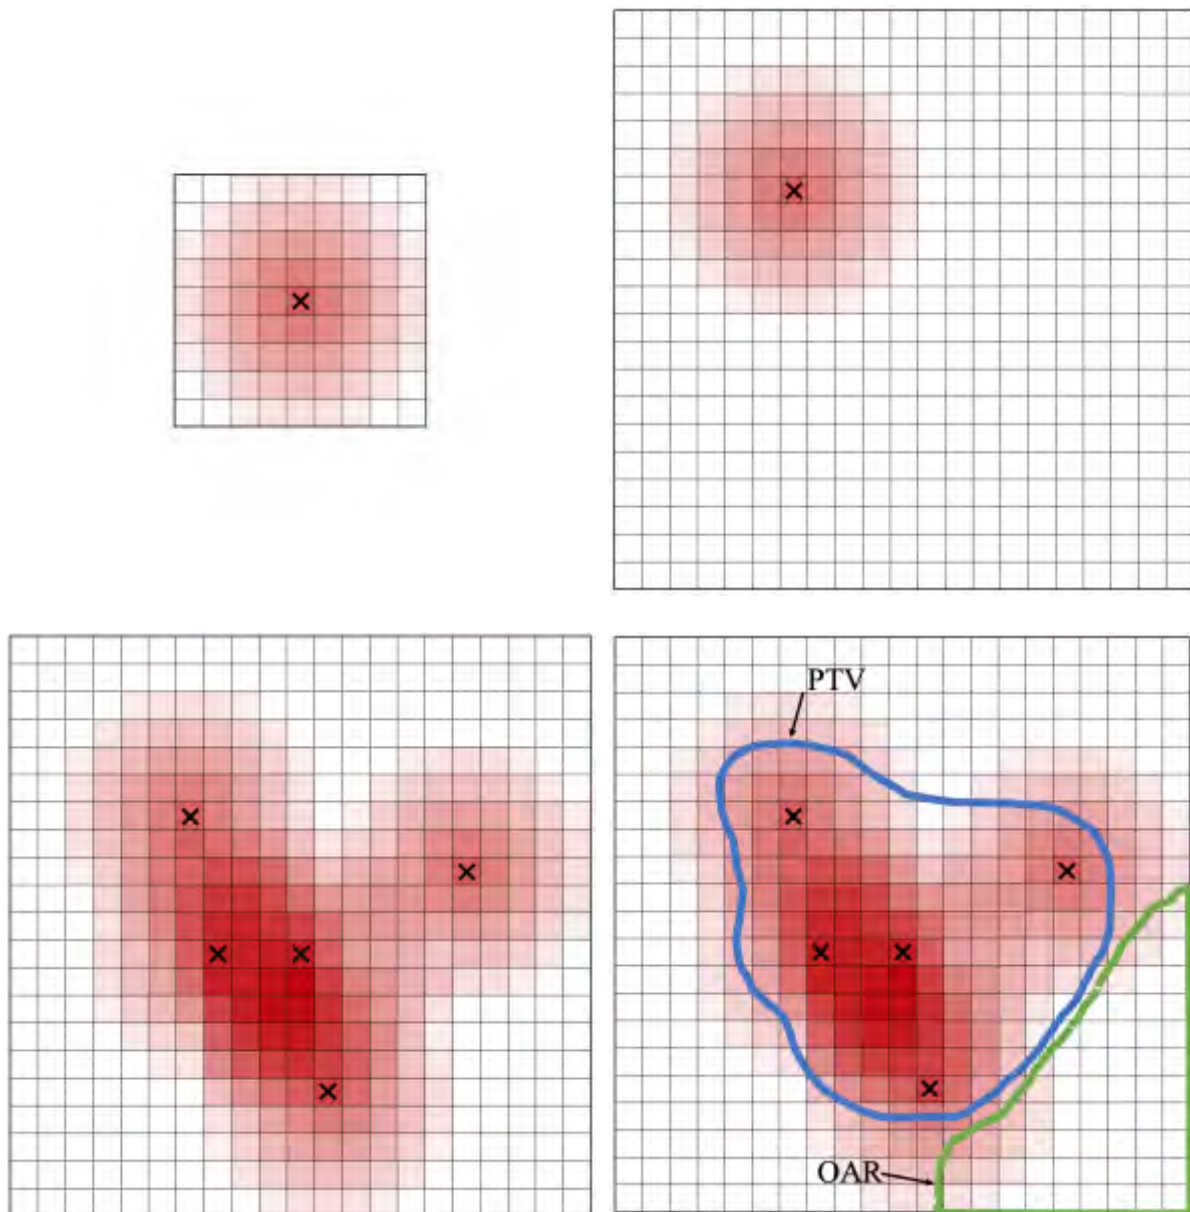


Figure 3.1: Illustration of the absorbed dose distribution for (a) an injection, (b) the injection placed on the dose grid, (c) multiple injections convolved on the dose grid as dictated by a treatment plan and (d) structures superposed to calculate structure-specific absorbed dose distributions.

The entirety of the dosimetric workflow that is presented in this research was programmed from scratch in C++. The code is discussed in greater detail in the Appendix.

The structure of this chapter is as follows:

- Section 3.2 discusses the generation of the kernel,
- Section 3.3 discusses convolving the kernel to calculate the patient's dose grid,
- Section 3.4 presents a methodology for structure-specific dosimetry,
- Section 3.5 explores certain assumptions and considerations that are taken into account during the dosimetry,
- Section 3.6 presents a worked example of the complete dosimetric workflow that will be used throughout this research, and
- Section 3.7 concludes this chapter.

3.2 Kernel generation

A simplified version of the MIRD formulation of the absorbed dose, \bar{D} , which is a function of the S-value, S , and the cumulated activity, \tilde{A} , is given below:

$$\bar{D} = S \cdot \tilde{A} \quad 3.1$$

Throughout this work, the S-value is used to refer to the mean absorbed dose deposited in a voxel for a single radioactive decay in the injected volume. The cumulated activity, is the integral of the injection's activity with respect to time:

$$\tilde{A} = \int_0^T A(t) dt \quad 3.2$$

As mentioned in Section 1.3, by incorporating the radioelements into microspheres, the biological half-life is quasi-infinite. Because the radioelement does not move, the integral can be completed to infinity (provided the radioelement's half-life is much less than the patient's life time, which is the case for all radioelements tested in this work). In this case, the cumulated activity is a function of the injection's initial activity, A_0 , and the radioelements half-life, $T_{1/2}$:

$$\tilde{A} = \int_0^{\infty} A(t) dt = \frac{T_{1/2}}{\ln(2)} \cdot A_0 \quad 3.3$$

3.2.1 S-values

The calculation of \tilde{A} is therefore a straightforward task, and all that remains to be able to calculate the absorbed dose distribution for an injection is the voxelised distribution of S-values. Once the S-value for each voxel in the kernel has been calculated, it is trivial to select an A_0 and instantaneously obtain the voxelised absorbed dose distribution for an injection.

S-values were calculated using Monte Carlo code. Monte Carlo codes exploit the law of large numbers: to take the example of radioactive decay, the principal is that, although it is impossible to predict the resulting energy deposition of an individual radioactive decay (since radioactive decay is a random process), it is possible to predict the distribution for a large number of radioactive decays.

3.2.2 GATE

The Monte Carlo code used to calculate the kernels was GATE 7.2 (Jan *et al.*, 2004; Sarrut *et al.*, 2014). GATE is an open-source, community-based collaboration for numerical simulations of medical imaging and radiotherapy, that is based on the simulation toolkit, Geant4 (Agostinelli *et al.*, 2003).

The results from GATE were expressed as Gy for the number of particles simulated. To convert these into S-values (which, to reiterate, are the average energy per radioactive decay), it is necessary to divide the voxel values by the number of simulated particles and multiply by the number of particles generated by a single decay of the simulated radioelement. The radioelement ^{166}Ho , for example, emits 8.19 particles per decay (of which 1.00 are β^- , 3.74 are γ and 3.45 are e^-).

3.2.3 Modelling the injection

Many options are available when modelling a radioactive source. The source could, for example, be considered a point source, and the absorbed dose could be calculated in a voxelised grid surrounding this point source. Such a kernel is referred to as a dose-point kernel (DPK) (Ahnesjö, Andreo and Brahme, 1987). Alternatively, the radioactive source could occupy a whole voxel, in which case the resulting distribution would be referred to as a voxel-dose kernel (VDK) (Dieudonné *et al.*, 2011).

The assumption was made in Section 1.3 that the injections would form spheres, containing a homogeneous distribution of radioelements. Hence, the radioactive source was modelled as an isotropic, spherical source in a voxelised grid. The centre of the sphere was placed in the centre of the central voxel of the grid.

It should be noted that the proposed setup is not the same as either a DPK or a VDK. As will be described in greater detail in Section 3.5.1, it was assumed that the irradiated volume (the kernel surrounding the radioactive source) was homogeneous soft tissue with a density of 1.04 g/cm^3 .

3.2.4 GATE parameters

Precision of kernels generated with GATE depend on five major factors:

- The number of simulated particles,
- The volume simulated,
- The voxel size,
- The chosen GATE physics list and energy cuts, and
- The spectra of the simulated radioelement.

The first three are closely related. For a given simulated volume and voxel size, increasing the number of simulated particles results in more precise kernels but longer simulation times. For a given number of particles and voxel size, increasing the simulated volume increases statistical uncertainties (because the solid angle of voxels far from the source is small, leading to relatively few particle interactions). Lastly, for a given number of particles and simulated volume, decreasing voxel size increases statistical uncertainties (since fewer particle interactions take place in each voxel). Hence, if one wishes to decrease voxel size, for example, one should increase the number of simulated particles accordingly.

Two voxel sizes were used in this work, 0.1^3 mm^3 and 0.5^3 mm^3 , which were simulated with 10^{10} and 10^9 particles, respectively. For the vast majority of this work, kernels with voxels of 0.5^3 mm^3 are used. The only exception is in Section 4.7, when a study is conducted into the effect of voxel size on the dosimetry.

The spherical source was placed in the centre of the central voxel of the kernel. It was empirically determined that a total kernel size of 24.5 mm was sufficiently large to contain voxels containing non-negligible absorbed doses, where absorbed doses beneath 0.01 Gy were considered negligible.

Hence, the 0.5^3 mm^3 kernels consisted of 49 voxels in each direction, corresponding to 24 voxels, then the injection centred on the 25th voxel, and then 24 more voxels. In total 0.5^3 mm^3 kernels therefore consisted of 117,649 voxels.

The kernels can be relatively small because the radioelements that are used throughout this work are β^- -emitters, which have a small interaction path and therefore irradiating a small volume. This will be discussed in greater detail in Chapter 5, in which a study of radioelements is conducted.

The physics list that was used was “emstandard_opt3”, which is considered the standard for radiation therapy applications with GATE. The physics lists controls probabilities of physics process, such as Compton scattering, photoelectric absorption, ionisation and Bremsstrahlung.

The following parameters were set to 1/10th of the voxel size (corresponding to 0.05 mm for kernels with voxel sizes of 0.5^3 mm^3):

- SetCutInRegion for gammas and electrons: when the range of the particle drops beneath this threshold, no more secondary particles will be created for charged particle processes, such as ionisation and Bremsstrahlung.
- SetMaxStepSizeInRegion for electrons: the maximum step size of a particle in a given region.
- setStepFunction for electrons: the step limit must be greater than this.

These parameters are described in greater detail in (Agostinelli *et al.*, 2003).

Lastly, emission spectra for the radioelements were taken from MIRD (Eckerman and Endo, 1989). All possible emitted particles were simulated: mono-energetic photons, (X-rays and γ -rays), β^- -minus, Auger electrons and internal conversion.

3.2.5 Simulation times

Shown below, in Table 3.1, is a summary of all kernels used throughout this research as well as the section in which they are used. As can be seen, simulation times vary enormously (between 8.6 hours and 39 days). A cluster of computers, consisting of 80 cores, was exploited to minimise simulation times as much as possible.

Radioelement	Injection volume (μL)	Voxel size (mm^3)	# particles simulated	Simulation time (d)	Used in Section
^{90}Y	5	0.5^3	10^9	10.13	Throughout
^{166}Ho	5	0.5^3	10^9	0.98	5.2
^{131}I	5	0.5^3	10^9	0.50	5.2
^{177}Lu	5	0.5^3	10^9	0.36	5.2
^{166}Ho	5	0.1^3	10^{10}	39.03	4.7
^{90}Y	10	0.5^3	10^9	10.09	5.3.2
^{90}Y	20	0.5^3	10^9	10.06	5.3.2

Table 3.1: Simulation times for kernels used throughout this work. Simulation times are shown for a single core of a 2.5 GHz MacBook Pro. The 0.1^3 mm^3 kernel of ^{166}Ho was simulated with 10^{10} particles because the decreased kernel size requires an increased number of primary particles, as described in Section 3.2.4.

3.3 Patient absorbed dose distributions

As was illustrated in Figure 3.1, calculating the resulting absorbed dose distribution for a treatment plan is as simple as convolving the injection kernel as many times as necessary in the positions dictated by the plan. This assumes that the irradiated volume is comprised exclusively of a homogeneous medium. This assumption is one of the considerations that is discussed in Section 3.5.

Injections were only allowed to be placed in the centre of voxels (since the kernels were calculated with the injections placed in this fashion). This means that treatment plans were optimised to the same precision as the kernel. That is to say, a kernel discretised with voxels of 0.5^3 mm^3 will result in a treatment plan with a precision of 0.5^3 mm^3 .

In this work, treatment plans could only propose to place injections in the PTV. If all the injections are confined to the PTV, then the smallest cube required to ensure that all of the energy emitted by a treatment plan would be contained within the dose grid is the cube surrounding the PTV plus a margin of 12.5 mm (the distance from the centre of an injection to the edge of the kernel, as discussed in Section 3.2.4). This is illustrated in Figure 3.2.




Figure 3.2: Illustration of the dose grid. Four injections and their associated kernels are shown at the edges of the PTV. It can be seen that the blue box is therefore the smallest possible size of the dose grid that is sure to contain the entirety of the radiation of all proposed treatment plans.

Lastly, it should be noted that since the kernel is convolved on the dose grid, the voxel size of the dose grid must match the voxel size of the kernel. If one wished to perform the optimisation with a dose grid with voxel sizes that do not match any pre-calculated kernels, it would of course be possible to resample the kernel prior to the start of the optimisation process. A resampled kernel would be subject to partial volume effects and so it would be preferable to produce kernels with voxel sizes that match the sizes that will be used in the treatment planning process. The effects of the kernel size on the dosimetry are the topic of discussion in Section 4.7.

3.4 Structure-specific dosimetry

Once the dose grid had been calculated, structure-specific dosimetry was performed by superposing structure masks over the dose grid. Masks are binary arrays that can be overlaid on the dose grid, revealing the absorbed dose distribution in a given structure. The masks were created using the C++ class, “`vtkStencillImage`”. The use of the visualisation toolkit (VTK) (Schroeder, Martin and Lorensen, 1996) is further explored in the Appendix.

The results of structure-specific dosimetry are often displayed in the form of cumulative DVHs. Such DVHs show the fractional volume of the structure receiving greater than or equal to a given absorbed dose.

Once structure-specific dosimetry has been performed, it is possible to evaluate dosimetric requirements. This will be further expanded upon during the worked example in Section 3.6.

3.5 Considerations

Below, in Sections 3.5.1 through 3.5.3, three aspects that should be taken into consideration are discussed. These aspects are the assumption of homogeneity, the attenuation of the microsphere's casing and calculation times.

3.5.1 Assumption of homogeneity

For the calculation of treatment plans' absorbed dose distributions, a crucial assumption was made: the entirety of the irradiated volume was a homogeneous medium. Specifically, it was assumed that the medium was soft tissue with a density of 1.04 g/cm^3 .

For the treatment of glioblastomas, this assumption was considered reasonable since density in the brain, generally speaking, is extremely homogeneous compared to abdominal tumours, where inhomogeneities due to air and bone are present. In the brain, as long as the tumour is more than a few centimetres from the skull, considering the dose grid to be composed entirely of soft tissue is a valid assumption.

This assumption would not be applicable if inhomogeneities such as air and/or bone were present, as is the case for the treatment of abdominal tumours. To be able to perform treatment planning in inhomogeneous media, it would be necessary to calculate and employ inhomogeneous kernels.

The calculation of inhomogeneous kernels would not be possible with Monte Carlo simulations as it would be far too time-consuming (Table 3.1 showed simulation times to the order of magnitude of days, whereas treatment planning needs to be completed in hours, at most).

The calculation of inhomogeneous kernels could therefore be done by using the homogeneous kernels that have already been calculated and distorting them as a function of the density of the voxels that are irradiated. The voxel densities would be used to calculate the radiologic distance (Siddon, 1985) – the equivalent distance travelled through water with regards to density (1.00 g/cm^3) – as shown in Figure 3.3. The radiological distance would then be used to modify the kernel as suggested in (Ahnesjö, Andreo and Brahme, 1987).




Figure 3.3: Illustration of the radiologic distance. The radiological distance is the equivalent distance traversed through water with regards to density. The line corresponding to air is not to scale; the density of air is so low that the radiological distance through air is very large compared to water and bone.

To be able to perform this method, it would be necessary to calculate the radiologic distance between a source voxel (i.e., the centre of an injection) and all irradiated voxels. Either an analytic function or a lookup table can then be consulted to convert the radiological distance to an absorbed dose. This would then have to be repeated for each injection. It can be seen that this method is extremely time-consuming.

To reiterate, this method of generating kernels for use in inhomogeneous media was not used because the particular problem happened to be in a homogeneous medium. If one wished to expand upon this treatment to perform dosimetric calculations in inhomogeneous media (such as for the treatment of abdominal tumours), these are considerations that would have to be taken into account.

3.5.2 Radioelement distribution and microsphere casing

The two points that are discussed in this subsection are related to unknown parameters of the microsphere. Firstly, the distribution of the radioelements throughout the injection is unknown. Secondly, the composition and structure of the microspheres are also unknown.

3.5.2.1 Distribution of radioelements

The first uncertainty is related to the distribution of radioelements throughout the injection. As was mentioned in Section 1.3.1, it was assumed that after the injection, the radioelements would be distributed homogeneously. In reality, this would not be the case because radioelements would be grouped tightly together in microspheres.

Section 5.3.2 conducts a study into the effect of the injected volume and finds that there is no noticeable difference between injections of 5 and 20 μL with the same number of radioelements.

Therefore, if spreading the radioelements over a volume that is four times bigger makes no noticeable difference, it is safe to assume that the inhomogeneous distribution of radioelements due to them being grouped together in microspheres is also negligible.

3.5.2.2 *Microsphere casing*

The second unknown parameter is the composition and structure of the microsphere, with particular focus on the microsphere casing. The radioelements are inside of the microsphere and therefore to be able to irradiate the surrounding tissue, the radiation must first escape the microsphere casing.

Microsphere casings can be composed of a wide range of materials. TheraSpheres and SIR-spheres are two examples of microspheres that are used in radiotherapy. These microspheres contain ^{90}Y and are used for SIRT of the liver (as described in Section 1.2.3). The casing of TheraSpheres is insoluble glass (Salem *et al.*, 2002) and SIR-spheres are composed of resin (Lim *et al.*, 2005). Microspheres of ^{166}Ho have also been developed from low-molecular-weight polymer microspheres (Zielhuis *et al.*, 2006).

Microsphere casing was not taken into account for two reasons. Firstly, specific composition of the microspheres is typically private property of the manufacturer. Secondly, even if the composition were known, this research aims to stay as general as possible and not create models for specific microsphere setups.

3.5.3 Calculation times

Each time a treatment plan's injections are altered, it is necessary to update the patient dose grid and then recuperate structure-specific dosimetric information (such as the DVH). This process must be extremely fast (to the order of milliseconds), as algorithms can change many treatment plans each iteration.

Attempts were made to speed up this process by minimising the number of voxels and by exploiting the graphics processing unit (GPU), as discussed in the two following sections.

With the following considerations, structure-specific dosimetry was performed in 4 ms for a treatment plan with 20 injections and the spherical PTV OARs, both with a radius of 10 mm. The process consisted of the convolution of the 20 injections on the dose grid, the superposition of the structures' binary masks and the extraction of their DVHs.

3.5.3.1 The number of voxels

The size of voxels has already been discussed in Section 3.2.4 and will be further explored in Chapter 5. Considering both the size of the kernel and the size of the voxels to be fixed, a potential method of reducing the number of voxels involved in the convolution of the kernel is by disregarding voxels containing negligible absorbed doses.

The injection's absorbed dose distribution is spherical, yet the kernel is cubic. The cube must be big enough so that along the smallest distance from the centre of the kernel to the edge (shown as A in Figure 3.4), all non-negligible absorbed dose is accounted for. This means that along the hypotenuses (shown as B), the absorbed dose must be considerably lower and there must therefore be volumes (shown as C) that contain negligible absorbed doses.




Figure 3.4: 2D illustration of a cubic kernel containing a spherical distribution. A is the minimum distance for all non-negligible absorbed doses to be recorded. B is the maximum distance from the injection centre to the edge of the kernel ($\sqrt{2}A$ in 2D and $\sqrt{3}A$ in 3D). Hence, all the voxels between A and B from the injection centre, shown by the area C (which is a volume in 3D), must contain negligible absorbed doses.

To use the 5 μL injection of ^{90}Y with an A_0 of 5 MBq as an example, the mean absorbed dose at the centre of the six faces of the cube (a distance of A) was 0.00340 Gy. At the eight vertices (a distance of B), however, the mean absorbed dose was 0.000840 Gy.

To circumvent the problem of using a cubic kernel to quantify a spherical absorbed dose distribution, a user-defined threshold was used to establish non-negligible voxels. The indices of all voxels receiving greater than this threshold (for example, 0.1 Gy) are stored in an array and only these voxels are convolved on the dose grid.

This means that the number of voxels was decreased by at least the ratio between a sphere and a cube, $\pi/6$, almost a factor of 2. This factor rises as the threshold is increased and more voxels are considered to contain negligible absorbed doses. The percentage of voxels that are included in the dosimetric calculation as a function of the user-defined threshold is given below, in Table 3.2.

Threshold (Gy)	Percentage of voxels included (%)
0	100
0.01	36.2
0.1	27.2
1	18.5

Table 3.2: The reduction of the number of voxels included in the convolution of the kernel as a function of the user-defined threshold that determines if voxels contain non-negligible absorbed doses.

3.5.3.2 GPU

Modern computing increasingly exploits the Graphical Processing Unit (GPU). In computing, GPUs were initially used exclusively for graphical rendering, such as rendering pixels on the screen. Rendering each pixel is a relatively fast process, but screens are composed of many pixels. Hence, the core of a central processing unit (CPU) is more powerful than the core of a GPU (~ 2.5 GHz compared to ~ 1 GHz), but GPUs contain many more cores than CPUs ($\sim 1,000$ compared to ~ 4). GPUs are therefore being used increasingly in modern computing for problems that can be easily fractionated, taking advantage of the hyper-threaded GPU.

Similarly, calculating the absorbed dose in each voxel in a patient is a repetitive task and one that is easily parallelised. It was therefore hoped that the GPU could be used to accelerate this calculation. CUDA was used throughout for this purpose. CUDA is an application programming interface that was developed by Nvidia, enabling general purpose GPU programming (Nickolls *et al.*, 2008).

This method was unfortunately deemed infeasible because data must always be copied from the CPU to the GPU so that the calculation can be performed, and then copied back from the GPU to the CPU. For example, if an algorithm alters 100 treatment plans per iteration, the dose grid of each plan would have to be copied over individually to the GPU, the calculation would be performed and then the results would be copied back to the CPU. This copying to and fro is required because the memory of the GPU is considerably smaller than that of the CPU and storing hundreds of dose grids on the GPU is not possible. This copying process was time-consuming, negating the increased calculation speeds.

3.6 A worked example

Shown below, is a worked example from start to finish of a dosimetric calculation. The PTV and OAR were the 10 mm-radius spheres that were described in Section 1.7. The treatment plan uses 10 injections of ^{90}Y . Each injection has a volume of 5 μL and an A_0 of 5 MBq.

Firstly, the kernel was generated with GATE. As was discussed, the results of GATE must be divided by the number of particles that were simulated (10^9 for this kernel) and then multiplied by the number of particles that are emitted on average per decay (1.00 for ^{90}Y).

Once the kernel had been generated, all of the voxels were multiplied by the cumulated activity, using the A_0 of 5 MBq and $T_{1/2}$ of 2.7 days (2.33×10^5 s). This meant that the kernel was now expressed in the form of absorbed dose. Figure 3.5 shows a 1D projection of the kernel, with the absorbed dose shown on the left axis and the corresponding S-value shown on the right axis.




Figure 3.5: 1D projection of the absorbed dose (and S-value) distribution for a 5 µL injection of ^{90}Y with an A_0 of 5 MBq. The limits of the 5 µL injection are also shown.

Figure 3.6 shows some slices through the kernel. Absorbed doses greater than or equal to 100 Gy are shown as white. A threshold was used for absorbed doses greater than this because, as can be seen in Figure 3.5, absorbed doses close to the centre of the injection are extremely elevated.




Figure 3.6: Slices through the absorbed dose distribution of the kernel of a 5 µL injection of ^{90}Y with an A_0 of 5 MBq. Slice distance from the centre of the injection is given beneath the distributions and are expressed in mm. All absorbed doses greater than or equal to 100 Gy are shown as white.

An example of the calculation of a dose grid is shown in the two following figures. Figure 3.7 shows 10 injections randomly distributed within the PTV, which in turn is placed next to the OAR. The limits of the dose grid are shown with the white cube. By convolving the kernel shown in Figure 3.6 at each of the 10 injection positions, the dose grid was calculated. Slices of this dose grid are shown in Figure 3.8.



Figure 3.7: 10 injections randomly placed within the PTV which is located next to an OAR. The limits of the dose grid can be seen as the edges of the white cube (a margin of 12.5 mm about the PTV).




Figure 3.8: Slices through the dose grid. The edges of the PTV and OAR are shown in green and cyan, respectively. Absorbed doses greater than or equal to 100 Gy are shown in white.

DVHs of the PTV and OAR were calculated for the absorbed dose distribution shown in Figure 3.8 and these DVHs are shown below.

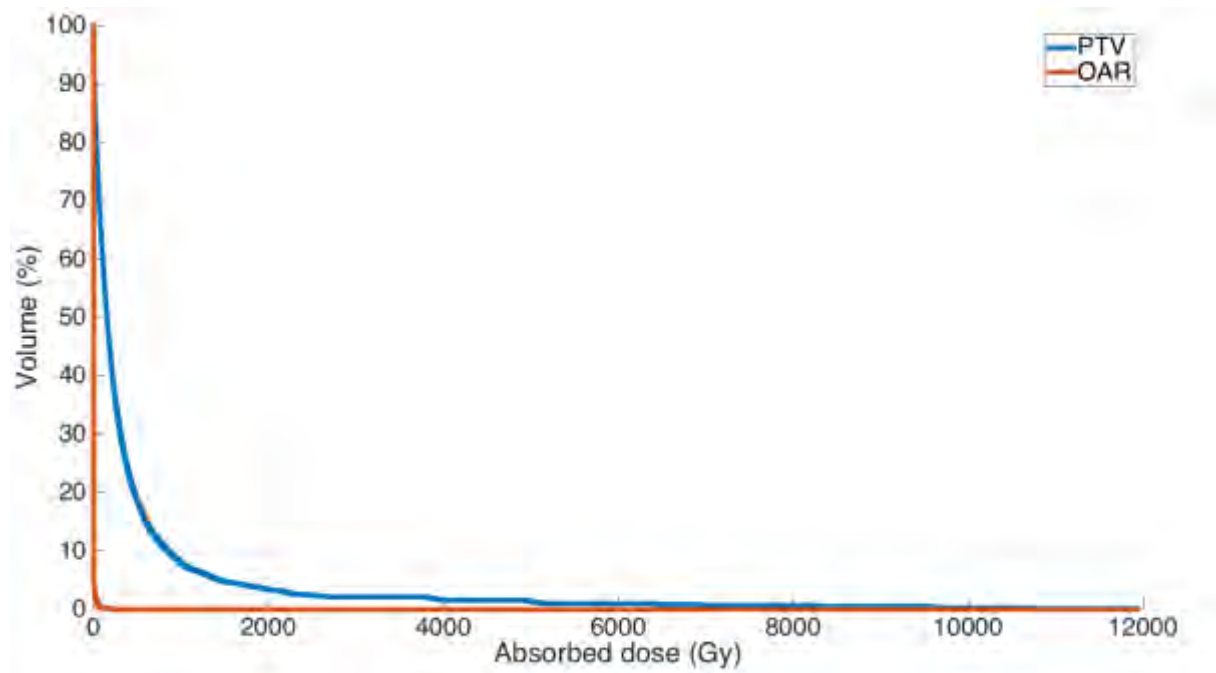


Figure 3.9: DVH of the PTV and OAR. The maximum absorbed dose delivered to the PTV was 11,934 Gy, meaning that the interesting range of the DVH from 0 to ~200 Gy, cannot be easily analysed. For this reason, DVHs shown throughout this work will be displayed up until a maximum of 200 Gy.

The DVH shown in Figure 3.9 extends up to the maximum absorbed dose found in the PTV. Most of the interesting information lies within the first few hundred Gy of this DVH. For this reason, all DVHs presented from now onwards in this work will be shown up to a maximum of 200 Gy, as can be seen in Figure 3.10.




Figure 3.10: DVH of the PTV and OAR, limited to a maximum of 200 Gy. It can be seen that the OAR requirement is respected because the OAR line passes underneath the black circle. Conversely, the PTV requirement is not respected because the PTV line does not pass over the black diamond.

Also shown in Figure 3.10 are the two treatment plan requirements. It can be seen that the PTV does not satisfy the inequality given in Equation 1.1 because the PTV DVH passes underneath the diamond at (95 Gy, 95%). It can be seen from the DVH that the volume receiving 95 Gy was actually only 63%.

The OAR, on the other hand, does satisfy the inequality given in Equation 1.2 because it passes under the point (10 Gy, 10%). From the graph, it can be seen that just 4% of the OAR receives greater than or equal to 10 Gy.

It can be seen that, compared to a typical EBRT DVH (such as the illustration shown in Figure 2.2), absorbed doses reach extremely high values with microbrachytherapy. Conversely, EBRT DVHs drop off rapidly after the target absorbed dose. As long as they are confined to the PTV, the extremely high absorbed doses that are observed in microbrachytherapy should not pose a problem.

3.7 Conclusion

For the dosimetric calculation of microbrachytherapy, two main methodologies have been proposed in this chapter: the first for generating the absorbed dose distribution for a single injection (known as a kernel) and the second for using this kernel to calculate the absorbed dose distribution of the patient (known as the dose grid) and performing structure-specific dosimetry.

The generation of kernels was done with GATE, by considering the injections to be spheres containing a homogeneous distribution of radioelements. Kernels were generated for various radioelements, injection volumes and voxel sizes. The number of particles simulated was between 10^9 and 10^{10} . Kernels are stored in S-values, meaning that they can easily be converted to absorbed doses by multiplying by the cumulated activity, which is simply a function of A_0 and the radioelement's half-life.

The creation of the dose grid was performed by creating a voxelised cuboid. Dose grids were created with the same size voxels as the kernel that would be convolved. The dose grid was the size of the PTV plus a margin of 12.5 mm in all directions. Since the injections are confined to the PTV, this dose grid must account for all irradiated tissue, regardless of whether the tissue be the PTV, the OAR or other healthy surrounding tissue.

The dose grid for a given treatment plan could then be calculated by simply convolving the injection kernel as many times as necessary in the position of each injection. This assumes that the irradiated volume consists entirely of a homogeneous medium (in this case, it is assumed to be soft tissue), an assumption that is deemed valid for microbrachytherapy in the brain, such as for the treatment of glioblastomas.

Structure-specific dosimetry was then performed simply by superposing binary masks corresponding to each structure on the dose grid. Once the dosimetry had been performed, it was possible to present the information in the form of DVHs. This allowed for comparison with the treatment requirements, such as those stated in Equations 1.1 and 1.2.

Attempts were made to speed up the calculation of the dose grid as much as possible. Excluding voxels containing negligible absorbed doses proved to be efficient, reducing the number of voxels by at least half, whereas GPU acceleration was found to improve calculation times.

4 COMPARISON OF OPTIMISATION ALGORITHMS

4.1 Introduction

This chapter consists of a comparison between the three algorithms that were presented in Chapter 2: NMS, BA and NSGSA-II. Such a comparison is not straight-forward since the algorithms are fundamentally different. NMS, for example, can only optimise a given number of injections at any one time. Furthermore, NSGA-II is an MO algorithm, so many results are presented together in the form of the Pareto front, as opposed to NMS and BA, which each produce a single treatment plan.

All algorithms that are compared throughout this chapter were written from scratch based on their descriptions in literature (as presented in Section 2.5). They were written in object-oriented C++, allowing for their incorporation into the microbrachytherapy TPS (which was also internally developed) that will be presented in the Appendix.

When necessary, the results of individual algorithms are presented. It is hoped that this will aid the reader to understand how each algorithm functions and manages to decrease the cost function, thereby finding treatment plans of increasing quality. Once the algorithms have been presented individually, they are then compared. Each comparison is performed by performing the same optimisation many times for each of the three algorithms and analysing the distribution of the final cost function.

Throughout this chapter, whenever plots of the cost function are shown, dashed lines are used to represent unacceptable treatment plans (treatment plans that violate either $Cost_{PTV} \leq 5\%$, $Cost_{OAR} \leq 10\%$, or both). Conversely, acceptable treatment plans are depicted using solid lines.

It should also be noted that when the cost is plotted as a function of the number of treatment plan updates, the graphs do not start at 0 updates. This is because treatment plan updates are made during the setup phase of each algorithm.

Throughout this chapter, the algorithm parameters remain unchanged. These empirically determined parameters are shown below, in Table 4.1.

Algorithm	Parameters
NMS	$\alpha=1, \beta=1, \gamma=1, \delta=1$
BA	$ns=50, ne=3, nre=10, nb=2, nrb=5, size_{patch}=10$
NSGA-II	$N=200, \text{probability of mutation}=1$

Table 4.1: Parameters used by the algorithms during the optimisation of the algorithms throughout this chapter. Parameter explanations are given in Sections 2.5.2 through 2.5.4.

The structure of this chapter is as follows:

- Section 4.2 describes injection parameters that will be used throughout this chapter.
- Sections 4.3 through 4.5 compare the algorithms in increasingly complex test cases:
 - Test Case 1 (Sections 4.3): optimise the irradiation of the PTV for a fixed number of injections without the presence of the OAR.
 - Test Case 2: (Section 4.4): optimise the irradiation of the PTV for a variable number of injections without the presence of the OAR.
 - Test Case 3 (Section 4.5): optimise the irradiation of the PTV and the OAR for a variable number of injections.
- Section 4.6 discusses the findings of the three test cases and identifies the best-suited algorithm for this project.
- Section 4.7 performs an analysis into the effect of voxel size on the accuracy of microbrachytherapy treatment planning.
- Section 4.8 demonstrates the optimisation of a treatment plan for a real pre-clinical case of a glioblastoma.
- Section 4.9 then concludes this chapter.

4.2 Injection parameters

During the development of a new therapy, many variables are under consideration at the same time. As was stated in the introduction to this thesis, this research consists of a comparison of optimisation algorithms and then comparisons of the ideal radioelement, initial activity and injected volume.

There is a circular dependency between the method of treatment planning and the injections:

- To be able to develop a method of treatment planning, injections are required.
- To be able to decide the best injection parameters, a method of treatment planning is required.

In this research, this circular dependency was broken by initially assigning arbitrary injection parameters, thereby permitting the comparison of optimisation algorithms in this chapter. Once the ideal optimisation algorithm has been developed, the following chapter will then be able to use this algorithm to compare injection parameters and find the ideal setup.

The injections used throughout this chapter were 5 μL injections of ^{90}Y , each with an A_0 of 5 MBq.

4.3 Test Case 1: Fixed number of injections and no OAR

The PTV used throughout this section is as described in Section 1.7. Since the number of injections is fixed and no OAR is present, both $Cost_{inj}$ and $Cost_{OAR}$ are not used. Therefore, the only cost function used in this section is $Cost_{PTV}$. Acceptable treatment plans are required as per usual to deliver $Cost_{PTV} \leq 5\%$.

Throughout this test, the number of injections was fixed at 20 (with the exception of Section 4.3.1.1, which uses just 2 injections to highlight the inner-workings of NMS).

4.3.1 NMS

4.3.1.1 With 2 injections

Shown below, in Figure 4.1, is the convergence of NMS for just 2 injections. As described in Section 2.5.2, the simplex is created with $3n_{inj} + 1$ treatment plans. Therefore, just 2 injections were chosen so as to minimise the number of plots on the graph as much as possible.

The evolution of the cost function for these 8 treatment plans (7 in the simplex plus the best found so far) can be seen. As explained in Section 2.5.2, normally just the worst treatment plan is altered. The exception for this is the reduction movement, where all treatment plans move towards the best solution.




Figure 4.1: Cost of NMS as a function of the number of updates. During each generation, the worst solution is altered. A reduction (in which all solutions move towards the best) can be seen at 38 updates. Reductions do not guarantee lowering the cost function of all solutions, and this is why some of the lines can be seen to increase at this point.

It can be seen that at 8 updates, treatment plan 7 was altered. This greatly improves the cost, becoming the best plan found so far, and so the blue ‘best’ line descends accordingly. Treatment plan 6 then becomes worst ranked, and at 14 updates it is altered, bettering the best treatment plan that was just found. The worst treatment plans are continually improved until the first reduction, which can be seen at 38 updates (this must be a reduction because all treatment plans are altered at the same time).

Unlike the other movements, reductions do not guarantee to reduce the cost function, as can be seen with treatment plans 3 and 6 at 38 updates. This is because all plans move towards the best solution, and it is possible that between the plan’s starting position and the position of the best solution, the cost function increases. However, as the size of the simplex decreases, all solutions eventually converge.

The optimisation converged in 26 iterations, which required 138 updates. The final cost was 81.1% (meaning that 81.1% of the PTV was under-irradiated). Since none of the plans found by this run of the algorithm were acceptable (under 5%), all lines are dashed. The optimisation process took a total of 2 s to complete.

4.3.1.2 With 20 injections

Clearly, since 81.1% of the PTV was under-irradiated at the end of the optimisation, significantly more than 2 injections are required to sufficiently irradiate the PTV. However, when the number of injections (and therefore the number of vertices of the simplex) increases, it becomes impractical to display the evolution of each treatment plan as a function of iteration number. Therefore, Figure 4.2 shows the evolution of the cost function, displaying only the best and worst costs for an optimisation of 20 injections.

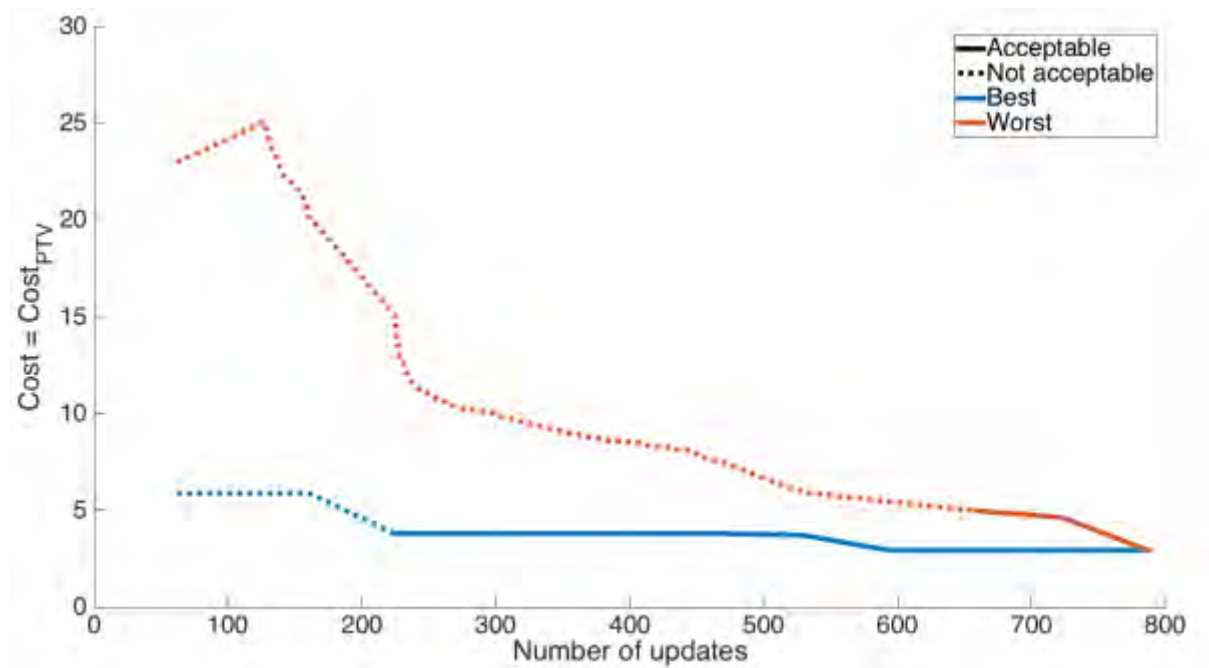


Figure 4.2: Evolution of the cost for NMS for the best and worst treatment plans for an optimisation of 20 injections. It can be seen that, because NMS mostly improves the worst solution, the best solution changes little over the course of the optimisation.

It can be seen in Figure 4.2 that, because NMS spends most of its iterations improving the worst solution, the best solution changes little over the course of the optimisation. The best solution starts at 5.86% and finishes at 2.94%. Despite this small improvement, the solution nonetheless passes from an unacceptable plan to one that is acceptable (and thus the dashed line becomes a solid line).

The optimisation finished in 89 iterations, corresponding to 787 updates. The time taken was 24 s. The increased number of updates before convergence (and therefore increased optimisation time) compared to Figure 4.1 is due both to the increased number of vertices of the simplex (an increased number of treatment plans to alter) and the increased number of degrees of liberty by increasing the number of injections from 2 to 20.

4.3.2 BA

Shown below, in Figure 4.3, is an example of the flower patch search for BA. The optimisation was performed with 20 injections. As shown in Table 4.1, BA was configured such that there were 50 scout bees, 3 elite flower patches and 2 best flower patches. Each elite flower patch was assigned 10 forager bees and each best flower patch was assigned 5 forager bees, corresponding to a total population of 90 bees.

First, the 50 scout bees were randomly distributed throughout the search-space (50 treatment plans were created, each with 20 randomly placed injections). These scout bees were ranked and the best 5 scout bees became flower patch centres.

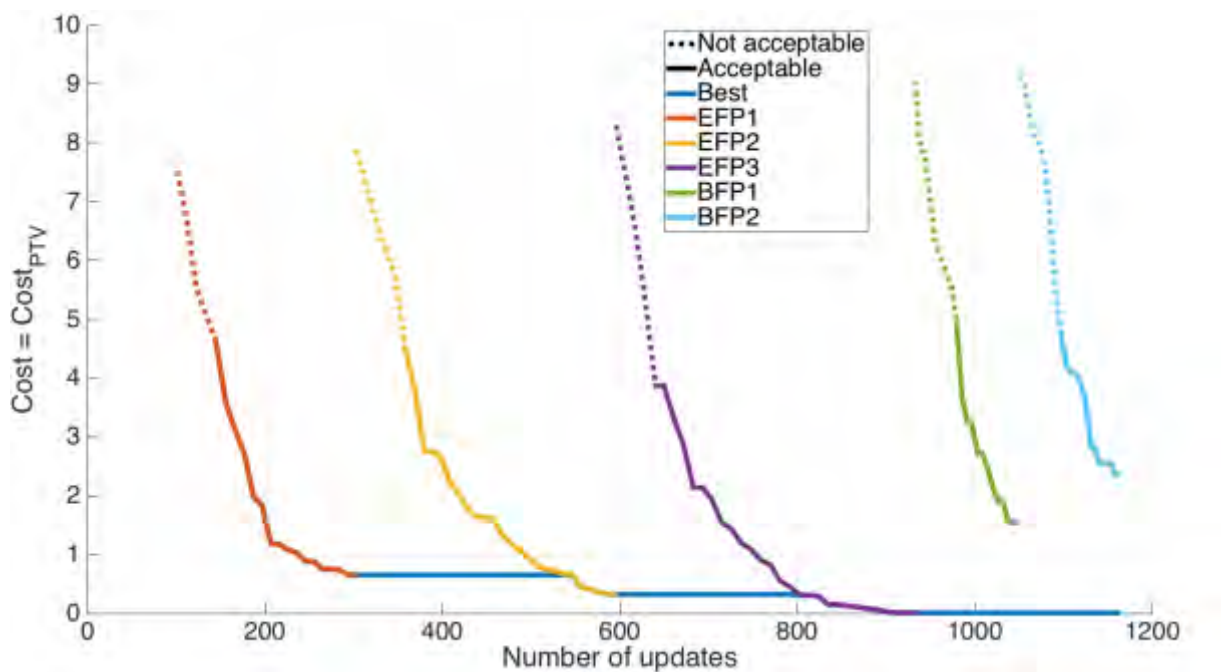


Figure 4.3: Evolution of the cost as a function of the number of treatment plan updates for “elite” and “best” flower patches with BA. The start of each line is the cost at the initial centre of the flower patch (i.e., the cost of the scout that was used to create the patch). As better solutions are found by forager bees in the surrounding search-space, they become the flower patch centre and the cost decreases accordingly. At the end of each line, the flower patch has passed beneath a user-defined minimum size and it is assumed that the local minimum has been found. The best line corresponds to the best solution found in all flower patches up until that point.

EFP1 (the first elite flower patch) was initialised at the position of the best scout, EFP2 at the position of the 2nd-best scout bee and so on, such that BFP2 was initialised at the position of the 5th-best scout bee. It is for this reason that the starting point of each line from EFP1 to BFP2 increases.

It can be seen that although the first elite flower patch was initially the most promising, the best treatment plan was actually found in the vicinity of the 3rd elite flower patch (because the EFP3 line finishes with the lowest cost). This highlights the importance that algorithms demonstrate both exploratory behaviour (searching areas even if they do not initially seem to be the most promising) and exploitive behaviour (searching in the vicinity of promising scout bees to find local minima).

The optimisation shown in Figure 4.3 was performed for just 1 iteration (so as not to clutter the figure with too many lines), but the process can be repeated for as many iterations as desired. A new iteration would start with redistributing the scout bees and so, each iteration is completely independent of the other iterations. The optimiser need only remember the best treatment plan found amongst all the flower patches of all the iterations so as to be able to propose this plan to the user at the end of the optimisation process.

The optimisation shown in Figure 4.3 finished after 1,162 updates, taking a total of 21 s. The best treatment plan had a cost of 0.00561% (meaning that just 0.00561% of the PTV was under-irradiated).

4.3.3 NSGA-II

To be able to demonstrate the production of new treatment plans via crossovers, Figure 4.4 shows the first five generations for a small population of just 10 treatment plans. The optimisation was performed with 20 injections. The mutation probability was 1, meaning that at as each treatment plan was created with a crossover, it would surely be mutated and one of the injection positions would be randomly changed. Normally, mutations can lead to an increase or decrease in the number of injections; this was not the case here, however, because the number of injections was fixed at 20.




Figure 4.4: Evolution of the cost as a function of generation (i.e., iteration) for a population of 10 treatment plans in NSGA-II. Solutions that were created by crossovers are joined to their parents via solid blue lines, whereas solutions surviving from previous generations are shown via red dashed lines.

In Figure 4.4, each solution created from a crossover is connected via blue lines to its two parents. It can be seen that better ranked solutions contribute more to future generations the lower ranked solutions. The best ranked solution in the 0th generation, for example, contributes to four crossovers (depicted by the four blue lines that connect it to solutions in the 1st generation), whereas the worst contributes to just one.

Truncation replacement (joining the current and previous generations together and keeping the best half, as discussed in Section 2.5.4.3) means that some parent solutions survive from previous generations. These solutions are shown with the red dashed line. Truncation replacement guarantees that the quality of each generation is at least as good as the previous, and this can be seen in the figure. If all the newly created children were of better quality than all the parents, then no parents would survive from the previous generation (there would be no red dashed lines). However, if all of the children were of worse quality than the parents, only the parents would survive (there would be only red dashed lines and no blue lines). In reality, it is normally somewhere in between, thus the mixture of red and blue lines.

The process was then repeated with a full-sized population of 200 treatment plans and the optimisation was performed over 50 generations. The results of the best solution as a function of the number of updates is shown in Figure 4.5. The optimisation finished in 1,560 updates, which took 23 s. The final cost function was 0.143%.

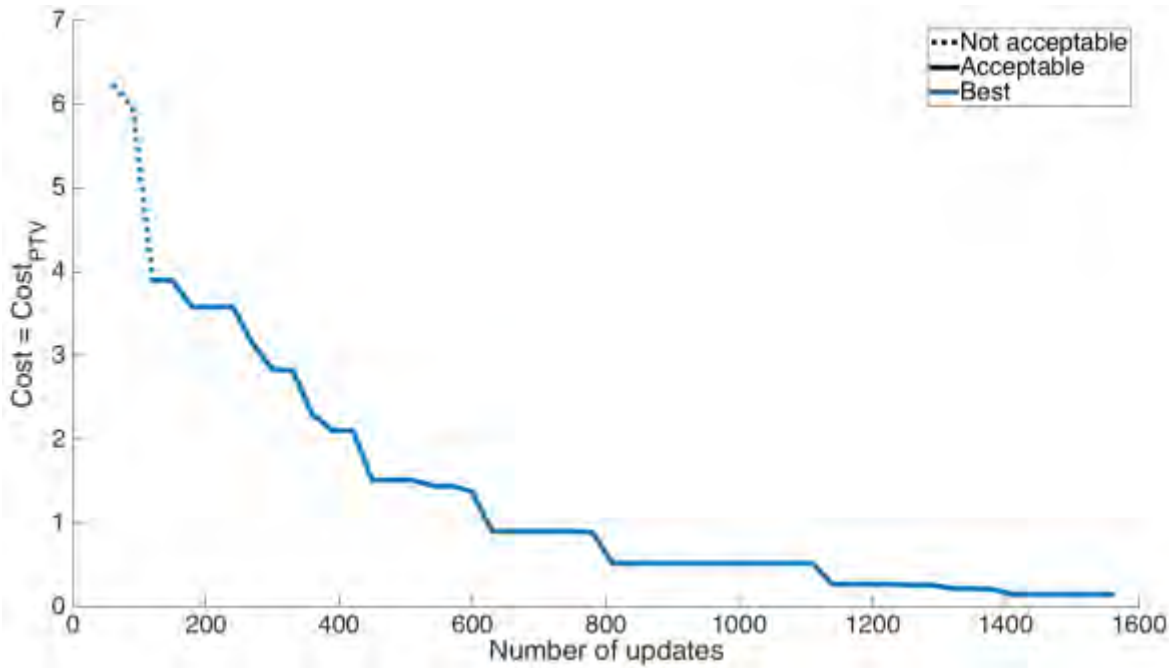


Figure 4.5: Cost as a function of treatment plan updates for NSGA-II.

4.3.4 Comparison

To compare the algorithms, each was fixed at 20 injections. The algorithms were launched 300 times each. Algorithms stopped either when they surpassed 1,000 updates or, in the case of NMS, converged.

A histogram of their final cost distributions is shown in Figure 4.6 and additional information is given in Table 4.2.




Figure 4.6: Frequency histogram for the final cost of the three algorithms, each launched 300 times with 20 injections.

	NMS	BA	NSGA-II
Minimum cost	1.29	0.00	0.104
Maximum cost	7.85	0.553	1.30
Mean cost	4.32	0.109	0.573
Standard deviation	1.20	0.0967	0.214
Acceptable plans (%)	72.3	100	100
Mean time (s)	11.4	18.5	12.1

Table 4.2: Minimum, maximum, mean and standard deviation of the final cost, percentage of acceptable plans and the mean optimisation time for the three algorithms. Each algorithm was launched 300 times with 20 injections.

4.3.5 Discussion

For this first test case, the results of the individual algorithms were initially presented. Figure 4.1, Figure 4.3 and Figure 4.4, in particular, show insight into the inner-workings of the three algorithms and how they manage to reduce the cost function.

Figure 4.1 shows how the worst solution is constantly improved and how all solutions eventually converge to a single point. Figure 4.3 shows the search of forager bees as they centre in on a local minimum. Lastly, Figure 4.4 shows how crossovers are used to create new treatment plans and how keeping the best results from previous generations guarantees that the quality of each generation never decreases.

Subsequently, a comparison was performed between the three algorithms. It can be seen, both in Figure 4.6 and in Table 4.2, that BA produces treatment plans that consistently have the lowest cost function. Both BA and NSGA-II produced acceptable treatment plans 100% of the time.

NSGA-II's principal drawback (as mentioned in Section 2.5.4.5) is that it does not explicitly search in the immediate surrounding area of promising areas. It is perhaps for this reason that BA produced better results.

NMS not only found treatment plans with the biggest costs, but it also has the largest spread in the quality of treatment plans found, indicating a lack of robustness, probably due to non-optimal initialisations and its deterministic nature.

4.4 Test Case 2: Variable number of injections and no OAR

Increasing the complexity over the first test case, this comparison is conducted by allowing the number of injections to vary between 1 and 20. The cost function, $Cost_{inj}$, is therefore introduced to minimise the number of injections. For the SO algorithms, the two costlet functions, $Cost_{PTV}$ and $Cost_{inj}$, were given relative weightings of 0.8 and 0.2, respectively. These weightings were empirically determined to give satisfactory results.

As with the previous test, the results of the individual algorithms are given in Sections 4.4.1 through 4.4.3, followed by the comparison between the three algorithms in Section 4.4.4.

As was mentioned in Section 2.3.4.3, feasible plans (i.e., those that violate none of the treatment requirements) are always ranked ahead of infeasible plans. Hence, it is possible that treatment plans with higher costs will be preferred over those with lower costs. Examples of this will be seen in Figure 4.7 and Figure 4.8.

Whilst for the first test case, treatment plans with costs of under 5% were deemed acceptable, this is no longer true for the second test case. This is because the cost function is now a weighted sum of two cost functions, so although acceptable plans are still required to have $Cost_{PTV} \leq 5\%$, they are not required to have a specific value for $Cost_{Total}$.

4.4.1 NMS

As mentioned in 2.5.2.2, NMS is only capable of optimising a fixed number of injections. For this reason, the optimisation is performed individually for each number of injections between 1 and 20. NMS was therefore launched 20 times and the evolution of the cost for each number of injections is shown in Figure 4.7.




Figure 4.7: The evolution of the total cost for each number of injections that were optimised individually. It can be seen that only by using 20 injections could a treatment plan be found that sufficiently irradiates the PTV (as this is the only solid line amongst all the dashed lines).

The only number of injections that resulted in an acceptable treatment plan (as depicted by the only solid line amongst all the other dashed lines) was 20, even though it can be seen that using 19 injections resulted in a lower cost function.

The total number of treatment plan updates for the optimisation was 8,600, taking a total of 113 s.

4.4.2 BA

Shown below, in Figure 4.8, is the evolution of the cost as a function of updates for a variable number of injections with BA.




Figure 4.8: Evolution of BA for an optimisation with a variable number of injections. It can be seen that at 626 updates, EFP2 find a solution with a lower total cost than the best solution. However, the best solution is acceptable, whereas the one proposed by EFP2 is not, and so this solution does not become the best. It is only at 668 updates that the solution proposed by EFP2 is both acceptable and lower than the best, and at this point it replaces the best solution.

It can be seen that at 626 updates, the cost of the 2nd elite flower patch falls beneath the best line. The best line, however, does not reduce with it because at this point the 2nd flower patch proposes an unacceptable treatment plan and the current best plan is acceptable. It is only at 668 updates that the plan found by EFP2 becomes acceptable and so the best line drops down accordingly.

As was the case with Figure 4.3, the best ranked scout did not produce the best treatment plan, which was actually found in the vicinity of the 5th ranked scout bee, again highlighting the importance of searching areas of the search-spaces, even if they are not the most promising initially.

4.4.3 NSGA-II

Shown below, in Figure 4.9, are the results for an MO optimisation of NSGA-II. The optimisation was performed with a population of 200 treatment plans over 500 generations.




Figure 4.9: Evolution of the Pareto front as a function of the number of generations. $Cost_{inj}$ is shown along the bottom, whereas the number of injections are shown along the top. The best SO plan (lowest total cost whilst respecting the treatment criteria) for each generation is marked with a black cross.

The 500 generations completed in a total of 463 s, requiring a total of 50,200 treatment plan updates. The final cost was 16.4, at which point $Cost_{PTV}$ was 1.74 and $Cost_{inj}$ was 75 (15 injections).

When regarding the best total costs (the black crosses) in Figure 4.9, an important distinction should be made between:

- Performing an MO optimisation and from that calculating the best SO plan, and
- Performing an SO optimisation by combining the individual costlets prior to the optimisation process (as described in Section 2.5.4.4).

Figure 4.9 is an example of running an MO optimisation and calculating the best SO treatment plan in the Pareto front. Figure 4.10, however, shows that if the user is concerned by the total cost (i.e., wants the best result from an SO optimisation), then the costlet functions should be combined prior to the optimisation process.




Figure 4.10: Evolution of the total cost as a function of the number of treatment plan updates. Clearly, when the user is concerned by the total cost (i.e., the SO cost), it is better to use NSGA-II as an SO algorithm. If NSGA-II is used in MO mode, it concentrates its resources on exploring the entirety Pareto front, which is not advantageous when the user is only concerned by one specific point of the Pareto front.

As one might expect, when the user is concerned in the total cost, performing an SO optimisation produces better results than extracting the best SO result from an MO optimisation. This is because an MO optimisation tries to optimise an ND surface (where N is the number of costlets) with the greatest spread possible between each of the solutions across this surface. An SO optimisation, on the other hand, concentrates on optimising a 1D surface (a single point) and so all of its resources are focused on reducing the total cost.

Hence, whenever a Pareto front is shown in this research, NSGA-II will be used in MO mode. However, whenever NSGA-II is used as an SO algorithm (for example, the comparison in Section 4.4.4), it will be used in SO mode.

4.4.3.1 Population diversity

As discussed in Section 2.5.4, mutations and the crowding distances are used by NSGA-II to promote diversity in the population of treatment plans.

Mutations are used to randomly alter the treatment plans. Each time a mutation occurs, a new random injection could be added or one of the existing injections could be deleted or altered.

The crowding distance, on the other hand, is used to favour unique solutions, encouraging the Pareto front to spread out as wide as possible. Moreover, using the crowding distances, solutions at the edges of the Pareto front are always ranked first, further encouraging the spreading out of solutions.

Shown below, in Figure 4.11, is the increase of diversity in the population thanks to mutations and the crowding distance. Some of the first 20 generations are shown for a population of 200 treatment plans. All plans were all initialised identically, each with 10 identical injections. The initial $Cost_{PTV}$ for these plans was 38.6%.




Figure 4.11: Example of increasing diversity in NSGA-II due to mutations and the crowding distance. All solutions were initialised identically with 10 injections, all with matching coordinates. It can be seen that, thanks to mutations and the crowding distance, the solutions rapidly spread out. After just 20 generations, solutions have been proposed that cover the full range of allowed numbers of injections (1 to 20).

It can be seen that, despite starting identically, the Pareto front rapidly spreads out, so as to cover the widest range possible. By doing so, the largest amount of search-space is explored with the aim of offering a wide range of Pareto-optimal solutions to the user at the end of the optimisation process.

4.4.4 Comparison

Given the increased complexity of the problem – the search-space has been increased by allowing the algorithms to alter the number of injections – the maximum number of treatment plan updates that was used as a stopping criterion was increased from 1,000 to 2,000. This could not be imposed for NMS (because it was run many times in parallel for each of the different number of injections), so it was allowed to run until convergence.




Figure 4.12: Frequency histogram for the final cost of the three algorithms, each launched 300 times with between 1 and 20 injections.

	NMS	BA	NSGA-II
Minimum cost	19.2	16.6	16.3
Maximum cost	23.8	19.6	20.9
Mean cost	21.8	18.4	17.9
Standard deviation	0.794	0.572	0.839
Acceptable plans (%)	46.0	99.0	99.7
Mean time (s)	107	37.6	23.2

Table 4.3: Minimum, maximum, mean and standard deviation of the final cost, percentage of acceptable plans and the mean optimisation time for the three algorithms, each launched 300 times with between 1 and 20 injections.

4.4.5 Discussion

The conclusion that can be drawn from the comparison is not the same as that for the first test case. From both Figure 4.12 and Table 4.3, it can be seen that, generally, NSGA-II finds treatment plans with lower costs than BA. BA is, however, more consistent, with a lower standard deviation.

One can conclude from this that, although NSGA-II does not explicitly search in the vicinity of promising solutions, enough generations will pass that at least one of the offspring will be of a high quality. This is not guaranteed, however, and this probably explains the increased standard deviation compared to BA.

With regards to MO versus SO, it can be seen in Figure 4.9, that the treatment plan with the lowest total cost is not necessarily the plan that sufficiently irradiates the tumour whilst using the minimal number of injections. The treatment plan with the lowest cost at the end of the 500th generation used 15 injections and delivered a total cost of 16.4 (corresponding to a $Cost_{PTV}$ of 1.74 and a $Cost_{inj}$ of 75). However, there is an acceptable plan that uses just 12 injections. This plan, despite having a higher total cost, is arguably better, since it sufficiently irradiates the PTV and uses less injections. This highlights the difficulty of choosing cost function weightings that can be trusted to produce a treatment plan the best reflects the user's idea of the best treatment plan. Therefore, the ideal algorithm may not be the algorithm that is the most efficient in lowering the cost function. It can be seen that MO algorithms, such as NSGA-II, offer more choice to the user, something which is clearly advantageous.

The simplicity of NMS is a clear disadvantage. The fact that each number of injections must be optimised individually means that, whilst NMS was the fastest to converge for a fixed number of injections, it is by far the slowest for variable numbers of injections. Furthermore, since NMS is a deterministic algorithm, the mean final cost is considerably higher than the other two algorithms.

4.5 Test Case 3: Variable number of injections and OAR

The setup for this comparison is the same as the previous test case, except that this time the OAR (as described in Section 1.7.2) is included. As a reminder, it is required that no more than 10% of the OAR receive over 10 Gy ($Cost_{OAR}$ should be less than or equal to 10%).

With the inclusion of the OAR, $Cost_{OAR}$ is added correspondingly to the total cost function. The weightings, w_{PTV} , w_{inj} and w_{OAR} , were 0.8, 0.1 and 0.1, respectively. These weightings were determined experimentally.

The methodology of NMS and BA remain unchanged, so their results are not shown prior to the comparison. Instead, the individual results of NSGA-II are shown, followed directly by the comparison.

4.5.1 NSGA-II

To demonstrate NSGA-II in MO mode with three costlet functions, a population of 600 treatment plans were optimised over 100 generations. The resulting 3D Pareto front is shown below, in Figure 4.13. Of the 600 plans in the population, 178 made up the final Pareto front after 100 generations. All of these treatment plans are considered acceptable since $Cost_{PTV} \leq 5\%$ and $Cost_{OAR} \leq 10\%$ for all plans.




Figure 4.13: 3D Pareto front for NSGA-II. The data points are coloured as a function of the number of injections used by that plan. A $Cost_{inj}$ of 70% corresponds to 14 injections, whereas a $Cost_{inj}$ of 100% corresponds to 20 injections, as described in Section 2.4.2. The plan with the lowest total cost is highlighted as “best”.

The lowest total cost function amongst all the Pareto-optimal solutions is highlighted. The total cost for this plan is 9.46, corresponding to $Cost_{PTV}$, $Cost_{inj}$ and $Cost_{OAR}$ of 0.688%, 85% (17 injections) and 4.10%, respectively.

The total optimisation time was 593 s, and a total of 30,600 treatment plan updates were performed.

4.5.2 Comparison

As with the previous sections, the comparison was performed by launching each of the algorithms 300 times. A maximum of 2,000 treatment plan updates was permitted for BA and NSGA-II, and NMS continued until convergence. The frequency distribution histogram of the final cost functions is shown in Figure 4.14 and additional information is given in Table 4.4.




Figure 4.14: Frequency histogram for the final cost of the three algorithms, each launched 300 times with between 1 and 20 injections.

	NMS	BA	NSGA-II
Minimum cost	11.5	9.69	9.34
Maximum cost	15.1	11.2	11.3
Mean cost	13.3	10.5	10.1
Standard deviation	0.779	0.305	0.320
Acceptable plans (%)	68.0	100	100
Mean time (s)	160	74.8	32.8

Table 4.4: Minimum, maximum, mean and standard deviation of the final cost, percentage of acceptable plans and the mean optimisation time for the three algorithms, each launched 300 times with between 1 and 20 injections.

4.5.3 Discussion

The Pareto front shown in Figure 4.13 shows that there are many acceptable treatment plans, each with different qualities. If a user was proposed all of these plans, they might choose to decrease the irradiation of the PTV therefore permitting them to decrease the irradiation of the OAR. From here, by increasing the number of injections, it would then be possible to keep the decreased irradiation of the OAR whilst increasing the irradiation of the PTV. This is the liberty that is provided by the Pareto front.

From the frequency histogram, the same conclusion can be drawn as from the second test case. NSGA-II produces treatment plans of a slightly higher quality than BA, and NMS is far behind, proposing satisfactory treatment plans just 68% of the time.

It can be seen in Table 4.4, that NMS is the slowest algorithm. This is unsurprising since the optimisation process is repeated many times for each of the number of injections. Curiously, however, NSGA-II is over twice as fast as BA. This is because the algorithms check at the end of each iteration to see if they have satisfied the stopping criteria. For BA, even though the maximum number of treatment plan updates was 2,000, the end of the iteration could occur much later, in one case occurring at 3,278 updates.

4.6 Discussion

The results of each algorithm were discussed at the end of each test case. Therefore, this section discusses general points that can be drawn overall from the three tests. This section discusses: the choice of relative costlet weightings, extending the number of cost functions above three, the maximum number of treatment plan updates as a stopping criterion, improving NMS and optimisation times. Lastly, in this discussion the best algorithm is decided upon, permitting the treatment planning of a real tumour in Section 4.8.

4.6.1 Costlet weightings

In the second test case, the relative weightings of 0.8 and 0.2 were proposed for $Cost_{PTV}$ and $Cost_{Inj}$, respectively. These values were established empirically. $Cost_{Inj}$ is normalised such that it is 100% when the maximum number of injections are used, something which occurs regularly. Conversely, although $Cost_{PTV}$ can vary between 0% and 100%, treatment plans with $Cost_{PTV} \geq 20\%$ considerably under-irradiate the PTV and are not worth considering. Since $Cost_{PTV}$ was generally less than $Cost_{OAR}$, w_{PTV} therefore had to be greater than w_{Inj} so that they contributed equally to the total cost.

This highlights the complexity of joining costlet functions into a single SO cost function and the difficulty in normalising them. The maximum number of injections was set at 20 because trial tests showed that this was a suitable maximum. However, if the maximum had have been set to 40, treatment plans would have been proposed with $Cost_{Inj} \approx 50\%$. If this had have been the case, it would have been necessary to increase w_{Inj} .

In short, there is no easy way to determine costlet importance for SO optimisations. The relative importance should be decided through trial-and-error. The user should pick relative weightings and perform the optimisation. If upon analysing the results, the user finds that the optimiser overly favoured a certain costlet, it would then be necessary to alter the weightings and restart the optimisation. This, needless to say, is a long process, but one that should be accelerated once the user has already performed similar treatment plans.

4.6.2 The Pareto front in 4D (or more)

In theory, MO algorithms can be used for an infinite number of cost functions. However, as the number of dimensions surpasses three, the visualisation of the Pareto front becomes challenging. Shown below, in Figure 4.15, is the optimisation of the same setup as the third test case (optimisation of the PTV, the OAR and the number of injections) with a fourth cost function.

The fourth cost function is the conformity index (CI), as described in Section 2.3.4.4, which is minimised when the PTV is sufficiently irradiated and the surrounding healthy tissue is not over-irradiated.

The optimisation for which the results are shown below was performed with a population of 600 over 100 generations. The optimisation finished in 652 s, making a total of 30,600 treatment plan modifications.




Figure 4.15: 2D projections of the 4D Pareto front. From the projections, it is possible to ascertain correlations (such as increasing $Cost_{Inj}$ decreases $Cost_{PTV}$), but other relationships are harder to analyse (such as $Cost_{OAR}$ and $Cost_{CI}$). For this reason, it would be better to display all Pareto-optimal results in the form of a table.

It is easy to draw certain conclusions from the previous figure. For example, it can be seen that there is a negative correlation between $Cost_{Inj}$ and $Cost_{PTV}$. This is to be expected since the portion of the PTV that is under-irradiated is likely to decrease as the number of injections increases.

Some aspects, however, are more difficult to analyse using the data presented in 2D projections of a 4D surface. For example, it is difficult to understand the relationship between $Cost_{OAR}$ and $Cost_{CI}$. For this reason, it would probably be more suitable to use a table to present the Pareto front for four or more dimensions. This would allow the user to more precisely observe the trade-offs between each objective.

Lastly, it should be noted that as the number of costlet functions increases, the number of dimensions that need to be populated by the Pareto front increases accordingly. For example, when passing from the 2D Pareto front in Figure 4.9 to the 3D Pareto front in Figure 4.13, the number of solutions in the Pareto front increased from 19 to 178. It can therefore be seen that as the number of costlets is increased, the population should be increased accordingly so as to be able to populate the search-space. This will, of course, slow down the optimisation process.

4.6.3 Maximum number of updates as a stopping criterion

Throughout this chapter, many figures have been presented that show the evolution of the cost as a function of either iteration number or the number of treatment plan updates. In all of these examples, it can be seen that the increase in quality of the best treatment plan is often initially rapid. As the optimisation continues, however, it gets increasingly difficult to find treatment plans of better quality. For BA and NSGA-II – two algorithms that do not converge, unlike NMS – they eventually reach a point of diminishing return.

The moment that this point of diminishing return occurs depends on many factors, such as the number of cost functions, the number of injections and the size of the tumour. Hence, the user should set the maximum-number-of-updates stopping criterion accordingly. This should get easier as the user's experience increases.

4.6.4 Improving NMS

Two methods are proposed below, which could potentially be used to improve the results produced by NMS. The first is to use a meta-optimiser, and the second is to use restarts.

4.6.4.1 Meta-optimiser

An algorithm, such as the golden-section search algorithm (Kiefer, 1953), could be used to improve optimisation times when the number of injections is variable. Such a meta-optimiser would mean that, instead of testing all numbers of injections (every number between 1 and 20 in these examples), the algorithm would intelligently select numbers of injections should be tested.

For example, it can be seen that in the second and third test cases, the optimal number of injections was typically between 15 and 20. There is little incentive, therefore, to test the number of injections between 1 and 12, say. The meta-optimiser would reduce the trialled number of injections by essentially plotting cost as a function of the number of injections and search for the minimum of this curve.

4.6.4.2 *Restarts*

Restarts can be used to improve the quality of results produced by NMS. Since NMS is heavily dependent on the initialisation of the simplex at the start of the optimisation process, recreating the simplex and re-running the algorithm multiple times, proposing the best treatment plan from all of these runs at the end, would minimise the impact of a poor initialisation.

NMS, however, was already shown to be the slowest algorithm. If 10 restarts were used each time, optimisation times would be increased by a factor of 10. Even with using a meta-optimiser to reduce optimisation times, it is likely that NMS would be considerably slower than the other algorithms.

4.6.5 *Optimisation times*

Despite the fact that a difference could be seen between the quality of treatment plans produced by the three algorithms, a positive conclusion can be drawn from all three: Optimisation times for all algorithms were extremely rapid. For the most complicated test case in this chapter, NMS – the slowest algorithm – took on average just 160 s to perform the treatment planning.

Clearly, the optimisation was performed on a small tumour and so optimisation times would increase as the volume increases and the number of injections increases accordingly (as will be seen in Section 5.4). However, the initial results are extremely satisfactory in terms of clinical timescales for treatment planning (the upper limit is about an hour) and so these algorithms should continue to perform rapidly, even for larger setups.

Optimisation times, however, could be further improved by parallelising the algorithms.

4.6.5.1 *Parallelised algorithms*

Certain algorithms lend themselves more readily to parallelisation. This means that computers can profit from the maximum capacity of multi-core computers via multi-threaded code.

BA, for example, is easily parallelised because the search of each flower patch can be effectuated independently of the other flower patches. Similarly, NSGA-II creates the new population by repeatedly creating two children from two parents; parent-child production can also be undertaken independently. Hence, BA and NSGA-II can easily be split into multi-threaded code. NMS, conversely, operates in a linear fashion, meaning that it is not suitable for splitting into multi-threaded code.

Of the three, NSGA-II was therefore programmed so as to be able to make use of multi-threaded CPUs, decreasing its simulation time. In order to be able to perform a fair comparison, all NSGA-II optimisations performed throughout this chapter were performed on a single core.

4.6.6 The chosen algorithm

From the three test cases, it was shown that for microbrachytherapy treatment planning, NMS is a poor choice. NMS is efficient for convex cost functions (only one global minimum and no local minima) and where a good initialisation can be assured. Neither of these two are true for this specific problem, where the cost function contains many local minima (there are many treatment plans in different areas of the search-space with of similar quality) and vertices of the simplex are initialised by random treatment plans (randomly initialising treatment plans does not guarantee the wide spread of the vertices that is required, as discussed in Section 2.5.2.1).

BA and NSGA-II have been shown to give similar, acceptable results. NSGA-II, however, not only performs well as an SO algorithm, but it also has the advantage of being an MO algorithm. This means that for problems where the relative weightings are unknown, one can perform treatment planning in MO mode and pick the best treatment plan from the non-dominated solutions that populate the Pareto front. For this reason, it was decided that, between the three algorithms that were tested, NSGA-II was the algorithm that was best suited for the problem of microbrachytherapy treatment planning.

Therefore, NSGA-II will be used for all future treatment plan optimisations in this research.

Since SO algorithms will no longer be used, $Cost_{inj}$ can now be expressed as the number of injections instead of the normalised value that has been used until now by dividing by the maximum number of injections.

4.7 Voxel sizes

Now that a method of treatment planning has been developed, and an algorithm has been selected for all future treatment planning optimisations, an important study can be conducted into the impact of voxel size on the accuracy of calculated dosimetry.

The absorbed dose profiles of β^- -emitting radioelements are very steep; the 1D kernel for a 5 μL injection of ^{90}Y with an A_0 of 5 MBq in Figure 3.5 shows that the absorbed dose drops from greater than 10^4 Gy to less than 1 Gy in less than 1 cm.

For this reason, the voxels that are used throughout the treatment planning process should be small enough that the dosimetry is performed with reasonable accuracy. It is therefore useful to study whether changing the voxel size has an effect on the Pareto front.

Throughout this section, injections of ^{166}Ho are used instead of ^{90}Y because, as was shown in Table 3.1, simulation times for ^{90}Y are considerably longer than for other radioelements, and so the production of kernels with small voxel sizes (such as 0.1^3 mm^3) would have taken unreasonably long. For this reason, the test was performed with ^{166}Ho , but the conclusions should be the same regardless of the radioelements, as all β^- -emitting radioelements have similarly steep gradients (as will be shown in Section 5.2.1).

In this section, a comparison is performed between 0.1^3 , 0.5^3 and 1.0^3 mm^3 . Shown below, in Figure 4.16, is the central slice through the absorbed dose distributions of 5 μL injections of ^{166}Ho with A_0 's of 5 MBq and for these three voxel sizes.

The 0.1^3 mm^3 kernel was calculated using 10^{10} particles (compared to 10^9 for the original 0.5^3 mm^3 kernel). The 1.0^3 mm^3 kernel was calculated by resampling the 0.5^3 mm^3 kernel.




Figure 4.16: Central slice of absorbed dose distributions for 5 μL injections of ^{166}Ho with A_0 's of 5 MBq and voxel sizes of (from left to right) 0.1^3 , 0.5^3 and 1.0^3 mm^3 .

The study into the effect of voxel sizes on the treatment planning will be conducted in two parts. First, the effect on the Pareto front will be studied. Secondly, the effect on individual treatment plans will be studied by selecting a single treatment plan and performing the dosimetry with different voxel sizes to see if the DVHs change.

Throughout this study, the PTV and the OAR were present, though only $Cost_{PTV}$ and $Cost_{inj}$ were used (so as to be able to present the Pareto front in 2D).

4.7.1 Optimisation

The optimisation was performed over 1,000 generations with a population of 80. The number of injections was constrained between 1 and 35. The maximum number of injections was increased from 20 to 35 because ^{166}Ho produces lower energy β^- particles than ^{90}Y that has been used previously (and as will be discussed in the following chapter).

The Pareto fronts for these optimisations are shown in Figure 4.17 and optimisation times for each of the voxel sizes is shown after Figure 4.17, in Table 4.5.




Figure 4.17: Pareto front the optimisation of the PTV (10 mm radius) with injections of ^{166}Ho for kernels with voxel sizes of 0.1^3 , 0.5^3 and 1.0^3 mm^3 .

Voxel size (mm ³)	Optimisation time (s)
0.1	33,837
0.5	245
1.0	68

Table 4.5: Optimisation times as a function of voxel spacing. The optimisations were performed over 1,000 generations.

4.7.2 DVHs

It can be seen in Figure 4.17 that voxel size – at least, for sizes under 1.0³ mm³ – has little impact on the Pareto front. This, however, does not confirm whether voxel size significantly affects a plan’s dosimetry, since the optimiser finding plans with matching cost functions does not necessarily mean that the dosimetry of plans with matching coordinates but different voxel sizes would be the same.

This is demonstrated in Figure 4.18, where the same injection coordinates are applied with the three kernels of different sized voxel. The plan was chosen from the 0.1³ mm³ Pareto front, as this is assumed to be the most accurate of the three voxel sizes. The chosen plan used 18 injections, which was the acceptable treatment plan that used the fewest injections.




Figure 4.18: DVHs for the same plan (the first that passes under 5% on the y-axis for the voxel size of 0.1³ mm³ in the previous figure) with different voxel sizes.

The DVHs shown in Figure 4.18 are very similar. For this reason, the residual between the DVHs, using 0.1^3 mm^3 as the reference (i.e., the DVH of 0.5^3 mm^3 minus the DVH of 0.1^3 mm^3 and the DVH of 1.0^3 mm^3 minus the DVH of 0.1^3 mm^3), are shown in Figure 4.19.




Figure 4.19: Differences in 0.5^3 and 1.0^3 mm^3 voxel size DVHs compared to the DVH of the 0.1^3 mm^3 voxel size.

From Figure 4.19, it can be seen that in this treatment plan, increasing the voxel size underestimates the irradiation of the PTV, whilst overestimating of the irradiation of the OAR. It can also be seen however, that differences between 0.5^3 mm^3 and 0.1^3 mm^3 are smaller than those between 1.0^3 mm^3 and 0.1^3 mm^3 .

The difference in the PTV at 95 Gy is -3.93% and -0.853% for 1.0^3 mm^3 and 0.5^3 mm^3 , respectively.

4.7.3 Discussion

The figures shown in this section demonstrate that for voxel sizes between 0.1^3 and 1.0^3 mm^3 , there is little difference in the optimal Pareto front found by the optimiser. However, a decrease of accuracy starts to become noticeable for voxels of 1.0^3 mm^3 .

The difference of -3.93% at 95 Gy for 1.0^3 mm^3 compared to 0.1^3 mm^3 could mean that where the dosimetry calculated with voxel sizes of 0.1^3 mm^3 might consider 98% of the PTV to be sufficiently irradiated (and therefore consider the treatment plan to be acceptable), the dosimetry calculated with voxel sizes of 1.0^3 mm^3 might consider 94% to be sufficiently irradiated (and so the plan would be deemed unacceptable). For this reason, it was decided that voxel sizes of 1.0^3 mm^3 did not provide enough accuracy. Conversely, with differences of less than 1.5%, it was decided that 0.5^3 mm^3 was sufficiently accurate.

Calculation times, as one might expect, increase significantly upon decreasing voxel size. This is because decreasing each side of the voxel increases the number of voxels cubically (for example, decreasing voxel sizes from 0.5^3 mm^3 to 0.1^3 mm^3 increases the number of voxels by a factor 125).

It should also be noted that the number of generations required before convergence also increases as voxel size decreases. This is because injections are always placed at the centre of a voxel, so increasing the number of voxels increases number of potential injection positions, giving the optimiser greater liberty when searching for the optimal plan.

Lastly, it can be seen that the computer memory required to perform treatment planning with voxel sizes of 0.1^3 mm^3 is extremely large. A cubic dose grid with sides of 45 mm (a spherical PTV with a radius of 10 mm has a diameter of 20 mm, plus the margin of 12.5 mm on all sides to ensure that all irradiation is accounted for) and voxel size of 0.1^3 mm^3 consists of 91.1 million voxels. For a population of 80 treatment plans, there is a total of 7.29 billion voxels. If each voxel is stored in double-precision floating-point format (8B), then the total memory required is over 54.3 GB compared to 0.435 GB for voxel sizes of 0.5^3 mm^3 . As larger tumours are treated, this problem would be further exemplified.

4.7.4 Conclusion

In conclusion, it was decided that 0.5^3 mm^3 was the correct compromise of voxel size; it is neither too large as to lead to imprecise dosimetry, nor too small as to unnecessarily slow down the optimisation process.

4.8 Optimisation of a real tumour

Now that the algorithm has been selected and the study of the best voxel size has been carried out, the optimisation of a real tumour can be performed. This optimisation was performed on a glioblastoma implanted in the skull of a mini-pig.

The validation of the optimisation process in the preclinical trials was particularly challenging. Mini-pigs are considerably smaller than humans and so their tumours are smaller too. This means that the percentage growth rate of tumours in the mini-pigs is much greater compared to humans.

Consequently, the CT imagery of the tumours needed to be performed just before the treatment; images taken just a few days prior to the treatment quickly became redundant because the tumour changed so rapidly. The CT scan was therefore performed ~30 minutes prior to the start of the treatment.

Treatment planning would have to take place in the time between the completion of the CT scan and the start of the treatment. This meant that the whole process – from the segmentation of the tumour, to the optimisation of the Pareto front, to the selection and validation of a treatment plan – needed to happen in under 30 minutes.

All visualisations that are shown in this section were programmed in C++ using VTK, an open-source library which is further explored in the Appendix.

4.8.1 Tumour segmentation

Below, in Figure 4.20, the CT of the skull of a mini-pig can be seen. A glioblastoma was induced in the mini-pig, and the segmentation of this tumour is shown in pink. This segmentation was performed using 3D Slicer.

The tumour in question is small compared to a typical glioblastoma, the GTV was just 2.28 cm³. It is, however, a real glioblastoma, and serves as a good starting point for this research.

The limits of the dose grid are also shown in Figure 4.20, and it can be seen that much of the dose grid extends beyond the skull of the mini-pig. As will be seen in the absorbed dose distribution of the optimised treatment plan, however, the irradiation with microbrachytherapy is so localised that the dose grid is larger than necessary. Furthermore, humans are considerably bigger than mini-pigs, and so a margin of 12.5 mm around the PTV will likely remain entirely inside of the brain for humans even if this is not the case for the mini-pig shown below.




Figure 4.20: 3D visualisation of the skull of a mini-pig as seen from three angles. The glioblastoma has been segmented and is shown in pink. The limits of the dose grid are also shown.

4.8.2 Treatment planning

Once the PTV had been segmented, treatment planning was performed. As with the second test case of this chapter, the number of injections was allowed to vary between 1 and 20. 5 μL injections of ^{90}Y with A_0 's of 5 MBq were used. NSGA-II was used in MO mode, so as to provide results giving the most information possible.

The optimisation was performed over a total of 100 generations with a population of 600. The process finished in 398 s, after a total of 30,600 treatment plan updates. The Pareto front for this optimisation is shown below, in Figure 4.21.




Figure 4.21: Pareto front for the treatment planning of a real tumour. The treatment plan that was deemed the best according to the user (because it uses the fewest injections amongst all the acceptable plans) is shown with the vertical cross, whereas the best SO plan is shown with the diagonal cross.

4.8.3 Results from the chosen treatment plan

From the presented Pareto front, one treatment plan must be chosen. The treatment plan that was selected was the plan that used the fewest injections whilst sufficiently irradiating the PTV (the first acceptable plan). This plan is highlighted with the vertical cross. The diagonal cross highlights the plan that has the lowest SO cost. This, again, highlights the extra choice that is given to the user by using an MO algorithm.

The selected treatment plan used 9 injections and left just 2.17% of the PTV under-irradiated. The results of this treatment plan are shown in Figure 4.22, Figure 4.23 and Figure 4.24.

Figure 4.22 shows the distribution of the 9 injections throughout the PTV. Figure 4.23 shows slices of the plan's absorbed dose distribution, with slices shown every 1 mm (every 2 slices). Finally, Figure 4.24 shows the DVH of the PTV for the selected treatment plan.




Figure 4.22: 3D distribution of injections throughout the PTV as seen from three angles.




Figure 4.23: Slices through the absorbed dose distribution. The contours of the skull are shown in white and the contours of the PTV are shown in red. Slices are shown every 2 slices (1 mm).




Figure 4.24: DVH for the PTV of the selected plan. The PTV's absorbed dose requirement is also shown with the black diamond. It can be seen that the plan corresponding to this DVH is acceptable because it passes over the black diamond.

4.8.4 Discussion

The slices of the absorbed dose distribution, shown in Figure 4.23, confirm the assumption that was made when it was seen that the dose grid protruded from the mini-pig's skull: the absorbed dose distribution from microbrachytherapy injections is extremely concentrated about the source, so much so that no visible amount of radiation escapes its skull.

It can, however, be seen that bone is irradiated. This means that the assumption of homogeneity that was discussed in the chapter on dosimetry, in Section 3.5.1, is not valid and therefore the displayed dosimetry within the skull is not accurate. In reality, bone is denser than soft tissue, and so the radiation would not have travelled as far as is displayed. In this example, the tumour is very close to the edge of the skull (just 1 mm on slices 56 through 66). For a human, the brain is considerably bigger and so it is less likely that the tumour would be so close to the skull. Therefore, the dosimetric inaccuracy that occurred by irradiating the skull was not considered a serious problem.

Lastly, it can be seen in Figure 4.23 that there is a non-negligible irradiation of healthy tissue surrounding the PTV. This is to be expected as no cost function was applied to optimise the irradiation of healthy tissue. This could have been done, for example, with the inclusion of a cost function on the CI.

It can be seen in the DVH, shown in Figure 4.24, that the selected treatment plan does sufficiently irradiate the PTV as it passes over the top of the point (95 Gy, 95%).

4.9 Conclusion

In Sections 2.5.2 through 2.5.4, the following algorithms were presented: NMS, BA and NSGA-II. This chapter then performed a comparison between these algorithms.

For the comparison, the PTV and OARs described in Section 1.7 were used. The algorithms were used to optimise the irradiation of the PTV, the irradiation of the OAR and the number of injections.

The injections used throughout this comparison were 5 μL injections of ^{90}Y with an A_0 's of 5 MBq. Chapter 5 consists of a comparison of radionuclides and A_0 's, however any combination will suffice for this chapter, since a good algorithm should be able to find optimal treatment plans, regardless of tumour size, radionuclide or A_0 .

It was shown that BA and NSGA-II produced satisfactory plans. BA seemed to be more efficient for a fixed number of injections, whereas NSGA-II was more efficient when the number of injections is allowed to vary. A possible explanation for this is that BA puts a lot of resources into exploring relatively few areas of the search-space; as can be seen in Figure 4.3 and Figure 4.8, it took around 200 treatment plan updates to explore each flower patch. However, NSGA-II explores the search-space in a more diverse manner, which is advantageous when the number of injections is variable and so the search-space is increased.

Treatment plans produced by NMS were of considerably lower quality. As discussed in Section 4.6.4, the quality of the results produced by NMS could be improved by using restarts (at the expense of optimisation times) and optimisation times could be improved by using a meta-optimiser so as not to test every number of injections.

Given NSGA-II's tendency to produce the best quality treatment plans, coupled with the fact that it is an MO algorithm and so can be used even when relative costlet weightings are unknown, NSGA-II was chosen as the best algorithm. It was decided that from thereon, NSGA-II would be used for all treatment plan optimisations.

Once NSGA-II had been selected as the best algorithm, a study was undertaken into the effect of voxel size on dosimetry accuracy. It was found that 0.5^3 mm^3 was the best compromise of voxel size; 0.1^3 mm^3 was too slow and required too much memory, whereas 1.0^3 mm^3 was not accurate enough.

Following this, the optimisation of a real glioblastoma implanted in the brain of a mini-pig was performed, permitting the validation of the TPS. At the end of an optimisation of 100 generations with a population of 600, a satisfactory treatment plan was found that sufficiently irradiated the PTV with just 9 injections. The results of the optimisation, including the 3D distribution of injections, the 2D slices of the absorbed dose distribution and the DVH were shown.

Now that a treatment planning method has been established, it is possible to develop certain aspects of the injection, such as the best radioelement, injection volume and initial activity. This will be the subject of Chapter 5.

5 COMPARISON OF INJECTIONS

5.1 Introduction

The microspheres used in microbrachytherapy are injected inside the tumour. The goal is therefore that the radioelements deposit their energy such that the entirety of the PTV is sufficiently irradiated but the irradiation of the surrounding healthy tissue, particularly the OARs, is minimised. Many parameters can affect the efficacy of the injections, and now that methods of dosimetry and treatment planning have been developed and put in place, these parameters can be evaluated.

The structure of this chapter is as follows:

- Section 5.2 discusses the choice of the radioelement;
- Section 5.3 discusses the initial activity and the volume of each injection;
- Section 5.4 uses the optimal radioelement, initial activity and injection volume to study the possibility of using microbrachytherapy to treat larger tumours; and
- Section 5.5 discusses modifying the form of the injections, passing from spherical to cylindrical injections.
- Section 5.6 concludes the findings of this chapter.

Throughout this chapter, whenever an injection parameter is analysed, the other parameters are held constant. For example, when the effect of the volume of each injection is being studied, both the radioelement and A_0 are held constant. With the exception of Section 5.4, the PTV used throughout this chapter is a 10 mm-radius sphere.

5.2 The choice of radioelement

The choice of radioelement is clearly of utmost importance for microbrachytherapy. A radioelement that deposits its energy too locally will not sufficiently irradiate the tumour (and so, too many injections will be required), whereas a radioelement that deposits its energy over a large volume will potentially over-irradiate the surrounding tissue.

The chosen radioelement would inevitably be a β^- -emitter; γ -emitters depositing their energy too far from the source and α -emitters depositing it too close. A non-exhaustive list of β^- -emitting radionuclides is shown in Table 5.1.

Radionuclide	Half-life	E_{β}^{mean}	Range(E_{β}^{mean})	$E_{\gamma}^{principal}$ (%)
^{90}Y	2.7	935	4.02	–
^{194}Ir	0.8	808	3.34	328 (13)
^{188}Re	0.7	795	3.28	155 (15)
^{32}P	14.3	695	2.75	–
^{166}Ho	1.1	666	2.60	80 (6)
^{89}Sr	50.5	583	2.18	–
^{165}Dy	0.1	438	1.47	95 (4)
^{109}Pd	0.6	360	1.10	88 (4)
^{111}Ag	7.5	350	1.06	342 (7)
^{186}Re	3.7	329	0.97	137 (10)
^{159}Gd	0.8	311	0.89	363 (8)
^{198}Au	2.7	311	0.89	412 (95)
^{77}As	1.6	228	0.55	239 (1.6)
^{153}Sm	1.9	225	0.54	103 (28)
^{105}Rh	1.5	190	0.41	319 (19)
^{131}I	8.0	181	0.38	364 (81)
^{47}Sc	3.3	162	0.32	159 (68)
^{199}Au	3.1	143	0.26	158 (37)
^{67}Cu	2.6	141	0.25	185 (49)
^{177}Lu	6.7	140	0.25	208 (11)
^{169}Er	9.3	111	0.17	–

Table 5.1: Non-exhaustive list of β^- -emitting radionuclides. E_{β}^{mean} from (Baum, 2014). Range calculated in water, with interpolation from (ICRU, 1984). Half-lives, energies and ranges are expressed in days, keV and mm, respectively.

A comparison of some of these radioelements will be conducted. Of all of the radioelements presented in Table 5.1, four were chosen to be compared – ^{177}Lu , ^{131}I , ^{166}Ho and ^{90}Y . The radioelements were chosen so as to represent a wide range of β^- -emitting radioelements. The properties that were taken into account during the selection process were: high E_{β}^{mean} , short half-life, γ -emissions of sufficiently high energy to enable imaging and the radioelement's abundance in the domain of medical physics.

The probability spectra of β^- -emissions of the four selected radioelements are shown below, in Figure 5.1.




Figure 5.1: Probability spectra of β^- -emissions for ^{177}Lu , ^{131}I , ^{166}Ho and ^{90}Y . Data from MIRD (Eckerman and Endo, 1989).

5.2.1 Kernel generation

The kernels for the four radioelements were calculated with GATE using the method described in Section 3.2.4. Barring the energy spectra, the input parameters for GATE were identical for each radioelement.

The injections were modelled as $5\ \mu\text{L}$ spheres. One-dimensional projections of the kernels with A_0 's of 5 MBq are shown in Figure 5.2.




Figure 5.2: Absorbed dose as a function of distance for single 5 μL injections with A_0 's of 5 MBq for the radioelements ^{177}Lu , ^{131}I , ^{166}Ho and ^{90}Y . The physical limit of the injection (a sphere of 5 μL has a radius of 1.06 mm) is also shown.

5.2.2 Pareto optimisation

A comparison of the radioelements was then conducted by optimising $Cost_{PTV}$ and $Cost_{Inj}$. 5 μL injections were used, each with an A_0 of 20 MBq.

The maximum number of generations was set to 1,000 and the number of injections was constrained between 1 and 100. The maximum was set at 100 because greater than this was deemed unfeasible in a clinical environment.




Figure 5.3: Pareto front for optimisation of the PTV with 5 μL injections of ^{177}Lu , ^{131}I , ^{166}Ho and ^{90}Y each with an A_0 of 20 MBq.

5.2.3 DVH comparison

In Figure 5.4, a DVH is shown for each of the four radioelements. The plans shown for ^{131}I , ^{166}Ho and ^{90}Y are the plans that used the fewest injections whilst sufficiently irradiating the PTV, corresponding to 28, 19 and 8 injections, respectively. ^{177}Lu did not satisfy the PTV requirement, so the DVH shown in Figure 5.4 uses the maximum number of injections, 100.




Figure 5.4: DVHs for the plans that use the fewest injections whilst sufficiently irradiating the PTV ($Cost_{PTV} \leq 5\%$). Since it was not possible to achieve this with the maximum number of injections with ^{177}Lu , the DVH for 100 injections is shown. For this reason, ^{131}I , ^{166}Ho and ^{90}Y pass over the diamond at (95 Gy, 95%), whereas ^{177}Lu passes underneath.

5.2.4 Discussion

Figure 5.2 shows that despite delivering high absorbed doses close to the injection centre, the irradiation of ^{177}Lu and ^{131}I are concentrated around the source. For an A_0 of 5 MBq, the distances at which the absorbed dose drops beneath 100 Gy (D_T) for ^{90}Y , ^{166}Ho , ^{131}I and ^{177}Lu are 4.51, 3.03, 2.01 and 1.70 mm, respectively.

It can be seen in Figure 5.3 that ^{177}Lu was the only radioelement out of the four that was incapable of sufficiently irradiating the PTV without surpassing the pre-defined maximum number of injections, 100. This is due to its low-energy β^- -emissions. Of the remaining radioelements, ^{90}Y sufficiently irradiated the tumour with 8 injections, whereas ^{166}Ho and ^{131}I required 19 and 28 injections, respectively.

It can be seen that, although three of the four radioelements are capable of sufficiently irradiating the PTV, using the radioelement with the highest energy β^- -emissions permits a reduction of the number of injections.

5.2.5 Conclusion

In conclusion, ^{90}Y , owing to its high energy emissions, could be considered the ideal radioelement. This would mean that the fewest injections would be required, and that surrounding healthy tissue would be little exposed.

The tested PTV was relatively small (a sphere with a radius of 10 mm), and so ^{90}Y is the only radioelement of the four tested that could be realistically scaled up to larger PTVs.

Furthermore, since ^{90}Y already exists in microsphere format for SIRT (see Section 1.2.3), it should be easily adapted to microbrachytherapy.

5.3 Initial activity and injected volume

5.3.1 Initial activity

Shown below, in Figure 5.5, are the 1D absorbed dose profiles of 5 μL injections of ^{90}Y with A_0 's ranging between 5 and 100 MBq. The absorbed dose profiles of the radioelements shown in Figure 5.2 are shown with a logarithmic y-axis, hiding the steepness of the gradient. For this reason, a linear y-axis is used in Figure 5.5.




Figure 5.5: 1D absorbed dose distribution for a 5 μL injection of ^{90}Y with various values of A_0 . The physical limit of the injection (a sphere of 5 μL has a radius of 1.06 mm) is also shown.

The same kernel was used for these 5 curves, since the S-values (and therefore the kernel) depend only on the radioelement and the volume of the injection. In Figure 5.5 it can be seen that as A_0 is increased, the total absorbed dose increases by the same factor. Hence, it can be seen that increasing A_0 increases the total irradiation of the patient.

It can also be seen, however, that because the gradient is so steep, the extra volume that is irradiated only increases slightly. This is better shown in Figure 5.6, which displays the same curves as in Figure 5.5, but with the maximum of the y-axis at 500 Gy.




Figure 5.6: 1D absorbed dose distribution for 5 μL injections of ^{90}Y with various values of A_0 . The y-axis of this figure is limited to a maximum of 500 Gy. The physical limit of the injection (a sphere of 5 μL has a radius of 1.06 mm) is also shown.

In Figure 5.6, it can be seen that by increasing A_0 twentyfold, from 5 to 100 MBq, the radial distance from an injection that would be sufficiently irradiated in the PTV only increases from 4.51 to 7.54 mm, a marginal increase of 3 mm.

Despite the fact that increasing A_0 causes only a marginal increase in sufficiently-irradiated volume of the PTV, this would nonetheless help decrease the number of injections required. Furthermore, if the irradiation is contained to the PTV, the increased total irradiation should not affect healthy tissue.

Shown below, in Figure 5.7, are the Pareto fronts for the optimisation of the PTV with 5 μL injections of ^{90}Y with A_0 's between 5 and 100 MBq. The number of injections was limited between 1 and 20. The optimisations were performed over 100 generations with a population of 100 treatment plans.




Figure 5.7: Pareto front as a function of A_0 . Clearly increasing A_0 reduces the number of injections that are required to sufficiently irradiate the PTV, shifting the Pareto front towards the bottom-left of the graph.

As one would expect, increasing A_0 decreases the number of injections accordingly; despite the small increase indicated in Figure 5.6, it is still enough to have an impact. By doubling the A_0 from 5 MBq to 10 MBq, the number of injections required to adequately irradiate the PTV drops from 15 down to 11. At 100 MBq, just 5 injections are required.

5.3.2 Injection volume

Following the analysis into the effects of the chosen radioelement and A_0 , a further study was conducted into the effect of the injected volume on the dosimetry. Shown below, in Figure 5.8, are the 1D absorbed dose profiles for 5 MBq injections of ^{90}Y with injection volumes ranging between 5 and 20 μL .




Figure 5.8: 1D absorbed dose profiles for single ^{90}Y injections of various volumes and A_0 of 5 MBq. The injection radii (1.06, 1.34 and 1.68 mm) are shown for each of the volumes.

One would expect that increasing the injected volume whilst guarding the activity constant would decrease the maximum absorbed dose (because the concentration of radioelements is decreased), and this can be seen in Figure 5.8.

Since the activity is constant, the integral of the three curves must be identical. This must mean that the irradiated volume must increase, and so an optimisation with injections of 20 μL should require less injections than with injections of 5 μL .

Hence, Figure 5.9 shows the Pareto front for an optimisation of the PTV with between 1 and 20 injections. The injections consisted of ^{90}Y with A_0 's of 5 MBq.




Figure 5.9: Pareto front for the optimisation of the PTV with between 1 and 20 injections of 5 MBq ^{90}Y . The volume of these injections varied between 5 and 20 μL . It can be seen that varying the injection volume has a negligible effect on the Pareto front.

5.3.3 Discussion

In Figure 5.6, it was shown that increasing A_0 only slightly increases the volume of the PTV that is sufficiently irradiated. Nonetheless, it was also shown, in Figure 5.7, that increasing A_0 does have a noticeable effect on the number of injections required.

It would therefore be advantageous to use injections with the highest A_0 possible. It is not, however, possible to increase the initial activity of the injections indefinitely, and so in reality A_0 should be as great as possible, subject to the realistic constraints of the activation method.

Following the analysis into the effect of altering A_0 , the impact of changing the injected volume was studied.

In theory, it should be beneficial to spread the radioelements over a large volume, decreasing the maximum absorbed dose in the centre of the injection and increasing the total irradiated volume (as shown in Figure 5.8). However, it was shown that doing this did not have a noticeable impact on the Pareto front shown in Figure 5.9.

Furthermore, the injected liquid will displace the tissue that occupied that volume prior to the injection. Hence, the greater the volume of the injection, the more tissue that is displaced. This displaced tissue will exert a restorative force on the injected liquid, potentially distorting its spherical distribution. If this occurs, the accuracy of the absorbed dose distribution that was calculated by the TPS will be reduced.

Conversely, although increasing the injected volume does not have a direct impact on the improvement of the Pareto front, it could be used indirectly to improve it. Since it is not possible to increase the concentration of an injection's A_0 indefinitely, increasing the volume could help increase A_0 . If the maximum attainable concentration of A_0 , for example, was 1 MBq/ μ L, increasing the injected volume would increase A_0 , decreasing the number of required injections.

It would therefore be pertinent to perform physical tests to evaluate the maximum volume that can be injected without distorting the spherical form of the injections (this could be done simply by injecting into chicken breast, for example, and analysing the injection forms with CT or MRI imagery).

Once this has been evaluated, the maximum injected volume and A_0 should be used so as to minimise the number of injections.

5.4 Larger tumours

Now that the ideal injection parameters have been discussed for microbrachytherapy, the feasibility of microbrachytherapy for larger PTVs can be explored. Specifically, treatment planning was performed, for spherical PTVs with radii of 10 mm (the original radius), as well as 20 and 30 mm.

The optimisations were performed with 20 μ L injections of ^{90}Y with A_0 's of 20 MBq. The number of injections were limited between 1 and 200, the size of the NSGA II population was 400 and the optimisations were performed for 200 generations. The resulting Pareto fronts are shown below, in Figure 5.10. Supplementary information is also given in Table 5.2.




Figure 5.10: The Pareto front for the optimisation of spherical PTVs with radii ranging between 10 and 30 mm. NSGA-II was used with 20 μL injections of ^{90}Y with A_0 's of 20 MBq. It can be seen that treatment plans can be found for PTVs smaller than 20 mm. However, 30 mm is too voluminous to be treated, even with the maximum number of injections, 200.

Radius (mm)	Volume (cm^3)	Optimisation time (s)
10	4.19	104
20	33.51	2,392
30	113.10	7,575

Table 5.2: Optimisation times as a function of the size of the PTV.

From Figure 5.10, it can be seen that the acceptable treatment plan that uses the fewest injections for the 10 mm sphere required 7 injections. For the 20 mm sphere, 50 injections were required. Lastly, for the 30 mm sphere, it was not possible to sufficiently irradiate the PTV, even with the maximum number of injections, 200.

Optimisation times increased linearly with the volume of the PTV. Although, compared to the PTV with a radius of 10 mm, it took considerably longer to optimise the PTV with a radius of 30 mm, an optimisation time of 2 hours is nonetheless acceptable.

It was discussed in (Dempsey, Condon and Hadley, 2005), that for a cohort of 70 patients, the mean volume of the glioblastoma is just greater than 34 cm^3 , however the largest among the patients was 131 cm^3 . The mean and maximum correspond roughly to the spheres with radii of 20 and 30 mm. It can therefore be seen that microbrachytherapy could be a viable treatment for average-sized glioblastomas, but the number of injections required for the largest tumours would probably make it a poor choice for those.

5.4.1 Conclusion

It was found that A_0 should be as large as possible, and that increasing the injected volume did not significantly affect the dosimetry. However, the injected volume should be as large as possible (without causing distortion of the injected sphere) so as to increase A_0 .

Lastly, treatment planning was performed for larger PTVs. It was found that microbrachytherapy could be a viable choice for average-sized glioblastomas but probably an inefficient choice for the largest glioblastomas.

5.5 Cylindrical injections

Until now, injections have been modelled as spheres. In this section, the model is developed such that cylindrical injections are used. Instead of performing the injection in a fixed position, the needle is withdrawn and a constant flow is administered from the syringe, such that the injected volume is spread over a line and the injection forms a cylinder.

By passing from spherical to cylindrical injections, the radioelements are more spread out, such that for a given number of injections, a larger volume can be treated. Furthermore, by spreading out the radioelement, the maximal absorbed doses close to injections should be decreased.

Another advantage that can be exploited by using cylindrical injections is that the volume of each injection can potentially be increased. This is because the length of the cylinder can be as long as desired without increasing the resistive pressure of the displaced tissue.

The injector pressure is controlled such that cross-sectional surface of the cylindrical injection is fixed at 0.68 mm^2 , which meant for a minimum volume of $5 \mu\text{L}$, a minimum injection length of 7.31 mm (rounded up to the nearest whole millimetre, 8 mm) was required.

5.5.1 The injector and the injections

It was assumed that injections would be performed from an “injector” needle protruding from a central “guide” needle. The injector extends from the guide, performing a cylindrical injection as it is retracted towards the guide. The trajectory of many blue injector needles from the green guide needle can be seen in Figure 5.11.

With the injector withdrawn inside of the guide, the guide can rotate and the injector can then extend to irradiate a new volume. In this fashion, a volume with the shape of the cone can be irradiated. Lastly, the depth of the guide can be altered, such that the total irradiated form vaguely resembles a Christmas tree.

Assuming that the depth of the guide needle can be controlled to the nearest 1 mm, that the injector needle rotates in increments of 20° and that the angle between the guide and the injector is 30° , all of the possible needle trajectories for the 10 mm-radius spherical PTV are shown in Figure 5.11.




Figure 5.11: 3D visualisation of all possible injector trajectories (blue) within a spherical PTV of radius 10 mm. The guide needle is shown in green.

Amongst all the proposed needle trajectories, the TPS would then need to decide which trajectories should be used and what the start and end depth of the injection should be along each trajectory. In this sense, the injection coordinates can be expressed in cylindrical coordinates: the depth of the guide, the rotation of the guide and, lastly, the extension of the start and the end of the injection.

5.5.2 Absorbed dose distribution of cylindrical injections

The dosimetry methodology had to be modified to take into account several changes brought on by passing over to cylindrical injections. Listed below are two potential problems that need to be addressed before being able to use cylindrical injections:

- The injections now have variable lengths, how should the injection kernels therefore be produced? Does a kernel need to be simulated (a time-consuming process) for each injection length?
- With spherical injections, injections were placed in the centre of voxels, meaning that the distribution never needed to be altered. Cylindrical injections will not necessarily run parallel to the voxels (the injections may start and end in arbitrary positions). How should the kernel be resampled to account for this?

The first question is addressed in Section 5.5.2.1 and the second in Section 5.5.2.2.

5.5.2.1 Generation of injection cylindrical kernels

The method of generating cylindrical kernels was much the same as those for spherical kernels, as described in Section 3.2. However, instead of modelling the radioactive source as a sphere, a cylinder with a cross-section of 0.68 mm^2 and a length of 20 mm was used.

20 mm was chosen as the length of the cylinder as this was the maximum extension of the injector. It can be seen in the 1D S-value profile (recall the simple conversion from S-value to absorbed dose in Equation 3.1) shown in Figure 5.12, that a plateau is reached just 1.5 mm after the start of the injection.




Figure 5.12: S-values through the central axis of a cylindrical injection, with cross-sectional area of 0.68 mm^2 and length of 20 mm (as shown by the black solid lines). It can be seen that the kernel plateaus ~ 1.5 mm after the start and before the end of the injection, as shown by the dashed lines. Therefore, shorter injections can be modelled by removing parts of the kernel that lie in the plateaued area.

Therefore, the central section, between 1.5 and 18.5 mm can be considered as an infinitely long cylinder, and can be removed to create injections of reduced length. For example, to create an injection of 8 mm, the section from 4 to 16 mm would be removed such that the section from -5 to 4 mm joins directly to the section from 16 to 25 mm.

By producing the kernel in this fashion, it was possible to perform the GATE simulation just once (as this is the time-consuming step) and obtain kernels for any required injection length.

5.5.2.2 Rotating and resampling the kernels

Once the kernel has been cut down to the correct size, it must then be translated and rotated, such that the injection is aligned with the desired trajectory of the injector needle (one of the blue lines in Figure 5.11). After the kernel has been translated and rotated, it is necessary to resample it, such that the voxels are aligned with the voxels of the dose grid and the injection can be convolved.

The resampling requires interpolation, which introduces the partial volume effect. This effect is highlighted in the illustrations shown in Figure 5.13.




Figure 5.13: Illustrations of the absorbed dose profile of (a) a spherical injection; (b) a cylindrical injection, centred on voxel centres; and (c) the absorbed dose distribution required for a diagonally placed cylindrical injection.

Resampling is a relatively slow process, so it was not possible to resample the kernel when required during the optimisation process. Instead, after the position of all possible injections had been calculated (shown with the blue lines in Figure 5.11), their resampled distributions were also calculated. This resampling was performed with the VTK class, “`vtkImageReslice`”, with the interpolation mode set to linear.

In this fashion, all the resampling was performed prior to the optimisation loop and then the optimisation algorithm could rapidly convolve the newly calculated distributions, much in the same way as it did with the spherical injections. The only difference for the algorithm is that before there was one kernel for all injections and now each injection has its own kernel (the resampled version of the original).

The workflow for the calculation of absorbed dose kernels prior to the optimisation loop was therefore: calculate the position of all possible injections, and then for each injection: cut down to the correct length, translate and rotate to the correct position, and then resample.

An example of a rotated and resampled kernel is given below in Figure 5.14. The kernel is shown both up to its maximum absorbed dose and capped at 100 Gy.




Figure 5.14: Absorbed dose distribution for a rotated and resampled cylindrical injection, both shown to its maximum absorbed dose and capped to 100 Gy.

5.5.3 Comparison with spherical injections

Below, in Figure 5.15, is the Pareto front for optimisations of the original PTV (radius of 10 mm) with injections of ^{166}Ho . The spherical injections had a volume of 5 μL , whereas the cylindrical injections had the form described above.

The number of injections were restricted between 1 and 100, the population was 200 and the number of generations was 100. The specific A_0 was 1 MBq/ μL . For spherical injections of 5 μL , this equates to 5 MBq. A_0 of the cylindrical injections depended on the length of each injection.




Figure 5.15: Pareto fronts for optimisations of the PTV with injections of ^{166}Ho in the form of spheres and cylinders. It can be seen that for few injections (for example, $\text{Cost}_{\text{Inj}} = 10$), using cylindrical injections is better at reducing the under-irradiated portion of the PTV (smaller Cost_{PTV}). However, for high values of Cost_{Inj} , Cost_{PTV} reaches an asymptote for cylindrical injections. This is because there are portions of the PTV that are out of reach of the cylindrical injections (for example, the upper-extremities of Figure 5.11).

5.5.4 Discussion

From Figure 5.15 it can be seen that, for fewer than 34 injections, cylindrical injections manage to better irradiate the tumour. For example, with a maximum of 10 injections, cylindrical injections would leave 20.8% of the PTV under-irradiated, whereas spherical injections would leave 66.5% under-irradiated.

However, the cylindrical curve becomes asymptotic as the number of injection passes ~ 50 . After this point, adding more injections does not significantly decrease the volume of the PTV that is under-irradiated. This is undoubtedly because, as can be seen in Figure 5.11, imposing a minimum injection length of 8 mm (which was imposed because the surface area was constant and the minimum injectable volume was set at $5 \mu\text{L}$), there are no injections that traverse the upper-extremities of the tumour, and nor are there injections in the bottom-centre.

For this reason, cylindrical injections are limited. Such a problem could, however, be circumvented. Possible solutions include:

- Decreasing the minimum injectable volume,
- Approach the tumour from multiple angles (multiple green lines in Figure 5.11),
- Increasing the cross-sectional area of the injection, and
- Sacrificing healthy tissue.

By allowing injections through healthy tissue, both at the top and the bottom of the PTV, injections could be proposed that cover the whole of the tumour. This would, of course, be at the expense of traversed healthy tissue, but this may be considered a worthwhile sacrifice.

Cylindrical injections are promising because, unlike the spherical injections proposed in this research, a realistic delivery method has also been discussed. If one could resolve the problem of not being able to access all of the tumour, they would also require less injections. This is, however, is a physical problem that goes outside the scope of this research.

5.6 Conclusion

It was found that ^{90}Y , with the highest energy of all β^- -emitting radionuclides, is the best suited radionuclide for microbrachytherapy.

A_0 for these injections should be the highest possible, so as to minimise the number of injections required to sufficiently irradiate the tumoural zone. A_0 's between 5 and 100 MBq were tested.

It was shown that increasing the injected volume did not have a noticeable impact on the Pareto front. However, increasing the injected volume could be used as a method of increasing A_0 , thus decreasing the required number of injections. Injections should therefore be as voluminous as possible whilst not causing excess pressure by the displaced tissue that the spherical shape is perturbed.

It was shown that, for 20 μL injections of ^{90}Y , with A_0 's of 20 MBq, it was possible to sufficiently irradiate realistically-sized glioblastomas. However, the largest glioblastomas seem to be too voluminous to be treated with microbrachytherapy.

Passing from spherical to cylindrical injections seems advantageous because the radioelements are spread over a larger volume and so the total number of required injections should diminish. However, it was seen that, for the discussed method of injection delivery, parts of the tumour would be left under-irradiated. If these physical problems can be resolved, cylindrical injections could be a promising method of treating larger tumours.

CONCLUSIONS AND FUTURE PERSPECTIVES

Conclusions

The main objective of this research was to develop a TPS that could be used to produce optimal treatment plans for microbrachytherapy. With that objective in mind, various steps were required. This research started with the development of dosimetry for this new treatment. Algorithms were then developed to find optimal microbrachytherapy treatment plans. These algorithms were compared, and the best-suited algorithm was selected. Injection parameter modifications were proposed with the aim of improving the efficacy of this modality. Finally, the TPS was created, permitting users to easily create microbrachytherapy treatment plans. These points are described in greater detail below.

Since a method of performing dosimetry did not previously exist for microbrachytherapy, this was the first aspect that was developed in this research. The dosimetry was performed by generating kernels with GATE, convolving these kernels into the desired injection positions and then superposing structures (the PTV and OARs) to extract DVHs.

With the dosimetric methodology in place, algorithms were developed in order to produce optimal treatment plans. These algorithms were designed to find treatment plans that sufficiently irradiated the PTV (at least 95% of the PTV should receive at least 95% of the target absorbed dose, 100 Gy) with the fewest injections possible. Constraints could also be added to protect OARs and to maximise the CI.

Once the algorithms had been developed, they could be compared. It was found that, given the advantages proposed by using an MO algorithm, NSGA-II was the best-suited for this application.

Modifications were then proposed to the injections parameters. Notably, comparisons were conducted into different radioelements, initial activities, injected volumes and the injection forms (spherical versus cylindrical). It was found that the best-suited radioelement was ^{90}Y ; the initial activity should be as great as possible; varying the injected volume (between 5 and 20 μL) has little effect on the dosimetry; and cylindrical injections are advantageous in that the pressure of the displaced tissue is reduced, but there is a risk of not being able to irradiate the upper extremities of the tumour.

With these optimised parameters, a study was conducted into the viability of microbrachytherapy for larger tumours. It was found that microbrachytherapy could be an interesting treatment for small-to-average sized glioblastomas, but too many injections would be required for it to be feasible for the largest glioblastomas.

Lastly, a TPS, named DosiToul, was developed that incorporated all the aspects that had been developed thus far in the research. As detailed in the Appendix, DosiToul was developed in C++ and permitted the user to easily generate treatment plans with the aid of a simple GUI.

Microbrachytherapy is in the preclinical testing phase and, as such, it was possible to test the TPS in a realistic environment on mini-pigs with implanted glioblastomas. This included the need to work to clinical time constraints, meaning that structure segmentation (from CT images) and treatment planning needed to be performed in the ~ 10 minutes between imaging the mini-pig and preparing the mini-pig for the treatment. In this setting, the reliability of the application and the rapidity of the algorithm were both tested, and it can be concluded that the TPS performed extremely satisfactorily, producing treatment plans within this time limit.

Future perspectives

With regards to the future perspectives of this research, the following questions are addressed below:

- What further experiments could be performed to compliment this research?
- Could microbrachytherapy be used for other tumours?
- What further developments could be made to DosiToul?
- Could the target audience of DosiToul be increased?
- How can the maintenance and documentation of DosiToul be ensured?

Physical experimentation

It should be noted that although the treatment planning methodology was tested in a preclinical setting on mini-pigs, there are certain aspects of this work that rest theoretic.

For example, once a radioelement has been selected (such as ^{90}Y), it is then necessary to find or develop it in microsphere form. This would be simple for ^{90}Y (since it is already in use for treatments with radioembolisation), but could prove to be more complicated for other radioelements. Once the radioelement is in the form of the microsphere, the maximum attainable activity should be determined. Subsequently, once the maximum activity is known, a study should be conducted into the maximum volume that can be injected whilst keeping control of the injected form (such as spherical). Increasing the volume is clearly helpful as it increases the activity of the injection, but controlling the form of the injection is imperative to assure accurate dosimetry.

Furthermore, with continued trials (both preclinical and clinical), the efficacy of microbrachytherapy will be better understood. This will enable the better establishment of the target absorbed dose that is required to destroy the tumour. Performing such trials would also increase our understanding of the radiobiological effects of microbrachytherapy, enabling the use of radiobiological cost functions.

Microbrachytherapy for abdominal tumours

Extending the methodology presented in this work such that microbrachytherapy could be proposed for abdominal tumours would certainly be an interesting future project. Glioblastomas are challenging for microbrachytherapy because they are so large. However, the dosimetric calculation is simplified by the fact that brains are extremely homogenous with regards to density.

If one wished to expand this new modality so that it could be used on a wider range of tumours, it would be necessary to take into account density inhomogeneities. This could be done with the methods of kernel distortion that were discussed in Section 3.5.1, however the calculation would have to be performed quickly. There was a wide range in optimisation times of NSGA-II throughout this research, ranging from 30 s to just over 2 hours (for the optimisation of the 30 mm-radius spherical PTV).

Realistically, optimisation times should not surpass ~10 hours at an absolute maximum. Therefore, the method of distorting the homogeneous kernels for inhomogeneous media should be rapid; perhaps with the aid of GPU acceleration this would be possible. At any rate, if a method of fast kernel distortion was put in place, the dosimetry could be performed for new areas of the body, and microbrachytherapy could potentially be used on a wider range of tumours.

Further development of DosiToul

Aside from expanding DosiToul to account for inhomogeneities, an exciting future project could be the verification of treatment plans with Monte Carlo simulations. A button could be integrated into the GUI enabled the user to perform a GATE simulation of a given microbrachytherapy treatment plan. This would certainly take a long time, but if for example cluster computing was available, the user could return the next day to compare the dosimetry of proposed treatment plan, as performed by both DosiToul and GATE. Presumably, the dosimetry performed by DosiToul is most accurate for spherical injections in homogeneous media and least accurate for cylindrical injections in inhomogeneous media (due to both the partial volume effect and the required distortion of the homogeneous kernel).

Increasing the target audience of DosiToul

Clearly, the target user of DosiToul is teams wishing to perform treatments of microbrachytherapy. To widen the demographic of people that could potentially use DosiToul, one of two things could take place: either DosiToul could be adapted for other forms of radiotherapy, or it could be made available to researchers to allow for open-source collaboration.

The former of the two could undoubtedly prove to be interesting for other developing treatments that, like microbrachytherapy, do not currently have an associated method of treatment planning. Naturally, more mature modalities of radiotherapy would already have their own form of treatment planning, and so DosiToul would be of limited interest to them.

Unfortunately, making the code available for other researches to promote collaboration would not be possible due to intellectual property issues. However, if this were not an issue, it is possible to imagine further development of DosiToul to propose different dosimetric methods, cost functions and improved optimisation algorithms, as well as other radiotherapy modalities.

DosiToul maintenance and documentation

DosiToul is currently in a state that means that it can be used by an end-user. Naturally, if modifications were made to the current delivery method of microbrachytherapy, modifications to the software also would be required. Given the private nature of the project, this would most likely be the work of a small group (such as post-doctoral project). Throughout the development of DosiToul, however, Doxygen files were created, enabling easy continuation by different future personnel.

Closing remarks

In conclusion, this research has successfully proposed a methodology of dosimetry and treatment planning for microbrachytherapy. Injection parameters have been studied and optimised. An application with a graphical-user interface was developed and the whole TPS was used in a preclinical environment under real-life time constraints. The TPS was shown to be capable of producing satisfactory treatment plans, which were then performed on numerous glioblastomas implanted in mini-pigs.

The invention of microbrachytherapy combined with developments shown in this work demonstrate the feasibility to efficiently treat inoperable tumours whilst protecting surrounding healthy tissue. It is hoped that with the aid of microbrachytherapy, tumours that currently have a poor prognosis might have better support in the near future.

CONCLUSIONS ET PERSPECTIVES FUTURES (FR)

Introduction

Ces travaux de thèse ont débuté avec l'objectif final de créer un TPS qui pourrait être utilisé pour produire les plans de traitement optimaux pour le domaine particulier de la micro-curiethérapie. L'objectif était donc de créer une application avec une interface graphique qui pourrait être facilement manipulée par l'utilisateur du produit final et qui ne demande aucune connaissance approfondie des processus dosimétriques ou algorithmiques de leur part. Le produit final devrait être simple, permettant à l'utilisateur d'effectuer le traitement avec des coordonnées d'injections proposées, tout en les aidant avec des visualisations claires.

La micro-curiethérapie est dans la phase de test préclinique et, en tant que tel, il a été possible de tester le TPS dans un environnement réaliste sur des glioblastomes implantés dans des cerveaux de mini-cochons. Dans ces conditions, la fiabilité, ainsi que la rapidité de l'algorithme ont été testées puis validées.

La structure de ce résumé est tel que :

- Premièrement, les aspects de la planification de traitement sont conclus.
- Deuxièmement, les paramètres d'injection qui pourraient être optimisés sont conclus.
- Troisièmement, la création du TPS est conclue.
- Quatrièmement, les perspectives futures de ce travail sont discutées.
- Dernièrement, l'intégralité des travaux de recherche est conclue.

La planification de traitement

Dans le cadre du développement du processus de la planification de traitement, les points suivants sont discutés :

- La dosimétrie de la micro-curiethérapie,
- Les fonctions coûts,
- Les algorithmes d'optimisation, et
- La génération d'un plan de traitement pour un vrai glioblastome.

La dosimétrie de la micro-curiethérapie

La dosimétrie de la micro-curiethérapie a été effectuée en générant un kernel caractérisant la distribution de dose absorbée pour une seule injection. Ce kernel a été créé avec GATE. Avec l'hypothèse que les injections irradiaient un volume homogène, il a été possible de faire des convolutions de ce kernel dans les positions déterminées dans le plan de traitement. Une fois les kernels en position, la dosimétrie a été effectuée en extrayant la distribution de dose absorbée du patient et les HDVs résultants.

Les simulations GATE ont été effectuées pour plusieurs radioéléments, tailles de voxel et volumes d'injection. Au total, 7 kernels ont été utilisés durant ces travaux de recherche, nécessitant plus de 71 heures de simulation sur un seul cœur d'un processeur de 2,5 GHz d'un MacBook Pro.

La convolution des kernels sur la distribution de la dose absorbée du patient (la « dose grid ») nécessitait, pour des raisons de facilité, que la taille des voxels du kernel et ceux de la dose grid soit la même. Ceci veut dire que la dosimétrie, pour la majorité des études réalisées dans ce travail, a été effectuée avec une précision de $0,5^3 \text{ mm}^3$.

Une fois la dose grid calculée pour un plan de traitement donné, la dosimétrie spécifique à chaque structure (le PTV et les OARs) a été effectuée en superposant des masques binaires de chaque structure sur la dose grid. Les structures pouvaient être importées au format « DICOM RT Structure » ou « STL ». Elles pouvaient également être créées en forme de sphères ou de coquilles ou générées à partir d'une image CT en utilisant les valeurs de densité. Après avoir superposé les masques binaires sur la dose grid, les HDVs ont été extraits pour chaque structure permettant l'analyse d'un plan de traitement

Le processus d'optimisation nécessitant un grand nombre d'analyse et une multitude de comparaison de plans de traitement, la rapidité du calcul dosimétrique est indispensable. Des efforts ont été faits pour accélérer ce calcul en excluant de la convolution des voxels contenant des doses absorbées négligeable. Il a été montré que la simulation de la distribution de dose absorbée d'une injection (sphérique) dans un kernel cubique, contient beaucoup de voxels dont les doses absorbées sont négligeables. C'est pourquoi il a été décidé de les exclure afin d'accélérer les temps de calcul d'un facteur $\sqrt{3}$.

Avec l'accélération du calcul dosimétrique, l'intégralité de ce processus pouvait être réalisée dans un temps relativement court. Par exemple, pour un plan de traitement de 20 injections de 5 μL avec ^{90}Y , une activité initiale de 5 MBq et contenant deux structures (un PTV et un OAR), la dosimétrie pour chaque structure est réalisée en 4 ms (la convolution de ces 20 injections sur la dose grid, la superposition des masques binaires de structures et l'extraction de leurs HDVs).

Les fonctions coûts

Avec la méthodologie de la micro-curiethérapie en place, les fonctions coûts ont été utilisées afin de quantifier les plans de traitement instantanément. Ceci veut dire qu'ils pouvaient être classifiés tout simplement en comparant quelques chiffres.

Il a été constaté que les fonctions coûts qui sont typiquement utilisées en radiothérapie externe ne sont pas adaptées au contexte de la micro-curiethérapie. Cela est dû à la mise en place des sources radioactives dans le patient. En effet, certaines parties de la tumeur se trouvent fortement irradiées. Ceci est souvent pénalisant en radiothérapie externe, où l'objectif est d'avoir une couverture la plus homogène possible dans la tumeur.

Afin d'éviter cela, quatre fonctions coûts ont été proposées, qui pénalisent soit la partie du PTV qui est sous-irradiée, soit la partie de l'OAR qui est sur-irradiée, soit le nombre d'injections, soit l'indice de conformité (une mesure à la fois de l'irradiation du PTV adéquate et de la sur-irradiation de tissus sain environnants).

Concernant le PTV, une fonction coût volume-dose a été utilisée pour pénaliser le pourcentage du volume recevant moins de 95% de la dose absorbée ciblée, soit 100 Gy ici. Il a été décidé qu'un plan de traitement convenable devrait laisser moins de 5% du PTV sous-irradiée. Ceci veut dire que, tant que la fonction coût (de sous-irradiation) du PTV était en dessous de 5%, le plan de traitement était convenable par rapport à l'irradiation du PTV.

Concernant l'OAR, une fonction coût volume-dose a été utilisée. Cette fonction coût pénalise le pourcentage de l'OAR recevant plus de 10 Gy. Il a été décidé qu'un plan de traitement sera considéré comme convenable si moins de 10% de l'OAR recevait plus de 10 Gy.

Les algorithmes d'optimisation

Par conséquent, des algorithmes pouvaient être développés afin de minimiser ces quatre fonctions coûts. Les fonctions coûts pouvaient être combinées ensemble dans une somme pondérée, pour être optimisées avec un algorithme simple-objectif (SO), ou optimisées simultanément, en utilisant un algorithme multi-objectifs (MO).

Trois algorithmes ont été développés et testés dans le cadre de ce travail : le « Nelder-Mead Simplex » (NMS) (SO et déterministe), le « Bees Algorithm » (BA) (SO et stochastique) et le « Non-Dominated Sorting Genetic Algorithm » (NSGA-II) (MO et stochastique). Trois cas de test ont été proposés pour l'optimisation du PTV sphérique avec un rayon de 10 mm. Le premier cas de test utilisait un nombre d'injections fixe, le second un nombre d'injections variable et le troisième un nombre d'injections variable qui incluait un OAR qui devait être protégé. L'OAR était également une sphère avec un rayon de 10 mm, placé de sorte que leurs surfaces soient en contact (le centre de l'OAR et celui du PTV sont séparés de 20 mm).

Il a été trouvé, que pour cette problématique, le NMS était le moins efficace parmi les trois algorithmes. Sa nature déterministe rend facilement influençable par une mauvaise initialisation. Comme l'initialisation était aléatoire, la qualité de cette initialisation ne pouvait donc pas être assurée. De plus, le NMS pouvait seulement optimiser un nombre d'injections fixe. Par conséquent, si l'utilisateur voulait faire une optimisation avec un nombre d'injections variable, il fallait lancer l'algorithme pour chaque nombre d'injection fixe (soit 20 fois pour une optimisation qui comprend 1 à 20 injections). Ceci veut dire que le NMS était plus lent que les autres algorithmes. Il serait possible de combiner le NMS avec un méta-optimiseur afin de réduire le nombre d'injections testé (diminuant ainsi les temps d'optimisation), mais il serait malgré tout plus lent que les deux autres algorithmes. Ceci, combiné avec la mauvaise cohérence de ses résultats a fait que le NMS n'a pas été retenu en tant qu'algorithme le mieux-adapté pour la planification de micro-curiethérapie.

Parmi les deux algorithmes restants, il a été montré qu'à la fois le BA et le NSGA-II produisaient des résultats satisfaisants. Par contre, le NSGA-II, étant un algorithme MO, montrait plus d'informations à l'utilisateur. Ces informations étaient visualisées en forme dite de « frontière de Pareto », qui est le compromis entre chaque fonction coût (permettant à l'utilisateur de choisir le meilleur plan parmi un grand nombre fixe de possibilités d'injections). En permettant une gestion de tous ces compromis, l'utilisateur est donc libre de décider quel plan de traitement est considéré comme optimal pour chaque patient. Avec des algorithmes SO, inversement, cette décision est prise par l'algorithme en fonction des pondérations relatives choisies entre les fonctions coûts.

Dans le troisième cas test (le plus compliqué parmi les trois), les algorithmes ont été lancés 300 fois chacun. Dans 100% des cas, le NSGA-II a trouvé des plans de traitement convenables (à la fois par rapport aux exigences du PTV et de l'OAR), avec un temps de calcul moyen de 32,8 s. C'est pourquoi le NSGA-II a été choisi comme étant l'algorithme optimal à utiliser par la suite pour ces travaux de recherche afin de réaliser les planifications de traitement.

Après avoir sélectionné le NSGA-II, une étude a été effectuée sur la taille idéale du voxel dans le cadre de la planification de traitement en micro-curiethérapie. Des voxels de $0,1^3$, $0,5^3$, et $1,0^3$ mm³ ont été testés et il a été trouvé que des voxels de $0,5^3$ mm³ donnaient le meilleur compromis entre la vitesse et la précision de la dosimétrie.

Les paramètres d'injection

Une fois la méthode de planification de traitement développée, les aspects du traitement lui-même ont pu être analysés. Dans cette section, les points suivants ont été discutés :

- Le radioélément injecté,
- Son activité initiale (A_0),
- Le volume de chaque injection, et
- La forme des injections : sphériques ou cylindriques.

Le radioélément injecté

Par rapport au radioélément injecté, une comparaison a été faite entre quatre radioéléments émetteurs β^- . Ces radioéléments ont été choisis afin de représenter une large gamme de radioéléments. Les critères définis sont les suivants : une combinaison d'une E_{β}^{mean} élevée, une demi-vie courte, des émissions γ d'énergies suffisamment élevées permettant l'imagerie et l'abondance du radioélément dans le domaine de la physique médicale. Les quatre radioéléments suivants ont été choisis : ^{90}Y , ^{166}Ho , ^{131}I et ^{177}Lu .

Il a été montré que l'émission des particules β^- de haute énergie était le facteur principal du choix du radioélément. Par exemple une comparaison a été effectuée avec le PTV (toujours une sphère avec un rayon de 10 mm) et des injections de 5 μL avec $A_0 = 5 \text{ MBq}$. Le ^{177}Lu émettant des particules β^- de plus basses énergies parmi les quatre radioéléments comparés, n'était pas capable d'irradier suffisamment le PTV avec moins de 100 injections, qui était le nombre maximal permis. Afin de pallier ce problème, il a été choisi un radioélément avec une émission d'énergies plus élevées, le ^{90}Y , dans le but de réduire ainsi le nombre d'injections à 8 pour ce même problème.

L'activité initiale

Par la suite, une étude a été effectuée pour évaluer l'influence de A_0 sur la dosimétrie. Il a été constaté que, l'augmentation de A_0 entraîne une augmentation du volume autour d'une injection qui est suffisamment irradiée car la dose absorbée est directement proportionnelle à A_0 . Cependant, il a également été démontré que le gradient de dose absorbée était tellement abrupt que, augmenter A_0 par un facteur 20 (de 5 à 100 MBq) avait peu d'influence sur la distance radiale (de 4,51 à 7,54 mm). Malgré cette petite augmentation, il a été montré que l'augmentation de A_0 réduisait le nombre d'injections requis pour suffisamment irradier le PTV. Ainsi il a été recommandé que A_0 soit la plus élevée possible.

Le volume injecté

Subséquentement, une étude a été effectuée sur l'effet d'augmenter le volume par injection. Les volumes testés étaient de 5, 10 et 20 μL , avec une A_0 de 5 MBq. Il a été relevé que, même si en théorie un volume plus important devait être suffisamment irradié par l'augmentation de ce volume d'injection (car les radioéléments sont étalés sur un volume plus important), l'amélioration de la frontière de Pareto était négligeable. Par conséquent, la seule raison qui justifie une augmentation du volume injecté est d'augmenter A_0 de l'injection (en supposant une concentration de A_0 par μL fixe).

C'est pourquoi, un volume conséquent doit être injecté (afin d'augmenter A_0) sans distorsion de la sphère injectée. Ce volume maximal devrait être déterminé expérimentalement.

La planification de traitement pour les tumeurs plus volumineuses

Parmi les trois sous-sections précédentes, il a été démontré que la configuration optimale doit utiliser les injections de ^{90}Y . Ces injections doivent utiliser A_0 la plus élevée possible et les injections doivent être les plus volumineuses possible.

Avec ces paramètres optimaux en place, la planification de traitement a été effectuée pour une série de tumeurs comprenant plusieurs volumes, allant de 10 à 30 mm de rayon. Un nombre maximal de 200 injections d' ^{90}Y de 20 μL et de 20 MBq a été utilisé pour ce test.

Il a été constaté que 7 et 50 injections sont nécessaires pour une couverture adéquate des PTVs de rayons 10 et 20 mm, respectivement. Cependant, il a été trouvé que même avec 200 injections (le maximum permis), il n'était pas possible de suffisamment irradier une tumeur qui possède un rayon de 30 mm.

Il a été noté qu'une sphère de 20 mm est la taille moyenne des glioblastomes, tandis que 30 mm représente les glioblastomes les plus volumineux. Pour cette raison, la micro-curiothérapie pourrait être envisagée pour les glioblastomes de tailles moyennes, mais probablement pas pour les glioblastomes les plus volumineux. Cela dépend, bien sûr, de l'efficacité de l'injecteur : en effet, s'il pouvait effectuer des injections rapidement, le nombre d'injections ne serait plus une limite.

Les injections cylindriques

L'étape suivante a consisté à optimiser la technique d'injection. Au lieu d'insérer l'aiguille en plaçant l'injection dans un seul endroit, l'injection était effectuée pendant le retrait de l'aiguille. Dans cette configuration, le volume injecté ressemblait plus à un cylindre qu'à une sphère.

Les injections cylindriques ont été proposées car cela veut dire que le volume injecté était répandu sur une surface plus large, diminuant ainsi les forces de résistance au sein des tissus déplacés. De plus, non seulement les forces de résistance diminuaient, mais en épargnant les injections, des volumes plus importants pourraient être irradiés par injection.

Il a été montré que, pour peu d'injections, les injections cylindriques étaient capables d'irradier un volume plus important. Cependant, cette technique d'injection rendait les extrémités supérieures du PTV inaccessible. Ainsi, lorsque le nombre d'injections augmentait, une « asymptote » était atteinte, sans augmenter significativement la partie du PTV que l'on souhaitait irradier.

Pour circonscrire ce problème, des injections plus courtes pourraient être utilisées, soit en augmentant la surface du cylindre, soit en diminuant le volume injecté minimale de 5 μL . Les injections pourraient donc être partiellement effectuées dans les tissus sains (en les sacrifiant), ou bien la tumeur pourrait être abordée par plusieurs angles avec l'aiguille principale.

La création du TPS

Finalement, avec des méthodes de dosimétrie et de la planification de traitement en place, et avec les paramètres d'injection optimisés, une application avec une interface graphique a été développée. L'application, codée dans son intégralité à partir de rien, est écrite en langage C++ orienté objet.

Le code utilisé pour créer l'application, appelée « DosiToul », est constitué de 133 fichiers source et il contient plus de 11,000 lignes de code. Le code exploite trois bibliothèques « open source » : Qt, VTK et GDCM. Qt a été utilisée afin de créer l'interface et gérer l'interaction avec l'utilisateur. VTK a été utilisée du rééchantillonnage des matrices 3D à la visualisation des données. Enfin, GDCM a été utilisée pour lire les fichiers « DICOM RT Structure ».

DosiToul utilise une base de données pour stocker les patients, les études et les séries des images DICOM. Chaque série peut contenir un ensemble de structures (par exemple le PTV ou les OARs) et un ensemble d'optimisations. L'utilisateur peut créer des nouvelles optimisations, les paramétrer avec un assistant de quatre pages, effectuer des optimisations et puis afficher les résultats.

Les résultats sont affichés de façon à ce que l'utilisateur puisse choisir son plan de traitement adapté parmi les solutions optimales au sens de Pareto, puis afficher la distribution de dose absorbée, les injections dans l'espace 3D (par rapport au PTV et aux OARs), les HDVs et une liste d'injections pour le plan de traitement. La base de données peut être enregistrée sur le disque dur, permettant à l'utilisateur de continuer son travail a posteriori.

Perspectives futures

Par rapport aux perspectives futures de ce travail, deux points sont discutés :

- Le traitement des tumeurs abdominale et
- L'expérimentation physique.

Le traitement des tumeurs abdominales

L'adaptation de la méthodologie présentée dans ce travail vers la prise en charge des tumeurs abdominales pourrait être un projet futur intéressant. Les glioblastomes sont difficiles pour la micro-curiethérapie car ils sont souvent volumineux. Par contre, le calcul de la dosimétrie est simplifié du fait que le cerveau est extrêmement homogène par rapport à la densité.

Pour développer cette nouvelle modalité afin de l'utiliser avec d'autres type de tumeurs, il est nécessaire de prendre en compte les inhomogénéités de densité. Cela pourrait être réalisé avec la méthodologie proposée dans ce travail de recherche, cependant il faudrait que ce processus soit rapide. Dans un contexte clinique, les temps d'optimisation ne devraient pas dépasser ~10 heures (cas extrême).

La méthode de distorsion du kernel homogène pour les milieux inhomogènes devrait donc être rapide (par exemple, en exploitant la technologie GPU). Dans tous les cas, si une méthode de distorsion de kernel rapide était mise en place, la dosimétrie pourrait être effectuée pour d'autres régions anatomiques, et la micro-curiethérapie pourrait être utilisée potentiellement sur d'autres tumeurs.

L'expérimentation physique

Enfin, il doit être noté que même si la méthodologie de planification de traitement a été testée dans un environnement préclinique sur des mini-cochons, il y a des aspects de ce travail restent théoriques.

Par exemple, une fois un radioélément sélectionné (tel que ^{90}Y), il est ensuite nécessaire de le développer en forme de microsphère. Cela devrait être simple pour ^{90}Y car il est déjà utilisé pour des traitements de radioembolisation, mais cela n'est pas le cas pour tous les radioéléments.

Une fois le radioélément en forme de microsphère, l'activité maximale atteignable devrait être donnée par le constructeur. Subséquemment, une étude devrait être effectuée sur le volume maximal qui pourrait être injecté sans la distorsion du liquide injecté.

En outre, avec plus d'essais (à la fois précliniques et cliniques), l'efficacité de la micro-curiethérapie sera mieux comprise. Cela permettra une optimisation de la détermination de la valeur de la dose absorbée ciblée nécessaire pour détruire la tumeur. Poursuivre de tels essais permettrait d'approfondir la connaissance des effets radio-biologiques de la micro-curiethérapie, permettant l'utilisation des fonctions coûts radio-biologiques.

Conclusion

Pour conclure, ces travaux de recherche ont permis de mettre en œuvre une méthodologie dosimétrique et une planification de traitement optimale pour la micro-curiethérapie. Des paramètres d'injection ont été étudiés et optimisés et une application avec une interface graphique utilisateur a été développée. L'intégralité de ce TPS a été utilisée et validée dans un environnement préclinique avec des contraintes de temps réelles, tout en étant capable de produire des plans de traitement convenables qui a permis le traitement de plusieurs glioblastomes, implantés dans des cerveaux de mini-cochons.

Cette nouvelle technologie associée aux développements réalisés au cours de ces travaux démontre la faisabilité de l'irradiation efficace des tumeurs inopérables, tout en préservant les tissus sains environnants, avec la micro-curiethérapie. De telles tumeurs, ont encore de mauvais pronostics aujourd'hui, et elles pourront sûrement être mieux prises en charge dans un proche avenir.

BIBLIOGRAPHY

Agostinelli, S. *et al.* (2003) 'GEANT4 - A simulation toolkit', *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 506(3), pp. 250–303. doi: 10.1016/S0168-9002(03)01368-8.

Ahnesjö, A., Andreo, P. and Brahme, A. (1987) 'Calculation and Application of Point Spread Functions for Treatment Planning with High Energy Photon Beams', *Acta Oncologica*, 26(1), pp. 49–56. doi: 10.3109/02841868709092978.

Astrahan, M. (2008) 'Some implications of linear-quadratic-linear radiation dose-response with regard to hypofractionation.', *Medical physics*, 35(9), pp. 4161–4172. doi: 10.1118/1.2969065.

Baum, R. P. (2014) *Therapeutic Nuclear Medicine*. doi: 10.1007/978-3-540-36719-2.

Bleeker, F. E., Molenaar, R. J. and Leenstra, S. (2012) 'Recent advances in the molecular understanding of glioblastoma', *Journal of Neuro-Oncology*, 108(1), pp. 11–27. doi: 10.1007/s11060-011-0793-0.

Bult, W. *et al.* (2012) 'Radioactive holmium acetylacetonate microspheres for interstitial microbrachytherapy: An in vitro and in vivo stability study', *Pharmaceutical Research*, 29(3), pp. 827–836. doi: 10.1007/s11095-011-0610-7.

Dempsey, M. F., Condon, B. R. and Hadley, D. M. (2005) 'Measurement of tumor "size" in recurrent malignant glioma: 1D, 2D, or 3D?', *AJNR American journal of neuroradiology*, 26(4), pp. 770–776. Available at: [papers3://publication/uuid/4EC4CDC3-5A54-4C62-8E32-2E3DEE8E326F](https://pubs.rsna.org/doi/10.3174/ajnr.A3266).

Dieudonné, A. *et al.* (2011) 'Fine-Resolution Voxel S Values for Constructing Absorbed Dose Distributions at Variable Voxel Size', 51(10), pp. 1600–1607. doi: 10.2967/jnumed.110.077149.Fine-Resolution.

Eckerman, K. F. and Endo, A. (1989) *MIRD: Radionuclide Data and Decay Schemes*. Society of Nuclear Medicine (Advancing molecular imaging and therapy). Available at: <https://books.google.fr/books?id=zEyo06YS5NMC>.

Gay, H. a. and Niemierko, A. (2007) 'A free program for calculating EUD-based NTCP and TCP in external beam radiotherapy', *Physica Medica*, 23(3–4), pp. 115–125. doi: 10.1016/j.ejmp.2007.07.001.

Hestenes, M. R. and Stiefel, E. (1952) 'Methods of Conjugate Gradients for Solving Linear Systems', *Journal of Research of the National Bureau of Standards*, pp. 409–436. doi: 10.6028/jres.049.044.

ICRU (1984) 'Stopping powers for electrons and positrons (Report 37)', pp. 206–207.

Jan, S. *et al.* (2004) 'GATE - Geant4 Application for Tomographic Emission: a simulation toolkit for PET and SPECT', *Physics in Medicine and Biology*, 49(19), pp. 4543–4561. Available at: <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC3267383/>.

Kiefer, J. (1953) 'Sequential minimax search for a maximum', *Proceedings of the American Mathematical Society*, 4, pp. 502–506. doi: 10.1090/S0002-9939-1953-0055639-3.

Lim, L. *et al.* (2005) 'A prospective evaluation of treatment with Selective Internal Radiation Therapy (SIR-spheres) in patients with unresectable liver metastases from colorectal cancer previously treated with 5-FU based chemotherapy', *BMC Cancer*, 5(1), p. 132. doi: 10.1186/1471-2407-5-132.

Malaterre, M. (2008) 'Grassroots DICOM'. Available at: <http://gdcm.sourceforge.net/gdcm.pdf>.

McGary, J. E., Grant 3rd. W., W. and Woo, S. Y. (2000) 'Applying the equivalent uniform dose formulation based on the linear-quadratic model to inhomogeneous tumor dose distributions: Caution for analyzing and reporting.', *Journal of applied clinical medical physics [electronic resource] / American College of Medical Physics*, 1(4), pp. 126–137. doi: 10.1120/1.1314839.

Nelder, J. A. and Mead, R. (1965) 'A simplex method for function minimization', *The computer journal*. Br Computer Soc, 7(4), pp. 308–313.

Nickolls, J. *et al.* (2008) ‘Scalable parallel programming with CUDA’, *AMC Queue*, 6(April), pp. 40–53. doi: 10.1145/1365490.1365500.

Pham, D. T. *et al.* (2006) ‘The bees algorithm—a novel tool for complex optimisation problems’, *Proceedings of the 2nd International Virtual Conference on Intelligent Production Machines and Systems (IPROMS 2006)*, pp. 454–459. doi: <http://dx.doi.org/10.1016/B978-008045157-2/50081-X>.

Salem, R. *et al.* (2002) ‘Yttrium-90 Microspheres: Radiation Therapy for Unresectable Liver Cancer’, *Journal of Vascular and Interventional Radiology*, 13(9), pp. S223–S229. doi: 10.1016/S1051-0443(07)61790-4.

Sarrut, D. *et al.* (2014) ‘A review of the use and potential of the GATE Monte Carlo simulation code for radiation therapy and dosimetry applications’, *Medical Physics*, 41(6), p. 64301. doi: 10.1118/1.4871617.

Schroeder, W. J., Martin, K. M. and Lorensen, W. E. (1996) ‘The design and implementation of an object-oriented toolkit for 3D graphics and visualization’, *Proceedings of Seventh Annual IEEE Visualization '96*, 1, pp. 93–100. doi: 10.1109/VISUAL.1996.567752.

Siddon, R. L. (1985) ‘Fast calculation of the exact radiological path for a three-dimensional CT array’, *Medical Physics*, 12(2), pp. 252–255. doi: 10.1118/1.595715.

Srinivas, N. and Deb, K. (1994) ‘Multiobjective optimization using nondominated sorting in genetic algorithms’, *Evolutionary Computation*, 2(3), pp. 221–248. doi: 10.1017/CBO9781107415324.004.

Weller, M. *et al.* (2005) ‘Radiotherapy plus concomitant and adjuvant temozolomide for glioblastoma’, *Cancer/Radiothérapie*, 9(3), pp. 196–197. doi: 10.1016/j.canrad.2005.05.001.

Zielhuis, S. W. *et al.* (2006) ‘Production of GMP-grade radioactive holmium loaded poly(L-lactic acid) microspheres for clinical application’, *International Journal of Pharmaceutics*, 311(1–2), pp. 69–74. doi: 10.1016/j.ijpharm.2005.12.034.

6 APPENDIX: CREATION OF THE TPS

6.1 Introduction

Throughout this research, the relevant building blocks have been put in place to be able to create a microbrachytherapy TPS. The TPS that was designed for this project was named “DosiToul”, and will be referred to as such throughout this appendix.

DosiToul was written in C++, consisting of 133 source files and over 11,000 lines of code. Three open-source libraries were used for the development of this program: Qt, VTK and Grassroots DICOM (GDCM) (Malaterre, 2008).

This appendix discusses the structure of DosiToul, the tools that were used to create it and how a user could use it to produce microbrachytherapy treatment plans.

The structure of this appendix is as follows:

- Section 6.2 discusses the structure of the principal objects that make up the TPS;
- Section 6.3 demonstrates how an end-product user might use DosiToul to create microbrachytherapy treatment plans, including screenshots demonstrating the usage of the application;
- Section 6.4 describes the programming tools that was used to create DosiToul as well as the 3rd party libraries that were used;
- Section 6.5 discusses the object-oriented classes that were created for DosiToul (this section is rather technical); and
- Section 6.5 concludes this appendix.

6.2 DosiToul structure

DosiToul operates with a top-down structure. The top-level object is the database. The database can contain multiple patients. Each patient can contain multiple studies and each study can contain multiple series. Each series contains a set of DICOM images. This structure is typical of most medical imaging applications.

However, as well as attributing DICOM images to each series, DosiToul can also attribute multiple structures and perform multiple optimisations. Structures are all 3D surfaces, such as the PTV, OARs, test objects (spheres, etc.) and the skull. For each series, optimisations can be created, edited and then saved in the database.

Each optimisation contains a kernel, a dose grid, a list of all possible injections (cylindrical injections are used in DosiToul, so the injections are calculated prior to the start of the optimisation process, as described in Section 5.5.2.2), a list of user-defined cost functions. If the optimisation process is finished, the optimisation also contains a set of Pareto-optimal treatment plans. The dose grid, the list of possible injections and the treatment plans are all created at the start of the optimisation process using user-defined parameters.

The treatment plans at the end of the treatment planning process will contain all the information required to be able to perform the given treatment – the injection coordinates, the resulting dosimetry for each structure, and several visualisations of the treatment.

An illustration of the object hierarchy of DosiToul is shown in Figure 6.1.

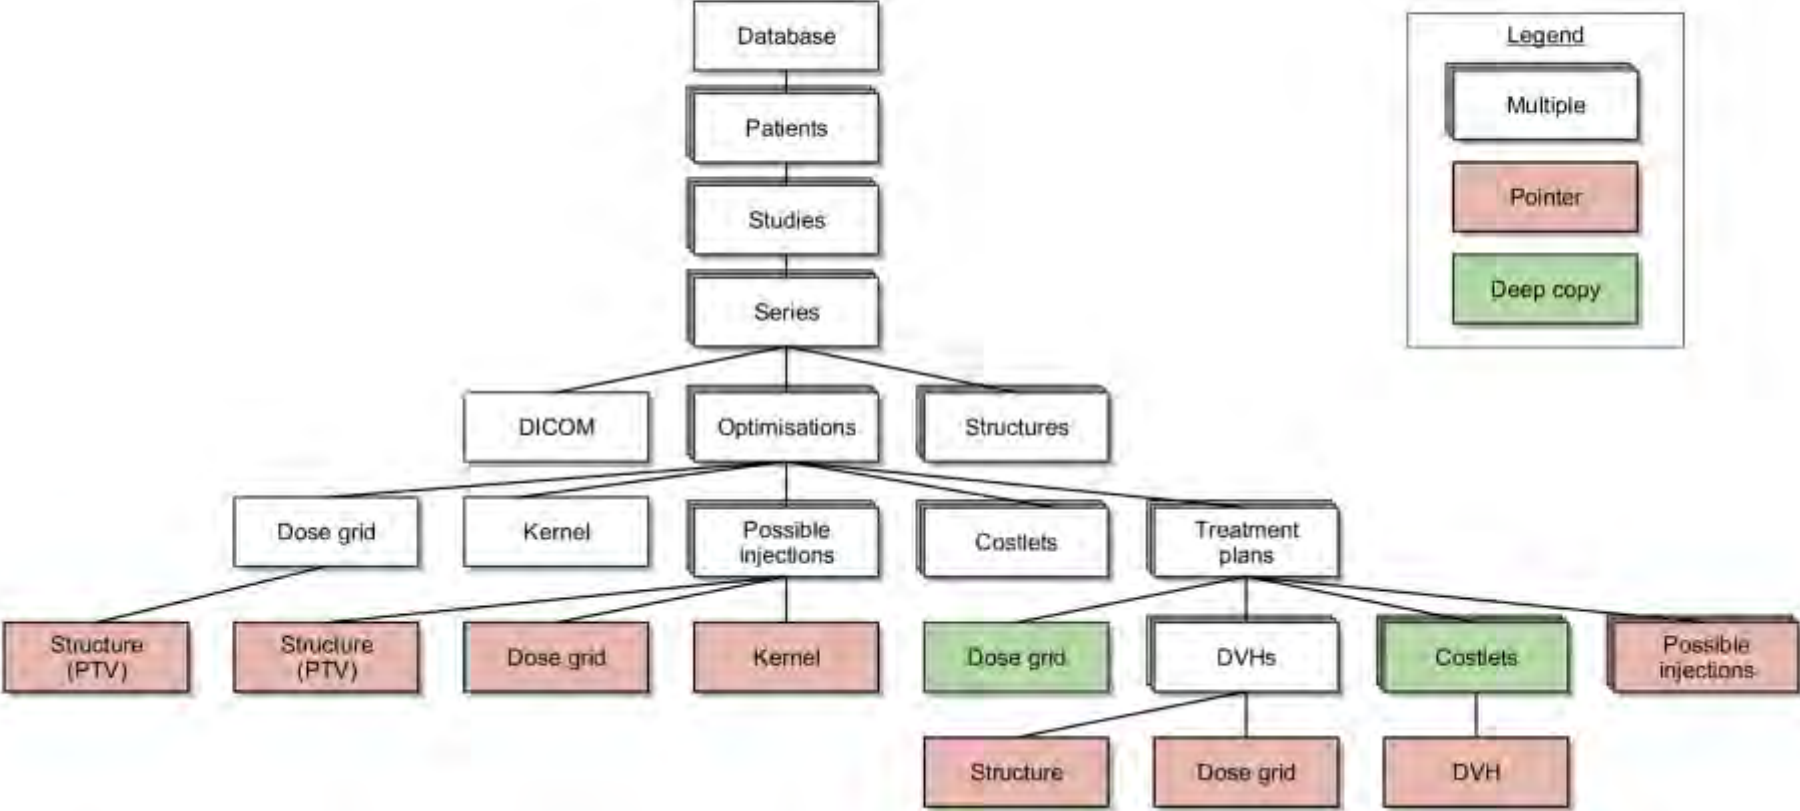


Figure 6.1: Illustration of the object hierarchy of DosiToul.

In Figure 6.1, it can be seen that certain objects are shared with each other as pointers. Pointers are used to minimise memory usage (by not copying objects), and so that changes to certain objects can be passed instantaneously throughout the program. For example, a pointer is used to pass the kernel throughout DosiToul. This is because the kernel does not change and so performing deep copies would be a waste of memory. An example of using pointers for immediate updates throughout the program are DVHs, which are shared with the cost functions. When the treatment plan is updated, the dose grid is updated accordingly and then so are the DVHs. The costlets need to have up-to-date versions to be able to instantaneously quantitate the quality of the updated treatment plan.

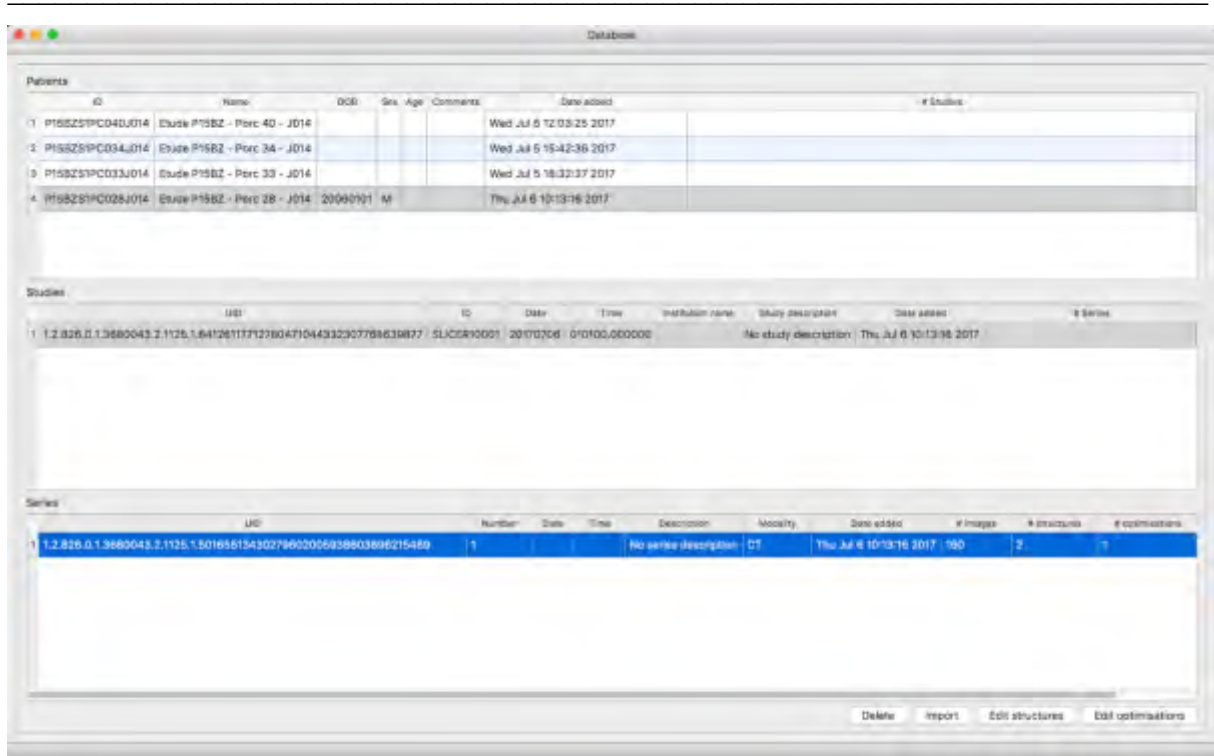
In other instances, deep copies of objects are required. For example, upon creating an optimisation, the user enters the parameters necessary to setup the dose grid (for example, the desired voxel size) and the dose grid is then created. Each treatment plan, however, needs to have its own dose grid to be able to produce dosimetry specific to that plan. For this reason, each treatment plan contains a deep copy of the dose grid.

6.3 Using DosiToul

6.3.1 Database

The database is the first page that confronts the user upon opening DosiToul. A screenshot of this database is shown below, in Figure 6.2.

Appendix: Creation of the TPS



The screenshot shows a software window titled "Database" with three main sections: Patients, Studies, and Series. Each section contains a table of data. The Patients table has columns for ID, Name, DOB, Sex, Age, Comments, Date added, and # Studies. The Studies table has columns for UID, ID, Date, Time, Institution name, Study description, Date added, and # Series. The Series table has columns for UID, Number, Date, Time, Description, Modality, Date added, # Images, # Structures, and # Optimizations. At the bottom of the window, there are buttons for "Delete", "Import", "Edit structures", and "Edit optimizations".

Patients							
ID	Name	DOB	Sex	Age	Comments	Date added	# Studies
1	P155251PC04J014	Etude P15BZ - Parc 40 - J014				Wed Jul 5 12:03:25 2017	
2	P155251PC034J014	Etude P15BZ - Parc 34 - J014				Wed Jul 5 15:42:36 2017	
3	P155251PC033J014	Etude P15BZ - Parc 33 - J014				Wed Jul 5 18:32:37 2017	
4	P155251PC025J014	Etude P15BZ - Parc 28 - J014	20090101	M		Thu Jul 6 10:13:16 2017	

Studies							
UID	ID	Date	Time	Institution name	Study description	Date added	# Series
1.2.826.1.1.3680043.2.1126.1.641261171278047104932307788638877	SLICR10001	20170706	010100.000000		No study description	Thu Jul 6 10:13:16 2017	

Series									
UID	Number	Date	Time	Description	Modality	Date added	# Images	# Structures	# Optimizations
1.2.826.1.1.3680043.2.1126.1.641261171278047104932307788638877	1			No series description	CT	Thu Jul 6 10:13:16 2017	150	2	1

Figure 6.2: Database window. At the top, a list of patients is shown. When a patient is selected, all studies for that patient are shown in the second list. When a study is selected, all series within that study are shown in the bottom box. Using the buttons along the bottom, patients, studies and series can be added; structures can be added to series; and optimisations can be performed.

The screenshot of the database reflects the hierarchy shown in Figure 6.1. The first table shows a list of all the patients in the database. Information is shown about the patients, such as their ID, date of birth and sex.

Once a patient has been selected, such as the fourth patient in the above example, all of the studies for that patient are shown. The studies list contains all relevant information about each study, and when one is selected, all of the series for the corresponding study are shown.

Patients, studies and series can be added to and removed from the database by using “Import” and “Delete”.

When the user clicks “Import”, a dialog box is displayed. The user selects the relevant folder and the database is updated. This is done by scanning the folder for all DICOM images, and using the class, “vtkDICOMDirectory”, to sort all images into the correct patient, study and series.

Once the user has added the desired directory to the database, the number of corresponding DICOM images should be shown under the series information. The example shown in Figure 6.2 contains 160 DICOM images.

To the right of the number of images, the number of structures and optimisations are shown. When a new series is added to the database, these are both zero. Before being able to create optimisations, the user must first add at least one structure. This is done by clicking “Edit Structures”.

6.3.2 Structure editor

When the user clicks on “Edit Structures”, the Structure Editor window opens, as shown in Figure 6.3.




Figure 6.3: Edit structures window. On the left, a 3D visualisation of all structures is shown. On the right, the corresponding table is displayed. Structures can be added, deleted and altered using the buttons along the bottom.

On the left of this window is a 3D visualisation of all the structures that have been added or created. On the right a table is displayed, giving information about each of the structures.

At the bottom of the page, controls enabling the user to manipulate the structures are shown. The first three buttons enable the user to add or create new structures. The user has five ways to add structures, as described in the following sections. The last three buttons enable the user to delete, change the name and change the colour of the selected structure.

6.3.2.1 DICOM RT Structure files

Using the GDCM class, “`vtkGDCMPolyDataReader`”, the user can import DICOM RT Structure files. This is the standard format for exporting structure segmentations in medical imaging. Since the user is capable of importing DICOM RT Structure files, they are free to perform image segmentation with the vast majority of image treatment stations that are commonly used in clinical environments, and then import these segmentations into DosiToul.

6.3.2.2 STL files

Standard triangle language (STL) is the standard format of computer-aided design (CAD) programs. It is also one of the exportation formats of segmented structures with 3D Slicer, an open-source image treatment program. 3D Slicer was used internally to perform segmentations and so this format was indispensable.

6.3.2.3 Spheres

Spheres have been used extensively throughout this research. The user may create spherical structures as desired by designating their radii and centre coordinates.

6.3.2.4 Shells

Shells can be created by selecting a pre-existing structure (such as the PTV) and setting the minimum and maximum thickness to create the shell about that structure.

Shells are useful for calculating the dosimetry around structures, especially around the PTV. Cost functions can be applied to these shells, forcing the optimiser to find treatment plans with little irradiation of tissue outside of the PTV.

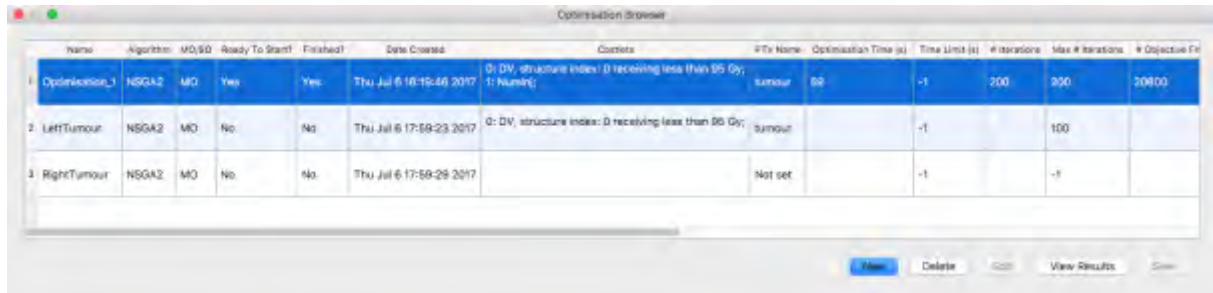
6.3.2.5 DICOM Hounsfield units

Lastly, structures can be created from the densities of the voxels in the DICOM CT scan, expressed in Hounsfield units (HU). Bone, for example, typically has HU between 700 and 1,300. Once the user has chosen the limits, the algorithm, “`vtkDiscreteMarchingCubes`”, is used to build the structure.

In the preclinical mini-pig tests, the skull was generated using this method, allowing the user to observe the hole cut through the mini-pig’s skull that would be used as the point of entry for the needle.

6.3.3 Optimisation browser

Once the user has added all of the desired structures to the series in the database window, they may now start to produce treatment plans by clicking “Edit Optimisations”. Clicking this button opens the Optimisation Browser window, as shown in Figure 6.4.



Name	Algorithm	MO/GO	Ready To Start	Finished	Date Created	Contexts	P/E Name	Optimization Time (s)	Time Limit (s)	# Iterations	Max # Iterations	# Objective F
Optimization_1	NSGA2	MO	Yes	Yes	Thu Jul 6 16:15:46 2017	0: DV, structure index: 0 receiving less than 95 Gy; 1: Normal	tumour	59	-1	200	300	20000
LeftTumour	NSGA2	MO	No	No	Thu Jul 6 17:59:23 2017	0: DV, structure index: 0 receiving less than 80 Gy;	tumour		-1		100	
RightTumour	NSGA2	MO	No	No	Thu Jul 6 17:59:26 2017		Not set		-1		-1	

Figure 6.4: Optimisation browser. Optimisations can be created and deleted using “New” and “Delete”. They can be parameterised using “Edit”. The optimisation can be executed using “Start” and then the results can be viewed with “View Results”.

The browser lists all optimisations that have taken place or are being set up (the fifth column states whether they are finished or not). With “New”, the user can create a new optimisation, give it a name, and then with “Edit”, they can parameterise the optimisation.

Once all the necessary information has been entered about the optimisation, “Start” becomes available. The user starts the optimisation process, and a new window opens, visualising the progress of the treatment planning.

When the optimisation is finished, the progress window can be closed and, in the optimisation browser, “View Results” becomes active. The user clicks this option, enabling them to choose from the Pareto-optimal solutions and view the corresponding dosimetry for that plan.

6.3.4 Setting up a new optimisation

Upon creating a new optimisation and clicking on “Edit”, a wizard opens with four pages. The user enters the necessary information page-by-page, and the treatment planning process is ready to start.

If at any point, the user wishes to save their progress and finish setting up the optimisation at a later point, the wizard can be closed with “Save and Close”.

6.3.4.1 Algorithm parameters

The first page of the wizard requires parameters related to the algorithm, as shown in the screenshot in Figure 6.5.

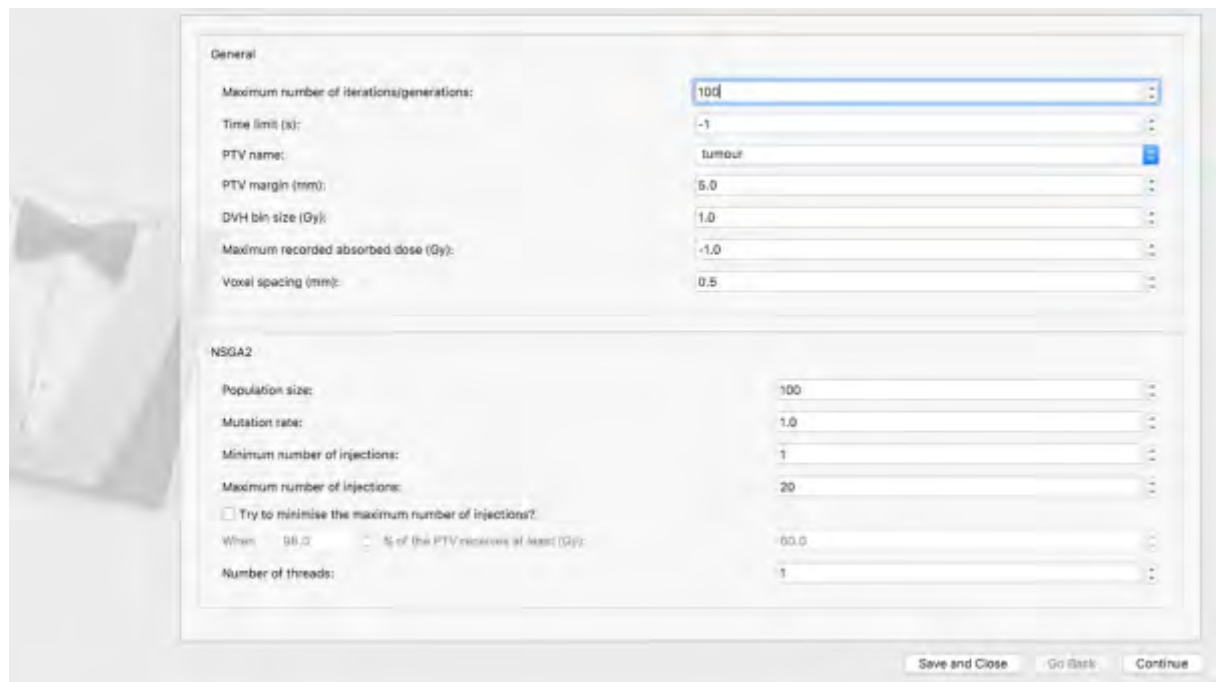


Figure 6.5: Algorithm settings. The top box requires information that is common to all algorithms (for example, the time limit), whereas the bottom half requires information that is specific to a given algorithm (for example, the population size for NSGA-II).

At the top, parameters that are common to all algorithms are shown. Setting values to -1 indicates that the value is unlimited. The required parameters are:

- The maximum number of iterations,
- The time limit,
- The structure that will be treated as the PTV for the optimisation (this is a drop-down list containing all structures that were present in the Structure Editor),
- The margin of the dose grid that should be added around the PTV,
- The DVH bin size,
- The maximum recorded absorbed dose, and
- The desired voxel spacing.

The bottom half of the page shows the parameters that are specific to a given algorithm. In the example shown in Figure 6.5, the parameters are shown for NSGA-II. The displayed parameters are:

- The population size,
- The mutation rate,
- The maximum and minimum number of injections,
- A parameter that reduces the maximum number of injections, and
- The desired number of threads that should be used (for multi-core CPUs).

Once acceptable values have been selected on this page (for example, leaving both the maximum number of iterations and the optimisation time limit as -1 is not permitted as this would mean that optimisations would continue indefinitely), “Continue” becomes available and the user can continue to the second page.

6.3.4.2 Kernel parameters

The second wizard page contains all the parameters required to create the injection kernel. Shown below, in Figure 6.6, are the parameters required to create a kernel for cylindrical injections.

Path:	/Users/CRCT_Richard/Documents/OneDrive/PhD/Code/GATE/TheraneM_cylinder/Kernel_Cylinder_r487_i20/output/kernel-Dose.vti	<input type="button" value="Browse"/>
Half-life (s):	96696	<input type="text"/>
Microparticle concentration (mg/ml):	800.0	<input type="text"/>
Fractional mass of Holmium in microparticle (unitless):	1.00	<input type="text"/>
Initial activity per mass (MBq/mg):	3.0	<input type="text"/>
Modelled injection radius (mm):	0.47	<input type="text"/>
Modelled injection length (mm):	20.00	<input type="text"/>

Figure 6.6: Kernel settings. Information, such as the path of the GATE simulation and the radioelement’s half-life, is required here.

The user is required to give the path to the GATE simulation that contains the voxelised S-values (as described in Section 3.2 and 5.5.2.1), the radioelement's half-life, the concentration of microspheres, the fractional mass of the radioelement within the microsphere (as discussed in Section 3.5.2, the microspheres often contain glass, polymers or resin casings), and the radius and length of the modelled cylinder (which as discussed in Section 5.5.2 were 0.47 and 20 mm, respectively).

Once all parameters have been correctly entered, “Continue” becomes active. For example, if the GATE file does not exist, “Continue” stays greyed out. When the user clicks on “Continue”, they advance to the third wizard page, the Costlet page.

6.3.4.3 Costlet parameters

Upon entering all the required parameters for the kernel, the user must enter all the desired costlets. The page displays a table that is initially blank. As costlets are created, they are added to the table. Costlets are added and removed using “Add” and “Delete” at the bottom of the page, as shown in Figure 6.7.

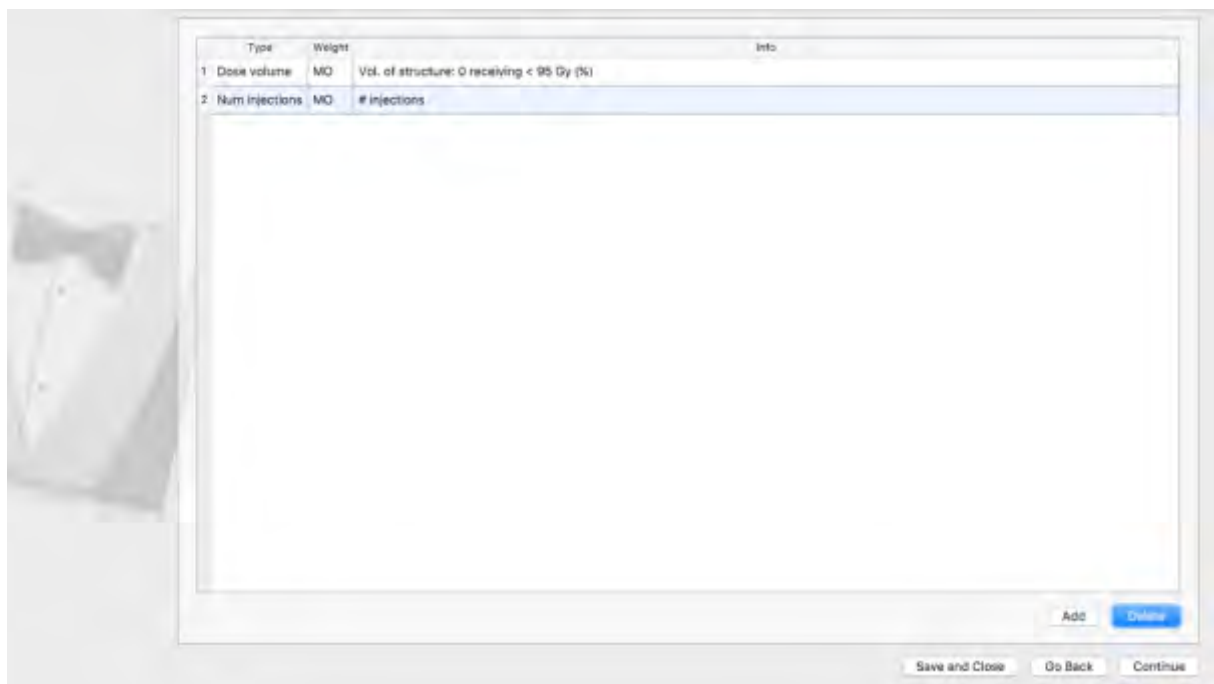


Figure 6.7: Costlet settings. A list of all costlets is shown and new costlets can be added with “Add”. Conversely, they can be removed with “Delete”.

When the user clicks “Add”, a small dialog box is displayed, the user chooses the desired costlet type and the relevant parameters for that choice of costlet. The example shown in Figure 6.8 is a dose-volume costlet that penalises the percentage of the tumour receiving less than 95 Gy.




Figure 6.8: Setting up individual costlets. Here, the parameters for a dose-volume costlet are shown. The user chooses the structure to be penalised, the inequality direction (“<”, “≤”, “>” or “≥”) and the reference absorbed dose. Upon clicking “OK”, the costlet is added to the list.

Once the user has added the desired costlets (at least 1 is required), the page should look similar to Figure 6.7, and the user can advance to the final page.

6.3.4.4 Injection parameters

The final wizard page requires the user to enter all the information pertaining to the injections. Figure 6.9 shows an example of the required information for cylindrical injections.



Figure 6.9: Injection parameters. On the left, all possible injections are shown relative to all the structures (in this case the skull and the PTV). On the right, all parameters required for the injections are entered. Once all parameters have been entered, all possible injections are computed using “Calculate”.

As can be seen, a lot of information is required on this page. The first seven parameters are described in Section 5.5.1. The negligible absorbed dose box is used to minimise the number of voxels in the kernel, as described in Section 3.5.3.1. The focal point, the X tilt and the Z tilt are used to describe the trajectory of the guide needle, shown in green.

The direction of the 0° rotation is used to give the surgeon a point of orientation. For example, 0° could be set towards the top of the operating table. Following this, the user can decide whether the injections will be clockwise or anti-clockwise.

Lastly, the user dictates the distance between the focal point (normally in the centre of the PTV) and the point that the program considers to be the zero-distance. This could, for example, be set at the top of the skull. The surgeon would place the guide so that it was at the level of the top of the skull, and then all depths of injections would be relative to this. Hence, a depth of 7 mm corresponds to 7 mm under the top of the skull.

Once all parameters have been entered, the user clicks “Calculate”. The program will then calculate all possible injections corresponding to those parameters and show them as the blue lines that can be seen in Figure 6.9.

If for example, on the first wizard page, the user decided that the maximum number of injections should be limited at 100, but only 20 possible injections were calculated on the last wizard page, an error will be displayed. This lets the user know that they should either alter the injection parameters (change the angle of approach, perform shorter injections, etc.) or decrease the maximum number of injections.

When the number of possible injections is greater than the maximum number of injections, the setting up of the optimisation is complete. The user may now click “Done”. The wizard closes and the user returns to the Optimisation Browser.

6.3.5 Performing the optimisation

The Optimisation Browser enables the user to see a list of all optimisations. If an optimisation has been fully set-up but not carried out, then the “Ready to Start?” and “Finished?” columns of Figure 6.4 will display the text, “Yes” and “No”, respectively. The user may select this optimisation and only then will “Start” become active. The user clicks “Start”, and the optimisation process begins.

A new window opens, showing the progress of the optimisation. An example of the progress of NSGA-II is shown below, in Figure 6.10.




Figure 6.10: Optimisation progress. Here, a real-time visualisation of the Pareto front is shown. The user may cancel the optimisation with “Stop” or click “Done” once the optimisation is complete.

The Pareto front is shown in green and all non-Pareto-optimal solutions are shown in pink. Once the treatment planning is complete, “Done” becomes active, the user clicks this and the window closes.

6.3.6 Viewing optimisation results

Once the treatment planning is complete, the text in the “Finished?” column of the Optimisation Browser becomes “Yes”. At this point, “Start” is greyed out and “View Results” becomes active. The user clicks this button and the page shown in Figure 6.11 is displayed.

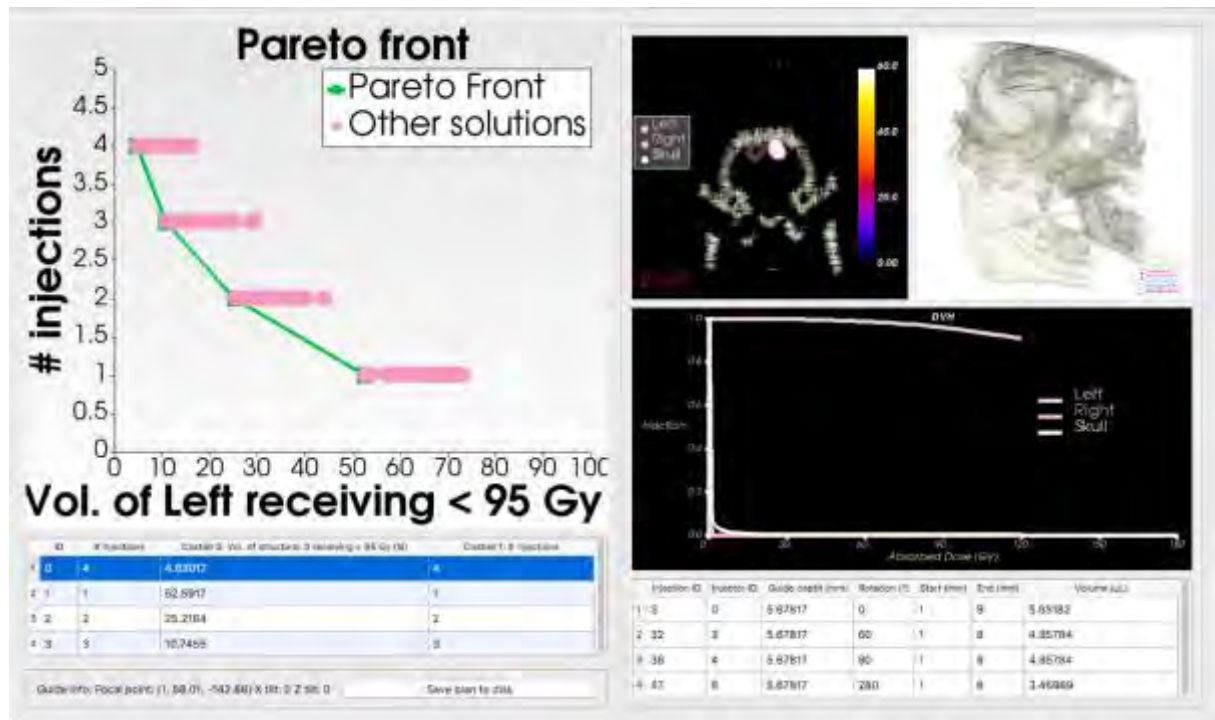


Figure 6.11: Optimisation results page. At the top-left, the final Pareto front is shown, at the bottom-left the Pareto-optimal solutions are shown in the table. Once a solution has been selected, the right-hand side updates automatically, showing slices of the absorbed dose distribution (top-left), the 3D visualisation of the injections (top-right), the DVH (middle) and the coordinates of the injections (bottom). The selected treatment plan can be saved using “Save plan to disk”.

At the top-left of the results page, the Pareto front can be seen. Below this is a table containing all the Pareto-optimal solutions. The user may choose one of these plans and the right side of the window changes accordingly. In the above example, there are only four Pareto-optimal solutions (four green points in the graph), and so the table only contains four entries.

The right side of the window contains 2D slices of the distribution of absorbed dose (top-left), a 3D visualisation of the injections relative to the skull and the PTV (top-right), the DVH for all structures (middle) and the coordinates for the injections (bottom).

Since the user has selected the treatment plan that uses four injections, the coordinates of these four injections are shown in the table. The optimisation was set up for cylindrical injections, and so the coordinates are given in cylindrical coordinate: the depth of the guide needle (relative to the zero distance that was given in Figure 6.9), the rotation of the injector, the start and end depths of the injection and the volume that will be injected over this distance.

Once the user has selected the treatment plan that they wish to execute, “Save plan to disk” can be used to export the given treatment plan to disk such that the treatment can be printed out and the visualisations viewed in programs such as 3D Slicer and Paraview.

When the user is finished with DosiToul and they close the application, the entirety of the database – the patients, studies, series, structures, optimisations and treatment plans – are all saved to disk so that upon loading the program the next time, they can continue from where they finished.

6.4 C++

As has been previously mentioned, DosiToul was created using object-oriented C++. In Section 6.5, descriptions of the classes used in DosiToul are given. To be able to understand these descriptions, it is necessary to first discuss some the keywords that are used in object-oriented programming:

- Class: grouping together all characteristics that are common to a type of object.
 - For example, all triangles have height, width, etc.
- Object: an instantiation of a class.
 - For example, a triangle named “A” might have a height of 6 cm and a width of 2 cm.
- Data member: values that belong to a class.
 - For example, the triangle class will contain members for the width and height of the triangle.
- Member function: a function that belongs to a class.
 - For example, a triangle class might have a member function to compute the area of the triangle and return the value to the user.
- Inheritance: the idea that classes can be grouped together.
 - For example, classes for equilateral and isosceles triangles might be created, both inheriting from a base class, describing triangles in general.
- Abstract classes: classes that contain some information about a type of object, but not enough to create an object.
 - For example, a general class might be created for shapes. This class would contain the variable, area, but the area cannot be calculated until the type of shape is known. A derived class, such as a triangle, must therefore be created.

- Polymorphism: the ability of an object to take multiple forms.
 - For example, using polymorphism, a vector of shapes containing both triangles and squares can be stored together.
- Pointers: share the address of an object without having to copy it.
- Shared pointers: the same idea as a pointer, except it keeps a count of the number of times it is referenced. When this counter becomes zero, the object calls its destructor, simplifying memory management.
- Multi-threading: splitting the code so that it runs asynchronously on multiple CPUs (executed using `std::thread`).

As well as C++ features, three open-source C++ libraries were also used: Qt, VTK and GDCM. The use of these three libraries is described in the following sections.

6.4.1 Qt

Qt was used to create the user interface. Within the main window (“QWindow”), widgets (“QWidgets”), tables (“QTable”), buttons (“QPushButton”), sliders (“QSlider”), file dialog boxes (“QFileDialog”) and a host of other classes were all used to allow the user to navigate throughout the TPS and carry out optimisations of treatments.

Signals and slots are a method, used by Qt, to send information throughout the code. When a button is pushed, for example, it will send out a signal. For each signal that is sent out, there can be any number of objects that are listening for that signal (via slots). For example, when the user clicks “Save plan to disk” a dialog box must open to allow the user to select the desired save location. The slot of the dialog box is therefore connected to the signal of the save button such that when the save button is clicked, it sends out a signal and the relevant slot is triggered and a resulting piece of code is executed.

6.4.2 VTK

VTK was used extensively throughout this project. 3D arrays, such as the kernel and the patient dose grid were stored using the class “`vtkImageData`”. Structures were stored as “`vtkPolyData`” (a series of points describing the surface of the structure). Patient-specific dosimetry was performed by creating binary masks of the structures with “`vtkStencilImage`” and superposing them on the dose grid with “`vtkImageAccumulate`”. Patient DICOM files were read with “`vtkDICOMReader`”. Resampling of 3D arrays was done with “`vtkImageReslice`”.

VTK was also used for the visualisation of data in DosiToul. The visualisations that were created include:

- 3D visualisations (such as the structures in the Structure Editor),
- 2D slices of 3D arrays (such as the patient dose grid in the Results Viewer), and
- 2D, 3D and ND graphs (such as the real-time Pareto front whilst the algorithm is running).

Rendering and interaction of these visualisations were performed using “vtkActor”, “vtkCamera”, “vtkRenderer”, “vtkRenderWindow” and “vtkRenderWindowInteractor”. Graphs were created with “vtkChartXY” and “vtkChartXYZ”.

These VTK visualisations were linked together with the Qt window by embedding “QVTKWidgets” (an inheritance of “QWidget”) in the main “QWindow”. An example of this is shown in Figure 6.12, which is an annotated version of the Structures Editor.

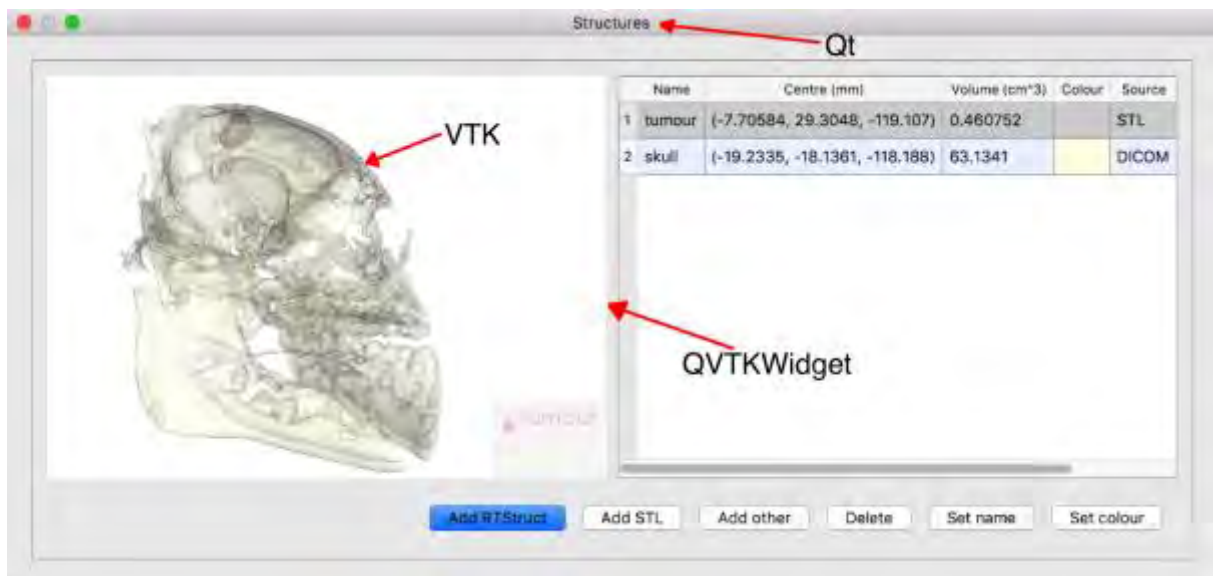


Figure 6.12: Example of a VTK widget embedded in a Qt window via a “QVTKWidget”.

6.4.3 GDCM

Only one class was used from GDCM, but this class was vitally important. The class, “vtkGDCMPolyDataReader”, was used to read DICOM RT Structure files, which were then converted to vtkPolyData to be used as structures.

6.5 DosiToul classes

The most important classes that were used in DosiToul are described below, including their data members and some of the most important member functions.

6.5.1 DT2_Database, DT2_Patient and DT2_Study

The database (“DT2_Database”) can be created either from scratch or by giving the path to a previously created database on the hard disk from which it can be loaded. New entries can be added to the database by providing the path to DICOM images on the hard drive. The DICOM images are sorted into patients (“DT2_Patient”), studies (“DT2_Study”) and series (“DT2_Series”), and all new entries are added to the database.

6.5.2 DT2_Series

The series contains the DICOM images (“DT2_DICOM”), a list of structures (“DT2_Structure”) and a list of optimisations (“DT2_Optimisation”).

6.5.3 DT2_DICOM

The DICOM (“DT2_DICOM”) class contains a list of filenames corresponding to each of the DICOM images in that series. When the “update” method is called, the DICOM images are read and stored as “vtkImageData” (a 3D array). The patient array (“vtkMatrix4x4”) can be extracted to be able to superpose structures on the DICOM for 3D visualisation.

A member function was also implemented to retrieve information stored in the DICOM’s header (via its tag) or to print the entirety of the header.

6.5.4 DT2_Optimisation

The optimisation class (“DT2_Optimisation”) is an abstract base class from which derived classes inherit, such as “DT2_Optimisation_NSGA2”.

Some of the members of “DT2_Optimisation” include:

- The maximum number of treatment plan updates allowed,
- The number of treatment plan updates already performed,
- The maximum time limit,
- The current time spent on the optimisation,
- A list of costlets (“DT2_Costlet”),
- The dose grid (“DT2_DoseGrid”),
- The kernel (“DT2_Kernel”), and
- A list of all possible injections (“DT2_Injection”).

6.5.4.1 *DT2_Optimisation_NSQA2*

Since “DT2_Optimisation_NSQA2” is a derived class, it contains all of the information in “DT2_Optimisation”, plus:

- The maximum and minimum number of injections,
- The mutation rate,
- The population size, and
- The population (a list of “DT2_TreatmentPlan_MO”).

6.5.5 *DT2_Array3D*

An abstract base class, “DT2_Array3D” was used to store voxelised 3D arrays in the form of “vtkImageData”.

The class contained member functions permitting the user to:

- Access array elements (either from their 1D or 3D index),
- Get the maximum or minimum values in the array,
- Get the voxel spacing,
- Get the arrays origin,
- Get the dimensions (the number of voxels in each dimension),
- Get the bounds (the limits of the array in 3D space),
- Get the extent, dimensions, voxel spacing, bounds, and
- Get the amount of memory that the array occupies (the number of voxels multiplied by the memory of the data type, such as 8B for a double).

6.5.5.1 *DT2_Kernel*

The kernel (“DT2_Kernel”) contains the absorbed dose distribution for an injection. The user sets the relevant information, such as the path to the GATE simulation that is saved on the computer, and the kernel is created.

6.5.5.2 *DT2_DoseGrid*

The dose grid (“DT2_DoseGrid”) is a derived class of “DT2_Array3D”. As well as containing the voxelised 3D array, it also contains:

- A pointer to the kernel,
- A pointer to the PTV’s structure,
- A margin to be added to the PTV, and
- The desired voxel spacing.

The dose grid is initialised by setting the PTV's structure and the margin that the user wishes to add around the PTV, and the voxelised grid is created with the desired voxel spacing.

A member function was implemented to allow the user to insert injections on the dose grid by convolving the kernel in the desired location.

6.5.6 DT2_TreatmentPlan

Each treatment plan ("DT2_TreatmentPlan") contains:

- Its own dose grid ("DT2_DoseGrid"),
- Its own list of costlets ("DT2_Costlet"),
- A list of DVHs ("DT2_DVH"), and
- A list of injections ("DT2_Injections").

The first two are created by performing deep copies of the versions that were created during the setup phase, and belong to "DT2_Optimisation". It is important to remember that deep copies were required so that each treatment plan could be independent and have its own corresponding dosimetric information.

6.5.6.1 DT2_TreatmentPlan_MO

The MO treatment plan ("DT2_TreatmentPlan_MO") is a class that is derived from "DT2_TreatmentPlan". Therefore, it contains all the members that are in "DT2_Treatment plan", plus MO information:

- The plan's Pareto front,
- The plan's crowding distance,
- The number of plans that this plan dominates, and
- The number of plans that dominate it.

6.5.7 DT2_Structure

The structure class ("DT2_Structure") is an abstract base class from which all other structure classes are derived. The base class contains:

- The points that constitute the structure's surface ("vtkPolyData");
- The actor that can be rendered ("vtkActor"); and
- The structure's name, colour and volume.

The derived classes (“DT2_Structure_DICOM”, “DT2_Structure_RTStruct”, “DT2_Structure_Shell”, “DT2_Structure_Sphere” and “DT2_Structure_STL”) differ only in their method of creating or reading the structures.

Each of the derived classes will contain information that is specific to their type. For example, the sphere class will contain the coordinates to the sphere’s centre as well as its radius. The DICOM RT Structure file will contain all the information that is stored in the DICOM’s header.

6.5.8 DT2_DVH

The structures’ DVHs (“DT2_DVH”) contain:

- A pointer to the treatment plan’s dose grid,
- A pointer to a structure,
- A binary mask that was created by superposing the structure on the dose grid, and
- A histogram (“vtkImageAccumulate”).

A method exists in “vtkImageAccumulate” that is used to set the bin size of the DVH.

6.5.9 DT2_Costlet

Costlets were created using the abstract base class, “DT2_Costlet”. Derived from this were the specific costlets, such as “DT2_Costlet_NumInjections”, “DT2_DoseVolume” and “DT2_ConformityIndex”. Each of the individual costlets contained different information (for example the costlet for the number of injections contains a pointer to the treatment plan’s injections and the dose-volume costlet contained a structure’s DVH).

This set of classes is an example of polymorphism that was briefly explained in Section 6.4; by creating a list of pointers to costlets, the list could contain a mixture of different types of costlets.

6.5.10 DT2_Line, DT2_Matrix3x3, DT2_Point and DT2_Vec3D

Various geometrical classes were created, such as lines (“DT2_Line”), 3x3 matrices (“DT2_Matrix3x3”), points (“DT2_Point”) and 3D vectors (“DT2_Vec3D”). These classes were used extensively when calculating all possible injections. The guide needle, for example, was represented using “DT2_Line”, the rotation matrix between the DICOM reference frame and the injection reference frame was stored using “DT2_Matrix3x3”.

6.5.11 DT2_ Injection

The cylindrical injections were more complicated than the spherical injections. Spherical injections needed just the coordinates to the centre of the injection.

Cylindrical injections were more complicated due to the double needle (the “guide” and “injector”) approach that was described in Section 5.5.1. Therefore, separate classes were created for the trajectory of the guide needle (“DT2_TrajectoryGuide”), the trajectory of the injector needle (“DT2_TrajectoryInjector”) and the resulting cylindrical injection (“DT2_InjectionCylindrical”).

6.5.11.1 DT2_TrajectoryGuide

The user gave the focal point of the guide needle’s trajectory, as well as the X- and Z-tilt relative to the DICOM vertical.

From here, the frame of reference was changed such that the guide travelled exclusively in the y-direction. This was done using Rodrigues’ rotation formula and the result was stored as a “DT2_Matrix3x3”.

Subsequently, the points at which the guide needle entered and left the PTV were the calculated by intersecting the guide line with the PTV using the class, “vtkOBBTree”.

The position of each injector trajectory was calculated by iterating over each depth of the guide needle (for example, 0 mm, 1 mm, 2 mm, and so on), and then iterating over the rotation of the injector needle (for example, 0°, 20°, 40°, and so on). For each injector needle trajectory, a corresponding “DT2_TrajectoryInjector” is created.

6.5.11.2 DT2_TrajectoryInjector

The trajectory injector starts at the position that the injector needle comes out from the guide needle and stops when the injector needle leaves the PTV. The point that the injector needle leaves the PTV was also calculated by intersecting the two using “vtkOBBTree”.

Once the start and end points of the injector have been calculated, all possible injections along this line are calculated (for example, an injection between 1 and 8 mm, then 1 and 9 mm, then 1 and 10 mm, then 2 and 9 mm, and so on). For each possible injection, a “DT2_Injection” is created.

6.5.11.3 DT2_Injection

The injection class (“DT2_Injection”) contains the start and end depths of the injection. By supplying “DT2_Injection” with a pointer to the cylindrical injection kernel, the kernel can be translated, rotated and resampled into the correct position so that the injection can be quickly placed on the dose grid.

6.5.12 DT2_Visualisation

A base abstract class, “DT2_Visualisation”, was used for all visualisations. From this, four derived visualisations were created:

- “DT2_Visualisation_Block3D”,
- “DT2_Visualisation_DVH”,
- “DT2_Visualisation_ParetoFront”, and
- “DT2_Visualisation_Structs3D”,

6.5.12.1 DT2_Visualisation_Block3D

All 3D arrays were visualised using “DT2_Visualisation_Block3D”, which displayed 2D slices of the array. The “vtkImageData” was supplied, as well as the maximum value (if the user wished to clip the data to, say, 100 Gy), the desired look up table (for example, greyscale or fire) and pointers to structures that the user wishes to superpose on the array (as can be seen in the top-centre of the results page shown in Figure 6.11).

The user could then navigate through the slices by using the mouse wheel, zoom by dragging with the right mouse button, change the level and contrast by dragging with the left mouse button, move the scene by dragging with the left mouse button whilst holding down the “Shift” key, or rotate by dragging with the left mouse button whilst holding down the “Cmd” key (on Mac).

The orientation of the visualisation (sagittal, axial or coronal) could be changed using the class, “vtkImageResliceMapper”.

6.5.12.2 DT2_Visualisation_DVH

The DVHs were visualised using “vtkChartXY” by supplying pointers to the DVHs and, if desired, a maximum absorbed dose (200 Gy was often used throughout this work).

6.5.12.3 DT2_Visualisation_ParetoFront

The class, “DT2_Visualisation_ParetoFront”, is actually an abstract base class with three derived classes:

- DT2_Visualisation_ParetoFront_2D,
- DT2_Visualisation_ParetoFront_3D, and
- DT2_Visualisation_ParetoFront_ND.

The 2D version was created using “vtkChartXY”, the 3D version was created using “vtkChartXYZ” and the ND version was created by displaying many 2D projections (“vtkChartXY”) of the ND Pareto front, as shown in Figure 4.15.

To create a Pareto front visualisation, the user need only supply a pointer to the list of treatment plans. The plans will be displayed as a function of their Pareto front and plotted in the relevant position according to their costlet functions.

6.5.12.4 DT2_Visualisation_Structs3D

The last visualisation, “DT2_Visualisation_Structs3D”, is used to display 3D structures. For this class, three types of 3D object can be visualised all together:

- The structures (the PTV and OARs);
- The injections, injectors trajectory, guide trajectory; and
- The bounding box of the dose grid.

An example of all three in the same visualisation can be seen in the top-right of the results page, shown in Figure 6.11.

6.6 Conclusion

Using the information that has been accrued over the entirety of this research, a TPS has been created that could be used by an end-product user. The TPS contains a database and can be used to create and store multiple structures for a given DICOM series.

Optimisations can then be created, edited, executed and reviewed. The creation of the optimisations was designed to be as simple as possible, and can be completed with just four wizard pages.

Once the optimisation has been completed, the results can be reviewed. From here, a treatment plan can be selected, and after validation, delivered to the patient. To aid the selection, much information is given on this page, including visualisations of the Pareto front and a table of all Pareto-optimal solutions. Once a plan has been selected, the treatment plan is displayed via 2D slices of the dose grid, 3D visualisation of the injections relative to the structures, DVHs of all structures and the coordinates of all injections.

The entirety of the database can be saved to disk so that it can be reused in the future. Objects that are saved include the patients, studies, series, DICOM images, structures, optimisations and treatment plans.

DosiToul is a large project, spanning 133 source files and over 11,000 lines of code. The code was implemented in object-oriented C++, making use of the open-source libraries, Qt, VTK and GDCM.

SCIENTIFIC PRODUCTIONS

Journal papers

R. Brown, M. Bardiès and X. Franceries “Treatment planning of microbrachytherapy with 3D” NSGA-II EMBEC & NBC 2017, IFMBE Proceedings, Vol. 65, pp. 450-453, 2017, Online: http://dx.doi.org/10.1007/978-981-10-5122-7_113

International conferences (posters)

R. Brown, M. Bardiès, X. Franceries “Treatment planning of microbrachytherapy with 3D” NSGA-II EMBEC & NBC 2017, Jun. 2017, Tampere, Finland

R. Brown, M. Bardiès, X. Franceries “Microbrachytherapy: even more localised dose profiles?” ESTRO 36, May 2017, Vienna

R. Brown, A. Petitguillaume, M. Bardiès, X. Franceries “Microbrachytherapy treatment planning via inverse problem-solving algorithms” 18th International Conference on the use of Computers in Radiation Therapy (ICCR), Jun. 2016, London

MICROBRACHYTHERAPY TREATMENT PLANNING

RICHARD BROWN

An innovative form of radiotherapy, microbrachytherapy, is under development. This therapy targets solid, inoperable tumours by performing injections of liquid containing radioactive microspheres in suspension.

Many injections are required to sufficiently cover the tumoural volume, and so to be able to deliver the position of these injections, a method of treatment planning has been developed and validated throughout this research.

Throughout this work, three main questions are addressed:

- How to perform the dosimetry for microbrachytherapy?
- How to perform treatment planning for this modality?
- What are the optimal injection properties to deliver the most efficient treatment?

Microbrachytherapy dosimetry was performed by calculating the absorbed dose distribution for an injection. This distribution was then convolved at each injection position within the tumour to calculate the patient's absorbed dose distribution. Dosimetry of the tumour and the organs at risk was performed by extracting and analysing dose-volume histograms (DVHs).

Once a method of dosimetry was put in place, optimisation algorithms were developed to generate patient-specific treatment plans. For this, three algorithms were tested and compared: Nelder-Mead Simplex, the Bees algorithm and the non-dominated sorting genetic algorithm-II. It was found that, thanks to its MO optimisation, the non-dominated sorting algorithm-II was the most flexible, and was used preferentially.

Lastly, a comparison of injection parameters was performed. It was found that between ^{90}Y , ^{166}Ho , ^{131}I and ^{177}Lu , optimal injections consisted of microspheres of ^{90}Y . Injection volumes of 5, 10 and 20 μL and initial activities of 5, 10 and 20 MBq were tested. It was found that 20 μL injections with 20 MBq were optimal because they minimise the number of injections required.

This new technology combined with developments shown in this work demonstrate the feasibility – that was validated on animals – the ability to inject liquid containing radioactive microspheres in suspension to efficiently treat inoperable tumours whilst protecting surrounding healthy tissue. Such tumours, despite still having a poor prognosis, will surely have better support in the near future.

PLANIFICATION DE TRAITEMENT EN MICRO-CURIETHERAPIE

RICHARD BROWN

Une nouvelle modalité de radiothérapie, la micro-curiéthérapie, est en cours de développement. Cette thérapie cible des tumeurs solides inopérables en effectuant des injections de liquide contenant des microsphères radioactives en suspension.

Plusieurs injections sont nécessaires pour suffisamment irradier la zone tumorale et donc, afin d'optimiser le positionnement de ces injections, une méthode de planification de traitement nécessaire a été développée et validée au cours de cette thèse.

Tout au long de ce travail, trois thèmes principaux seront discutés :

- Comment réaliser la dosimétrie particulière de cette micro-curiéthérapie ?
- Comment effectuer la planification de traitement pour cette modalité ?
- Comment optimiser le plan de traitement afin qu'il soit le plus efficace possible ?

La dosimétrie en micro-curiéthérapie a été réalisée en calculant la distribution de dose absorbée pour une injection. Cette distribution a été convoluée à la position des autres injections dans la tumeur pour calculer la distribution de dose absorbée dans le patient. Pour effectuer la dosimétrie spécifique dans la tumeur et les organes à risque, les histogrammes dose-volume (HDV) ont été extraits et analysés.

Une fois la méthode de dosimétrie établie, nous avons développé une méthodologie de planification de traitement pour développer et optimiser le plan pour chaque patient. Pour cela, nous avons testé et comparé trois algorithmes : la méthode de Nelder-Mead, l'algorithme des abeilles et l'algorithme « Non-Dominated Sorting Genetic Algorithm-II » (NSGA-II). Nous avons montré que, grâce à l'optimisation multiobjectif, le NSGA-II donne plus de liberté à l'utilisateur ; c'est pourquoi il a été utilisé par la suite.

Enfin, nous avons effectué une comparaison entre les paramètres d'injection. Nous avons mis en évidence qu'entre les radio-isotopes ^{90}Y , ^{166}Ho , ^{131}I and ^{177}Lu , les injections de ^{90}Y sont optimales. Nous avons testé des injections de 5, 10 et 20 μL et des activités initiales de 5, 10 et 20 MBq. Nous avons trouvé que des injections de 20 μL avec 20 MBq sont optimales car celles-ci permettent de minimiser le nombre d'injections requis.

Cette nouvelle technologie associée aux développements réalisés dans ses travaux démontre la faisabilité, qui a pu être validée sur animal, de pouvoir injecter un liquide contenant des microsphères radioactives en suspension afin de pouvoir traiter efficacement, tout en préservant les tissus sains environnants, des tumeurs inopérables encore de mauvais pronostic aujourd'hui, mais sûrement mieux prises en charge dans un proche avenir.