THE UNIVERSITY OF HULL

Feasibility of Selective Multiple Boosting in the Planning of Intracranial Radiation Therapy

being a Thesis submitted for the Degree of

PhD

in the University of Hull

by

Charlotte Anne Walker BSc(Hons)

August 2007

Acknowledgements

Firstly, my thanks go to the Institute of Clinical Biosciences, University of Hull who generously provided the main funding for this research. A special mention goes to Dr John Greenman, Director of the Institute, for his enthusiasm and optimism as well as initial backing of the project. The Princess Royal Hospital (Hull & East Yorkshire Hospitals NHS Trust), as the sponsoring establishment, provided the rest of the funds, as well a base for my work and myself for the last three years.

Thanks go to my supervisors: Professor Andy Beavis, Dr Gary Liney and Professor Roger Philips, for their expertise, advice, encouragement and direction.

My heartfelt thanks go to the radiation physics department at the Princess Royal, especially to the physics trainees who have kept me sane – you know who you are! I am also grateful to the radiographers at the Centre for MR Investigations, University of Hull, for being so helpful and accommodating during their busy schedules.

Most loving thanks to Simon, for always being there and for giving me the inspiration and motivation to finish this thesis.

Finally, I dedicate this PhD to my parents
for your unquestioning love and support

"Anyone who has never made a mistake has never tried anything new."
Albert Einstein

Abstract

IMRT provides unprecedented means to sculpt radiation dose in three-dimensions, pledging potential to improve local disease control via conformal dose escalation as well as better normal tissue sparing via conformal avoidance. In this thesis, the feasibility and practicality of delivering high doses to intelligently-defined multiple tumour sub-volumes is verified for intracranial radiation treatments. The term 'selective multiple boosting' (SMB) is adopted for the proposed planning scheme. Physical control over local dose deposition is characterised and quantified through the design of two pseudo-anatomy models. The models show that intra-structural optimisation is easily implemented within a standard IMRT planning module. For concentric, spherical boost volumes, regional dose can be controlled at 10 and 5 mm resolution, where a dose differential of 5 and 3 Gy respectively is readily achievable, whilst incrementally boosting neighbouring volumes is more difficult. The limitations of functional imaging techniques are discussed in this context and resolution issues investigated for magnetic resonance spectroscopic imaging (MRSI). Interpolation experiments show that coarse resolutions are not a barrier to using data effectively for defining boost volumes. However, a need for caution and further research on the interpretation and reproducibility of advanced imaging modalities is highlighted. The idea of a composite or hybrid target volume is presented, consisting of multiple tumour contours delineated on different imaging sequences and amalgamated into a probability map of tumour existence. Software has been written for this purpose and demonstrated viable for SMB planning on real patient data. An IMRT plan evaluation toolbox has also been developed, implementing both existing and novel means of characterising the dosimetry and biological effect of these advanced and complex plans. Whilst it is recognised that further work and understanding is required, it is concluded that SMB is not only feasible but also a valuable pursuit in terms of potential clinical gain.

Table of Contents

List of Figures

List of Tables

List of Abbreviations

ADC:	apparent diffusion coefficient
aMR:	advanced magnetic resonance
BDR:	boost dose ratio
BOLD:	blood oxygen level dependent contrast
BTV:	biological target volume
BVR:	boost volume ratio
Cho:	choline
CI:	conformity index
CMRI:	Centre for MR Investigations, The University of Hull
CMS:	Computerised Medical Systems, Inc., St. Louis, MO
CN:	conformation number
CNS:	central nervous system
Cr:	creatine
CSF:	cerebro-spinal fluid
CT:	(x-ray) computed tomography
CS:	conformity score
CTV:	clinical target volume
dDVH:	differential dose-volume histogram
DICOM:	digital imaging and communications in medicine
DRD:	dose ratio distribution
DRVH:	dose-ratio volume histogram
DSC:	dice similarity coefficient
DTI:	diffusion tensor imaging
DVH:	dose-volume histogram
DWI:	diffusion-weighted imaging
EPI:	echo-planar imaging
EUD:	equivalent uniform dose
FDG:	^{18}F-fluorodeoxyglucose
FFT:	fast Fourier transform
FID:	free induction decay
FLAIR:	fluid attenuated inversion recovery
fMRI:	functional magnetic resonance imaging
FOV:	field-of-view
GE:	gradient echo

GBM:	glioblastoma
gEUD:	generalised equivalent uniform dose
GTV:	gross target volume
GUI:	graphical user interface
HI:	homogeneity index
HS:	homogeneity score
HU:	Hounsfield unit
HyMRI:	hybrid magnetic resonance images/ing
IC:	inhomogeneity coefficient
ICRU:	international commission on radiation units and measurements
IPEX:	IMRT plan examination
IMRT:	intensity modulated radiation therapy
LQ:	linear-quadratic
MI:	matching index
MLC:	multi-leaf collimator
MR:	magnetic resonance
MRI:	magnetic resonance imaging
MRS:	magnetic resonance spectroscopy
MRSI:	magnetic resonance spectroscopic imaging
NAA:	N-acetyl aspartate
NB:	no boost
NTCP:	normal tissue complication probability
OAR:	organ at risk
PB:	pencil beam
PET:	positron emission tomography
PID:	pseudo isodose distribution
ppm:	parts per million
pOAR:	pseudo organ at risk
pPTV:	pseudo planning target volume
PRESS:	point-resolved spectroscopy
PTV:	planning target volume
PTVa:	planning target volume annulus
PWI:	perfusion-weighted imaging
QA:	quality assurance
RF:	radiofrequency

ROI:	region of interest
RTOG:	Radiation Therapy Oncology Group
SE:	spin echo
SIB:	simultaneous integrated boost
SMB:	selective multiple boosting
SNR:	signal-to-noise ratio
SPECT:	single photon emission computed tomography
STEAM:	stimulated echo acquisition mode
SUV:	standard uptake value
TC:	target coverage
TCP:	tumour control probability
TE:	echo time
TEP:	tumour existence probability
TPS:	treatment planning system
TR:	repetition time
TVR:	treatment volume ratio
VOI:	volume of interest
WHO:	World Health Organisation

Chapter 1 - Introduction

1.1. State-of-the-Art Radiotherapy

Around 13,000 central nervous system (CNS) tumours are diagnosed each year in the UK. Of these, around 4,000 are primary brain tumours, where the cancer originates in the brain itself, the remaining being metastatic tumours that have spread (metastasised) from a primary disease site somewhere else in the body, most typically from a lung or colorectal cancer. Disease control rates for malignant brain tumours are significantly lower than for other cancers and survival of patients with malignant brain tumours has barely improved over the last fifteen years compared to most other cancer sites. Additionally, brain tumour patients rank highest for the burden of cancer to the individual patient, in terms of years of life lost to adult cancer [1]. High-grade primary tumours of the brain in particular represent a major clinical challenge. The main prognostic factors in brain tumour patients are: origin of the tumour (primary, metastatic), histology (type, grade), size, location, patient age, and Karnofsky performance status. Treatment pathways depend substantially on these factors, where radiotherapy has a role for symptomatic unresectable tumours, recurrent tumours or highly aggressive tumours.

Radiation treatment utilises high energy radiation, most commonly photons, to treat cancer by sterilising tumour clonogens and stopping them from dividing. The probability of doing this and therefore controlling the tumour is proportional to the radiation dose delivered to the patient. Damage caused by this radiation is likely to occur in surrounding normal tissues and whilst these can more readily repair themselves

than tumour cells, planning of treatment regimes nonetheless requires a compromise between delivering a high radiation dose to the tumour volume whilst limiting dose to surrounding normal tissue. A therapeutic ratio is defined as the ratio between tumour lethal dose and normal tissue tolerance i.e. the trade-off between disease control and radiation-induced side effects. Optimisation of the therapeutic ratio must therefore consider the dose distribution in healthy tissue as well as in the tumour.

In the past, radiotherapy consisted of irradiating large volumes of the brain, typically inducing unpleasant side effects [2-4]. More recently, expedited by the advent of the computer-controlled multi-leaf collimator (MLC), limited-field exposures have become canonical treatments for primary brain tumours and solitary metastases. Intensity modulated radiation therapy (IMRT) is the ultimate technique in dose manipulation, providing unprecedented means to accurately plan and deliver complex, conformal 3D dose distributions where these distributions may be designed or 'sculpted' around arbitrary 3D tumour volumes. In this way, the potential for both dose escalation and normal structure avoidance promises improvements in treatment tolerance and outcome.

There is negligible motion associated with the brain and patient immobilisation can be reproduced inter-fractionally and maintained intra-fractionally to better than 2 mm [5]. Also, the brain is comprised of very homogeneous tissues with an almost uniform electron density, making dose calculations relatively simple and accurate. These consistencies, combined with the many critical structures present in the brain, make it a good anatomical site for researching the incorporation of both local boosts and precise avoidance strategies in order to fully realise the capabilities of IMRT.

One approach for maximising the conformality of IMRT is to reduce treatment margins added to tumour volumes, assuming the disease can be accurately defined. Malignant brain tumours however, are particularly infiltrative in nature, leading to an increased risk of marginal misses and recurrences if margins on treatment volumes are reduced too much – in fact 80 % of recurrences occur within 2-3 cm of the original tumour bed. Delineation of target volumes must consider not only the current extent of the tumour but also the pattern of tumour spread at diagnosis and the tumour recurrence pattern after radiotherapy. Unfortunately, larger margins make full target dose escalation unsafe; an increased amount of normal tissue is included in the treated volume, hence limiting dose by risk of complications and thereby negating the conformal gain of IMRT. Target volume definition and margin definition must be adequate for an individual tumour and a philosophy of margin reduction will reach a limit of benefit even with minimal motion and set-up uncertainties coupled with the accurate localisation of sub-clinical disease.

This thesis proposes to maintain current margin expansions, and boost only a sub-region of the target through the use of small-field high precision IMRT. Currently, it is only in cases of very large tumours that multiple volume dose prescribing is considered. Tumour characteristics vary across their extent, with parameters affecting radiosensitivity such as clonogen cell density and oxygenation, distributed non-uniformly. Higher doses delivered to these regions through some kind of tumour profiling, will effectuate an increase in the probability of disease control. This thesis examines the targeting of tumour sub-regions that would benefit from a higher dose, where this relationship is derived from some pre-determined index.

Previous attempts at dose escalation with radiotherapy have focused on phased treatments e.g. shrinking field techniques, whereby a boost is delivered at the end of a course of therapy, consecutive to the majority target dose. Contemporary intensity-modulated beams mean that multiple prescription dose levels can be delivered simultaneously i.e. during the same fraction, resulting in intentionally inhomogeneous dose distributions [6]. This method of boost dose escalation is proposed then as an increase in fraction size to a sub-region of tumour. Since the boost dose is delivered at the same time as the prescribed tumour doses, overall treatment time and the number of fractions remains constant. This approach results in steeper dose gradients (better normal tissue sparing) outside the target volume when compared with uniform dose escalation and even sequential boost delivery [7]. It also allows a sole IMRT plan to be employed for the entire treatment, increasing simplicity and minimising the scope for errors.

Novel inverse planning schemes, including the establishment of optimum dose values as well as improved target delineation must evolve in order to maximise the accuracy and efficacy of simultaneous integrated boost (SIB) methods. Layered targets are sometimes used in current clinical practice, typically derived from an incremental margin definition of the target volume e.g. a boost dose may be prescribed to the gross tumour volume (GTV) or clinical target volume (CTV), from which the planning volume is grown. Although motion is less of an issue in brain therapy, inevitably there is some uncertainty associated with the treatment planning of any anatomical site. Although it is difficult to estimate the contribution to this from some parameters, it is reasonable to assume that the biggest uncertainty lies in the definition of the target volume and critical structures. The development and verification of novel and accurate imaging methodologies are therefore integral in the pursuit of small field high precision

IMRT. Of course other uncertainties arise from the dose calculation model, treatment delivery by MLC and various other elements in the planning & delivery chain.

This thesis approaches the rationale of SIBs in a systematic and pragmatic manner, addressing in particular these issues of uncertainty in the context of the feasibility and utility of novel imaging methods for defining SIBs. In particular, the issue of planning these effective boosts is investigated and feasible planning strategies proposed and proven. The technique shall be henceforth referred to as Selective Multiple Boosting (SMB).

1.2. Dose Escalation in Radiotherapy

Treatment pathways and corresponding dose levels are entirely tumour and patient specific. A typical regimen for primary brain tumours for post-surgical radiotherapy may comprise 60 Gy in 1.8-2 Gy fractions, variable according to tumour histology, location and patient history. For brain metastases, where only a single metastasis is evident, the tumour may be surgically removed or radiosurgery performed. Similarly then, IMRT could be used to deliver this high, conformal dose to the metastasis. If margins are kept small for this initial treatment, whole brain irradiation may still be safe later if required. Alternatively, whole brain irradiation could be delivered with a boost to the evident tumour.

Loco-regional control of cancer is a prerequisite for long-term survival and experience with stereotactic radiosurgery as well as recent conformal therapy trials shows that large radiation boosts can yield good local control. Tanaka *et al.* found a significant survival benefit for patients with malignant glioma treated with 80-90 Gy over conventional 60 Gy treatments, without increases in normal-tissue complications [8]. Shaw *et al.* safely

escalated gross target volume (GTV) dose to 70-75 Gy in fourteen glioblastoma (GBM) patients using IMRT SIBs at 2.5 Gy per fraction and also showed that small volumes irradiated up to 90 Gy are well-tolerated [9].

The predominant failure pattern in high-grade gliomas is local. In fact, most recurrences present within the original GTV, implying that high doses would be best deposited here and lower doses are likely sufficient to control microscopic disease in the rest of the target volume. The GTV itself is heterogeneous in terms of radio-resistance and therefore local dose control within the GTV itself may offer further benefits. Even where conformal dose escalation has been investigated, recurrence is still prolific [10]. 'Intelligent' guidance of the escalated dose may serve to improve the local failure rate, affirming a place for the development of an SMB methodology.

As well as evidence gleaned from clinical trials for dose escalation, Tomé and Fowler showed through radiobiological calculations that selective boosting can return an important gain in tumour control probability (TCP) [11]. Patently, a boost volume that is a greater proportion of the target volume for a tumour of any size will yield a higher increase in TCP. However, Tomé and Fowler also demonstrated that the number of cells i.e. tumour size, influences the effectiveness of simultaneous boosting more than radiosensitivity alone i.e. greater advantage for smaller tumours [11]. The gain in TCP achieved through selective boosting will be different for different tumour characteristics due to their inherent biology. Any dose response curve reaches a plateau at high doses where the TCP is 100 % and further dose escalation would be futile. More resistant tumours i.e. those with a lower baseline TCP may see a higher gain with boosts than less-resistant tumours with less steep dose-response curves. Gains would of course be less if boosted sub-volumes contained a higher than average number of radio-resistant

cells. Mathematical modelling of tumour growth and patient survival is a useful tool in predicting gains afforded by dose escalation and in planning relevant clinical trials [12, 13].

Commonly, malignant gliomas have genetic abnormalities associated with significant resistance to both chemotherapy and radiotherapy and it may be that higher doses of radiation will better control these radio-resistant tumours. Patients with brain metastases are often given radiotherapy with a palliative intent, typically whole brain therapy followed by some targeted treatment on the individual tumour/s. It has been shown that patients with absent or controlled extracranial tumour may gain an actual survival benefit from higher doses of radiation to solitary metastases [14]. Therapeutic radiation then, plays a significant and effective role in the treatment of all different types of brain tumours.

In order to characterise the SMB technique, two metrics are introduced; the boost dose ratio (BDR), the ratio of the boost dose to the prescribed dose to the rest of the target volume, and the boost volume ratio (BVR), the proportion of the target that is boosted. Lee *et al.* and Chan *et al.* for example, excluded patients from dose escalation where the final boost target exceeded one-third of the brain volume [10, 15]. A maximum, safe dose deliverable to arbitrary small volumes has not been proposed. Any such recommendations are improbable even in the future, since dose prescriptions and treatment strategies must always be patient-specific given tumour type, grade, size and location. In terms of dose escalation in contemporary SMB then, a philosophy of 'bigger is better' is to be adopted, assuming of course, adequate normal tissue tolerance.

1.3. Imaging in Radiotherapy

Despite attempts at dose escalation, most failures occur within the irradiated volume. Chan *et al.* showed that even with doses as high as 90 Gy, the predominant failure pattern in high-grade gliomas remains local [10]. In the past, treatments have been limited not only by the capability to plan and deliver conformal treatments but also by the limits of imaging to define the true extent of a tumour. It has been proven that dose escalation with 3D conformal therapy and IMRT is well-tolerated, the quest now then is for improvements in target delineation. Intentional heterogeneity in the target volume will be beneficial in terms of TCP but would be disastrous if the position and extent of the high dose cannot be planned to coincide with the tumour and is deposited in normal tissues.

IMRT dose sculpting is guided by 3D imaging and defining an effective target volume assumes that the image from which the treatment is planned is both accurate and complete. Discrete volumes are binary by nature i.e. voxels are either inside or outside the contoured volume. Imaging modalities such as MRI and CT are central to any conformal therapy but provide only limited information on tumour viability. CT images have excellent geometric accuracy and provide direct electron density figures useful for predicting absorbed dose in the patient but have poor soft tissue contrast and are sub-optimal in assessing tumour extension. MRI provides excellent anatomic detail and is the preferred modality for brain tumours, but no direct electron density information is obtainable. In order to optimise the prognosis for brain tumour patients, it would be beneficial if any therapy could be directed towards the active tumour region. Advanced imaging techniques are capable of non-invasively providing metabolic and functional information to this end and when combined with anatomical images, '4D' data may be extracted for both tumour and critical structures.

Nuclear medicine approaches such as positron emission tomography (PET) or single photon emission computed tomography (SPECT) can be used for this purpose but are expensive and require patients to undergo additional examinations involving intravenous injections of radioactive agents. Also, in order to image activity within brain tumours, radiopharmaceuticals with short half-lives are required, which is problematic for institutions that do not have their own cyclotron/radiopharmacy facility. Evolving advanced magnetic resonance (aMR) techniques however, can be acquired on a clinical MR scanner immediately following a routine imaging study and are totally non-invasive. Multi-voxel proton magnetic resonance spectroscopic imaging (MRSI) for example, is a well-established technique in providing the metabolic status of tissue and can be used to map metabolic activity across an anatomic volume. Using this differential information, it is potentially feasible to better identify regions of metabolically active tumour and deliver boost doses to these areas, resulting in a biologically conformal, patient-specific treatment plan. The idea of incorporating functional imaging into treatment planning is not new. The concept of a 'biological target volume' (BTV) has already been proposed by Ling *et al.* and various efforts have been made towards using such information in the planning process [16-24].

There are many questions surrounding uncertainties in functional imaging, and understanding of the role of these in radiotherapy planning is in its early stages [25]. These include issues of resolution (voxel-size), post-processing techniques and image registration, as well as uncertainties inherent to any specific modality. In order for any functional imaging modality to be applied effectively and reproducibly to a goal of individualised IMRT planning, three major themes must be addressed. Firstly, uncertainties must be quantified and their effects accounted for in some way e.g.

margins. Secondly, a robust interpretation of the biological meaning of the data, accompanied by a correspondence to dose escalation must be developed e.g. standard uptake values (SUV) and finally, novel planning techniques to include both plan development and evaluation will need to evolve.

Whilst much attention is being paid to the potential of functional imaging techniques in radiotherapy planning, robust methodologies for how such images should be interpreted and utilised in the planning remain unclear. Traditionally, planning margins have been used to account for uncertainties due to patient motion, both inter- and intra-fractionally. With the advent of new and more specific target delineation, margins for other uncertainties may become necessary to account for sensitivity and technical issues associated with the imaging technique. Implementation of SMB then must be carried out with great caution and insight to ensure that the gain from the use of the state-of-the-art imaging techniques is not lost due to inadequate or poorly understood data and processing methods e.g. resolution, image registration. Assuming the accurate 3D maps of radiobiological parameters can be constructed and planning strategies devised, consideration must be given as to how accurate and well understood the information is and whether the physical and technical limitations of planning and delivery allow the information to be used in a meaningful way so that any additional scans/procedures the patient undergoes add real value to their treatment. Reliable and reproducible target definition also has implications in multi-centre trials and research.

1.4. Thesis Outline

This thesis proposes the concept of 'Selective Multiple Boosting' (SMB); the use of localised, precisely targeted dose escalation in explicitly delineated target volumes to improve chances of local disease control in intracranial tumours. The research probes

issues surrounding effective target delineation and multiple boost planning methodologies in the context of feasibility, practicality and appropriateness for generating truly individualised and meaningful patient plans. A synopsis of each thesis chapter follows.

Chapter 2 – Intracranial Radiotherapy

Introduces the field of radiotherapy in the context of treatment for brain tumours. Begins with the fundamental theory of radiation therapy and how this has developed into state-of-the-art IMRT techniques, including a detailed discussion on inverse planning and optimisation. Also includes a brief introduction to neuroanatomy.

Chapter 3 – Imaging for Radiotherapy

Discusses the role of 3D imaging in radiotherapy treatment planning. Encompasses background information on CT and MRI, as well as various functional imaging techniques. Highlights current technologies and their usefulness in the radiotherapy setting.

Chapter 4 - IPEX: A Plan Evaluation Environment for IMRT

Detailed introduction to the importance and derivation of plan evaluation tools, from the simple e.g. DVH to the complex e.g. TCP, gEUD. Presents a new plan evaluation environment; 'IPEX', which brings together both existing and novel evaluation tools into one easy to use interface.

Chapter 5 – Controlling Local Dose for Selective Multiple Boosting

Aims to characterise what is physically achievable in terms of dose deposition for a given TPS (CMS[§] XiO®). Planning experiments on 'pseudo' boost anatomies are designed to demonstrate and quantify control over local dose deposition for different theoretical SMB scenarios.

[§] Computerised Medical Systems, Inc., St. Louis, MO

Chapter 6 – Significance of Image Resolution

Examines the issue of coarse resolution associated with certain functional imaging modalities. Investigates the effects of different interpolation methodologies applied to MR spectroscopy defined malignant volumes in terms of volume size, shape, location and resulting dosimetry when they are utilised as boost volumes within a treatment plan.

Chapter 7 – Guiding Selective Multiple Boosting with Hybrid Images

Proposes that different image types can be amalgamated to provide optimal information for boost volume definition. In-house software 'HyMRI' is presented, which combines image sequences into a DICOM parameter map of tumour existence probability (TEP). A five patient study is included, showing that sub-regions where the tumour is known to exist with a greater certainty can be successfully assigned a higher dose, albeit with appropriate patient selection.

Chapter 8 – Conclusions

Concludes that selective multiple boosting (SMB) techniques are feasible in terms of both dose deposition and existing imaging technology for target delineation. Further investigation into the uncertainties surrounding functional imaging techniques is recommended. It is predicted that better understanding will further expedite progress in more intelligent and efficacious target delineation.

Chapter 2 - Intracranial Radiotherapy

2.1. Neuroanatomy and Adult Brain Neoplasia

The central nervous system (CNS) is made up of the brain and spinal cord. The brain is protected by the skull and three thin sheets of tissue, the meninges, which reside inside the skull and further help to shield the brain. At the largest anatomical scale, the brain itself can be divided into three sections; the forebrain (cerebrum), midbrain (brainstem) and hindbrain (cerebellum), as shown in *Figure 2.1*. These can be sub-divided into different regions and sub-regions according to anatomy and function. The cerebrum consists of the right and left hemispheres, which communicate with each other through a thick band of nerve fibres called the corpus callosum. Each hemisphere is further divided into four lobes. The frontal lobes form the anterior region of the cerebrum, and are involved in planning, problem-solving, and higher cognitive functions associated with thought and emotions. The temporal lobes sit posteriorly to the frontal lobes, approximately level with the ears. These are responsible for sound, smell and short-term memory. The parietal lobes are superior to the temporal lobes and posterior to the frontal lobes, at the top of the brain. The primary sensory cortex resides here, as do the visuo-spatial and language processing centres. Finally, the occipital lobes are posterior to the parietal and temporal lobes and it is here that all visual information is processed. The brainstem sits inferior to the cerebrum, anterior to the cerebellum. It connects the brain to the spinal cord and is responsible for vital automatic functions such as breathing and heart rate. Most cranial nerves come from the brainstem. The cerebellum sits most posteriorly, inferior to the cerebrum and controls balance and co-ordination.

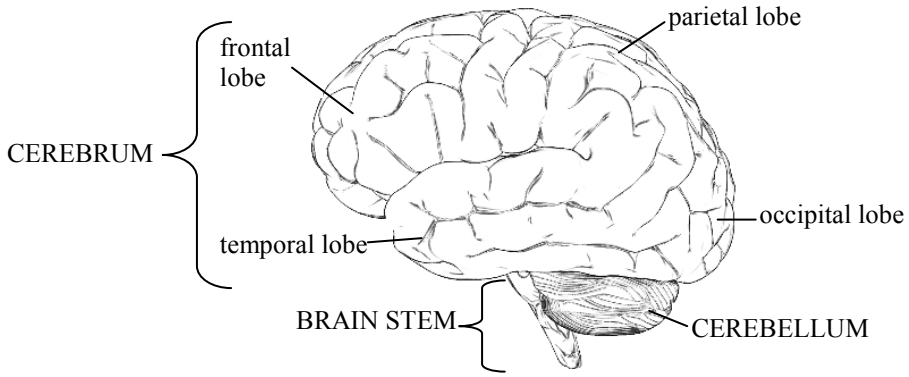

Figure 2.1: Left lateral schematic view of the large-scale anatomy of the brain.

Other significant structures include the pituitary and pineal endocrine glands and the ventricles; cerebrospinal fluid (CSF) filled spaces in the brain connected to the spinal cord and meninges. Most adult brain tumours grow in the forebrain or meninges.

Brain tumour nomenclature is derived either from the type of cell the abnormality develops from or the area in which it is growing. A meningioma for example, is a tumour developed from the meninges. Nearly one hundred different types of brain tumour exist, although around half of all primary brain tumours are some kind of glioma, which develop from the glial cells or neuroglia in the brain. These are also the most common primary CNS tumour seen for radiotherapy treatment. There are three main types of glioma, the most prevalent being astrocytoma.

Tumour grade is a measure of how fast a tumour is likely to grow or spread and refers to how abnormal the cancer cells are at biopsy. Low grade (WHO grades I and II) are the slowest growing brain tumours. High grade (WHO grades III and IV) are the fastest growing and most malignant tumours. Astrocytomas are named according to their grade: grade I astrocytomas are pilocytic, grade II diffuse, grade III anaplastic and the most malignant and fast-growing grade IV astrocytomas: glioblastoma (GBM). The latter two, anaplastic and GBM, are the most common in this group. Ependymomas,

which develop from ependymal cells around the ventricles and spinal cord and oligodendrogliomas, which develop from oligodendrocyte cells, are the remaining two glioma types. Mixed gliomas can also occur, which contain more than one type of cell. In this case, the treatment and prognosis is influenced by the cell type of the highest grade within the tumour. Gliomas may be focal or diffuse i.e. there may or may not be a clear boundary between the tissue and normal brain tissue. In fact, malignant gliomas tend to infiltrate surrounding tissue quite significantly, making them difficult to treat effectively. The second most common primary adult brain tumour is the meningioma, accounting for around 25 % of this group. These are typically low-grade and can often be completely removed with surgery alone.

The median survival of patients with the most common grade of malignant glioma, GBM, is a dismal 12 months with surgery and radiotherapy alone. The addition of Temozolomide has been shown to increase median survival to around 14 months as well as two year survival from 10 to 26 % [26]. However, even with these multi-modality treatments, consisting of some combination of surgery, radiation and chemotherapy, tumour recurrence is virtually inevitable. Gliomas thus present a considerable clinical challenge.

2.2. Radiation Therapy

Radiotherapy uses ionising radiation to treat solid cancers by sterilising tumour clonogens (dividing cells). Early treatments were carried out over a century ago and it remains today, in various forms, a key treatment for many malignancies. Technology in the field continues to improve, whereby advances in computing power and engineering are leading rapidly to more sophisticated and efficacious treatments. Whilst different types of radiation may be used such as photons, electrons and protons, the work

presented here focuses on the most typical; external beam photon therapy. This utilises linear accelerators (linacs), to generate high energy (typically 4-25 MV) bremsstrahlung x-rays by the deceleration of fast-moving electrons in the usual way.

Photons are a type of indirectly ionising radiation, characterised by uncharged particles that interact with matter to produce charged particles (electrons in this case), which go on to ionise atoms and deposit energy within the material, thus attenuating the photon beam. Photons are very penetrating and having no definite range, demonstrating exponential absorption. Secondary electrons are produced via one of three interaction processes: the photoelectric effect, Compton effect or pair production. The predominance of the phenomena depends on both incident photon energy and the atomic number, Z, of the absorbing material. At therapeutic energies, the most dominant interaction of photons with the patient is the Compton effect (Compton scattering). In this mechanism, the bombarding photon having quantum energy $h\nu_0$, gives up some of this energy to an atomic electron, which is emitted at an angle, θ. The photon, of depleted energy $h\nu'$, is scattered at another angle, ϕ. For this to occur, the energy of the incident photon must be much larger than the binding energy of the atomic electron i.e. the electron must be 'free'. Because of this, the Compton interaction mode is entirely independent of Z (the number of electrons) and depends only on the photon energy. At higher therapeutic energies, interaction may occur by means of pair production. In this case, the photon gives up all of its energy and is absorbed in creation of both a negative (e^-) and positive (e^+) electron, the resulting pair having a kinetic energy of ($h\nu$ - 1.02) MeV. A threshold photon energy of 1.02 MeV is required in order for an electron pair to be emitted in this way, since the rest mass energy of one electron is 0.51 MeV.

Once electrons have been generated via one of the processes described above, the charged particles go on to impart energy to the matter by way of Coulomb-force interactions. Collisions between the electron and nucleus of an atom generate excitation and sometimes ionisation of the atom, which result in a loss of some or all of the electron's energy. Electrons undergo multiple collisions and changes in direction of motion through the material, as they gradually deposit their kinetic energy.

Incident therapeutic x-ray beams consist of a large number of photons with a variety of energies and are characterised in terms of fluence. For a number of photons N, entering a sphere of cross-sectional area A, fluence and energy fluence are given by

$$\Phi = \frac{dN}{dA} \tag{2.1}$$

$$\Psi = \frac{dE}{dA} \tag{2.2}$$

respectively, where E is the sum of energies of the N photons. Intensity is given by the energy fluence rate, and is also a useful term in characterising radiation fields. There is a need to link incident energy fluence to dose received by the patient in order to quantify the radiation received and signify the biological effects initiated by it. Absorbed dose, D, results from the energy flux (due to photons) through the patient and is derived from the energy absorption per mass. D is defined as the total energy imparted to tissue

$$D = \frac{d\varepsilon}{dm} \tag{2.3}$$

where $d\varepsilon$ is the mean energy transferred to material of mass dm by the ionising radiation. The SI unit for absorbed dose (henceforth referred to simply as 'dose') is the gray (Gy), equivalent to 1 Jkg^{-1}. Dose can be estimated in three dimensions using complex calculation systems which model radiation transport through different

geometries and materials within different patients. Many algorithms exist that account for differing processes, and these vary in accuracy and computation time, which must be traded against one another. An accurate dose calculation tracks both primary and scattered radiations in three dimensions, whilst accounting for density differences i.e. tissue inhomogeneities.

Radiotherapy is administered to patients with one of two intents: radical therapy aims to achieve curative control of the tumour and palliative therapy provides pain relief and increased quality of life for those with advanced disease. Curative radiotherapy aims to eradicate, by sterilisation, all clonogenic tumour cells within a defined treatment volume without damaging close-lying tissues and organs. It works by damaging the DNA molecules of the tumour stem cells to induce cell death or toxicity (cytotoxicity) and stop them dividing. This is done most effectively if both DNA strands are broken close to one another i.e. a double strand break, as single strand breaks can repair themselves more easily. If multiple double strand breaks are induced, there is even less likelihood of repair.

Clonogenic cells proliferate via a four-phase cycle (*Figure 2.2*) and their variant sensitivity to radiation (radiosensitivity) is distributed around this. For this reason, treatment regimes are planned in repeated small dose 'fractions' delivered over a period of time, typically around thirty days. In general, normal (i.e. non-cancerous) tissues exhibit better repair fidelity and so fractionation schemes also serve to give healthy tissues chance to repair themselves, thus limiting radiation-induced toxicity.

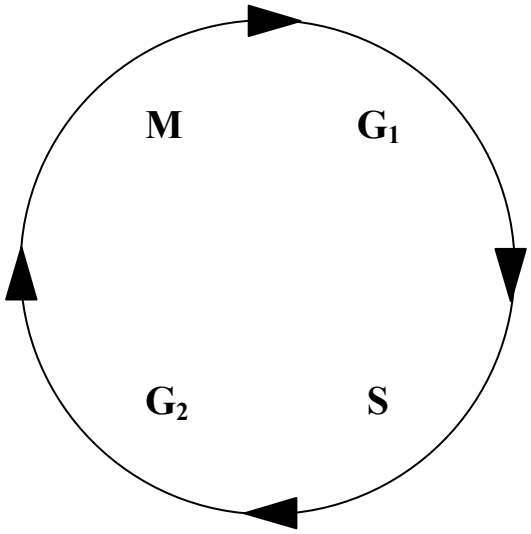

Figure 2.2: The four-phase cell cycle; G_1 phase: first phase of cell growth, S phase: DNA synthesis/replication, G_2 phase: second phase of cell growth, M phase: mitosis/cell division.

Radiosensitivity is affected by other factors as well as the cell cycle, some cells may be hypoxic (oxygen deficient) for example, making them particularly resistant to radiation. Assuming all clonogenic cells are irradiated, a failed treatment i.e. local recurrence of the tumour, occurs when some of the dividing stem cells repair themselves and repopulate. The probability of achieving local tumour control is a function of how many tumour cells have been sterilised, which is proportional to the radiation dose received. The probability of controlling the tumour then, is proportional to the radiation dose delivered to the patient. The deliverable dose is limited however by the tolerance of surrounding normal anatomy, since the photon interactions do not discriminate between healthy and malignant tissue (*Figure 2.3*).

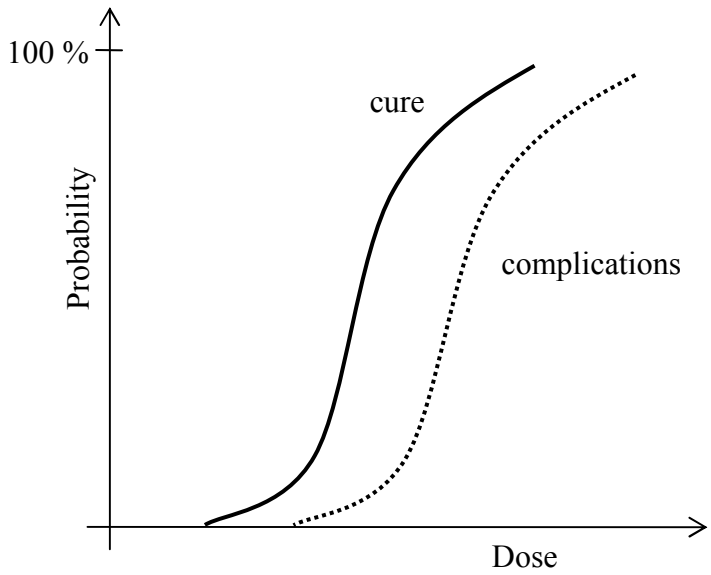

Figure 2.3: Increasing radiation dose to assure high probability of disease control has adverse effects on surrounding normal tissues. Radiotherapy treatments are a compromise between achieving a tumourcidal dose and minimising side effects. A 'therapeutic ratio' is defined as the ratio between tumour lethal dose and normal tissue tolerance.

Radiotherapy is a critical modality in the treatment pathway of brain tumours, where it is used post surgically for high malignancy and/or incomplete resection and sometimes instead of surgery altogether when this is not viable.

2.3. Intensity Modulated Radiation Therapy (IMRT)

Intensity modulated radiation therapy (IMRT) is a state-of-the-art highly conformal radiotherapy technique, providing unprecedented means of generating complex, conformal dose distributions in three dimensions. The main planning and delivery methods for IMRT were already established by 1995, although commercial products have only become available in the last six or seven years. Many UK centres still do not have the capability or expertise to offer IMRT at all, and for many others its use in its infancy. The Princess Royal Hospital, Hull, UK has been treating suitable patients with IMRT since January 2002, and the process here is now well established. IMRT

planning is undertaken with the Computerised Medical Systems, Inc.[§] (CMS) XiO® 4.2 IMRT module.

An intensity modulated beam is one whose fluence cross-section (profile) is highly non-uniform, enabling the radiation to be 'sculpted' around a particular tumour shape (*Figure 2.4*). These beams are most commonly delivered by a multileaf collimator (MLC), a computer controlled device for delimiting the radiation beam. MLCs have jaws which are subdivided into independent, abutting (usually tungsten) leaves, enabling the direct creation of complex field shapes. A series of these shaped fields can then be superimposed to produce the required intensity modulation across the beam. In this way, IMRT is an effective treatment for tumours with close-lying sensitive structures that may limit the dose deliverable by conventional means as well as providing the capability to deliver simultaneous dose escalation or 'boosts' to small sub-volumes within a tumour and attempt better tumour control. That is, IMRT effectuates a dual-benefit, in that it may be utilised for conformal therapy and conformal avoidance both independently and simultaneously.

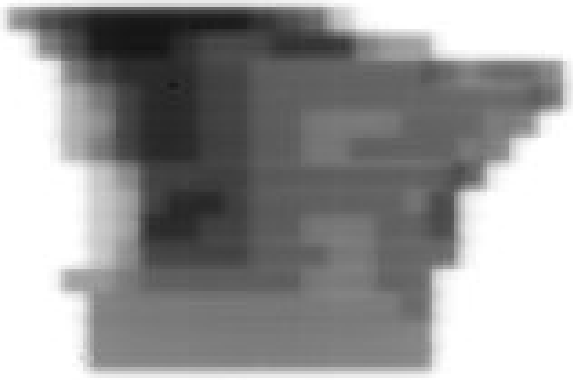

Figure 2.4: Greyscale representation of the cross-section of an intensity modulated beam.

[§]St. Louis, Missouri

IMRT treatment plans are designed using an 'inverse planning' methodology, whereby the radiation beams themselves are derived from the desired dose distribution via a mathematical search. The planning is described as inverse, since the treatment plan is designed 'backwards' from what dose distribution is actually desired, as opposed to the conventional 'forward' trial and error approach. Many beam combinations are tried and beam design depends on the patient, tumour shape, existence and proximity of organs at risk (OARs) and the dose levels prescribed by the clinician. In order to iteratively test and improve these beam combinations, an optimisation algorithm operates under the guidance of an appropriate objective or cost function. Due to the linear relationship between fluence and dose, local fluence is employed as the optimisation variable. This relationship only holds of course to the first approximation i.e. dose at a point. Beyond this, scatter complicates the relationship and primary doses must be convolved with dose spread functions. Although inverse IMRT planning is defined and solved as an optimisation problem, this does not mean that resulting plans are optimal in any clinical sense. The perfect solution i.e. high dose gradient across PTV boundary, homogeneous PTV dose and zero dose to OARs and everywhere else is practically impossible – the goal is to find a physical solution from a pool of feasible solutions that best meet the given requirements.

The cost function simply represents the conditions, goals and limitations defined by the oncologist and treatment planner and the planned dose distribution is generated within these parameters. Constraints imposed may include a minimum acceptable dose to the tumour, maximum acceptable dose to proximal OARs and relative importance weightings that rank these trade-offs. The oncologist's prescription defines both dose and dose-volume requirements i.e. doses prescribed to each structure and also how much of a structure should receive a specified level of dose. The treatment planner pre-

selects the beam configuration (number and direction of fields) and translates the clinician's prescription into optimisation objectives. The best possible (within the constraints of set goals, patient geometry and algorithm limitations) dose distributions are returned by the optimiser and once the planner has accepted these, delivery instructions e.g. MLC segments, can be calculated. The simplified flow of the inverse planning process can be seen in *Figure 2.5* below.

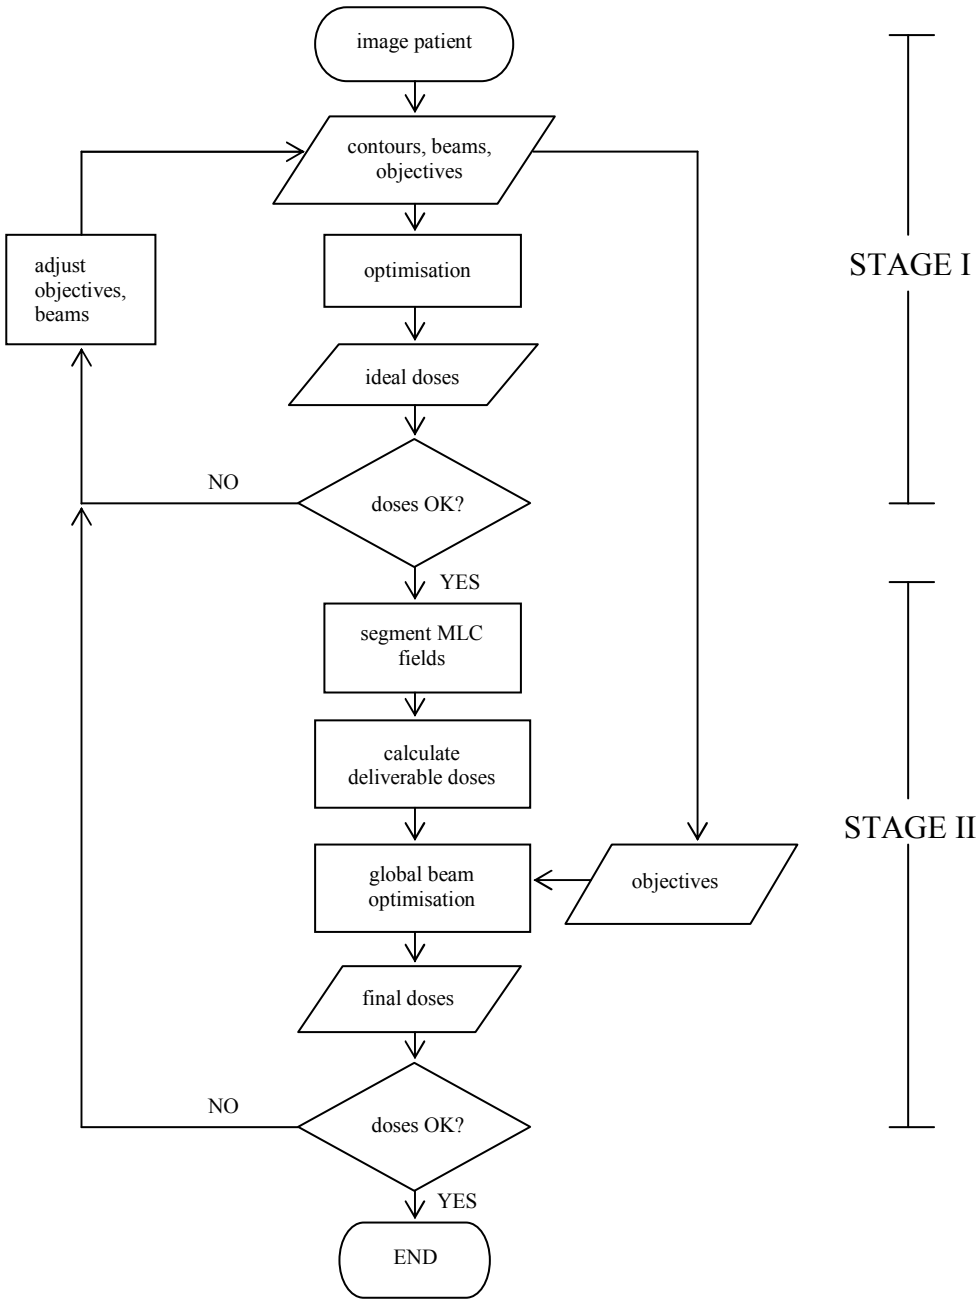

Figure 2.5: XiO® IMRT inverse planning workflow; stage I optimisation generates the ideal dose distribution, stage II calculates the physically deliverable dose distribution.

2.4. Contouring

The patient volume is discretised into voxels and all anatomy structures included in the optimisation must be contoured prior to optimisation. These contours are used to identify voxels claimed by different objectives and the XiO® optimiser uses this information exactly. For any 3D planning approach, a set of axial images are acquired (most commonly CT) of the patient, on which the region for treatment is delineated. Any relevant OARs and an external contour representative of the skin (or immobilisation mask) must also be drawn on every image slice. According to the ICRU Report 50, the target volume should compose of three individual volumes and be contoured on the image dataset by the radiation oncologist [27]. In this context, the tumour itself, as visible on the images, is termed the gross target volume (GTV) and the clinical target volume (CTV) encompasses the GTV, plus a margin to account for microscopic spread of the disease. A review of margin recipes that consider uncertainties in volume determination and location is given by van Herk [28].

Margin sizes are tumour specific - low grade and slow growing tumours such as some meningiomas may be suited to treatment with smaller margins and in fact, for benign tumours (assuming non-infiltrative), a CTV is sometimes not necessary at all. Similarly, metastases tend to be well localised with minimal invasiveness and margins for sub-clinical spread can also be kept small in this case. Margin reduction for high grade, fast proliferating tumours is not feasible due to their infiltrative nature. A margin of several cm of normal appearing brain tissue is typically added to visible malignant tumour to try and prevent recurrence at the edge of the irradiated volume. The final planning target volume (PTV) besets the CTV with a margin added to include motion and setup uncertainties. Given standard calculation methods, the PTV is designed to enclose the CTV with the 95 % isodose on 90 % of treatment occasions and dose is

planned to conform to this envelope. The PTV is treatment specific i.e. dependent on anatomical site (organ motion) and setup reproducibility. For a patient immobilised with a thermoplastic mask undergoing 3D radiotherapy of the brain, a CTV-to-PTV margin of approximately 0.5 cm is typical. The ICRU later recommended sub-dividing the PTV into an internal margin (IM) plus a set-up margin (SM), accounting for size, shape, position variations of the CTV and patient-beam positioning uncertainties respectively [29]. In this same report, an analogous PTV enveloping OARs is proposed to consider movement, changes in shape and/or size as well as set-up uncertainties termed the planning organ at risk volume (PRV). The role of such explicit margin definition lies in sites such as the prostate, where both organ motion and set-up uncertainties can be significant and in cases where OARs may be highly serial or mobile and to a much lesser extent in the brain, where both motion and set-up variations are minimal. Where appropriate, the term target volume in the work presented here refers to the PTV; the full extent of the volume to be treated. In brain treatment planning, OARs inspected generally include one or more of the following depending on tumour location: eye orbits, lenses, optic nerves, optic chiasm, brainstem, pituitary and parotid glands.

XiO® provides standard contour types such as patient and target volumes and the user can define further contours of class 'Int' (internal). In general, these are typically the OARs deemed important for a particular site. However, it is straightforward to redefine 'Int' contours as targets, whereby multiple boost volumes can be defined within the PTV itself and used within the optimisation. Imaging and accurate delineation of target and OAR geometry are a critical part of the inverse planning process, since voxels lying within the bounding contour of a particular structure are assigned as belonging uniquely to that structure within the optimisation, from which the treatment plan is designed.

2.5. Optimisation

The goal of any optimisation is to find a solution corresponding to the extremum (typically the minimum) of some objective function, which formulates goals that should be met in order to achieve the desired, final solution. Mathematically, the problem in *n*-dimensional Euclidean space is stated as: given $f : \Re^n \mapsto \Re$,

$$\begin{aligned} &\text{minimise } f(\underline{x}) \\ &\text{subject to } \underline{x} \in \Omega \subset \Re^n,\ \underline{x} = \left[x_1,....,x_n\right]^T \in \Re^n \end{aligned} \tag{2.4}$$

Here, x_1,…..,x_n are the decision variables or system parameters, $f(\underline{x})$ is the objective function returning a scalar quantity and $\Omega \subset \Re^n$ is the constraint or feasible set. The 'best' vector $\underline{x}$, i.e. the one that gives rise to the smallest value of f, is called the minimiser of f over Ω and each set of adjustments to $\underline{x}$ is referred to as an iteration. This is a constrained optimisation problem, since $\underline{x}$ is constrained to lie in Ω. For $\Omega = \Re^n$, the problem is said to be unconstrained.

During IMRT optimisation, the fluence profile of each radiation beam is modelled as a collection of rectangular sub-fields called beamlets, the individual intensities of which are variable. Beamlets can be thought of as 1D intensity profiles whereby the number of photons travelling along each beamlet direction can be controlled. Changes in fluence are simulated by changes in beamlet 'weight', meaning the number of monitor units (dose units) per field. These individual beamlet weights are the system variables, x_1,…..,x_n, within the optimisation scheme. It is usual to constrain beamlet weights to be non-negative, otherwise certain solutions may equate to negative fluence. Beamlet dimensions depend on the device used for delivering the irregular fields. For delivery by a multileaf collimator (MLC), the beamlet cross section is defined by the physical MLC leaf width in one dimension and a measure in the direction of leaf movement.

Typical beamlet size is just 5 x 5 mm^2. With such a large number of variables ($10^3 \ldots 10^4$), the number of physically feasible IMRT plans is huge and as such any optimisation algorithm is computationally intensive. For 1000 beamlets for example, the search is in 1000-dimensional space, where each different combination defines a point in space i.e. each point in the search space is a vector of beamlet weights.

2.5.1. Objective Function

3D dose distributions proposed by the optimisation algorithm are tested against that desired using the objective function, which produces a scalar value, to score each plan generated for a given set of parameters. Generally, the objective function is sought to be minimised whereby a lower cost corresponds to a better plan. The objective function, along with pre-selected beam parameters such as directions, energy etc. dictates the search domain for the optimisation.

The cost function is made up of either objectives or hard constraints depending on whether unconstrained or constrained optimisation methods are used. In constrained optimisation, the cost function value is forced to be at or below the prescribed boundary, whereas in unconstrained optimisations, the value of the cost function is minimised during the optimisation, with the different objectives competing against one another under some priority scheme. In unconstrained optimisation, the objective function models a particular clinical objective and is usually composed of sub-objectives corresponding to individual anatomic structures and different goals for each structure. The most common objective function to represent this is the quadratic or least-squares function. This is a convex function i.e. the value at the midpoint of any interval in its domain is lower than the average of its values at the ends of the interval. This means that any optimisation problem based on such a function has only one local minimum;

the global minimum. The quadratic objective function minimises the sum over all voxels of the squared difference (or ratio) between prescribed and actual dose. It is minimised when dose is uniform in a target volume, but is relatively insensitive to a few low dose voxels. This may be a problem if dose is so low that there is a chance of clonogenic cells surviving in that voxel. Also, it may not distinguish between a very low dose voxel and a very high dose voxel with the same dose difference (dose ratios may be preferable in this regard). Similarly, the objective function may be insensitive to a few hot voxels in a given OAR, which potentially may cause clinical complications. Any objective function based on a mean square deviation lacks spatial information and can therefore result in inhomogeneous dose distributions. Xia *et al.* showed that using exponents greater than the usual power of two can improve homogeneity [30].

Clinical objectives in IMRT optimisation may be expressed in terms of limits on dose, dose-volume combinations (volume of tissue receiving a certain dose) or in terms of biological response. Most physical objective systems use a combination of dose and dose-volume objectives. Both maximum and minimum dose criteria should be specified to limit hot spots as well as cold spots. Dose-volume criteria provide more flexibility for the optimisation process and better control over dose distributions in some cases than simple dose-based goals. Wu and Mohan (2000) showed that optimisations with dose-based objectives converge faster in fewer iterations than optimisations with dose-volume criteria, but the resulting plan may not be as close to that desired [31].

In principle, the most clinically relevant optimisations should be based on biological objective functions that model the radiobiological response of the involved structures. However, the biological response of a tumour to radiation is not well understood. Lian and Xing incorporated statistical uncertainties associated with radiobiological models

into the optimisation process and found this improved final solutions via lessening the influence of inter-patient variation of these biological characteristics [32]. Jones and Hoban investigated the use of a biological optimisation method compared to a physical dose method and found that biologically based optimisations had less homogeneous dose distributions [33]. Because of the uncertainty in dose-response models, and because of the variability of response between patients, practical plan optimisation is usually based on physical dose objectives, although attempts have been made to combine the two in a hybrid physico-biological approach [34].

Unconstrained optimisation leads to objectives competing against one another resulting in a trade-off between prescriptions. Clinical objectives are, by the nature of the problem, incompatible with one another and so a method of controlling the trade-offs must be incorporated. This is achieved through the use of importance weights, which serve to scale the influence of different structures and objectives on the optimisation. The weighting factors are normally set by the treatment planner based on a combination of experience and trial-and-error, since the values do not correspond to any physical quantity. The process is subjective and not intuitively obvious and as such, attempts have been made at determining them as an optimisation component, and at updating them automatically according to optimisation results [35, 36]. Constrained optimisation uses hard constraints to meet dose goals, without the need for importance factors at all.

2.5.2. Search algorithms

Search methods iteratively reduce the objective function value by carrying out ‘tests’ near an estimate of the solution to determine the direction in which to search. The intensities of the beamlets are then changed iteratively according to some update sequence until predefined convergence criterion are met, which may be the number of

iterations or some relative value of the objective function [37]. For multi-variable optimisation ($f : \Re^n \mapsto \Re$), as in IMRT, the search is most commonly undertaken with gradient methods. Gradient descent methods, a simple type of deterministic search algorithm, are most common in IMRT optimisation and are used by most commercial systems [31]. These tend to reach a reasonable solution faster and in fewer iterations than do stochastic methods such as simulated annealing.

Gradient-based approaches to unconstrained optimisation are the steepest descent, conjugate gradient and Newton's methods. Steepest descent is the fundamental first-order method and requires that the objective function be quadratic. The method of steepest descent uses only first-order derivatives i.e. the gradient, to select suitable search directions, since the gradient denotes the variation of the objective function in the solution space (*Figure 2.6*). $-\underline{\nabla} f(\underline{x})$ points in the direction of the maximum rate of decrease of f at $\underline{x}$ and is therefore the best direction to search in to find a minimum. Given a point $\underline{x}_{(i)}$ a movement to the point $\underline{x}_{(i+1)}$ is given by

$$\underline{x}_{(i+1)} = \underline{x}_{(i)} - \alpha \underline{\nabla} f\left(\underline{x}_{(i)}\right) \tag{2.5}$$

where the positive scalar $\alpha_i > 0$ is called the step size. At each step, starting from the point $\underline{x}_{(i)}$ a line search is conducted in the direction of $-\underline{\nabla} f\left(\underline{x}_{(i)}\right)$ until a minimiser $\underline{x}_{(i+1)}$ is found. The descent takes place in orthogonal steps and the procedure is repeated iteratively, the gradient tending to zero as the minimiser is approached. If, for some *i*, $\underline{\nabla} f\left(\underline{x}_{(i)}\right) = \underline{0}$, then the point $\underline{x}_{(i)}$ satisfies first-order necessary conditions and the algorithm stops. In practice, this rarely happens and a practical stopping or convergence criteria is to check the norm

$$\left\| \underline{\nabla} f\left(\underline{x}_{(i)}\right) \right\| < \varepsilon \tag{2.6}$$

where ε is a small specified tolerance. Alternatively, the norm between two successive iterations may be computed

$$\left\| \underline{x}_{(i+1)} - \underline{x}_{(i)} \right\| \tag{2.7}$$

and the algorithm halted if it is less than some tolerance. The method of steepest descent may have a fixed or variable step size, chosen to achieve the maximum amount of decrease of the objective function at each iteration. Algorithms using larger step sizes undergo fewer gradient evaluations than a more direct approach with smaller step-sizes but may result in a more zigzag route to the minimum, slowing the rate of convergence compared to other algorithms. The conjugate gradient method involves searching along each of n mutually conjugate (orthogonal) directions instead of along axes directions. These conjugate direction vectors are generated from gradient evaluations and the minimum is found in exactly n iterations.

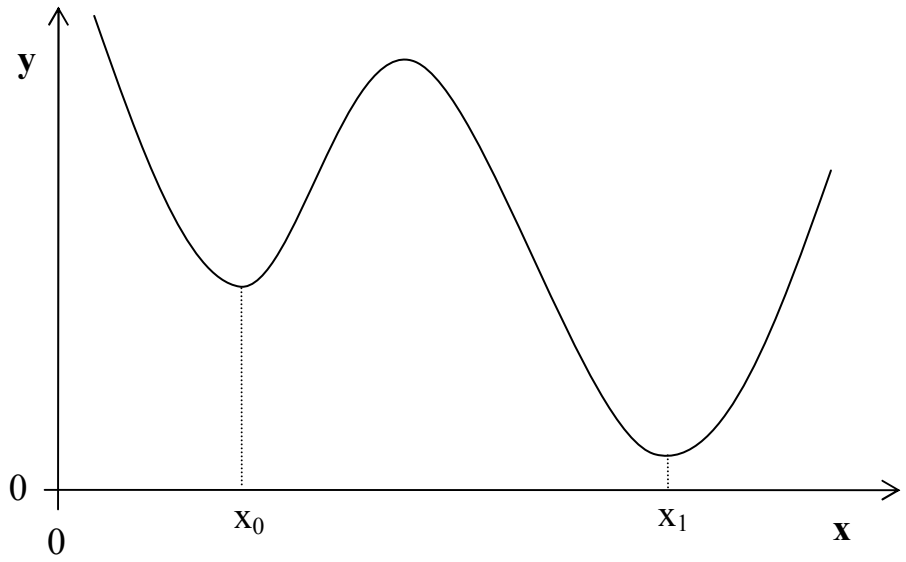

Figure 2.6: Example of a scalar, real-valued function, f, of a single variable $f : \Re \mapsto \Re$ *without constraint, having two minima; a local minimum at* x_0 *and the global minimum at* x_1.

A given objective function may have multiple local minima, as seen in *Figure 2.6*, and a deterministic search program like the gradient method may become trapped in one, corresponding to a sub-optimal score. The use of dose-volume constraints in particular, as well as dose-response objective functions may introduce multiple local minima

although any gradient-based algorithm may potentially get trapped in a local minimum [38, 39]. The existence of local minima is determined by the shape of the objective function but it is the choice of search algorithm that determines whether the traps can be avoided. Wu and Mohan (2002) investigated local minima traps and found that the number of minima depends on the parameters of the objective functions and the consequences of them are dependent on the number of beams, but conclude that multiple local minima do not affect the clinical outcome of optimisation [40]. Similarly, others have shown that local minima traps are not usually impediments to finding satisfactory solutions in IMRT optimisation [41, 42]. In any convex programming problem, every local minimum is a global minimum and so for a least square/quadratic objective function, there's only one, single global minimum. Indeed, even if a global minimum is reached, the plan may not be satisfactory clinically, it is only optimum according to the given criteria. Additionally, there is a degeneracy in that a solution may not be unique – different beam weight combinations may produce the global minimum since the result depends on the initial weights selection.

XiO® optimisation is carried out in two stages, as shown in *Figure 2.5*. Stage I generates 'ideal' intensity maps; a greyscale representation of the dose distribution, and these must be accepted by the user in order to proceed to stage II, which converts the ideal solution to a deliverable form (MLC segments). The dose is then recalculated for each individual segment and this final dose distribution can then be checked for approval.

2.5.3. Practical IMRT Optimisation

The cost function is a sum of objective functions which are anatomic structure-specific functions of dose. The general XiO® cost function is

$$f = \sum_{\sigma=1}^{m} \frac{w_\sigma}{n} \sum_{i=1}^{n} \left(D_i - D_0 \right)^2 \tag{2.8}$$

where f is the weighted average cost per voxel, m is the number of objectives, w_σ is an objective-specific importance weighting, n is the number of voxels in the structure and $\left(D_i - D_0 \right)^2$ represents the cost per voxel at dose voxel D_i for objective σ. This form is volume normalised to eliminate large structures being favoured by the optimisation process. One problem with this formulism is that a few far outlying dose points might contribute to the cost as much as a larger number of violating dose points nearer to the reference dose. One advantage is that convex cost functions such as these arrive at a single, global minimum.

The planner formulates the objective function by specifying dose and/or dose-volume goals for both targets and OARs. Targets may have minimum and maximum dose objectives and OARs can have maximum and dose-volume objectives. A goal or threshold dose may be stated but these are not themselves objectives. The planner must ensure that the PTV and any other tissues to be evaluated are fully enclosed by the calculation volume. A ranking system determines how voxels are shared when contours overlap i.e. which objectives are affected. Regions where contours overlap may be designated as belonging to the PTV, the OAR or neither, depending on clinical requirements. Of course, when two or more structures are in contact with each other, none of them will completely achieve its dose objectives, since that would require an impossible instantaneous change in dose i.e. an infinite gradient. A relative importance weight with respect to other objectives can be specified and a power, representing the magnitude of penalty applied to dose outside acceptable ranges. For any unconstrained optimisation, there is no predictor of how high an importance weight must be before an objective is satisfied; inverse planning is iterative, experimental and rarely intuitive.

Since an OAR is, by its definition, in the immediate neighbourhood of the PTV, dose cannot be forced to very low levels immediately outside the target without compromising tumour dose and so a match may be best in the least squares sense but not clinically acceptable. To overcome this, an 'optimisation margin' may be introduced around the transient region to minimise infinite gradient problems whereby points within the margin are free from any constraints and do not contribute to the objective function. MLC leaves do not cover this transition zone, allowing a rim of beamlets around the outside of the target (*Figure 2.7*). Choosing ineffective margin widths however, could lead to a compromise in the dose conformity. If the margin is zero, the algorithm optimises only those beamlets that intersect the target. The optimisation margin serves to get rid of 'cold spots' that would otherwise occur close to the target boundary. The default value is 0.5 cm.

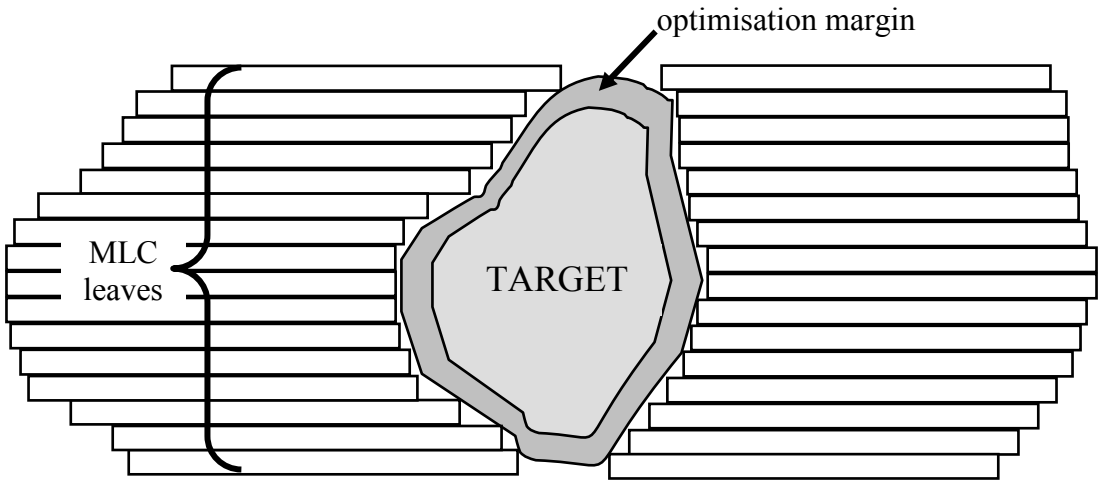

Figure 2.7: Optimisation margin in a beam's eye view.

The final dose distribution is highly dependent on the mode of delivery as well as the objectives and optimiser used to design it. For delivery by MLCs, the jaws are conformed to the PTV before the optimisation is initiated. This is done with a margin to account for the penumbra of the radiation beam i.e. the dose fall-off at the beam edges.

For delivery by MLCs, the user specifies the resolution in the direction of leaf movement; the step increment. The beamlets are defined as square or rectangular in this way. It is typical to define square beamlets for a step-and-shoot mode of delivery. A small leaf step increment may result in some beamlets having no dose voxels and therefore zero intensity and generate more MLC segments than a larger increment. The number of segments relates to issues of treatment time whereby fewer segments correspond to a shorter treatment time, as it takes time for the leaves to move into each new position. The user also sets the scatter extent distance, as the distance to which a beamlet contributes dose beyond its geometric edge. A larger scatter distance results in a longer time to complete optimisation, but the optimised dose distribution is usually closer to the final dose distribution. This may be of particular importance when conformity is critical e.g. proximal radiosensitive OAR.

XiO® uses an unconstrained conjugate gradient optimisation [43, 44]. Gradient descent methods use the gradient of the objective function with respect to the input parameters, the beamlet weights in this case. The gradient is then a normalised vector of partial derivatives of the objective function with respect to beamlet weights. The initial search direction is thus given by

$$-\nabla f = -\left|\frac{\partial f}{\partial w_1}\right|\hat{w}_1 - \left|\frac{\partial f}{\partial w_2}\right|\hat{w}_2 ..... - \left|\frac{\partial f}{\partial w_n}\right|\hat{w}_n \qquad (2.9)$$

where f is the cost function, $w_1, w_2 ..... w_n$ are beamlet weights and $\hat{w}_1, \hat{w}_2 ..... \hat{w}_n$ are unit vectors along $w_1, w_2 ..... w_n$ i.e. the 'axes' in search space. The gradient is multiplied by -1 to define the search direction, since the problem is one of function minimisation.

Prior to stage I optimisation, initial beamlet weights are set to 1.0 and beamlet voxel doses normalised so that the combined dose at the isocentre is equal to the target goal

dose. Each iteration consists of a gradient calculation and then a line search in the direction of the negative gradient. A line search is simply a search for the minimum in a specific direction within the function space where each line search direction is orthogonal to the previous direction. The line search magnitude, the normalised distance in line search direction from some starting point in the search space, is referred to as the step size and controls the precision of the result. The line search updates beamlet weights and evaluates the effect of the changed doses on the cost function. If the value of f does not decrease, the optimiser stops.

Step size, the number of iterations and convergence criterion are controlled by the user. Convergence is stated in terms of a percentage where

$$\% \text{ convergence} = \frac{\text{previous cost} - \text{current cost}}{\text{initial cost}} \times 100 \tag{2.10}$$

The user specifies convergence criterion both as (2.10) and as a maximum number of iterations. Optimisation is complete when it reaches one or the other, whichever occurs first. This is the same for both stage I and stage II optimisation. The rate of progress to minimisation is asymptotic and approximately 95 % of all improvement occurs in the first thirty iterations.

At the end of stage I optimisation, the intensity map for each optimised beam is generated. This is a graphical representation of the IMRT beam used to show the optimised intensity i.e. the ideal fluence of the beamlets projected to the isocentre. XiO® uses a sliding window leaf segmentation algorithm to generate MLC segments from the ideal intensities, user-specified number of discrete intensity levels and a minimum allowed segment size. A segmented MLC intensity map is then produced, showing the relative intensities of the beamlets, derived from the segmented MLC leaf segment weights and shapes. After segmentation and once the dose calculation is

complete, the leaf segments and sequence may be reviewed. XiO® uses a pencil beam (PB) based fast 3D superposition for dose calculation during optimisation and superposition or convolution for calculating the final, deliverable dose distribution.

With IMRT dose optimisation there is generally no net gain – improvement in the dose to a structure is often accompanied by dosimetrically adverse effects at other points in the same or different structure. The heart of planning lies in the fact that some dose distributions will be more clinically acceptable and desirable than others. Treatment plans are evaluated and competing plans ranked with various quantitative and visual tools. Isodose lines connect points of equal dose to provide a homogeneity metric, and 3D dose coverage can be displayed graphically to give an indication of conformity. Dose statistics and the dose-volume histogram (DVH) invariably form part of the dose distribution analysis. The integral (cumulative) DVH is most commonly used, which is a plot showing the fraction of a structure volume receiving a specified dose or greater. 3D dose can be converted to biological outcomes using tumour control probability (TCP) and normal-tissue-complication probability (NTCP) models. These are predictors of treatment success and side effects respectively, in fact *Figure 2.3* actually depicts the sigmoidal TCP and NTCP as dose-response curves.

All work presented in this thesis is planned on an XiO® IMRT workstation to deliver 6 MV photons on a Varian Medical Systems, Inc.[§] Clinac 600C equipped with a 120 leaf MLC (60 leaf pairs). With this setup, the maximum field at isocentre is 40 x 40 cm^2, with a leaf width that projects to 0.5 cm (at isocentre) in the central 20 cm of the field (outer 20 cm of field = 1.0 cm) i.e. so long as the diameter of a given PTV is smaller than 20 cm, the MLC leaf width is always 0.5 cm.

[§]Palo Alto, California

Chapter 3 - Imaging for Radiotherapy

3.1. Introduction

Contemporary 3D imaging methods are fundamental to the deployment of conformal radiotherapy. Serving to drive correct target volume definition and accurate, reproducible patient setup, quality imaging is a prerequisite to safe IMRT. Although imaging plays an extensive and multifaceted role in radiotherapy i.e. diagnosis, treatment positioning, planning, verification dosimetry and follow-up, the work presented in this thesis looks only at the use of advanced imaging for treatment planning, specifically target delineation.

Axial 3D imaging data, most commonly computed tomography (CT) or magnetic resonance imaging (MRI), is stored digitally on a computer in a standard DICOM format and can be transferred electronically over a network link to the treatment planning system (TPS). Here, the targets and relevant normal structures can be outlined on each individual 2D image, from which the TPS renders 3D volumes. Dose is then planned to conform to these boundaries precisely. In this way, treatment planning is heavily dependent on structure delineation, the sensitivity and specificity of which is determined in most part, by the imaging modality.

3.2. X-Ray Computed Tomography (CT)

CT became popular for diagnostic purposes in the UK in the early 1980s and over the last ten years, its use in radiotherapy has become widespread, with CT presently the standard imaging modality for 3D treatment planning. Patient anatomy is visualised in

transverse slices as a map of x-ray absorption variations in the different tissues (across the field of view). Fast, high resolution spiral (or helical) CT is the preferred method, where the patient is moved through a continuously rotating x-ray beam. Radiation from these multiple x-ray projections is measured as it exits the patient by detectors in the gantry, and the internal structure of the patient reconstructed from specific attenuation values assigned to each voxel; the Hounsfield (or CT) number. These values are displayed on a scale of Hounsfield Units (HU), where each number represents some shade of gray. The maximum is depicted white and the minimum shown as black, with the numbers normalised to the attenuation of water (*Figure 3.1*).

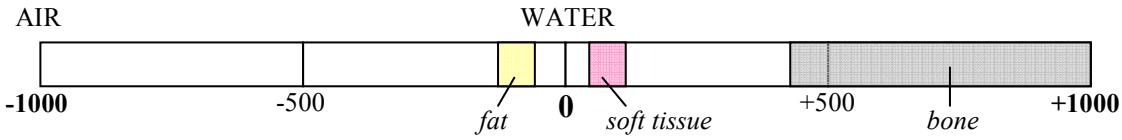

Figure 3.1: Principle of CT imaging: the Hounsfield scale of CT numbers for mapping anatomy as a function of x-ray attenuation.

There are two distinct advantages of using CT for treatment planning. Firstly, the patient can be CT scanned in the treatment position and patient contours easily extracted from resulting data. This is due to the intrinsic geometric stability and accuracy of CT images and because of this robustness, other images can be effectively registered to them. Secondly, correlation of Hounsfield numbers to electron density information can be made with look-up tables and these densities can be used to more accurately calculate absorbed dose in the patient across the region of interest. This is important since a given therapeutic beam may traverse any combination of fat, bone, muscle and air, with such inhomogeneities affecting both absorption of the primary beam and secondary electron fluence. Furthermore, CT usually best evaluates bony involvement of tumours, particularly in cases where the skull base is involved and haemorrhage and calcification are also best defined on CT. Unfortunately however, even with modern

helical CT scanners providing good image quality at high resolutions, in general, the soft tissue contrast is still poor for accurately assessing tumour extension in the brain (*Figure 3.2*). Exogenous contrast agents are available to enhance CT image contrast, but the use of them for planning scans is limited by induced temporal changes in electron density, which may affect the dose calculation. Additionally, high density objects such as joint replacements and dental amalgam can produce image artefacts that significantly obscure the required anatomy. This can have considerable consequences on the dose calculation, where air cavities are distorted and artificial white streaking may be interpreted as bone. Appropriate use of window levelling can enable CT images to be extensively interpreted by radiology experts. However, such expertise is not normally available within an oncology department. Finally, only physical information can be gleaned, i.e. no functional measures are possible, although a form of dynamic contrast enhancement is possible.

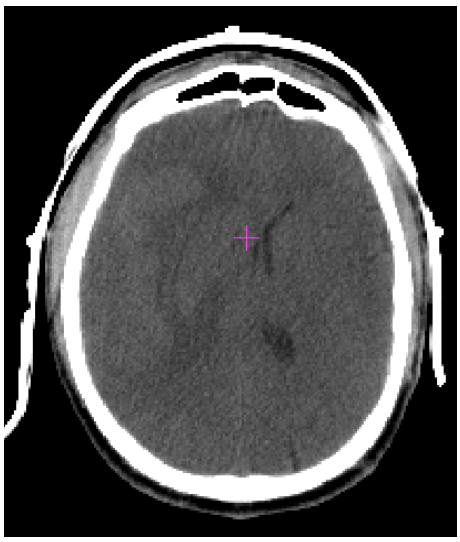

Figure 3.2: Axial CT image of the brain for an astrocytoma patient (scanned in immobilisation mask), showing bone, air and (limited) soft tissue contrast.

3.3. Magnetic Resonance Imaging (MRI)

Availability of magnetic resonance imaging (MRI) in the clinical setting came slightly later than CT. By the mid to late 1990s MR images were becoming available to oncology departments for radiotherapy treatment planning. MRI does not use ionising

radiation as CT does, rather it is derived from paramagnetic properties of the nucleus; nuclei possess angular momentum (spin) and thus a magnetic moment. MR images are formed by measuring radiofrequency (RF) signals created by nuclear spin state transitions (of protons) in the presence of a strong magnetic field. Transitions are induced by applying some pulse sequence to the patient, which is a combination of RF pulses and magnetic gradients. Resulting signals depend on the proton density and relaxation times, which in turn are determined by the environment the protons are in. MR acquisition sequences are designed to modulate the influence of these relaxation processes for tailoring of image contrast to a particular clinical need. MR images can be acquired in multiple planes (2D) or as a volumetric dataset (3D) and in any imaging plane directly.

In quantum physics, total angular momentum can only take certain (discretised) values, and when nuclei are placed in an external magnetic field, the number of spin states is given by the magnetic quantum number, m_z. m_z can take any of the (2I+1) values in the series (–I,-I+1….I-1, I), where I is the spin quantum number, which may be zero, half integral or integral. The hydrogen nucleus is a single proton with I = ½ and so for protons, two spin states are created; m_z = -½ and m_z = +½. These are a higher energy level where the magnetic moments are opposing the external magnetic field and a lower energy level in which the nuclei are aligned with the magnetic field. Ordinarily the spin states are randomly oriented, but with the application of an external magnetic field, a net magnetisation is created in the direction of the main magnetic field. The spins are distributed according to the Boltzmann distribution and from this, the population difference and hence the net magnetisation, M_0 may be calculated. This image signal can be enhanced by increasing the main field strength, hence a strong magnetic field affords better sensitivity.

In terms of classical physics, each magnetic moment precesses about the applied magnetic field, B_0, in a conical motion. This induces an angular velocity, the frequency of precession, ω_0, given by the Larmor equation,

$$\omega_0 = \gamma B_0 \tag{3.1}$$

This is the resonant frequency i.e. the frequency at which spins may be stimulated to induce transitions. The constant of proportionality, γ, is the magnetogyric ratio and is nucleus-specific. The hydrogen nucleus has the highest γ of any element (42.576 $MHzT^{-1}$) and so as well as having 99.98 % abundance in the human body, fortuitously also provides the best MR signal.

In order to detect M_0, a second magnetic field, B_1 is introduced at the resonant frequency. This is achieved using RF coils and serves to tip the spins out of alignment with B_0 i.e. out of the longitudinal plane. A 90˚ RF pulse for example, flips the spins into the transverse plane and equalises spin populations and an 180˚ pulse corresponds to a full population inversion. Once the RF pulse is turned off, the signal, called the free induction decay (FID), can be detected by some receiver coil. At this point, the signal peaks and quickly decays due to relaxation processes, whereby the spins return to the Boltzmann distribution. Relaxation is divided into simultaneous transverse and longitudinal processes; T_2 decay and T_1 recovery respectively. T_2 (spin-spin) decay is the process where all the spins begin to precess at slightly different rates and a loss of phase coherence occurs, without energy loss. Inhomogeneities in B_0 speed up the dephasing and so T_2 decay is sometimes characterised to include this factor by means of a second parameter, T_2*. Conversely, T_1 (spin-lattice) recovery is the process whereby spins return to equilibrium under the influence of B_0, incurring an exchange of energy between spins. Changes in magnetization during excitation and relaxation are described

by the Bloch equations, a set of differential equations yielding each component of magnetization as a function of time [45].

In order to construct an image, it is necessary to spatially encode the acquired signal and localise the protons in 3D space. This is done via the application of magnetic field gradients, produced by gradient coils, in three orthogonal planes. Gradients are simply linear changes in space made to the magnetic field which make the resonant frequency of the spins spatially dependent. This predictable range of Larmor frequencies across the patient can then be used to determine pixel intensities in all three-dimensions via slice selection, frequency encoding and phase encoding. Slice selection involves applying a gradient along the axis of the desired imaging plane where slice thickness and position may be altered with different gradient strengths or RF bandwidths. Frequency and phase encoding are used to encode the remaining two in-plane directions respectively for localisation within the image slice. Imaging acquisition time is determined by the repetition time, TR (time between successive pulse sequences), and the number of phase encoding steps used. Phase and frequency information is stored in a raw data matrix called k-space which is reconstructed to become the final image by 2D or 3D Fourier Transform (FT).

Image contrast is determined by the distribution of protons and sequence timing and whilst different MR sequences are designed to be application-specific, most use some variation of the spin echo (SE) or gradient echo (GE). The echo time, TE, in both these cases is the time between the start of the first RF pulse and production of maximum signal (the echo). Repetition time (TR) is the amount of time between successive pulse sequences applied to the same slice. Free fluids e.g. CSF, appear dark on a T_1-weighted image (short TE and short TR) and bright on a T2-weighted image (long TE and long

TR). In gradient sequences, changes may also be made to the angle of excitation - flip angle - to further manipulate contrast.

In order to maximally enhance MR image contrast, an intravenous injection of some paramagnetic agent, typically based on the gadolinium atom can be administered to the patient, which shortens the T_1 and T_2 processes. The T_1 shortening is most used clinically, where a typical imaging protocol acquires static T_1-weighted images before and after the contrast injection. Small increases of gadolinium generate a huge increase in signal in areas where the contrast agent can 'leak out'. The protective blood-brain barrier in the brain is impermeable to these agents and so uptake is seen only in regions where the barrier is damaged. Tumours, due to their increased vascularity, usually exhibit rapid preferential uptake, which can be differentiated from surrounding normal tissue.

MRI is the preferred imaging modality in diagnostic radiology for most intracranial disease (*Figure 3.3*). Tumours appear hypointense (dark) on T_1-weighted and hyperintense (bright) on T_2-weighted images. Contrast-enhancement (post injection) can suggest malignancy. Although its soft tissue determination is unparalleled, spatial accuracy of MRI data does not match that of CT. Geometric accuracy is a function of the linearity of magnetic field gradients and magnetic susceptibility of human tissue and there is no way of deriving electron density information from MR images for dose calculation purposes. To resolve these issues, MR data can be easily registered to CT. CT-MR image registration is achieved by superimposing the (secondary) MR dataset on the (primary) CT dataset via a series of translations, rotations and scaling. Target volumes can then be drawn on the more detailed MR images and superimposed on the CT images for planning. Coincidentally, distortions are minimal in head imaging (small

field of view) since homogeneity is maintained at the centre of the magnet, where the brain is positioned. Moreover, the brain is relatively homogeneous with regard to electron density. In fact others have shown it is feasible to plan radiotherapy with MRI alone [5, 46].

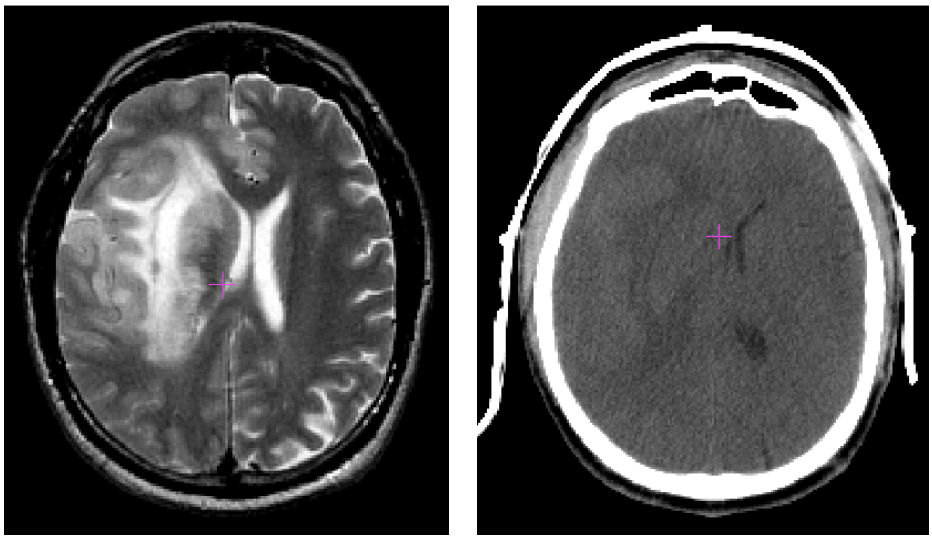

a. *b.*

Figure 3.3: Comparing MR and CT. a. axial MR image of the brain for astrocytoma patient, showing improved soft tissue contrast when compared to b. CT image of the same patient.

3.4. Functional Imaging

State-of-the-art functional imaging provides physiological information not discernable on either CT or conventional MR images. Specialist MR and radioisotope imaging techniques can both be used to this end, where image contrast is derived from different biological processes in the patient. Functional imaging may be useful in radiotherapy for identifying indices of disease control e.g. cell density, proliferation. Different functional imaging modalities may be characterised in terms of their sensitivity and specificity. A very sensitive test can be trusted when no anomaly is apparent i.e. good resolution, whilst an abnormality seen on a very specific test can similarly be relied upon to represent disease. In this way, a given functional imaging technique having perfect sensitivity and specificity wholly and reliably depicts any disease that it is designed to detect. In reality, each has its limitations, and functional imaging, at its

dawn of implementation in radiotherapy, is cautiously employed in conjunction with rather than as a replacement for, standard CT/MR images.

3.4.1. Advanced Magnetic Resonance (aMR) Techniques

Magnetic Resonance Spectroscopic Imaging (MRSI)

Magnetic resonance spectroscopy (MRS) is a technique for acquiring metabolic information non-invasively on a patient. Multi-voxel MRS can be obtained across a 3D matrix of voxels, which is sometimes called chemical shift imaging (CSI) or more frequently magnetic resonance spectroscopic imaging (MRSI). The resonant frequencies of nuclei depend on their chemical environment and MRSI works on the principle that these differences can be measured. The shift occurs because protons in different environments experience shielding of the magnetic field by the electron clouds of neighbouring atoms. The spectra of any nucleus can be obtained if the RF coil is tuned to the specific resonant frequency. In proton MRSI (^{1}H-MRSI), currently the most useful, fat and water peaks are suppressed in order to facilitate inspection of far lower concentration metabolites. This is achieved by the use of frequency dependent pulses before acquisition to saturate the required signal. The chemicals to be inspected have concentrations approximately eight thousand times less than water and so these would be scaled out if water were included in the examination. Phosphorus (^{31}P) spectroscopy can provide information about tumour energy status and is used mainly for examining muscle metabolism. All further references to MRSI in this thesis are to ^{1}H-MRSI.

Chemical shifts are expressed as frequency differences from some reference compound in order to identify the chemical species. Spectra are produced depicting amplitude vs. frequency on a ‘parts per million’ (ppm) scale, where the frequency separation of peaks

depends on the field strength and the homogeneity of the magnetic field. The reference compound most commonly used in MRS is tetramethylsilane (TMS), assigned a value of zero ppm. Separation of the peaks in absolute frequency units increases directly with field strength, therefore MRS at 3.0 T should be better resolved than at 1.5 T, for example. The position of the peaks within the spectrum provides information on what kind of proton gave rise to that signal. Peak heights depend on the spectroscopy sequence and the TE and TR used and these are chosen to maximise the signal-to-noise ratio (SNR) of the spectrum peaks. Either a STEAM (stimulated echo acquisition mode) or PRESS (point-resolved spectroscopy) sequence is used for acquisition.

Multiple voxel MRSI uses phase encoding techniques to encode information in-plane so that the location of all voxels used for the 3D MRS acquisition is known i.e. where each signal came from. This also enables spectra from multiple voxels to be acquired in one pass. This information can be precisely overlaid on the corresponding anatomical images for spatial correlation, forming the so called 'metabolite map'.

The value of MRSI in discriminating between solid tumour, necrosis and normal tissue is well established [47-49]. Molecules that can be studied with MRSI include choline (Cho), N-acetyl aspartate (NAA), creatine (Cr), citrate and lactate. For brain spectra, the metabolites of interest are NAA, indicating neuronal function and Cho, representing cell density. A typical tumour spectrum shows an increased Cho peak and a decreased NAA peak (*Figure 3.4*). Raised choline typically indicates malignancy (high grade) and NAA is lowered when normal brain cells have been destroyed.

Absolute quantification of metabolite concentration is difficult and so an increase in the Cho/NAA ratio is used to index malignancy. Peak ratios are affected by the TE used for

acquisition, whereby short TEs are useful for attaining more peaks, since each metabolite has a different T_1 and T_2. A statistical measure of spectra abnormality can be determined by calculating the number of standard deviations of difference between Cho and NAA in a given voxel and control or 'normal', healthy voxels. This value is termed the 'abnormality index', and is now a standard approach [21]. In Hull, the threshold of abnormality is defined at an NAA/Cho ratio of 0.6, derived in-house from a cohort of healthy volunteers [50]. Voxels in the metabolite map, overlaid on anatomical images, are coloured to represent values on the abnormality scale, where darker voxels represent a higher deviation from the normal population.

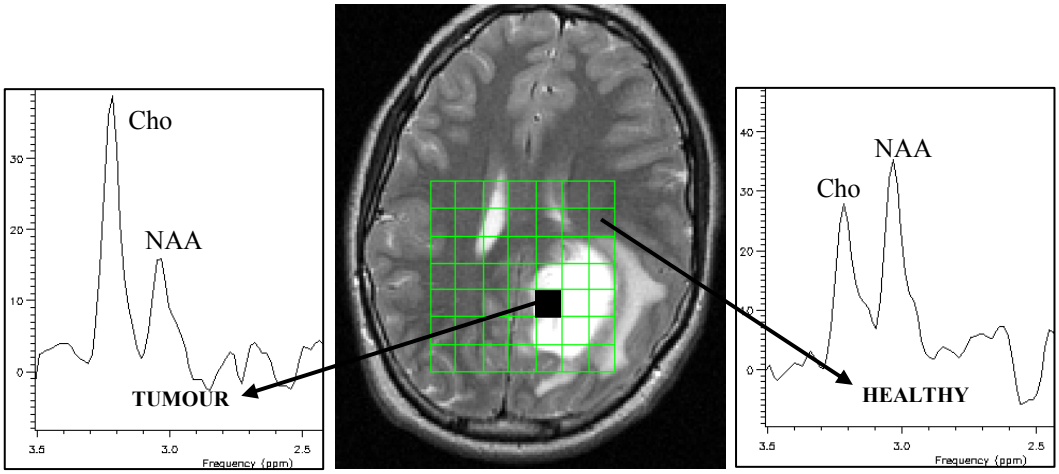

Figure 3.4: Voxels with abnormal spectra (defined as NAA/Cho ratio<0.6) are depicted as abnormal (shaded) to produce metabolite map. This is overlaid on the T_2-weighted anatomic images.

The potential of MRSI in radiotherapy is already recognised in the literature [21, 51-53] and the process for acquiring MRSI data in the brain is well-defined in Hull, where a 12 x 12 x 8 cm^3 volume of spectra can be acquired in approximately 20 minutes with a typical voxel size of 1 x 1 x 1 cm^3.

Other aMR Techniques

There are many eloquent cortices in the brain, each responsible for specific functions, which may lie close to a tumour. If damaged, cognitive impairment can have detrimental effects on a patient's quality of life. These functional regions are found in

the same relative areas for most people, but inter-person variations can be significant, especially in pathological brains, and their exact positions cannot be obtained from morphological imaging alone. Functional MRI (fMRI) utilises the blood oxygen level dependent contrast (BOLD) effect to monitor the brain's haemodynamic response to stimulation and map brain activation in the relevant area. Speech, visual and many sensory and motor centres can be localised in this way, where the BOLD effect causes a 2-3 % increase in MR signal.

In fMRI examinations, the patient alternates between directed rest and task known as the paradigm (such as a finger movement sequence for motor cortex localisation), while MR images are acquired. Deoxyhaemaglobin is paramagnetic, meaning that deoxygenated blood has a shorter T_2* and a lower MR signal therefore blood oxygenation arising from neuron activation can be detected as a T_2* signal increase. The signal trace from each voxel can be analysed and correlated to the stimulus pattern given by the paradigm to create a parameter map highlighting the area of the brain that responded to the task. The fMRI parameter map can be registered to the anatomic dataset, as shown in *Figure 3.5*.

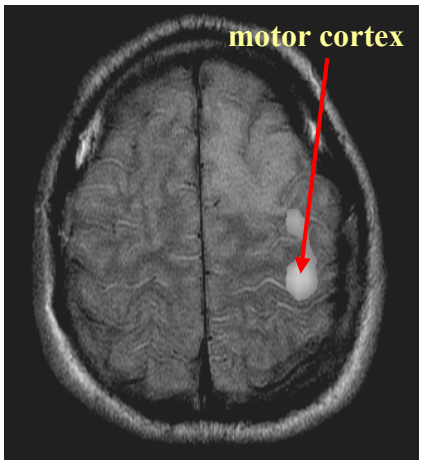

Figure 3.5: BOLD activation in an fMRI examination, for localisation of the motor cortex.

The feasibility of reducing irradiation of a functioning cortex by using fMRI in an IMRT plan has been proven, whereby the primary motor cortex for a patient was identified with fMRI and successfully defined as an OAR in the planning system [16].

Diffusion is the random motion of molecules along a concentration gradient and diffusion-weighted MRI utilises the signal changes produced by the motion of spins. A special gradient sequence is used so that the signal attenuation depends on the degree of the diffusion. The range of proton motion due to diffusion processes is in the range of microns to millimetres therefore acquisition techniques have to be extremely rapid to overcome physiological motion as well as patient motion. Echo-planar imaging (EPI) is usually used, which can acquire images every 50-100 ms. This consists of a 90° and 180° pulse with two large gradient pulses distributed symmetrically about the 180°. If the protons are motionless, the gradient pulses cancel one another out, but for protons moving, a loss of signal intensity occurs, proportional to the rate of diffusion. The gradient and timing is described by the b-value, where higher b-values result in greater sensitivity to diffusion (the relationship is exponential). The raw diffusion images produced by the scanner show unrestricted diffusion as dark due to exponential signal loss and areas where protons are restricted from diffusion are bright where cells swell. T_2 shine-through can pose a problem in DWI due to the required long TE in order to apply large gradient pulses. This can be overcome by the use of apparent diffusion coefficient (ADC) maps calculated from more than one b-value, where the DW image is combined with an image without gradients. ADC depends on the structure of tissues and these maps are opposite in contrast to the raw DW images. DWI is particularly useful in assessing and diagnosing acute stroke and also serves as a means to differentiate tissue characteristics such as oedema and necrosis. These standard sequences however, measure diffusion in one direction only.

Diffusion-tensor imaging (DTI) is a more sophisticated form of DWI that is sensitive to anisotropic diffusion i.e. diffusion in multiple directions, and can therefore demonstrate white matter tract anatomy or tractography. This can be useful in looking at disruption and infiltration of tumour for delineation purposes, as well as having potential for defining tracts as OARs in the planning process and predicting patterns of glioma recurrence [19, 54-57]. Other aMR techniques include perfusion-weighted imaging (PWI) and R_2* mapping for identifying hypoxic tumour regions and molecular imaging for evaluating tumour biology at a cellular level.

All MR and aMR data presented in this thesis was acquired at the Centre for MR Investigations (CMRI), The University of Hull.

3.4.2. Radioisotope Imaging

Radioisotope imaging is undertaken within a nuclear medicine department, where a radioactive labelled tracer is injected into the patient. These tracer compounds are designed to be selectively taken up in tissues of interest and the patient can then be imaged with a detector sensitive to the emitted radiation. In positron emission tomography (PET), radioisotopes that decay via positron emission are chosen, where these positrons annihilate on contact with electrons in the patient. Each annihilation produces two gamma rays of equal energy (511 keV) travelling in opposite directions. Detection of these two coincident photons at 180° to one another forms the basis of PET imaging. The final image is constructed using standard tomographic techniques showing tracer distribution through the body. Single photon emission tomography (SPECT) operates on a similar principle, but with the use of radioisotopes that emit single gamma rays per nuclear disintegration.

^{18}F-fluorodeoxyglucose (FDG) is the most common tracer used in PET imaging of the brain, which is preferentially taken up by cells with raised metabolism. Tumour cells exhibit increased metabolism due to their abnormal rate of growth and therefore demonstrate increased uptake. FDG-PET can be used in this way to identify tumour regions of rapid cell proliferation. Different tracers can be employed to detect various biological attributes e.g. hypoxia. Radioisotope imaging is extremely sensitive, since it can detect even single atomic decays, although specificity is dictated by the biological uptake of the individual tracer. Unfortunately, these techniques provide very limited information on normal anatomy and typically have long acquisition times. Like most other functional techniques, resolution is generally coarse; typically 8-10 mm for SPECT and 5 mm for PET, although dedicated brain PET scanners can afford resolutions as high as 2-3 mm [58, 59]. Currently, radioisotope images are mainly used in diagnosis and staging rather than planning, although combined PET-CT scanners are now commercially available [60].

Chapter 4 - IPEX: A Plan Evaluation Environment for IMRT

4.1. Introduction

The ideal scenario in conformal therapy of any kind shapes the prescribed isodose precisely around the target and deposits the dose uniformly throughout the volume. In reality however, limitations to achieving this are imposed by irregularity of target volumes, close proximity of dose-limiting structures and demands on the delivery device. Plans must be accurately evaluated then against some acceptable criteria and competing plans ranked to this same end.

Traditionally, treatment plan evaluation, including the ranking of competing plans, is realised dosimetrically with the visual inspection of 2D isodose distributions, dose-volume statistics in the form of the dose volume histogram (DVH) and the reporting of minimum, maximum and mean dose levels where appropriate. Biological evaluation in terms of clinical outcome predictors such as tissue control probability (TCP) and normal tissue complication probability (NTCP) can add value to a dosimetric analysis although no complete and exact model exists for calculating these. Typically, current practice accepts a plan, if isodose lines and DVH information indicate good target coverage and normal tissue dose criteria are met. However, the physical dose distributions in IMRT plans vary in complexity, conformity and homogeneity compared with conventional treatments, and with this in mind, conventional evaluation tools alone are becoming inadequate in assessing the quality of the vast amount of 3D dose information inherent in these advanced plans. Additionally, IMRT treatments are designed using inverse

planning techniques, which use optimisation algorithms to satisfy user-specified objectives for target and OAR dosimetry. The number of competing inverse plans may be significantly higher than the options afforded by conventional planning, where several unique plans may all be deemed satisfactory under these terms. The task of comprehensive plan evaluation is further complicated with the introduction of the selective multiple boosting (SMB) technique, as the dosimetry of multiple targets must be simultaneously assessed. IMRT plan evaluation therefore, requires a great deal more vigilance than the interpretation of conventional plans, and the development of novel, effective quality factors, for interpreting plans, as well as the comparison or ranking of competing plans is essential in order to maximise and assess the efficacy of the SMB technique.

4.2. Treatment Plan Characterisation

Many different properties of a treatment plan may be represented by specific analysis tools, encompassing dose and biological attributes in both visual and quantitative form. Dose tools are important to verify acceptable target coverage and OAR sparing whilst biological tools aim to predict likely clinical outcomes.

4.2.1. The DVH

Dose statistics and the dose-volume histogram (DVH) invariably form part of any 3D dose distribution evaluation. The dose calculation engine in the planning system returns a 3D array of dose points that cover the patient's anatomy and from this dose matrix, the dose to each voxel within a given structure can be obtained, from which key dose statistics may be computed. These may include minimum, maximum and mean doses to a given structure volume. In order to extract dose-volume statistics however, e.g. volume irradiated to at least 95 % of the prescribed dose, a DVH must be employed.

The differential DVH (dDVH) is a simple representation of the frequency of dose values within a particular structure. The percent of the total structure volume is plotted on the ordinate as a function of dose on the abscissa. An optimum target dDVH would be an infinitely narrow peak at the prescription dose, showing that 100 % of the target is receiving the prescribed dose, with perfect homogeneity. It follows that the optimum OAR dDVH is an infinitely narrow peak at 0 Gy, indicating 100 % of the OAR is receiving no dose at all (*Figure 4.1a*). Of course in reality, this cannot be realised, and physical dDVHs look something like the examples in *Figure 4.1b*, the target having a broad peak around the prescription dose and the OAR displaying several peaks, since different parts of the organ receives different doses.

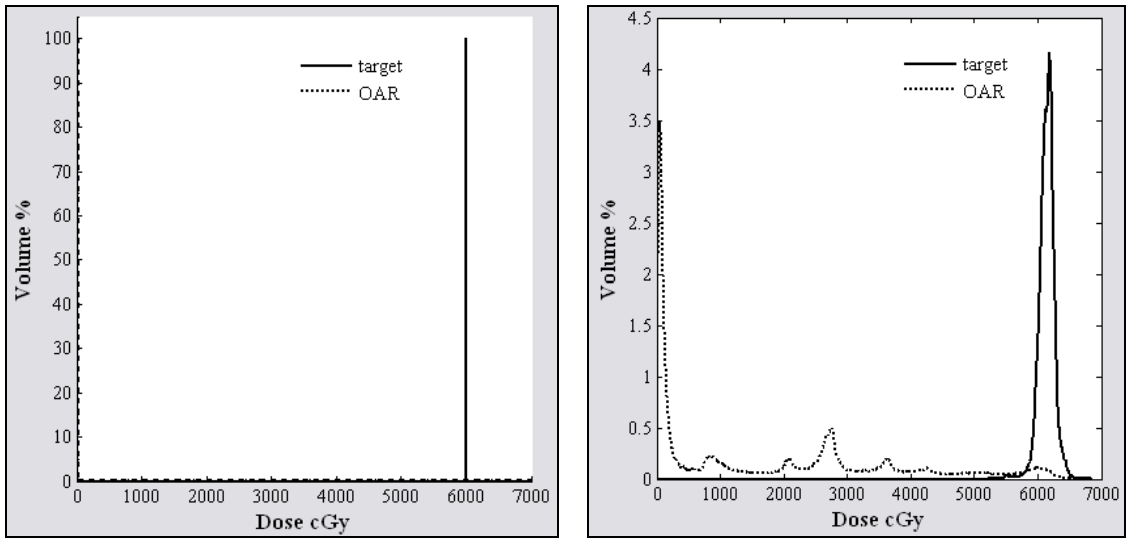

a. *b.*

Figure 4.1: The differential DVH (dDVH). a. Infinitely narrow peaks of the ideal dDVH indicating complete, uniform target coverage and total OAR sparing. b. A brain IMRT plan for the target (PTV), the broader peak representing realistic heterogeneity around the prescribed dose of 60 Gy and the OAR (normal brain tissue), where different parts of the organ receives different doses.

The integral or cumulative DVH (referred to simply as DVH) is most commonly used, being a histogram representing the fractional volume of a structure receiving at least a given dose. This is the accepted form, since it readily resolves dose-volume criteria such as the percentage volume of the target covered by a given isodose. DVH plots

start with 100 % of the volume at 0 Gy, indicating that all of the structure receives at least zero dose (*Figure 4.2*). The DVH also serves to represent the dose prescription in plan optimisation.

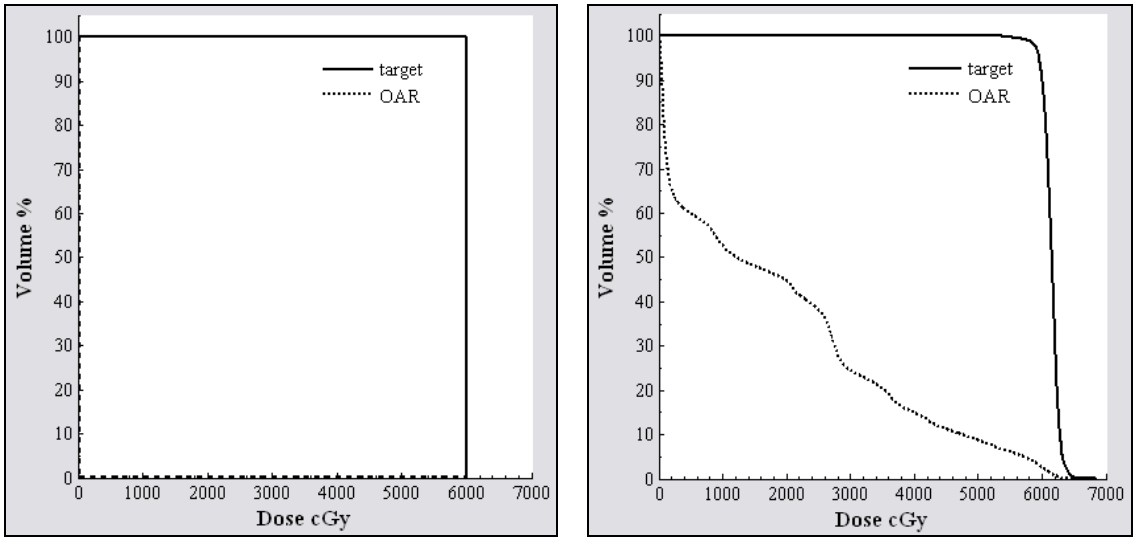

a. *b.*

Figure 4.2: The cumulative DVH. a. ideal target and OAR DVH showing complete, uniform target coverage and total OAR sparing. b. DVH of the target (PTV) and OAR (normal brain tissue) in a brain IMRT treatment (as Figure 4.1b).

4.2.2. Visualising Dose Distributions

IMRT can create hot and/or cold spots in unexpected locations, which are not identifiable on DVHs. It follows then, that inspection of isodose distributions over the entire image volume is required for evaluation of the spatial dose distribution. 2D isodose lines can be superimposed on the anatomical images within the planning system over all three planes in the patient. It is usual to review at least the 95 %, 100 % and 107 % of the prescription dose levels, in line with ICRU recommendations, although a maximum of fifteen isodose levels can be displayed simultaneously on the current planning system (CMS[§] XiO® 4.2) [27]. Isodose lines can also be generated on orthogonal planes reconstructed from the original axial data for isodose inspection in the sagittal and coronal planes.

[§] St. Louis, MO.

Isodoses may also be mapped in three dimensions to construct a 3D isodose 'cloud' for a given level of dose. This isosurface can then be displayed on a 3D surface rendering of the target and/or other structure contours (*Figure 4.3*). The structure-isosurface volumes can be visualised freely in all directions, which is useful for assessing target coverage and provides good spatial information with regard to the location of hot/cold spots. No quantitative information can be deduced from this however, such as the distance between the isosurface and the anatomical structure i.e. conformity.

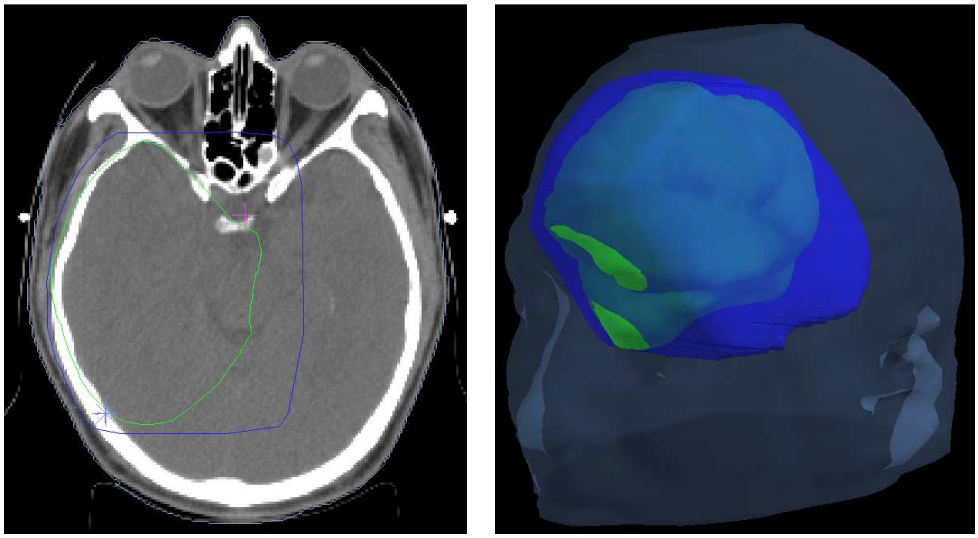

a. *b.*

Figure 4.3: Qualitative dose distribution evaluation. a. 2D isodose lines showing the PTV (green) enclosed by the 95 % isodose (blue). b. 3D rendering of the same 95 % isodose (blue) and PTV (green), demonstrating incomplete target coverage.

4.2.3. Conformity

Dose and dose-volume statistics and visualisation are important and flexible plan evaluation tools but do not readily permit the simultaneous and quantitative inspection of dose both inside and outside the target volume. This can be characterised instead by the 'conformity' of a plan, which refers to the relative volumes of the target and surrounding normal tissues receiving the treatment dose or higher. Good PTV conformity for example, implies that the PTV is well enclosed by a therapeutic dose, whilst adjacent normal tissue is well spared. Wagner *et al.* showed that conformity correlates to clinical outcome and with one of the attractions of IMRT being its ability

to accurately sculpt or conform dose, it is evident that some quantitative measure of conformity may be useful [61]. Many indices have been proposed to this end; which has led to a variety of different formulae and yet none of which has become an accepted standard. Some methods are equivalent under certain conditions, but in general, different indices will generate different results for a given plan. Only a brief summary of their relative methods, merits and utility is appropriate here, for an in-depth review of conformity indices, the reader is referred to Feuvret *et al.* [62].

Analysis of the conformity of a given plan typically involves visual inspection of the target volume, V_t, with respect to some isodose volume, V_{iso}, the volume of (any) tissue receiving a specified dose. Qualitatively, this may include the use of 2D isodose lines or 3D dose visualisation and comparisons may be made in terms of the relative size, shape and/or location in space of the target contour with respect to the isodose volume. A therapeutic isodose volume approximating 150 % of the PTV has been suggested, although the size of the target itself may influence conformity, whereby larger targets exhibit better conformity [63-65]. Larger targets however, are typically associated with an increased toxicity risk, thereby opposing the importance of in-depth conformity analyses for smaller lesions.

The first measuring index for this type of plan interrogation was proposed by Nedzi *et al*, who define a treatment volume ratio (TVR), as the simple ratio of V_t to V_{iso}, the isodose level specified to achieve desired treatment outcome [66]. The ICRU later published a conformity index, CI, as the inverse of TVR, whereby a lower result represents better conformity [29]. Both CI and TVR however, imply that the isodose or treated volume totally encompass the target, since they do not utilise spatial information regarding the relative location of the two volumes. A totally unacceptable plan, where

the two volumes are offset for example, could still produce a perfect score as long as they were equal in size. CI and TVR are useful then only when the target volume is known (or assumed) to be covered by the isodose surface. This information can be easily read off the DVH, but ratios such as these, which index 'global' volumetric data only, should be employed with caution and understanding in order to avoid misinterpretation. Van't Riet *et al.* addressed this issue with the introduction of a conformation number, CN,

$$CN = \frac{V_{t \cap treat}}{V_t} \times \frac{V_{t \cap treat}}{V_{treat}} \tag{4.1}$$

where $V_{t \cap treat}$ is the volume of target enclosed by V_{treat} [67]. Here, the degree of target volume coverage and the volume of normal tissue receiving dose greater than or equal to the reference dose are combined into a single number, a higher value representing better conformity. The first component in the expression is identical to the volumetric quantity derived from the target DVH. Using this two-component approach, a geographic miss will fail to return a perfect score, although the same score may be derived from clinically different situations, of course; full coverage of the target volume with poor normal tissue avoidance, and some target underdosing with complete normal tissue avoidance, for example. For completeness therefore, it is sensible to quote some coverage index alongside CN. Lomax and Scheib define a coverage factor, TC, as the percentage ratio of $V_{t \cap treat}$ to V_t [65]. This is equivalent to both the target DVH and the first component in the formula for CN (4.1).

A novel approach to quantifying conformity could be to employ similarity coefficients such as those used in reproducibility studies e.g. dice (Sorenson's) or matching index, DSC and MI respectively [25, 68]. These quantify relative spatial/volume overlap in terms of set theory and a problem may occur in that structure-specific information is not evident and a target enclosed by an isodose and an isodose enclosed by a target could

produce the same score. Such an index could be further combined with some coverage coefficient in a product or mean for example in order to provide the relevant insight.

Given that the various indices are derived in different ways, any given clinical scenario may of course lead to differing scores. A simple 2D isodose-target model was developed in order to generate evidence data for comparison and characterisation of available formulae. The simulation was constructed in MATLAB[§] in order to typify a variety of indices as functions of both isodose-target volume separation and isodose-target volume ratios. This was achieved by generating two circles, one representing the target and the other representing the volume covered by some arbitrary isodose. These were moved along the x-axis with respect to one another (keeping the centre of the circles at the same point on the y-axis) in order to modulate target-isodose separation. Next, each volume was grown concentrically with respect to the other to investigate size variations. Results were obtained by summing relevant pixels to determine relative volumes (areas) and both the union and intersection of the model, from which a selection of conformity indices were computed.

Figure 4.4a shows a selection of indices as functions of V_{iso}-V_t volume separation. This demonstrates how the simple volume ratios TVR and CI give an optimal score even when the isodose volume does not cover the target, providing the volumes are of equal size. DSC and TC are essentially equivalent under this constraint, represented by a linear relationship with separation. CN and MI are both exponential functions of separation and are roughly equivalent in that MI is approximately equal to (CN-1). MI is the only increasing function with separation distance but all have a maximum score of unity when the volumes are equally sized.

[§] Mathworks, Inc.

Figure 4.4b illustrates the relationship between the indices and the V_{iso}-V_t volume ratio when the volumes are concentric. It can be seen that TC, CN, MI and DSC all reach a global maximum or minimum when V_t=V_{iso} and that TC stays at its maximum value of 100 thereafter. The volume ratios CI and TVR both have values much greater than unity when the relative volumes change, being increasing and decreasing functions respectively. CN is equivalent to TC when V_{iso}<V_t, exhibiting a linear relationship over this range and the DSC curve is the same shape as CN, both reaching a maximum value of unity when V_t=V_{iso} and decreasing thereafter with similar derivatives. CN is equivalent to TVR when the volume ratio is greater than unity i.e. when V_{iso}>V_t. As suggested initially by the separation experiments and again here, MI approximates (1-CN).

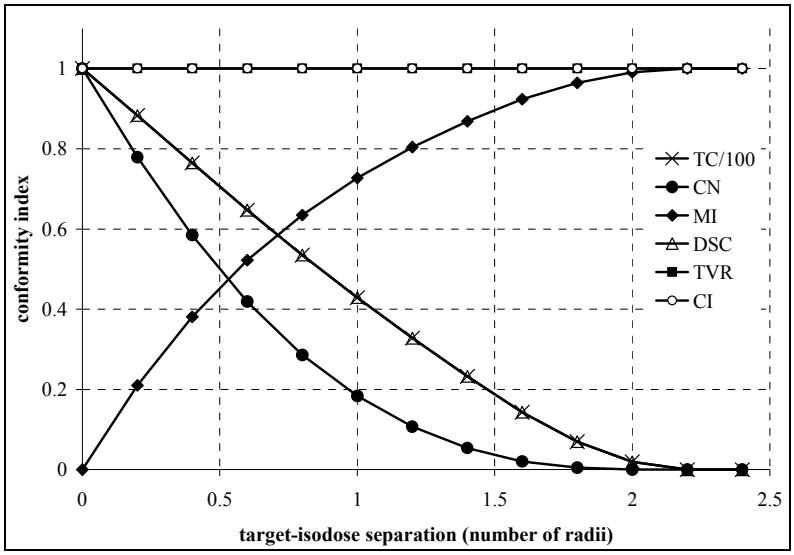

a.

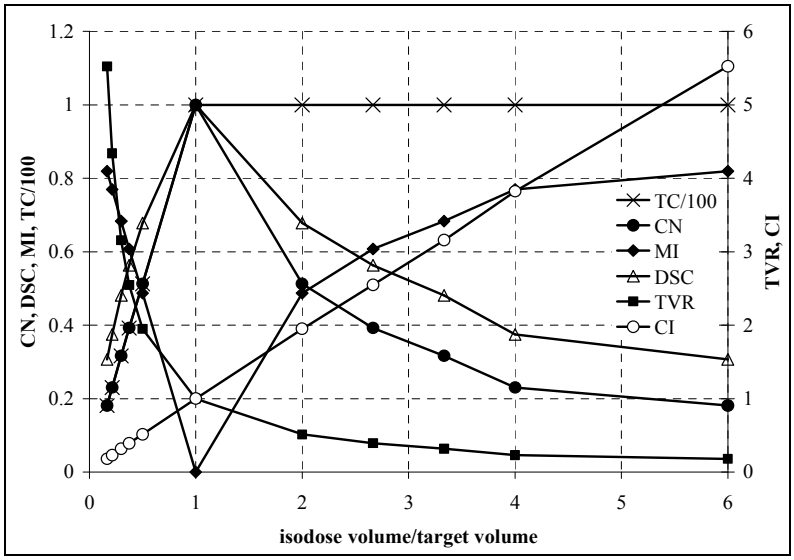

b.

Figure 4.4: 2D model to characterise different conformity indices. a. conformity as a function of target-isodose separation. b. conformity as a function of isodose to target volume ratio.

From the shape of the curves in *Figure 4.4*, it is possible to characterise the different metrics by defining three main categories; firstly the simple ratios TVR, CI are seen to be the inverse of one another, secondly, MI, DSC and CN functions are the same shape, where MI ~ (1-CN) and finally TC stands alone, equivalent to DSC when $V_t=V_{iso}$ but unique when the volumes are varied, being the only function to remain at its maximum level over an infinite range of $V_{iso}>V_T$.

The term ‘conformal’ therapy is itself associated with unresolved ambiguity as to the criteria that define a conformal treatment. Several attempts have been made in establishing a baseline for a conformal plan. Leung *et al.* describe a plan as conformal if and only if the ratio of the partial volume of target enclosed by the prescription isodose to the target volume is greater than or equal to 0.95 and its associated TVR is less than or equal to 2 [69]. Van’t Riet *et al.* reserve the term conformal therapy only for those plans having CN > 0.60 [67]. These tolerances then, should be aimed for, and ideally exceeded in any conformal plan, especially IMRT.

4.2.4. Homogeneity

Homogeneity refers to the degree of dose uniformity within a target. IMRT dose distributions are, by their intent, highly complex and generally much more heterogeneous than traditional conformal plans. Although Nakamura *et al.* did not find dose heterogeneity to be a significant risk factor for toxicity, a more homogeneous target dose is, in general, desirable [64]. The increased heterogeneity may be caused by steep dose gradients between target and adjacent OAR, concavity of the required dose distribution, minimal distances between target and critical structure or a limitation on beam directions. In external beam radiotherapy, maximum heterogeneity of -5 % to + 7 % of the prescribed dose is generally accepted [27]. The dDVH and DVH are currently the best tools for complete understanding of the dose uniformity in a target, although quantifying dose uniformity in a single index can also be useful in comparative research studies.

Dose uniformity may be characterised in terms of dose levels (maximum, minimum) or relative volumes receiving a given dose. Weber *et al.* use a normalised expression in an inhomogeneity coefficient, IC, constructed from dose levels,

$$IC = \frac{D_{max} - D_{min}}{D_{min}} \tag{4.2}$$

where D_{max} and D_{min} are the maximum and minimum doses in the plan respectively [70]. Any approach using absolute minimum or maximum doses in it formulism is potentially problematic in that the particular dose level quoted may exist in only a few or even one voxel/s only, making the method very sensitive to target delineation, dose calculation parameters and furthermore, concealing the size of any hot spot. However, the global minimum in a target is significant, since it takes only one clonogenic cell (smaller than one voxel) to survive for disease control to be potentially compromised. Any index must of course be interpreted with caution and on a case-specific basis regardless of how it is derived. If a target is adjacent to an OAR for example, target underdosage in a few voxels near the boundary may be accepted as a compromise in order to stay within the tolerance dose of the critical structure, but this scenario may give a low homogeneity score.

The Radiation Therapy Oncology Group (RTOG) proposed a homogeneity index as the ratio of the maximum target dose to the prescribed dose in order to measure target overdosing, where a value of less than 2 is deemed acceptable [71]. Wu *et al* .(2003) report a homogeneity analysis designed specifically for IMRT boost treatments [72]. Their homogeneity index, HI, is quoted for the boost volume (GTV) as

$$HI = \frac{D_{2\%} - D_{98\%}}{D_p} \times 100\% \tag{4.3}$$

where D_p is the prescribed dose, 2 % of the target receives dose $D_{2\%}$ or greater, 98 % of the target volume receives $D_{98\%}$ or higher and a lower HI represents a more uniform target dose. Most recently, Kinhikar *et al.* reported the homogeneity of a plan as the percentage of the PTV receiving dose higher than 95 % and lower than 107 % of the prescribed dose [73].

Until recently, homogeneity indices have almost exclusively been employed for brachytherapy or radiosurgery applications. As IMRT becomes more routine however, it may be desirable and indeed essential, to address and quantify the increased heterogeneity it incurs. Homogeneity indices are particularly useful if a specific hot or cold area is suspected.

4.2.5. EUD

The equivalent uniform dose, EUD, is a tool for evaluating heterogeneous dose distributions and is useful therefore, in analysing concomitantly boosted plans, which are intentionally non-uniform. The concept of EUD was first proposed by Niemierko in 1997 as the single uniform dose which, if applied across a given tumour, produces the same biological end-point as the given non-uniform distribution i.e. any two dose distributions are equivalent then if they cause the same radiobiological effect [74]. The formula is based on probabilities of cancerous cells surviving irradiation and takes into account all factors that could possible affect the tumour's response including proliferation rates, repair rates and clonogenic cell densities. The original formula is complex, and many of the biological parameters cannot be accurately derived. The generalised EUD (gEUD), based on the power-law response of biological systems, was proposed by the same author in 1999 as a simple approximation to the full mechanistic EUD [75]. The gEUD can be calculated (in Gy) from the dose calculation points within a tumour, assuming these are distributed evenly within the target volume, using

$$\mathrm{gEUD} = \left(\frac{1}{N} \sum_{i=1}^{N} D_i^a \right)^{\frac{1}{a}} \tag{4.4}$$

where N is the number of voxels in the structure of interest, D_i is the dose in the i^{th} voxel and a is a tissue specific parameter describing the dose-volume effect, obtained

from clinical outcome data [76, 77]. gEUD is simply the generalised mean value of the non-uniform dose distribution and therefore is always some value between the minimum and the maximum dose of a given dose distribution. Li *et al.* showed that the gEUD can be derived from EUD and Wu *et al.* (2002) demonstrated that gEUD is a good fit to phenomenological data [78, 79]. The gEUD then, is a good approximation to the EUD and is generally more useful due to its relative simplicity and the fact that only one tissue-dependent factor is required.

EUD can be applied to both tumour and normal tissues. An optimal plan has a maximised EUD in the target/s and a minimised EUD in any critical structures. For normal tissues, the gEUD is usually calculated in terms of n, a seriality parameter, which describes how sensitive a structure is to hot spots, where n=1/a. It can be deduced from equation (4.4) that where a=1, gEUD is the arithmetic mean of the dose distribution, a=∞ and a=-∞ return the maximum and minimum dose respectively. It follows that in general, the value of a for aggressive tumours is large and negative, so that EUD is mainly determined by the lowest dose values and for OARs, a is positive so that values between the mean and highest dose values have the biggest impact on the EUD result. Although based on radiobiological modelling, gEUD measures the equivalency of a physical dose distribution and as such does not actually predict clinical outcome.

4.2.6. TCP & NTCP

3D dose can be converted to biological outcomes using tumour control probability (TCP) and normal-tissue-complication probability (NTCP) models [80, 81]. These are measures of treatment success and side effects respectively and both require either

experimental data or some kind of theory to model the response of both healthy and pathological tissue to a known radiation dose.

The concept of TCP uses the linear-quadratic (LQ) model of cell survival. A one-to-one relationship exists between TCP and EUD, since they are both derived from the LQ model; the TCP for a non-uniform dose distribution is equal to the TCP at the EUD on the dose-response curve. Different versions of the TCP formula can be derived in order to incorporate a range of effects such as non-uniform dose distributions, tumour proliferation, non-uniform clonogenic cell density and inter-patient heterogeneity of radiosensitivity. In its simplest form, i.e. for uniform clonogenic cell density and a uniform dose, for a single patient, TCP is given by

$$TCP = e^{-N_s} \tag{4.5}$$

where N_s is the number of surviving clonogenic cells calculated from

$$N_s = N_0 e^{-\alpha D} \tag{4.6}$$

with α, the linear radiosensitivity term from the LQ model for the given patient, D the uniform dose to the tumour and N_0, the starting number of clonogenic cells, found from

$$N_0 = \rho V_t \tag{4.7}$$

where ρ is the uniform density of clonogens in the tumour and V_t is the tumour volume (cm^3). It is clear then, that a TCP of 100 % is achieved when the number of surviving clonogens is zero. This form assumes not only uniform cell density, uniform dose and exact knowledge of the radiosensitivity of the patient, but also uniform radiosensitivity across the tumour i.e. no hypoxic cells. Also, this version does not account for any temporal effects e.g. proliferation.

For a fixed value of α, the dose response curve is too steep when compared to data observed clinically. Incorporating a distribution of α to represent radiosensitivity

variability across a patient population leads to a less steep curve which better matches clinical data. For a Gaussian distribution of α and a uniform dose distribution, TCP is given by

$$TCP = \frac{1}{K}\sum_{j=1}^{K} e^{-N_s} \tag{4.8}$$

$$N_s = N_0 e^{-\alpha_j D} \tag{4.9}$$

where K is the number of patients representing the population and α_j is the radiosensitivity of the j^{th} patient. This means that any final TCP result is the average value of patients with different radiosensitivities. For a non-uniform dose distribution, (4.7) and (4.9) become

$$N_{0_i} = \frac{\rho V_t}{n} \tag{4.10}$$

$$N_s = \sum_{i=1}^{n} N_{0_i} e^{-\alpha_j D_i} \tag{4.11}$$

respectively, where n is the number of tumour voxels, D_i is the dose to voxel i (uniform across voxel) and N_{0_i} is the starting number of clonogenic cells in voxel i. Since the biological sensitivity of tumours to radiation cannot be explicitly defined and TCP is particularly sensitive to its parameters, TCP is most useful in predicting relative rather than absolute clinical response i.e. comparing planning strategies or in the ranking of similar plans [82].

NTCP is defined by a sigmoid relationship between a uniform radiation dose given to a specified volume of an organ and the probability of a complication. The model uses the 'effective volume' method, which transforms a DVH into one of uniform dose delivered to an effective, fractional volume. For a homogeneously irradiated organ, NTCP is given by

$$NTCP = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{t} \exp\left(\frac{-t^2}{2}\right) dt \tag{4.12}$$

$$t = \frac{D - TD_{50}}{mTD_{50}} \tag{4.13}$$

where D is the uniform, total dose delivered to the volume under consideration, m is a fitting parameter that determines the steepness of the NTCP curve and the parameter TD_{50} represents the dose level (in Gy) leading to 50 % NTCP i.e. the dose at which 50 % of the population will exhibit complications [76, 77].

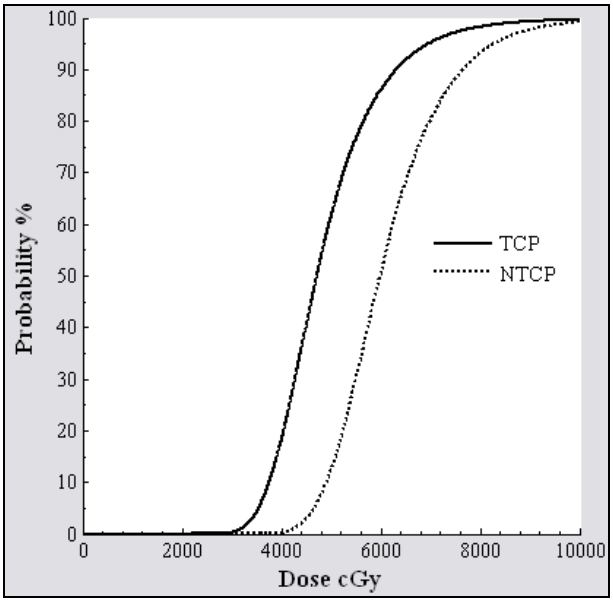

Figure 4.5: TCP and NTCP are sigmoidal curves that predict clinical outcome. These form the principle of the therapeutic ratio, defined as the ratio between tumour lethal dose and normal tissue tolerance at a given level of response. The further the NTCP curve is to the right of the TCP curve, the larger and more desirable the therapeutic ratio.

An optimum dose distribution for a given tumour is that which maximises the TCP whilst minimising the NTCP (*Figure 4.5*). A good radiotherapy treatment should aim for a TCP greater than or equal to 50 % and a corresponding NTCP of less than 5 %. In fact, a prescription dose may be chosen in part from the TCP curve by reading off the dose which produces a probability of control of 50 %. Unfortunately, modelling tissue response to therapeutic radiation is difficult, not only because it's difficult to measure *in vivo* but also due to the variability of response between patients. Because of this,

practical plan evaluation is usually based on physical measures, with a biological aspect used as a secondary tool. However, as functional imaging develops and becomes integral in the planning process, TCP and NTCP will no doubt have a more prominent role in assessing these biologically heterogeneous plans.

4.3. IPEX Software

Current plan evaluation can be rather cumbersome when several tools are required within the planning system. The aim of developing software in-house was to reduce the multi-process approach into a convenient series of analysis tools and establish a standardised framework for evaluating complex plans, whilst also making the process more manageable and user-friendly. The software, IPEX (*I*MRT *P*lan *EX*amination), has been written in MATLAB 7.1 and facilitates comprehensive, automatic and systematic IMRT plan evaluation. The graphical user interface (GUI) is divided into three main sections: dose-volume, dose-ratio and biological, which the user launches sequentially (*Figure 4.6*). The program requires the input of a number of plan data files generated by XiO® during the optimisation and dose calculation stages of planning. The user must transfer each of these via a network from the UNIX planning station to their MATLAB directory on a Windows desktop. The relevant files are:

- Files containing DVH data for all required structures.
- Dose plane matrices (*.txt) for every image slice in the dataset.
- Wire contour files (*.WC) containing co-ordinates of all structure contours in the patient frame of reference.

The planning system calculates dose only within some specified region; the user must record the height, width and central co-ordinates of this dose calculation box for the given plan. A discussion of the algorithms and functionality of the program follows.

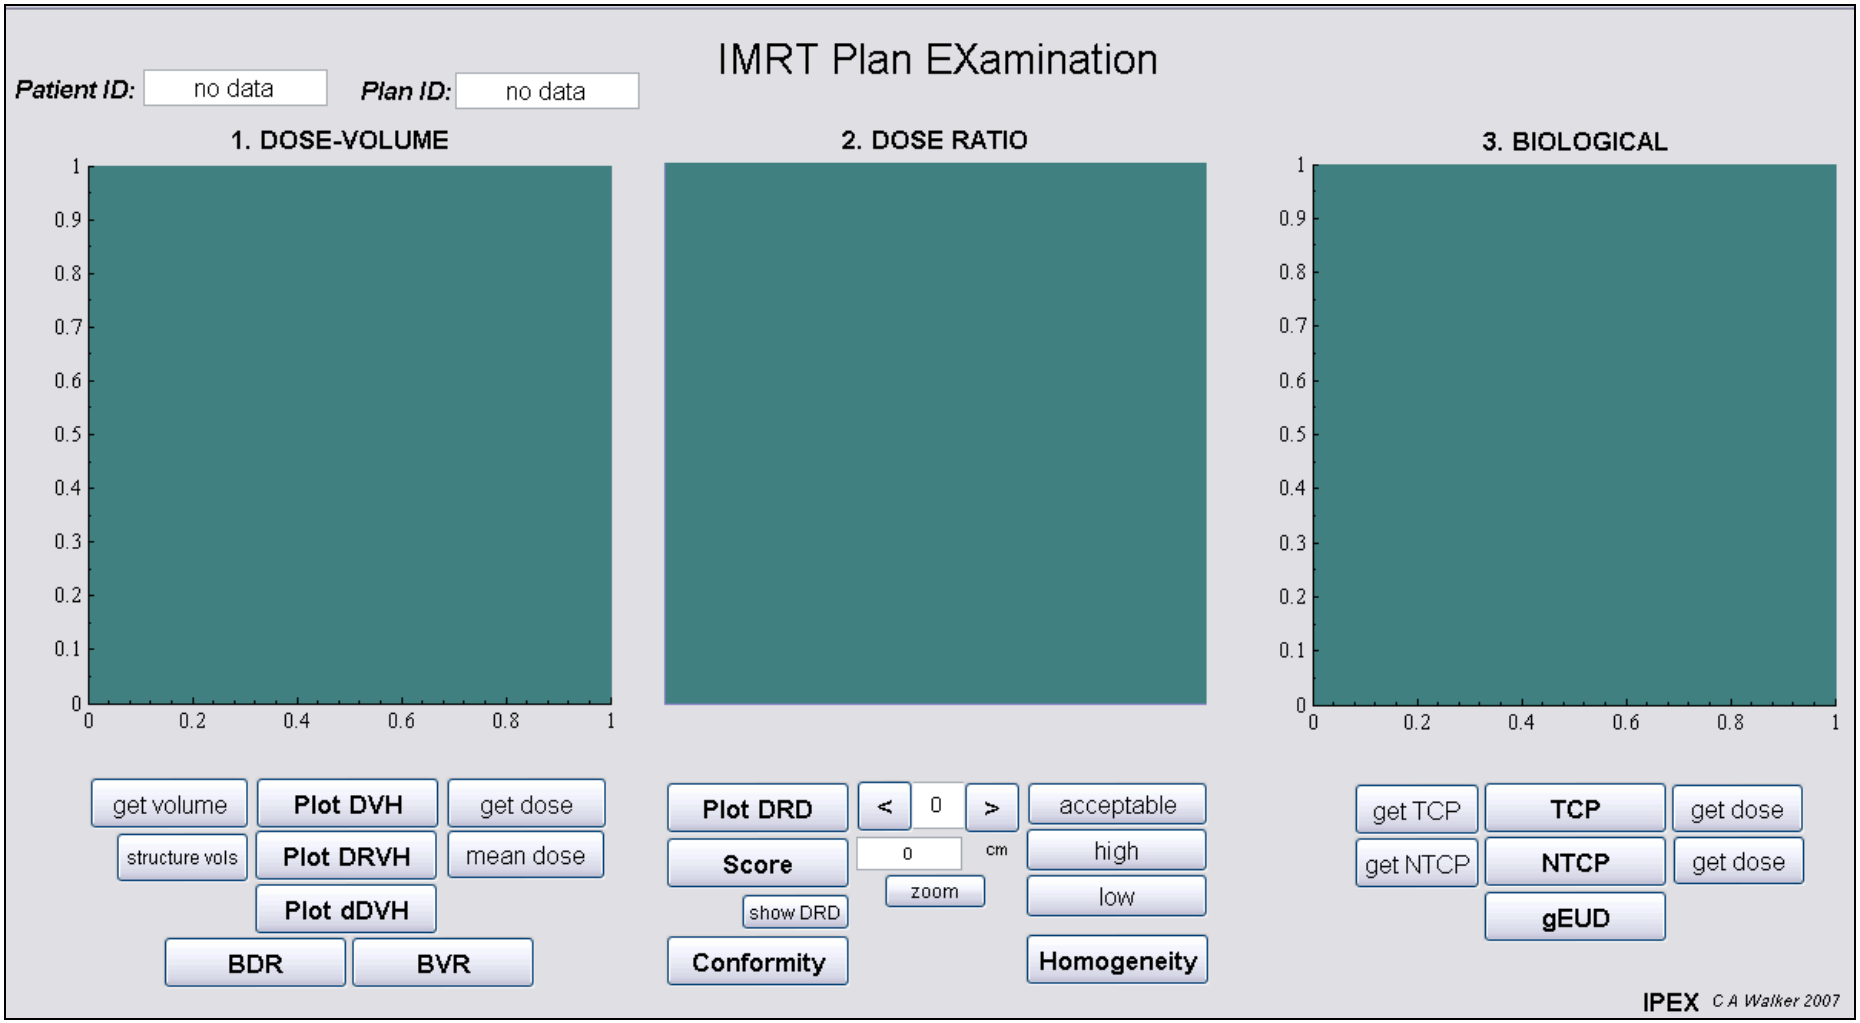

Figure 4.6: The IPEX (IMRT Plan EXamination) GUI for dosimetric and biological evaluation of IMRT plans generated with the CMS XiO® planning module.

4.3.1. Dose-Volume Analysis

The dose-volume analysis in IPEX incorporates a DVH function, from which dose and dose-volume statistics can be obtained, as well as dDVH plots (*Figure 4.7*). For a given target, it is useful to report the percentage of the target/s receiving 95 %, 100 % and 107 % of the prescription dose [27]. For both targets and OARs, the mean dose and $D_{98\%}$, $D_{2\%}$ to represent the minimum and maximum doses are also relevant. Where the DVH function in XiO® is limited to displaying seven structures at any one time, IPEX can plot a maximum of fourteen structures simultaneously. The program reads the RTOG files produced by the planning system in order to plot the corresponding DVH. The user can then obtain dose for a specified volume or get a volume for a specified dose via a single button press. As well as the most commonly used integral DVH, the user can plot the dDVH. The volume and mean dose of a structure are also given. The program calculates the boost dose ratio (BDR) as the ratio of the boost dose level to the prescribed dose to the rest of the target volume and the boost volume ratio (BVR), the proportion of the target volume that is boosted, which are useful in characterising the SMB technique.

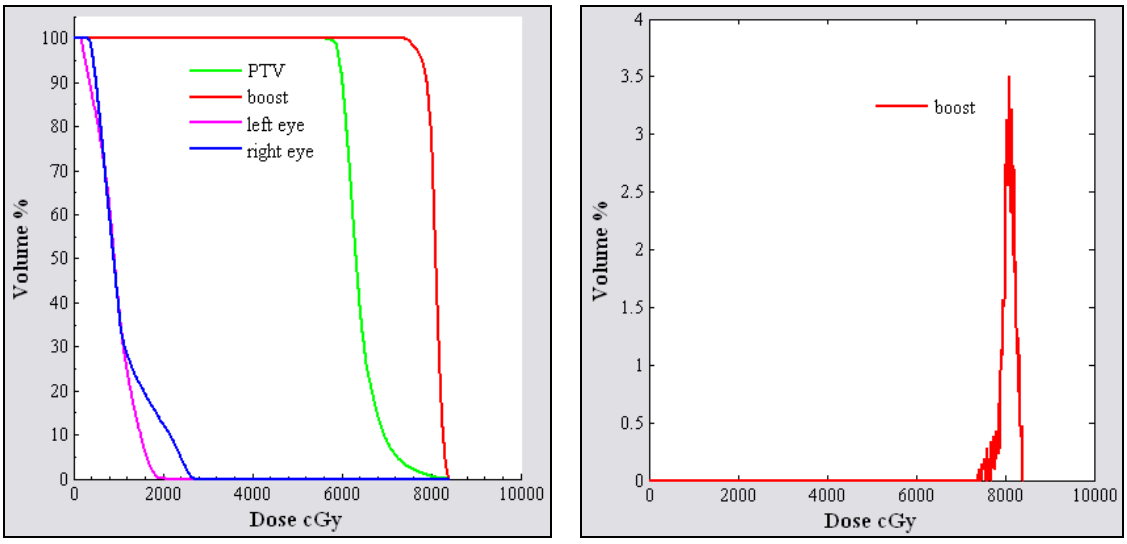

a. *b.*

Figure 4.7: Dose-volume analysis in IPEX. a. The standard DVH, from which dose and volume values may be read and mean structure doses and volumes calculated. Example is a brain IMRT plan with one boost volume. b. The differential DVH (dDVH) for the boost volume, prescribed 80 Gy.

4.3.2. Dose Ratio

Spatial dose information is important in identifying and evaluating hot/cold spots. Isodose lines and 3D isosurfaces can be visualised in XiO®, but a colourwash displaying all dose levels is not available. IPEX provides capability for displaying the 'dose ratio distribution' (DRD) of a plan, which is calculated from a user-specified prescription dose as the dose ratio D_i to D_p, where D_i is the dose in voxel i of the given structure and D_p is the prescription dose. The DRD then, is a spatial distribution of the fraction of the prescription dose received by every voxel within the dose calculation region. The dose-ratio in every voxel is depicted by its own colour, with a colourbar serving as a quantitative reference. Structure contours are overlaid on the DRD, enabling the location of hot/cold spots with respect to a particular target or OAR to be easily identified. The user can switch between a display of the entire DRD and those regions where dose is within some acceptable range or unacceptably high or low (*Figure 4.8*). The user can scroll through the entire patient anatomy to view the DRD on every slice, as well as zoom in on a particular region. The y-axis location is given for each DRD slice, so that this can be readily corresponded to the data in the planning system. A dose-ratio volume histogram (DRVH) can also be plotted, generated simply by dividing the dose points of the standard DVH by a given prescription dose. This is as useful in estimating the volume of particular hot spots, as well for easy extraction of the standard -5 /+7 % dose-volumes.

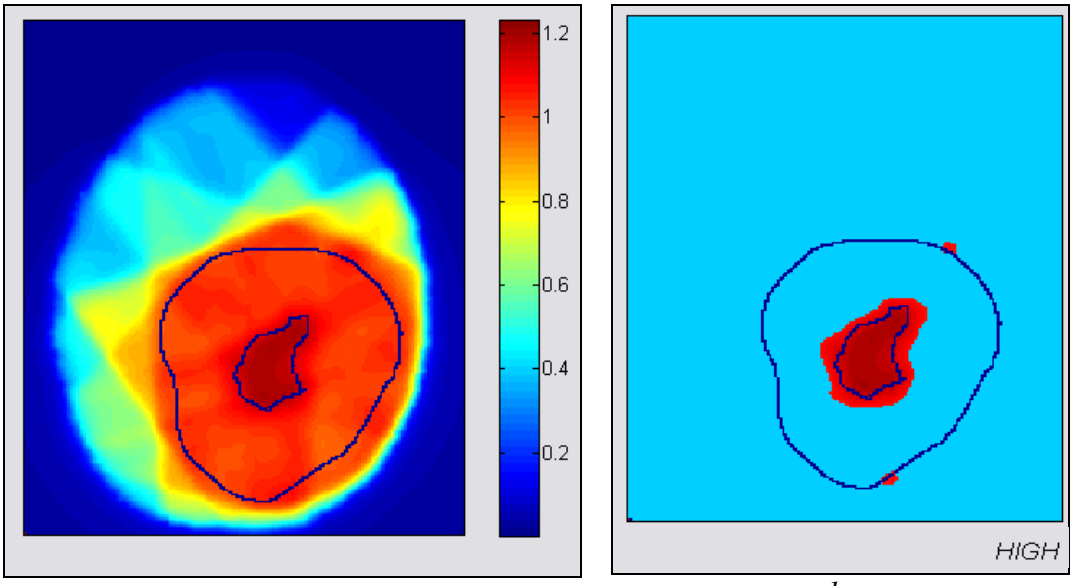

a. b.

Figure 4.8: Dose-ratio distribution visualisation in IPEX. a. The dose ratio distribution for the PTV of an IMRT brain treatment, with PTV and boost volume contours overlaid. b. Dose ratio distribution showing the hot spots (more than +107 % of prescription dose) for the PTV – the higher dose of the boost volume is also highlighted.

4.3.3. Scoring

Quantitative tools are easy and concise to report, and are useful both as measures of a particular plan characteristic and in ranking similar, competing plans.

Conformity Scoring

Any conformity index should be conveniently derived from existing plan data and easily interpreted in order to be of any real practical use. Taking this into consideration, the indices implemented in IPEX are the CI, defined simply as the ratio of the isovolume to the target volume and CN given by equation (4.1). CN is redefined as a conformity score (CS), where in this form it is presented as a percentage. Reporting of the two indices provides information on the relative size of the isodose volume as well as the target coverage, resulting in a comprehensive conformity analysis.

Homogeneity Scoring

IPEX provides a novel homogeneity analysis in terms of both minimum and maximum doses in a given target; HS_{min} and HS_{max} respectively, as well as an overall homogeneity score, HS. These are calculated from the local minimum dose, D_{min} and maximum dose, D_{max} relative to D_p in a given structure as

$$HS_{min} = \frac{D_{min}}{D_p} \tag{4.14}$$

$$HS_{max} = \frac{D_p}{D_{max}} \tag{4.15}$$

$$HS = \left(HS_{min} \times HS_{max}\right) \times 100\% \tag{4.16}$$

where a higher score corresponds to better homogeneity. HS (4.16) is equivalent to D_{min} as a percentage of D_{max}.

Objective Function as a Dose Score

IMRT optimisation iteratively minimises a function of dose objectives to find the best possible dose distribution. In this way, the objective function – a mathematical combination of desired dose objectives – serves to score a given distribution of dose. It follows then that an objective function can also be employed as an evaluation tool for a final treatment plan. This concept has been proposed elsewhere as the 'dose quality factor' (DQF) [83]. The general case is given by

$$\text{score} = \sum_{\sigma=1}^{m} \frac{w_\sigma}{n} \sum_{i=1}^{n} \left(\frac{D_i}{D_0} \right)^{p_\sigma} \tag{4.17}$$

with m being the number of objectives, w_σ the objective specific importance weighting, n the number of voxels in the structure, D_i the dose in voxel i, D_0 the desired dose level and p_σ the exponent of objective σ. IPEX uses a score function for multiple targets and OARs based purely on desired dose ranges and does not take into account dose-volume

or biological objectives. Unlike the DQF, the IPEX tool scores the whole target or organ rather than selected sub-volumes. The function uses dose ratios rather than dose differences (used in the optimisation scheme) in order to yield a score of unity or less, where unity corresponds to a plan which perfectly matches the given objectives. The score tool does not include any sort of priority tool to account for overlapping structures; the algorithm assumes that all voxels within a given contour are owned by that structure. The score of a target thus comprises of three objectives: a minimum accepted dose, a prescribed dose and a maximum acceptable dose.

$$\text{score}_{\text{target}} = \frac{w_{\min}}{n}\sum_{i=1}^{n}\left(\frac{D_i}{D_{\min}}\right)^{p_{\min}} + \frac{w_p}{n}\sum_{i=1}^{n}\left(\frac{D_i}{D_p}\right)^{p_p} + \frac{w_{\max}}{n}\sum_{i=1}^{n}\left(\frac{D_i}{D_{\max}}\right)^{p_{\max}} \tag{4.18}$$

where w_{min}, w_p and w_{max} are the weighting factors for the minimum, prescribed and maximum doses respectively, the sum over which must be unity, D_{min}, D_p and D_{max} are the minimum acceptable, prescribed and maximum acceptable doses and p_{min}, p_p, and p_{max} the exponents of the minimum, prescribed and maximum dose objectives respectively. The weighting factors and exponents can be used to distinguish between each dose objective by modulating the influence each objective has over the final score. In the case of $D_i > D_{min}$, D_i/D_{min} is set to an optimal value of unity, as is D_i/D_p if $D_i > D_p$ and D_i/D_{max} where $D_i < D_{max}$. Target and OAR scores are calculated separately, due to the differing objectives. IPEX computes the score of an OAR based on a maximum dose objective only, and as with the target objectives, D_i/D_{max} is set to unity in the case of $D_i < D_{max}$.

$$\text{score}_{\text{OAR}} = \frac{w_{\max}}{n}\sum_{i=1}^{n}\left(\frac{D_i}{D_{\max}}\right)^{p_{\max}} \tag{4.19}$$

For this single objective approach for OARs, w_{max} is set to unity, as there are no other objectives to rank its importance against. An overall dose score for a given collection of structures is computed from the individual target and OAR scores using

$$\text{score} = \sum_{t=1}^{T} w_t \text{score}_{\text{target}} \times \sum_{o=1}^{O} w_o \text{score}_{\text{OAR}} \tag{4.20}$$

where T and O are the number of targets and OARs, w_t and w_o weighting factors for each individual target, t and OAR, o. The sum must be unity over all target weighting factors and OAR weights respectively. IPEX uses default values of m^{-1} for objective weighting factors, T^{-1} for individual target weightings and 1/O for individual OARs i.e. equal weighting of all objectives and targets/OARs. The default values of all scoring parameters used by IPEX are shown in *Table 4.1*.

Structure type	**Parameter**	**IPEX default value**
TARGET		
objective 1 (minimum dose)	D_{min} w_{min} p_{min}	95 % of D_p 0.6 2
objective 2 (prescription dose)	D_p w_p p_p	user-specified 0.3 1
objective 3 (maximum dose)	D_{max} w_{max} p_{max}	107 % of D_p 0.1 -1
OAR		
objective 1 (maximum dose)	D_{max} w_{max} p_{max}	50 Gy 1 -2

Table 4.1: Default values in IPEX for the dose 'score' function for both targets and OARs. The target minimum dose objective has the biggest influence over the final score, having the highest weighting factor and power i.e target underdosing is more significant than target overdosing. The use of a negative exponent for maximum dose objectives ensures that overdose rather than underdose is penalised. These are default values only - all parameters can be individually set by the user.

4.3.4. Biological

gEUD

The XiO® planning module does not provide EUD capability, but the data available via 'RTOG files' provides the dose information necessary to calculate the gEUD for any of the plan structures within IPEX using

$$\mathrm{gEUD} = \left(\sum_{i=1}^{n} v_i D_i^a \right)^{\frac{1}{a}} \tag{4.21}$$

where v_i is the fractional volume receiving dose D_i and n is the number of dose bins used to generate the DVH. The user specifies the value of the tissue-specific parameter, a, where the default value is a=-8. This is an average value for brain tumours, for normal brain tissue, a=+5 is typical [84].

TCP & NTCP

IPEX uses equations (4.5), (4.10) and (4.11) to calculate the TCP of a given target. This version of the model accounts for non-uniform dose distributions and inter-patient heterogeneity but no temporal effects e.g. cell proliferation. In order to model radiosensitivity variability, the algorithm generates a random Gaussian distribution with a user-specified mean, α_{mean}, and standard deviation, σ_α, over a user-specified population, K. The same (random) distribution is reproduced for each calculation by resetting the state of the random number generator to some fixed value at each call. If α_{mean} is very small, negative numbers may be randomly generated, which are meaningless. In order to avoid the problem of these non-physical negative values of α, the program shows an error message that such a distribution of patients does not exist. The TCP for a number of arbitrary, uniform dose levels is calculated using (4.7), (4.8)and (4.9), from which a dose-response curve is plotted of TCP as a function of

dose. The user can return the TCP for a specified dose and a dose for a given TCP value.

IMRT plans utilising SMB at different doses per fraction can have different radiobiological consequences than traditional plans delivered with a uniform dose per fraction. Additionally, the longer treatment times of some IMRT treatments may also be radiobiologically relevant and so it may be important to include time effects in a clinical plan evaluation scenario. However, TCP is used from this point forward as a comparative measure between other plans of identical fractionation regimes and so time response mechanisms were not deemed necessary. For a realistic prediction of TCP, these should be modelled, especially since high-grade tumours such as GBMs have accelerated repopulation, which is considered a contributory factor in local failure.

IPEX also generates a dose-response curve for normal tissue and OARs, using the NTCP model in (4.12) and (4.13) from which the corresponding NTCP for a given dose can be obtained. This assumes that the whole organ, of uniform functionality, is irradiated uniformly. Default values for the biological parameters are given in *Table 4.2* [76, 77]. It is important to recognise that current dose-response data is limited, especially for atypical treatments such as IMRT and specifically, SMB. All available data on the radiosensitivity of different tissues was acquired before the era of conformal therapy and indeed before 3D dose calculations were possible. However, until more up to date data becomes available, they serve as a starting point for radiobiological evaluations.

Model	Parameter	IPEX default value
gEUD	a	-8
TCP	α_{mean}	0.35 Gy^{-1}
	σ_α	0.05 Gy^{-1}
	ρ	10^6 cells/cm^3
	K	10^4
NTCP	TD_{50}	60 Gy
	m	0.15

Table 4.2: Default parameters for biological plan evaluation in IPEX using gEUD, TCP and NTCP. Parameters are set to user-specified values during program use.

4.4. Software Testing & Validation

It is critical to test any piece of software before use in order to validate its accuracy and applicability. As well as significant checking of the standalone functions during their creation, IPEX was further validated by generating a pseudo isodose distribution (PID) in XiO® file formats and running this through the software. Because the simulated dose distribution was exactly known, the corresponding plan evaluation results could also be manually calculated, against which the results given by IPEX could be verified. The PID was a single plane (i.e. 2D) 200 pixel square matrix, with square isodoses, equispaced at 5 % increments. The prescription dose (100 %) was set to 60 Gy and the central 10 x 10 square of the PID set to 105 % (63 Gy). A pseudo PTV, denoted pPTV, was overlaid to encompass the 95 % isodose exactly, which represented optimal PTV dosimetry i.e. 100 % of the volume receives 95 % or greater and zero volume receives 107 % or greater. A pseudo OAR, pOAR, was also overlaid in a known, arbitrary position in the region of interest.

Figure 4.9 shows how IPEX successfully read the created files for the DRD (*Figure 4.9a*), with the isodoses represented correctly and the contours overlaid in the correct positions. *Figure 4.9b* shows the DVHs of the pPTV and pOAR, where the dose-

volumes were manually calculated from the pixels in the dose matrix and compared against the DVH values. Conformity indices were returned as their maximum values; CI=1 and CS=100 % for the ideal PID and homogeneity scores were also recorded, with HS_{min}=0.95, HS_{max}=0.95 and HS=90.3 % . Because the PID was so simple, it was easy to check the gEUD results against manual calculations from the pseudo RTOG file, which confirmed values of 58.4 Gy and 44.3 Gy for the target and OAR respectively. IPEX calculated the results as exactly those calculated manually from the PID in all cases, and this was taken as adequate validation of the software.

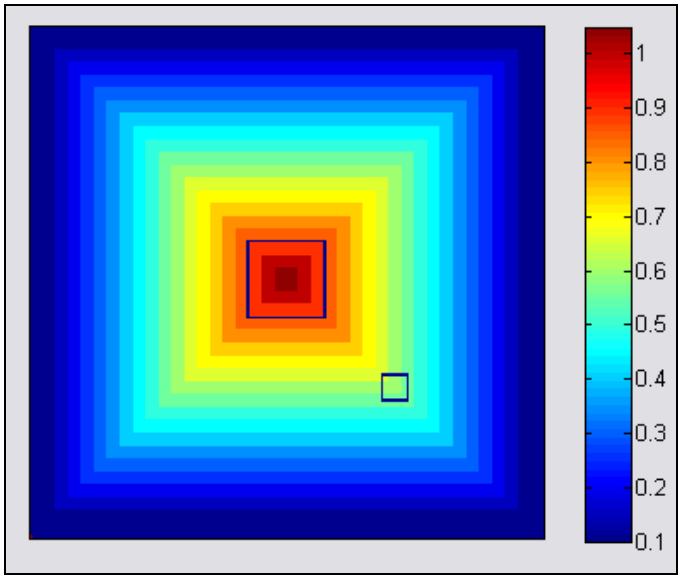

a.

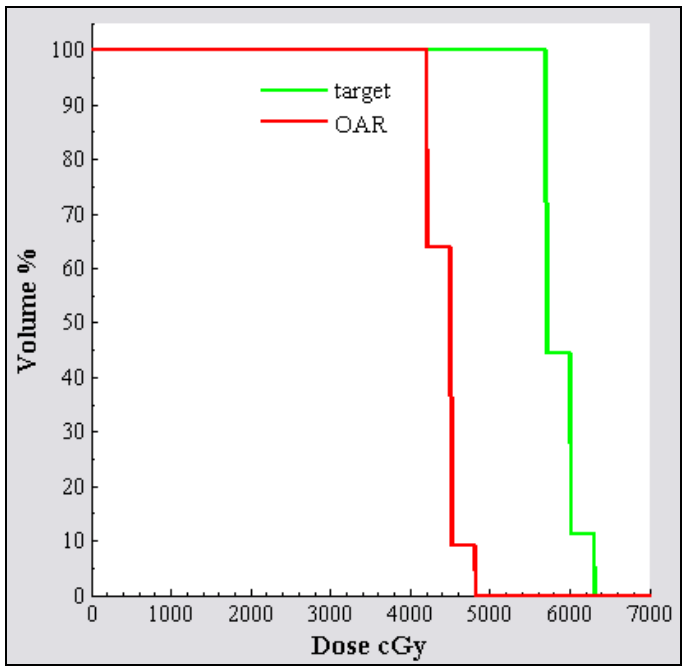

b.

Figure 4.9: Validation of IPEX with a simple pseudo dose distribution (PID). a. DRD showing the square PID, with pPTV and pOAR contours overlaid. b. DVH of the pseudo PTV (pPTV) and pseudo OAR (pOAR).

4.5. Conclusions

A plan-analysis toolbox, IPEX, has been developed, whereby numerous dosimetric and biological methods have been successfully implemented and verified. This is a useful environment for plan evaluation, as it includes novel tools not available within the planning system, and each one is easily launched from one simple GUI, resulting in

tidy, systematic and user-friendly plan evaluation. The combination of plan interpretation methods available in IPEX serves as a comprehensive means of reporting IMRT treatments. The amount of pre-processing required renders the utilisation of every tool in all routine plan evaluation unrealistic. However, it may be that a small selection of tools are appropriate for evaluation of an individual plan or that all tools are deemed valuable in evaluating some particularly complex or critical plan.

Most methods are sensitive to target delineation and all of the evaluation tools assume the spatial distribution of radiosensitivity (i.e. biological significance) of a structure is uniform. It is clear that real-life tumours do not exhibit biological homogeneity and with the development of functional imaging techniques, this biological distribution may soon have to be incorporated in some way for accurate plan evaluation [17, 85-88]. Additionally, in cases where a very steep dose fall-off at the target surface is sought, metrics characterising dose gradient information may be useful [61, 89].

A single quantity for global plan characterisation or 'score', although desirable, will never prevail, for two reasons; firstly, due to the competing objectives of high target dose and minimal normal tissue dose and secondly because it infers a reduction in dimensionality i.e. a 1D representation (single scalar value) of a 3D dose distribution and a loss of spatial information is therefore inevitable. Ideally, plan evaluation of complex dose distributions should combine conventional tools with new measures in order to be truly descriptive. Any new tool must still be used in conjunction with traditional evaluation methods, in particular DVH plots, in order to fully realise informed and reliable plan evaluation. IPEX combines all of these approaches to plan evaluation in one easy to use interface that can be executed on any Windows PC. The program is vendor specific in that it is designed exclusively to read CMS XiO® plan

files, although it would be straight forward enough to adapt the tools and layout to handle other file formats, as long as similar data could be extracted from the given TPS.

Chapter 5 - Controlling Local Dose for Selective Multiple Boosting

5.1. Introduction

Intensity modulated radiotherapy (IMRT) provides extraordinary capability to accurately plan and deliver complex, conformal 3D dose distributions. The concept of selective multiple boosting (SMB) uses IMRT to design and deliver multiple, simultaneous dose escalation or 'boosts' to sub-volumes within a tumour, which may potentially improve the probability of disease control. Advanced imaging techniques are capable of non-invasively providing metabolic and functional information to target with SMB, the regions that will benefit most from higher doses [21-23]. Before implementation of these advanced delineation techniques can evolve, it's imperative to understand what is practically achievable in terms of control over exactly where and how much dose can be deposited in SMB planning schemes.

Current optimisers are not designed to easily read or interpret parameter maps derived from different imaging modalities. New approaches are therefore required, for defining the objective function and in understanding the influence of optimisation criteria in the context of intra-structure trade-off. The purpose of the SMB technique is not to escalate dose across whole volumes, but to focus discrete boost levels on those regions that will respond best to each dose. Care must be taken then that dose levels do not 'merge' when boost volumes are in close proximity to one another, as this may contravene the multiple-boost philosophy.

Evidence for the feasible domain of the application of advanced imaging to deliverable, clinically desirable radiotherapy plans is limited and no quantitative restrictions on this capability have been investigated. This study aimed to characterise a commercial IMRT module (CMS[§] XiO®) by investigating optimisation parameters firstly for multiple, concentric spherical boost volumes and secondly for cubic pseudo-voxels. The pseudo spherical anatomy represents a standard boosting scenario whereby boost volumes are enclosed within larger planning volumes, with spheres being the simplest geometry for feasibility testing. The pseudo cubic boosts were intended to simulate large voxels of functional imaging such as magnetic resonance spectroscopic imaging (MRSI). These experiments should go some way to answering what sort of influence over dose deposition is achievable and therefore on what kind of resolution it is feasible to use biological data for radiotherapy planning, regardless of what may be radiologically possible.

5.2. Methods

5.2.1. Spherical Pseudo-Boosts

Plan parameters

A 21 year old female patient with a low-grade tumour was CT scanned and the images loaded into a CMS XiO® planning station. Here, 'pseudo' concentric spherical boosts, located centrally within the oncologist-defined PTV, were created using an auto-contouring tool (*Figure 5.1*).

[§]St. Louis, Missouri

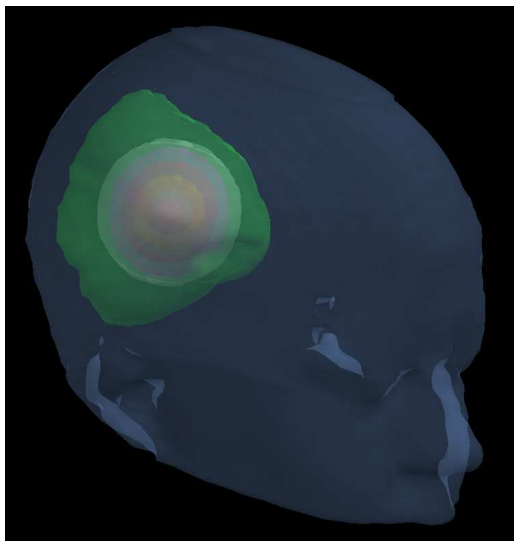

Figure 5.1: Patient 3D geometry for spherical experiments; oncologist-defined PTV (green) encompassing six spherical pseudo-boost volumes.

The spherical contours are systematically referred to as Boost 1, Boost 2, Boost 3, Boost 4, Boost 5 and Boost 6, where Boost 1 is the largest boost volume with radius, R = 3.0 cm and therefore closest to the PTV surface. The radii of each boost volume decreased in 0.5 cm increments resulting in the smallest volume, Boost 6, having R = 0.5 cm.

A typical but arbitrary IMRT beam configuration was selected since the dose distributions were to be compared relatively and experimentally, not clinically (the patient was treated with conventional radiotherapy methods). This utilised five coplanar beams at gantry angles[§] 0°, 72°, 135°, 216°, 288°. Dose calculations were carried out with the FFT convolution algorithm on a 2 x 2 x 2 mm grid, based on a step and shoot MLC mode of delivery. A baseline dose of 55 Gy was prescribed to the PTV throughout, and all prescribed doses entered into the optimisation as simple minimum and maximum dose objectives, both equal to the given prescribed level. The optimiser was set to converge at thirty iterations or a difference of 0.01 % between the current and previous cost function value. The cost function value (normalised to 100 at start of search) and number of iterations at which each plan converged were recorded to give an indication of how closely and easily the optimiser achieved the set goals.

[§] IEC convention

Treatment plans were evaluated in terms of dose coverage and conformity to each individual volume and dose level using IPEX (Chapter 4). Quantitative dose tools only were employed, to facilitate simple comparisons of the large number of plans. Target coverage, TC, was assessed in terms of the percentage volume of a given structure receiving 95 % or higher of its prescription dose level. Conformity was determined using the conformity number, CN, given by equation (4.1). The degree of conformity increases as CN tends to unity and optimal target dose coverage occurs when TC = 100 %. Finally, the generalised equivalent uniform dose (gEUD) and corresponding tumour control probability (TCP) values for the PTV were computed, in order to give some indication of the dosimetric and biological benefit of the different boost scenarios.

Characterising practical dose deposition

Initially, Boosts 1, 3 and 5 only were included in the planning, representing $\Delta R = 1.0$ cm between boost volumes. Dose objectives, importance weights and powers (objective function exponents) were then investigated as independent controlling parameters. Firstly, the dose prescribed to each boost volume was varied, based on a constant dose differential, ΔD. ΔD is the difference in prescribed dose, D_p, between sequential volumes i.e. Boost 5 was invariably prescribed the highest dose where

$$\Delta D=(D_{p\_PTV}-D_{p\_Boost1})=(D_{p\_Boost1}-D_{p\_Boost3})=(D_{p\_Boost3}-D_{p\_Boost5}) \quad (5.1)$$

ΔD was increased over the range 1-12 Gy in 1 Gy increments, and based on the dual-parameter dosimetric evaluation method (TC-CN), a maximum achievable dose differential, ΔD_{max}, was deduced. This was then selected as the dose objective for intra-PTV importance scaling. This was achieved firstly by varying the importance weight and secondly the power of each structure individually; PTV, Boost 1, Boost 3 and Boost 5. Resulting plans were evaluated as before.

The experiment of varying the dose difference between sub-volumes, ΔD, was repeated to include all the pseudo boost volumes (Boost 1, Boost 2, Boost 3, Boost 4, Boost 5 and Boost 6), representing ΔR = 0.5 cm between each volume in order to simulate finer resolution functional data and further test the capability of the optimiser. ΔD was varied over the range 0.25-8 Gy. Smaller values of ΔR were not considered relevant, given that the best available MLC leaf width for these experiments was 5 mm.

Finally, the use of rectangular rather than square beamlets was investigated. The beamlet dimension in the x-direction is set by the physical MLC leaf width (5 mm), but the z-dimension can be varied by changing the leaf step increment i.e. the incremental distance each leaf can move. XiO® allows values between 0.3–2.0 cm, in 0.1 cm increments. Decreasing the step below 5 mm results in an increase in dose resolution in one direction compared to the conventional square beamlet scenario, potentially therefore, giving finer control over dose deposition. A larger leaf step creates larger beamlets and hence results in a decrease in dose delivery resolution. The investigation was carried out on all six boost volumes rather than the three-boost model, as the dosimetry of closer-lying structures should show the most measurable benefit from better dose resolution. For all other experiments, conventionally square, 5 x 5 mm^2 beamlets were used in optimisations, reflecting the resolution of the MLC leaves.

5.2.2. Cubic Pseudo-Voxels

Eight individual, contiguous, cubic volumes were contoured on the same CT dataset using the auto-shape contouring tool in the centre of the oncologist-defined PTV to simulate voxels of functional data e.g. MRSI (*Figure 5.2*).

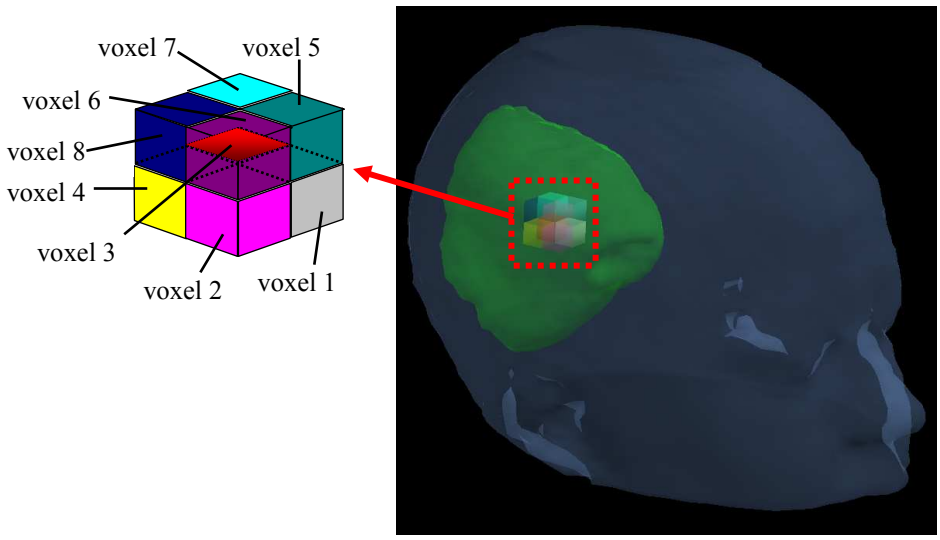

Figure 5.2: 3D patient geometry for cubic experiments; oncologist-defined PTV encompassing eight cubic pseudo-boost 'voxels'.

A boost dose, ΔD_1, was assigned to pseudo voxels 2-8 individually and a higher boost, ΔD_2, prescribed to voxel 1. If D_{PTV}, $D_{vox\_2\text{-}8}$ and D_{vox_1} represent the dose objectives for the PTV, voxels 2-8 and voxel 1 respectively then

$$\Delta D_1 = D_{vox\_2\text{-}8} - D_{PTV} \tag{5.2}$$

$$\Delta D_2 = D_{vox\_1} - D_{vox\_2\text{-}8} \tag{5.3}$$

ΔD_1 was set to a constant value of 10 Gy i.e. prescribed dose was 55 Gy to PTV and 65 Gy to voxels 2-8. ΔD_2 was then varied from 1 to 20 Gy and the plans assessed in terms of optimisation results, individual target dose coverage and conformity as previously. Homogeneity was also evaluated using the IPEX homogeneity score (4.16). All other plan parameters were as the spherical boost experiments.

ΔD_2 was varied using just 2 volumes in the optimisation; voxel 1 and voxel_combined, an amalgamation of voxels 2-8 (*Figure 5.3*). This was to investigate whether the optimiser was affected by instruction to sculpt dose around a larger volume rather than individual smaller volumes covering the same volume in space. This could be an important issue when defining the process for using real functional data in planning. This two-volume scenario was then used to investigate the power function at a constant

$\Delta D_2 = 5$ Gy by altering the power of voxel 1 and voxel_combined independently. Importance weights were not investigated after the spherical boost experiments showed that these had only a subtle impact on dose distributions.

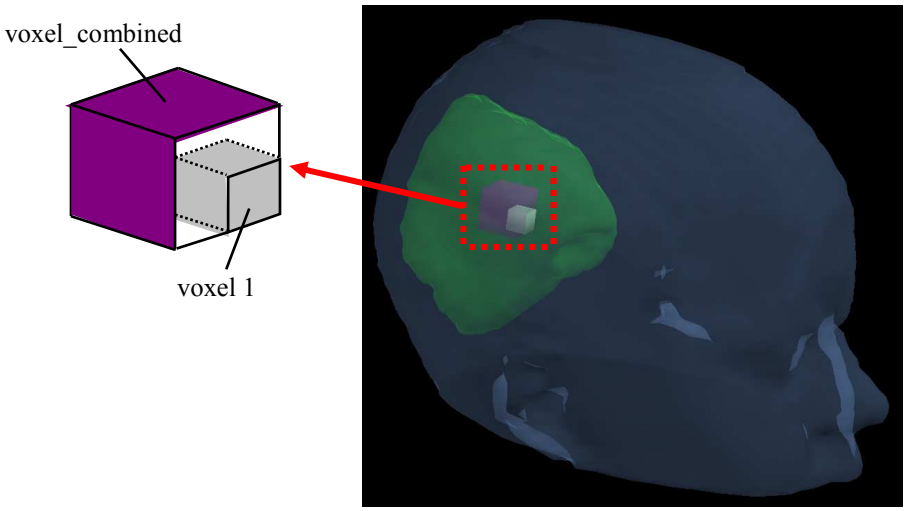

Figure 5.3: 3D patient geometry for the two-volume study, where voxel_combined covers the same volume as voxels 2-8.

A final experiment using a new parameter, ΔD was carried out. Here, ΔD represents the dose difference between each voxel sequentially i.e. each voxel is prescribed a different dose where $\Delta D = D_{vox\_2} - D_{vox\_1} = D_{vox\_3} - D_{vox\_2}$….etc. The choice of dose was entirely arbitrary, the purpose of which was to investigate the feasibility of optimising on a 1 x 1 x 1 cm^3 voxel-by-voxel basis. Dose objectives for the PTV and voxel 1 were kept constant at 55 and 60 Gy respectively.

5.3. Results & Discussion

5.3.1. Spherical Pseudo-Boosts

Maximum achievable dose differential, ΔD_{max}

Objective function scores were better for ΔR = 0.5 cm than ΔR = 1.0 cm when ΔD was the varying parameter. This is shown in *Figure 5.4* by convergence in fewer iterations and to a lower cost function value. This suggests that the optimiser more quickly and closely met the required goals for the finer resolution scenario.

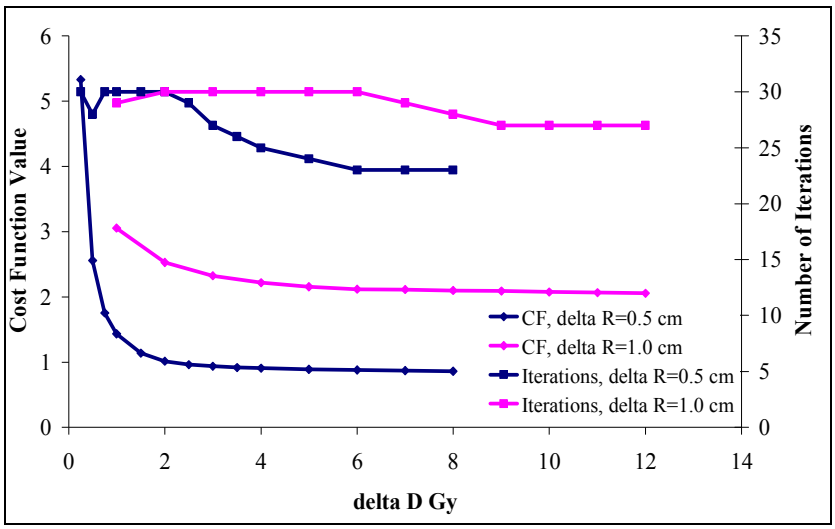

Figure 5.4: Optimisation scores at convergence as a function of ΔD; number of iterations and normalised cost function value for ΔR=1.0 cm and ΔR=0.5 cm.

Target dose coverage (95 % of prescription level), TC, of the boost volumes was optimal (100 %) at small ΔD, dropping off as ΔD increased (*Figure 5.5*). This implies that only values below some maximum, denoted ΔD_{max}, are viable in terms of an achievable dose differential. PTV coverage was of course 100 % throughout, since the higher dose boost volumes are contained within the PTV itself and the dose prescribed to the PTV is constant. At ΔR = 1.0 cm (*Figure 5.5a*), TC of all the boost volumes began to fall off at around 5 Gy, reaching a minimum of 84 % coverage (Boost 3, middle sized) at 12 Gy. Therefore, to attain 100 % boost coverage of 95 % of the individual prescription dose, $\Delta D \leq 5$ Gy is required i.e. $\Delta D_{max} = 5$ Gy. It is worth noting that a dose difference of 1-2 Gy is meaningless since this level of inhomogeneity is likely to occur and be accepted in a typical IMRT plan anyway and therefore does not constitute boosting methodology. Excluding the PTV, Boost 5 (smallest) maintained the highest coverage across the range of ΔD.

At ΔR = 0.5 cm (*Figure 5.5b*), TC began to decrease at 3 Gy, suggesting that only values $\Delta D \leq 3$ Gy are feasible in order to maintain dose coverage of the boosts to their individual levels i.e. $\Delta D_{max} = 3$ Gy on a 5 mm scale. At ΔD = 8 Gy, the lowest TC was

89 % (Boost 4) compared to 94 % (Boost 3) at ΔR = 1.0 cm. This shows, as expected, that it is easier to irradiate boost volumes defined on coarser resolutions i.e. larger volumes.

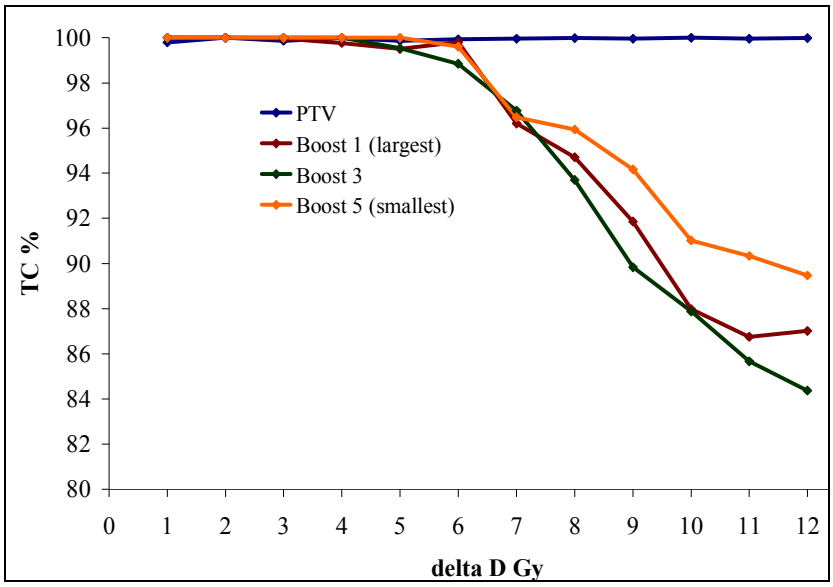

a.

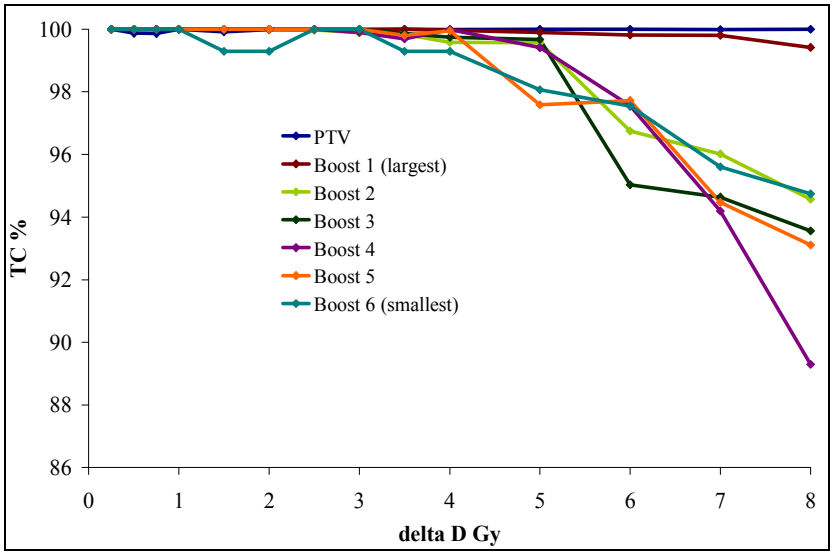

b.

Figure 5.5: 95 % dose target coverage, TC, as a function of ΔD for a. ΔR = 1.0 cm and b. ΔR = 0.5 cm.

Coverage of the target volumes by the prescription dose as opposed to 95 % of this value did not show any such trend and in fact remained fairly constant across the range of ΔD, except in the case of the smallest boost volumes (*Figure 5.6*).

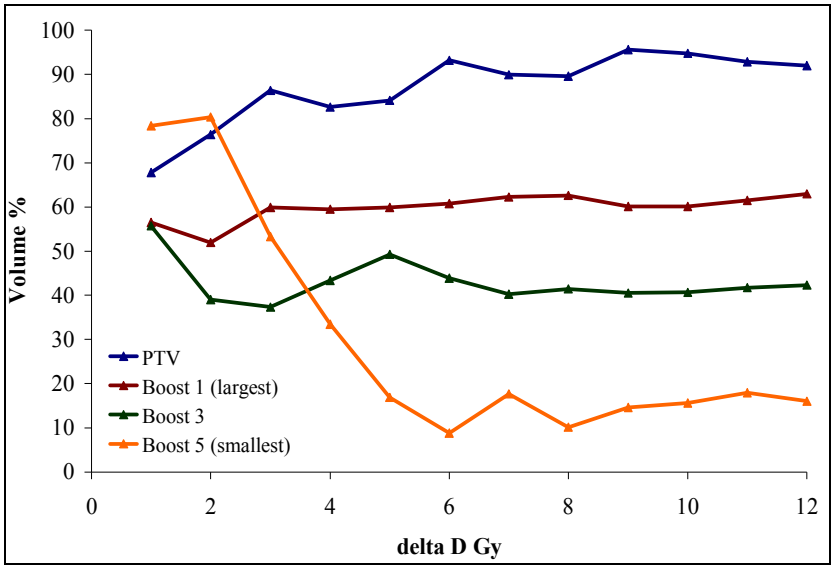

a.

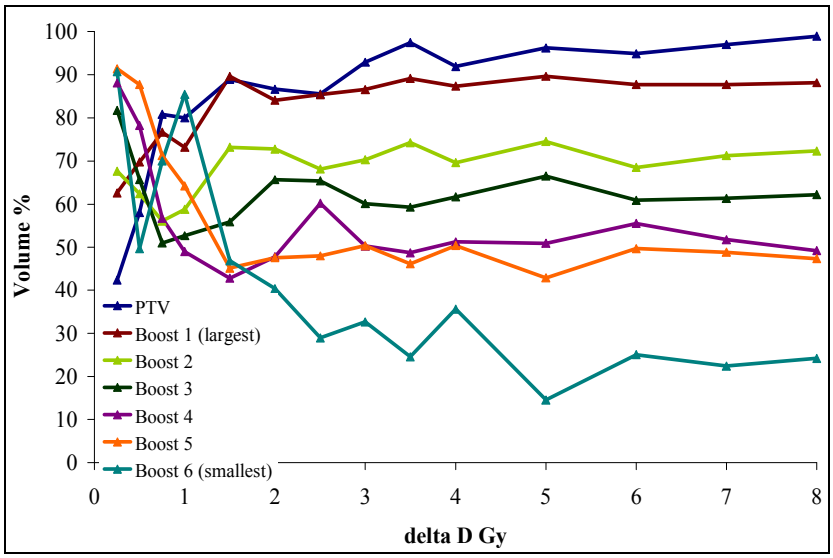

b.

Figure 5.6: Prescription dose coverage as a function of ΔD for a. ΔR = 1.0 cm and b. ΔR = 0.5 cm.

Conformity, CN, improved initially with increasing ΔD (*Figure 5.7*), but since better conformity arises from smaller isodose volumes, the improvement reaches a maximum threshold and the isodose volume decreases to produce sub-optimal target coverage thereafter. It is important to note therefore that although CN appears to continue to show improvement beyond ΔD_{max}, this does not represent better dose distributions in a general sense, only that the 95 % isodose volume is getting smaller. PTV conformity remained constant (CN ≈ 0.7) for both resolution settings. Interestingly, the conformity

plots of other boost volumes seem to intersect this line at approximately ΔD_{max}, suggesting that for perfect target coverage, i.e. TC = 100 % for each volume, the PTV has the best conformity of all. In both experiments, the smallest volume exhibited the worst conformity over the whole range of ΔD. It is known that larger targets tend to demonstrate better conformity, implying that optimising to control dose on an extremely small scale e.g. voxel-by-voxel may not be feasible [65].

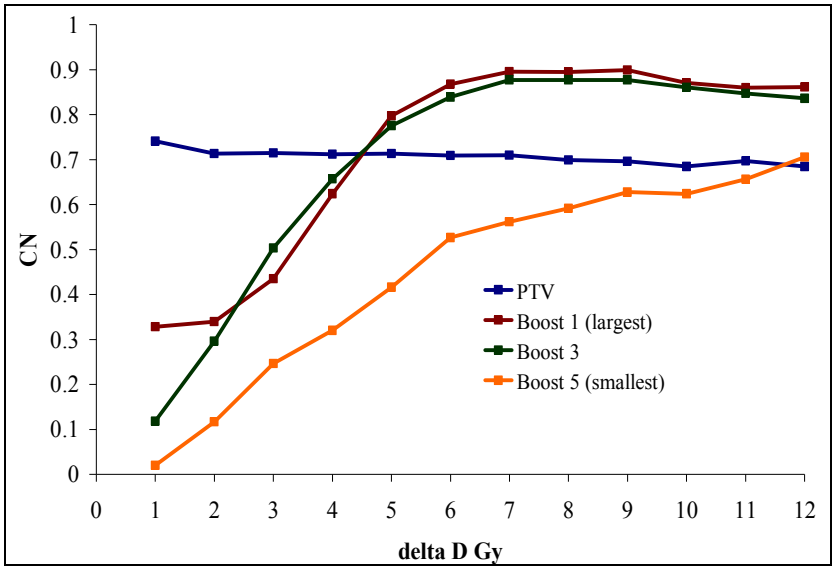

a.

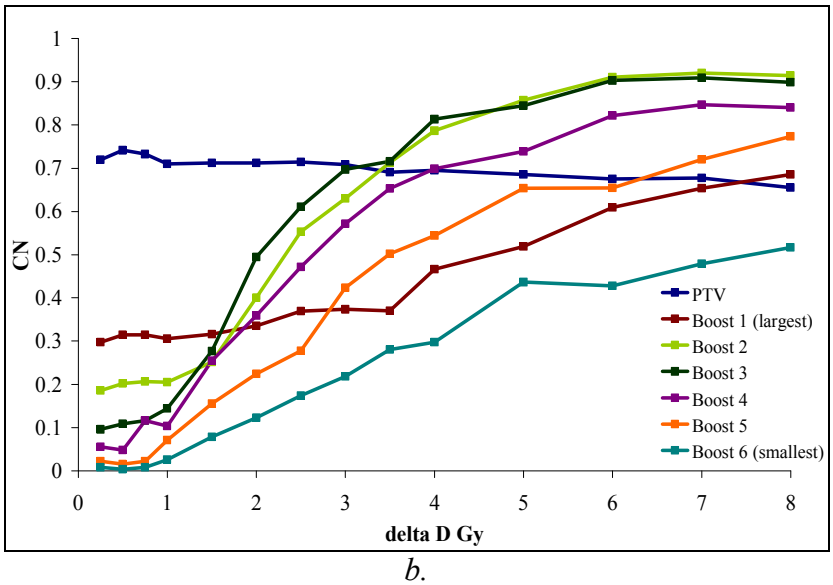

b.

Figure 5.7: Conformity as a function of ΔD for a. ΔR = 1.0 cm and b. ΔR = 0.5 cm.

Higher values of ΔD yielded higher gEUD and TCP results, as the PTV effectively received a higher dose when the boosts were intensified (*Figure 5.8*). Both biological

models approximate a linear increase with ΔD. For the range 1-8 Gy of ΔD, a gain in PTV gEUD was observed of 3.2 Gy when ΔR = 1.0 cm and 4.9 Gy for ΔR = 0.5 cm. This corresponded to an increase in TCP of 11.7 % and 19.4 % respectively. The mean gain in TCP per 1 Gy increment was 1.5 % when ΔR = 1.0 cm, almost doubling to 2.8 % with the additional boosts at ΔR = 0.5 cm. Even with six boosts enclosed by the PTV, which afforded a maximum boost dose ratio (BDR) of 187 %, the equivalent uniform dose delivered to the PTV was still only 4.9 Gy more than when a maximum BDR of just 111 % was applied. This suggests that the boost scenario yields only a small dosimetric gain in this instance, although it is of course further dependent on the boost volume ratio (BVR) as well as the radiosensitivity of the given tissues.

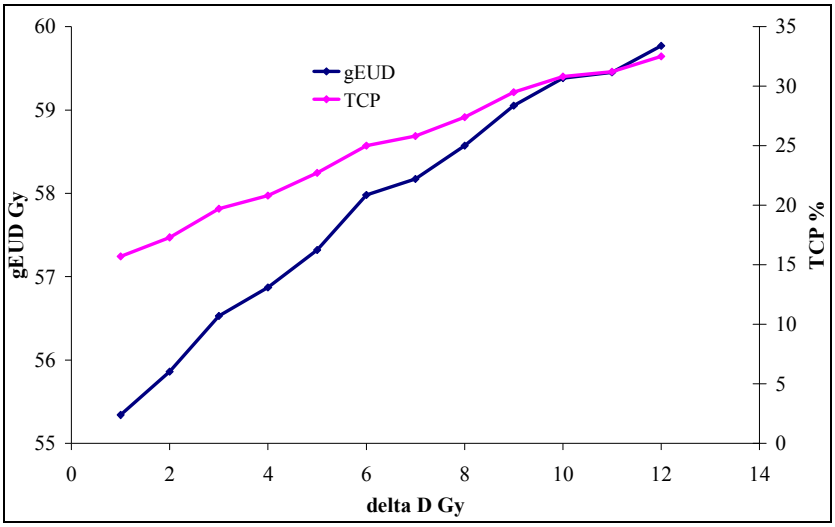

a.

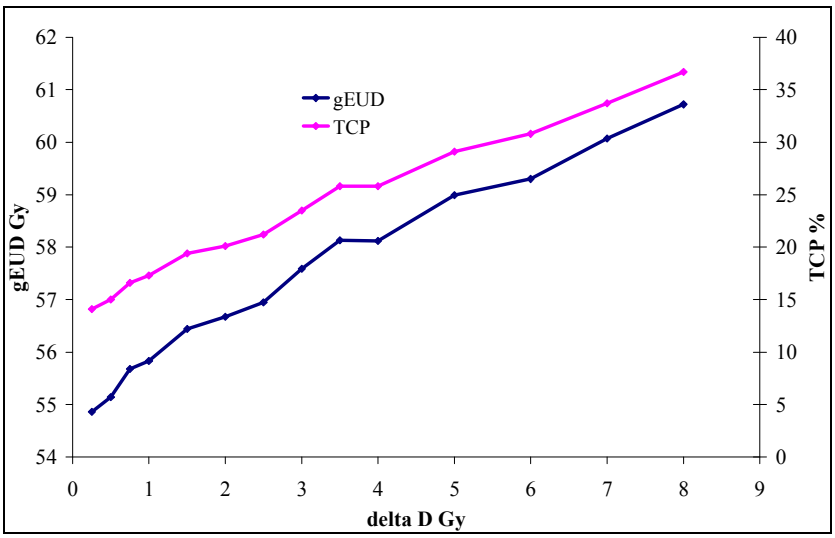

b.

Figure 5.8: Biological indices (gEUD and TCP) as a function of ΔD, illustrating dosimetric and control probability gains for larger values of ΔD. a. ΔR = 1.0 cm. b. ΔR = 0.5 cm.

Intra-structure trade-off

The changing of importance weights of individual structures for the triple boost scenario appeared to have very little impact on dose distributions. *Figure 5.9* shows the middle sized boost (Boost 3) importance weight as a function of conformity with the only notable difference being for the smallest boost (Boost 5), where CN increases with importance weight. This is most likely because the higher dose isodose volume gets smaller as the encompassing lower dose structure takes higher importance. TC values for all structures were 99 % or greater across the range of Boost 3 importance weights

except at values below 0.25 when the TC of boost 3 dropped to 82 %. In all cases, optimal dose distributions were observed when a given importance weight was unity i.e. all structures equally weighted.

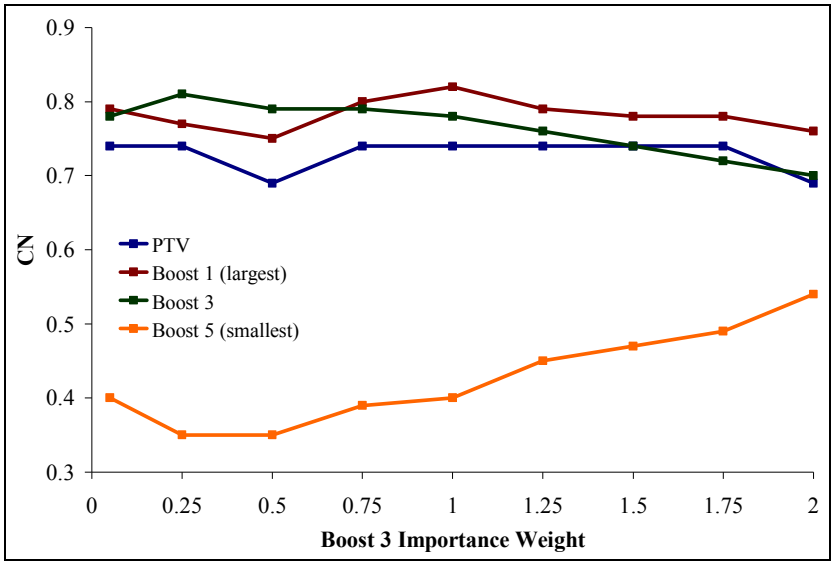

Figure 5.9: Scaling the influence; middle-sized boost (Boost 3) importance weight, given here as a fraction of other structure weightings, as a function of conformity.

Scaling the intra-PTV influence using 'powers' (cost function exponents) had a much greater effect on dose distributions than did using importance weights. *Figure 5.10* illustrates the variation of conformity as a function of middle boost (Boost 3) power. For each structure power variation, TC was less than 100 % for the largest volume contained within the given test volume e.g. Boost 5 (smallest) for Boost 3 experiment, as the power was increased from the default value of 2. The power tool therefore, is important for getting structure doses into acceptable ranges and importance weightings useful for tuning dose distributions once adequate coverage has been achieved.

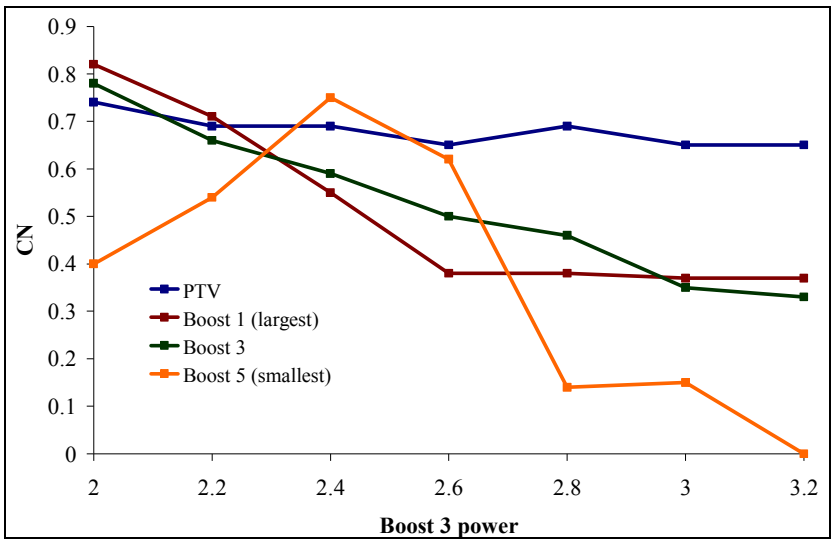

Figure 5.10: Scaling the influence; Boost 3 power as a function of conformity.

The power experiments caused the greatest fluctuations in optimisation scores of all the tests. *Figure 5.11* demonstrates this for Boost 3 power variation. As the power is increased initially the cost function value minimises to a more optimum value. Above 2.8 however, the value surges to its maximum before decreasing again. The number of iterations is also a minimum when the power is 2.8, suggesting that the optimiser finds it easiest to meet the set goals at this point. Normally, it would be expected that when the power is low, the cost function might increase as the calculation struggles to achieve the trade-off and once it is set very high this goal takes priority over everything else and the optimiser quickly converges to a low cost function value. This was only observed for Boost 5 (smallest). This could be due to the fact that Boost 5 has no additional volumes contained within it to be included in the optimisation therefore producing more predictable results. The other volumes all have structures contained within them which amplifies trade-off issues.

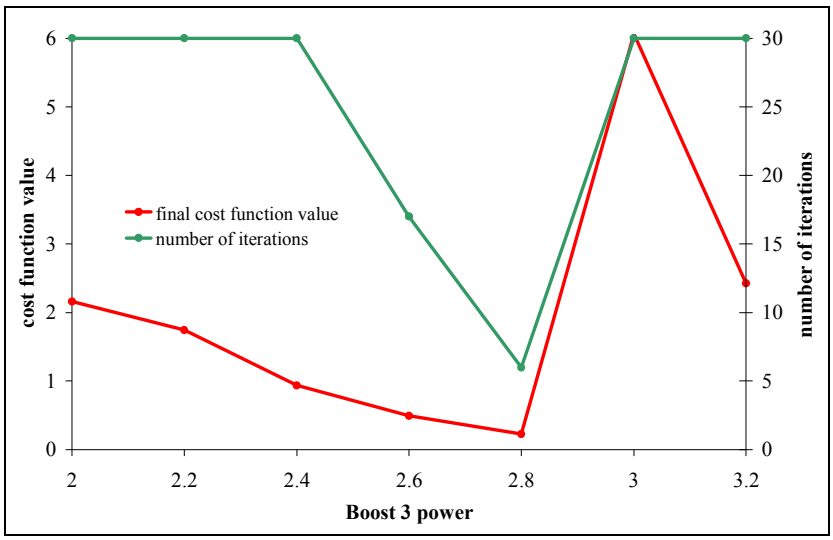

Figure 5.11: Optimisation scores at convergence as a function of Boost 3 power; number of iterations and normalised cost function value.

Delivery resolution

The use of variable rectangular beamlets did not have a significant impact on dose distributions. *Figure 5.12a* shows that the optimiser found it more difficult to fulfil dose objectives with larger (> 10 x 5 mm^2) beamlets, represented by a higher (less optimal) cost function value at convergence. For leaf step increments between 3 and 10 mm however, the optimisation score seemed to be unaffected. The larger beamlet optimisations converged more quickly, confirming an increase in calculation time for smaller step increments due to fewer optimisation variables (beamlets). All plans converged in 27-30 iterations. The effect on individual target dose coverage values is displayed in *Figure 5.12b*, decreasing as larger beamlets are used. This is to be expected, as the optimiser has less control over how the dose is deposited and therefore cannot meet the dose objectives set on such small scales as the six boost model. The most important result to note is that reducing the step increment below 5 mm does not result in better dose distributions i.e. square 5 x 5 mm^2 beamlets are indeed as effective as 3 x 5 mm^2 beamlets.

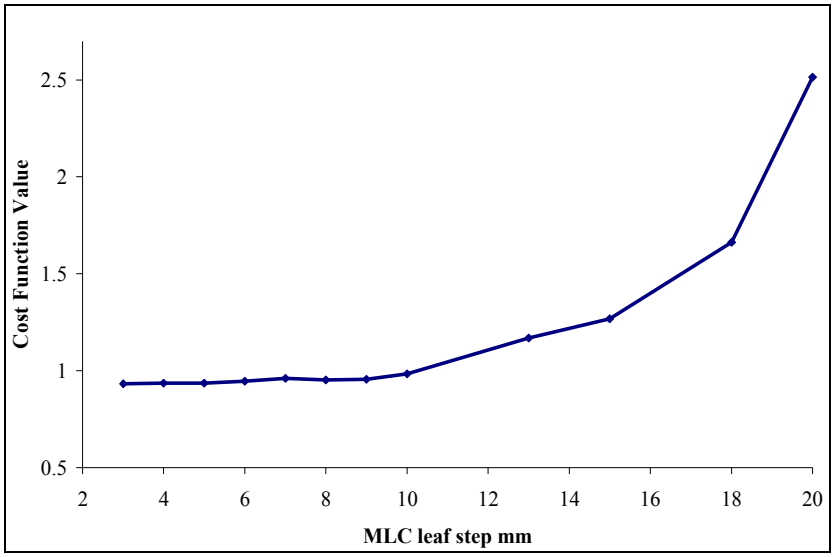

a.

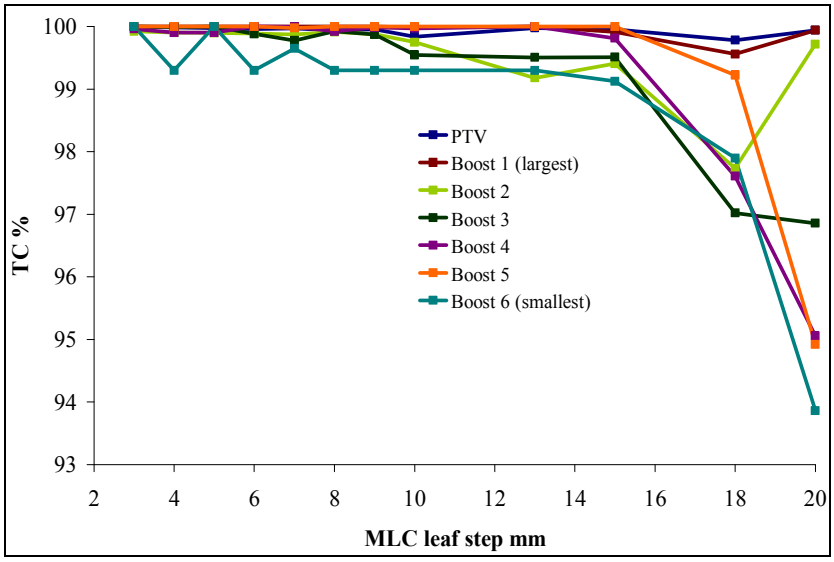

b.

Figure 5.12: Increasing dose deposition resolution with rectangular beamlets; a. optimisation scores at convergence and b. 95 % dose target coverage, TC, both as a function of MLC leaf step increment.

It was not possible to investigate leaf step increments any smaller than 3 mm as this is the minimum value afforded by XiO® and anyway, the smaller the beamlet size, the longer the optimisation, calculation and delivery time due to more variables to optimise and more segments to deliver, which may offset any small dosimetric benefit. It has been shown elsewhere that smaller physical MLC leaf widths of 2-5 mm can be beneficial, but any smaller do not further improve results due to the range of electrons [90]. Also, smaller MLCs implicate accuracy and suitability issues of a given dose

calculation engine e.g. if the step increment is too small, assuming dose is typically calculated on a 2 x 2 x 2 mm^3 grid, some beamlets may have no dose voxels, and therefore no intensity.

5.3.2. Cubic Pseudo-Voxels

Two models for dual-dose boost

The cost function value at which the eight-voxel system converged increased as ΔD_2 increased in an exponential manner (*Figure 5.13*). This implies that the optimiser finds it more difficult to satisfy all the objectives when a bigger dose differential is required on this sort of scale. The shape of the plot is similar when voxels 2-8 are replaced with voxel_combined, but the gradient of the curve is shallower (less change), although the cost is higher over the range of ΔD_2. This suggests that when ΔD_2 is less than 9 Gy, the optimiser find it easier to meet the dose objectives in terms of 8 individual volumes but above this level the lines on the graph intersect and the two volume scenario results in a lower cost. All plans converged in 30 iterations.

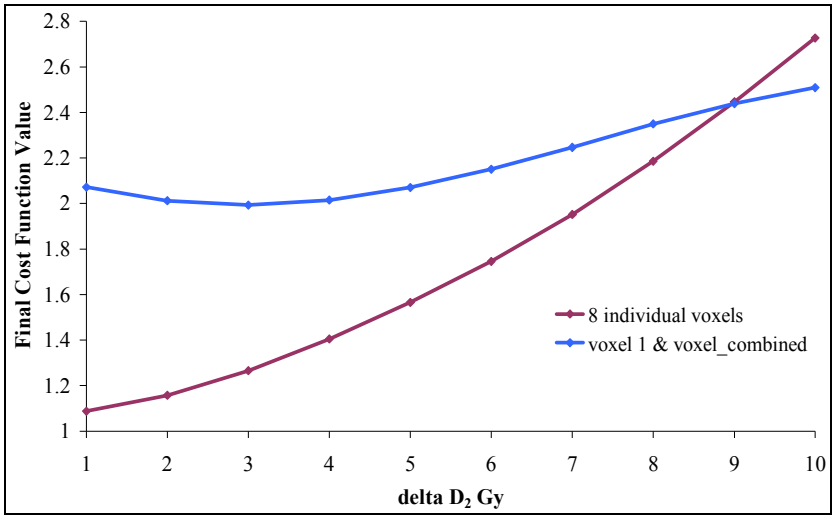

Figure 5.13: Normalised cost function values at convergence as a function of ΔD_2; comparison of 8 individual voxels and voxel 1 with voxel_combined.

Figure 5.14 shows the dosimetric evaluation of voxel 1 for the 7 voxel and voxel_combined experiments. Coverage of voxel 1 by 95 % of its prescribed dose, TC,

starts to fall below 100 % at very low values of ΔD_2 (*Figure 5.14a*). TC is worse for the eight voxel set up than using voxel_combined over the entire range of ΔD_2. Conformity appears to be the same for both experiments, with CN increasing as ΔD_2 increases (*Figure 5.14b*). This occurs as the isodose volume gets smaller, which eventually leads to sub-optimal coverage as previously discussed.

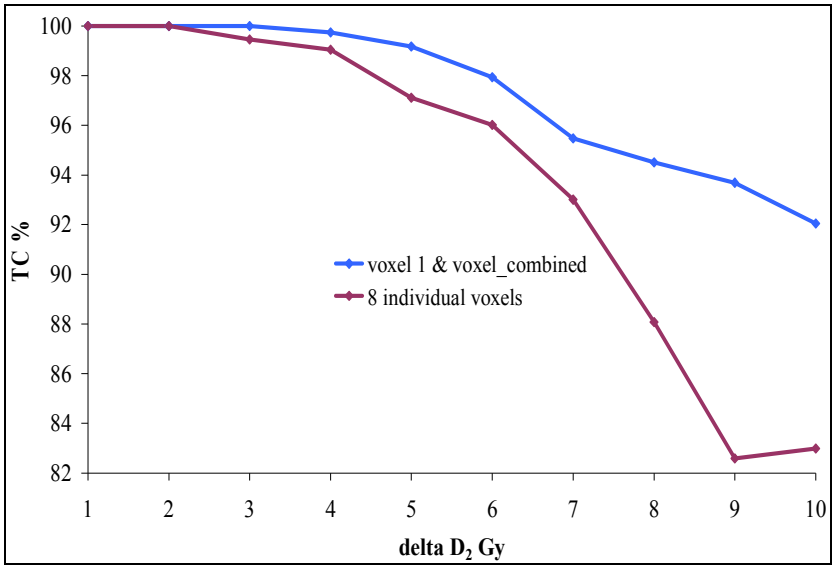

a.

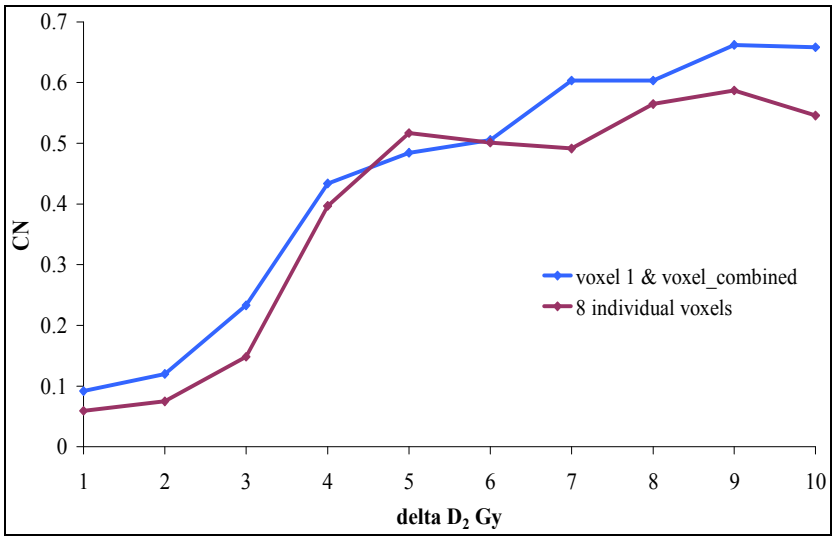

b.

Figure 5.14: Pseudo voxel 1 dosimetric evaluation as a function of ΔD_2 for 8 individual voxels and voxel 1 with voxel_combined; a. dose coverage and b. conformity.

Dose evaluation of volumes other than voxel 1 showed quite the opposite. When 8 individual voxels were used, coverage by the 95 % dose level, TC, was 100 % for most voxels over the entire range of ΔD_2. The exceptions were voxels 2 and 3, the TC of

which dropped to a minimum of 99.6 % for both. When the plan was optimised on voxel_combined however, dose coverage of the larger volume was worse, decreasing to 89.4 % at ΔD_2 = 9 Gy.

The volume of each structure receiving its individual prescription dose does not appear to show any kind of trend (*Figure 5.15*). Fluctuations are notable however, with voxel 1 having the most constant prescription isodose volume over the range of ΔD_2. The volume of voxel_combined receiving the prescription dose shows greater stability than when the seven individual voxels were used. Interestingly, the volume of the PTV receiving its prescription dose (55 Gy) clearly dips at ΔD_2 = 6 Gy in both experiments.

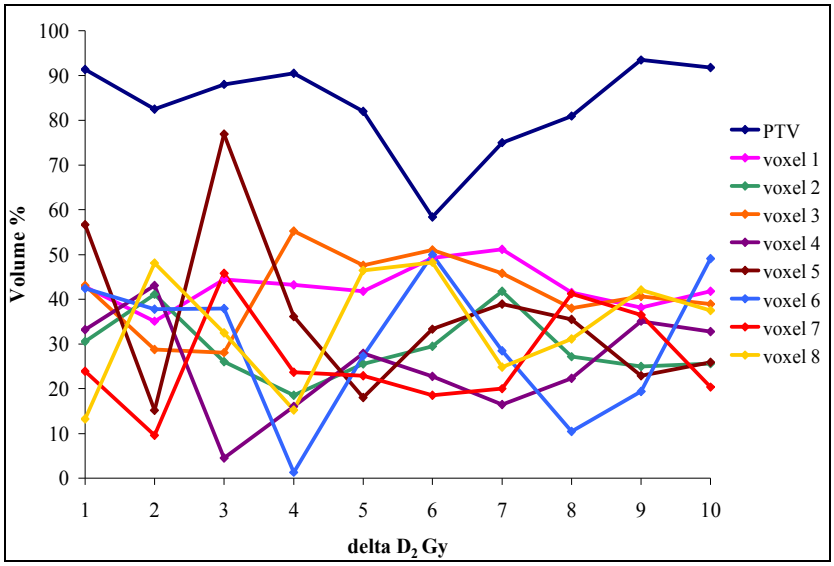

a.

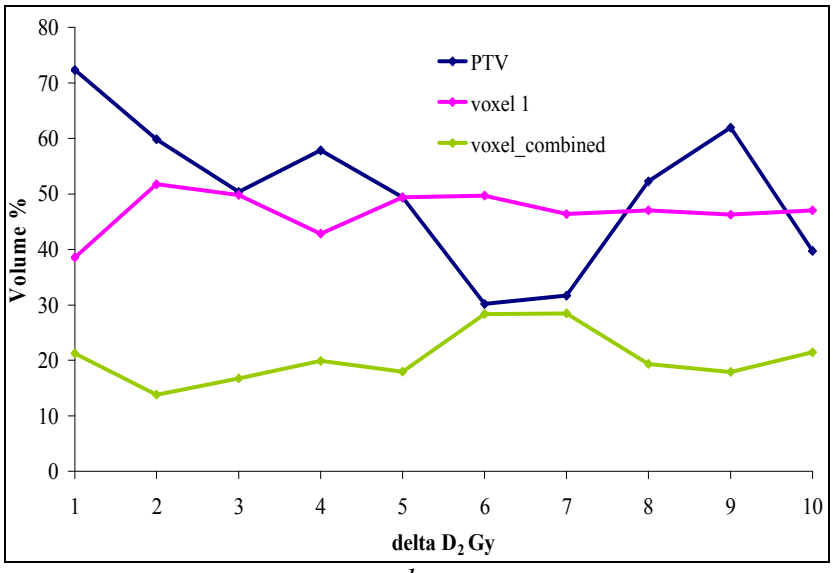

b.

Figure 5.15: Prescription dose coverage, as a function of ΔD_2 for a. 8 individual voxels and b. voxel 1 & voxel_combined.

The homogeneity score, HS, of all pseudo voxels degrades as ΔD_2 increases, with voxel 7 and 8 remaining the most stable (*Figure 5.16*). The effect is worse for voxels 2, 3 and 5 since these all have faces adjoining voxel 1, where the highest dose is being deposited. Voxel 1 showed the biggest decrease in homogeneity of all the voxels. HS is lowest for the PTV where poor homogeneity is expected by the nature of the experiment i.e. the higher dose boosts are within the PTV and these regions are therefore included in the PTV dose calculations. When the plan was optimised on only voxel 1 and

voxel_combined, the homogeneity of voxel 1 improved slightly at higher values of ΔD_2. HS of voxel_combined was lower overall than the individual voxels 2-8. Homogeneity analysis also showed that the difference between maximum and minimum doses could be as much as 3 Gy within a particular voxel. This suggests that a dose differential of less than 3 Gy is meaningless with the intention of dose escalation.

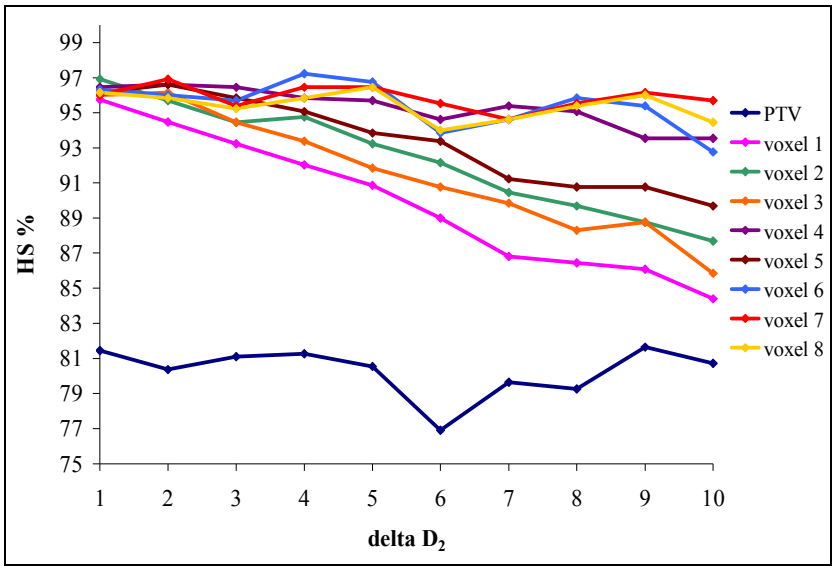

a.

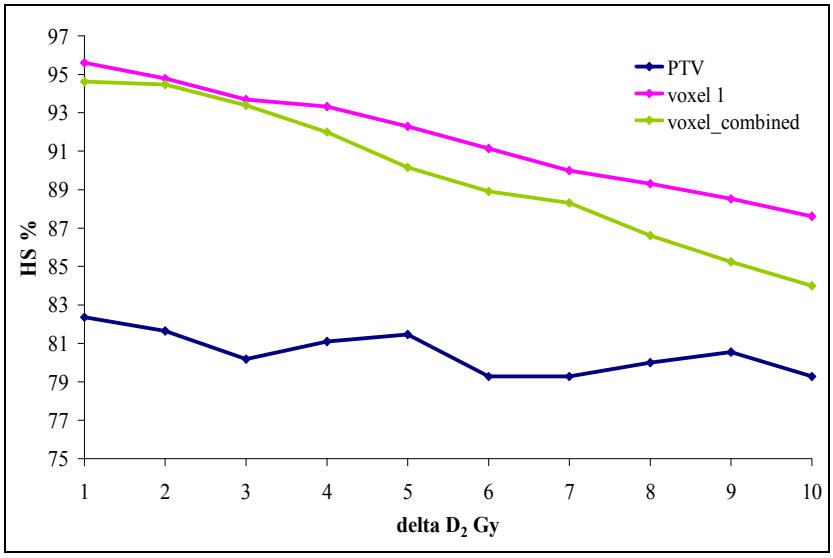

b.

Figure 5.16: Homogeneity as a function of ΔD_2 for a. 8 individual voxels and b. voxel 1 with voxel_combined.

Based on these results, a fixed value of ΔD_2 = 5 Gy was selected to investigate the role of the XiO® 'power' function on the voxel model. This maximised achievable dose

coverage, conformity and homogeneity whilst recognising that a value less than this would not be above ordinary inhomogeneities seen in IMRT plans. Based on *Figure 5.14a*, the power experiments were carried out on voxel 1 with voxel_combined rather than the seven individual voxels as this set up appeared to maximise target coverage.

Intra-structure trade-off

Optimisation scores varied dramatically as a function of both voxel 1 power and voxel_combined power (*Figure 5.17*). The optimisation converged at the maximum 30 iterations at low power values but converged in fewer iterations at powers of around 2.3 in both cases. Whilst the number of iterations continued to decrease when voxel_combined power was increased, at a voxel 1 power of 2.5 the number of iterations rose back up to 30, staying there across the remaining values. The cost function values at convergence was very low at small powers, decreasing to practically zero when the voxel_combined power was 3.0, implying that the optimiser was able to meet the set goals very closely. In the case of voxel 1 however, the cost soared to more than 25 when the power was 2.5 and decreased in an exponential manner to a near-zero value when the power was 4. This suggests that the optimiser could meet the requirements when voxel 1 was assigned either a very low or very high power weighting, but struggled for values in between.

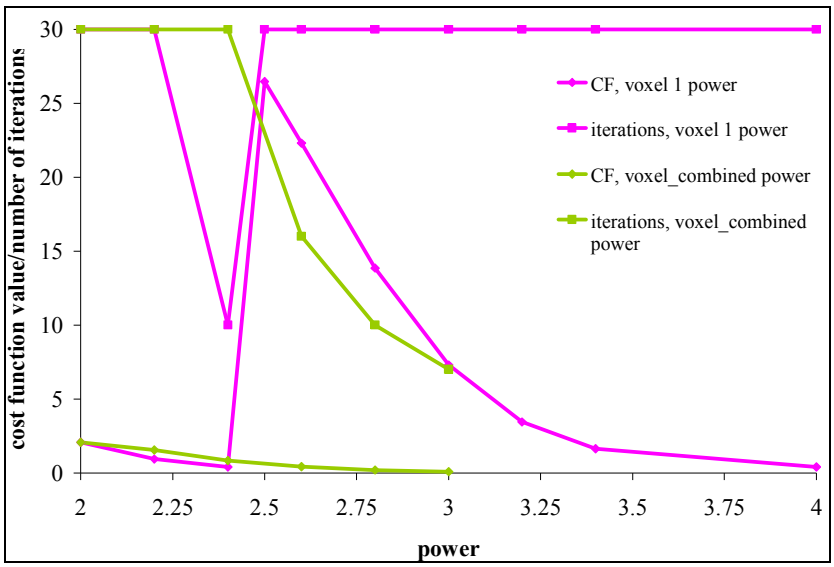

Figure 5.17: Optimisation results for changing the power of voxel 1 and voxel_combined independently.

Increasing the power of voxel 1 ensured that 100 % of it was always covered by at least 95 % of the prescription dose (*Figure 5.18a*). As the power of voxel 1 was increased, even though TC of voxel 1 was already 100 %, TC of voxel_combined continued to decrease. When the power of voxel 1 was set to 4.0, coverage of voxel_combined was just over 60 %. As the power of voxel 1 was increased, CN of both volumes decreased almost simultaneously (*Figure 5.18b*).

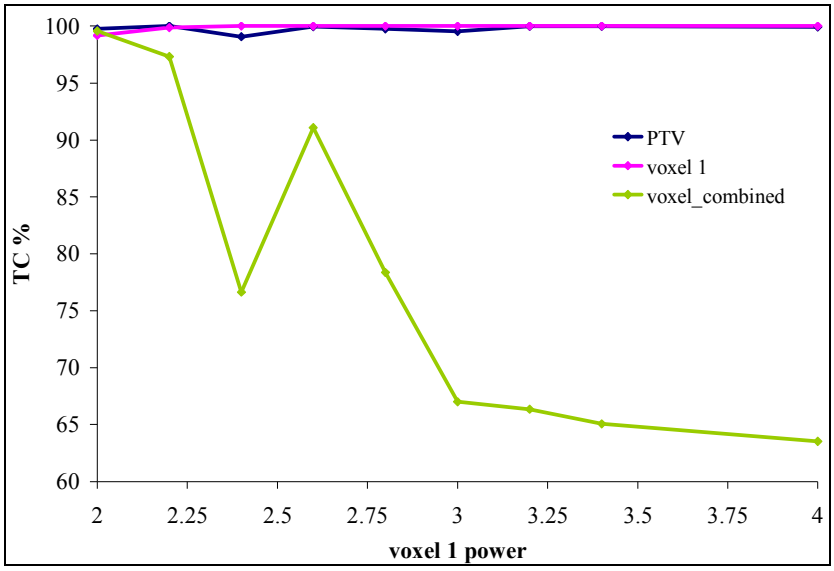

a.

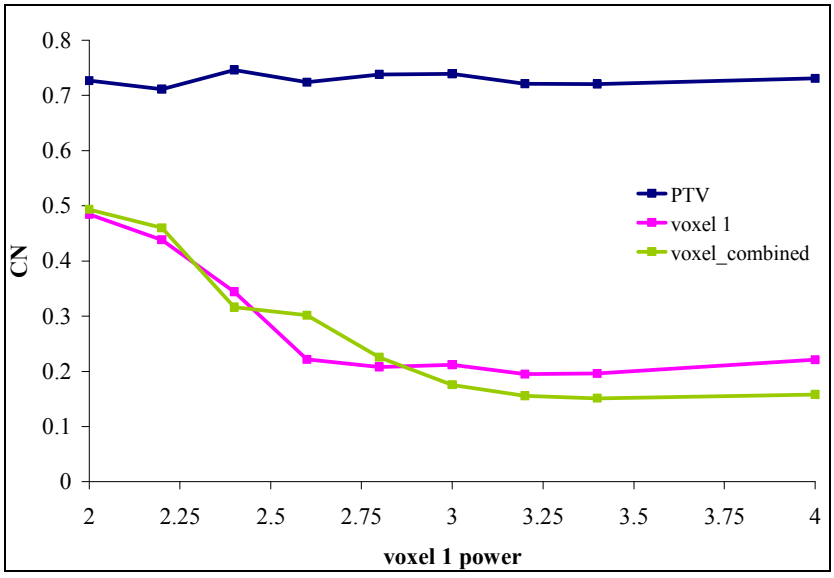

b.

Figure 5.18: Dosimetric evaluation as a function of voxel 1 power; a. dose coverage and b. conformity.

Increasing the power of voxel_combined enforced a constant TC of 100 % on voxel_combined, but the coverage of voxel 1 dropped off quickly, reaching zero when voxel_combined power was just 3.0 (*Figure 5.19*). The conformity of voxel_combined worsened slightly with a power increase as the optimiser deposited more dose to ensure full target coverage.

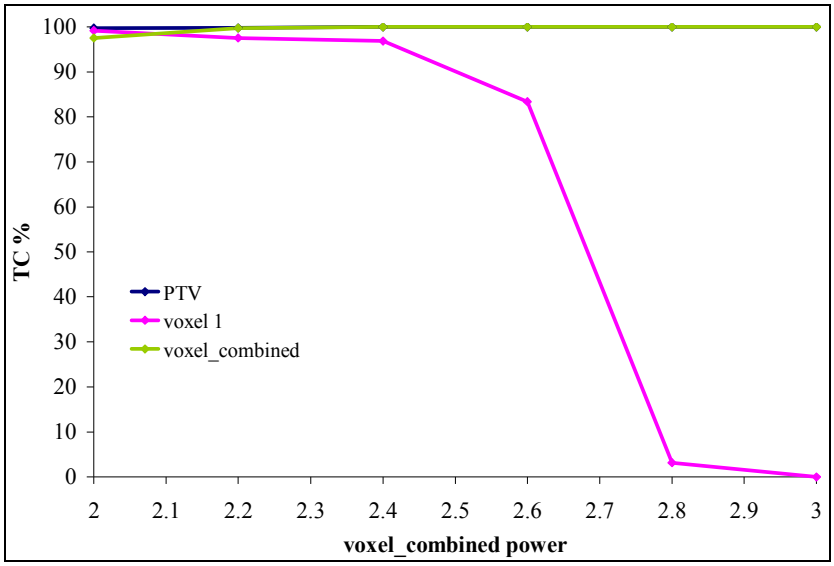

a.

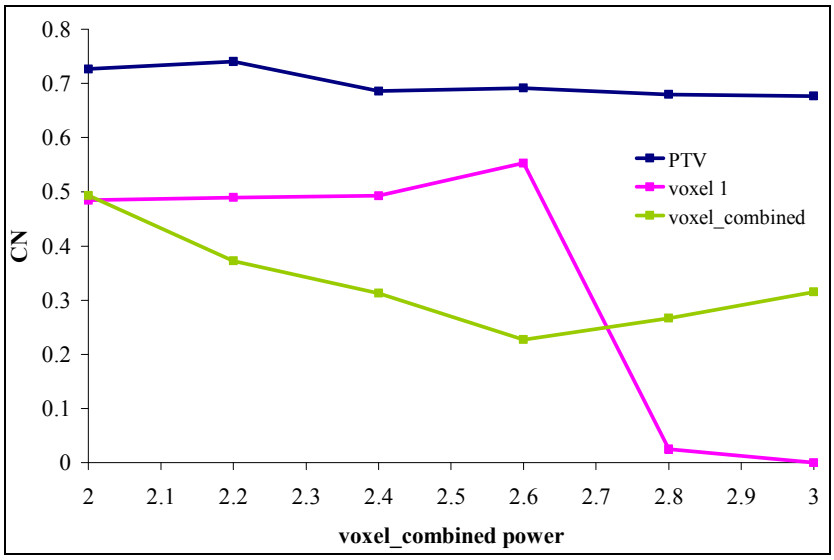

b.

Figure 5.19: Dosimetric evaluation as a function of voxel_combined power; a. dose coverage and b. conformity.

Figure 5.18 and *Figure 5.19* show how changing the power of voxel_combined had a much more dramatic effect on dose distributions than voxel 1 power. This is most probably due to the difference in prescribed dose. Voxel 1 is prescribed a high dose and so when it is given a lot of importance, the planning system puts more high dose in to ensure optimal coverage. By the physical nature of dose deposition, surrounding areas will also receive some of this high dose hence coverage of voxel_combined by its lower prescribed dose does fall to zero. Conversely, when the influence of the lower dose voxel_combined is set very high, the optimiser will simply avoid putting any high dose

in at all and so voxel 1 fails to receive any of its required dose. Cost functions in XiO® are volume normalised to make objectives independent of structure size, thus the difference between voxel 1 and voxel_combined (1:7 ratio) should not have affected the results.

Increasing the power of a given structure improved homogeneity within it (*Figure 5.20*). The increase in HS for voxel 1 was at the expense of a decreasing HS for voxel_combined but the two have similar homogeneity when voxel_combined power is increased. This is because the two volumes are receiving similar dose levels at this point.

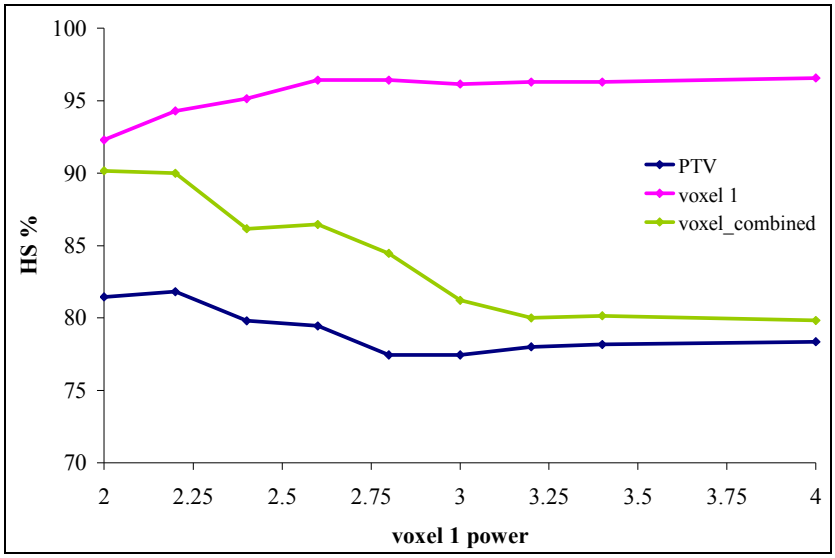

a.

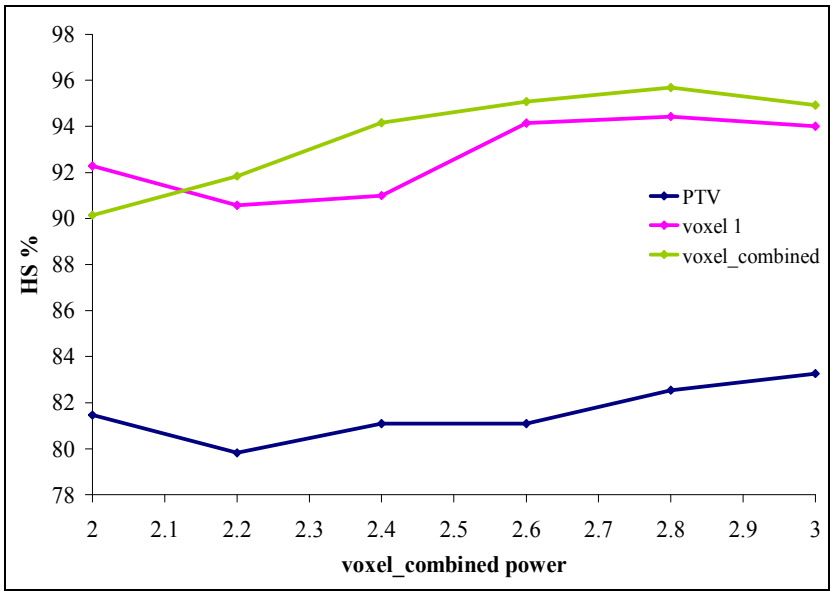

b.

Figure 5.20: Homogeneity as a function of individual structure power; a. voxel 1 power and b. voxel_combined power.

Voxel-by-voxel dose prescribing

Delivering an incremental boost, ΔD, to eight separate voxels proved much harder than the previous two-boost scenario. Even at ΔD = 1 Gy, the voxels were not fully covered by 95 % of their prescription doses and up to ΔD = 5 Gy, TC of some of the voxels fell to as little as 70 % (Figure 5.21). Even at only ΔD = 3 Gy, the highest TC was under 90 %. This suggests that optimising on this resolution may not be feasible when the required dose to each 'voxel' is different. Previous homogeneity analysis has already

shown that the difference between maximum and minimum doses within a structure can be as much as 3 Gy under normal circumstances and so it is fruitless to consider a differential prescription any less than this. Interestingly, although coverage by the 95 % isodose levels decreases so remarkably, the volume of the structures receiving the prescription level remains fairly constant (*Figure 5.22*). *Figure 5.21b* shows poor conformity for all voxels, with the highest dose regions exhibiting the best conformity. It is well known that smaller volumes have worse conformity than larger volumes and it seems this becomes problematic when considering optimisation on an intra-structural scale. Van't Riet *et al.* defined CN > 0.6 as 'conformal' therapy and the CN values of most voxels barely reach 0.2 over the range of ΔD [67].

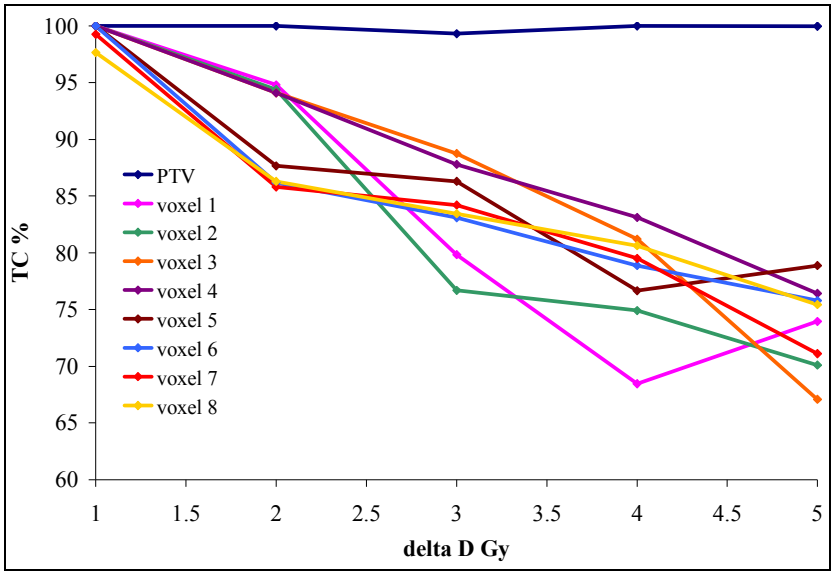

a.

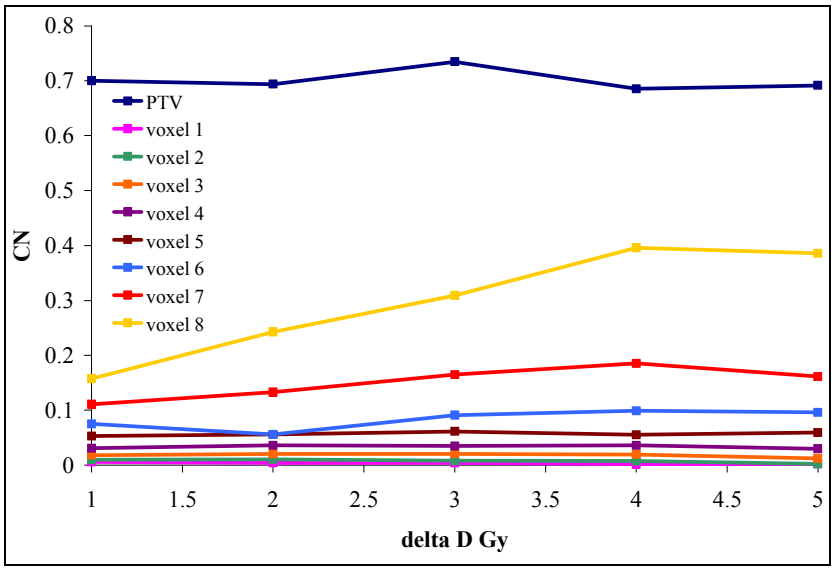

b.

Figure 5.21: Prescribing a different dose to every voxel; a. 95 % dose coverage and b. conformity as a function of ΔD.

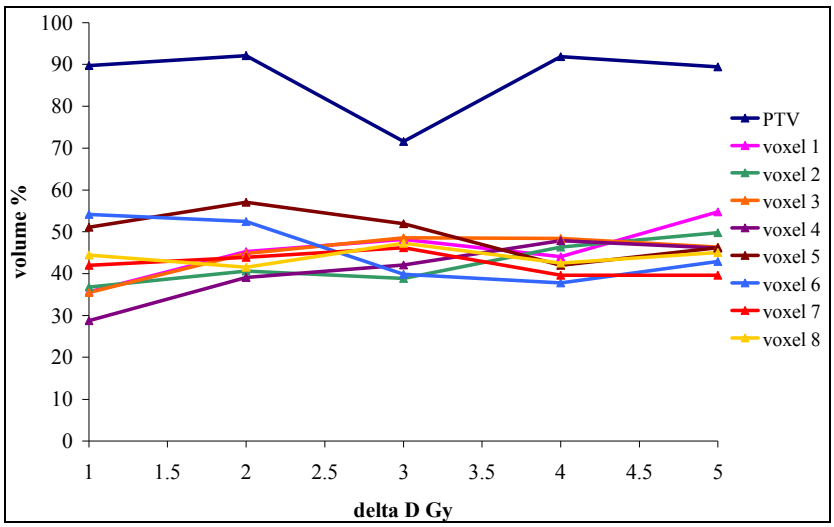

Figure 5.22: Prescribing a different dose to every voxel; PTV and pseudo voxels receiving appropriate prescription dose as a function of ΔD.

Homogeneity in each voxel decreases almost linearly with ΔD (*Figure 5.23*), with the inhomogeneity (difference between the minimum and maximum dose) in all voxels rising to over 15 Gy at ΔD = 5 Gy.

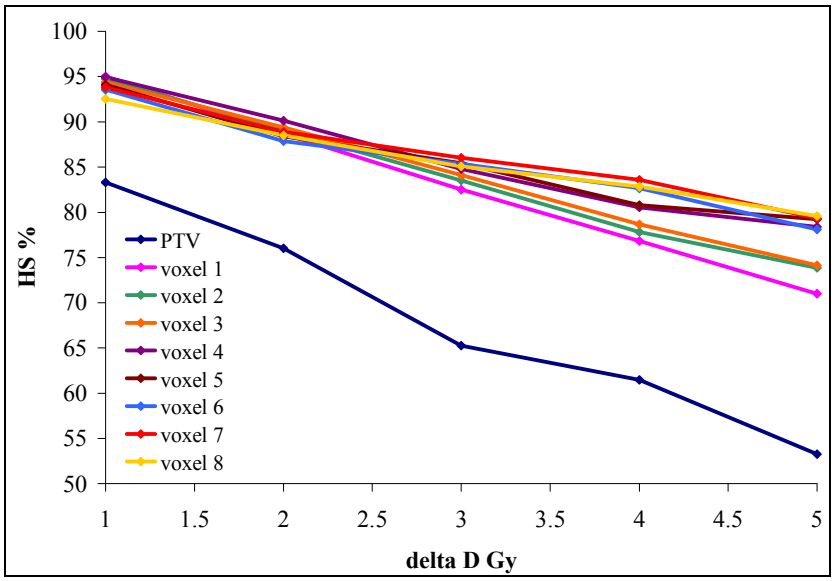

Figure 5.23: Prescribing a different dose to every voxel; homogeneity as a function of ΔD.

5.4. Conclusions

Owing to the detailed and specialised nature of the investigations in this chapter, a summary of the work is given below:

- Simulation of three pseudo-boost anatomies; three concentric spheres corresponding to ΔR=1.0 cm, six concentric spheres corresponding to ΔR=0.5 cm and eight

contiguous cubic boosts occupying 0.9 x 0.9 x 0.9 cm^3, with each of the three scenarios being contained within an identical, authentic PTV.

- ΔD defined as some constant dose increment between specified pseudo-boost volumes. Treatment plans created over various ranges of ΔD for all three target scenarios and the greatest achievable value denoted ΔD_{max}.
- ΔD_1 defined as the boost dose increment between the PTV and seven out of eight of the cubic pseudo-boosts.
- ΔD_2 defined as the boost dose difference between these seven boosts and the remaining one of eight cubic pseudo-boosts.
- Intra-PTV influence characterised by varying importance weights and powers (objective function exponents) of the individual pseudo-boosts.
- Plans evaluated based on multiple target coverage and conformity, optimisation scores, gEUD and homogeneity where appropriate.
- Dose deposition resolution investigated by generation of rectangular beamlets of various dimensions.

The experiments show that regional dose can be controlled at 10 mm and 5 mm resolution, where an incremental dose differential of 5 and 3 Gy respectively is readily achievable for a boost-within-boost scenario. Boosting small, neighbouring volumes is more difficult however. To put this in context, it's important to appreciate that the various imaging resolutions, margins, calculation and delivery times, delivery device and patient immobilisation methods all influence what is achievable, accurate and therefore valuable in terms of precision dose sculpting.

Deliverability of the planned dose distributions here was not verified. It may be useful to look at what is physically deliverable in terms of dose deposition by checking

delivered dose against predicted dose with film and/or gel dosimetry. IMRT dose distributions are built up from combinations of smaller fields and the SMB technique utilises many small fields, which are difficult to accurately QA.

Before the implementation of SMB can be realised, it's important to understand exactly what is achievable in terms of dose deposition. This study has shown that regional dose can be controlled at 10 mm and 5 mm resolution for certain anatomies and that importance scaling is also effective on this scale. Considering that resolution of most functional imaging techniques is no better than this - in fact typical voxel size for an MRSI examination is 10 x 10 x 10 mm - it is reasonable to conclude that functional imaging for defining regions that would benefit from dose escalation is a viable pursuit. Prescribing a different dose to each neighbouring voxel in a functional dataset may not be realistic however and larger boost volumes may therefore be unavoidable in SMB plans.

Chapter 6 - Significance of Image Resolution

6.1. Introduction

In order to optimise the prognosis for brain tumour patients it is essential that any therapy is directed to the active, proliferating tumour region. Gliomas in particular, exhibit low radiation sensitivity and are associated with a poor prognosis, with the majority of failures occurring within the irradiated volume. It has been shown elsewhere that dose escalation may improve the results of radiotherapy for these patients [8]. Even better outcomes may be expected if this is coupled with improved tumour delineation. Intensity-modulated radiation therapy (IMRT) provides unparalleled means to deliver highly conformal treatments and there is great interest in its capability to deliver precise concomitant boosts. The selective multiple boosting (SMB) approach may allow for further dose escalation as delivering the higher doses to only a small volume rather than the entire planning volume will reduce toxicities in surrounding tissue.

As well as being radioresistant, gliomas are particularly infiltrative and as such may be poorly defined on conventional images such as MR and CT. In spite of this, accurate delineation is essential in planning focussed treatments such as IMRT, particularly if dose escalation is to be considered. Functional imaging modalities can provide the necessary insight to identify regions that would benefit most from escalated dose and studies to implement this information in radiotherapy planning are already underway [17, 19-22, 51, 52].

Multi-voxel proton MR spectroscopic imaging (MRSI) is now a well-established, non-invasive technique for evaluating the metabolic status of tissue and can be carried out on a clinical MRI scanner at the time of imaging. The study of appropriate metabolites can indicate tumour presence i.e. facilitate target delineation, as well as characterising the tumour cells; in terms of clonogen density, for example. MR spectra may be acquired from a three-dimensional grid of voxels simultaneously and statistical techniques applied to characterise abnormality in terms of the concentrations of various intracellular metabolites. Adult brain tumours in particular have received much attention from this technique in recent years [49, 53, 91, 92], henceforth quality and reproducibility of data has improved. Tumours generally exhibit decreased *N*-acetylaspartate (NAA) accompanied by an increase in choline (Cho) and as such a ratio of the two peaks may be used as an indicator of abnormality. Fourier transform reconstruction and apodisation analysis techniques of the MRSI data are used to generate quantitative maps of the spatial distribution of these metabolites and can be represented as intensity contrasts on the MR image dataset for anatomic reference; the metabolite map.

The resolution of functional imaging techniques varies greatly. The typical resolution in PET for example, is 5 mm, whereas pixel size in fMRI is as good as 1 mm [25, 58]. MRSI has relatively low spatial resolution; typically 10 x 10 x 10 mm^3 voxels. If IMRT boost doses are to be planned to these spectral abnormalities, it is crucial that the volumes are known to be both complete and correct for accurate depiction of dose requirements. It was shown in Chapter 5 that local dose can be controlled at 5 mm resolution and so target variations on this scale could affect where higher doses are deposited and the efficacy of these boosts. To reduce the total error i.e. the positional

and volumetric uncertainty in a high-dose volume, there is a need to address all possible sources of error. In conventional radiotherapy planning, margins are added to target volumes for motion and set-up uncertainties, and so these must next be extended to encompass uncertainties associated with a given advanced imaging technique if SMB is to be realised.

Because of its inherently low resolution, raw MRSI acquisition data is often up-sampled so that metabolite maps overlaid on anatomic images correspond to the finer resolution of the image dataset (*Figure 6.1*). Different approaches lead to a range of volumes and it is usual for the smoothed data to be published without interpolation methodology. The aim of this study was to investigate volume variability of spectra metabolite-map interpolation and the direct impact of this on IMRT boost volumes and dosimetry.

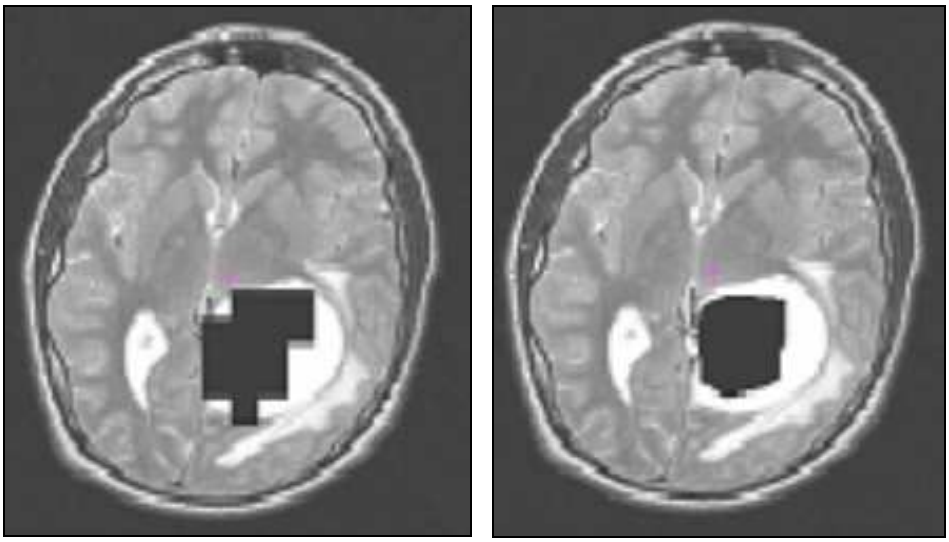

a. *b.*

Figure 6.1: Metabolite map interpolation. Example of a. non-interpolated metabolite map (in-plane resolution: 10 x 10 mm^2) and b. 'smooth' interpolated metabolite map (in-plane resolution: 1 x 1 mm^2).

6.2. Methods

A patient with a diagnosed astrocytoma underwent an MRSI examination directly after an MR imaging study on a clinical 1.5 Tesla General Electric[§] Signa scanner equipped with a single channel quadrature (birdcage) RF head coil. Whole-brain axial T_2-weighted contiguous slices of 3 mm thickness were acquired followed by a multivoxel 3D-MRSI scan using a PRESS sequence with brain-optimised spectral-spatial excitation pulses with TE/TR: 90/1000 ms [93]. The MRSI field of view (FOV) was 12 x 12 x 8 cm^3, nominal spectroscopic resolution (voxel dimensions) was 1 x 1 x 1 cm^3 and the volume of interest (VOI) was 7 x 7 x 5 cm^3 in the x, y and z planes respectively. Very selective spatial saturation (VSS) pulses and spatial selective saturation bands were employed to optimise definition of the spectral volume and minimise lipid contamination. This protocol resulted in a scan time of nineteen minutes.

6.2.1. MRSI analysis and interpolation

Spectral data was processed off-line at the CMRI, using the Spectroscopy Analysis GE™ (SAGE) software and tables of voxel-by-voxel peak amplitudes of Cho and NAA generated. From these peak heights, T_2-weighted datasets and spectra header files, three metabolite maps were generated with in-house software SIRAMAS [94]: a nearest neighbour interpolation through-plane (10 x 10 x 3 mm^3) followed by a linear interpolation in-plane (1 x 1 x 3 mm^3) and cubic spline in-plane (1 x 1 x 3 mm^3), the resolution of the latter two being defined by the acquired anatomical image data sets. A parameter map of the MRSI data, depicting a binary abnormality level i.e. positive (tumour) or negative (non-tumour) was produced whereby an NAA/Cho ratio of < 0.6 was assigned tumour status and a ratio ≥ 0.6 designated non-tumour. This threshold is derived from a study of healthy volunteers at the CMRI [50]. Fusion of this to the

[§] Milwaukee, WI, USA

anatomical images provided useful data visualisation. Once the two data sets were fused, voxels in the anatomy images that were designated tumour-like i.e. NAA/Cho ratio < 0.6, in the MRSI analysis, were altered to be coloured black.

6.2.2. MRSI contouring and registration

The fused MRSI parameter maps and anatomy MR images were stored as standard DICOM MR objects and exported to a CMS[§] XiO® planning station. Here, each of the three MRSI datasets was registered to the planning CT images in order to allow their use in the treatment planning process. This was done using the CMS Focal Fusion software, which provides an automatic registration process by maximising mutual information [95, 96]. The program creates an internal file with a transformation matrix to map MRI data to the CT co-ordinate frame of reference. This was used to precisely standardise the registration process by manually inserting the initial transformation matrix into each MRSI-CT registration, avoiding alignment error and operator inconsistencies. This was possible given the knowledge that each of the MRSI datasets had common coordinate systems as they were generated from a single MRSI investigation.

Anatomical volumes; the gross target volume (GTV), clinical target volume (CTV) and the planning target volume (PTV) were contoured by a consultant clinical oncologist from the conventional MRI and CT images alone. The oncologist was not involved in defining the MRSI volumes, these were delineated by the author on the MRSI-CT registered dataset with an automatic (threshold based) contouring tool (*Figure 6.2*). This final registration proved to introduce through-plane ambiguities into the metabolite map representation, as it involved registering images from datasets defined on frames of

[§] Computerised Medical Systems, Inc., St. Louis, MO

reference with different grid resolutions and with slightly offset angular orientation. This induced a gray scale effect on the MRSI-positive volume, creating uncertainty in what exactly should be contoured i.e. what was anatomically or pathologically correct. To reflect this, three different contours were drawn to represent the raw spectra data: 'rawThresh' using automatic thresholding, 'rawAll', a manually drawn contour around all MRSI-positive regions and 'rawDef', a manually drawn contour around only the definitely MRSI-positive data i.e. lowest intensity regions (*Figure 6.2b*). This resulted in three representations of the raw data and two interpolations, totalling five different MRSI volumes. A 3D automatic margin tool (rolling-ball) was used to 'grow' the smallest MRSI volume (linear) by 3 mm, as suggested by Graves *et al.* and this volume termed linear+margin [51]. Furthermore, the linear+margin volume was then grown incrementally in order to wholly cover each of the larger volumes in space.

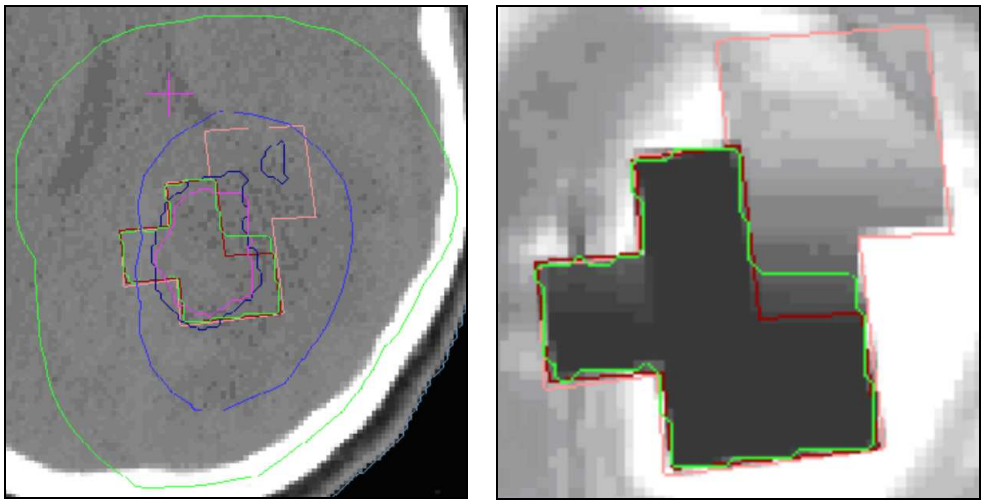

a. *b.*

Figure 6.2: MRSI contours. a. rawThresh (green), rawAll (pink), rawDef (red), linear (purple), cubic (blue) and planning contours; PTV (green), GTV (blue). b. three contours to represent 'raw' MRSI data owing to uncertainties from MRI-CT registration; rawThresh (green), rawAll (pink) and rawDef (red).

6.2.3. IMRT planning on MRSI contours

IMRT plans were produced to deliver 55 Gy to the planning target volume (PTV), as defined by the oncologist, with a simultaneous boost of 65 Gy to the MRSI volume, corresponding to a boost dose ratio (BDR) of 118 %. A set of plans were developed,

one for each of the boost volumes. A 'control' or standardised set of plan parameters (*Table 6.1*) were established by producing an acceptable plan where the boost volume was defined by the (smooth) linearly interpolated MRSI data.

structure	type	objective	dose cGy	volume %	relative importance	power
MRSI*n*	target	max	6500	0	100	2.0
		goal	6500	100		1.0
		min	6500	100	150	2.0
PTV	target	max	5600	0	100	2.1
		goal	5500	100		1.0
		min	5500	100	130	2.3
patient	OAR	max	5500	0	10	2.0
		dose vol	4000	20	5	2.0
		dose vol	2000	50	5	2.0

Table 6.1: IMRT planning parameters.

The plan utilised 6 coplanar beams (5 equally spaced beams with additional field added to improve an otherwise 'cold' region) with gantry angles[§] of 36˚, 60˚, 108˚, 180˚, 252˚, 324˚. The IMRT parameters were copied to each subsequent treatment plan wherein each 'MRSI volume' was used in turn and the optimisation run for each to generate five separate treatment plans. A plan was also developed for the linear+margin volume to inspect the potential of the optimiser to conform to shape as well as size. For all plans, all relevant OARs (orbits, lenses, optic nerves, optic chiasm, brainstem, pituitary) were kept within standard tolerances [76, 77].

6.2.4. Plan evaluation

Plans were evaluated in terms of visual inspection of isodose lines, dose volume histograms (DVHs), homogeneity, conformity and biological tools within IPEX

[§] IEC convention

(Chapter 4). Target dosimetry was assessed in line with ICRU recommendations i.e. in terms of the prescription dose -5/+7 % [27]. Homogeneity of the MRSI boosts was assessed in terms of both cold (HS_{min}) and hot (HS_{max}) regions with the IPEX tools given by equations (4.14) and (4.15), where a higher index in both cases correlates to a more homogeneous dose distribution. A HI_{min} of less than unity represents the extent of target underdosing and a HI_{max} of less than unity relates to target overdosing. The homogeneity score, HS, given by equation (4.16), gives an overall indication of how closely the dose distribution meets the dose level prescribed i.e. 65 Gy to the MRSI boost volume.

The quantity and quality of 3D dose coverage was evaluated by inspecting the relative volumes of the boost and surrounding tissue receiving 95 % of the prescription dose. A standard conformity index, CI, was reported, given by

$$CI = \frac{V_{95}}{V_t} \tag{6.1}$$

where V_{95} is the volume of tissue irradiated to at least the 95 % level and V_t is the target volume [29]. A value between 1 and 2 is consistent with good conformity, 1 being the perfect result in terms of volume-size correlation. Using the formula given by van't Riet *et al.* (equation (4.1)), the conformation number, CN, was also recorded; based not only on the size but also on the relative position of the 95 % isodose volume with respect to the target volume [67]. To this end, a third volume, $V_{t\cap 95\%}$, the volume of target inside the 95 % isodose volume (i.e. the intersection of the two volumes), was determined. A higher CN indicates a better degree of conformity. Finally, the generalised equivalent uniform dose (gEUD) and corresponding tumour control probability (TCP) were computed within IPEX to highlight any dosimetric or radiobiological benefit afforded by the different volumes.

6.3. Results

6.3.1. Volume differences

In order to achieve best positioning within the constraints of limited resolution and given the angularity afforded by the saturation bands, it was not possible to get spectra from the entire enhancing region; 89 % of the T_2 hyperintense region was covered by the MRSI VOI. The largest MRSI volume arose from the raw data; rawAll, and the smallest was the linearly interpolated volume (*Table 6.2*). The mean MRSI volume was 33.95 cm^3 and the standard deviation over all volumes was ± 4.93 cm^3. The corresponding boost volume ratio (BVR) ranged from 5.88 to 8.63 %, with mean 7.09 %.

MRSI Contour	**Contour Description**	**Volume cm^3**	**BVR %**
§rawAll	No in-plane interpolation, manually drawn contour around 'all' MRSI-positive regions.	41.96	8.63
actualRaw	*Actual volume of raw spectra data extracted from interpolation software.*	*37.00*	*7.61*
§rawThresh	No in plane interpolation, threshold-based contour.	33.99	6.99
§rawDef	No in plane interpolation, manually drawn contour around 'definitely' MRSI-positive data only.	33.17	6.82
Cubic	Cubic interpolation.	32.08	6.60
linear	Linear interpolation.	28.57	5.88
σ		***4.93***	***0.86***

Table 6.2: Volumes covered by MRSI interpolations.
§volumes arising due to CT-MRSI registration issues, not directly from interpolations.

Adding a margin of 3 mm to the smallest volume did surpass the largest volume in size (volume) but did not encompass it spatially (*Figure 6.3*). A maximum 3D isotropic margin of 1.0 cm was required for the linear contour to entirely cover the volume in space occupied by the largest volume but this increased the size (volume) by over 100%.

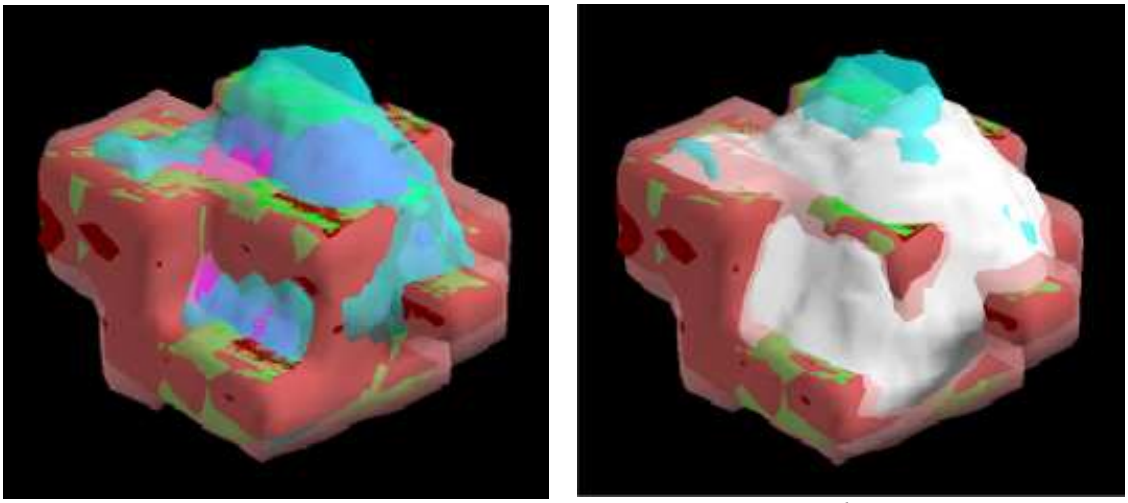

a. *b.*

Figure 6.3: 3D visualisation of MRSI contours. rawAll (pink), rawThresh (green), rawDef (red), cubic (blue), linear (purple) and in b. including the linear volume plus a 3 mm margin (white).

6.3.2. Planning (dosimetric) differences

Variations in MRSI volumes affected how much tissue received higher doses (*Figure 6.4*) and whereabouts this high dose was deposited. As boost volume increased, more of the PTV received the prescription dose but more also received higher doses. This is expected since the planning system treats the boost as a subset of the PTV. Simultaneously, less of the boost volume received 95% of the boost prescription and 100 % of the boost prescription dose. ICRU Report 50 prescribes dose exclusively to the PTV however, and for this study, the PTV dose does not degrade in any significant way in accordance with changes in the boost [27].

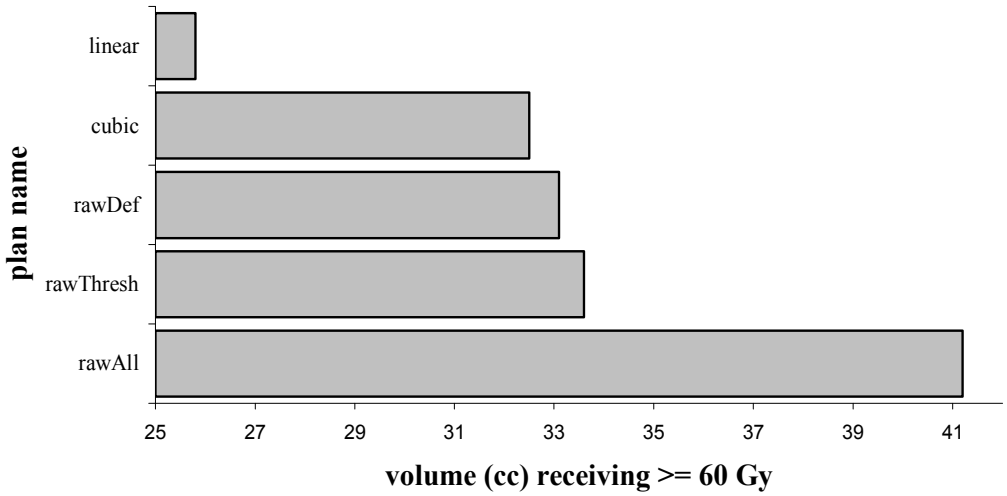

Figure 6.4: Irradiated volume differences. Difference in volume irradiated to greater than or equal to 60 Gy (i.e. greater than PTV $V_{107\%}$) caused by interpolation differences.

Generally, more hot spots (particularly inside boosts) were observed when the boosts were larger. Plan_rawAll (largest MRSI volume) had a 60 Gy hot spot actually outside the PTV and Plan_linear+margin had three 60 Gy hot spots outside the PTV. OAR dosimetry was equally good for all plans and the range of mean dose in the normal brain tissue over all plans was only 0.20 Gy.

Inspection of the DVHs showed that the largest boost volume (rawAll) influenced dose distributions the most. Homogeneity in terms of maximum dose (i.e. hot spots) was very consistent whilst homogeneity in terms of minimum dose (i.e. cold spots) degraded slightly as the boost volume increased. This is most likely due to the physical limitations on maintaining a lower dose in the PTV whilst delivering escalated dose to the larger boost volume. The overall homogeneity, HS, did not show much of a trend, other than the fact that the smallest volume (Plan_linear) scored highest. Interestingly, the score for Plan_linear+margin was the lowest of them all (HS = 84.1 %), indicating a worse homogeneity even than the angular volume, Plan_rawAll, of the same size (*Table 6.3*).

Plan Name	HS %	CI	CN
Plan_rawAll	87.1	1.16	0.762
Plan_rawThresh	86.5	1.27	0.742
Plan_rawDef	86.5	1.31	0.721
Plan_cubic	87.9	1.20	0.791
Plan_linear	88.9	1.31	0.749
σ	*0.9*	*0.06*	*0.023*
Plan_linear+margin	84.1	1.14	0.837

Table 6.3: Homogeneity and conformity analysis. Volumes ranked in decreasing size (top to bottom).

Quantitative examination of plan conformity did not reveal any notable variability of quality with changes in volume, the standard deviation of CN over the range of volumes being only ± 0.02 (*Table 6.3*). The exception was the 'smooth' large volume, Plan_linear+margin, which exhibited far better conformity than the others, with a CN of 0.837, suggesting the optimiser can more easily conform dose to smoother, rather than angular, volumes. This was confirmed visually using the XiO® 3D dose reconstruction tool (*Figure 6.5*). It has been observed elsewhere that, in general, larger volumes tend to exhibit better conformity [65].

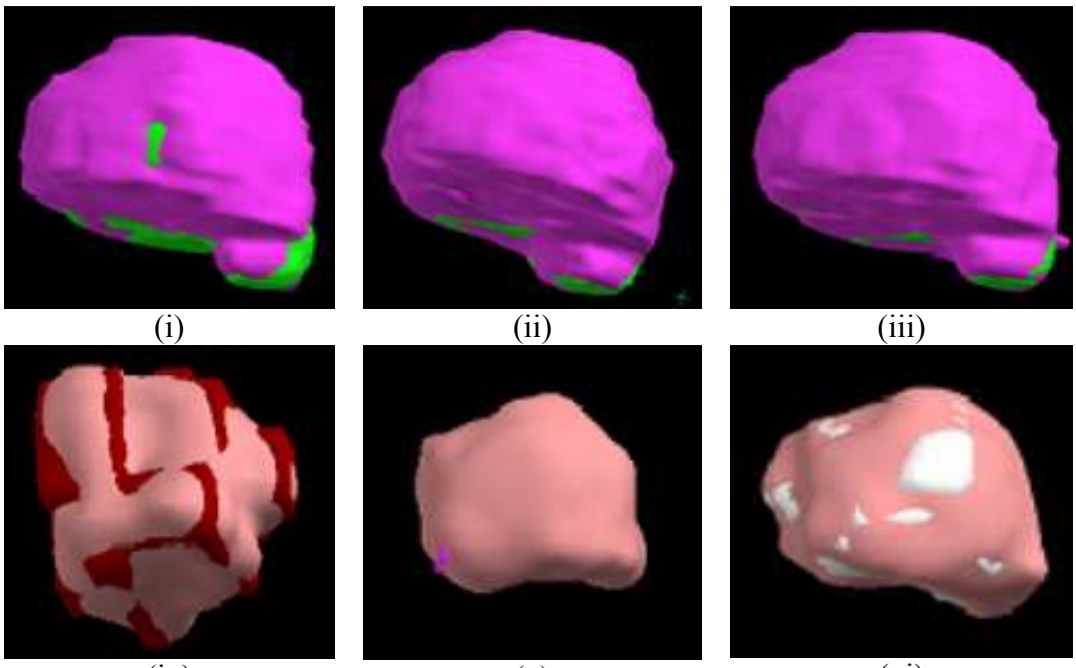

Figure 6.5: 3D dose conformity. 95% PTV coverage: Plan_rawAll (i), Plan_linear (ii) and Plan_linear+margin (iii). 95% boost coverage: Plan_rawAll (iv), Plan_linear (v), Plan_linear+margin (vi).

The gEUD of the PTV encompassing the different MRSI volumes ranged from 56.2 Gy for the smallest boost (Plan_linear) to 56.8 Gy for the largest boost (Plan_rawAll). This corresponded to a range over the TCP values of just 1.5 %, suggesting that the slightly larger boosts yield a minimal benefit in control probability. Overall, the largest, angular boost volume influenced dose distributions the most, albeit insignificantly in terms of both benefit and detriment.

6.4. Discussion

Whilst work has been carried out on using MRSI data in treatment planning, previous studies have not yet investigated the radiation therapy consequences of MRSI volume variability due to interpolation of the metabolite maps. It can be seen that in this case, the largest difference in spectra volume arose from registration issues (through-plane), not from in-plane interpolation as was expected. This highlights unresolved contouring issues. The best approach to this problem in the future will most likely be the trial-and-error use of an automatic contouring tool based on a pixel intensity threshold that draws

a volume approximating the size of the actual volume of the original data, which may be extracted from our metabolite-mapping software. Of course this introduces an unexpected uncertainty, adding to partial volume effects (individual spectra voxels may contain a mixture of tumour, necrosis and normal tissue) and any inherent spatial uncertainty in the acquisition data. It is possible to design phantoms as quality assurance tools for the spectral data acquisition, although it is difficult to improve on current resolutions due to spectral 'bleed' between voxels and in keeping scan times a reasonable length [97]. The registration issue would pose a particular problem if the abnormality levels i.e. NAA/Cho ratios, were to be separated into different greyscale levels, and dose levels prescribed to the differing degrees of spectral abnormality, as the greyscale information would be blurred in the image registration. Other groups have resolved this issue with in-house treatment planning software, whereby the contours may be drawn on the MRSI-MRI fused images and then transferred to the CT images at the registration stage [98]. This greyscale problem should also diminish if patients are MR-scanned in the same immobilisation mask used for treatment, as this would correct the angulation differences. Head and neck radiotherapy patients have been scanned in this way at the CMRI, using a body coil in order to allow room for the immobilisation mask [96]. This is not a feasible set-up for spectroscopy acquisition however, as in this case only the small head coil produces an adequate SNR, and of course the immobilisation mask would not fit inside this smaller coil.

Other uncertainties that affect the MRSI volumes also arise in the processing as well as acquisition stages. Defining a suitable NAA/Cho threshold for tumour presence for example, also affects the size of the spectroscopy-defined volumes [50]. Henceforth there is a chain of uncertainties, any one of which may affect the reliability of the contours. Additionally, inter- and intra-observer variability may affect the size and/or

location of the oncologist-defined target volumes (CTV and therefore PTV), which may increase the relevance and influence of metabolite map variations. The accuracy of the registration algorithm has been investigated in detail elsewhere and is of the order ±1-3 mm [96, 99, 100]. At higher field strengths e.g. 3 T, spectroscopy acquisition becomes more reliable and robust and these high-field scanners promise a dawn of improved functional imaging data. Finally, MRSI in this study was determined by T_2-weighted images i.e. the volume of interest (VOI) was positioned over the T_2 hyperintensity and it could be that other sequences afford better results. The variation in regional dosimetry resulting from the choice of interpolation method must be considered therefore in the context of these other causes of dosimetric variability in clinical practice.

The MRSI volumes are very nearly enclosed by the anatomical GTV (*Figure 6.2*), the spectroscopy-derived volumes being smaller than the oncologist-defined target. This discrepancy is difficult to characterise, since the methodology used for processing the spectra means that the tumour can rarely be fully covered by the MRSI FOV at acquisition. This is a common problem in MRSI techniques, although attempts have been made to compare CTV definition using MRSI against MRI alone, despite this incompleteness [53]. If the metabolic volume is indeed smaller than the classical CTV, as is suggested by much of the published data and our own experience, this means that a smaller volume, which may benefit more from the higher dose, would receive the boost dose level if it were prescribed to the MRSI volume rather than if it was prescribed to the conventional CTV. This would afford an advantage to the patient in terms of normal tissue dose minimisation if a smaller volume receives the high dose and could potentially lead to GTV-CTV margin reduction.

The XiO® IMRT treatment planning module uses an unconstrained conjugate gradient optimisation algorithm [43, 44]. Optimiser capability was tested by planning to the angular volume at the raw spectroscopic resolution and a 'smooth' interpolated volume of the same size. Examination of dose conformity to both the PTV and MRSI boost implies that the optimiser finds it easier to sculpt dose around a smooth volume. It could be said that it is advantageous to use an interpolated MRSI volume to maximise optimiser performance regardless of other differences incurred by the interpolation step. The angular volumes are artificially produced by the low spatial resolution of the MRSI and the true metabolism based tumour volume is of course not likely to be this shape, in fact it would have a resolution as fine as the molecular size of the metabolites. Unfortunately however, the true metabolic volume cannot be exactly identified and localised in a given patient and must therefore be represented in the most correct and accurate way achievable with current knowledge and technology. Whilst the effect of volume size and shape on dose conformity in IMRT is a general planning issue, it is important to understand the reasons why volume differences may arise and the implications of any such uncertainties.

The boost dose of 65 Gy used here is probably not sufficient to improve results by dose escalation alone but the study goes some way to defining a framework for boost dose methodology. If the region of active tumour can be more accurately delineated (as MRSI shows the potential to do), the probability of control may improve, even with such modest dose escalation strategies. The work here is presented as a case study and as such the findings are based solely on one patient. It is difficult to draw hard and fast conclusions from a study of one patient, although it may be said that relatively small variations in target volume definition and coverage are likely to have only a minor impact on treatment outcome anyway. Smaller volumes may incur larger uncertainties

as is the case in any volume delineation method. It is important to note however, that patient geometry i.e. size and position of spectra ROI, size and position of PTV and the proximity of any OARs will further affect the parameters investigated.

6.5. Conclusions

MRSI is a useful tool in radiation therapy planning where the metabolic information may be used in target delineation strategies or for defining small regions of aggressive tumour that would benefit from high boost doses. For this technique to become robust, reproducible and effective, the radiotherapy community must be mindful of any uncertainties and variability in MRSI volumes in order to be as sure as possible exactly where this extra dose is being deposited.

Recent studies incorporating such data into the planning process have used interpolations of the original acquired spectroscopic volumes and whilst a margin of 2-3 mm has been proposed to account for these uncertainties, it is not clear what effect this up-sampling has on the treatment planning in terms of the geometric differences in the MRSI volumes. The effects of interpolating metabolite maps via three different methods on the size of the resulting volumes and dose distributions in using the spectroscopic data to plan IMRT boost treatments have been investigated. It was found that the difference in volume size arose from registration issues associated with the functionality of the planning system and not from the initial interpolation. An uncertainty of ±2-3mm in MRSI volume size was confirmed but may not account for spatial variations. Given the dose deposition capabilities established in Chapter 5, these variations are not significant. It was also shown that plan optimisation is affected by not only the size but also the shape of the boost volume.

For this case study, no clinical degradation in using different interpolated metabolite maps in IMRT plans was noted. While differences between treatment plans were observed, the clinical significance of these is negligible for this example study and it is therefore concluded that metabolite map interpolation issues are not detrimental to generating accurate and desirable IMRT plans.

In the future, as MRSI becomes more routine in radiotherapy planning, it may be useful to generate some sort of 'confidence map' in addition to the usual metabolite map. A visual representation of the numeric uncertainties would be useful for case-specific planning. Alternatively, a standard margin could be added to MRSI volumes taking these uncertainties into account. Until these values have been derived over a large number of study sets, it's difficult to make explicit recommendations. In the mean time, MRSI may be useful as a supplementary tool, to be employed only in conjunction with conventional imaging techniques and target delineation methods.

Chapter 7 - Guiding Selective Multiple Boosting with Hybrid Images

7.1. Introduction

Presently there is much interest in functional imaging for delineating boost volumes in radiotherapy planning, where it is desirable to keep any boost volume as small as possible to minimise side effects. However, as discussed in Chapter 6, planning volumes defined by advanced functional techniques may in fact be comparable in size to conventional modalities once uncertainties and resolutions have been accounted for with relevant margins. Additionally, although MRSI volumes often reside within the GTV, likewise they can extend beyond the conventional T_1 post-contrast-enhancing lesion (often used to delineate the GTV) thus leading to a much bigger planning boost volume than using conventional imaging alone [53]. Although boosting the larger, more explicit volume may be preferable; in reality it may not be feasible depending on the location of the tumour and surrounding normal tissue tolerances.

Conventional MR and CT images give an indication of tumour presence to a given sensitivity, rather than characterising some functionality or proliferation measure i.e. they have inferior specificity to most functional modalities. MR images however, have excellent (sub mm) resolution and can be readily utilised in routine planning without the need for specialised processing. High field scanners (3T and above) with optimal functional imaging capabilities are not readily available and most centres are likely to have a 1.5T scanner at best. Acquisition times are also much shorter when anatomical imaging alone is undertaken. Furthermore, as well as availability, variability and

reproducibility issues, not all patients are suitable for functional imaging studies. MRSI for example, requires that a patient lie very still for approximately twenty minutes in the scanner, which may not be possible due to pain or cognitive impairment.

Relaxation times can be manipulated in traditional MRI studies to optimise contrast for a given application or for characterisation of certain anatomies. In this way, information can be gathered from different imaging perspectives in order to try and improve the specificity of the morphological techniques. The most familiar exploitation of contrast takes the form of T_1-weighted and T_2-weighted images. T_1 images have excellent anatomic detail and paramagnetic contrast agents administered intravenously increase both sensitivity and specificity. The contrast-enhancing (hyperintense) region on a T_1 image reflects areas where there has been a breakdown of the blood-brain barrier, which can be indicative of tumour infiltration. This alone may not be a reliable gauge of active tumour however, due to the presence of non-enhancing tumour tissue or contrast-enhancing necrosis. T_2-weighted images also exhibit good contrast between abnormal and normal tissue, whereby solid tumours appear as hyperintense signal. Nevertheless, non-specific processes such as inflammation or oedema may also appear hyperintense, making it difficult to determine what is and what is not tumour. The volume of abnormal signal on T_2 images is usually larger than T_1 for this reason [53]. It is proposed that the combined T_1 post contrast (T_1c) and T_2 abnormality i.e. $T_1c \cup T_2$ should then encompass all tumour cells.

Fluid-attenuated inversion recovery (FLAIR) can be used instead of a T_2-weighted sequence and often better defines the extent of a lesion. FLAIR sequences suppress the signal from fluids (cerebro-spinal fluid (CSF), oedema) so that the tumour is more readily defined. This may be useful when a lesion is close to the periphery of the

hemispheres or near to the ventricles, proximal to CSF. Additionally, lesion to background contrast is generally higher on FLAIR images compared with T_2s. FLAIR images at the Centre for MR Investigations, Hull (CMRI), are acquired via a propeller sequence, which minimises motion artefacts associated with patient head rotation. FLAIR images are acquired routinely at the CMRI as part of the standard diagnostic tumour protocol and so it would be useful to utilise these somehow in the radiotherapy planning process, since the patient would not have to undergo any additional scans. For this purpose, T_2-weighted and FLAIR images are discussed interchangeably, since FLAIR images are themselves heavily T_2-weighted.

Hull was the first centre in the UK to install a whole body capable 3.0-Tesla (3T) MR scanner; the highest field-strength in clinical use at the time of writing and twice the field-strength of most conventional scanners. The 3T scanner has a greater signal-to-noise ratio, resulting in higher resolution images with better detail and clarity. This represents significant progress in neuro-imaging, where smaller abnormalities (very small tumours for example) may be seen, which may be harder to visualise on an equivalent image acquired at 1.5T.

7.2. HyMRI: Image Amalgamation Software

In-house software was written in MATLAB[§] for generating Hybrid Magnetic Resonance Images (HyMRI). HyMRI (a term used here to describe both the technique and the piece of software) is based on the premise that different imaging perspectives can be tailored to give more complete information by combining them in some way. The program reads in a series of standard images generated via different MR sequences. Here, the user guides an automatic region-based contouring tool to delineate the region

[§] The Mathworks, Inc.

of abnormality on both datasets independently. The program then compares the contours slice by slice and amalgamates them to form a parameter map of tumour existence probability (TEP), derived from set theory analysis of the different contours. Finally, the parameter map is overlaid on some anatomical image and the fused image exported in a specific DICOM convention to the planning system. This dataset can subsequently be used for contouring and planning, specifically for delineating boost volumes (*Figure 7.1*).

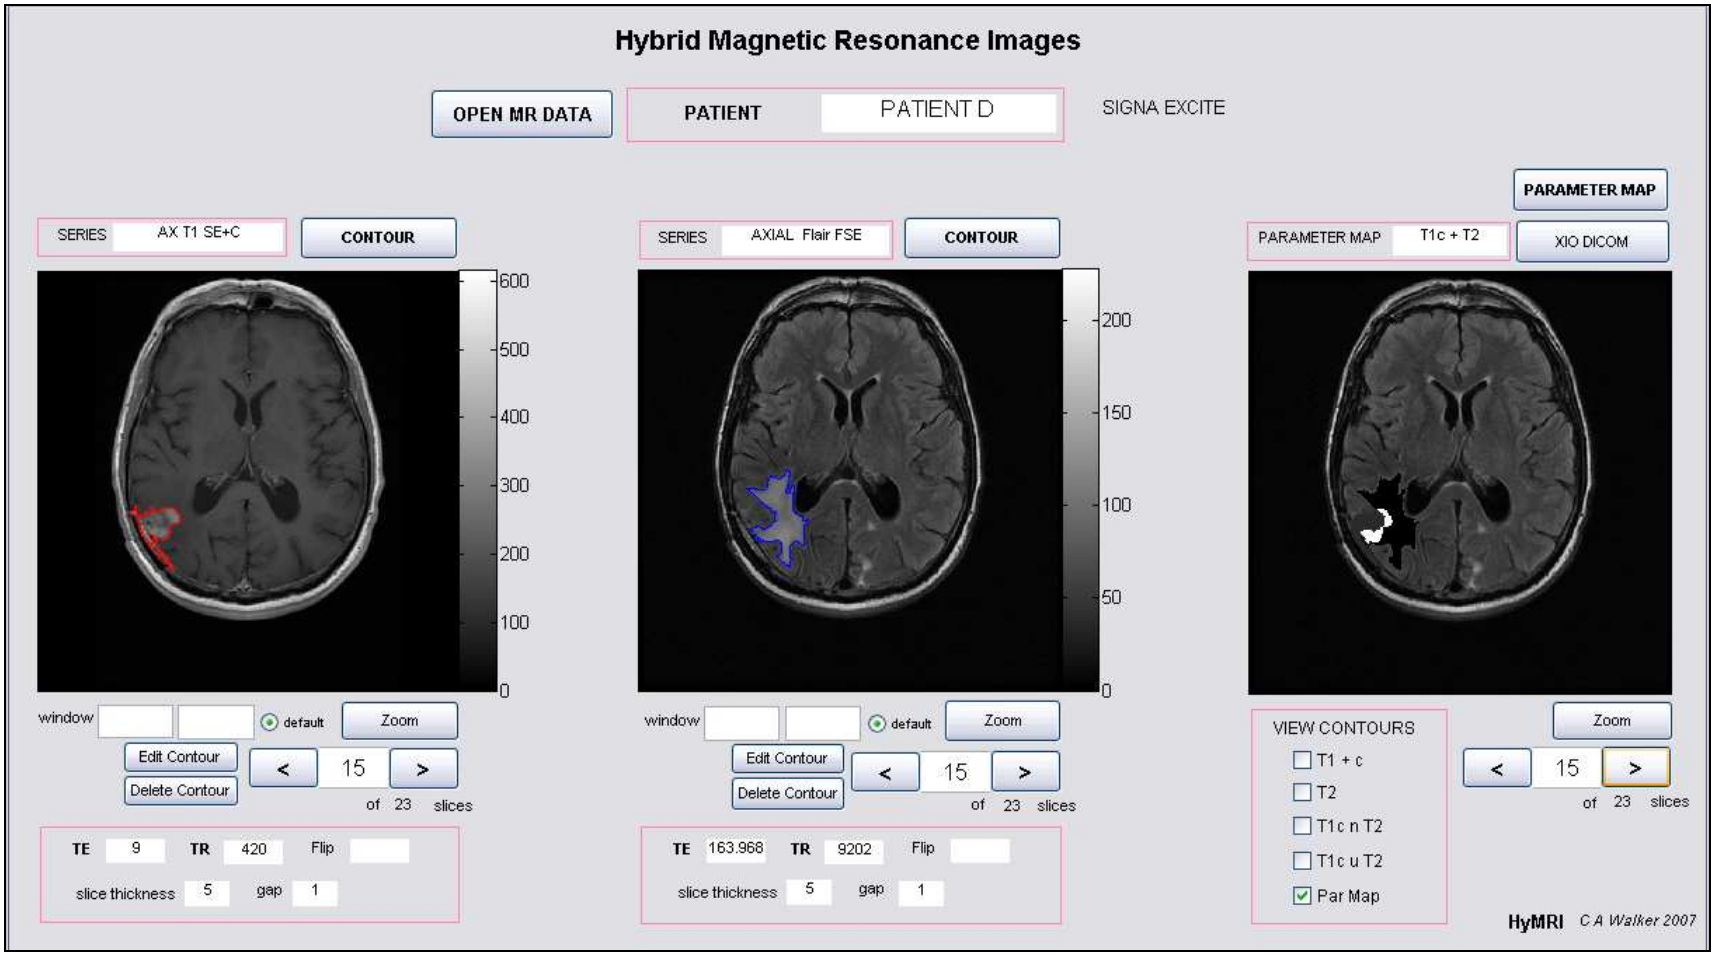

Figure 7.1: HyMRI (Hybrid Magnetic Resonance Images) GUI for amalgamating different MR sequences. The resulting parameter map of tumour existence probability (TEP) fused to an anatomical reference image (far right), can be exported to the XiO® planning system.

HyMRI is designed at present for processing a T_1-weighted post contrast (T_1c) and a T_2-weighted or FLAIR dataset acquired on either a 1.5 or 3.0-Tesla International General Electric[§] Signa scanner. The software is manufacturer specific only in terms of recognising certain field names in the DICOM image header and this function could easily be extended to read images generated by other manufacturers. Furthermore, the software could very easily be adapted to process more than two datasets, incorporating different sequences and/or functional techniques. However, the simple example of a T_1c-T_2 amalgamation is used from hereon in to demonstrate the methodology.

The relevant DICOM images can be pulled from the scanner and copied onto a normal PC to be read by the HyMRI software. For each image dataset, the user specifies a minimum and maximum window level at which the image is thresholded. The relevant region of interest (ROI) can then be selected from the resulting binary image. In general, it is most effective to use a larger window range and cut away any unwanted parts using the 'Edit Contour' tool. The user can readily scroll through the image sets, as well as zoom in on particular regions of an image to facilitate the contouring process.

A single button click generates the parameter map fused to the T_2 image set. Three grey levels within the parameter map are correlated to an arbitrary tumour existence probability, denoted TEP. A higher TEP value represents a high probability that the tumour exists there, with a lower TEP value corresponding to a lower probability of the presence of tumour cells. For the case of two image datasets, there are simply three probability levels: TEP1, TEP2 and TEP3. Based on this system, TEP1 is that region where voxels enclosed by the T_2 contour exists uniquely i.e. ($T_2 - T_1$c), since this carries a relatively low chance of containing tumour; the hyperintense T_2 signal representing

[§] Milwaukee, WI, USA

mostly oedema. TEP2 is defined as the volume encompassed by the T_1c contour exclusively i.e. ($T_1c - T_2$), as the contrast-enhancement delineates the bulk of the tumour volume, thus indicating a higher probability of tumour existence. Finally, TEP3 is defined as the intersection of the two contours i.e. ($T_1c \cap T_2$), the area where tumour is most likely to exist. The user can scroll through and zoom in on the amalgamated contours as well as the parameter map on each image slice. Once they are satisfied that the final parameter map best represents the abnormality depicted on the images, a final button press saves the fused images in DICOM format and in a given nomenclature readable by XiO®.

HyMRI handles differences afforded by real imaging data by interpolating images of different matrix dimensions. For example, pure T_2-weighted images are stored by the scanner as 1024 x 1024 matrices, whereas T_1-weighted and FLAIR images are generated as 512 x 512 matrices. In order for the different MR datasets to be compared on a pixel-by-pixel basis, the larger T_2 matrices are interpolated to be of equal size to the T_1s. Also, the final dataset, to which the parameter map is fused, is resized to its original matrix, as XiO® detects this original value and displays the data accordingly.

7.3. Methods

Five brain tumour patients were randomly selected for a study investigating the feasibility of the HyMRI software and it's usefulness in planning selective multiple boosts. The data was taken from diagnostic MR scans and as such the type of tumour is, for this purpose, arbitrary. All images were acquired on either a 1.5 or 3 T clinical General Electric Signa scanner, equipped with an eight channel quadrature (birdcage) RF head coil. Of course patients undergoing diagnostic tumour imaging would not necessarily go on to have a course of focussed radiotherapy; if they had multiple

metastases or slow-growing tumours for example. To reflect this, all data was processed and planned in an identical manner (in terms of margins and doses) and each case subsequently discussed in the context of the standardised methodology.

Since the patients were not undergoing radiotherapy treatment, no planning CT was available and so treatment plans were generated using MRI as the primary planning data. This was possible due to the relative uniformity of tissue density in the brain, which reduces the need for electron density information provided only by CT images. This was implemented simply by applying a bulk density correction at the dose calculation stage i.e. no heterogeneity correction. Also, the small field of view required at acquisition minimises distortion effects which can be problematic in using MRI for planning and others have shown that it is reasonable to plan brain radiotherapy with MRI alone [5]. Finally, since the treatment plans are intended to serve as proof of principle and not to be used for patient treatment, the robustness of the images in terms of dose calculation and geometric stability were not deemed essential. If the HyMRI technique were to be tested on the scans of patients undergoing radiotherapy treatments, the MR scans could be registered very easily to the planning CTs in line with current practice.

Brain tumours are difficult to segment because they have a wide range of appearances and effects on surrounding structures. They may enhance fully, partially or not at all with MR contrast and vary considerably in image intensities on MR images. These detailed radiological characteristics of tumours were not integrated into the contouring here at either the HyMRI or treatment planning stage. This would require an experienced clinical oncologist to delineate the tumour volumes, introducing issues of intra- and inter-observer variability. In order to simplify and standardise the planning

study then, contours within HyMRI were derived from a simple intensity thresholding and planning volumes were grown from these TEP volumes using a fixed margin approach. In reality, the tumour would be delineated based on experience, expertise and many different image characteristics and would, of course be histology and patient-specific.

Parameter	Patient A	Patient B	Patient C	Patient D	Patient E
Field strength	3 T	3 T	1.5 T	1.5 T	3 T
Image sequence					
i.	T_1c	T_1c	T_1c	T_1c	T_1c
ii.	T_2	FLAIR	FLAIR	FLAIR	T_2
TE ms					
i.	8	9	240	10	8
ii.	104	139	240	163	100
TR ms					
i.	400	400	420	520	480
ii.	3860	9500	9202	9202	4600
Pixel spacing mm					
i.	0.469	0.469	0.469	0.469	0.469
ii.	0.234	0.430	0.469	0.469	0.273

Table 7.1: MR image acquisition parameters for patients A-E.

In order to acquire the T_1c images and as part of their routine diagnostic scans, patients were injected intravenously at a standard dose of 0.1 mmol/kg of a gadolinium-based (Gd-DTPA) contrast agent. Either FSE T_2-weighted or FLAIR images were also acquired as part of the same examination, at exactly the same slice locations and in line with standard tumour imaging protocol (*Table 7.1*). All images were acquired at 5 mm slice thickness with 1 mm slice gap. The two datasets were then processed in the HyMRI software to generate the TEP parameter maps and these exported to the XiO® planning system.

Within the planning system, the TEP parameter map was used to define a clinical target volume (CTV) straight away, rather than growing a CTV from a gross target volume (GTV). This is because the TEP volumes represent not only the bulk of the tumour

volume but also the surrounding fluid, which may or may not harbour microscopic spread as represented by the concept of a CTV. The CTV was therefore defined as the union of the two contours i.e. (TEP1∪ TEP2) with a +1.5 cm 3D isotropic margin added using an automatic margin tool within XiO®. The PTV was outlined as the CTV plus a +0.5 cm 3D isotropic margin to account for uncertainties in planning and patient set-up as recommended in ICRU Report 50 [27]. These are fairly average margin sizes when considering malignant tumours. In fact, a margin of +2 cm to the T_2 hyperintensity volume is typical for conformal therapy PTV definition for high-grade tumours [53, 101, 102]. In clinical practice, benign tumours for which radiotherapy does have a role may only have a margin of 0.5-1 cm added to the visible tumour. The CTV was edited to be bounded by the skull i.e. where disease spread would not be possible, and the PTV bounded by the skin to account for movement and setup uncertainties in the CTV. In some cases the PTV was also edited around critical OARs such as the eyes and brainstem. The patient outline was contoured in every case, along with the normal brain and OARs were defined as the eye orbits, lenses, optic nerves, optic chiasm, brain stem and pituitary gland. All contouring was undertaken by the author for consistency and to eliminate multi-observer error for fair comparisons. Margins were not added to the TEP volumes, since by the nature of their definition, this would result in substantial overlap issues. The PTV-annulus (PTVa) is defined as the PTV minus the TEP planning volumes i.e. $\mathrm{PTVa} = \mathrm{PTV} - (\mathrm{TEP2} + \mathrm{TEP3})$.

The regions with the highest probability of tumour existence i.e. highest TEP, were planned to receive the highest doses. A conventional dose of 60 Gy was prescribed to the PTV and boost doses of 75 Gy and 80 Gy to the TEP2 and TEP3 volumes respectively, representing ambitious boost dose ratios (BDR) of 125 % and 130 %. The TEP1 volume ($T_2 - T_1c$) was not assigned a boost dose as this did not harbour a high

probability of tumour existence and also because it received a tumourcidal dose anyway, being part of the region from which the PTV was grown. In order to prioritise any overlapping voxels, the highest dose boost volume - TEP3 - was ranked the highest in the optimisation, with TEP2 next and PTV ranked lowest of the three targets. Figure 7.2 shows the HyMRI workflow from raw data acquisition through to the final SMB plan.

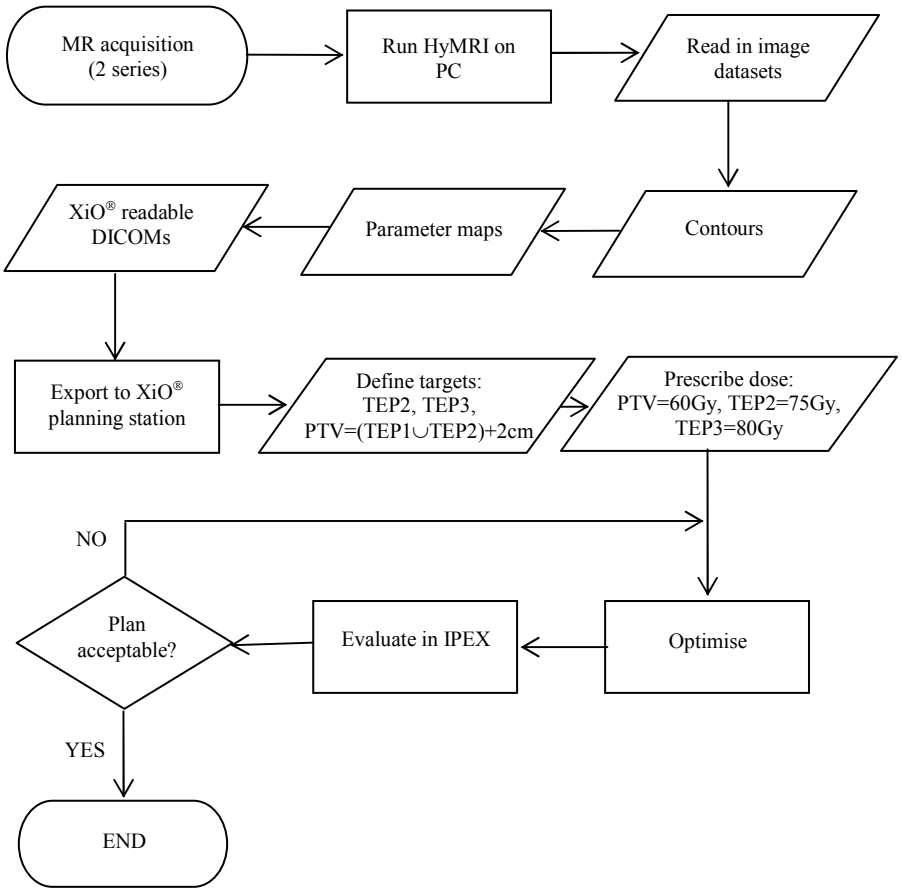

Figure 7.2: Workflow for SMB planning using HyMRI software. Target volumes and dose levels are specific to this study and may be adjusted accordingly.

Prior to commencement of planning, an experiment was set-up to determine an optimal optimisation margin for each anatomy. Eight plans for a given patient anatomy were optimised using an incremental optimisation margin of between 0 and 10 mm and evaluated in terms of the resulting DVHs. An optimisation margin of 3 mm produced the most optimal results in terms of target coverage and so this value was applied to all

further planning. Plans were designed so that all isocentres were placed at the centre of each individual PTV and only isocentric sets of coplanar beams were considered. Final dose was calculated on a 2 x 2 x 2 mm^3 grid using the Fast Fourier Transform (FFT) convolution algorithm.

Two IMRT plans were produced for each patient: a plan with TEP boosts as outlined above (SMB) and a plan with no boost at all (NB). This allowed for both control probability and normal tissue dose comparisons to be drawn between a baseline plan and the new technique. Plan evaluation was carried out in IPEX (Chapter 4) primarily through inspection of the DVHs of both target volumes and OARs. 3D dose distributions were optimised to ensure best possible coverage of the three target volumes (PTV, TEP2, TEP3) with their individual 95 % isodose surfaces. In order to evaluate this, the variable $V_{95\%}$ was recorded; the volume of the target receiving at least 95 % of its prescribed dose. All involved OARs were evaluated using both DVHs and dose statistics. For the optic nerves, orbits, brainstem, optic chiasm and normal brain, the maximum dose was given by $D_{2\%}$ i.e. the dose received by at least 2 % of the structure. For the smaller lenses and pituitary, the global maximum was observed. Finally, the mean dose of each OAR was calculated. OAR dose constraints are given in *Table 7.2*.

Structure	**Maximum dose Gy**
Orbits	60
Lenses	12
Optic nerves	50
Pituitary	45
Optic chiasm	55
Brainstem	60

Table 7.2: Maximum dose criteria imposed on OARs for treatment plan optimisation and evaluation.

The generalised equivalent uniform dose, gEUD, was calculated for the PTV of each plan [75]. The gEUD provides the uniform dose that would have to be delivered to the whole PTV to achieve the same biological outcome as the concomitantly boosted PTV and serves to give some indication of the dosimetric benefit afforded by the SMB plans over a NB IMRT treatment. The tumour control probability, TCP, was also computed to this same end.

7.4. Results

The following cases are presented as planning studies only and as such do not represent patients who would necessarily be referred for or benefit from radiation therapy as part of their treatment regime. The purpose is to explore the feasibility of guiding boost doses with the hybrid MR parameter maps on real brain tumour anatomies. The prescribed doses therefore are chosen arbitrarily, albeit realistically.

7.4.1. Patient datasets

Patient A is a 56 year old female with a right sided deep temporal lesion five years after excision of lung cancer, including radiotherapy. She had undergone a non-diagnostic biopsy and the radiology report, although inconclusive, suggested a solitary metastasis. Patient B is a 68 year old male with a solitary lesion in the right occipital lobe, most likely to be a lung metastasis. A processing artefact arose on the FLAIR images for this patient, whereby the patient within the image looked to have been cropped anteriorly in a circular fashion (*Figure 7.3*). Missing anatomy implicates inaccuracies in the dose calculation and if this had been a patient coming for treatment, the MR images would not have been suitable for treatment planning. Fortunately, since this was a feasibility study only, the data was still valid here, as the tumour depiction was intact and enough

normal anatomy was visible in order to complete the contouring process. The source of the artefact cannot be obviously deduced although it is likely to have arisen as a complication of the PROPELLER option used in the FLAIR acquisition, which is used to correct for patient motion. This works by acquiring k-space in 'blades' and it seems as if the error seen here is some specific artefact associated with this different type of reconstruction.

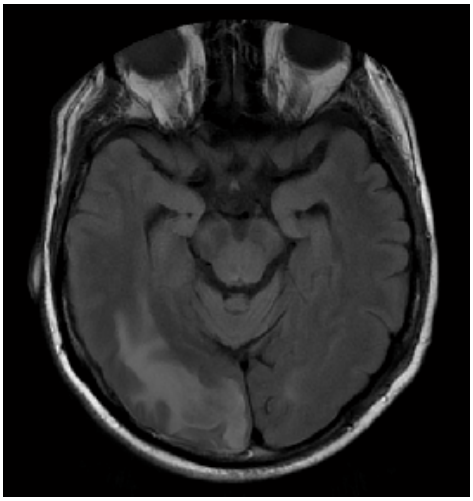

Figure 7.3: Artefact associated with FLAIR data of Patient B, where anterior anatomy has been cut-off due to some processing complication most likely owing to the propeller option on the scanner's FLAIR sequence.

Patient C is a 63 year old male with a high grade glioma in the right temporo-parietal region with considerable associated vasogenic oedema and Patient D is a 63 year old female patient with a solitary breast metastasis in the right hemisphere, also surrounded by a great deal of oedema.

Patient E is a 54 year old female with a sizeable recurrent olfactory groove meningioma. There were several technical problems associated with Patient E, highlighting a need for caution and selection in applying the HyMRI process. In this case, some parameter related to the field-of-view (FOV) was set differently for T_1c and T_2 acquisition, resulting in the patient being a different size within the image matrix i.e. an apparent magnification (same matrix). Fortunately, the FOVs were both centred at the same

location and so a simple scaling could be implemented. The exact difference was extracted from the 'ReconstructionDiameter' field in the DICOM header, being 240 for the T_1c images and 280 for the T_2. This was addressed by enlarging the smaller T_2 images by (280/240) and 'cropping' the resulting matrix to the standard 1024 x 1024.

Furthermore, although the T_1c and T_2 images were acquired at exact corresponding slice locations, the T_2 dataset contained two additional images. These were deleted and the new number of slices assigned. Finally, the patient exhibited significant restlessness and forgetfulness, which meant she was unable to remain still for the duration of the scan. As a consequence, there is quite a bit of movement artefact in both image datasets, most notably in the T_1c images. Owing to this through-plane patient movement, there arose a variable uncertainty in slice location with respect to patient anatomy. This was accounted for by visually inspecting the images and manually rearranging them accordingly (*Figure 7.4a*). Replacing image 14 with image 15 in the T_1c dataset gave the best approximation to the correct anatomical order. The maximum error associated with this type of extreme patient movement is therefore estimated to be ± 1 image slice i.e. ± 6 mm in the z-direction.

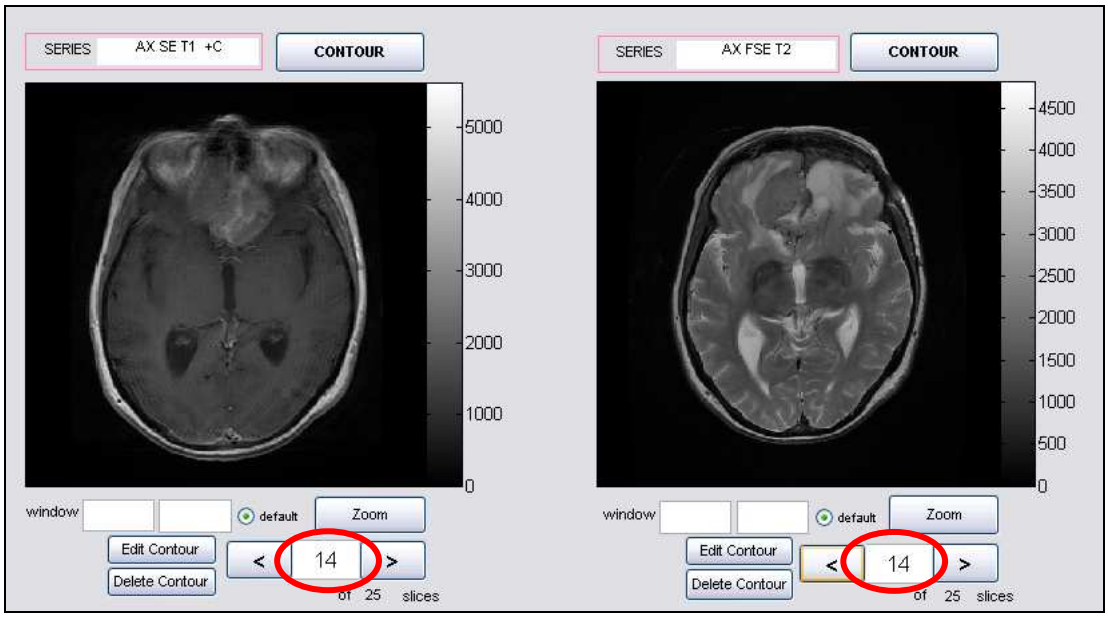

a.

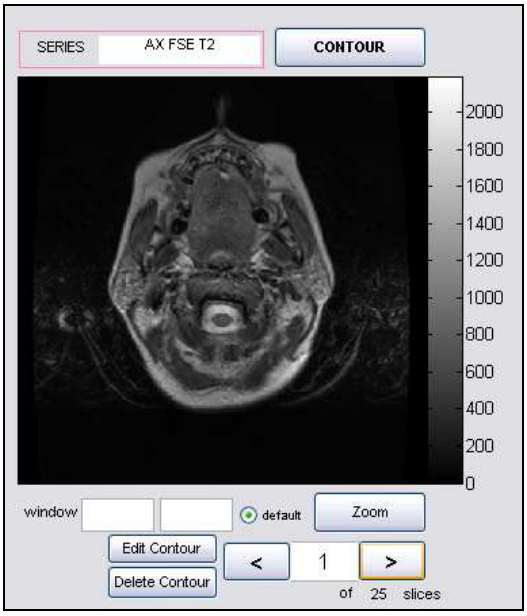

b.

Figure 7.4: Screenshots taken from HyMRI showing significant movement artefacts associated with data of Patient E. a. Patient motion in the z-direction identified as different visible anatomies at the same slice location. b. Ghosting artefact attributable to patient motion in-plane.

Owing to the extensive movement artefact, the data from Patient E was not taken forward for treatment planning. In a case where there is already so much uncertainty in the image representations of the patient's anatomy, it would not be appropriate to consider planning precisely targeted therapy. This is an important consideration when selecting patients for the SMB, or indeed any IMRT method.

7.4.2. Parameter maps

TEP parameter maps were successfully generated within HyMRI for Patients A-D (*Figure 7.5*) and exported in DICOM format to the TPS.

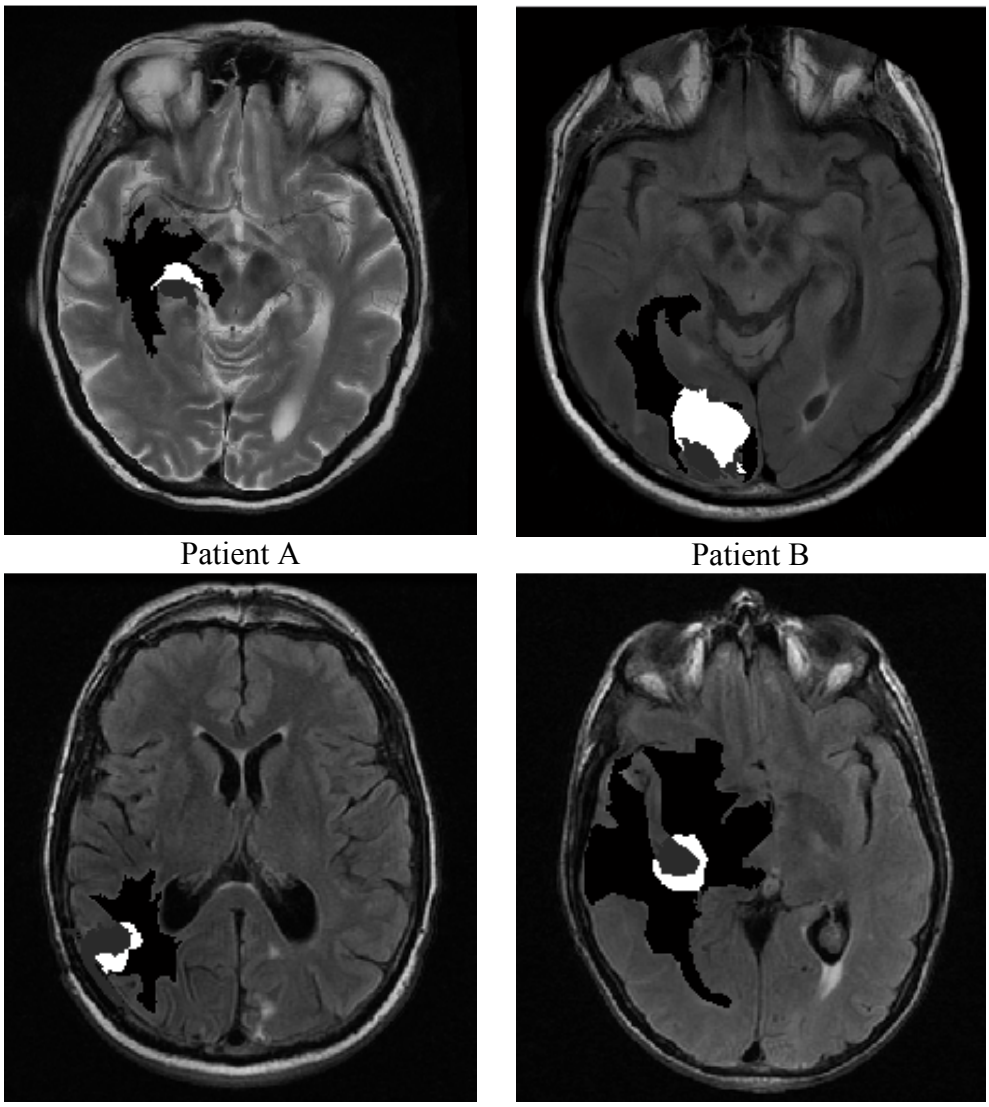

Figure 7.5: Parameter maps overlaid on T_2-weighted images for the final four patients selected for planning. The parameter map was generated in HyMRI software where three grey levels represent probabilities of tumour existence: lowest probability TEP1 (black), TEP2 (grey) and highest probability TEP3 (white).

7.4.3. Contours

Out of the four patients that went on to the planning stage (Patient A-D), Patient B had the largest PTV and Patient A the smallest, with mean 594 cm^3 and range 517 cm^3 (*Table 7.3*). The boost volume ratio (BVR), taken as the proportion of PTV that is

boosted was calculated for both the individual TEP2 and TEP3 volumes as well as for the two boosts combined. The relative size of the TEP2/TEP3 contours was evenly divided; half of the patients (Patients A and C) had a larger TEP3 contour compared to the TEP2 contour. Patients B and D had a larger TEP2 contour than TEP3, although the differences were smaller in these cases. The mean BVR for TEP2 (BVR_{TEP2}) was 0.79 % with range 1.04 %. The mean BVR for TEP3 (BVR_{TEP3}) was 1.17 % with range 2.47 %. The mean BVR for the boosts combined i.e. TEP2 $\cup$ TEP3 (BVR_{total}) was 1.96 % with range 2.89 %. The largest BVR_{total} was that of Patient B, at over 3 % of the PTV. This is still small compared to the MRSI defined boosts in Chapter 6, where the BVR ranged from 5.88 to 8.63 %. A 3D visualisation of the planning contours is shown in *Figure 7.6*.

Structure	**Volume cm^3**			
	Patient A	**Patient B**	**Patient C**	**Patient D**
PTV	332	849	418	778
PTVa	328	817	411	766
TEP2 ***BVR_{TEP2}***	0.63 0.19 %	6.21 0.73 %	5.15 1.23 %	8.02 1.03 %
TEP3 ***BVR_{TEP3}***	2.11 0.64 %	25.36 2.99 %	2.18 0.52 %	4.06 0.52 %
BVR_{total}	0.83 %	3.72 %	1.75 %	1.55 %

Table 7.3: Size of planning volumes, including the boost volume ratio (BVR) for individual and combined boosts.

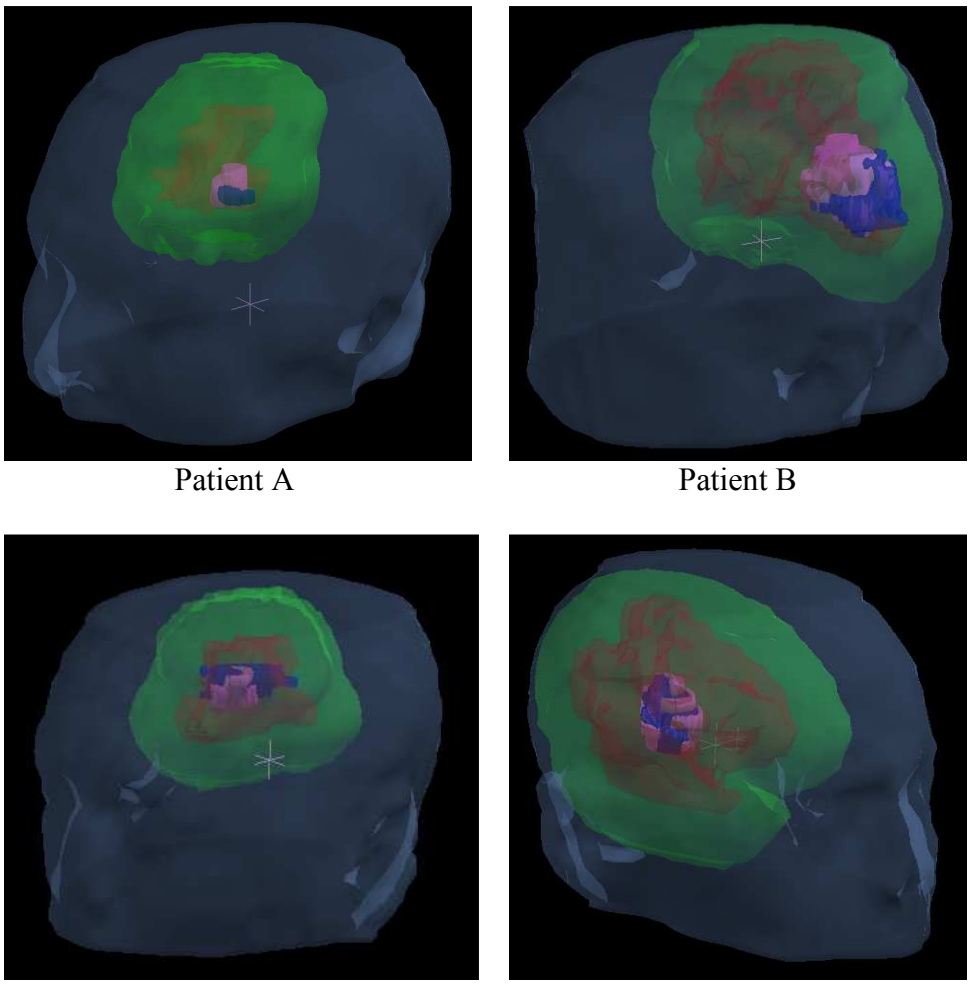

Figure 7.6: 3D rendering of planning volumes generated from the parameter maps; two boost volumes TEP2 (blue) and TEP3 (pink) and the PTV (green). TEP1, from which the PTV is derived, is shown in red.

7.4.4. Dosimetry

The SMB plan of Patient A maintained a $V_{95\%}$ of at least 97.9 % for all three targets; PTVa, TEP2 and TEP3 (*Figure 7.7a*). $V_{107\%}$ of the two TEP boosts was less than 1 %. All OAR dosimetry was within the tolerances set by *Table 7.2* except the brainstem, where $D_{2\%}$ was unacceptably high at 75.7 Gy (*Figure 7.7b*). This arose due to the proximal location of both the PTV and TEP boosts. The optic chiasm was also too hot, with $D_{2\%}$=57.9 Gy. For this anatomy therefore, the SMB technique would not be appropriate due to the radiosensitivity and proximity of the brainstem. For the NB plan, with PTV $V_{95\%}$=97.2 %, $D_{2\%}$ of the brainstem reduced to 63.3 Gy, which is still higher than desired. The optic chiasm was better spared in the NB plan, with $D_{2\%}$=52.3 Gy as

well as the left optic nerve ($D_{2\%}$=50.8/39.1 Gy) and pituitary gland (D_{max}=48.9/40.3 cGy). Dose-volume analysis may still render the plan unacceptable as the volume of the brainstem receiving at least 40 Gy, V_{40Gy}=81.0 %. However, a clinical decision would have to be made if this was a real case in order to establish the trade-off between target coverage and brainstem sparing. Considering the damaging doses received by various OARs in the SMB plan, the gain in gEUD across the PTV afforded by the boosts is only +1.85 Gy, corresponding to a TCP gain of +8 % (*Table 7.4*). This suggests that boost volumes so close to a sensitive structure are neither viable nor especially beneficial, even when the volumes are small; Patient A had both the smallest PTV and BVR of the group.

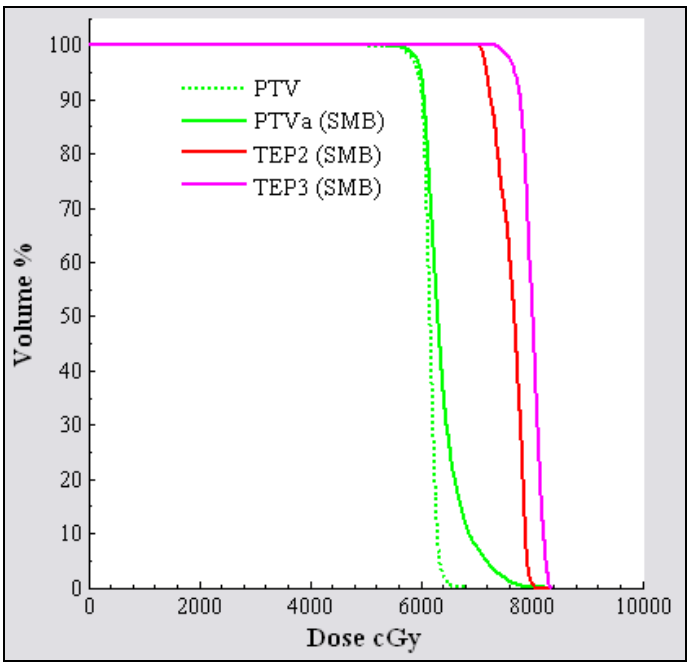

a.

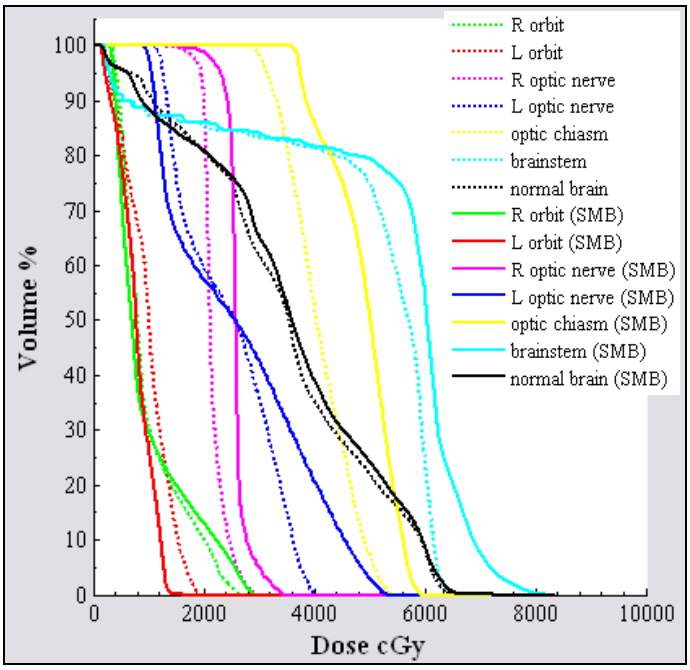

b.

Figure 7.7: DVHs for Patient A. Both plans utilised five coplanar beams at gantry angles[§]: 72°, 144°, 216°, 288° and 306°. a. Target DVHs for both the SMB and non-boosted plan. b. OAR DVHs for both the SMB and non-boosted plan.

An SMB plan was produced for Patient B that maintained $V_{95\%}$ of between 97.9 and 99.7 % over the three boost volumes, with $V_{107\%}$ of the two boosts both less than 1 % (*Figure 7.8a*). Dose-volume analysis of the OARs showed that all were well within set tolerances and the brainstem was actually better spared in the SMB plan than when no

[§] IEC convention

boosts were used (*Figure 7.8b*). Not only was the SMB plan feasible and well tolerated, it also afforded a PTV gEUD gain of +5.0 Gy, corresponding to an overall TCP gain of +20 % (*Table 7.4*). This suggests that neither PTV nor boost size is a barrier to safe and valuable SMB planning, since Patient B had both the largest PTV and BVR of the group.

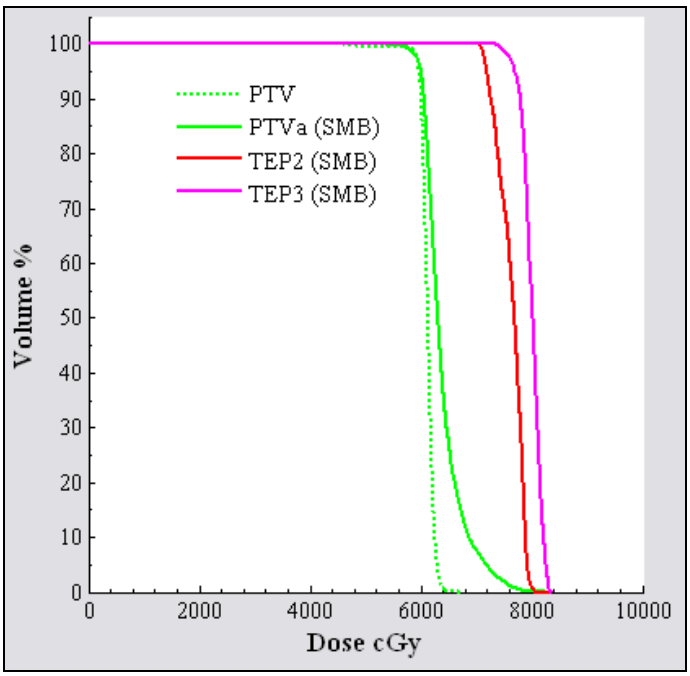

a.

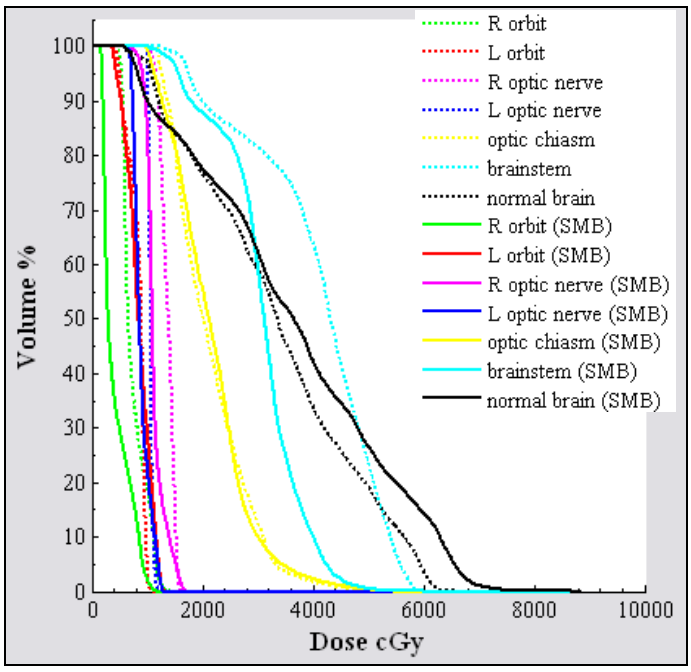

b.

Figure 7.8: DVHs for Patient B. The SMB plan utilised five coplanar beams at gantry angles: 72°, 144°, 216°, 288° and 306°. The non-boosted plan consisted of these plus two additional coplanar beams at gantry angles: 270° and 180°. a. Target DVHs for both the SMB and non-boosted plan. b. OAR DVHs for both the SMB and non-boosted plan.

In the SMB plan of Patient C, $V_{95\%}$ of the three targets ranged from 96.3 to 99.4 % with $V_{107\%}$=7.4 and 2.9 % for TEP2 and TEP3 respectively (*Figure 7.9a*). Once again, all OARs were safely within the specified tolerances and the differences in OAR dosimetry

between the SMB and NB plan were negligible (*Figure 7.9b*). A gain in TCP of +20 % was observed for the SMB plan, where the boosts provided an increase in PTV gEUD of 4.8 Gy (*Table 7.4*).

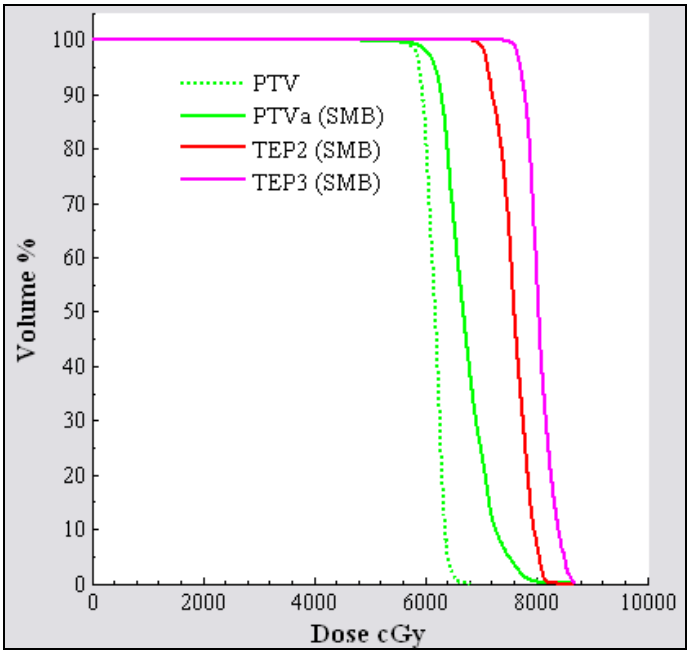

a.

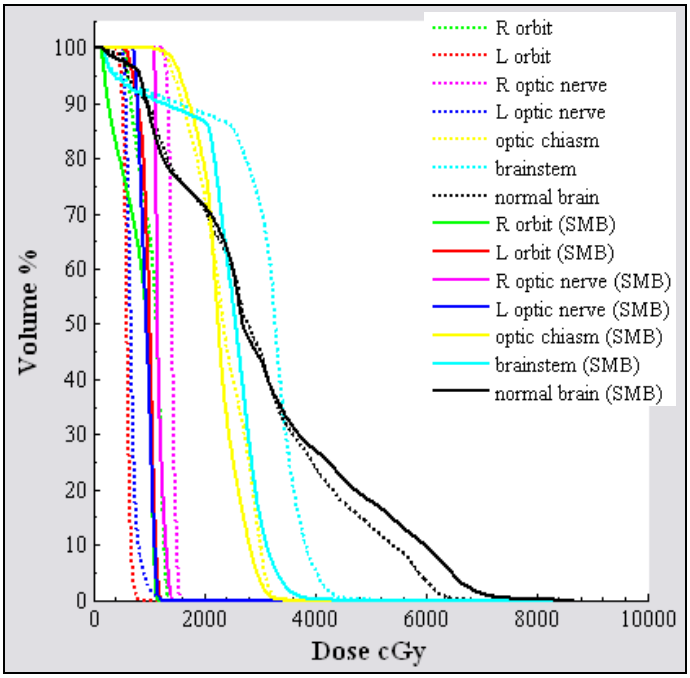

b.

Figure 7.9: DVHs for Patient C. Both plans utilised five coplanar beams at gantry angles: 72°, 144°, 216°, 288° and 306°. a. Target DVHs for both the SMB and non-boosted plan. b. OAR DVHs for both the SMB and non-boosted plan.

Patient D was the most difficult case to plan. The best and final SMB plan provided $V_{95\%}$ ranging from 96.0 to 99.0 % for the three targets, but this was at the violation of

numerous sensitive structure tolerances (*Figure 7.10*). For the right optic nerve, $D_{2\%}$=65.6 Gy with V_{45Gy}=100 %. D_{max} of both lenses was outside of the specified tolerances (14.7 and 14.6 Gy), as well as the optic chiasm ($D_{2\%}$=62.8 Gy) and pituitary (D_{max}=53.6 Gy). An acceptable dose distribution could not be produced, even for the NB plan and even experimenting with an arrangement of nine beams. As with Patient A, this dilemma is for clinical resolution and in this case, neither the SMB nor the NB plan protocol was feasible; the planning volumes were simply too close to the eye structures to physically achieve simultaneous target coverage and OAR sparing. In terms of comparison between the SMB and NB plan, the SMB yielded a +6.5 Gy gain in PTV gEUD, corresponding to a 27 % gain in TCP – the highest of the group (*Table 7.4*). Although it is unlikely that either plan would be clinically acceptable, it is interesting to note that OAR dosimetry for the SMB plan is no worse than the NB plan for such a significant gain in control probability.

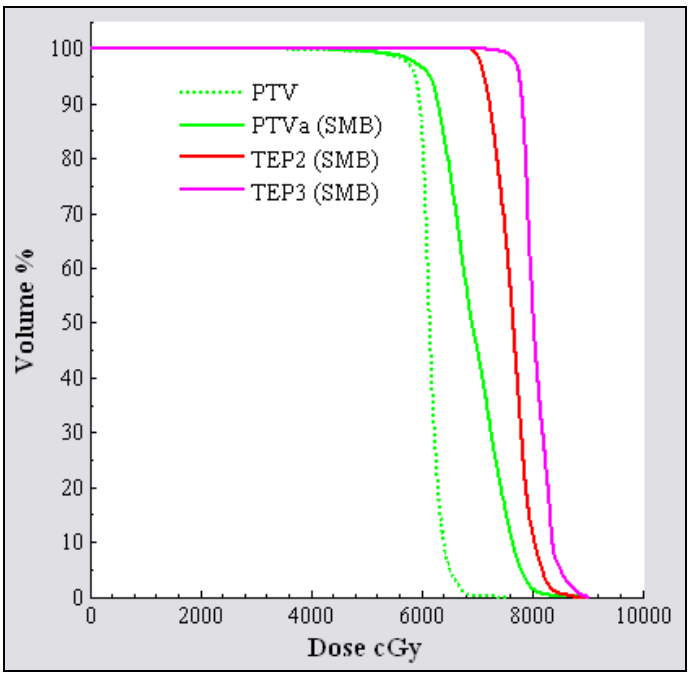

a.

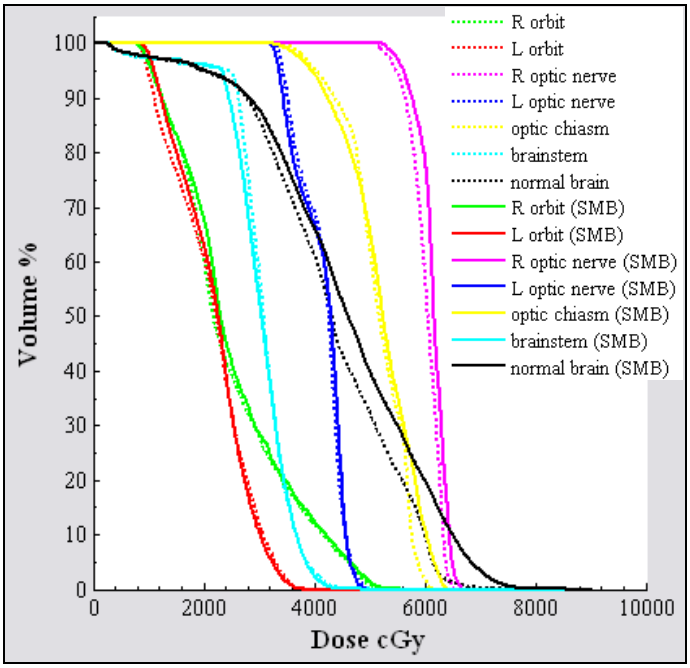

b.

Figure 7.10: DVHs for Patient D. Both plans utilised seven coplanar beams at gantry angles: 72°, 144°, 180°, 216°, 270°, 288° and 306°. a. Target DVHs for both the SMB and non-boosted plan. b. OAR DVHs for both the SMB and non-boosted plan.

The mean dose to normal brain tissue was negligibly higher in all four patients for the SMB plan than the NB plan, but the differences were all under 3 Gy. The mean of normal brain mean doses for the SMB plans was 37.3 Gy compared with 35.3 Gy for the NB plans.

	gEUD cGy		**TCP %**	
	PTV_{SMB}	**PTV_{NB}**	**PTV_{SMB}**	**PTV_{NB}**
Patient A	6322	6137	46	38
Patient B	6583	6087	45	25
Patient C	6624	6144	55	35
Patient D	6680	6026	51	24

Table 7.4: TCP and gEUD gains for SMB plans over NB plans.

7.5. Discussion

Although different imaging techniques can be employed to define regions of tumour with differing levels of certainty, this is not enough to reduce planning margins due to the infiltrative and recurrent nature of malignant brain tumours. The combined HyMRI-SMB technique proposes to combine these different imaging modalities to derive parameter maps of TEP and hence escalate dose only to small volumes, where the tumour is known to some given certainty to exist, within the conventional PTV. It has been shown that in-house software HyMRI can be used effectively to generate these parameter maps from which SMB plans can be designed. The HyMRI parameter maps were derived from two MR sequences in these examples; a T_1-weighted post contrast injection and a T_2-weighted sequence. These were selected for the feasibility study since both offer sub mm resolution, are well established, have short scan times, are easy to interpret (no specialist post-processing/registration required) and can therefore be incorporated most straight-forwardly into routine treatment planning.

HyMRI methodology assumes that no inter- nor intra-examination movement occurs i.e. the different imaging sequences do not require registration to one another, a pixel to pixel correlation is assumed. The method also requires adherence to a strict imaging protocol i.e. all imaging sequences must be acquired at exactly the same slice locations and FOV. Differences in slice locations and FOVs between examinations and/or patient movement artefacts can render the images useless, as with Patient E, where patient movement created unsalvageable inaccuracies in the imaged anatomy. This illustrates how some patients may not be suitable for long scans such as MRSI and may require alternative boost volume definition methods, although the accuracy of even conventional imaging can be compromised in such circumstances. Image registration may help in these cases.

HyMRI methodology is entirely dependent upon the integrity of the imaging sequences used. In using T_1c images, it is assumed that tumours are either ring enhancing or fully enhancing with contrast agent, which is not always the case. Many tumours do not enhance at all. In general, low-grade gliomas do not show contrast-enhancement as the presence of contrast enhancement generally indicates high-grade (malignancy). Some astrocytomas and typically all oligodendrogliomas do not enhance, and some high grade astrocytomas that do enhance often show enhancement in only a region of the tumour.

The same MR sequences can also be utilised in different ways to glean the most reliable information from them. Patients would be most likely to come for radiotherapy post surgery if the tumour is resectable. In these cases, T_1 images pre-contrast may be a useful addition, as they demonstrate post-operative haemorrhage as high signal. Post-contrast T_1 images could then be compared with pre-contrast T_1 to subtract necrotic areas from the contrast enhancing region, since this represents a combination of

haemorrhage and tumour. Spectroscopy may also be useful in this case, since it can discriminate between necrosis (characterised by an absence of all metabolites except lipids and lactate) from tumour and viable normal tissue. Thinner slices, possibly without a gap, although increasing scan time, would further provide a more accurate depiction of tumour volumes. A CT or PET derived target volume could serve as another TEP level although an advantage with the demonstrated method is that it abates the need for any image registration e.g. MR-CT.

SMB plans were successfully derived from HyMRI parameter maps for two out of four patients, both yielding gains in TCP of +20 % and at no cost to surrounding normal tissues when compared with the NB plans. In the two examples where a suitable SMB plan could not be reached, it happened that an acceptable NB plan was not viable either for the given contours. In these non-standard cases, it would be worth pursuing a non-coplanar beam solution, as this would afford superior flexibility in OAR avoidance. Should this also fail to provide a solution, it would then be the responsibility of the clinician to instruct editing of the contouring, adjustment of the dose prescription or to relax normal tissue tolerances. In such cases where dose is so critical and radiosensitive structures so proximal, it is unlikely that an SMB approach will ever be feasible. This highlights a need for individual patient selection if this kind of boosting strategy is to be adopted.

The TEP volumes, by the nature of their definition, are contiguous. This immediately creates the requirement of an instantaneous change in dose i.e. an infinite dose gradient when a different dose is prescribed to each volume. This is, of course, physically impossible and so the dose distribution across the TEP volumes will never be ideal. As demonstrated in Chapter 5, it's easier for the optimiser to achieve a larger dose

differential when the given volumes are enclosed within one another than neighbouring volumes. Given this, in all four examples, even though the TEP volumes were small (0.63-25.36 cm^3), a dose differential of 5 Gy between the boosts was feasible, given 5 x 5 mm^2 beamlets, showing that the HyMRI data can be used effectively. There is an argument for adding a planning margin to the TEP volumes themselves as these are, by definition, individual PTVs. However, given that the TEP volumes generally appear to be small, even a minimal margin would result in the boost volumes overlapping quite extensively (since they are already adjacent) and therefore depositing a different dose in each would be both unattainable and meaningless.

7.6. Conclusions

It has been shown that the HyMRI method is a simple and feasible means of defining concomitant boosts in brain IMRT, however, caution must be exercised when selecting patients for such specialised planning. Not every patient will be suitable for the SMB method given the proximity of certain sensitive structures relative to both the PTV and TEP derived boost volumes. However, for certain anatomies (two out of four cases in this instance), the SMB planning technique can be very effective, yielding a gain in TCP of 20 % against a non-boosted plan and this with insignificant differences in normal tissue dosimetry.

This is a dosimetric study only, carried out on a small sample of patients, the next step will be to explore oncologist-defined contours on a larger population undergoing radiotherapy as part of their treatment with real dose prescriptions. Finally, as the uncertainties associated with functional imaging data become better understood, this too could be incorporated into the HyMRI derived parameter map. A greater number of

more meaningful TEP levels will then be inferred and truly individualised ‘intelligent’ IMRT dose sculpting realised.

Chapter 8 - Conclusions

8.1. Conclusions

This thesis investigated and verified the feasibility of planning simultaneous high doses to multiple, carefully delineated sub-volumes within a tumour, for which the term selective multiple boosting (SMB) has been adopted. Experimental findings are summarised thus:

Chapter 4 – IPEX: A Plan Evaluation Environment for IMRT

- Novel software 'IPEX' developed in order to comprehensively evaluate and rank complex IMRT plans.
- Tools include DVH, dose ratio distribution (DRD), conformity, homogeneity and overall dose scores as well as biological evaluation methods encompassing gEUD, TCP and NTCP.
- Presents both standard tools and novel quality factors not available in commercial systems. Facilitates independent evaluation of plans on any standard PC, without the need for a specialist planning station thereby freeing up time on busy clinical workstations.

Chapter 5 – Controlling Local Dose for Selective Multiple Boosting

- Physical limits of dose deposition characterised for a given planning and delivery system by way of extensive planning experiments on two pseudo-boost anatomy models.
- Achievable dose gradients in multiple target planning quantified, suggesting that planning of concentric boost volumes on 5 mm scale is readily achievable, while

boosting neighbouring volumes e.g. voxel-by-voxel dose prescribing, is not realistic.

- Implies limits of value on multiple boost dose planning and delivery.

Chapter 6 – Significance of Image Resolution

- Functional imaging data often has poor resolution. MR-spectroscopy volumes are typically interpolated to a better resolution, which alters the MRS-defined boost volume.
- Effects of interpolation on these spectroscopy defined boost volumes investigated as an indicator of resolution importance given the proven dose deposition capabilities in Chapter 5.
- Issues of coarse resolution and subsequent interpolation do not pose a problem in the effective deployment of functional imaging defined boost volumes given current delivery technology.

Chapter 7 – Guiding Selective Multiple Boosting with Hybrid Images

- Image amalgamation software 'HyMRI' written for combining information gleaned from multiple image sequences.
- Parameter maps of tumour existence probability (TEP) can be derived within this environment and exported to the planning system, upon which differential dose prescribing may be effectively applied.
- May be useful in characterising uncertainties associated with a particular functional (or conventional) imaging technique for better individualised target volume definition.

Radiation therapy comprises an ongoing trade-off between tumour control and normal tissue complication. Given that disease control and both short and long term toxicity are dependent not only on dose, but on volume and fractionation as well, novel

treatment regimes in any of these regards are welcomed. The SMB methodology promises in fact to address all three of these influencing factors. A higher dose delivered to the boost volumes is proposed, which is proven to yield gains in disease control probability. Conversely, these high dose volumes are kept as small as possible in order minimise toxicity. For this to be effective, the small, high dose volumes must be delineated explicitly. Indeed, with the advent of conformal radiotherapy, specifically IMRT, target definition has become one of the main issues in improving the efficacy of radiation treatments. Advanced imaging techniques, in addition to a novel way of combining image information have been presented here to this end. Finally, the fractionation scheme is unusual in that the boost doses are delivered simultaneously to the standard target dose. This means not only that the most prolific tumour region is receiving a higher dose per fraction right from the start of treatment, but also that the treatment lasts no longer than a conventional regime.

Various SMB experiments presented here have predicted useful gains in gEUD and subsequent TCP values. Of course, when claims are made regarding gains in TCP or improved likelihood of local disease control, it is important not to take this out of the limited context for which it holds. That is, each step, be it target delineation or dose escalation is just one link in the radiotherapy chain, whereupon radiotherapy is just one link in the treatment chain of any one patient. Moreover, the true biological sensitivity of a given tumour to radiation is not well understood. Steeper dose response curves will naturally yield bigger gains from SMB plans. Paradoxically, steeper dose response curves arise from greater cell radiosensitivity and it is in fact radioresistant tumours that pose the greatest challenge to treatment strategies.

8.2. Future Work

8.2.1. Imaging

Functional imaging may augment conventional CT and MRI for defining tumour margins, and despite unresolved issues surrounding interpretation and robustness, this is strongly encouraged; there is no gold standard in tumour imaging and the more information available, the better the approximation. Further studies correlating image contrast to histology should be carried out before any dramatic change in target volumes could really be adopted. Different techniques are likely to develop synergistically for addressing tumour targeting, for which HyMRI will prove a valuable tool in deriving composite targets.

Any imaging technique has multiple uncertainties associated with it in terms of reproducibility, completeness, post processing and/or registration issues. Where such uncertainties arise, it would be useful to incorporate these into the optimisation loop somehow, such as by the development of probability-based contouring. Also, in the context of a spatial biological map derived from functional imaging, contour-based planning can become cumbersome. A probability based approach would also resolve issues of overlap between target volumes and/or OARs. Changes in tumour volume occur as growth or shrinkage during any course of radiotherapy. Any given image, however accurate and whatever the resolution, is only one snapshot in time. Inter-fraction imaging may become essential for consistent, effective and precise targeting throughout the treatment course.

8.2.2. Planning & delivery

Whilst finer delivery resolution may be possible with micro-MLCs, a consequence of this is longer treatment times and increased leakage radiation. Besides, as more precise treatments do evolve, other issues such as the dose calculation algorithm and type of optimisation engine become more important. The applicability of extremely precise, multiple boosts will always be limited by other factors anyway including inter- and intra-fraction motion, physical and biological changes in tumour volume during treatment course and contouring errors, as well as numerous physical and technical limitations of IMRT. An ongoing concern with IMRT is that a change in patient position could bring some region of the target volume outside the high dose region, or a nearby critical structure into the high dose, emphasising the need for both margins and caution. It will take the implementation of new technologies such as adaptive radiotherapy and proton facilities to make even more precise delivery a useful pursuit. As treatments become ever more precise and complex, plan evaluation must develop in a similar vein. Simple tools such as the conformity index are useful but still not used much in clinical practice – commercial planning systems must evolve and adapt to offer these new functions.

8.2.3. Clinical decision-making

It has been shown in this thesis that it is quite possible to define regions that would potentially benefit from escalated dose and deliver this escalated dose via photon IMRT. The next remaining question is what dose level should be prescribed to these regions, which is a biological rather than physical problem. More clinical data is therefore needed in order to assess the rate of both local control and complication as a function of dose for these SMB treatments. Current dose response data is limited to older, conventional treatments and very little information has been collated for any conformal

therapy or IMRT. SMB plans further exacerbate this limitation. Various metrics will necessarily evolve in order to quantify biological status based on imaging and the corresponding dose-response relationship, which will require experimental determination like the empirical dose recommendations currently used for different disease sites. High grade brain tumours, specifically glioblastoma (GBM) have accelerated repopulation, proposed as a reason for local failure. In addition to dose-response studies, by considering and understanding proliferation effects, certain tumours may be found to have higher TCPs with non-standard fractionation schemes e.g. hyperfractionation.

In preparing complex SMB treatments, there is not only patient time to consider but also staff hours required for both training and the safe deployment of treatment. IMRT planning is itself an involved and time-consuming process and deriving planning volumes from advanced imaging methods will further add to this burden. It is imperative that the benefit of these state-of-the-art treatments is sure to outweigh the increase in both patient scan times and expert man hours required to deliver these treatments. Furthermore, not all patients will be suitable for SMB planning, in terms of enduring longer scans or because of the location or size of their tumour. Careful patient selection criteria must be adopted then in order to reap the maximum benefit of the SMB methodology.

Whilst new methods of imaging and dose optimisation have been presented both here and elsewhere in the literature, medical decision-making must next be integrated for any real improvement in cancer management to be realised. In terms of real patient benefit in improving both local control and survival, further multi-modality approaches are anticipated, with molecular targeting playing a key role alongside radiation dose

escalation and new chemotherapy drugs. The importance of multimodality treatments e.g. chemical sensitisers, given the limits of radiation dose escalation, cannot be emphasised enough.

Further work then is envisaged involving both functional imaging improvements and correlation of this to dose-response models in order to take the SMB framework forward into useful clinical practice.

References

1. Burnet NG, Jefferies SJ, Benson RJ, Hunt DP*, et al.* Years of life lost (YLL) from cancer is an important measure of population burden--and should be considered when allocating research funds. *Br J Cancer* 2005;**92**(2):241-5.

2. Vick NA and Paleologos NA. External beam radiotherapy: hard facts and painful realities. *J Neurooncol* 1995;**24**(1):93-5.

3. Balducci M, Mattiucci GC, Dinapoli N, Bavasso A*, et al.* Impact of dose and volume on the tolerance of central nervous system. *Rays* 2005;**30**(2):189-95.

4. Levegrun S, Ton L, and Debus J. Partial irradiation of the brain. *Semin Radiat Oncol* 2001;**11**(3):259-67.

5. Beavis AW, Gibbs P, Dealey RA, and Whitton VJ. Radiotherapy treatment planning of brain tumours using MRI alone. *Br J Radiol* 1998;**71**(845):544-8.

6. Suzuki M, Nakamatsu K, Kanamori S, Okumra M*, et al.* Feasibility study of the simultaneous integrated boost (SIB) method for malignant gliomas using intensity-modulated radiotherapy (IMRT). *Jpn J Clin Oncol* 2003;**33**(6):271-7.

7. Dogan N, King S, Emami B, Mohideen N*, et al.* Assessment of different IMRT boost delivery methods on target coverage and normal-tissue sparing. *Int J Radiat Oncol Biol Phys* 2003;**57**(5):1480-91.

8. Tanaka M, Ino Y, Nakagawa K, Tago M*, et al.* High-dose conformal radiotherapy for supratentorial malignant glioma: a historical comparison. *Lancet Oncol* 2005;**6**(12):953-60.

9. Shaw EG, Stieber V., Tatter S., Ellis T., Hinson W., Kearns W., Bourland J.D., Munley M., Lesser G., Stanton C. A Phase I Dose Escalating Study of Intensity Modulated Radiation Therapy (IMRT) for the Treatment of Glioblastoma Multiforme (GBM). in *Annual ASTRO Meeting*. 2002. New Orleans, LA.: Int J Radiation Oncol Biol Phys.

10. Chan JL, Lee SW, Fraass BA, Normolle DP*, et al.* Survival and failure patterns of high-grade gliomas after three-dimensional conformal radiotherapy. *J Clin Oncol* 2002;**20**(6):1635-42.

11. Tome WA and Fowler JF. Selective boosting of tumor subvolumes. *Int J Radiat Oncol Biol Phys* 2000;**48**(2):593-9.

12. Kirkby NF, Jefferies SJ, Jena R, and Burnet NG. A mathematical model of the treatment and survival of patients with high-grade brain tumours. *J Theor Biol* 2006.

13. Burnet NG, Jena R, Jefferies SJ, Stenning SP*, et al.* Mathematical modelling of survival of glioblastoma patients suggests a role for radiotherapy dose escalation and predicts poorer outcome after delay to start treatment. *Clin Oncol (R Coll Radiol)* 2006;**18**(2):93-103.

14. Mehta MP, Tsao MN, Whelan TJ, Morris DE*, et al.* The American Society for Therapeutic Radiology and Oncology (ASTRO) evidence-based review of the role of radiosurgery for brain metastases. *Int J Radiat Oncol Biol Phys* 2005;**63**(1):37-46.

15. Lee SW, Fraass BA, Marsh LH, Herbort K*, et al.* Patterns of failure following high-dose 3-D conformal radiotherapy for high-grade astrocytomas: a quantitative dosimetric study. *Int J Radiat Oncol Biol Phys* 1999;**43**(1):79-88.

16. Garcia-Alvarez R, Liney G, and Beavis A. Use of functional magnetic resonance imaging for planning intensity modulated radiotherapy. *Journal of Radiotherapy in Practice* 2003;**3**:66-72.

17. Alber M, Paulsen F, Eschmann SM, and Machulla HJ. On biologically conformal boost dose optimization. *Phys Med Biol* 2003;**48**(2):N31-5.

18. Ling CC, Humm J, Larson S, Amols H*, et al.* Towards multidimensional radiotherapy (MD-CRT): biological imaging and biological conformality. *Int J Radiat Oncol Biol Phys* 2000;**47**(3):551-60.

19. Jena R, Price SJ, Baker C, Jefferies SJ*, et al.* Diffusion tensor imaging: possible implications for radiotherapy treatment planning of patients with high-grade glioma. *Clin Oncol (R Coll Radiol)* 2005;**17**(8):581-90.

20. Hamilton RJ, Sweeney PJ, Pelizzari CA, Yetkin FZ*, et al.* Functional imaging in treatment planning of brain lesions. *Int J Radiat Oncol Biol Phys* 1997;**37**(1):181-8.

21. Pirzkall A, McKnight TR, Graves EE, Carol MP*, et al.* MR-spectroscopy guided target delineation for high-grade gliomas*. Int J Radiat Oncol Biol Phys* 2001;**50**(4):915-28.

22. Xing L, Cotrutz C, Hunjan S, Boyer AL*, et al.* Inverse planning for functional image-guided intensity-modulated radiation therapy*. Phys Med Biol* 2002;**47**(20):3567-78.

23. Yang Y and Xing L. Towards biologically conformal radiation therapy (BCRT): selective IMRT dose escalation under the guidance of spatial biology distribution*. Med Phys* 2005;**32**(6):1473-84.

24. Liney G and Beavis A, in *Recent Research Developments in Radiology*. 2003, Transworld Research Network. p. 39-61.

25. Garcia-Alvarez R, Liney GP, and Beavis AW. Repeatability of functional MRI for conformal avoidance radiotherapy planning*. J Magn Reson Imaging* 2006;**23**(2):108-14.

26. Stupp R, Mason WP, van den Bent MJ, Weller M*, et al.* Radiotherapy plus concomitant and adjuvant temozolomide for glioblastoma*. N Engl J Med* 2005;**352**(10):987-96.

27. ICRU, Prescribing, Recording and Reporting Photon Beam Therapy. 1992, International Commission on Radiation Units and Measurements: Washington.

28. van Herk M. Errors and margins in radiotherapy*. Semin Radiat Oncol* 2004;**14**(1):52-64.

29. ICRU, Prescribing, Recording and Reporting Photon Beam Therapy. 1999, International Commission on Radiation Units and Measurements: Washington.

30. Xia P, Yu N, Xing L, Sun X*, et al.* Investigation of using a power function as a cost function in inverse planning optimization*. Med Phys* 2005;**32**(4):920-7.

31. Wu Q and Mohan R. Algorithms and functionality of an intensity modulated radiotherapy optimization system*. Med Phys* 2000;**27**(4):701-11.

32. Lian J and Xing L. Incorporating model parameter uncertainty into inverse treatment planning*. Med Phys* 2004;**31**(9):2711-20.

33. Jones L and Hoban P. A comparison of physically and radiobiologically based optimization for IMRT. *Med Phys* 2002;**29**(7):1447-55.

34. Stavrev P, Hristov D, Warkentin B, Sham E*, et al.* Inverse treatment planning by physically constrained minimization of a biological objective function. *Med Phys* 2003;**30**(11):2948-58.

35. Llacer J, Deasy JO, Portfeld TR, Solberg TD*, et al.* Absence of multiple local minima effects in intensity modulated optimization with dose-volume constraints. *Phys Med Biol* 2003;**48**(2):183-210.

36. Shaw E, Arusell R, Scheithauer B, O'Fallon J*, et al.* Prospective randomized trial of low- versus high-dose radiation therapy in adults with supratentorial low-grade glioma: initial report of a North Central Cancer Treatment Group/Radiation Therapy Oncology Group/Eastern Cooperative Oncology Group study. *J Clin Oncol* 2002;**20**(9):2267-76.

37. Xing L, Hamilton RJ, Spelbring D, Pelizzari CA*, et al.* Fast iterative algorithms for three-dimensional inverse treatment planning. *Med Phys* 1998;**25**(10):1845-9.

38. Chui CS and Spirou SV. Inverse planning algorithms for external beam radiation therapy. *Med Dosim* 2001;**26**(2):189-97.

39. Deasy JO. Multiple local minima in radiotherapy optimization problems with dose-volume constraints. *Med Phys* 1997;**24**(7):1157-61.

40. Wu Q and Mohan R. Multiple local minima in IMRT optimization based on dose-volume criteria. *Med Phys* 2002;**29**(7):1514-27.

41. Mattia M, Del Giudice P, and Caccia B. IMRT optimization: variability of solutions and its radiobiological impact. *Med Phys* 2004;**31**(5):1052-60.

42. Zhang X, Liu H, Wang X, Dong L*, et al.* Speed and convergence properties of gradient algorithms for optimization of IMRT. *Med Phys* 2004;**31**(5):1141-52.

43. Wiesmeyer M. Effective cost functions as a means of improving user interaction in constrained gradient-based inverse planning. in *ICCR*. 2004. Seoul, Korea.

44. Bortfeld T, Stein J, and Preiser K. Clinically relevant intensity modulation optimization using physical criteria. in *XII International Conference on the Use of Computers in Radiation Therapy*. 1997. Madison, WI: Medical Physics Publishing.

45. Bloch F, Hansen W, and Packard M. Nuclear induction. *Phys Rev* 1946;**69**:127.

46. Lee YK, Bollet M, Charles-Edwards G, Flower MA*, et al.* Radiotherapy treatment planning of prostate cancer using magnetic resonance imaging alone. *Radiother Oncol* 2003;**66**(2):203-16.

47. Bruhn H, Frahm J, Gyngell ML, Merboldt KD*, et al.* Noninvasive differentiation of tumors with use of localized H-1 MR spectroscopy in vivo: initial experience in patients with cerebral tumors. *Radiology* 1989;**172**(2):541-8.

48. Fulham MJ, Bizzi A, Dietz MJ, Shih HH*, et al.* Mapping of brain tumor metabolites with proton MR spectroscopic imaging: clinical relevance. *Radiology* 1992;**185**(3):675-86.

49. Pirzkall A, Nelson SJ, McKnight TR, Takahashi MM*, et al.* Metabolic imaging of low-grade gliomas with three-dimensional magnetic resonance spectroscopy. *Int J Radiat Oncol Biol Phys* 2002;**53**(5):1254-64.

50. Garcia-Alvarez R, Use of advanced magnetic resonance techniques for radiotherapy planning of intracranial tumours, in *Postgraduate Medical Institute*. 2006, University of Hull: Hull.

51. Graves EE, Pirzkall A, Nelson SJ, Larson D*, et al.* Registration of magnetic resonance spectroscopic imaging to computed tomography for radiotherapy treatment planning. *Med Phys* 2001;**28**(12):2489-96.

52. Nelson SJ, Graves E, Pirzkall A, Li X*, et al.* In vivo molecular imaging for planning radiation therapy of gliomas: an application of 1H MRSI. *J Magn Reson Imaging* 2002;**16**(4):464-76.

53. Pirzkall A, Li X, Oh J, Chang S*, et al.* 3D MRSI for resected high-grade gliomas before RT: tumor extent according to metabolic activity in relation to MRI. *Int J Radiat Oncol Biol Phys* 2004;**59**(1):126-37.

54. Price SJ, Burnet NG, Donovan T, Green HA*, et al.* Diffusion tensor imaging of brain tumours at 3T: a potential tool for assessing white matter tract invasion? *Clin Radiol* 2003;**58**(6):455-62.

55. Price SJ, Pena A, Burnet NG, Pickard JD*, et al.* Detecting glioma invasion of the corpus callosum using diffusion tensor imaging. *Br J Neurosurg* 2004;**18**(4):391-5.

56. Price SJ, Pena A, Burnet NG, Jena R*, et al.* Tissue signature characterisation of diffusion tensor abnormalities in cerebral gliomas. *Eur Radiol* 2004;**14**(10):1909-17.

57. Price SJ, Jena R, Burnet NG, Carpenter TA*, et al.* Predicting patterns of glioma recurrence using diffusion tensor imaging. *Eur Radiol* 2007;**17**(7):1675-84.

58. Minn H. PET and SPECT in low-grade glioma. *Eur J Radiol* 2005;**56**(2):171-8.

59. Heiss WD, Habedank B, Klein JC, Herholz K*, et al.* Metabolic rates in small brain nuclei determined by high-resolution PET. *J Nucl Med* 2004;**45**(11):1811-5.

60. Townsend DW, Beyer T, and Blodgett TM. PET/CT scanners: a hardware approach to image fusion. *Semin Nucl Med* 2003;**33**(3):193-204.

61. Wagner TH, Bova FJ, Friedman WA, Buatti JM*, et al.* A simple and reliable index for scoring rival stereotactic radiosurgery plans. *Int J Radiat Oncol Biol Phys* 2003;**57**(4):1141-9.

62. Feuvret L, Noel G, Mazeron JJ, and Bey P. Conformity index: a review. *Int J Radiat Oncol Biol Phys* 2006;**64**(2):333-42.

63. Knoos T, Kristensen I, and Nilsson P. Volumetric and dosimetric evaluation of radiation treatment plans: radiation conformity index. *Int J Radiat Oncol Biol Phys* 1998;**42**(5):1169-76.

64. Nakamura JL, Verhey LJ, Smith V, Petti PL*, et al.* Dose conformity of gamma knife radiosurgery and risk factors for complications. *Int J Radiat Oncol Biol Phys* 2001;**51**(5):1313-9.

65. Lomax NJ and Scheib SG. Quantifying the degree of conformity in radiosurgery treatment planning. *Int J Radiat Oncol Biol Phys* 2003;**55**(5):1409-19.

66. Nedzi LA, Kooy HM, Alexander E, 3rd, Svensson GK, *et al.* Dynamic field shaping for stereotactic radiosurgery: a modeling study. *Int J Radiat Oncol Biol Phys* 1993;**25**(5):859-69.

67. van't Riet A, Mak AC, Moerland MA, Elders LH, *et al.* A conformation number to quantify the degree of conformality in brachytherapy and external beam irradiation: application to the prostate. *Int J Radiat Oncol Biol Phys* 1997;**37**(3):731-6.

68. Dice LR. Measures of the amount of ecologic association between species. *Ecology* 1945;**26**:297-302.

69. Leung LH, Chua DT, and Wu PM. A new tool for dose conformity evaluation of radiosurgery treatment plans. *Int J Radiat Oncol Biol Phys* 1999;**45**(1):233-41.

70. Weber DC, Bogner J, Verwey J, Georg D, *et al.* Proton beam radiotherapy versus fractionated stereotactic radiotherapy for uveal melanomas: A comparative study. *Int J Radiat Oncol Biol Phys* 2005;**63**(2):373-84.

71. Shaw E, Kline R, Gillin M, Souhami L, *et al.* Radiation Therapy Oncology Group: radiosurgery quality assurance guidelines. *Int J Radiat Oncol Biol Phys* 1993;**27**(5):1231-9.

72. Wu Q, Mohan R, Morris M, Lauve A, *et al.* Simultaneous integrated boost intensity-modulated radiotherapy for locally advanced head-and-neck squamous cell carcinomas. I: dosimetric results. *Int J Radiat Oncol Biol Phys* 2003;**56**(2):573-85.

73. Kinhikar RA, Deshpande SS, Mahantshetty U, Sarin R, *et al.* HDR brachytherapy combined with 3-D conformal vs. IMRT in left-sided breast cancer patients including internal mammary chain: comparative analysis of dosimetric and technical parameters. *J Appl Clin Med Phys* 2005;**6**(3):1-12.

74. Niemierko A. Reporting and analyzing dose distributions: a concept of equivalent uniform dose. *Med Phys* 1997;**24**(1):103-10.

75. Niemierko A. A generalized concept of equivalent uniform dose (EUD). *Med Phys* 1999;**26**(6):1100.

76. Emami B, Lyman J, Brown A, Coia L*, et al.* Tolerance of normal tissue to therapeutic irradiation. *Int J Radiat Oncol Biol Phys* 1991;**21**(1):109-22.

77. Burman C, Kutcher GJ, Emami B, and Goitein M. Fitting of normal tissue tolerance data to an analytic function. *Int J Radiat Oncol Biol Phys* 1991;**21**(1):123-35.

78. Li XA, Wang JZ, Stewart RD, and DiBiase SJ. Dose escalation in permanent brachytherapy for prostate cancer: dosimetric and biological considerations. *Phys Med Biol* 2003;**48**(17):2753-65.

79. Wu Q, Mohan R, Niemierko A, and Schmidt-Ullrich R. Optimization of intensity-modulated radiotherapy plans based on the equivalent uniform dose. *Int J Radiat Oncol Biol Phys* 2002;**52**(1):224-35.

80. Webb S and Nahum AE. A model for calculating tumour control probability in radiotherapy including the effects of inhomogeneous distributions of dose and clonogenic cell density. *Phys Med Biol* 1993;**38**(6):653-66.

81. Lyman JT. Complication probability as assessed from dose-volume histograms. *Radiat Res Suppl* 1985;**8**:S13-9.

82. Ebert MA. Viability of the EUD and TCP concepts as reliable dose indicators. *Phys Med Biol* 2000;**45**(2):441-57.

83. Kron T, Grigorov G, Yu E, Yartsev S*, et al.* Planning evaluation of radiotherapy for complex lung cancer cases using helical tomotherapy. *Phys Med Biol* 2004;**49**(16):3675-90.

84. Ma L, Kwok Y, Chin LS, Simard JM*, et al.* Concomitant GRID boost for Gamma Knife radiosurgery. *Med Phys* 2005;**32**(11):3419-23.

85. Marks LB, Sherouse GW, Munley MT, Bentel GC*, et al.* Incorporation of functional status into dose-volume analysis. *Med Phys* 1999;**26**(2):196-9.

86. Miften MM, Das SK, Su M, and Marks LB. Incorporation of functional imaging data in the evaluation of dose distributions using the generalized concept of equivalent uniform dose. *Phys Med Biol* 2004;**49**(9):1711-21.

87. Das SK, Miften MM, Zhou S, Bell M*, et al.* Feasibility of optimizing the dose distribution in lung tumors using fluorine-18-fluorodeoxyglucose positron emission tomography and single photon emission computed tomography guided dose prescriptions. *Med Phys* 2004;**31**(6):1452-61.

88. Yang Y and Xing L. Inverse treatment planning with adaptively evolving voxel-dependent penalty scheme. *Med Phys* 2004;**31**(10):2839-44.

89. Bova FJ, Meeks SL, Friedman WA, and Buatti JM. Stereotactic plan evaluation tool "the UFX index". in *41st Annual ASTRO Meeting*. 1999. San Antonio, Texas: Int J Radiation Oncology Biol Phys.

90. Bortfeld T, Oelfke U, and Nill S. What is the optimum leaf width of a multileaf collimator? *Med Phys* 2000;**27**(11):2494-502.

91. Nelson SJ. Multivoxel magnetic resonance spectroscopy of brain tumors. *Mol Cancer Ther* 2003;**2**(5):497-507.

92. McKnight TR, Noworolski SM, Vigneron DB, and Nelson SJ. An automated technique for the quantitative assessment of 3D-MRSI data from patients with glioma. *J Magn Reson Imaging* 2001;**13**(2):167-77.

93. Star-Lack J, Vigneron DB, Pauly J, Kurhanewicz J*, et al.* Improved solvent suppression and increased spatial excitation bandwidths for three-dimensional PRESS CSI using phase-compensating spectral/spatial spin-echo pulses. *J Magn Reson Imaging* 1997;**7**(4):745-57.

94. Garcia-Alvarez R, Liney G, Manton D, and Beavis A. Spectro Image Registration And Metabolite-mApping Software, a new open source research tool. in *ISMRM*. 2005. Miami, Florida, US.

95. Maes F, Collignon A, Vandermeulen D, Marchal G*, et al.* Multimodality image registration by maximization of mutual information. *IEEE Trans Med Imaging* 1997;**16**(2):187-98.

96. Moore CS, Liney GP, and Beavis AW. Quality assurance of registration of CT and MRI data sets for treatment planning of radiotherapy for head and neck cancers. *J Appl Clin Med Phys* 2004;**5**(1):25-35.

97. Hunjan S, Adalsteinsson E, Kim DH, Harsh GR, *et al.* Quality assurance of magnetic resonance spectroscopic imaging-derived metabolic data. *Int J Radiat Oncol Biol Phys* 2003;**57**(4):1159-73.

98. Chang J, Thakur S, Perera G, Kowalski A, *et al.* Image-fusion of MR spectroscopic images for treatment planning of gliomas. *Med Phys* 2006;**33**(1):32-40.

99. Meyer CR, Boes JL, Kim B, Bland DH, *et al.* Demonstration of accuracy and clinical versatility of mutual information for automatic multimodality image fusion using affine and thin-plate spline warped geometric deformations. *Medical Image Analysis* 1996/7;**1**(3):195-206.

100. Maes F, Vandermeulen D, and Suetens P. Comparative evaluation of multiresolution optimization strategies for mulitmodality image registration by maximization of mutual information. *Medical Image Analysis* 1999;**3**(4):373-86.

101. Manon R, Hui S, Chinnaiyan P, Suh J, *et al.* The impact of mid-treatment MRI on defining boost volumes in the radiation treatment of glioblastoma multiforme. *Technol Cancer Res Treat* 2004;**3**(3):303-7.

102. Narayana A, Yamada J, Berry S, Shah P, *et al.* Intensity-modulated radiotherapy in high-grade gliomas: clinical and dosimetric results. *Int J Radiat Oncol Biol Phys* 2006;**64**(3):892-7.